BANK OF ENGLAND

# Working Paper No. 449

# Misperceptions, heterogeneous expectations and macroeconomic dynamics

Richard Harrison and Tim Taylor

May 2012

BANK OF ENGLAND

# Working Paper No. 449
# Misperceptions, heterogeneous expectations and macroeconomic dynamics

Richard Harrison[1] and Tim Taylor[2]


## Abstract

We investigate the extent to which misperceptions about the economy can become self-reinforcing and thereby contribute to time-varying macroeconomic dynamics. To do so, we build a New Keynesian model with long-horizon expectations and dynamic predictor selection. Because agents solve multi-period optimisation problems (households maximise expected lifetime utility and firms maximise the discounted flow of future profits), their current decisions are influenced by expectations of the distant future and cannot in general be characterised by the familiar Euler equations that represent the rational expectations equilibrium of these models. We assume that agents have access to a set of alternative predictors that can be used to form expectations and choose among them based on noisy measures of their recent performance. This dynamic predictor selection generates endogenous fluctuations in the proportions of agents using each predictor, contributing to macroeconomic dynamics. We explore the behaviour of our model when agents have access to two simple predictors. One of the predictors is consistent with a mistaken belief that macroeconomic variables are more persistent than implied by the fundamental shocks hitting the economy. We show that the presence of a 'persistent predictor' can lead to changes in beliefs which are self-reinforcing, giving rise to endogenous fluctuations in the time-series properties of the economy. Moreover, we show that such fluctuations arise even if we replace the 'persistent predictor' with learning under constant gain.


**Key words:** Expectations, macroeconomic dynamics, heuristics.

**JEL classification:** E17, D82, D84.

(1) Bank of England. Email: richard.harrison@bankofengland.co.uk
(2) Bank of England. Email: tim.taylor@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. The authors wish to thank Charlie Bean, Spencer Dale, Cars Hommes, Tom Sargent, Peter Sinclair, Mike Woodford, Tony Yates, an anonymous referee and seminar participants at de Nederlandsche Bank and the Bank of England for useful discussions on earlier versions of this paper. All errors are ours. This paper was finalised on 20 February 2012.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/Pages/workingpapers/default.aspx

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk


ISSN 1749-9135 (on-line)

**Contents**

**Summary**

An important question for economic policy makers is the extent to which the expectations of key decision makers in the economy affect – and are affected by – economic outturns. In particular, it is possible that mistaken beliefs about the behaviour of the economy can influence the behaviour of households and firms in a self-fulfilling manner. For example, a belief that inflation will be more persistent could influence price-setting behaviour so that actual inflation turns out to be more persistent. Such a feedback could reinforce the initial belief causing more households and firms to believe that inflation will be persistent.

This type of mechanism is illustrated in the following quote from the Bank of England's February 2008 *Inflation Report*: 'If households' and businesses' medium-term inflation expectations are heavily influenced by their recent experience, then repeated above-target outturns may cause them to place weight on the assumption that inflation will be persistently above [the inflation target of] 2%. If those expectations were built into higher wages and prices, that would raise medium-term inflationary pressures.'

To investigate this phenomenon, we build a small macroeconomic model in which the decisions of households and firms depend on their expectations for future income and costs, so that spending and price-setting decisions depend on expectations extending into the distant future. We assume that, to form their expectations, households and firms have access to a small set of alternative 'predictors'. These predictors are simple forecasting equations for relevant variables (for example, future inflation could be forecast by inputting recent observations for inflation into a simple equation). Households and firms choose between these predictors based on their recent forecasting performance. So a predictor that has forecast (say) inflation very well over the last few quarters will tend to be used more than a predictor with a worse forecasting record.

This 'dynamic predictor selection' creates the possibility of a feedback process between beliefs about the behaviour of the economy and its actual behaviour. We find that it is straightforward to generate this type of effect in our model under the assumption that households and firms choose between two predictors. The first predictor has very good properties when used by all households and firms. Its forecasting performance is close to the best possible predictor (the 'rational

expectation'). The second predictor is a 'misperceptions predictor' which embodies a mistaken belief that inflation is more persistent. When we simulate the model, we are able to generate occasional periods of high, volatile and persistent inflation. This occurs when (random) shocks generate enough persistence in the inflation rate observed by households and firms to lead more of them to choose expectations based on the misperceptions predictor.

## 1 Introduction

The notion that expectations may be a source of economic fluctuations has a long history, dating back at least to Pigou (1929) and Keynes (1936). But rational expectations is the dominant paradigm in modern macroeconomics and removes the possibility that endogenous changes in the expectations formation process may be an important source of economic dynamics.[1] The sharp changes in the time-series properties of macroeconomic data that we observe are difficult to account for in standard rational expectations models.[2]

Of course, academics have long debated the extent to which rational expectations can be considered reasonable *a priori*. Lucas (1986, page S402) argues that:

> the question whether people are in general 'rational' or 'adaptive' does not seem to me worth arguing over. Which of these answers is most useful will depend on the situations in which we are trying to predict behavior and on the experiences the people in question have had with such situations.

Monetary policy makers recognise the importance of this issue too. Bernanke (2007) asks: 'What is the right conceptual framework for thinking about inflation expectations in the current context?'[3] He argues that models of learning are useful deviations from the assumption of rational expectations, which seems inappropriate when there is structural change or uncertainty about policy objectives. Indeed, there is now a vast literature exploring models in which agents update the way they form expectations in the light of new data, capturing learning about the economy.[4]

[1]There is of course an extensive literature that examines the behaviour of 'indeterminate' rational expectations models. A recent example, applied to monetary policy issues, is Lubik and Schorfheide (2004). Models in which rational behaviour gives rise to 'higher-order beliefs' have also been examined (see for example, Morris and Shin (2003)).

[2]Within the rational expectations paradigm the focus has largely been to investigate structural change, reductions in the variance of the shocks hitting the economy ('good luck') or changes in policy behaviour to explain these properties.

[3]Of course, many policymakers have voiced concerns about the possible sensitivity of the economy to alternative assumptions about the way firms and households form expectations (see, for example, King (2005)).

[4]A key reference is Evans and Honkapohja (2001) and early contributions include Bray (1982), Marcet and Sargent (1989) and Sargent (1993). Evans and Honkapohja (2008) provide an excellent recent literature review, with a focus on monetary policy applications. Sargent (2007) relates the lessons from learning models to a variety of historical episodes.

The presence of non-rational expectations creates the potential for dynamic feedback between macroeconomic outcomes and expectations, as discussed, for example, in the Bank of England's February 2008 *Inflation Report* (page 45):

> If households' and businesses' medium-term inflation expectations are heavily influenced by their recent experience, then repeated above-target outturns may cause them to place weight on the assumption that inflation will be persistently above [the inflation target of] 2%. If those expectations were built into higher wages and prices, that would raise medium-term inflationary pressures.

In this paper we explore the dynamic feedback between outcomes and expectations in a benchmark New Keynesian model with long-horizon expectations and dynamic predictor selection. Specifically, we introduce dynamic predictor selection into the model of Harrison and Taylor (2012). That model solves the decision rules of households and firms conditional on their expectations for future events that are outside of their control. This gives rise to spending and price-setting decisions that depend on 'long-horizon expectations' as noted by Preston (2005). Households' consumption decisions depend on the discounted sum of expected future income and real interest rates and firms' prices depend on the discounted sum of their expected wage costs, productivity and inflation. This approach differs from the widely used 'Euler equation' approach, which simply replaces the one period ahead expectations terms in the Euler equations describing optimal behaviour with a non-rational expectations term. Harrison and Taylor (2012) find that, as actual expectations deviate further from rational expectations, the 'long-horizon' and 'Euler equation' representations of a New Keynesian model can generate rather different macroeconomic dynamics. While both approaches can be defended as descriptions of (distinct) forms of boundedly rational behaviour (see, for example, Branch and McGough (2006b)), we regard the description implied by the long-horizon expectations approach as more consistent with the underlying microfoundations.

The feedback between outcomes and expectations is generated by dynamic predictor selection: the notion that agents choose between a small set of forecasting rules (or 'predictors') based on noisy observations of past performance. This approach has several advantages. First, it can bring about time-variation in the persistence and variance of macroeconomic variables, as shown by Brazier, Harrison, King and Yates (2008). Second, it is consistent with experimental evidence about the way expectations are formed, and how the use of these rules varies over time (see, for example, Anufriev and Hommes (2006)). And third, it allows different expectational rules to coexist: heterogeneity appears to be a feature of the survey data on inflation expectations.[5] We introduce predictor choice into our long-horizon expectations model using the techniques developed by Brock and Hommes (1997), Brock and de Fontnouvelle (2000) and Branch and Evans (2006), who applied the discrete decision, multinomial logit models set out in Manski and McFadden (1981).

We explore the interaction between the beliefs of firms and households and the properties of our model economy. We show that fluctuations in the expectations formation process – the dynamic predictor selection which occurs endogenously in response to shocks – can bring about changes in the properties of the model that act (at least temporarily) to reinforce the initial change in beliefs. Misperceptions can for a time be self-reinforcing and this mechanism is a candidate explanation for time-variation in the moments of data. Our approach is to investigate the behaviour of our model when agents have access to a predictor that (wrongly) anticipates that shocks to inflation will be long-lasting. The quote from the *Bank of England Inflation Report* above suggests that such a possibility may be of practical policy interest. We show that misperceptions can indeed be temporarily self-reinforcing, leading to marked time-variation in the time-series properties of the data. Moreover, this result survives even when agents' misperceptions are generated by a constant-gain learning algorithm.

Our paper touches on several (large) strands of the literature, as noted above, though several recent papers are closely related to our analysis. Branch and McGough (2006b) introduce heterogeneous expectations in the familiar setting of a New Keynesian model. In subsequent work, Branch and McGough (2006a) endogenise the proportion of agents using each predictor and examine the properties of a set of monetary policy rules. de Grauwe (2008) also analyses dynamic predictor selection in a New Keynesian model and in particular allows agents to hold overly optimistic or pessimistic expectations of future output. Our main innovations are that we examine predictor choice in the context of a model with long-horizon expectations and we allow

[5]For example, Lombardelli and Saleheen (2003) document the heterogeneity among respondants to the Bank of England/NOP survey. And Appendix B reports heterogeneity between households and firms.

one of the predictors to be updated over time via constant-gain learning.

There has been recent interest in estimating models of this type. Milani (2007) estimates a model in which learning rules with different gain parameters are dynamically selected (though he does not use a model with long-horizon expectations). Milani (2006) estimates a long-horizon expectations model under constant-gain learning, but under the assumption that expectations are homogenous across agents. Our model allows for both heterogeneity of individual agents (eg across different households) and between types of agents (eg firms and households may have access to different sets of predictors). Recently, Slobodyan and Wouters (2009) have estimated a DSGE model in which expectations are formed as a weighted average of a small set of simple predictors (though the predictor choice is based on Bayesian information criteria). However, the non-rational expectations version of their DSGE model has an Euler equation representation. Estimation of a model like ours (with both long-horizon expectations and dynamic predictor selection) would be a welcome addition to the literature.

The key contributions of this paper are twofold. First, we analyse the interaction of heterogeneous expectations among both household and firms. Second, we assess the implications of specific misperceptions about macroeconomic dynamics (eg the persistence of inflation) that are likely to be of interest to policymakers.

Section 2 summarises the model set out in Harrison and Taylor (2012), which incorporates standard New Keynesian assumptions. We also describe the assumptions about expectations formation and the parameterisation of the model. Section 3 presents our results, demonstrating that the presence of misperceptions about the persistence of macroeconomic data can be self-reinforcing. Section 4 concludes.

## 2 The model

We use the baseline New Keynesian model set out in Harrison and Taylor (2012), which provides a more detailed derivation. Here we simply present the maximisation problems of the agents in the model and the associated log-linearised decision rules.

### *2.1 Households*

There is a continuum of households of unit mass, indexed by $i \in (0, 1)$. We assume that household $i$ solves:

$$\max \tilde{E}_{i,t-1} \sum_{s=t}^{\infty} \beta^{s-t} \left[ \frac{c_{i,s}^{1-\sigma} - 1}{1-\sigma} - \chi \frac{h_{i,s}^{1+\gamma}}{1+\gamma} \right]$$

subject to

$$b_{i,s} - \frac{R_{s-1}}{\Pi_s} b_{i,s-1} - w_s h_{i,s} - d_s + c_{i,s} = 0 \qquad \textbf{(1)}$$

where $c$ is consumption and $h$ is hours worked and the parameters $\sigma$ and $\gamma$ are both strictly positive.[6] The budget constraint **(1)** is written in real terms and $b$ represents the household's holdings of (one-period) bonds, $w$ is the real wage, $R$ measures the nominal interest rate paid on bonds, $\Pi$ is the inflation rate and $d$ is a collection of transfers (from government and firms). We use the notation $\tilde{E}_i$ to denote the expectations of household $i$. The '~' notation signals that the expectation is not rational and the $i$ subscript makes it clear that the expectation is specific to the individual household. We impose the assumption that when making decisions about consumption at date $t$ and beyond, the household has access to data up to and including date $t-1$.

Harrison and Taylor (2012) show that the log-linearised consumption and labour supply decisions are given by:

$$\begin{aligned} \hat{c}_{i,t} = \tfrac{1-\beta}{1-\psi_g+\frac{\sigma(1-\mu)}{\gamma}} \tilde{E}_{i,t-1} \textstyle\sum_{s=t}^{\infty} \beta^{s-t} & \begin{pmatrix} (1-\mu)\left(1+\gamma^{-1}\right)\hat{w}_s \\ +\beta^{-1}\tilde{b}_{i,t-1} + \left(\mu - \psi_g\right)\hat{d}_s \end{pmatrix} \\ - \tfrac{\beta}{\sigma} \tilde{E}_{i,t-1} \textstyle\sum_{s=t}^{\infty} \beta^{s-t} \left(\hat{R}_s - \hat{\Pi}_{s+1}\right) & \end{aligned} \qquad \textbf{(2)}$$

$$\hat{h}_{i,t} = \gamma^{-1}\hat{w}_t - \sigma\gamma^{-1}\hat{c}_{i,t} \qquad \textbf{(3)}$$

where we use the notation $\hat{x}_t \equiv \ln(x_t/x)$ for each variable $x_t$ to denote its log-deviation from its steady-state value, $x$ and $\tilde{x}_t \equiv (x_t - x)$ to denote the absolute deviation.

Equation **(2)** looks very much like a consumption function: current consumption depends on existing asset holdings plus the expected stream of future net income. So long-horizon

[6] This rules out the case of linear disutility of work in utility ($\gamma = 0$). Given our informational assumptions (to be discussed), this case is problematic because households are unable to forecast their total labour income when their labour supply is demand determined.

expectations matter.[7] The decision rules above therefore determine the household's choice variables as a function of the expected path of the variables outside of their control. In the following period, new shocks will have arrived, expectations will be updated and the household constructs a new consumption plan.

Of course, the Euler equation always describes the optimal relationship between current and future consumption. But to compute the future consumption that it expects to enjoy, the household must factor in forecasts of future net income and real interest rates and ensure that the consumption plan is consistent with the intertemporal budget constraint. In the case of rational expectations, the Euler equation is sufficient to describe the optimal consumption plan because the expectations operator is model consistent: the restrictions of expected consumption from the budget constraint facing the household are taken into account. Evans, Honkapohja and Mitra (2003) note that if households have access to a *subset* of the information required for rational expectations, then **(2)** can be expressed as an Euler equation containing a non-rational forecast of consumption. This information manifest itself in terms of restrictions on the household's forecasting rules for net income and real interest rates.

The labour supply relationship is given by equation **(3)**. Harrison and Taylor (2012) note that this equation is consistent with the assumption that at the beginning of the period, the household forms a consumption plan before splitting into a 'shopper' and a 'worker'. The shopper enters the goods market and purchases the consumption decided upon in the plan. The worker enters the labour market and supplies labour according to the intratemporal optimality condition relating consumption and labour supply, given the market wage.[8]

Following Harrison and Taylor (2012), we assume that there are a finite set of 'predictors' available to each group of agents. We use the index $i$ to denote the decisions of a household using predictor $i \in \{1, ..., I\}$. This is appropriate if all agents that use the same predictors make

[7] The assumption of non-rational expectations is not the only way in which long-horizon expectations may matter. For example, Rotemberg and Woodford (1999) assume that the information delay differs across firms so that *rational* expectations based on information at dates $t-1$ and $t-2$ are relevant for pricing behaviour at date $t$. This assumption means that the conventional representation of the New Keynesian Phillips Curve cannot be uncovered – so long-horizon expectations remain in their aggregate supply curve – see equation 22 on page 65 of Rotemberg and Woodford (1999).

[8] This timing assumption is necessary to ensure that the labour supply responds to meet demand. Given our assumptions about price-setting behaviour – explained in Section 2.2 – if the household sets either the nominal wage or the amount of labour supplied based solely on date $t-1$ information, then the real wage is unable to move to the level required to clear the labour market. So the equation in the text differs from the planned labour supply at the start of the period (given by $\hat{h}^p_{i,t} = \gamma^{-1}\tilde{E}_{i,t-1}\hat{w}_t - \sigma\gamma^{-1}\hat{c}_{i,t}$) because the real wage adjusts to clear the labour market.

the same decisions, which is the case if there is no dependence of current decisions on past decisions. But since households have access to financial assets that can be carried between periods, care is required when aggregating across household decisions. Harrison and Taylor (2012) show that both aggregate consumption and the consumption of households using each predictor can be recovered if one assumes that the group of agents using each predictor at each date is randomly drawn from the population. Our approach to dynamic predictor choice explained in Section 2.4 is consistent with this assumption.

We assume that the mass of households using predictor $i$ at date $t$ is given by $n_{i,t}$ where

$$\sum_{i=1}^{I} n_{i,t} = 1$$

so that aggregate consumption is given by:

$$\begin{aligned} \hat{c}_t &= \textstyle\sum_{i=1}^{I} n_{i,t}\hat{c}_{i,t} \\ &= \frac{1-\beta}{1-\psi_g+\frac{\sigma(1-\mu)}{\gamma}} \textstyle\sum_{i=1}^{I} n_{i,t}\tilde{E}_{i,t-1} \sum_{s=t}^{\infty} \beta^{s-t} \left[(1-\mu)\left(1+\gamma^{-1}\right)\hat{w}_s + \left(\mu-\psi_g\right)\hat{d}_s\right] \\ &- \frac{\beta}{\sigma} \textstyle\sum_{i=1}^{I} n_{i,t}\tilde{E}_{i,t-1} \sum_{s=t}^{\infty} \beta^{s-t}\left(\hat{R}_s - \hat{\Pi}_{s+1}\right) \end{aligned}$$

where we use the fact that bond market clearing requires:

$$\sum_{i=1}^{I} n_{i,s}\tilde{b}_{i,s} = 0$$

for all $s$.

### *2.2 Firms*

We assume that there are two types of firms – retailers and producers – and we consider each of them in turn.

Retailers are perfectly competitive and operate a production technology that combines the inputs of producers using a Dixit-Stiglitz technology:

$$y_t = z_t\left[\int_0^1 x_{j,t}^{1-\mu}\,dj\right]^{\frac{1}{1-\mu}}$$

where $x_j$ is the quantity of output purchased from producer $j$, $z$ is a productivity shock and $\mu > 0$. We assume that the log-linearised behaviour of the shock $z$ is

$$\hat{z}_t = \rho_z\hat{z}_{t-1} + u_t^z \tag{4}$$

where $u^z$ is an iid shock and $\left|\rho_z\right| < 1$.

Retailers sell their output to consumers and the government at nominal price $P$. Denoting the price of output purchased from producer $j$ as $P_j$, the cost-minimising price index for output is

$$P_t = z_t^{-1} \left[ \int_0^1 P_{j,t}^{\frac{\mu-1}{\mu}} dj \right]^{\frac{\mu}{\mu-1}}$$

and the associated demand curve for the output of producer $j$ is:

$$x_{j,t} = z_t^{\frac{1-\mu}{\mu}} \left[ \frac{P_{j,t}}{P_t} \right]^{-\frac{1}{\mu}} y_t$$

The set of producers $j \in (0, 1)$ produce differentiated products that form a Dixit-Stiglitz bundle consumed by households and the government. They produce using a constant returns in the single input (labour):

$$x_{j,t} = a_t h_{j,t}$$

where $a_t$ is a stochastic aggregate productivity term (common to all producers) and we assume that log-linearised productivity follows a simple AR(1) process:

$$\hat{a}_t = \rho_a \hat{a}_{t-1} + u_t^a \qquad \textbf{(5)}$$

where $u^a$ is an iid shock and $|\rho_a| < 1$.

Aggregating the production function across producers, combined with the production function of retailers, and log-linearising gives:[9]

$$\hat{y}_t - \hat{z}_t = \hat{a}_t + \hat{h}_t \qquad \textbf{(6)}$$

The real profit of producer $j$ is:

$$\Delta_{jt} = \frac{P_{jt}}{P_t} x_{jt} - w_t h_{j,t} = \left( \frac{P_{jt}}{P_t} - \frac{w_t}{a_t} \right) z_t^{\frac{1-\mu}{\mu}} \left( \frac{P_{j,t}}{P_t} \right)^{-\frac{1}{\mu}} y_t$$

We assume that price-setting is subject to a Calvo (1983) technology: with probability $0 < \alpha \leq 1$ a producer $j$ is given the chance to reset its price. Such a producer will choose the price to solve:

$$\max \bar{E}_{j,t-1} \sum_{s=t}^{\infty} \Lambda_s \left(\beta\alpha\right)^{s-t} \left( \frac{P_{jt}}{P_s} - \frac{w_s}{a_s} \right) \left( \frac{P_{j,t}}{P_s} \right)^{-\frac{1}{\mu}} z_s^{\frac{1-\mu}{\mu}} y_s$$

where the expectations of producer $j$ is denoted as the (non-rational) expectation $\bar{E}_{j,t}$.

[9]The Dixit-Stiglitz aggregator for output is not equivalent to the simple sum of production functions across firms: there is a wedge between the two measures. Our linearisation makes use of the fact that the distortion is second order so can be ignored when considering a linear approximation to the model. See Christiano, Evans and Eichenbaum (2005) for a discussion.

The first-order condition is:

$$\bar{E}_{j,t-1}\sum_{s=t}^{\infty}\Lambda_s(\beta\alpha)^{s-t}\left(\frac{\mu-1}{\mu}\frac{p_{j,t}}{\Pi_{t,s}}+\frac{1}{\mu}\frac{w_s}{a_s}\right)\left(\frac{p_{j,t}}{\Pi_{t,s}}\right)^{-\frac{1}{\mu}}z_s^{\frac{1-\mu}{\mu}}y_s=0$$

where we define the price set by producer $j$ relative to the *previous* period's aggregate price level as:

$$p_{j,t}\equiv\frac{P_{j,t}}{P_{t-1}}$$

and the relative inflation factor

$$\Pi_{t,s}\equiv\frac{P_s}{P_{t-1}}=\Pi_s\times\Pi_{s-1}\times...\times\Pi_t \text{ for } s\geq t$$

where we normalise by the aggregate price level from the previous period because this is contained in producers' information set.[10]

Harrison and Taylor (2012) show that the linearised pricing equation is:

$$\hat{p}_{j,t}=(1-\beta\alpha)\,\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\left(\hat{w}_s-\hat{a}_s\right)+\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\,\hat{\Pi}_s \qquad \textbf{(7)}$$

which makes it clear that if individual producers have different expectations about future costs and future inflation, then they will set different prices even when free to set them simultaneously.

Analogous to the treatment of households, we use the index $j$ to denote the decisions of a firm using predictor $j\in\{1,...,J\}$ and the mass of firms using predictor $j$ at date $t$ is given by $m_{j,t}$ where

$$\sum_{j=1}^{J}m_{j,t}=1$$

so that the average reset price of producers is

$$\begin{aligned}\hat{p}_t^* &= \textstyle\sum_{j=1}^{J}m_{j,t}\hat{p}_{j,t}\\ &=(1-\beta\alpha)\textstyle\sum_{j=1}^{J}m_{j,t}\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\left(\hat{w}_s-\hat{a}_s\right)+\sum_{j=1}^{J}m_{j,t}\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\,\hat{\Pi}_s\end{aligned}$$

and inflation is given by:

$$\hat{\Pi}_t=(1-\alpha)\,\hat{p}_t^*-\hat{z}_t$$

[10]Conventional treatments usually define the relative price of a firm $j$ in terms of the current aggregate price level: $p_{j,t}\equiv P_{j,t}/P_t$. There is no loss of generality in following our approach.

### *2.3 Government and market clearing*

The government budget constraint is:

$$B_t^g = R_{t-1} B_{t-1}^g + G_t - P_t \tau_t$$

where $B^g$ is nominal government debt (one-period bonds), $R$ is the nominal interest rate, $G (= P \times g)$ is nominal spending and $P \times \tau$ is nominal tax revenue.[11] In real terms:

$$b_t^g = \frac{R_{t-1}}{\Pi_t} b_{t-1}^g + g_t - \tau_t$$

and we assume that the government issues no debt:

$$B_t^g = b_t^g = 0$$

for all periods $t$. This means that the government runs a balanced budget each quarter and government spending is financed by (lump-sum) tax revenue. The log-linearised expression for government spending is:

$$\hat{g}_t = \rho_g \ln \hat{g}_{t-1} + u_t^g \quad \textbf{(8)}$$

where $0 \leq \rho_g \leq 1$ and $u_t^g$ is iid.

Monetary policy is conducted using a Taylor rule with interest rate smoothing and an iid shock, which has the log-linearised representation:

$$\hat{R}_t = (1 - \phi_r) \left( \phi_\pi \hat{\Pi}_{t-1} + \phi_y \hat{y}_{t-1} \right) + \phi_r \hat{R}_{t-1} + u_t^R \quad \textbf{(9)}$$

which implies that the monetary policy maker has a similar information set to private agents and so sets interest rates on the basis of lagged outturns for output and inflation.

Market clearing dictates that all output is consumed by households or government:

$$\hat{y}_t = \left(1 - \psi_g\right) \hat{c}_t + \psi_g \hat{g}_t \quad \textbf{(10)}$$

where $\psi_g$ is a parameter denoting the steady-state share of government expenditure in output.

[11] We consider our model as the 'cashless limit' (Woodford (2003)) of an economy in which households demand fiat money, the issuance of which generates seignoreige for the government. We do so for analytical convenience since the inclusion of money would create an additional choice variable and associated decision rule for households.

## 2.4 Predictors and predictor proportions

In this section we describe the backward-looking predictors used by households and firms of the general form assumed by Harrison and Taylor (2012). We also describe how households and firms choose which predictors to use: this allows us to complete the description of the model by providing an account of how the predictor proportions $n_i$, $(i = 1, ..., I)$ and $m_j$ $(j = 1, ..., J)$ are determined. In the rest of this paper, and without loss of generality, we assume that $I = J = 2$. For notation convenience for households we denote $n_1 = n$ so that $n_2 = 1 - n$. Similarly for firms we use the notation $m_1 = m$ so that $m_2 = 1 - m$.[12]

We assume that the fraction using each predictor evolves according to the observed forecast errors for that predictor. In particular the proportion $n$ at date $t$ is determined by:[13]

$$n_t = \frac{1}{2}\tanh\left[-\frac{\theta'_h}{4}\left(\Xi^h_{1,t} - \Xi^h_{2,t}\right)\right] + \frac{1}{2} \tag{11}$$

where $\theta_h$ is a $4 \times 1$ vector of 'intensities of choice' and $\left(\Xi^h_{1,t} - \Xi^h_{2,t}\right)$ is the $4 \times 1$ vector of the differences of 'fitness measures' for the variables of interest for households ($\hat{w}$, $\hat{d}$, $\hat{R}$ and $\hat{\Pi}$). The fitness measures are defined as geometric averages of past squared errors:

$$\Xi^h_{i,t} = \kappa_h SE_{i,t-1} + (1 - \kappa_h)\,\Xi^h_{i,t-1}\ , i = 1, 2$$

where $\kappa_h > 0$ specifies the extent to which past mean squared errors are discounted.

We similarly assume that the choice of $m$ is governed by:

$$m_t = \frac{1}{2}\tanh\left[-\frac{\theta'_f}{3}\left(\Xi^f_{1,t} - \Xi^f_{2,t}\right)\right] + \frac{1}{2}$$

where $\theta_f$ is a $3 \times 1$ vector of intensities of choice and $\left(\Xi^f_{1,t} - \Xi^f_{2,t}\right)$ is the $3 \times 1$ vector of the fitness measure for $\hat{w}$, $\hat{a}$ and $\hat{\Pi}$. This evolves according to

$$\Xi^f_{j,t} = \kappa_f SE_{j,t-1} + \left(1 - \kappa_f\right)\Xi^f_{j,t-1}\ , j = 1, 2$$

Following Harrison and Taylor (2012), we assume that households and firms use a VAR to form expectations of the variables that they care about. We define the vector of endogenous variables

[12] An alternative approach is to view the model as a representative agent model in which decisions are based on weighted forecasts of future variables relevant to those decisions. In this interpretation, $m$ and $n$ represent the weights that firms and households place on the available predictors.

[13] This formulation is used by Branch (2004) and maps a $4 \times 1$ vector to the unit interval.

as $\xi_t = [\hat{c}_t\ \hat{y}_t\ \hat{\Pi}_t\ \hat{p}^*_t\ \hat{\omega}_t\ \hat{R}_t\ \hat{d}_t\ \hat{a}_t\ \hat{g}_t\ \hat{z}_t]'$ and the vector of (iid) shocks plus a constant as $\zeta_t = [u^a_t\ u^g_t\ u^R_t\ u^z_t\ 1]'$. Now we assume that firms' and households' expectations of endogenous variables are generated by a general VAR forecasting model of the following form:

$$\xi_t = F\xi_{t-1} + G\zeta_t$$

This form permits a wide variety of forecasting models. For example, if we assume that the fourth column of $G$ is equal to a zero vector and specify $F$ so that the largest eigenvalue is less than 1 in magnitude, then expectations will ultimately converge to the steady state. Setting both the $F$ and $G$ matrices to zero generates the 'steady-state predictors' analysed in Brazier *et al* (2008).

This approach means that, for firms, we have:

$$\bar{E}_{j,t}\xi_{t+s} = F^s_{f,j}\left(F_{f,j}\xi_{t-1} + G_{f,j}\zeta_t\right) \quad \text{for } j = 1, ..., J$$

and for households:

$$\tilde{E}_{i,t}\xi_{t+s} = F^s_{h,i}\left(F_{h,i}\xi_{t-1} + G_{h,i}\zeta_t\right)$$

The current period decisions for consumption and prices depend on discounted sums of expected future outturns. We can transform our VAR forecasting model to perform these summations. So for example, for some arbitrary $F$, $G$ and discount rate $\delta \in (0, 1)$ we have:

$$\begin{aligned}\sum_{s=0}^{\infty}\delta^s\tilde{E}_{t-1}\xi_{t+s} &= \left(\sum_{s=0}^{\infty}\delta^s F^s\right)\left(F\xi_{t-1} + G\zeta_t\right) \\ &= (I - \delta F)^{-1}\left(F\xi_{t-1} + G\zeta_t\right)\end{aligned}$$

which is valid as long as the eigenvalues of $F$ are all less than $\delta^{-1}$ in absolute magnitude.

Two equations in the model contain terms with expectation operators: firms' price equation and households' consumption equation. Starting with firms, we have:

$$\begin{aligned}
\hat{p}_t^* &= \sum_{j=1}^{2} m_{j,t}\hat{p}_{j,t} \\
&= \sum_{j=1}^{2} m_{j,t}\left((1-\beta\alpha)\,\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\left(\hat{\omega}_s-\hat{a}_s\right)+\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\,\hat{\Pi}_s\right) \\
&= \sum_{j=1}^{2} m_{j,t}\left(\begin{array}{c} (1-\beta\alpha)\,\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\,\hat{\omega}_s-(1-\beta\alpha)\,\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\,\hat{a}_s \\ +\bar{E}_{j,t-1}\sum_{s=t}^{\infty}(\beta\alpha)^{s-t}\,\hat{\Pi}_s \end{array}\right)
\end{aligned}$$

Now we can use the VAR representation of expectations and perform the summations as above using the VAR coefficient matrices. We use selector matrices – denoted $S^{\cdot}$ – to pick out the summation of the forecast for the variable of interest: thus $S^x$ is the matrix that selects the forecast error for the variable $x$. So the pricing equation can be written as:

$$\hat{p}_t^* = \sum_{j=1}^{2} m_{j,t}\left[\begin{array}{c} (1-\beta\alpha)\,S^{\omega}(I-\beta\alpha F_{f,j})^{-1}\left(F_{f,j}x_{t-1}+G_{f,j}z_t\right) \\ -(1-\beta\alpha)\,S^{a}(I-\beta\alpha F_{f,j})^{-1}\left(F_{f,j}x_{t-1}+G_{f,j}z_t\right) \\ +S^{\pi}\left(I-\beta\alpha F_{f,j}\right)^{-1}\left(F_{f,j}x_{t-1}+G_{f,j}z_t\right) \end{array}\right]$$

Collecting terms, and defining $V_{f,j} \equiv (I-\beta\alpha F_{f,j})^{-1}$ gives:

$$\hat{p}_t^* = \sum_{j=1}^{2} m_{j,t}\left[\left((1-\beta\alpha)\left(S^{\omega}-S^{a}\right)+S^{\pi}\right)V_{f,j}\left(F_{f,j}x_{t-1}+G_{f,j}z_t\right)\right]$$

Analogous arguments can be applied to the consumption equation of households giving:

$$\begin{aligned}
\hat{c}_t &= \sum_{i=1}^{2} n_{i,t}\hat{c}_{i,t} \\
&= \sum_{i=1}^{2} n_{i,t}\left[\left(k_1 S^{\omega}+k_2 S^{d}-\frac{\beta}{\sigma}S^{R}+\frac{\beta}{\sigma}S^{\pi}F_{h,i}\right)V_{h,i}\left(F_{h,i}x_{t-1}+G_{h,i}z_t\right)\right]
\end{aligned}$$

where, $V_{h,i} \equiv \left(I-\beta F_{h,i}\right)^{-1}$ and

$$k_1 = \frac{(1-\beta)(1+\gamma)}{(1-\psi_g)\frac{\gamma}{1-\mu}+\sigma};\; k_2 = \frac{(1-\beta)\left(\mu-\psi_g\right)}{\left(1-\psi_g\right)+(1-\mu)\frac{\sigma}{\gamma}}$$

### *2.5 Model parameters*

The parameters can be divided into two groups. First, the so-called 'deep' parameters of the model describe the preferences and constraints of households and firms and must be specified in order to solve the model under rational expectations. Second, the parameters that are specific to the non-rational expectations version of the model. We consider each group in turn.

*2.5.1 'Deep' parameters*

The deep parameters in the model are calibrated following Harrison and Taylor (2012), based on the assumption that each period corresponds to one quarter. The table below documents the choice of these parameters. Harrison and Taylor (2012) discuss the motivation for the chosen values. As explained by Harrison and Taylor (2012) , the interest rate smoothing parameter is set to a relatively low value to ensure that the model does not impose a high degree of persistence by default. This allows for the possibility that expectations formation may become an important determinant of the persistence in the model. The same reasoning leads us to calibrate the persistence of the shock processes to be relatively low.

| | | | |
|---|---|---|---|
| $\alpha$ | 0.75 | $\beta$ | 0.99 |
| $\gamma$ | 0.50 | $\sigma_z$ | 0.50 |
| $\sigma$ | 2.26 | $\sigma_g$ | 2.14 |
| $\frac{1}{\mu}$ | 10.0 | $\sigma_r$ | 0.97 |
| $\psi_g$ | 0.22 | $\sigma_a$ | 1.79 |
| $\phi_\pi$ | 1.50 | $\rho_g$ | 0.30 |
| $\phi_y$ | 0.125 | $\rho_z$ | 0.10 |
| $\phi_r$ | 0.25 | $\rho_a$ | 0.60 |

*2.5.2 'Expectations' parameters*

The expectations parameters are those that govern the predictor choice mechanism. The 'intensities of choice' determine the extent to which agents switch to better-performing predictors. So, for households, the proportion of households choosing predictor 1 (of the two available) is given by equation **(11)** introduced in Section 2.4. The vector $\theta_h$ represents the intensities of choice for the household. If each element of $\theta_h$ tends to infinity, then all agents immediately switch to the best-performing predictor (as measured by the average difference in mean squared error). Conversely, if each element of $\theta_h$ tends to zero, agents are indifferent between the two predictors, regardless of their relative performance so that $n = 1 - n = 0.5$. The mapping described by this function can be motivated by appealing to a story in which agents observe the true relative predictor performance with noise. The extreme of infinitely large $\theta_h$ elements corresponds to the case in which the variance of noise is zero and the extreme of $\theta_h = 0$ corresponds to a situation in which the variance of the noise is infinitely large. As described in Section 2.4, there is an analogous function that governs the choice of predictors by firms and the intensities of choice in this function ($\theta_f$) are calibrated jointly with those for households.

Because the intensities of choice affect the choice of predictors by households and firms, the choice of these parameters is clearly important for the behaviour of the model. Our approach to calibrating them is based on the view that we want the elements of the $\theta_h$ and $\theta_f$ to compensate for the fact that some variables that agents forecast may be far more volatile than others. For example, the household must forecast dividend payments, which could be significantly more volatile than other variables that the household forecasts such as inflation. Our strategy proceeds in two steps. First, we set the elements of $\theta \equiv \{\theta_h, \theta_f\}$ so that each element of an 'average fitness measure' vector is equalised

$$\begin{bmatrix} n\,\Xi^h_{1,t} + (1-n)\,\Xi^h_{2,t} \\ m\,\Xi^f_{1,t} + (1-m)\,\Xi^f_{2,t} \end{bmatrix} = \iota \qquad \textbf{(12)}$$

where $\iota$ is the unit vector. Appendix A describes the details of the algorithm, which uses an asymptotic representation of the model in which the predictor proportions $n$ and $m$ are determined according to the predictor choice equations. Because the first step is based on an asymptotic representation of the model (in which the fitness measure coincides with the mean squared error), it can only pin down the relative size of the elements in $\theta$. So the second step is to rescale the entire $\theta$ vector to ensure that there is sufficient predictor switching in dynamic simulations of the model. While this step is somewhat judgemental, we argue in Section 3.1.1 that our calibration ensures that dynamic predictor choice is governed by relative forecast performance rather than the properties of the predictor choice mechanism (eg equation **(11)**) *per se*.

Our calibration procedure is clearly specific to the particular application, because of the dependence of equation **(12)** on the fitness measures. For the experiments discussed in

Section 3.1.2 below, our procedure leads to the following settings for the intensities of choice:

| | $\theta_h$ | $\theta_f$ |
|---|---|---|
| $\Pi$ | 7.4 | 7.4 |
| $w$ | 0.1 | 0.1 |
| $R$ | 1.7 | – |
| $d$ | 0.002 | – |
| $a$ | – | 0.9 |

which shows that, as expected, the forecast errors on relatively volatile series (for example dividends) are markedly lower than the weights placed on less variable series.

The other parameters that are important for the dynamic behaviour of the model are those that control the extent to which agents discount past mean squared errors when computing the fitness measure. We set these parameters to $\kappa_h = 0.05$ and $\kappa_f = 0.1$, based on gain parameters estimated using survey measures of inflation expectations. The estimation procedure is detailed in Appendix B.

## 3 Self-reinforcing misperceptions

The idea that changes in beliefs may become self-reinforcing is one that is of obvious concern to policymakers. In particular, policymakers are alert to the possibility that particularly elevated or depressed inflation expectations could become embedded in the wage-setting and price-setting processes. This is illustrated by the quote from the Bank of England's February 2008 *Inflation Report*, cited in the Introduction, which recognises the potential for beliefs about the economy (expectations of persistently high inflation) to lead to changes in behaviour (higher wage demands by households and increases in the prices charged by firms) that reinforce that belief (higher medium-term inflationary pressure).

This type of argument motivates our analysis in this section. We explore the extent to which, if agents begin to use predictors that embed assumptions of greater macroeconomic persistence, these beliefs may become self-reinforcing. And we study whether this type of mechanism can materially affect the behaviour of our model economy, and generate time-variation in the moments of the data generated by the model. In particular, we explore the interaction between the beliefs of firms and households and the properties of our model. We show that fluctuations in the expectations formation process – the dynamic predictor selection that occurs endogenously in response to shocks – can bring about changes in the properties of the model that act (at least temporarily) to reinforce the initial change in beliefs. Misperceptions can for a time be self-reinforcing; and this mechanism is a candidate explanation for time-variation in the moments of data.

Our approach is to investigate the behaviour of our model when agents have access to a predictor that (wrongly) anticipates that inflation shocks will be long-lasting. We start by describing how we have implemented misperceptions before studying their effect on the model. We then contrast the results with the case in which agents have access to a constant-gain learning algorithm.

### *3.1 Arbitrary misperceptions*

To place some structure on the beliefs that agents in our model hold and in line with our focus on non-rational expectations, we limit all agents in the model to forecast using AR(1) forecasting rules. This type of forecasting rule is simple to compute and use and will generate reasonable forecasts in many models. For these reasons, AR(1) forecasting rules are often used as a benchmark in studies of non-rational expectations.[14]

We assume that agents choose between two alternative forecasting rules. The first forecasting rule is one that brings about a form of 'restricted perceptions equilibrium'.[15] We call this forecasting rule the 'restricted perceptions predictor' (RPP). When all agents use the RPP, the data generating process of the model is such that, subject to their restricted perceptions, agents' beliefs are confirmed. That is, the AR(1) coefficients of the RPP are optimal in the sense that they provide the best AR(1) forecasts when all agents use the RPP to forecast. Computing the parameters of the RPP involves solving for the fixed point between agents' forecast rule coefficients and the AR(1) coefficients implied by the model.[16]

The second forecasting rule is one that embodies the misperception that cost-push shocks are

[14] See, for example, Branch and McGough (2006b).

[15] See, for example, Branch (2004).

[16] We use a simple function iteration scheme to compute the fixed point (see, for example, Fackler and Miranda (2002)). Loosely, we guess values for the AR(1) coefficients, solve for the VAR representation of the model and compute the asymptotic AR(1) coefficients for each variable. We then refine our guess of the forecast rule coefficients until it coincides with the asymptotic coefficients generated by the model.

more persistent than in fact they are. We call this forecasting rule the 'misperceptions predictor' (MP). The MP has the following properties. As before, agents use AR(1) models to forecast each variable. But in this case, the AR(1) coefficients are computed so that they are consistent with a version of our model in which cost-push shocks are more persistent than in the actual model. We compute the AR(1) coefficients of the MP as follows. We first recalibrate the model so that the cost-push shock is much more persistent than in the true model (we set $\rho_z = 0.9$ in place of $\rho_z = 0.1$). We then assume that agents forecast each variable as an AR(1) process and find the coefficients of these processes that are consistent with the behaviour of the economy when agents use those forecasting rules. So the misperceptions predictor corresponds to the restricted perception predictor of a *misspecified* model.

**Table 1: AR(1) coefficients describing RPP and MP forecasting rules**

| | RPP | MP | memo: RE |
|---|---|---|---|
| $z$ | 0.10 | 0.80 | 0.10 |
| $\Pi$ | 0.28 | 0.90 | 0.07 |
| $R$ | 0.32 | 0.86 | 0.25 |
| $c$ | 0.28 | 0.79 | 0.44 |
| $y$ | 0.29 | 0.58 | 0.31 |
| $d$ | 0.55 | 0.63 | 0.35 |
| $w$ | 0.47 | 0.62 | 0.04 |
| $a$ | 0.60 | 0.60 | 0.60 |
| $g$ | 0.30 | 0.30 | 0.30 |

This approach allows us to capture the notion that households and firms are unsure about the persistence of the variables on which their decisions depend. But it also imposes a certain consistency on their forecasting rules, by making the MP consistent with the restricted perceptions equilibrium of a misspecified model. By focusing on AR(1) predictors, we can also easily compare the two forecasting rules. The alternative beliefs are reported in Table 1, which lists the AR(1) coefficients that characterise the RPP and MP forecasting rules. We can see that the MP attributes higher persistence to all variables. The third column of Table 1 reports the asymptotic AR(1) coefficients generated by the rational expectations version of the model. In general, these coefficients are quite close to the coefficients of the RPP, which indicates that this predictor is likely to generate reasonable forecasts when used by all agents.

### *3.1.1 Implications for predictor choice*

Before examining the simulation properties of the model, we consider the behaviour of the fractions of households and firms using the RPP (denoted $n$ and $m$ respectively) in a particular limiting case. Specifically, we analyse the asymptotic behaviour of the model under the assumption that households and firms use the entire history of data generated by the model to evaluate the two predictors. This corresponds to the case in which the gain parameters are set to $\kappa_h = \kappa_f = t^{-1}$ and $t \to \infty$, where $t$ is the sample size of data. Such analysis is useful for examining the equilibrium properties of models with predictor choice, as shown by Branch (2004), and our analysis is inspired by the 'T-maps' considered in that literature.

Our approach is as follows. We first compute the optimal fractions of households using the RPP, denoted $n^*$, as the fraction of firms using the RPP ($m$) varies. This function traces out the fixed points of the following mapping

$$n^* = T_h(m, n^*)$$

as $m$ varies. At each value of $m$ we solve for the value of $n^*$ generated by the predictor choice equation **(11)** evaluated using the asymptotic moments of the model when the fractions of households and firms using the RPP are held at $n^*$ and $m$ respectively. An analogous procedure is applied to the proportion of firms using the RPP and traces out the fixed points of the following mapping

$$m^* = T_f(m^*, n)$$

as $n$ varies.

The resulting functions for $n^*$ and $m^*$ are plotted in Chart 1. The blue line plots $n^*$ as a function of $m$ and the red line plots $m^*$ as a function of $n$. We see from the blue line that as the fraction of firms using the RPP increases, the optimal fraction of households using the RPP also increases. But the red line shows that as the fraction of households using the RPP increases, the optimal proportion of firms using the RPP *falls*. This is driven by the behaviour of the real wage in the model: as we move along the red line, a higher fraction of households using the RPP is associated with greater real wage persistence, which is sufficient to attract some firms to the MP. Because the real wage moves to clear the labour market (as explained in Section 2.1, page 10), the increase in real wage persistence is driven by the interaction of labour demand and consumption behaviour.[17] As more households use the RPP, their consumption decisions are based on forecasts of wealth that are less persistent. Consumption becomes less persistent and more variable. This fall in persistence and increase in volatility carries over to labour demand. The net effect on the properties of the real wage depend on the joint properties of consumption and labour demand. For our parameterisation, real wage persistence increases.

**Chart 1: Asymptotic fixed point mappings for predictor choice proportions**

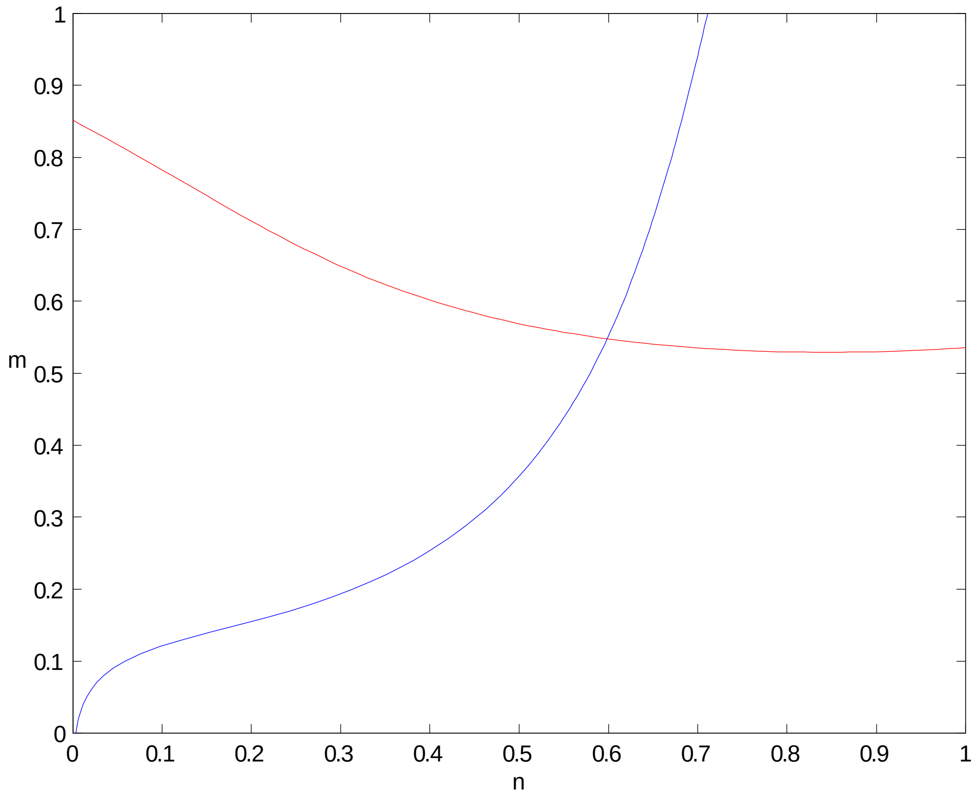


The two lines intersect once and the intersection point – which we label $(n^*, m^*)$ – can be thought of as the long-run equilibrium position of the model if agents were permitted to

[17] From equation **(3)**, we see that the properties of the real wage depend upon

$$\sigma c_t + \gamma h_t$$

accumulate an infinitely long history of data.[18] This analysis confirms that the RPP can be considered the 'preferred' predictor in the sense that most households and firms would choose to use it in the long run. We would also expect the predictor proportions to spend most of the time in the upper right corner of the chart when the model is simulated. Of course, the actual distribution of predictor choices will be determined by the dynamics of the model over finite samples, which we examine in the next subsection.

Finally, Chart 1 provides a cross-check on our procedure for calibrating the parameters governing the 'intensity of choice'. Recall from the discussion in Section 2.5.2 that we can think of the predictor choice mechanism as capturing the notion that agents observe noisy measures of relative predictor performance. Because the RPP is – by construction – the best predictor in the class of AR(1) forecasting models, we know that if agents observed predictor performance without noise, then $n^* = m^* = 1$. And if observations of relative forecasting performance were entirely dominated by noise, then the two predictors would appear to forecast equally well and $n^* = m^* = 0.5$. Because our calibration implies that the proportions of households and firms using the RPP is closer to the $n^* = m^* = 1$ case, we can be reasonably confident that the dynamics of predictor choice are driven primarily by signals about relative forecast performance rather than observation noise.

### *3.1.2 Simulation results*

We run 500 simulations of 4,000 periods each, using the model calibration detailed in Section 2.5. We allow for dynamic predictor selection, so that agents choose between the RPP and MP forecasting rules based on their recent forecast performance as described in Section 2.4. We discard the first 2,000 periods from each simulation and examine the behaviour of the model over the remaining periods.

We start by examining the nature of dynamic predictor selection in the model. Chart 2 summarises several aspects of predictor selection behaviour. The solid lines represent the contours of the joint empirical density of $m$ and $n$. We see that the bulk of the mass of the distribution is in the upper right corner of the chart and indeed lies to the North East of the point

[18] The point at which the lines cross is stable, in the sense that $(n_i, m_i)$ sequences generated by the recursions $n_i = T_h(m_{i-1}, n_i)$, $m_i = T_f(m_i, n_i)$ (or $m_i = T_f(m_i, n_{i-1})$, $n_i = T_h(m_i, n_i)$) will converge to $(n^*, m^*)$.

$(n^*, m^*)$ which is represented by the red star symbol. However, there is a cluster of observations in the bottom left corner of the chart, representing periods during which households and firms both use the MP predictor.

The arrows represent a 'vector field' based on the modal dynamics of the model. Specifically, the arrows are computed by first collecting all observations of $(n, m)$ pairs in a small interval and recording the $(n, m)$ position of each pair after 10 periods. These data are then used to compute the most likely direction and distance travelled by an $(n, m)$ pair over the following 10 periods. We can see that there is a general tendency for the predictor proportions to return to the top right corner of the chart. And in general $(n, m)$ pairs that are further away from the bulk of the mass of the empirical distribution are attracted back to it more rapidly over the following 10 periods. An exception to this behaviour is the bottom left portion of the chart, where there appears to be some attraction to the bottom left corner (where households and firms all use the MP predictor).

Chart 3 sheds some light on this issue by recording the actual positions of a group of randomly selected $(n, m)$ pairs from our simulations over time. The top left panel shows that we have collected a random sample of $(n, m)$ observations in small intervals around the points $(0.8, 0.8)$ and $(0.01, 0.01)$ These samples are represented by the collections of grey circles and black squares respectively. The top left panel shows that, by construction, in the initial period $(t = 0)$ all of the observations are collected in small intervals around the starting positions. The top middle panel shows a scatter plot of the $(n, m)$ pairs after one period $(t = 1)$ and the remaining panels depict scatter plots after $t = 5, 10, 50, 100$ periods. We can see that even though the $(n, m)$ pairs that start at $(0.01, 0.01)$ are well-dispersed after 10 periods, most of them remain in the bottom left corner (bottom left panel). Because our vector field is based on the modal positions of $(n, m)$ pairs over time, it detects a weak movement towards the bottom left corner for $(n, m)$ pairs in this region.[19] Our scatter plots also reveal that, in general, the predictor choices of firms change more rapidly than households over the short term (consistent with the fact that firms use a higher gain parameter in the evaluation of the recent forecast performance of the two predictors). This suggests that the interaction of household and firm predictor choice could be important in determining macroeconomic dynamics.

[19] We have verified that redrawing our vector field on the basis of the *mean* direction shows that all arrows point towards the top right corner of the chart.

**Chart 2: Contour chart for predictor choice distribution and ten-period modal vector field**

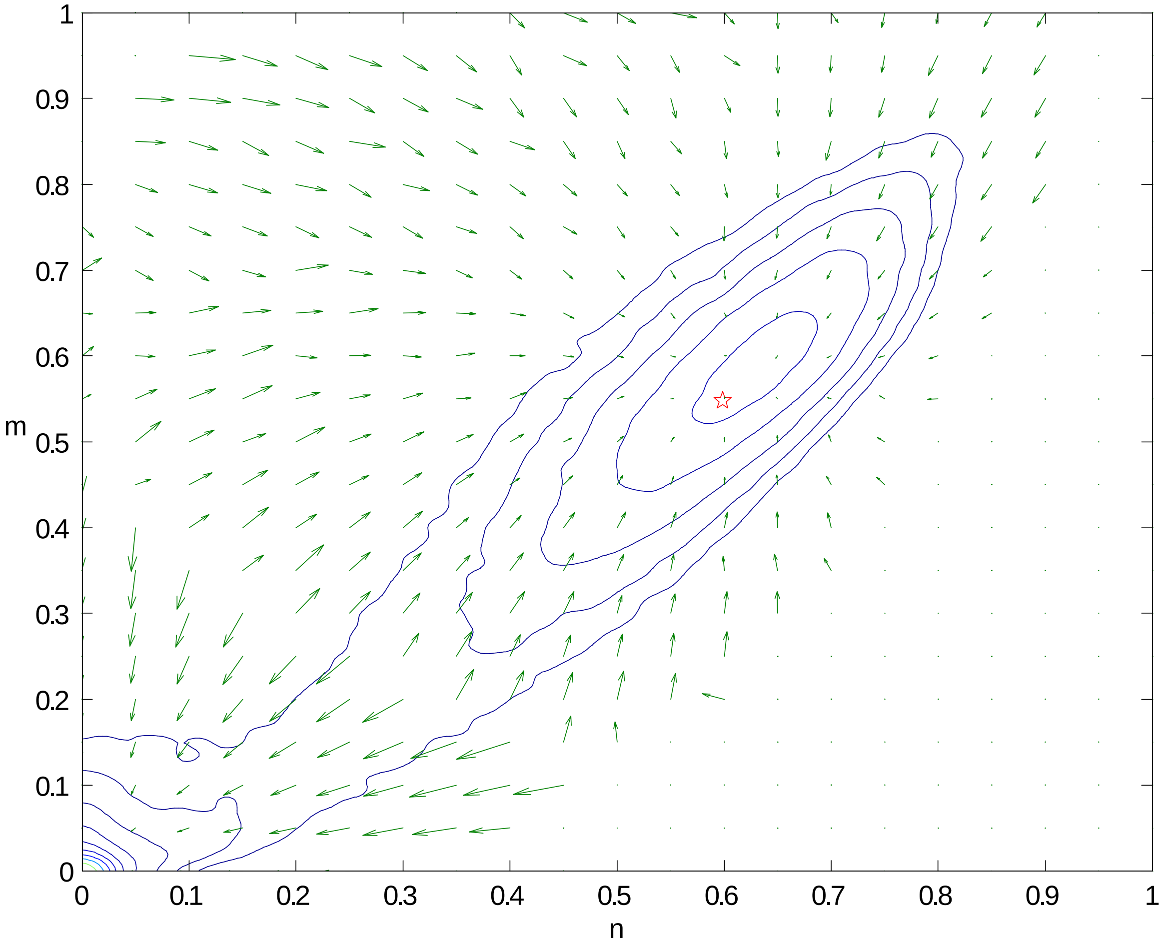


We now examine the time-series properties of the endogenous variables in the model. First, we study the distribution of variance and persistence estimates computed on non-overlapping samples of length 40 periods. We focus on inflation, interest rates and output. We compare these statistics to the equivalent statistics derived from simulations in which agents share common beliefs: in one case rational expectations (RE), and in another the benchmark restricted perceptions predictor (RPP). This illustrates that when agents can choose between the two forecasting rules, the model generates time-variation in the moments of the data.

Table 2 reports summary statistics for the distribution of variance estimates (the 10th, 50th and 90th percentiles). The left panel of Chart 4 plots a kernel density estimate of the distribution of variance estimates for inflation. Similarly, Table 3 and the right panel of Chart 4 characterise the distribution of persistence estimates. The simple measure of persistence used here is the

## Chart 3: Scatter plots of predictor choice pairs over time

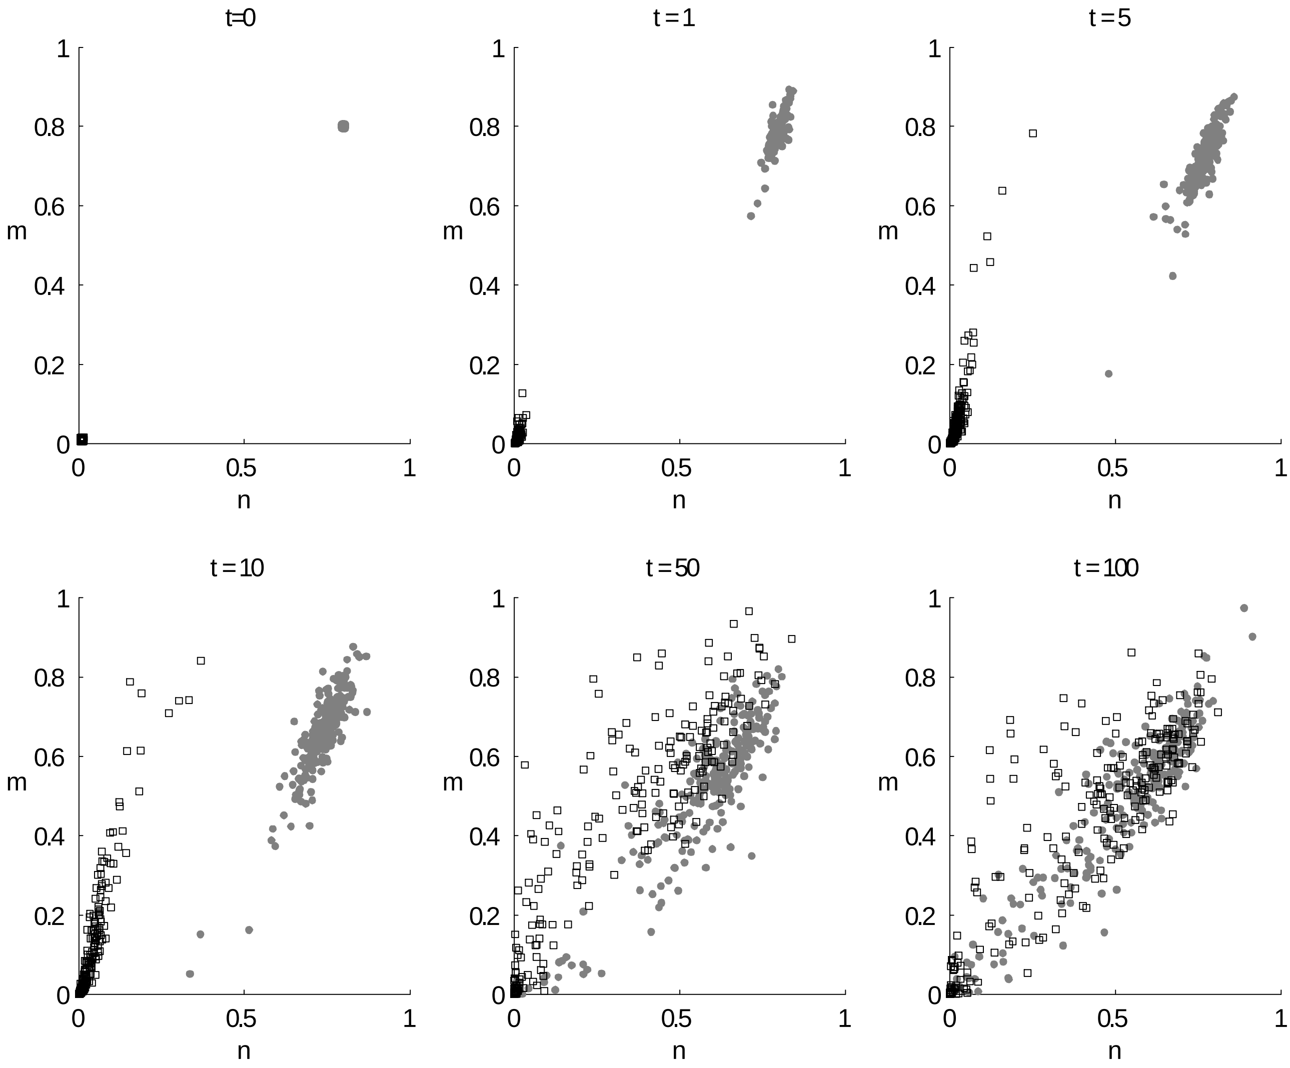

**Chart 4: Distributions of inflation variance and persistence for RE model (black), RPP model (green), MP/RPP switching model (blue) and learning/RPP switching model (red)**

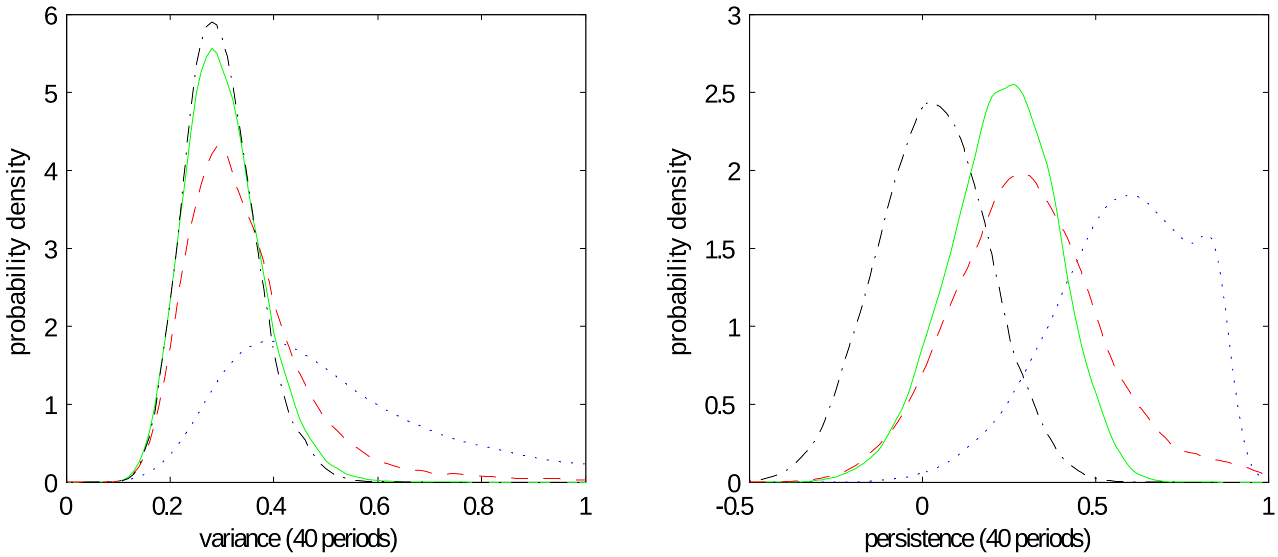


first-order serial correlation coefficient (allowing for a constant term in the regression).

For inflation and interest rates, the median estimates of variance and persistence are markedly higher in the misperceptions case than when expectations are rational. Output is more variable but no more persistent. Moreover, the distributions are positively skewed, indicating the model's ability to generate periods of particularly volatile or persistent outcomes. This is a direct consequence of the non-linearity that the predictor switching mechanism entails.

This can also be seen in Chart 5 which plots the joint density of our small sample estimates of the variance and persistence of inflation. The left panel plots the joint distribution for simulations in which agents switch between the RPP and the MP predictor. The right hand panel shows the joint distribution from simulations (using the same sequences of shocks) for a rational expectations version of the model. We see that, when agents switch between forecasting rules, there is a strong positive correlation between inflation persistence and the variance of inflation. This suggests that episodes such as the 'Great Inflation' experienced by many economies in the 1970s can be generated by the predictor switching in our model.

**Chart 5: Joint distributions of variance of inflation and persistence of inflation for switching model (left) and RE model (right)**

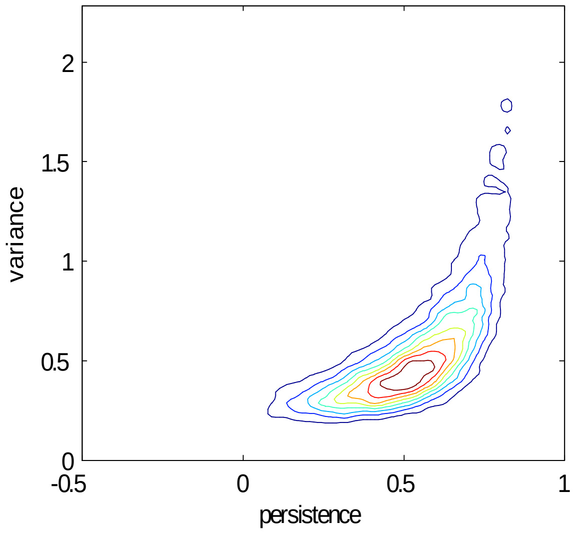


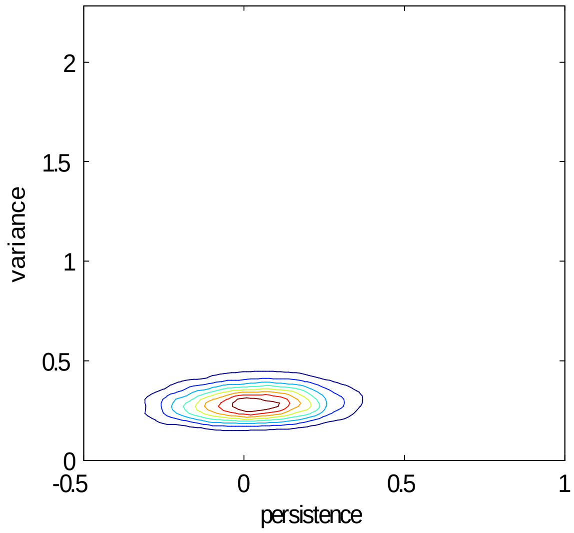


**Table 2: Summary statistics for variances**

| | RE | | | RPP | | | Switching: RPP/MP | | | Switching: RPP/learning | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | 10 | 50 | 90 | 10 | 50 | 90 | 10 | 50 | 90 | 10 | 50 | 90 |
| $\Pi$ | 0.21 | 0.29 | 0.39 | 0.21 | 0.30 | 0.40 | 0.31 | 0.55 | 2.41 | 0.22 | 0.33 | 0.54 |
| $r$ | 0.98 | 1.35 | 1.82 | 1.01 | 1.41 | 1.93 | 1.33 | 2.14 | 6.49 | 1.05 | 1.51 | 2.24 |
| $y$ | 0.24 | 0.34 | 0.46 | 0.19 | 0.27 | 0.37 | 1.02 | 1.95 | 8.80 | 0.21 | 0.30 | 0.50 |

The top panel of Chart 6 shows an illustrative run of simulated data, comparing in the blue and green lines respectively the paths for inflation when all agents' beliefs coincide with the RPP, and when agents are allowed to switch between the RPP and the MP. At the start of the period the paths for inflation are little different under the alternative assumptions about beliefs. The second panel of Chart 6 plots the predictor choice distribution of households and firms. For the purposes of this chart, we plot $1 - n$ and $1 - m$ on the y-axis (solid lines depict $1 - n$ and dashed lines plot $1 - m$). So, for example, the dashed green line shows that, at the start of the simulation, around half of firms are using the *RPP*. After about 40 periods a succession of large negative draws of

**Chart 6: Inflation behaviour under alternative assumptions about beliefs: top panel is annual inflation; middle panel, fraction of agents using misperceptions/learning predictor; bottom panel, estimated inflation persistence in learning predictor**

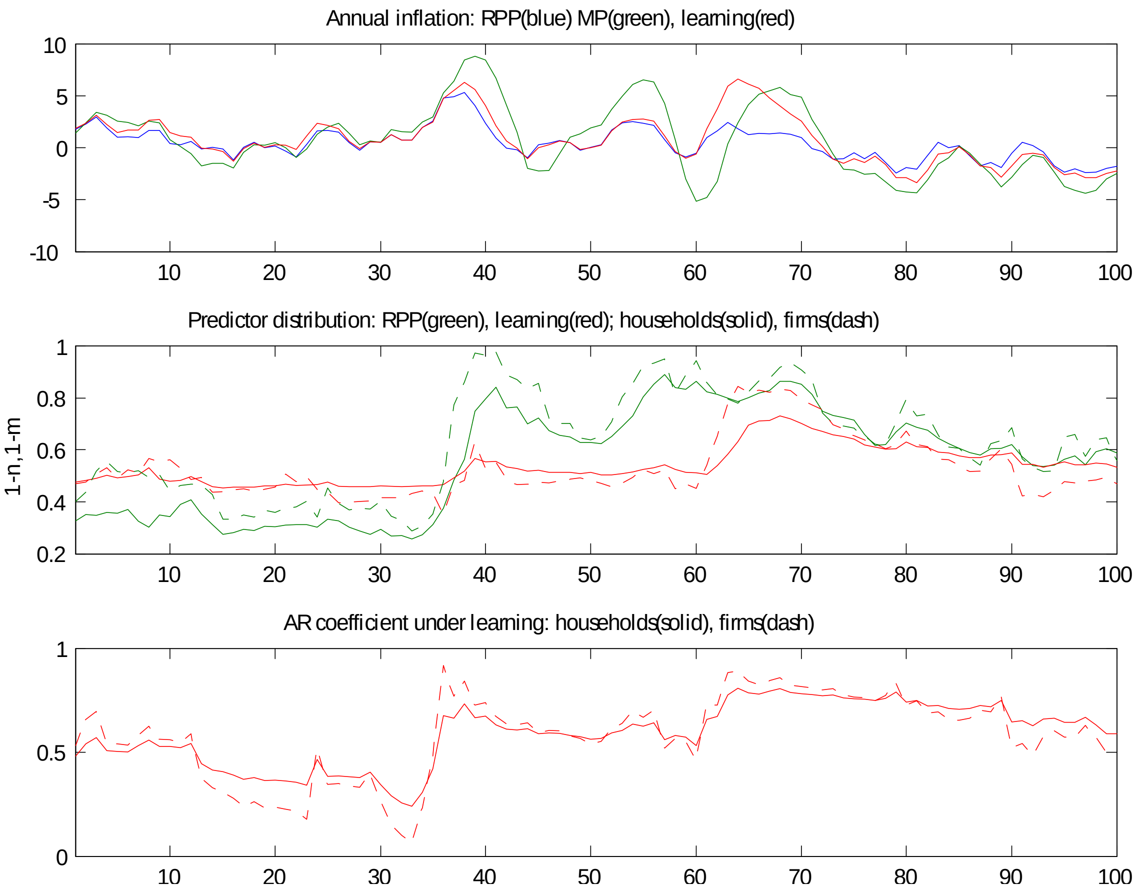


the productivity shock begins to push inflation up. At this time the behaviour of inflation is more accurately forecast by the MP, which embodies the misperceived belief that the cost-push shock, and hence inflation, is highly persistent. Large forecast errors for the RPP cause households and firms to switch to the MP: the green solid and dashed lines increase to unity. With more persistent beliefs about the inflation process, lagged inflation gains higher weight in the price-setting process, and actual inflation is also more persistent. In this sense, the shift towards the misperceived beliefs is self-reinforcing. The simulation demonstrates, again, that in our model firms react more quickly than households in changing their beliefs.

### *3.2 Misperceptions through constant-gain learning*

Our experiments so far have allowed for agents' beliefs to be consistent with an arbitrary misperception. Although we attempt to specify our misperceptions predictor in a sensible way, it remains an arbitrary approach. And there are plausible mechanisms through which such misperceptions may arise endogenously in response to realisations of data. One example that has received much attention is adaptive learning (see Evans and Honkapohja (2001)). In this section we explore whether agents in our model choose to use a constant-gain learning process to inform their expectations, when they have a choice between learning and the RPP that we have described already. We find that at times they do, and that this has qualitatively similar but quantitatively smaller effects on the time-series properties of the model economy compared to the arbitrary misperceptions case we described earlier.

We repeat exactly the same set of experiments as in the previous section, now allowing agents to switch between the RPP and a forecasting rule informed by a constant-gain estimation process. The new forecasting rule is set up as follows. We suppose that agents continue to forecast using AR(1) models for each variable. However, they update their AR(1) parameter estimates using a constant-gain least squares procedure. The gain parameters for households and firms are set to the same values as the parameters $\kappa_h$ and $\kappa_f$ that are used to define how households and firms discount past forecast errors when comparing predictors. As is common in the literature on constant-gain learning we make use of a 'projection facility' to ensure that the model is stable.[20] Our setup implies that the RPP dominates the 'adaptive' forecast rule (based on constant-gain estimation) in long samples (since the RPP is generated using asymptotic moments). So, as in the case of the MP predictor considered earlier, the adaptive forecast rule is in some sense dominated by the RPP.

The results are detailed in the fourth columns of Tables 2 and 3. Comparing them to the benchmark case where all households and all firms use the RPP, we note that there are small increases in the variability of 40-period estimates of persistence and volatility which are most noticeable in the upper tail of the distribution of inflation persistence estimates. This is clear from the red line in the right panel of Chart 4.

[20] For an introduction to constant-gain learning and projection facilities, see, for example, Carceles-Poveda and Giannitsarou (2007).

For much of the time, the learning rule generates beliefs that are similar to the RPP. This is an intuitive result. Suppose the economy is behaving as if agents use the RPP. Then the learning rule is close to the RPP and agents are indifferent between it and the RPP. In this case, expectations generated by both forecasting rules will be very similar, so data outturns will validate both the RPP and the learning rule. Returning to the top panel of Chart 6, this is the situation at the start of the illustrative simulation. But the remainder of the simulation shows that large shocks can cause significant oscillations in inflation relative to the baseline simulation in which all agents use the RPP. Again, this result is driven by the predictor switching in the model (see the red lines in middle panel of Chart 6).

The third panel of Chart 6 shows the AR(1) coefficient for inflation estimated under the learning processes used by firms and households. As inflation rises, so does its estimated persistence. The performance of the learning and RPP predictors diverge, and firms and households switch to using the learning predictor (as illustrated in the second panel of the chart). This tends to increase the persistence of inflation; and so on.

In general, learning mitigates the fluctuations in inflation that occur in the case where the alternative to the RPP is the (fixed) MP. However, at times learning can result in beliefs that deviate from the RPP sufficiently to noticeably alter the behaviour of the economy. So we cannot rule out that endogenously generated fluctuations in beliefs can be self-reinforcing, and generate time-variation in economic dynamics.

**Table 3: Summary statistics for persistence**

| | RE | | | RPP | | | Switching: RPP/MP | | | Switching: RPP/learning | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | 10 | 50 | 90 | 10 | 50 | 90 | 10 | 50 | 90 | 10 | 50 | 90 |
| $\Pi$ | -0.18 | 0.03 | 0.23 | 0.03 | 0.24 | 0.42 | 0.31 | 0.59 | 0.83 | 0.04 | 0.29 | 0.58 |
| $R$ | 0.01 | 0.21 | 0.40 | 0.08 | 0.29 | 0.46 | 0.27 | 0.50 | 0.78 | 0.10 | 0.32 | 0.53 |
| $y$ | 0.06 | 0.27 | 0.46 | 0.05 | 0.25 | 0.44 | -0.01 | 0.23 | 0.52 | 0.06 | 0.27 | 0.46 |

## 4 Conclusions

This paper explores the dynamic feedback between outcomes and expectations in a benchmark New Keynesian model with long-horizon expectations and dynamic predictor selection. The feedback between outcomes and expectations is generated by dynamic predictor selection: the notion that agents choose between a small set of forecasting rules (or 'predictors') based on noisy observations of past performance.

We explore the interaction between the beliefs of firms and households and the properties of our model economy. We show that fluctuations in the expectations formation process – the dynamic predictor selection which occurs endogenously in response to shocks – can bring about changes in the properties of the model that act (at least temporarily) to reinforce the initial change in beliefs. Misperceptions can for a time be self-reinforcing; and this mechanism is a candidate explanation for time-variation in the moments of data. Our approach is to investigate the behaviour of our model when agents have access to a predictor that (wrongly) anticipates that shocks to inflation will be long-lasting. We show that misperceptions can indeed be temporarily self-reinforcing, leading to marked time-variation in the time-series properties of the data. Most importantly, this result survives even when agents' misperceptions are generated by a constant-gain learning algorithm.

A final point to note from our analysis is the finding that heterogeneity between the expectations formation processes of households and firms may be important for macroeconomic dynamics. Our model is calibrated on the assumption that firms discount recent data more heavily than households. This means that firms switch between alternative predictors more quickly in response to new data. We believe that there would be merits in investigating further the empirical properties of inflation expectations processes of households and firms and examining the implications for macroeconomic dynamics in models such as ours.

## Appendix A: Calibration of the intensities of choice

The procedure for setting the elements of $\theta$ is a simple iterative one. We start with an initial guess for $\theta$ and solve for the predictor proportions $(m^{**}, n^{**})$ that represent the fixed point of the following mappings:

$$m^* = T_f\left(m^*, n\right)$$
$$n^* = T_h\left(m, n^*\right)$$

which solve for the optimal asymptotic values of $m$ and $n$ as a function of the pair $(m, n)$.[21] To derive the mapping $T_f$ we perform the following experiment. For each value of $n$ we find the $m^*$ which would be generated asymptotically by the predictor choice mapping (that is if the predictor proportions remained at $n$ and $m^*$) forever. The $T_h$ mapping is constructed analogously and the point $(m^{**}, n^{**})$ represents the intersection of these curves in $(m, n)$ space.[22] With this information in hand we examine whether

$$\begin{bmatrix} n^{**}\,\Xi_1^h\left(m^{**}, n^{**}\right) + \left(1 - n^{**}\right)\Xi_2^h\left(m^{**}, n^{**}\right) \\ m^{**}\,\Xi_1^f\left(m^{**}, n^{**}\right) + \left(1 - m^{**}\right)\Xi_2^f\left(m^{**}, n^{**}\right) \end{bmatrix} - \imath < \varepsilon \cdot \imath$$

where $\varepsilon > 0$ is a small number and the mean squared errors are evaluated (as the notation suggests) using the asymptotic representation of the model with predictor proportions $(m^{**}, n^{**})$. If this condition is satisfied we stop the process and accept the vector $\theta$. If it is not satisfied we try another guess for $\theta$. The sequence of guesses are guided by a numerical optimisation procedure.

This approach, though arbitrary, ensures that predictor choice is not dominated by the behaviour of a subset of very volatile variables. There are other ways that we could calibrate the intensities of choice which we intend to examine in further work. One approach might be to replace the $\imath$ vector with a vector of weights representing the relative importance of each variable to household's utility or firm's profits. Another would be to use evidence from predictor choice models that have been fitted to behaviour observed in laboratory experiments as in Anufriev and Hommes (2006).

[21] These mappings are similar to the 'T-maps' considered in the predictor choice literature – see for example Branch (2004).

[22] Our approach assumes that this point is unique, though it need not be. For our applications, however, there is indeed a unique point.

## Appendix B: Estimation of gain parameters

In this appendix, we provide details of the procedure for estimating the gain coefficients that households and firms use when selecting predictors. Specifically, our aim is to obtain estimates for the parameters $\kappa_h$ and $\kappa_f$ used to construct the weighted forecast performance measures:

$$\Xi^h_{i,t} = \kappa_h SE_{i,t-1} + (1-\kappa_h)\,\Xi^h_{i,t-1}\,, i = 1, 2$$

$$\Xi^f_{j,t} = \kappa_f SE_{j,t-1} + \left(1-\kappa_f\right)\Xi^f_{j,t-1}\,, j = 1, 2$$

introduced in Section 2.4.

Our strategy is to use survey measures of inflation expectations to estimate how quickly firms and households discount information about past data when forming their expectations. We do so by estimating the gain parameters for constant-gain learning models of inflation expectations. We believe that this would provide a reasonable proxy for the rate at which past forecast errors might be discounted when evaluating alternative forecasting rules. And the parameters we estimate have a direct analogue to the case – examined in Section 3.2 – in which households and firms behave when they can switch to using a forecasting rule estimated using a constant-gain algorithm.

Our estimates of the gain parameters are constructed as follows. We first posit a forecasting rule of the form

$$\pi^e_{t,t+1} = \alpha_t + \rho_t \pi_{t-1} + \varepsilon_t$$

where $\pi$ is quarterly inflation and $\pi^e$ is the expected quarterly inflation rate. The parameters of the model are $\alpha$ and $\rho$ which are assumed to be estimated using a constant-gain OLS algorithm. We use surveys of annual inflation expectations which we match to our model-based estimate given by

$$\Pi^e_{t,t+4} = \sum_{j=1}^{4} \pi^e_{t,t+j}$$

where the terms on the right hand side are constructed by projecting forward the forecasting equation assuming the coefficients are fixed at the estimates prevailing at date $t$.

Our estimation procedure minimises the sum of squared deviations between the observed survey measure of inflation expectations and the measure produced by the model. The minimisation is with respect to three objects: the initial values of the parameter vectors ($\alpha_0$ and $\rho_0$); the initial value of the covariance matrix of the parameter estimates ($R$) and the constant-gain parameter ($\kappa$). We restrict the sequence of forecasting coefficients to be stable (so that $\rho_t$ is less than 1 in absolute magnitude for the entire sample). Minimisation is performed using a numerical search algorithm. We can regard this approach as similar to a maximum likelihood estimation in which the initial state vector is estimated jointly with the parameters of the model. This seems appropriate in our case because of the relatively short sample of data.

For firms we used the BASIX UK inflation expectations series (stripping out the contribution of 'general public' expectations). These surveys cover a range of professional economists, business leaders and trade unions. We fit our model using quarterly RPIX inflation. Our point estimate for the gain parameter is $\kappa_f = 0.0647$. For households we attempted to repeat the exercise using BASIX general public expectations, but the estimation procedure was not stable. We repeated our procedure using US data (Michigan inflation expectation and PCE deflator inflation) and obtained a very low estimate ($\kappa_h < 0.001$). While our estimation results for household inflation expectations are far from satisfactory, we believe they indicate that households are somewhat slower to adjust their inflation expectations in response to changes in data outturns. This leads us to choose the parameters $\kappa_f = 0.06$ and $\kappa_h = 0.03$ for the calibration used in the paper. The latter figure is simply half of the estimated gain parameter for firms, which seems to generate sufficiently different discounting of past data, without generating extreme outcomes.[23]

[23] As is well known, constant-gain algorithms often generate many explosive simulations when the gain parameter is set to very low levels.

## References

**Anufriev, M and Hommes, C H (2006)**, 'Evolution of market heuristics', *CeNDEF Working Paper No. 07-06*, University of Amsterdam.

**Bernanke, B (2007)**, 'Inflation expectations and inflation forecasting', speech delivered at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute.

**Branch, W (2004)**, 'Restricted perceptions equilibria and learning in macroeconomics', in Colander, D (ed), *Post Walrasian macroeconomics: beyond the Dynamic Stochastic General Equilibrium model*.

**Branch, W and Evans, G (2006)**, 'Intrinsic heterogeneity in expectations formation', *Journal of Economic Theory*, Vol. 127, No. 1, pages 264–95.

**Branch, W and McGough, B (2006a)**, 'Dynamic predictor selection in a New Keynesian model with heterogeneous expectations', *mimeo*.

**Branch, W and McGough, B (2006b)**, 'A New Keynesian model with heterogeneous expectations', *mimeo*.

**Bray, M (1982)**, 'Learning, estimation and stability of rational expectations', *Journal of Economic Theory*, Vol. 26, No. 2, pages 318–39.

**Brazier, A, Harrison, R, King, M and Yates, T (2008)**, 'The danger of inflating expectations of macroeconomic stability: heuristic switching in an overlapping generations monetary model', *International Journal of Central Banking*, Vol. 4, No. 2, pages 219–54.

**Brock, W and de Fontnouvelle, P (2000)**, 'Expectational diversity in monetary economies', *Journal of Economic Dynamics and Control*, Vol. 24, No. 6, pages 725–59.

**Brock, W and Hommes, C (1997)**, 'A rational route to randomness', *Econometrica*, Vol. 65, pages 1,059–95.

**Calvo, G A (1983)**, 'Staggered prices in a utility-maximizing framework', *Journal of Monetary Economics*, Vol. 12, pages 383–98.

**Carceles-Poveda, E and Giannitsarou, C (2007)**, 'Adaptive learning in practice', *Journal of Economic Dynamics and Control*, Vol. 31, No. 8, pages 2,659–97.

**Christiano, L, Evans, G and Eichenbaum, M (2005)**, 'Nominal rigidities and the dynamic effect of a shock to monetary policy', *Journal of Political Economy*, Vol. 113, No. 1, pages 1–45.

**de Grauwe, P (2008)**, 'DSGE modelling when agents are imperfectly informed', *mimeo*.

**Evans, G and Honkapohja, S (2001)**, *Learning and expectations in macroeconomics*, Princeton University Press.

**Evans, G and Honkapohja, S (2008)**, 'Expectations, learning and monetary policy: an overview of recent research', *CDMA Working Paper No. 08/02* .

**Evans, G, Honkapohja, S and Mitra, K (2003)**, 'Notes on agents' behavioral rules under adaptive learning and recent studies of monetary policy', *mimeo*.

**Fackler, P and Miranda, M (2002)**, *Applied computational economics and finance*, MIT press, Cambridge.

**Harrison, R and Taylor, T (2012)**, 'Non-rational expectations and the transmission mechanism', *Bank of England Working Paper No. 448*.

**Keynes, J M (1936)**, *The general theory of employment, interest and money*, MacMillan, New York.

**King, M (2005)**, 'Monetary policy: practice ahead of theory, the Mais Lecture 2005', *Bank of England Quarterly Bulletin*, Summer, pages 226–36.

**Lombardelli, C and Saleheen, J (2003)**, 'Public expectations of UK inflation', *Bank of England Quarterly Bulletin*, Autumn, pages 281–90.

**Lubik, T and Schorfheide, F (2004)**, 'Testing for indeterminacy: an application to US monetary policy', *American Economic Review*, Vol. 94, No. 1, pages 190–217.

**Lucas, R E (1986)**, 'Adaptive behaviour and economic theory', *Journal of Business*, Vol. 59, No. 2, pages S401–S426.

**Manski, C and McFadden, D (1981)**, *Structural analysis of discrete data with econometric applications*, MIT press, Cambridge.

**Marcet, A and Sargent, T (1989)**, 'Convergence of least squares learning mechanisms in self-referential stochastic models', *Journal of Economic Theory*, Vol. 48, No. 2, pages 337–68.

**Milani, F (2006)**, 'A Bayesian DSGE model with infinite-horizon learning: do 'mechanical' sources of persistence become superfluous?', *International Journal of Central Banking*, Vol. 2, No. 3, pages 87–106.

**Milani, F (2007)**, 'Learning and time-varying macroeconomic volatility', University of California, *mimeo*.

**Morris, S and Shin, H (2003)**, 'Social value of public information', *American Economic Review*, Vol. 92, No. 5, pages 1,521–34.

**Pigou, A C (1929)**, *Industrial fluctuations, 2nd Ed.*, MacMillan, London.

**Preston, B (2005)**, 'Learning about monetary policy rules when long-horizon expectations matter', *International Journal of Central Banking*, Vol. 1, No. 2, pages 81–126.

**Rotemberg, J and Woodford, M (1999)**, 'Interest rate rules in an estimated sticky price model', in Taylor, J (ed), *Monetary policy rules*, NBER.

**Sargent, T (1993)**, *Bounded rationality in macroeconomics*, Clarendon Press.

**Sargent, T (2007)**, 'Evolution and intelligent design', *mimeo*.

**Slobodyan, S and Wouters, R (2009)**, 'Estimating a medium-scale DSGE model with expectations based on small forecasting models', *mimeo*.

**Woodford, M (2003)**, *Interest and prices*, Princeton University Press.