



Working Paper no. 303

The danger of inflating expectations of macroeconomic stability: heuristic switching in an overlapping generations monetary model

Alex Brazier, Richard Harrison, Mervyn King and Tony Yates

August 2006

Bank of England

The danger of inflating expectations of macroeconomic stability: heuristic switching in an overlapping generations monetary model

*Alex Brazier**

Richard Harrison[†]*

*Mervyn King**

and

*Tony Yates**

Working Paper no. 303

* Bank of England.

[†]Corresponding author: richard.harrison@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Monetary Policy Committee. A version of this paper was presented at ‘Monetary policy: a journey from theory to practice’ an ECB colloquium held in honour of Otmar Issing, Frankfurt, 16–17 March 2006. We are grateful for insightful discussions with Charles Bean, James Proudman, Tim Taylor, Gertjan Vlieghe and Fabrizio Zampolli. We also thank Bill Branch, Buz Brock, George Evans, Cars Hommes, Frank Smets and two anonymous referees for comments. This paper was finalised on 25 July 2006.

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/workingpapers/index.htm.

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.

Contents

Abstract	3
Summary	4
1 Introduction	6
2 The model	9
3 Rational expectations and optimal policy	11
4 Modelling the choice of heuristic	14
5 Model properties under rational expectations and a single heuristic	17
6 Model properties under heuristic switching	19
7 Monetary policy under heuristic switching	23
8 Model properties after the introduction of an inflation target heuristic	30
9 Conclusions	33
References	35

Abstract

The volatility of inflation and output has fallen in most advanced economies in the 1990s and 2000s. We use a monetary overlapping generations model to discuss the cause and durability of this macroeconomic change. In that model, agents' decision rules require them to make forecasts of future inflation, which, because of shocks to productivity, is uncertain. Agents make forecasts of inflation using two rules of thumb or 'heuristics'. One is based on lagged inflation, the other on an inflation target announced by the central bank. They switch between those heuristics based on an imperfect assessment of how each has performed in the past. The way the economy propagates productivity shocks into inflation depends on the proportion of agents using each. Movements in that proportion generate fluctuations in small sample measures of economic volatility. We use this simple model of heuristic switching to contrast the performance of monetary policy rules. We find that, relative to the rule that would be optimal under rational expectations, a rule that responds to both productivity shocks and inflation expectations better stabilises the economy but does not prevent agents switching between heuristics. Finally, we study the impact of introducing an explicit inflation target, which can be used by agents as a simple heuristic, into an economy that did not previously have one. Depending on the heuristics agents have access to before the introduction of the target, this can result in reduced inflation volatility.

Summary

In the past decade we have witnessed a step change in macroeconomic performance. Both output growth and inflation have been much more stable than they were in the 1970s and 1980s.

Policymakers – keen that this development should be durable – have tried to understand the causes of this ‘Great Stability’.

Research has tended to place explanations into two groups: good luck in the form of fewer, smaller shocks; or better stewardship of the economy by governments and central banks. But as Ben Bernanke, Chairman of the Federal Reserve, has pointed out, apparently smaller shocks might actually be the result of better stewardship of the economy through anchoring of expectations about the future. With clear objectives for monetary and fiscal policy, expectations of the future need not be guided by what happens to the economy today. They can instead be guided by those clear objectives. When decisions made today by businesses and households depend on their expectations of the future, their actions will be less sensitive to past developments.

So thinking about the way expectations are formed is likely to be an important step towards understanding the ‘Great Stability’. Standard economic models tend to assume that people form expectations using detailed knowledge about the way the economy works and the shocks that hit it. What if, instead, they use simple rules of thumb – or ‘heuristics’ – to form their expectations?

This paper explores that question using a very simple model in which output and inflation today depend on expectations of inflation tomorrow. Although the model is too abstract to explore questions about current monetary policy, it allows us to explore the more general issue of the role that expectations play in shaping economic performance.

In this model, predictions about future inflation affect decisions about what to do today. So if, other things being equal, inflation is expected to be high tomorrow, it will start to pick up today. Predicting future inflation is difficult because the economy is subject to temporary, unpredictable shocks. So people are assumed to use one of two heuristics. The first is to assume that inflation tomorrow will be the same as inflation yesterday – the ‘lagged inflation’ heuristic. The second is to assume that inflation tomorrow will equal a target announced by the central bank – the ‘inflation target’ heuristic.

It seems reasonable to assume that people will only adopt a particular heuristic if it would have predicted inflation well over the past. That opens up the possibility that people might switch between them. Sometimes, when inflation has been close to target, they are likely to use the inflation target heuristic. When they do, we know that it helps to keep inflation stable because it acts as an anchor for expectations. Sometimes, however, shocks will move inflation away from the target. If people then switch to using a lagged inflation heuristic, there will not be a firm anchor for expectations. This means that there are periods like ‘Great Stabilities’ in which inflation is very stable but these are interspersed with periods of greater volatility.

In our experiments, the inflation target is only used as a heuristic if it would have performed better than others in the past. There is no guarantee that it will be used. But the announcement of an inflation target at least opens up the possibility of more stable periods of economic performance. And, as such, inflation in an economy with an inflation target tends to be more stable than in an economy without a target.

That illustrates the importance of the monetary policy framework in this model. Given that framework, what does this simple model say about how a central bank should operate? It is not possible for us to draw conclusions for the conduct of monetary policy in the real world because the model is so abstract. But we illustrate how, in this model, monetary policy can better stabilise the economy by responding to inflation expectations. This contrasts with many standard economic models, in which inflation expectations contain no information about the state of the economy that is not already apparent in other indicators. But in our model, inflation expectations do contain information about the state of the economy. They indicate which heuristic people are using and, therefore, how the economy will respond to shocks.

1 Introduction

The United Kingdom has experienced a period in which the volatility of both real and nominal variables has fallen. In the ten years to 2005, the standard deviation of output growth was less than a third of its value in the ten years to 1985; the standard deviation of inflation was less than a tenth of its value in the ten years to 1985. Inflation persistence has also fallen dramatically. Similar developments are apparent in other advanced economies such as the United States and the euro area. These changes have found various names: the ‘Great Stability’, ‘Great Moderation’ or the ‘NICE’ (non-inflationary consistently expansionary) decade.⁽¹⁾ Policymakers face a challenge in judging how to react to these changes because their causes, and therefore their durability, are uncertain, as Velde’s (2004) lucid survey of the research so far makes clear.

There are two types of explanation for these changes. The first is that the reduction in volatility is due to better monetary and fiscal policy. The second is that it reflects either smaller shocks, or changes in the way those shocks are propagated into output and inflation volatility. Thus far, econometric studies have tended to put most of the improvement down to what Velde described as policymakers having a ‘good hand’ rather than engaging in ‘good play’: witness the line of work including Stock and Watson (2002), Sims and Zha (2004), Cogley and Sargent (2005) and many others. But Bernanke (2004) suggested that what is counted as good luck in such studies includes the effect of better monetary policy in anchoring inflation expectations.

Our paper presents a model in which the link between fluctuations in the time-series properties of inflation and expectations-formation is explicit. We work with a monetary overlapping generations model, in which we assume agents form expectations by choosing amongst simple rules of thumb or ‘heuristics’. Agents work when they are young and sell their output to the old in exchange for money, which is the only store of value available to them. They consume using that money when they are old. Young agents seek to minimise the disutility from working when young and maximise the utility they will gain from consuming when old. In doing this, they face the problem of forecasting the future purchasing power of the money balances they accumulate when young. In other words, they need to forecast the change in the price level. Uncertainty about future inflation is generated by our assumption that the productivity of young agents is subject to shocks.

(1) See, for example, Bernanke (2004) and King (2003).

We contrast the rational expectations equilibrium with that which emerges when agents use a finite set of heuristics to make their forecasts of inflation. They choose between the heuristics on the basis of their performance in forecasting inflation in the recent past. We assume they observe that performance with some noise but, the better the true past performance of a heuristic, the greater chance there is that an agent uses it to make the next period's forecast. These heuristics, as Gigerenzer *et al* (1999) and others have noted, are both fast to compute and frugal in their information requirements. Model-consistent expectations are attractive devices for those who work with model economies, but it may actually not be rational for agents to have acquired them, given the informational and computational costs of doing so. Our agents choose between two heuristics: one that sets forecast inflation equal to the steady-state value, which we term loosely an 'inflation target' heuristic, and one in which forecast inflation is set to the latest realisation of inflation, which we term the 'lagged inflation' heuristic.

Our model is closed by a process for nominal money growth, which characterises central bank behaviour. We use two such processes to study the dynamics of inflation: one in which the central bank follows the rule that would be optimal in the event that expectations were rational; and another that assumes the central bank attempts to take account of heuristic behaviour.

Our strategy is to use a model of heuristics to explain the Great Stability. We are therefore exploring an idea put forward by Branch and Evans (2005). And in combining a monetary overlapping generations model with heuristics, we are borrowing from Brock and de Fontnouvelle (2000), who did this in their quest to see whether heuristic behaviour could sustain equilibria in which paper money is valued.

When agents switch between inflation-forecasting heuristics, the time-series properties of inflation change over time. On average, the majority of agents use the inflation target heuristic. But there are times when everyone does, and times when no-one does. The way the economy propagates productivity shocks into inflation depends on the proportion using each heuristic. Because this proportion fluctuates, so does the way shocks are propagated into inflation. So the volatility of inflation is higher than in a rational expectations version of the model. It also means that there are greater *fluctuations* in the volatility of inflation and in the persistence of inflation. This model, for either of the money processes we use, exhibits pronounced episodes of high, followed by low inflation volatility and persistence. When agents use the inflation target heuristic, inflation tends

to be less variable and less persistent than when more agents use the lagged inflation heuristic.

We contrast the money process that would be optimal under rational expectations with one that attempts to take account of heuristics. We do so with the usual caveats that must accompany welfare analysis in overlapping generations models. Our welfare criterion is the unconditional expectation of the sum of the welfare of the old and young in any time period. This is equivalent to maximising the average level of welfare over all generations.

Under rational expectations, the optimal policy is for money growth to respond to the level of productivity. Such a rule eliminates both the volatility of labour supply, which is costly to the young, and the volatility of consumption, which is costly to the old. The success of monetary policy under rational expectations can be attributed to its leverage over expectations. By committing to future policy actions, monetary policy has extra leverage over current labour supply and inflation.

That leverage is not available when agents use heuristics so we investigate how policy might adapt in those circumstances. The model under heuristics is highly non-linear. There is no analytical expression for optimal policy available, so we confine ourselves to a search for a rule that responds linearly to two important state variables in the model: productivity and expected inflation. The best rule – according to our welfare criterion – amongst this ‘two-pillar’ class of rule increases money growth when productivity is high, and by more than under rational expectations; and it reduces money growth when inflation expectations rise. The welfare benefits from shifting away from the rational expectations policy are greater during periods when agents are using the backward-looking heuristic. Despite a monetary policy that attempts to take account of heuristics, heuristic switching still occurs and so there are still fluctuations in inflation volatility and inflation persistence. At the same time this model generates fluctuations in the estimated disturbances to linear autoregressive equations for inflation, echoing the findings of econometricians on real data.

The message from the paper to this point is that very stable macroeconomic outturns should not be taken for granted. But we go on to explore the notion that the widespread adoption of explicit inflation objectives by central banks can be modelled as the provision of a heuristic to which agents did not previously have access. When we introduce an inflation target heuristic to agents, we find that at least some adopt it immediately, and that subsequently the volatility of inflation is

lower, despite the heuristic switching that ensues. We illustrate how the impact of the introduction of the inflation target depends on how poorly performing is the heuristic with which agents start out.

2 The model

Our model is an overlapping generations model with money. It is deliberately stylised and was chosen as the simplest possible model in which agents must forecast future inflation.

Agents live for two periods. They work when young and consume when old. Young agents minimise the disutility from work (L) when young and maximise the expected utility from consumption (C) when old.⁽²⁾ Their output is produced with a linear technology, denoted $A_t L_t$, where A_t is productivity, known at time t when young agents determine their labour supply. Their output is sold at price P_t . Young agents accumulate nominal money balances (M_t) to the value of their output. Their consumption when old is determined by the real value of those same money balances $\left(\frac{M_t}{P_{t+1}}\right)$. We denote expectations formed by agents using the operator E_t . In some cases that will refer to rational expectations and in others will refer to a heuristic. At each stage we will make clear how agents are forming their expectations.

Formally, young agents solve the following problem:

$$\max_{L_t} E_t \left[-\frac{L_t^{1+\eta}}{1+\eta} + \frac{C_{t+1}^{1-\alpha}}{1-\alpha} \right] \quad \eta > 0, 0 < \alpha < 1 \quad (1)$$

subject to:

$$M_t = A_t L_t P_t \quad (2)$$

$$C_{t+1} = \frac{M_t}{P_{t+1}} \quad (3)$$

The problem that old agents solve is degenerate. They maximise utility by spending all their real balances on consumption goods. The young accumulate money from the old and from the government. The government's budget constraint implies that the nominal money stock evolves according to:

$$M_t = M_{t-1} + P_t D_t \quad (4)$$

where $D_t > 0$ is output purchased from the private sector in exchange for money. We assume that

(2) Note that, for simplicity, we assume that there is no discounting of future consumption.

government purchases are used for purposes that do not yield private utility.⁽³⁾ The instrument of monetary policy is the growth rate of the nominal money stock, G :

$$M_t = M_{t-1} + G_t M_{t-1} = (1 + G_t) M_{t-1}$$

so that, since $P_t D_t = G_t M_{t-1}$, the nominal value of government purchases equals the increase in the nominal money supply: there is no distinction between fiscal and monetary policy in this model.

The young consumer's problem can now be written as:

$$\max_{L_t} E_t \left[-\frac{L_t^{1+\eta}}{1+\eta} + \frac{1}{1-\alpha} \left(A_t L_t \frac{P_t}{P_{t+1}} \right)^{1-\alpha} \right]$$

Denoting inflation as $\Pi_{t+1} = \frac{P_{t+1}}{P_t}$, the first-order condition for labour supply is given by:

$$L_t^{\eta+\alpha} = E_t (A_t \Pi_{t+1}^{-1})^{1-\alpha}$$

This equation makes it clear that young agents have to make forecasts. If expected inflation tomorrow is high, agents expect the value of any money balances they accumulate by working when young to be eroded when they are old. Their demand for money balances will be lower.

Uncertainty about the future price level is introduced by a simple, stochastic process for productivity (A_t):

$$A_t = A_{t-1}^\rho Z_t \tag{5}$$

where Z_t follows a lognormal distribution.

For ease of exposition, we proceed by taking a first-order approximation around the non-stochastic steady state. Using lower case letters to denote log deviations from steady state, the (log-linearised) first-order condition for labour supply is:

$$l_t = \frac{1-\alpha}{\eta+\alpha} a_t - \frac{1-\alpha}{\eta+\alpha} E_t \pi_{t+1} \tag{6}$$

We use m_t to denote the log deviation of *real* money balances, $\frac{M_t}{P_t}$, from steady state. The *real*

(3) We could, analogously, assume that government purchases are redistributed back to agents, and that these redistributions enter utility in a way that was additively separable from other components. Our marginal condition for labour supply would be identical in this model, although consumption and mean levels of welfare would not be. Dropping the simplification used here would not affect the impact of heuristic switching on the dynamics of macroeconomic outcomes.

money demand condition is:

$$m_t = a_t + l_t$$

$$m_t = \frac{1 + \eta}{\eta + \alpha} a_t - \frac{1 - \alpha}{\eta + \alpha} E_t \pi_{t+1} \quad (7)$$

The linearised version of the government budget constraint is given by equation (8) below, where we denote the steady-state inflation rate as Π and use g_t to denote the absolute (note, not log) deviation of the growth rate of nominal money from its steady-state level.⁽⁴⁾

$$m_t = \Pi^{-1} g_t + m_{t-1} - \pi_t \quad (8)$$

We linearise around a positive steady-state inflation rate ($\Pi > 1$) to ensure that the frequency of negative government spending levels D implied by money growth g is negligible: we do not regard such outcomes as economically meaningful.

Linearising the productivity process gives:

$$a_t = \rho a_{t-1} + \zeta_t, \zeta_t \sim N(0, \sigma^2) \quad (9)$$

where lower case ζ_t is the log deviation of the disturbance Z_t from its steady-state value, 1.

To summarise the model: to maximise their expected utility, young agents must forecast inflation. Uncertainty about future inflation is introduced by fluctuations in the demand for real money balances arising from shocks to productivity. If those movements are not matched by equal movements in the nominal money stock, inflation will fluctuate. In the next section we calculate the monetary policy that maximises welfare when agents form rational expectations of inflation.

3 Rational expectations and optimal policy

The model is described by equations (6), (7), (8) and (9) together with an equation for money growth, g_t . We assume that monetary policy is characterised by the design of a rule for money growth to which the policymaker commits. The rule is designed to maximise a particular measure of welfare. It is designed before any realisation of productivity is observed so although money growth can respond to realisations of productivity, the policy rule itself is invariant to changes in productivity.

(4) The coefficient on g results from the fact that $(1 + g) = \Pi$ in steady state.

Our welfare measure is the sum of the utility of the young and old agents:

$$W_t \equiv -\frac{L_t^{1+\eta}}{1+\eta} + \frac{C_t^{1-\alpha}}{1-\alpha}$$

This differs from the utility function of a young agent (equation **(1)**) because it adds the utility of today's old to the disutility of work experienced by today's young. We assume that policy is designed to maximise the unconditional expectation of welfare. This maximises the average level of welfare across all generations and across all possible realisations of productivity.⁽⁵⁾

We assume that monetary policy maximises welfare taking the steady-state level of money growth as given. In this model, there would be welfare improvements from lowering the mean level of money growth and the associated government purchases (which do not yield private utility). We abstract from that component of policy to focus on the stabilisation role of monetary policy. Hence, the curvature of the welfare function means that, by stabilising the economy, we maximise the average level of welfare. Note that, conditional on a level of productivity that is known and different from the steady-state level of productivity, agents will not prefer steady-state levels of labour supply and future consumption. But, before the value of productivity is revealed, they will prefer stable over variable labour supply and consumption because of the curvature in utility. Welfare is maximised when labour supply and consumption do not deviate from their steady-state levels.

Our welfare function is

$$E[W_t - W] = -\frac{\eta}{2}E[l_t^2] - \frac{\alpha}{2}E[c_t^2] \quad \mathbf{(10)}$$

which we derive as the second-order Taylor approximation to the welfare measure. Policy maximises the unconditional expectation of a weighted sum of the variances of young agents' labour supply and old agents' consumption. Note that the linear terms that are anticipated in a second-order Taylor expansion drop out: the unconditional expectation of linear terms in log deviations from the steady state are zero.

Under rational expectations, we now demonstrate that monetary policy can stabilise labour supply and consumption completely by committing to a rule for money growth that feeds back from the

(5) Our procedure is similar to the practice of maximising 'period utility' in the monetary policy design literature that uses representative agent models.

model's driving variable, productivity:

$$g_t = \chi a_t \quad (11)$$

It is straightforward to show that, for an arbitrary value of χ , the rational expectations solutions for real money balances and inflation are given by:

$$m_t = a_t + l_t = \frac{1 + \eta - (1 - \alpha) \rho \frac{\chi}{\Pi}}{1 + \eta - (1 - \alpha) \rho} a_t \quad (12)$$

and:

$$\begin{aligned} \pi_t = & \left[\frac{\chi}{\Pi} \rho + \frac{1 + \eta - (1 - \alpha) \rho \frac{\chi}{\Pi}}{1 + \eta - (1 - \alpha) \rho} (1 - \rho) \right] a_{t-1} \\ & + \left[\frac{\chi}{\Pi} - \frac{1 + \eta - (1 - \alpha) \rho \frac{\chi}{\Pi}}{1 + \eta - (1 - \alpha) \rho} \right] \zeta_t \end{aligned} \quad (13)$$

Policy can completely stabilise employment when $m_t = a_t$. From equation (12), this is the case when $\chi = \Pi$. Under this rule there are no welfare costs to young agents from macroeconomic volatility. But what happens to the volatility of inflation and (hence) the utility of old agents? We know that the consumption of the old generation is determined by their accumulated money balances adjusted for subsequent inflation:

$$c_t = m_{t-1} - \pi_t$$

and, when $\chi = \Pi$, the equilibrium inflation equation (13) can be simplified to:

$$\pi_t = a_{t-1}$$

We already know that real money balances equal productivity because labour supply is stabilised:

$$m_t = a_t \Rightarrow m_{t-1} = a_{t-1}$$

so that:

$$c_t = m_{t-1} - \pi_t = 0$$

A policy rule in the form of equation (11), setting $\chi = \Pi$, eliminates all of the welfare costs of macroeconomic instability. Such a rule generates movements in inflation in the next period that are equal to the realisation of productivity in the current period. This strategy means that the real

value when old of any money balances accumulated when young is unaffected by realisations of productivity. Anticipating this, the young have no incentive to change their labour supply in response to changes in productivity. With labour supply constant and the impact of productivity on real money balances offset by inflation, the consumption of the old is constant. The key to the success of monetary policy in stabilising both labour supply and consumption is its leverage over not only the current money stock but also over anticipated future inflation. Indeed, it is clear from (6) that monetary policy can stabilise labour supply in the face of productivity disturbances *only* through its leverage over inflation expectations.

To re-emphasise: note that complete stabilisation of consumption and employment is optimal because of the curvature of agents' utility (a feature preserved by our quadratic approximation). Note too that monetary policy does not prevent agents responding to productivity shocks; it simply creates conditions that means that it is optimal for agents not to.

4 Modelling the choice of heuristic

So far we have assumed model-consistent expectations to provide a benchmark against which to compare subsequent departures from that assumption. Many have argued that in reality agents would find it too costly, or would not have the means to collect the information and carry out the computations required for a rational expectations equilibrium to be achieved. The route we choose is to adopt a model in which agents may have heterogeneous expectations and in which those expectations are based on simple heuristics.

4.1 The heuristic choice literature

The literature on heuristics is itself now very large and ably surveyed by one of its recent leaders in Hommes (2005). He charts the history of this strand of thought from the suggestion by Keynes (1936) that fluctuations in sentiment would influence the macroeconomy; through Simon (1957), who explained that agents were 'boundedly rational' in the face of costs of collecting information and computing the outcomes of their decisions. Another landmark is the emergence of experimental evidence that agents use simple heuristics to make decisions, culminating in Kahneman's (2003) Nobel lecture. This led to a large research programme exploring why it may have proven beneficial for nature to endow us with such heuristics: a topic that occupies, for

example, Gigerenzer *et al* (1999). We use a model in which agents choose between a finite set of heuristics based on noisy observations of past forecast performance. The papers from which we draw most inspiration in this respect are Brock and Hommes (1997), Brock and de Fontnouvelle (2000) and Branch and Evans (2005, 2006) who in turn ground their decision-making model in the discrete decision, multinomial logit models set out in Manski and McFadden (1981).⁽⁶⁾

We are not the first to combine a monetary overlapping generations model with a model of heuristic expectations formation. Brock and de Fontnouvelle (2000) do just this. But their concern is very different. Early students of rational expectations, monetary overlapping generations models noted that these models generated equilibria in which money had value and equilibria in which it did not. This was a source of discomfort since paper money in reality is pervasive, and yet there was no guide as to which amongst the model's equilibria should or would be selected. Brock and de Fontnouvelle (2000) is an effort to see whether heuristic behaviour can lead to monetary equilibria: they find that it can.

4.2 *Heuristic choice in our model*

Our agents select from two heuristics described by:

$$E_{1,t}\pi_{t+1} = \pi_{t-1}$$

$$E_{2,t}\pi_{t+1} = 0$$

The first predictor ($E_{1,t}\pi_{t+1}$) sets expected inflation equal to the latest observed outturn. We term this the 'lagged inflation' predictor. This predictor is based on lagged inflation (π_{t-1}) and not current inflation (π_t) which will itself depend on agents' expectations and will not be realised at the time agents are forming their expectations. The second predictor ($E_{2,t}\pi_{t+1}$) sets expected inflation equal to the target (since π represents the deviation of inflation from target we have

(6) See also de Grauwe and Grimaldi (2006). They show how exchange rate dynamics, and fluctuations in the performance of fundamentals models of the exchange rate, are affected by heuristic switching, embedding the Brock and Hommes approach, using the same model of predictor choice that we employ.

$E_{2,t}\pi_{t+1} = 0$). This we term the ‘inflation target’ predictor.⁽⁷⁾ This particular set of predictors includes plausible models for agents to use to forecast, but is itself arbitrary. For most of our analysis, exactly what is in this set of predictors is not important. What is important is that there are different predictors and that switching amongst them will generate changes in the way the model propagates shocks. However, later in the paper, we interpret the inflation target predictor as one that can be added to the set of available predictors if the central bank declares an explicit inflation objective. At that point it will be crucial to consider predictor sets that initially exclude, and later include, the inflation target predictor, so our predictor set must be taken more literally.

Agents in our model differ from those embedded within adaptive learning models. In those models, the tools that agents use to forecast encompass the true model. In variants where agents have access to the entire history of data, they may eventually learn the true coefficients. Our agents’ models are both misspecified, and agents have a fixed window for evaluating their predictors that prevents the apparent performance of these predictors converging over time.

We follow our predecessors in this literature and assume that the heuristics are selected according to their recent forecast performance. Specifically, we define the objective function as:

$$F_{i,t} = -\frac{1}{H} \sum_{j=1}^H [\pi_{t-j} - E_{i,t-j-1}\pi_{t-j}]^2 \quad (14)$$

for $i = 1, 2$. The term on the right-hand side is the ‘mean squared error’ of the heuristic, calculated over the previous H periods. This captures the ability of the heuristic to match the behaviour of inflation in the recent past. The objective can be thought of as some form of ‘utility function’: agents prefer heuristics with higher F scores.⁽⁸⁾

The proportion of agents choosing each predictor, $n_{i,t}$ is determined by the following function:

$$n_{i,t} = \frac{\exp(\theta F_{i,t})}{\sum_{j=1}^2 \exp(\theta F_{j,t})} \quad (15)$$

where the parameter $\theta > 0$ is referred to in previous work as the ‘intensity of choice’. Brock and de Fontnouvelle (2000) note that in this model θ is related to the amount of noise in observing the forecast error function F .⁽⁹⁾ The larger is θ , the more accurately agents observe the past forecast

(7) Mojon and Diron (2005) document how using the central bank’s stated target as a forecast rule of thumb can perform well relative to alternative models.

(8) The thought experiment that agents are conducting here is flawed, and highlights the difference between their behaviour and that under rational expectations: the performance of a heuristic in forecasting actually depends on how many agents use it for forecasting. Agents neglect this fact when they compute F from recent observations on π .

(9) The authors steer the reader to the unabridged (1996) version of this paper, *University of Wisconsin Working*

performance of the heuristics, and the more the portion of agents using each heuristic responds to forecast performance. The limit of $\theta = \infty$ represents the case in which all agents observe perfectly – and hence choose – the best heuristic in each period. As θ approaches zero, we approach a situation in which the noise in observing predictor performance is so large that predictor choice is entirely non-systematic. To emphasise, with a finite θ , the presence of measurement error means that agents will not always pick the best-performing heuristic. But the probability that they will pick a particular heuristic will increase with its past forecasting performance. The share of the population using each of the two heuristics will equal the probability that any individual picks that heuristic.

Aggregating across young agents, we have the following:

$$E_t \pi_{t+1} = n_{1,t} \pi_{t-1}$$

Thus the real money demand relation under heuristics is given by:

$$m_t = \frac{1 + \eta}{\eta + \alpha} a_t - \frac{1 - \alpha}{\eta + \alpha} n_{1,t} \pi_{t-1} \quad (16)$$

5 Model properties under rational expectations and a single heuristic

We simulate the model comprising the equation for n_1 , the portion using the lagged inflation heuristic, (15), and the linearised equations for real money demand, the government budget constraint, and the productivity and money processes (equations (16), (8), (9) and (11) respectively).

We use the following parameter values: $\eta = 0.2$, $\alpha = 0.41$, $\Pi = 1.02$, $\theta = 100000$, $\rho = 0.925$, $\sigma^2 = 0.000075$ and $H = 50$. Critically assessing the suitability of these parameters is difficult, given the highly stylised structure of the model. We emphasise simply that we are using this model in the hope that it can say something interesting about the dynamics of an economy over business cycle frequencies, and be of interest to monetary policy makers who have to design a policy to stabilise the economy over such time periods.

Nevertheless, some discussion of our chosen parameters is warranted. Our choices for η and α imply that the elasticity of real money demand to expected inflation (equal to $\frac{1-\alpha}{\eta+\alpha}$) is close to

Paper no. 9624, for a complete account of this interpretation (and others) of the model.

unity, which means that real money balances are relatively responsive to expected inflation. Marcet and Nicolini (2003) use parameter values that imply that real money demand is relatively less responsive to changes in expected inflation (their parameters would imply a slope $\frac{1-\alpha}{\eta+\alpha}$ of around 0.15), but simulations under this type of parameterisation are qualitatively similar to those we present here.

Our choice of Π implies that the steady-state inflation rate is 2% per period, which matches the rate chosen by some central banks if we interpret a period as one year. This choice bounds our choice for the variance of the productivity disturbance: this, together with the design of the process for monetary policy, g , will govern the frequency with which the implied level of government spending is negative, which we want to keep to a minimum. The degree of persistence in the shocks affects the chance of lagged inflation proving to be a good forecaster of future inflation, and therefore of agents using it as a heuristic. The variance of productivity implied by our assumed values for σ^2 and ρ is of a similar order of magnitude to cyclical output variations.⁽¹⁰⁾

The ability of the model to generate switches in heuristic use is also determined by the evaluation horizon H and the intensity of choice θ (which we prefer to interpret as the accuracy with which heuristic performance is observed). The shorter the evaluation horizon, the larger the fluctuations in observed forecast performance. The greater the intensity of choice, the larger the response of heuristic choice to movements in forecast performance. The important thing for the story in this paper is that some economically significant degree of heuristic switching occurs.

The table below records some time-series properties of four versions of our overlapping generations model. We report variances as an index for which 100 equals the rational expectations case. In each case the model is solved under the money process that is optimal under rational expectations. The first column reports the rational expectations version of the model discussed in Section 3. The variance and autocorrelation of inflation are calculated from the equivalent moments of the forcing process, productivity. For the other cases, statistics are computed from a 20,000-period simulation with all variables initialised at steady-state values. ‘Lagged inflation’ refers to a model in which agents are restricted to the heuristic that inflation

(10) The standard deviation of log productivity (a) is given by $\sqrt{\sigma^2 / (1 - \rho^2)} \approx 0.023$. The variance of residuals from a regression of annual UK (log) GDP on a time trend is around 0.03.

tomorrow is equal to inflation yesterday. ‘Inflation target’ refers to a model in which they are restricted to the inflation target heuristic.

Table A: Time-series properties of the rational expectations and single heuristics models

	Rational expectations	Lagged inflation	Inflation target
var(Π)	100	959	135
var(var(Π))	100	23220	108
$\rho(\Pi)$	0.92	0.55	0.72

Note: Variances relative to rational expectations case (=100).

These results serve as a benchmark against which we compare our model when agents switch between the two heuristics. They also provide some intuition about what happens to the time-series properties of variables as the number using each heuristic switches between the extremes implied by these first simulations. The first row of Table A shows the variance of inflation, which is about ten times larger when all agents use the lagged inflation heuristic compared with the rational expectations benchmark. The second row shows the variance of the variance of inflation. This is computed by first forming a time series of a rolling 50-period variance of inflation and then calculating the variance of that. When all agents use the lagged inflation heuristic, this measure is about 200 times larger than in the rational expectations case. The final row shows the coefficient from a first-order autoregression of inflation. This illustrates how the estimated time-series behaviour of inflation depends on the method with which agents are forecasting inflation. The results for the ‘inflation target’ model are similar to those for ‘rational expectations’.

6 Model properties under heuristic switching

In this section we report the results from simulating the model when agents switch between the two heuristics depending on their past forecasting performance. As a benchmark, we continue to assume that money growth follows the process that would be optimal if agents formed rational expectations. The summary statistics are shown in Table B. Once again, these simulations are over 20,000 periods. We continue to normalise all variances to equal 100 in the rational expectations case. Overall, the variance of inflation in this heuristic-switching economy is higher

than when all agents were forced to use the inflation target heuristic but lower than in the economy where all agents used the lagged inflation heuristic. The same is true of fluctuations in the small-sample variance of inflation.

Table B: Time-series properties of heuristic-switching model

$\text{var}(\Pi)$	190
$\text{var}(\text{var}(\Pi))$	1336
$\rho(\Pi)$	0.70

See note to Table A.

In Chart 1 below, we extract from our 20,000-period simulation an illustrative subsample of 1,000 consecutive periods. The top panel of the chart shows how the portion using the lagged inflation heuristic, n_1 fluctuates. It sometimes reaches the upper bound of 100%, but is generally close to zero. On average, the number using the lagged inflation heuristic is about 30%. Switching between the two heuristics is an important determinant of the time-series behaviour of variables.⁽¹¹⁾

Chart 2 is an alternative – histogram – representation of these movements in n_1 . It shows that the distribution of n_1 is bimodal. If the intensity of choice (θ) was infinite, then we would expect the observations to be either $n_1 = 0$ or $n_1 = 1$ as agents are able to perfectly observe the best performing predictor. But since θ is finite (though large) there are some observations between these extremes.

While the model spends most of the time in a region where the majority of agents are using the inflation target heuristic, there are episodes where almost all are using the lagged inflation heuristic. These results reflect the fact that agents in our model use a finite sample of recent data to evaluate predictor performance: in the jargon of the learning literature, they assess forecast performance using ‘constant gain’. If instead we allowed agents in the model to learn with ‘decreasing gain’ (that is, using the entire history of the data), the model would generate a histogram centred around a single, interior value of n_1 . This is because our model exhibits what has been called ‘negative feedback’ from heuristic use to heuristic performance. These aspects of

(11) Indeed, when plotting inflation alongside the series for productivity (a_t) it is difficult to discern by eye how the productivity shocks are transmitted into inflation outcomes. The reason is simply that heuristic switching changes the coefficients in the model equations: that is, the mapping from exogenous shocks to endogenous variables.

Chart 1: Heuristic switching under rational expectations policy (first and second-order autocorrelations are regression coefficients on first and second lags of inflation)

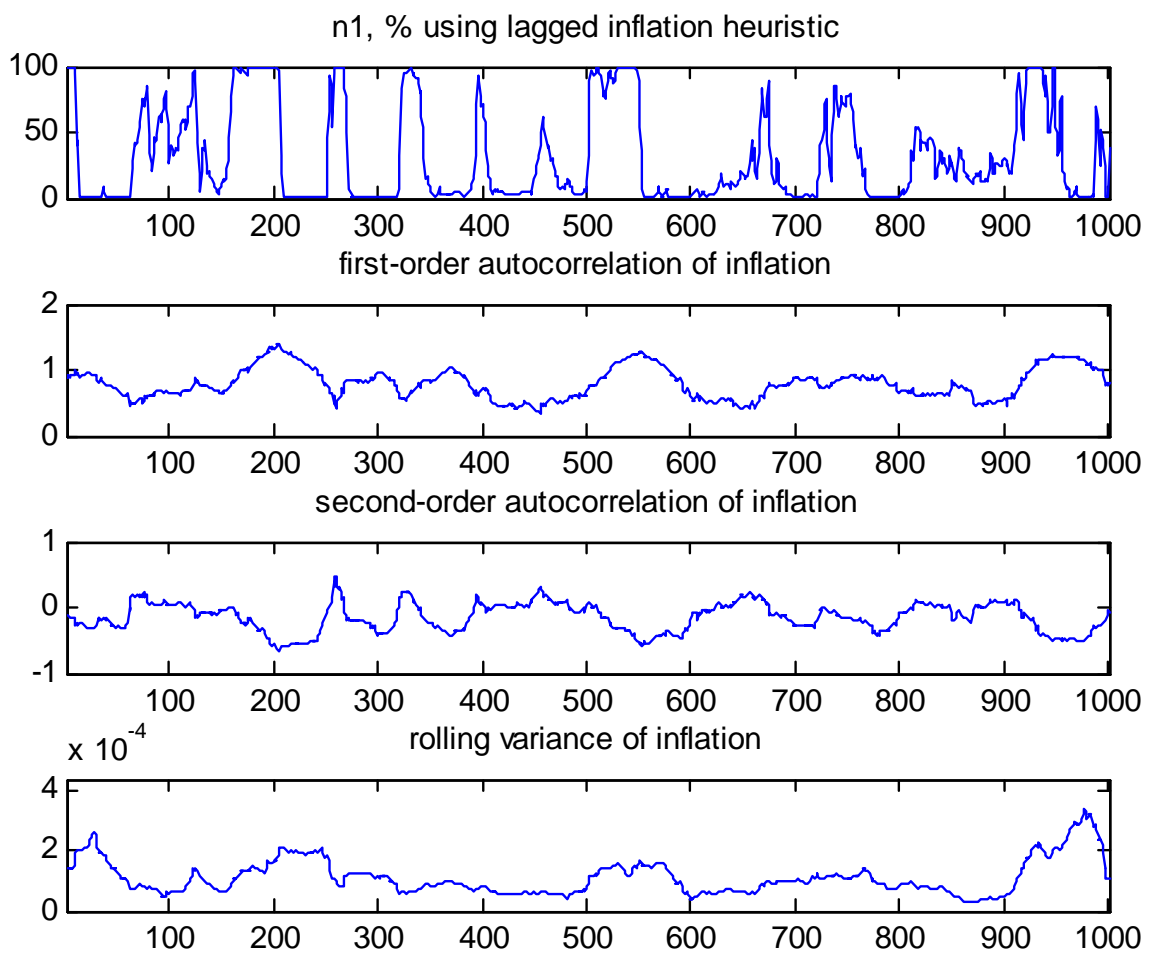
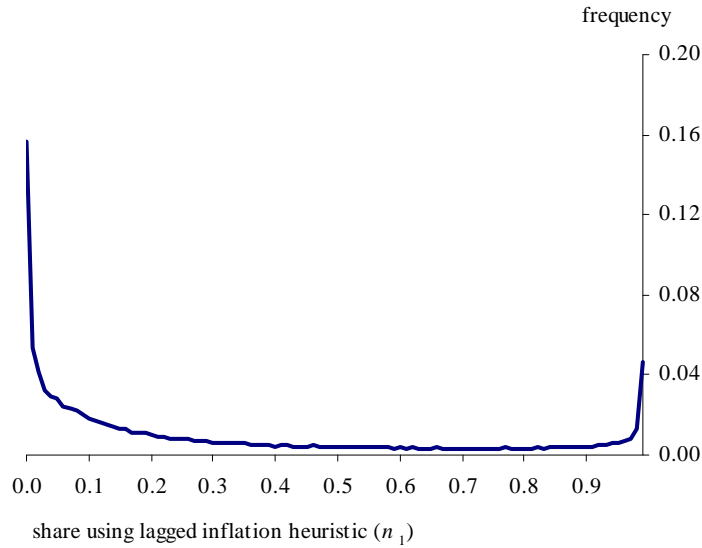


Chart 2: Share (n_1) using lagged inflation heuristic under rational expectations policy



macroeconomic models with predictor choice are discussed in Branch and Evans (2005) who suggest that this negative feedback effect may be relatively uncommon in macroeconomic models. Instead, they construct a simple model with ‘positive feedback’, characterised by multiple equilibria, some of which are unstable. At such equilibria, disturbances that, for example, increase the proportion of agents using a given predictor improve the relative performance of that predictor, further increasing the proportion, and so on.

Positive feedback and multiple equilibria can be generated in our model under suitable parameterisations for the productivity process and the conduct of monetary policy. For example, we found that the monetary reaction function⁽¹²⁾

$$\Pi^{-1}g_t = -m_{t-1} + 0.5a_t - 0.25n_{1,t}\pi_{t-1}$$

was able to generate these properties when we set $\rho = 0.6$. But under policy that is optimal when agents form rational expectations, and indeed, under the policy that attempts to take account of heuristic switching that we derive below, we have negative feedback between heuristic use and performance.

The bottom three panels of Chart 1 illustrate how heuristic switching generates small-sample

(12) The coefficient on the lag of real money balances is suggested by the form of the reaction function used by Branch and Evans (2005).

fluctuations in the time-series properties of inflation. The panels labelled first and second-order autocorrelation report rolling coefficients from a regression of inflation on two lags of itself. The bottom panel plots the variance of inflation. These moments are calculated over a horizon of 50 periods. When the proportion of agents using the lagged inflation heuristic is high for a sustained period, so is the variance of inflation; at these times the coefficient of the first lag of inflation in an autoregression of inflation is high, and the coefficient on the second lag is low. We gain some insight into these fluctuations by fixing n_1 and writing the reduced form for inflation:

$$\pi_t = \Pi^{-1}g_t - \frac{1 + \eta}{\eta + \alpha}a_t + \frac{1 + \eta}{\eta + \alpha}a_{t-1} + \frac{1 - \alpha}{\eta + \alpha}n_1\pi_{t-1} - \frac{1 - \alpha}{\eta + \alpha}n_1\pi_{t-2}$$

As we see in the simulations, so in this reduced-form equation for inflation we notice that the higher is n_1 the higher is the coefficient on π_{t-1} and the lower is the corresponding coefficient on π_{t-2} .

These fluctuations in the autocorrelation function for inflation echo the debates about what has caused the fluctuations in inflation persistence, documented by, amongst others, Benati (2004) and Levin and Piger (2004). That debate has thrown up two broad answers: that changes in inflation persistence have come about because of structural change; or that they reflect changes in monetary policy making and the introduction of inflation targeting. Our model generates changes in small-sample moments of inflation that reflect neither, but instead are the result of heuristic switching.

7 Monetary policy under heuristic switching

So far we have worked with the money growth process that would be optimal under rational expectations. We now consider if the central bank can improve on this process in the light of its knowledge about expectations formation. There are two motivations. From a positive standpoint, we can check that the heuristic-switching explanation for the appearance (and possible disappearance) of low inflation volatility is robust to cases in which the central bank follows a more sensible policy. From a normative standpoint, we can highlight the cost of the central bank incorrectly assuming that expectations are rational.

In Section 3, we showed that, under rational expectations, a rule for money growth that responded

to productivity could stabilise labour input and consumption. It did so through its impact on anticipated future money growth and inflation. When agents use heuristics, commitment to a policy rule no longer delivers any direct leverage on expected future inflation. Policy only affects expectations indirectly through past inflation. The lack of direct leverage over expectations means that, unlike the rational expectations case, policy cannot offset all the welfare losses arising from productivity shocks. It needs to adapt to the use of heuristics.

Additional complications arise in attempting a study of the welfare consequences of policy under heuristics. Heuristic switching makes the model non-linear, even when the individual decision rules are linearised.⁽¹³⁾ This non-linearity causes two problems.

The first problem is that we cannot derive an optimal monetary policy analytically, even when we use the quadratic approximation to welfare explained above. So we have to resort to numerical methods. We define a class of candidate monetary policy processes, and then simulate the model under each rule within that class, compute welfare, and look for the rule that scores the highest. The particular non-linear nature of our model means that we have to simulate for millions of periods to get reliable estimates of our welfare function. So we must confine our search across alternative policy rules to make the exercise manageable. We will work with the following class of rules for money growth:

$$g_t = \chi_1 a_t + \chi_2 E_t \pi_{t+1}$$

This process allows the policymaker to respond to productivity and to data on expected inflation. In that sense it operates a ‘two-pillar’ strategy. We assume that policymakers receive data on expected inflation, but do not attempt directly to internalise the interaction between policy, endogenous inflation outcomes and n_1 . (Indirectly, policymakers will choose the combinations of χ_1 and χ_2 that generate the most beneficial paths for n_1 , the proportion using the lagged inflation heuristic.) We search for the values of χ_1 and χ_2 that deliver the best welfare for our agents, defined by our criterion in equation **(10)**.

The second problem caused by the non-linearity of the model is that alternative policy rules will generate small differences in the mean rates of inflation. These will cause the average levels of utility to differ according to the policy rule, as the government budget constraint means that higher average inflation implies higher average government spending and higher resource destruction.

(13) The fraction (n_1) of agents using the lagged inflation heuristic affects the coefficients of the decision rules. And n_1 itself varies over time, in response to the behaviour of the economy.

The differences in means will not affect the welfare criterion we have chosen, which is defined on variances. So it must be stressed that our search can rank policy rules only according to their stabilisation properties, and not their effect on means.⁽¹⁴⁾

We focus on rules that respond to productivity and inflation expectations for two reasons. First, this class of rules allows us to nest the optimal policy under rational expectations, which responds to the only state variable in that model, productivity. Second, it also allows the policymaker to respond to another state variable in the heuristic-switching model, expected inflation. And that happens to echo the concerns of policymakers in reality.⁽¹⁵⁾

We can get some intuition for why a rule like this is likely to work by considering an extreme case that the policymaker will face: one in which all agents use the inflation target heuristic. When everyone is using the inflation target heuristic ($n_1 = 0$), the labour supply function (6) collapses to:

$$l_t = \frac{1 - \alpha}{\eta + \alpha} a_t$$

Fluctuations in labour supply are inevitable. The average expected welfare of young agents is lower than when agents have rational expectations and policy responds optimally. Under heuristics, monetary policy is powerless to influence this. But monetary policy can help old agents. The consumption of old agents at date t is:

$$c_t = m_{t-1} - \pi_t$$

and the evolution of real money balances is given by:

$$m_t = \Pi^{-1} g_t + m_{t-1} - \pi_t$$

so that the policymaker can fully stabilise c_t by committing to the policy rule:

$$\begin{aligned} g_t &= \Pi m_t \\ &= \Pi \frac{1 + \eta}{\eta + \alpha} a_t \end{aligned}$$

(14) These small differences in mean inflation will also have a small effect on the performance of the inflation target heuristic under the alternative policy rules. The higher the mean inflation rate, the worse the (zero) inflation target heuristic performs, and the smaller the portion of agents who use it.

(15) Note that expectations-based rules have been argued to have benefits in other contexts. For example, Evans and Honkapohja (2003) have recommended them as devices for implementing monetary policy to ensure that the rational expectations equilibrium is stable under least-squares learning.

which, since $\frac{1+\eta}{\eta+\alpha} > 1$, implies a stronger response to productivity shocks than under rational expectations.

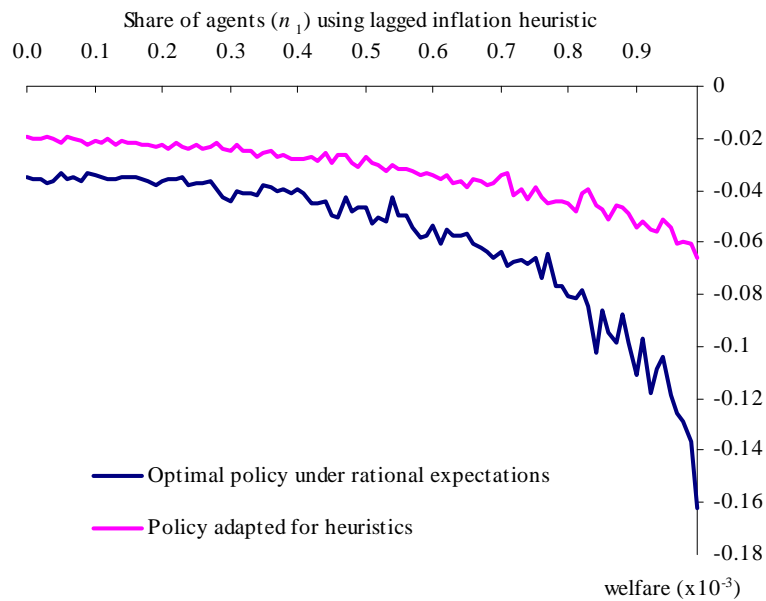
In the general case, where some agents use the lagged inflation heuristic, labour supply and the demand for real money balances depend on inflation expectations, which in turn depend on lagged inflation. In that case, even in the absence of a current productivity shock, labour supply and output can fluctuate. Without any policy action, inflation will move to bring the real value of money balances into line with output. These fluctuations are costly so monetary policy might do better by responding to inflation expectations as well as to productivity. Of course, one thing this discussion reveals is that the ideal response to productivity and inflation expectations should itself depend on n_1 . However, to make the analysis more tractable, we stick to rules that involve constant, independent values of χ_1 and χ_2 .⁽¹⁶⁾

The best rule in our grid search is one with values of $\chi_1 = 2$ and $\chi_2 = -1.75$. This policy shares a feature with the optimal policy under rational expectations in that money growth is expanded when productivity is unusually high. A positive shock to productivity reduces the price level; a positive money growth response by policy therefore acts to offset that. The policy response under heuristics is to respond more aggressively (recall that under rational expectations, χ equals Π , the steady-state rate of inflation, which is 1.02). We believe that this response allows the policy to perform well when few agents believe the inflation target: as described above, in this setting, an aggressive response to productivity can help to stabilise the consumption of old agents. The heuristics policy also suggests that money growth should fall when expected inflation rises. When expected inflation rises, labour supply and demand for real balances fall. Monetary policy can stabilise inflation by contracting the money supply.

The rule considered here generates higher welfare than arbitrary persistent processes for money growth, fixed money growth, and the policy that would be optimal under rational expectations (derived in Section 3). The welfare surface appeared well behaved in the space used for the grid search. Chart 3 below shows how welfare differs under the two policy rules at different values of n_1 , the portion using the lagged inflation heuristic. We arrange the simulated periods according to their associated value of n_1 and calculate average welfare at each value of n_1 .

(16) Using this short cut naturally raises the issue of whether it would be appropriate to build a model of heuristic policy design on the part of the central bank to go with the heuristic expectations-formation on the part of agents in the model. We leave that issue for future research.

Chart 3: Welfare generated by alternative policy rules as the share of agents using the lagged inflation heuristic varies



As we can see, when the central bank tries to take account of heuristics, it delivers higher welfare than the rational expectations policy at all values of n_1 . The welfare improvement achieved by the heuristics-adapted policy is greater for larger values of n_1 : the more agents are using the lagged inflation heuristic, the greater the benefit of following the policy adapted for heuristics, or, put another way, the greater the cost of policymakers mistakenly following the policy that would be appropriate under rational expectations.⁽¹⁷⁾

Table C shows summary statistics to compare with earlier vintages of the model, once again based on simulations of 20,000 periods. We continue to report moments of inflation as an index where 100 is the value for the model under rational expectations and the associated optimal policy.

(17) We have calculated that the minimum value for these costs, when few or no agents are using the lagged inflation heuristic, is still more than ten times the welfare cost of mistakenly pursuing the heuristics policy when agents actually have rational expectations. This is an indication that if policymakers were unsure how agents arrived at their forecasts a safe policy would be to assume that agents did not have rational expectations. This contrasts somewhat with Gaspar *et al* (2006), who found that the optimal RE policy does quite a good job of replicating the optimal policy in a model where agents are doing adaptive learning.

Table C: Time-series properties of heuristic-switching model

	Policy process:	
	Rational expectations	Heuristics
var(Π)	190	163
var(var(Π))	1336	255
$\rho(\Pi)$	0.70	0.62

See note to Table A.

When agents switch between heuristics, the variance of inflation is somewhat lower when policy adapts to that, rather than following the policy that would be optimal under rational expectations. Under the policy that adapts to heuristics, the variance of small (50 period) sample estimates of the variance of inflation is about a quarter that under rational expectations optimal policy. But note that it is still more than two and a half times the figure we observe for the model under rational expectations. Note too that inflation is a little less persistent.

Chart 4 plots data from 1,000 consecutive periods of the simulation used to compute the figures in Table C. Notice that the fluctuations in n_1 , the proportion of agents using the lagged inflation heuristic, are, to the eye, as pronounced as those under the policy that would be optimal under rational expectations. Chart 5 compares the histograms for n_1 that are generated in the heuristic-switching economy when policy follows both the rational expectations optimal rule and when it adapts to the use of heuristics. Relative to the rational expectations optimal policy, the heuristics-adapted policy reduces the probability mass at both extremes of n_1 and increases it slightly at interior values.

Under both policies heuristic switching generates small-sample fluctuations in the time-series properties of inflation. We can see this from the volatility in the coefficients on lagged inflation in an autoregression for inflation. The bottom panel of Chart 4 plots the variance of the residuals from a rolling 50-period regression for inflation on its own lags. This variance is clearly moving over time and tends to be high when the variance of inflation is high, and *vice versa*. We plot this time series to link our analysis to the econometric studies that report that large fractions of recent declines in macroeconomic volatility are down to ‘good luck’.⁽¹⁸⁾ Here, very loosely, when the

(18) See, for example, Stock and Watson (2002) and Cogley and Sargent (2005).

Chart 4: Heuristic switching under a heuristics policy (first and second-order autocorrelations are regression coefficients on first and second lags of inflation)

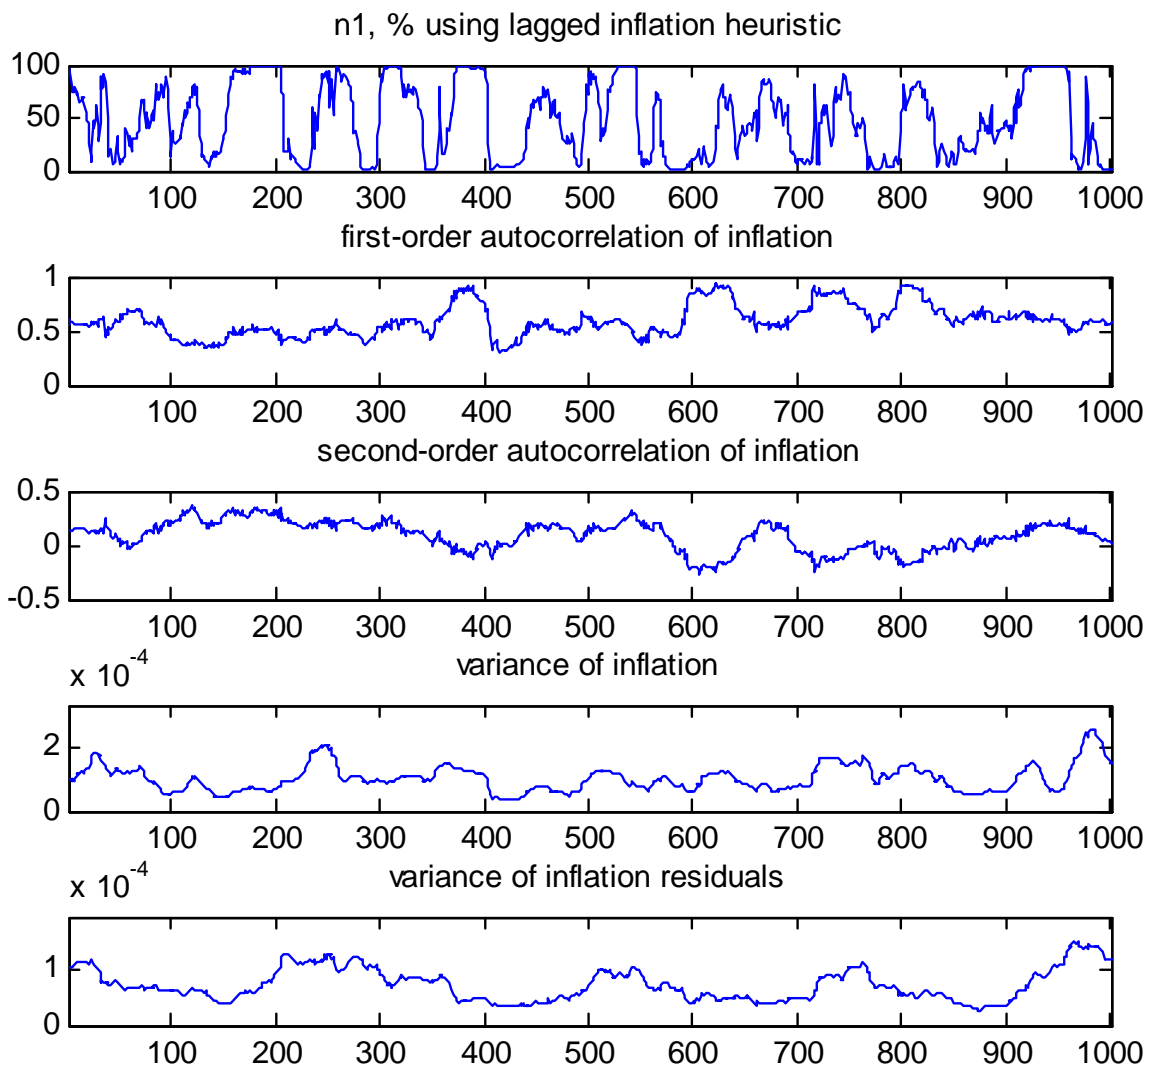
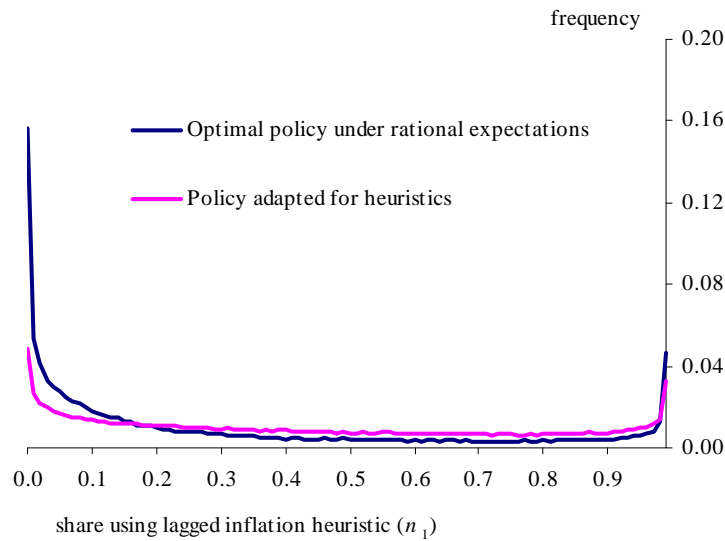


Chart 5: Share of agents using lagged inflation heuristic under alternative policy rules



number using the inflation target heuristic is low, the variance of inflation is low, and the variance of the shocks in a simple autoregression is low. In the language of the applied literature on the Great Stability, the econometrician estimates there to have been a period of good luck, when the true variance of the disturbances to our model economy is unchanging.

8 Model properties after the introduction of an inflation target heuristic

Thus far, we have investigated whether switching amongst heuristics can generate fluctuations in small-sample estimates of the volatility of inflation that are consistent with the marked reduction in volatility seen in recent decades. And our contention is that it can. These fluctuations occur regardless of whether monetary policy adopts a different rule. So far we have considered the set of heuristics as something beyond the control of policymakers. In this section, we assume that the monetary policy framework can influence the set of heuristics from which agents choose, and consider what happens when agents are given access to an inflation target heuristic, that was not previously available to them. We suggest that this may be a way of formalising what happened when many central banks adopted numerical objectives for inflation. This exercise is related to one conducted by Orphanides and Williams (2005). They interpret the introduction of a numerical objective for the central bank as equivalent to giving agents knowledge of the constant in the inflation process, knowledge that they show improves agents' estimates of the dynamics of

that process.

Table D: Impact of introducing the inflation target heuristic

Policy process:	Rational expectations		Persistent	
Heuristic:	Lagged inflation	Self confirming	Lagged inflation	Self confirming
before target:				
var(Π)	100	14	144	18
var(var(Π))	100	0	248	0
$\rho(\Pi)$	0.55	0.72	0.53	0.65
after target:				
mean(n_1)	31	42	15	69
n_1 impact	0	43	0	89
var(Π)	18	14	17	16
var(var(Π))	1	0	1	0
$\rho(\Pi)$	0.71	0.72	0.57	0.60

See note to Table A.

Table D presents simulations of the introduction of an inflation target heuristic into four different models. The four models correspond to the table columns, and comprise two different initial lagged inflation heuristics, derived under two different processes for monetary policy. Under the columns headed ‘rational expectations’ we have results that use our baseline process for money that would be optimal under rational expectations. Within this we use two heuristics. The first, ‘lagged inflation’ is our familiar lagged inflation heuristic. The column headed ‘self confirming’ refers to a model in which expectations of inflation are determined by the projection of inflation tomorrow on inflation yesterday which the model would generate as a self confirming equilibrium. Specifically, we assume that agents set $E_t(\pi_{t+1}) = \rho_h \pi_{t-1}$. We determine $E_t(\pi_{t+1})$ by the following process. First, agents collect all data to time $t - 1$ and run a regression $\pi_s = \rho_h^{ols} \pi_{s-2}$ for $s = \{1, \dots, t - 1\}$. Second, agents use ρ_h^{ols} to form $E_t(\pi_{t+1})$. Third, another data point for time t is generated. Agents add this to their data set, and return to the second step. The value of ρ_h used to compute numbers under the ‘self confirming’ column in Table D is the number to

which this iterative process converges.⁽¹⁹⁾

The two columns under ‘persistent’ repeat this analysis, but using a persistent process for money growth where the persistence and variance are set equal to the values chosen for the productivity process (and with no correlation between the two). Results are computed from a 20,000-period simulation in which the target is introduced half way through.

We report several details. First, in the top rows, we give statistics for the economy before the introduction of the inflation target heuristic into the set of heuristics from which agents choose. These are: the variance of inflation (row labelled ‘var(Π)’); the variance of short-sample estimates of that variance (‘var(var(Π))’); and the persistence of inflation (‘ $\rho(\Pi)$ ’). For the second half of the simulation, after the introduction of the inflation target, we report these same statistics, but with two additions. First, we report the average value of n_1 in the five periods immediately following the introduction of the target, and label this row ‘ n_1 impact’. Second, we report the mean of n_1 over the life of the rest of the simulation (labelled ‘mean(n_1)’). In this table, we normalise variances and the variance of variances relative to those computed for the top left-hand case in this table, the case where agents have a single, simple lagged inflation heuristic, and policy is conducted according to the rule that would be optimal under rational expectations.

The basic message is that the immediate impact effect of the introduction of the inflation target heuristic is maximal when, prior to that, agents use only the lagged inflation heuristic. In both the ‘lagged inflation’ simulations, n_1 , the number using the lagged inflation heuristic, drops to zero in the period immediately following introduction of the inflation target (albeit rising again thereafter). This is shown by the zeros recorded in the row labelled ‘ n_1 impact’. It turns out that in our model, if we exogenously impose that $n_1 = 1$, it greatly worsens the forecast performance of that heuristic, which is why when agents are free to choose between two heuristics, they jump to using the inflation target for a while.

This begs the question of why agents were content to use only the lagged inflation heuristic prior to the introduction of the target. It is beyond the scope of this paper to model the complete process that specifies the evolution of the set of heuristics that agents use. But for comparison, we have the simulations where agents start out life using a lagged inflation heuristic based on an

(19) The point to which this iteration converges might also be referred to as a restricted perceptions equilibrium. Subject to the restricted perceptions of the inflation process that agents have, their projections are optimal.

optimal projection of inflation tomorrow on inflation yesterday (the ‘self confirming’ simulations). Able to use such a projection, one which performs better than the simple lagged inflation heuristic, the effect of the new target heuristic is more muted: this is true under both our ‘rational expectations’ and ‘persistent’ monetary policy processes.

Similarly, we see that when agents are constrained to use the simple lagged inflation heuristic, the introduction of the inflation target has its largest effect on the time-series properties of inflation, reducing the variability of inflation and the fluctuations in small-sample estimates of this variability.⁽²⁰⁾

9 Conclusions

In the past decade both inflation and output growth seem to have become more stable in advanced economies. This coincided with the convergence of inflation expectations on inflation targets. We have illustrated how an economy populated by agents who choose amongst heuristics for forecasting inflation can generate fluctuations in the variance of inflation. There are periods in which agents use the inflation target heuristic, and there are periods when many agents choose to use a heuristic based on lagged inflation. In the former, a given shock will generate less variability in inflation. But a sequence of shocks that reduces the ability of the inflation target heuristic to match inflation in the past can lead agents to switch to the lagged inflation heuristic.

We asked how monetary policy might adapt to agents’ use of heuristics. Under rational expectations, a rule for money growth that responded to productivity could stabilise completely labour supply and consumption. It did so through its leverage over expectations. When agents use heuristics, monetary policy has no direct leverage over inflation expectations, which are determined entirely by the past behaviour of inflation. Relative to the policy that would be optimal under rational expectations, a money growth rule which reacts to both productivity and to inflation expectations can better stabilise the economy. Even under such a policy, agents switch back and forth between heuristics and the time-series properties of inflation tend to fluctuate.

Our final exercise was to simulate the introduction of an inflation target heuristic. When we did

(20) We repeated the simulation many times and found that the main determinant of the impact effect was the assumption about the heuristic that agents used before the introduction of the inflation target. This was more important than, for example, the recent history of productivity shocks in the periods preceding the target introduction.

this, there was some evidence that the introduction of this heuristic improves macroeconomic outcomes by reducing the volatility of inflation. By how much, and to what extent agents use the new heuristic depends on the performance of the heuristics they had before. These results suggest that some of the improvements seen in the United Kingdom and elsewhere could be locked in, at least if the inflation targeting regime can be thought of as having made available the simple heuristic that ‘inflation will equal the target’.

References

- Benati, L (2004)**, ‘UK monetary regimes and macroeconomic stylised facts’, *Bank of England Working Paper no. 290*.
- Bernanke, B (2004)**, ‘The Great Moderation’, remarks delivered at the meetings of the Eastern Economic Association, Washington, DC, 20 February.
- Branch, W and Evans, G (2005)**, ‘Model uncertainty and endogenous volatility’, unpublished manuscript.
- Branch, W and Evans, G (2006)**, ‘Intrinsic heterogeneity in expectations formation’, *Journal of Economic Theory*, Vol. 127(1), pages 264–95.
- Brock, W and de Fontnouvelle, P (2000)**, ‘Expectational diversity in monetary economies’, *Journal of Economic Dynamics and Control*, Vol. 24, pages 725–59.
- Brock, W and Hommes, C (1997)**, ‘A rational route to randomness’, *Econometrica*, Vol. 65, pages 1,059–95.
- Cogley, T and Sargent, T J (2005)**, ‘Drifts and volatilities: monetary policies and outcomes in the post WWII US’, *Review of Economic Dynamics*, Vol. 8(2), pages 262–302.
- Evans, G and Honkapohja, S (2003)**, ‘Adaptive learning and monetary policy design’, *Journal of Money, Credit and Banking*, Vol. 35, pages 1,047–72.
- Gaspar, V, Smets, F and Vestin, D (2006)**, ‘Adaptive learning, persistence and optimal monetary policy’, *Journal of the European Economic Association*, Papers and Proceedings, May.
- Gigerenzer, G, Todd, P and the ABC Research Group (1999)**, *Simple heuristics that make us smart*, New York, Oxford University Press.
- de Grauwe, P and Grimaldi, M (2006)**, ‘Exchange rate puzzles: a tale of switching attractors’, *European Economic Review*, Vol. 50, pages 1–33.
- Hommes, C (2005)**, ‘Heterogeneous agent models in economics and finance’, in Judd, K and Tesfatsion, L (eds), *Handbook of computational economics (Volume 2)*, Elsevier.
- Kahneman, D (2003)**, ‘Maps of bounded rationality: psychology for behavioural economics’, *American Economic Review*, Vol. 93, pages 1,449–75.
- Keynes, J M (1936)**, *The general theory of employment, interest and money*, New York, MacMillan.
- King, M (2003)**, ‘The Governor’s speech at the East Midlands Development Agency/Bank of England dinner’, *Bank of England Quarterly Bulletin*, Winter, pages 476–78.

- Levin, A and Piger, J (2004)**, ‘Is inflation persistence intrinsic in industrial economies?’, *ECB Working Paper no. 334*.
- Manski, C and McFadden, D (eds) (1981)**, *Structural analysis of discrete data with econometric applications*, MIT press, Cambridge.
- Marcet, A and Nicolini, J P (2003)**, ‘Recurrent hyperinflations and learning’, *American Economic Review*, Vol. 93(5), pages 1,476–98.
- Mojon, B and Diron, M (2005)**, ‘Forecasting the central bank’s objective is a good rule of thumb’, *ECB Working Paper no. 564*.
- Orphanides, A and Williams, J (2005)**, ‘Imperfect knowledge, inflation expectations, and monetary policy’, in Bernanke, B and Woodford, M (eds), *The inflation-targeting debate*, Chicago: University of Chicago Press, pages 201–34.
- Simon, H (1957)**, *Models of man*, Wiley, New York.
- Sims, C and Zha, T (2004)**, ‘Were there regime switches in US monetary policy?’, *Federal Reserve Bank of Atlanta Working Paper no. 14*.
- Stock, J and Watson, M (2002)**, ‘Has the business cycle changed and why?’, in Gertler, M and Rogoff, K (eds), *NBER Macroeconomics Annual 2002*, MIT Press.
- Velde, F (2004)**, ‘Poor hand or poor play? The rise and fall of inflation in the United States’, *Economic Perspectives*, Vol. 28(1), Federal Reserve Bank of Chicago.