

THE LONDON SCHOOL OF ECONOMICS
AND POLITICAL SCIENCE

Learning, Monetary Policy and Asset Prices

ALBERTO LOCARNO

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy.

London, April 2012

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work, save where it is clearly indicated that it is the work of others, and that the extent of any work carried out jointly by me and any other person is clearly specified.

Copyright is the author's. Quotation from the thesis is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without the prior written consent of the author.

I warrant that this authorization does not, to the best of my knowledge, infringe the rights of any third party.

I declare that my thesis consists of (about) 37670 words.

Statement of conjoint work

The chapter “Monetary policy uncertainty and the stock market” was written jointly with Professor Massimo Massa (INSEAD). I wrote the empirical sections and most of the survey of the literature, Professor Massa set up the theoretical model; the introduction and conclusions were written jointly. My own share of the work for this chapter can be put at 55 per cent; the remainder is to be attributed to Professor Massa.

Abstract

The dissertation examines several policy-related implications of relaxing the assumption that economic agents are guided by rational expectations. A first, introductory chapter presents the main technical issues related to adaptive learning. The second chapter studies the implications for monetary policy of positing that both the private sector and the central bank form their expectations through adaptive learning and that the central bank has private information on shocks to the economy but cannot credibly commit. The main finding of this chapter is that when agents learn adaptively a bias against activist policy arises. The following chapter focuses on large, non-linear models, where no unambiguous linear approximation eligible as perceived law of motion exists. Accordingly, there are heterogeneous expectations and the system converges to a misspecification equilibrium, affected by the communication strategies of the central bank. The main results are: (1) the heterogeneity of expectations persists even when a large number of observations are available; (2) the monetary policymaker has no incentive to be an inflation hawk; (3) partial transparency enhances welfare somewhat but full transparency does not. The final chapter adopts a model in which agents are fully informed and use Bayesian techniques to estimate the hidden states of the economy. The monetary policy stance is unobservable and state-independent, generating uncertainty among agents, who try to gauge it from inflation: a change in consumer prices that confirms beliefs reduces stock risk premia, while a change that contradicts beliefs drives the risk premia upward. This may generate a negative correlation between returns and inflation that explains the Fisher puzzle. The model is tested on US data. The econometric evidence suggests: (1) that a mimicking-portfolio proxying for monetary policy uncertainty is a risk factor priced by financial markets; and (2) that conditioning on monetary uncertainty and fundamentals eliminates the Fisher puzzle.

Acknowledgements

I thank Dr. Margaret Bray for her guidance in the preparation of this thesis. I also thank prof. Charles Goodhart, who assisted me in the initial stages of the research, and two anonymous referees of the International Journal of Central Banking. I benefitted from discussing the chapters of the thesis with Fabio Buseti, Giuseppe Ferrero and Claudia Miani at the Banca d'Italia and with the participants to the Learning Week 2008 at the Federal Reserve Bank of St. Louis and to the 2011 Oslo Annual Congress of the European Economic Association. I am responsible for all errors. The views expressed in this thesis are those of the author and do not involve the responsibility of the Banca d'Italia.

Contents

1	Three essays on learning and monetary policy	9
1.1	Determinacy and stability under learning	10
1.2	Learnability as a pre-requisite for full rationality	13
1.3	Learnability as an equilibrium selection device	14
1.4	Learning in misspecified models	16
1.5	Heterogeneous learning	17
1.6	Learning and evolutionary dynamics	18
1.7	Speed of convergence, transitional dynamics and perpetual learning	19
1.8	Escape dynamics	21
1.9	Learning as an alternative to rational expectations	22
1.10	A summary review of the dissertation	23
2	Adaptive learning and the bias against activist monetary policies	34
2.1	Introduction	34
2.2	The model	37
2.2.1	The structure of the economy and the delegation problem	37
2.2.2	The central bank's loss function	39
2.2.3	Signal extraction and the rational expectations equilibrium	41
2.3	Adaptive learning and monetary policy regimes	46
2.3.1	Private sector learning	48
2.3.2	Private sector and central bank learning	51
2.4	Imperfect knowledge and policy effectiveness	56
2.5	Conclusions	63
2.6	Appendix	64
3	Monetary policy with misspecified, heterogeneous and ever-changing expectations	78
3.1	Introduction and motivation	78

3.2	The literature	82
3.3	The model	87
3.3.1	The learning mechanism	90
3.3.2	Heterogeneous expectations and predictor selection	92
3.3.3	The role of expectations	95
3.4	Simulation results	102
3.4.1	Optimal monetary policy under rational expectations	103
3.4.2	Optimal monetary policy under learning	104
3.4.3	Perpetual learning	115
3.5	Sensitivity analysis	116
3.5.1	Experiment #1: the number of replications	117
3.5.2	Experiment #2: the size of Γ	118
3.5.3	Experiment #3: specification of the welfare function	118
3.5.4	Experiment #4: stochastic initial conditions for model proportions x_t^i	120
3.6	Conclusions	121
4	Monetary policy uncertainty and the stock market	138
4.1	Introduction	138
4.2	The literature	140
4.3	A simple model	145
4.3.1	Monetary policy regimes	146
4.3.2	Investors' behavior and stock returns	147
4.4	The role of monetary policy uncertainty	149
4.4.1	Pricing implications	149
4.4.2	New Perspective on the Fisher Puzzle	153
4.5	A proxy for monetary policy uncertainty	155
4.5.1	Monetary policy regimes	156
4.5.2	A market-based measure of monetary policy uncertainty	158
4.6	Evidence on the market price of policy uncertainty	160
4.6.1	Inflation and stock returns: the Fisher puzzle	160
4.6.2	Evidence of Pricing	161
4.6.3	Inflation and stock returns: the Fisher relationship revived	163
4.7	Conclusion	165
4.8	Appendix	166

List of Tables and Figures

Chapter 1

Chapter 2

Table 2a - Least squares learning and the volatility of output and inflation (<i>unconstrained estimator - decreasing gain sequence</i>)	72
Table 2b - Least squares learning and the volatility of output and inflation (<i>unconstrained estimator - constant gain sequence</i>)	73
Table 3a - Least squares learning and the volatility of output and inflation (<i>constrained estimator - decreasing gain sequence</i>)	74
Table 3b - Least squares learning and the volatility of output and inflation (<i>constrained estimator - constant gain sequence</i>)	75
Table 4a - Dynamic response to contractionary shocks (<i>decreasing gain sequence</i>) ...	76
Table 4b - Dynamic response to contractionary shocks (<i>constant gain sequence</i>)	77

Chapter 3

Figure 1 - Social welfare under rational expectations	123
Figure 2 - Predictor proportions and payoff function	124
Table 1 - Impact on welfare of the volatility of the main macroeconomic variables .	125
Table 2a - Mon. Pol. effectiveness under no transparency (<i>decreasing gain learning</i>)	126
Table 2b - Mon. Pol. effectiveness with transparency (<i>decreasing gain learning</i>) ...	127
Table 2c - Mon. Pol. effectiveness with full transparency (<i>decreasing gain learning</i>)	128
Table 3a - Mon. Pol. effectiveness under no transparency (<i>constant gain learning</i>) .	129
Table 3b - Mon. Pol. effectiveness with transparency (<i>constant gain learning</i>)	130
Table 3c - Mon. Pol. effectiveness with full transparency (<i>constant gain learning</i>) .	131
Table 4 - Sensitivity analysis: number of replications (<i>decreasing gain learning</i>) ...	132
Table 5 - Sensitivity analysis: initial conditions (<i>constant gain learning</i>)	133

Table 6a - Sensitivity analysis: specification of the welfare function (<i>no transp.</i>) ...	134
Table 6b - Sensitivity analysis: specification of the welfare function (<i>partial transp.</i>)	135
Table 6c - Sensitivity analysis: specification of the welfare function (<i>full transp.</i>) ..	136
Table 7 - Sensitivity analysis: impact of initial conditions on predictor proportions	137

Chapter 4

Figure 1 - Excess Returns and the value of the monetary policy parameter δ	175
Figure 2 - Excess Returns and the share of dividends over consumption	176
Figure 3 - Perceived regimes of monetary policy and uncertainty	177
Table 1 - Probability of the Monetary Policy Regime and Business Cycle	178
Table 2 - Economic tracking portfolio for I_t : regression coeff. and portfolio weights	179
Table 3 - Summary Financial Statistics	180
Table 4 - Expected Inflation and Excess Returns	181
Table 5 - Evidence of pricing	182
Table 6 - Expected Inflation and monetary uncertainty	185

Chapter 1

Three essays on learning and monetary policy

The concept of rational expectations (RE) rests on two pillars: individual rationality and mutual consistency of perceptions about the environment. The true stochastic process of the economy is assumed known, with unpredictable random shocks as the only source of uncertainty. RE are regarded by most researchers as the most appropriate hypothesis for economic analysis, since one necessary condition for optimisation is that individuals eliminate any systematically erroneous component of their behaviour, including in the formation of expectations; further, from a policy perspective, this assumption rules out policies designed to exploit patterns of suboptimal expectations.¹ In other respects, however, the RE hypothesis is unappealing, since it clashes with the principle of cognitive consistency, as it implies that agents within the model are much smarter and have much more knowledge than economists/econometricians, faced with problems of estimation and inference.² Besides, in most situations, there is no sufficient incentive to upgrade from bounded to unbounded rationality, since the costs may be consider-

¹See McCallum (2008) on this point.

²According to the definition in Evans and Honkapohja (2008), the principle of cognitive consistency is the requirement that private agents and policymakers in the economy behave like applied economists and econometricians.

able, even astronomical, while the benefits tend to be small. Finally, the RE assumption begs a crucial question that becomes vital in the presence of structural or policy changes: how is it that economic players can have rational expectations, if they do not initially know the exact nature of the equilibrium in which they find themselves?

For these reasons, macroeconomic theorists have been gradually moving beyond the strict RE framework to develop models in which agents have imperfect information and use simple and misspecified forecasting equations to form expectations. Though people's rationality is bounded, they learn gradually through practice; in their efforts to improve their knowledge of the stochastic process of the economy, agents adjust their model in the course of time, as new information becomes available, in the end acting as if they were unboundedly rational.

Three main approaches have been taken to modelling learning: the eductive learning approach assumes that agents engage in a process of reasoning – taking place in logical time – about the possible outcomes, knowing that other agents engage in the same process; the rational learning approach replaces full knowledge of economic parameters with priors and Bayesian updating under a correctly specified model, including common knowledge that all agents share this knowledge; and the adaptive learning approach, the most common, views agents as econometricians, who adjust their model over time as information becomes available, re-estimating the parameters of their perceived law of motion.

1.1 Determinacy and stability under learning

A model with adaptive learning has two main ingredients: (i) an equation describing agents' beliefs on the dynamics of economic variables and (ii) a temporary equilibrium of the system generated by the interaction between

expectations and the structure of the economy.

Let us assume that the economy is described by the following linear system:

$$\begin{aligned} y_t &= A_1 y_{t-1} + A_2 E_{t-1} y_{t+1} + B x_{t-1} + u_t \\ x_t &= F x_{t-1} + v_t \end{aligned} \quad (1.1)$$

where y_t and x_t are the vectors of endogenous and exogenous variables respectively. Agents' beliefs are described by a forecasting model, the so-called perceived law of motion (PLM), which usually has the same functional form as the (minimum state variable) solution of the RE equilibrium:

$$\hat{E}_{t-1} y_t = a_{1,t-1} y_{t-1} + b_{t-1} x_{t-1} \quad (1.2)$$

The operator \hat{E} refers to subjective beliefs, which may vary across individuals, and does not coincide with conditional expectations. The coefficients of the forecasting model are re-estimated in every period by recursive least squares (RLS).³ The learning process is described by the following set of recursive equations:

$$\begin{aligned} \theta_t &= \theta_{t-1} + \gamma_t R_t^{-1} z_{t-1} (y_t - \theta_{t-1}^T z_{t-1}) \\ R_t &= R_{t-1} + \gamma_t (z_{t-1} z_{t-1}^T - R_{t-1}) \end{aligned} \quad (1.3)$$

where the gain sequence $\{\gamma_t\}_{t=k}^{\infty}$ is equal to $\{t^{-1}\}_{t=k}^{\infty}$, and $\theta_t = (a_{1,t}, b_t)^T$ and $z_t = (y_t^T, x_t^T)^T$. Given the forecasts, the economy attains a temporary equilibrium, the so-called actual law of motion (ALM), which is equal to:

$$\begin{aligned} y_t &= A_1 y_{t-1} + A_2 \left(a_{1,t-1} \hat{E}_{t-1} y_t + b_{t-1} \hat{E}_{t-1} x_t \right) + B x_{t-1} + u_t \\ &= (A_1 + A_2 a_{1,t-1}^2) y_{t-1} + (B + A_2 a_{1,t-1} b_{t-1} + A_2 b_{t-1} F) x_{t-1} + u_t \\ &= T(\theta_{t-1})^T z_{t-1} + u_t \end{aligned} \quad (1.4)$$

³Econometric learning can be alternatively modelled using a (generalised) stochastic gradient (SG) updating rule. Equation (3) modifies to $\theta_t = \theta_{t-1} + \frac{1}{t} \Gamma z_{t-1} (y_t - \theta_{t-1}^T z_{t-1})$, where the time-invariant matrix Γ is usually set equal to either the identity matrix or $E(z_t z_t^T)^{-1}$.

where $T(\theta_{t-1})^T = T(a_{1,t-1}, b_{t-1}) = (A_1 + A_2 a_{1,t-1}^2, B + A_2 a_{1,t-1} b_{t-1} + A_2 b_{t-1} F)$ is the mapping that describes the evolution of the RLS estimator θ_t . Once the ALM is substituted for y_t in (3), the dynamics of the system is fully described by the recursive least squares equations:

$$\begin{aligned}\theta_t &= \theta_{t-1} + \frac{1}{t} R_t^{-1} z_{t-1} z_{t-1}^T (T(\theta_{t-1}) + u_t - \theta_{t-1}) \\ R_t &= R_{t-1} + \frac{1}{t} (z_{t-1} z_{t-1}^T - R_{t-1})\end{aligned}\tag{1.5}$$

With the shift $S_{t-1} = R_t$, (5) becomes a stochastic recursive algorithm (SRA), whose behaviour is well approximated by an ordinary differential equation (ODE)

$$\frac{d\phi}{d\tau} = h(\phi) = T(\phi) - \phi\tag{1.6}$$

where $\phi_t \equiv \text{vec}(\theta_t, S_t)$ and $h(\phi)$ is obtained by computing the asymptotic limit of the expectation of the 2nd term (the updating function) on the right-hand side of (5): the zeros of the ODE represent the only possible limit points of the SRA and the corresponding equilibria are stable if the (real part of the) eigenvalues of the Jacobian of $h(\phi)$ are negative. When $(a_{1,t}, b_t) \rightarrow (A_1, B)$, i.e. when the PLM comes to coincide with the ALM, an RE equilibrium is attained and agents have learnt the rational expectations equilibrium.

Two features of the model stand out: it is self-referential and the agents are not fully rational. “Self-referential” means that when individuals learn adaptively, the dependency between outcomes and beliefs is bidirectional, since the expectations that drive the temporary equilibrium change as new observations become available. This property affects the law of motion of the variables, which becomes non-stationary and keeps on changing until agents’ subjective beliefs eventually converge to the objective distribution of the variables. The absence of full rationality follows from the assumption that in estimating their forecasting model agents treat the economy as having constant parameters. But this is true only in the RE equilibrium; outside it, the PLM is misspecified, though the bias may eventually vanish, when the

PLM nests the RE solution and the estimates converge asymptotically.

1.2 Learnability as a pre-requisite for full rationality

Academic interest in learning was originally prompted by the idea that it might justify the RE hypothesis. An equilibrium cannot be produced out of thin air; there must be forces at work that propel the economy toward it. How is an RE equilibrium achieved? How can agents eventually become fully rational, when at first they are not? The early answers to these questions were hardly encouraging. Frydman (1982) proposed a proof of the impossibility of rational learning when individuals cannot determine the average of other agents' forecasts; DeCanio (1979) claimed that "... direct computation of rational expectations by flesh and blood agents in an actual market situation is impossible in practice": he acknowledged that an evolutionary learning procedure could lead to rationality, but was sceptical about its practical relevance. However, the tide was slowly turning. First, an influential paper by Bray (1982) showed how to prove stability under learning of the RE equilibrium in an asset market model; then Evans (1985) introduced the notional time concept of expectational stability. Finally, Marcet and Sargent (1989a,b) proposed to adopt, as a plausible learning concept, recursive least squares, and showed how stochastic approximation theory could be applied to prove the learnability of the RE equilibria. Their findings were subsequently revised and extended by Evans and Honkapohja, who present an exhaustive survey of the work on recursive learning in their 2001 book. The main message of the literature today is that many macroeconomic models posit that the RE equilibrium is learnable, provided that (i) all agents use the same mechanical learning rule to form expectations and (ii) the forecasting model is well-specified. However, as Bullard (1997) noted, pre-coordination on the learning rule simply pushes the issue of how agents can ultimately share the

same expectations-formation mechanism one step further back. McCallum (2007) makes the same point, reaffirming the opinion voiced by Lucas (1980):⁴ after a structural change, reliable analysis should pertain to the economy's behaviour after it has settled into a new dynamic stochastic equilibrium.

1.3 Learnability as an equilibrium selection device

In addition to assessing the plausibility of RE equilibrium, learning can serve as a selection criterion when models have multiple solutions. As an example, Evans and Honkapohja (2001) considered the non-stochastic Cagan model with government spending financed by seignorage; the model has two steady-state solutions, one with low inflation and one with high, and only the first equilibrium is learnable. In more general models, learning does not necessarily select a unique RE equilibrium, but the set of plausible solutions is usually significantly smaller than the set of all solutions.

A more complex problem is faced when an infinite number of solution paths converge on a single steady-state equilibrium, i.e. when there is indeterminacy. Unfortunately, there is no strict relation between learnability and determinacy of RE equilibria. McCallum (2007, 2008) showed that determinacy is a sufficient – but not a necessary – condition for learnability if agents can use current endogenous variables in their forecasting equations; if there are information lags, then the connection is severed and learnability can be achieved only under special assumptions. The same findings are presented in Bullard and Eusepi (2008).

Since E-stability does not imply determinacy, there must exist models where agents with bounded rationality can learn an indeterminate RE equilibrium. Indeed, in Evans and Honkapohja (2001) it is shown that the Taylor real-balance model can be at once indeterminate and E-stable. However, indeter-

⁴McCallum uses the expression “startup problem” to indicate this theoretical ambiguity.

minacy can imply the existence not only of multiple fundamental equilibria but also of non-fundamental sunspot solutions; for these latter to be learnable, agents would need to coordinate on a PLM that includes a variable that has no direct effect on the economy but becomes a driving force of the equilibrium outcome solely because agents believe it matters. The practical importance of this issue was suggested by Clarida et al. (2000), who claimed that the volatile inflation and output of the 1970s may have been due to sunspot phenomena.

On theoretical grounds, Woodford (1990) found that, in an overlapping-generations model, finite-state Markov sunspots may be stable under learning. Subsequent analyses, however, failed to find supporting evidence for Woodford's results. Honkapohja and Mitra (2004) studied learnability of sunspot equilibria in New-Keynesian models, where the prime source of indeterminacy is violation of the Taylor principle. They found that the private sector cannot coordinate on any of the non-fundamental equilibria induced by a passive monetary policy. At first sight, the non-repeatability of Woodford's results appeared to be due to the ad-hoc nature of the stochastic process assumed for the sunspot variable, suggesting that a sunspot solution with continuous support would never be stable under learning. Yet this notion does not correctly characterise the situation: Evans and McGough (2005) showed for a model with sunspots that a given equilibrium may have several consistent representations and that stability under learning is representation-dependent. In a later paper,⁵ they demonstrated that Woodford's findings are more general than is normally thought and apply to a large set of models, provided that the forecasting equations are expressed in terms of a common factor representation: whenever finite-state Markov sunspots are stable under learning, all sunspot equilibria will be stable.

⁵See Evans and McGough (2008).

1.4 Learning in misspecified models

What if the PLM is not correctly specified?⁶ Adaptive learning may be a more realistic way of modelling expectations than RE, but to posit that everyone selects the correct specification is quite far removed from reality: econometricians cannot include all lags and exogenous variables in their models, due to lack of degrees of freedom, informational uncertainty, limited computing skills or processing costs. Accordingly, it may be useful in modelling learning to move further away from rationality by assuming that agents may choose to form expectations using possibly misspecified forecasting models. If the PLM does not include variables that are relevant for the dynamics of the system, then the learning process cannot possibly converge on the RE solution, but it may nevertheless achieve a different equilibrium, which may be called a *restricted perceptions equilibrium* (RPE).⁷ In an RPE the ALM does not belong to the same class of function as the PLM and for the ODE to be properly defined, the ALM must be projected onto the space spanned by the PLM. The conditions for E-stability are therefore different, but as Guse (2008) showed for the case of the New Keynesian IS-AS model, they are not in general more restrictive.

In an RPE agents are permitted to fall short of rationality in failing to recognise certain patterns or correlations in the data. Hommes and Sorger (1998) proposed a variant of the RPE, called a consistent expectations equilibrium

⁶Here misspecified means that the PLM is underparameterised and does not nest the RE solution, which precludes the convergence of the model on the RE equilibrium. An overparameterised forecasting rule involves additional complexities, since the steady-state equilibrium is unaffected and the conditions for E-stability, though usually different, are not necessarily more restrictive than those prevailing under correct specification. An RE equilibrium is said to be strongly E-stable if it is locally E-stable even for a specified class of overparameterised PLMs.

⁷Not only may the RPE be different from the RE equilibrium, but the number of equilibria too may change. Evans and Honkapohja (2001, chapter 13) present a model featuring two equilibria when the PLM is correctly specified and only one when the forecasting equation does not include the lagged endogenous variable among the regressors.

(CEE), which requires that the sample average and the sample autocorrelations of the realisations of the ALM and of the PLM be equal.

1.5 Heterogeneous learning

Allowing for PLMs that differ from the MSV solution is the preliminary step to study heterogeneous expectations. It turns out that whether or not the presence of heterogeneous expectations matters depends on whether the model exhibits structural heterogeneity.⁸ Giannitsarou (2003) studied three types of heterogeneity. Agents may (i) have different priors (initial perceptions), (ii) have different degrees of inertia in updating or (iii) follow different learning rules.⁹ She found that the conditions for local convergence of heterogeneous and homogeneous learning are always identical when individuals have different priors, but the conditions for general convergence differ. However, it turns out that the representative learner may well be a good approximation of the population of the economy. Giannitsarou's results were confirmed by Honkapohja and Mitra (2006), who also found that when permanent divergences in expectations formation are combined with structural heterogeneity, the conditions for convergence of learning become significantly more stringent and instability can arise. Another natural way to introduce heterogeneity in the learning process is to assume that different agents have different types of forecasting models. Berardi (2007) considered two groups of individuals, one with the correct PLM and the other with an underparameterised model. He showed that the second group cannot possibly learn the

⁸Structural heterogeneity is present when the expectations of different groups affect the economy in different ways. It occurs for instance in the New Keynesian model: private expectations affect the economy through the parameters of the IS equation and the Phillips curve, while the central bank affects the equilibrium outcome through the policy rule.

⁹In terms of equation (3), (i) implies that θ_0 is different across (classes of) agents, (ii) implies that the gain sequence is $\delta_i t^{-1}$, with δ_i varying across individuals and (iii) implies that some people use RLS as a learning algorithm, while others adopt a (generalised) SG updating rule.

true process that drives the endogenous variables, so the RE equilibrium can never be achieved. Nevertheless, the system can still converge on an equilibrium in which all beliefs are confirmed by the data: the E-stability conditions are different, but not necessarily more restrictive than under homogeneous expectations.

1.6 Learning and evolutionary dynamics

A limitation of Berardi's analysis is that he does not allow for time variation in the share of people adopting the correct forecasting equation or for model switching. By contrast, Brock and Hommes (1997) consider dynamic predictor selection: in their model, agents adapt their beliefs over time by choosing from a finite set of different expectation functions on the basis of costs and of a measure of fit that is publicly available. The proportion of agents using predictor j is given by the multinomial logit ratio

$$n_{j,t} = \frac{\exp(\beta\pi_{j,t})}{\sum_k \exp(\beta\pi_{k,t})}$$

where β measures the intensity with which agents choose predictors with better fit and $\pi_{j,t}$ represents model- j goodness of fit. Brock and Hommes found that high values of β can lead to high-order cycles and chaotic dynamics. The rationale is straightforward: when agents use cheaper and less accurate predictors, the steady-state equilibrium is unstable, whereas the costly, sophisticated forecasting models are stabilising. Near the steady state it pays to use the cheap predictors, but this pushes the economy away from the steady state. For high enough intensity of choice, this tension leads to local instability and complex overall dynamics. Branch and Evans (2006), working on a similar model, found very different results. Assuming that agents use predictive performance to choose from a set of costless, misspecified econometric models, they obtained conditions under which there is an equilibrium

with agents heterogeneously split among the misspecified models even as the intensity of choice becomes arbitrarily large. This finding must be treated with caution, however, since it depends heavily on the criterion function used to discriminate across models: as is shown in Waters (2007), when is finite the multinomial logit never excludes predictors that have uniformly inferior performances.

1.7 Speed of convergence, transitional dynamics and perpetual learning

Unlike the asymptotic properties, the transitional dynamics of learning processes has not attracted much interest. Among the few analytical studies of the transition to the RE equilibrium is that of Benveniste et al. (1990), who related the speed of learning convergence to the eigenvalues of the Jacobian of the associated ODE and derived the conditions for root- t convergence of the parameters of the forecasting equation.¹⁰ Marcet and Sargent (1995) subsequently suggested a simple numerical procedure, based on model simulations, to estimate the rate at which the PLM approaches the ALM. As stressed by Bullard (2006), under RE, once a determinate equilibrium is shown to exist nothing else really matters, whereas with learning anything that affects how fast the private sector learns the RE equilibrium may affect social welfare. This issue was studied by Ferrero (2007), who showed that learning times depend upon the Taylor rule parameters.

In addition to a few regularity conditions on the exogenous processes and on the updating function, asymptotic convergence of the learning algorithm depends on the (positive and non-stochastic) gain sequence $\{\gamma_t\}_{t=k}^{\infty}$ being

¹⁰Root- t is the speed at which, in classical econometrics, the mean of the distribution of the least square estimator approaches the asymptotic value; under root- t convergence, the effects of initial conditions die out at an exponential rate.

such that $\sum_{t=k}^{\infty} \gamma_t = \infty$ and $\sum_{t=k}^{\infty} \gamma_t^2 < \infty$. These assumptions are necessary to avoid convergence of ϕ_t to a non-equilibrium point and to ensure the asymptotic elimination of all residual fluctuations.

Some authors have departed from this framework, studying the implications for the equilibrium outcomes of a constant gain sequence, i.e. of setting $\gamma_t = \gamma$ for all values of t . In theory constant gain learning precludes the convergence on the RE equilibrium: as long as the solution is stable, agents' expectations are correct on average but keeps on fluctuating around rather than at the equilibrium. This happens because observations are not assigned equal weight: those far in the past are discounted at an exponential rate, so that information does not accumulate fast enough to completely remove the randomness in the data.¹¹ A constant γ is justified when agents suspect that the economy is undergoing structural changes. Although in principle they might attempt to model structural change, this would call for an amount of knowledge comparable to that needed for RE; a reasonable alternative is to recognise, in adjusting the parameter estimates, that the more recent observations convey more accurate information on the economy's laws of motion.

Most of the studies that assume permanent learning relate to monetary policy. Orphanides and Williams (2005) analysed the impact of constant-gain learning on the effectiveness of central bank strategies. They worked with a two-equation system consisting of a modified Lucas supply curve and an aggregate demand relation, supplemented by a loss function describing the policymaker's preferences. Their analysis produced four main conclusions: first, the "naïve" choice, i.e. the policy that assumes RE on the part of agents, can be highly inefficient, generating higher volatility of both inflation and output; second, learning leads to a bias towards more "hawkish"

¹¹Orphanides and Williams refer to constant gain learning as *perpetual learning*, to stress the fact that full information about the structure of the economy is never achieved.

policies; third, persistent deviations of inflation expectations from target can arise following a sequence of unfavourable shocks; and fourth, if the inflation target is credibly announced to the public, the policy frontier is more favourable. According to these authors, what alters the policy response is the non-linear nature of the learning process. Since the forecasting model is estimated recursively, a positive price shock passes through to the intercept of the PLM, raising both expected and hence actual inflation in the next period. Unless the policy response is prompt and forceful, the persistence and volatility of inflation increase and the monetary authority fails to maintain a firm grip on the value of money. In Orphanides and Williams (2007), the authors considered the case of uncertainty about the natural rates of interest and unemployment and assessed not only the efficiency but also the robustness of alternative Taylor-type monetary rules. They confirmed their previous findings and further showed that the best policies are characterised by two features: aggressive response to inflation and a high degree of inertia. Indeed, difference rules, which disregard natural rates, appear to be robust to misspecification of private sector learning and to the magnitude of variation in the natural rates of interest and unemployment.

1.8 Escape dynamics

Closely related to permanent learning is the notion of “escape” dynamics, popularised in Thomas Sargent’s book on American inflation (1999), which constructed a model designed to study the implications for the equilibrium outcome of positing that the government believes that there is a non-vertical Phillips curve, when the real economy is actually governed by neo-classical natural rates. In Sargent’s model, policymakers estimate the unemployment-inflation trade-off using a constant-gain RLS algorithm, which eventually drives the economy to a particular kind of imperfect RE solution, called

the self-confirming equilibrium (SCE), where misspecified beliefs and realisations are mutually consistent. Notably, Sargent found that periodically the economy would escape from the SCE, settling for a prolonged period on a low-inflation equilibrium path. The escape dynamics would be triggered by a sequence of unusual shocks that changes the correlation between output and inflation. McGough (2006) extended Sargent’s results, showing that an escape-like path can be produced by even a small exogenous shock to the natural rate of unemployment. Bullard and Cho (2005) found that large deviations, generating a non-equilibrium outcome characterised by near-zero interest rates and persistently low inflation, can also arise in a microfounded New Keynesian model.

1.9 Learning as an alternative to rational expectations

If interest in the theoretical foundations of adaptive learning has mounted steadily over the past three decades, there has been an absolute boom in studies on its economic applications, covering such diverse topics as hyperinflation episodes, liquidity traps, currency crises, the forward premium puzzle, supply-side reforms, macroeconomic persistence and the fiscal theory of the price level. The bulk of this applied research, though, is in monetary economics.¹² The initial focus was on the uniqueness and learnability of equilibria, but scholars’ attention gradually moved to other subjects, such as the design of robust and optimal policy, the role of fundamentals in Taylor-type rules, the benefits from transparency, and the coordination of fiscal and monetary policies.

Notwithstanding its successes, the theory of econometric learning is not yet a viable alternative to rational expectations by any means. The enhancements

¹²Evans and Honkapohja (2008) give a comprehensive survey of recent studies on the monetary policy implications of assuming that agents are not fully rational and perfectly informed.

scored in terms of cognitive consistency are more than offset by the continuing indeterminacy of the choice of the learning algorithm, the specification of the PLM, the adoption of a specific gain sequence, and the selection criteria for the competing forecasting equations. Dropping the RE hypothesis eliminates jointly the assumption of individual rationality – which solves the problem of choosing the forecasting equation – and that of consistency of beliefs – which ensures that expectations are homogeneous. As stressed by Sargent (1993), economists embraced the hypothesis of rational expectations because if perceptions about the environment and other people’s behaviour are left unrestricted, and behaviours depend on perceptions, models can produce so many possible outcomes that they are useless as instruments for generating predictions: it is unfortunately undeniable that learning is not yet capable of providing the required restrictions on expectations.

1.10 A summary review of the dissertation

The essays included in this dissertation are mostly in applied and not theoretical economics, and all deal with learning and monetary policy; the first two assume that agents learn adaptively, whereas the third, which is based on Bayesian learning, treats monetary policy not as the subject matter, but as an input to study the real return on equities.

The first chapter (“Imperfect knowledge, adaptive learning and the bias against activist monetary policies”) studies the implications that abandoning the assumption of rational expectations has for the effectiveness of monetary policymaking. Effective policymaking requires that the monetary authorities commit to a systematic approach to policy. However, for policymakers to succeed in steering expectations and reaping the benefits of commitment, agents must be able to fully anticipate the future impact of monetary decisions, which is feasible only if the economic environment is stationary and

expectations are rational. When, on the contrary, the economy is subject to recurrent structural shifts and knowledge is imperfect, agents must rely on other methods for anticipating future events, which may alter the policy trade-offs dramatically. When knowledge is incomplete, an additional complication arises from the fact that committing to a systematic policy becomes problematic if not impossible (because the private sector cannot verify whether the central bank is delivering on its promises) and the gains from commitment are severely reduced (because expectations, being backward-rather than forward-looking, cannot be manipulated to increase policy effectiveness).

The economy is described by a Lucas-type equation, relating output to inflation surprises and supply shocks; the central bank controls price dynamics and the output-inflation trade-off is an exogenous time-varying stochastic process. In order to improve the discretionary equilibrium, society - the principal - assigns a loss function to the central bank - the agent - that may differ from society's preferences, if this is effective in increasing welfare. The assumption underlying this principal-agent approach is that it is possible to commit the monetary authority to a particular loss function, whereas the minimisation of the loss function occurs under discretion. Society has standard quadratic preferences on output and inflation and the central banker is endowed with either a quadratic or a lexicographic preference ordering. The literature on monetary policy almost always posits quadratic preferences, but this type of loss function does not reflect the task assigned to the monetary authority of most developed countries and is unable to account for the various aspects of actual policymaking. The assumption of lexicographic preferences overcomes a few of these shortcomings: it captures a hierarchical ordering of alternatives, which is typical of the mandate of virtually every inflation-targeting central bank, and allows the policymaker to focus on different policy objectives under different circumstances.

In this model both the private sector and the monetary policymaker have incomplete knowledge of the working of the economy and rely upon adaptive learning to form expectations (the private sector) or to estimate unknown parameters (the central bank) that form part of the optimal policy rule. The economy is subject to recurrent unobserved shifts and the monetary authority, which has private information on the shocks, cannot credibly commit.

The original contribution of this work is to extend the findings of Orphanides and Williams (2002) in three different ways. First, I assume that not only private agents but also the policymaker has imperfect knowledge; under this framework, policy effectiveness ends up depending both on inflation and output variability, so that the existence of a bias towards conservatism, if confirmed, cannot be attributed to the limited role of output volatility in the model. Second, I test whether society can increase welfare by appointing a policymaker whose preference ordering is lexicographic, which is rare in the literature and introduces non-linearities. Third, I test the claim that imperfect knowledge has a negative impact on economic stabilisation under a set of alternative learning mechanisms.

The research reaches three main findings. First, it is confirmed that when agents do not possess complete knowledge of the structure of the economy and rely on an adaptive learning technology, a bias towards conservatism arises, suggesting that society is better off designating a policymaker whose degree of inflation aversion is greater than its own: the rationale is that agents' and policymakers' attempts to learn how the economy works introduce inertia into the system and induce prolonged deviations of output and inflation from target, thereby raising the costs for the central bank of not responding promptly and forcefully to shocks. Second, what matters for society's welfare is that the inflation aversion of the monetary authority is strong enough to prevent expectations from fluctuating too much: the specific form of the loss function of the central bank is of second-order importance, so that it makes

little difference whether monetary authority is delegated to a central banker with lexicographic or quadratic preferences, though the former ordering tends to be associated with slightly better outcomes. Third, the bias against stabilisation policies and towards conservatism, and the relative efficiency of alternative monetary strategies, do not depend on whether the memory of the learning process is finite or infinite.

The second chapter (“Monetary policy in a model with misspecified, heterogeneous and ever-changing expectations”) studies the interaction between imperfect knowledge and monetary policy when the economy is complex and macroeconomic relations are described by a large system of non-linear equations. The vast literature on adaptive learning is overwhelmingly focused on small linear models, where the perceived law of motion (PLM) coincides with the minimum state variable (MSV) solution of the corresponding rational expectations equilibrium (REE). This is a convenient simplification, which allows straightforward analysis of the asymptotic properties of the learning algorithm and avoids the complexities of agents’ having heterogeneous information sets and possibly facing a multitude of alternative forecasting equations. With non-linear models these simplifications disappear, since a closed-form MSV solution does not exist; and if the model is medium-sized or large, no univocal linear approximation is available either, given the large number of state variables that in principle could be included in the PLM. If agents act like econometricians, who look at the data to choose the correct specification of a regression equation, the larger the model economy, the larger the set of options among which to select a PLM: different agents end up picking different forecasting models and no one sticks to the same PLM indefinitely, preferring to switch from one to another on the basis of observed predictive performance.

In a self-referential model, the lack of a univocal PLM implies that the equilibrium is not unique. Finally, since the PLM is misspecified and underpa-

parameterised it no longer nests the REE, and the learning algorithm converges to a limit point that is indeterminate and depends on the specific form of the expectations equations. With an unknown limit point, the issue of the convergence of the learning algorithm becomes somewhat vague and ill-defined. If one drops the simplifying assumptions of linearity and small size, most of the findings of the recent literature on adaptive learning become doubtful and their validity needs to be proven within a more general framework. Not only analytic issues but also policy prescriptions get intricate. For monetary policymaking, Orphanides and Williams (2007) have shown that when agents learn adaptively, the incentives and constraints facing monetary authorities are different and hence their strategies should change as well: compared with rational expectations, imperfect knowledge (i) reduces the scope for stabilisation of the real economy, (ii) demands more inflation-averse policies and (iii) increases the inertia in interest rate setting. The problems are compounded by unobserved structural changes (e.g. in natural rates), which call for quasi-difference rules in a quest for policy robustness. According to Orphanides and Williams, it is the non-linear nature of the learning process that dictates the policy response: upon re-estimation of the inflation forecasting model, a positive price shock passes through to the intercept of the forecasting equation, raising next-period expected and hence actual inflation; unless the policy response is prompt and firm, the persistence and the volatility of price changes increase and the monetary authority fails to keep a firm grip on the value of money.

While introducing learning in an otherwise linear system as in Orphanides and Williams changes the model's nature and equilibrium properties, it is not clear what happens when the model is intrinsically non-linear. Expectations heterogeneity introduces additional uncertainties, as stabilisation policies may interact with agents' choice among forecasting models. Several questions arise. Does the switch from full rationality to imperfect knowledge

amplify the impact of learning or does it just bring in an additional source of inertia, whose effects on the working of the economy and on monetary strategies are negligible? How does the speed of learning change because of the more complex structure of the economic environment? Does monetary policy affect the degree of heterogeneity of expectations, or vice versa? Can the central bank enhance welfare by providing information to households and firms, or should it exploit its private information to generate inflation surprises?

As this last question suggests, another relevant policy issue is whether the central bank can benefit from being transparent when private agents are not fully rational. The answer is not trivial, because when expectations are backward-looking, as in the case of adaptive learning, the monetary authority's influence on agents' beliefs is severely reduced. In a large-size model, where non-linearities abound, the flow of information from the monetary authority to the private sector is potentially very rich and the role of central bank communication highly relevant: agents do not know which variables are factored into the policy rule, whether the monetary authority targets lagged, current or future variables, or what the degree of policy inertia is. Since each of these aspects of the policy strategy affects the system's transitional dynamics and steady state growth path, to a large extent the monetary authority can decide the amount of information to make public so as to steer agents' choice of forecasting models and in so doing affect expectations and equilibrium outcomes.

The chapter is concerned with these issues. The first objective is to validate the findings of Orphanides and Williams on the influence of imperfect knowledge on the features of the monetary policy rule; in this respect, it is certainly a valuable asset to have a large and detailed model that identifies the channels through which monetary impulses are transmitted and the way expectation errors affect the persistence and the volatility of output and in-

flation, because this helps in understanding the interaction between policy measures, expectations and economic outcomes. Its second purpose is to measure the benefits associated with more transparent policies, comparing the impact on social welfare of alternative communication strategies.

This work is original in several respects. First, it analyses learning in an economy where expectations have a pervasive role that is unmatched in the literature. The very few papers that study bounded rationality in large non-linear models introduce learning only in the exchange rate equation;¹³ here, by contrast, learning affects not only the value of the domestic currency but also the term structure of interest rates, the conduct of monetary policy and the wage setting behaviour of unions and firms. Second, drawing from the literature on evolutionary game theory, it models the process through which agents switch from one forecasting model to another in order to form expectations. Third, it studies the consequences on social welfare of varying degrees of monetary policy transparency. When the REE cannot be achieved, due to the complexity of the economic environment, the disclosure of information from the central bank affects the selection of market participants' PLMs and determines the restricted perceptions equilibrium to which the economy converges. The description of the channels of monetary policy transmission helps to see whether the effect of central bank communication is mainly to reduce the noise or instead to provide distorted incentives to market participants, forcing them to attach too much weight to central bank communication and too little to their own information.

Three main results emerge. First, expectations heterogeneity is an intrinsic feature of the economy: regardless of the monetary policy in place, no PLM succeeds in ruling out all the other forecasting models, though the most inaccurate ones are eventually dismissed. Second, it is no longer true that when agents are boundedly rational, the monetary policymaker has an incentive to

¹³An exception is Dieppe *et al.* (2011).

adopt more inflation-averse policies; on the contrary, too strong a reaction to price shocks increases both inflation and output volatility and tends to make the model unstable and non-learnable. At first sight, this outcome may seem quite counterintuitive: a central bank that is more strongly committed to taming inflationary pressures should be more credible and more effective in anchoring long-run inflation expectations and bond yields. However, this connection is not present in the model and credibility depends on outcomes, not intentions: agents learn from the data and what matters is whether monetary policy makes the economy more or less stable. Third, more transparent (but not fully transparent) policies seem to be welfare-enhancing, though mildly so, mostly because they twist the slope of the yield curve and diminish the variability of short-term interest rates; besides, the degree of transparency alters the form of the optimal policy rule, by increasing inflation aversion. But the benefits of transparency are not monotonic: under full transparency, social welfare is marginally less than under opaqueness.

The final chapter (“Learning monetary policy regimes and asset pricing”) is somewhat different from the previous two. Agents are no longer modelled as econometricians but behave as fully-informed economists, using Bayesian techniques to filter unobserved variables and updating their estimates of the hidden state of the economy whenever new observations are available. Moreover, monetary policy is no longer the focus of the analysis but an exogenous input to the model that generates uncertainty among agents and is responsible for the negative correlation between stock returns and inflation, i.e. inducing the violation of the so-called Fisher effect, the one-to-one relationship between nominal returns and inflation.

This issue has originated a vast literature, both empirical and theoretical. The Fisher effect is usually viewed as a minimum requirement for market efficiency and agents’ rationality. Yet various studies have shown that common stocks, contrary to the conventional wisdom, do not insure investors

against inflation but tend to move in the wrong direction in the face of price shocks. The evidence on unconditional correlation is anything but convincing, however: it is generally spurious and suffers from omitted-variable bias. A number of explanations have been put forward for the negative correlation between stock returns and inflation, citing such factors as fiscal policy, wealth effects, central bank objectives, and beliefs about future output. None is entirely convincing, however, and my co-author Massimo Massa and I have sought to shift the focus to the information role of inflation, which is the most reliable signal of the intentions of the monetary policymaker when policy is opaque and investors are uncertain about the real objectives of the central bank.

The way stock returns react to inflation is strictly connected with the stance and the long-run targets of monetary policy. When the policymaker focuses almost exclusively on price stability, a rise in inflation engenders a monetary squeeze, which impinges on future growth and hence on expected dividends. That is, high inflation today foreshadows low cash flows tomorrow implying low stock prices today. On the contrary, when the central banker is a “dove”, aiming at moderating fluctuations in output and unemployment, inflation signals a cyclical upturn and anticipates favourable growth prospects. Accordingly, an acceleration in consumer prices is associated with higher future cash flows and stock prices.

The fact that investors cannot directly observe the monetary policy stance generates uncertainty. Any information that may help them infer the prevailing monetary policy regime reduces uncertainty and affects risk premia. Inflation, by offering the investors an opportunity to study the central bank’s reaction, provides such information. In particular, it may either reinforce or challenge established beliefs in a particular type of monetary policy. In the former case, by reinforcing investors’ assessment inflation reduces informational uncertainty, hence the risk premium and the required rate of return.

If this reduction in returns is large enough, it may induce a "puzzlingly" negative correlation between inflation and stock returns.

Some original contributions of the work are worth mentioning. Besides providing a link between asset pricing and monetary economics, it suggests a new channel, not yet properly explored, through which the central bank affects financial markets. Second, it sheds new light on the "rules versus discretion" debate, quantifying the cost in terms of risk premia of the lack of transparency and discretionary behaviour of the monetary policymaker. It would be interesting to extend this analysis to the international dimension to see whether countries characterized by different degrees of disclosure of the central bank's operating procedures and targets also display different impacts of monetary policy uncertainty on prices.

The research provides both an analytical and an empirical basis for the role of uncertainty in driving the equity risk premium and the correlation between stock returns and inflation. The empirical section includes an estimate - via a Markov-switching VAR model - of an index of monetary policy stance. This is then used to construct a tracking portfolio mimicking monetary policy uncertainty. The final step is a direct test of whether such a portfolio is in fact priced: monetary uncertainty is used in addition to the three Fama-French factors to explain stock returns in two types of investment portfolio (i.e. (i) industry and (ii) size and book-to-market portfolios). The results strongly support our working hypothesis: the coefficient of monetary policy uncertainty is positive and strongly significant. The coefficient remains highly significant even after the application of the White correction to control for the problem of generated regressors. In particular, the tracking portfolio is positive and strongly significant in 24 out of 25 cases for the size and book-to-market portfolios and in 13 out of 17 cases for the industry portfolios. The significance of the additional regressors survives a number of robustness checks.

The work points to the neglected role of learning as a potential source of uncertainty. Since inflation is a signal used by imperfectly informed investors to infer the stance of monetary policy, the learning process may raise the level of uncertainty and so increase the risk premium. The puzzling relationship between inflation and stock returns depends on the way in which the information conveyed by changes in the inflation rate modifies investors' uncertainty. If the signal embedded in inflation is consistent with investors' beliefs on the degree of "hawkishness" of the policymaker, an increase in inflation may reduce uncertainty and therefore the equity premium, thus leading to a negative correlation between inflation and returns.

Chapter 2

Adaptive learning and the bias against activist monetary policies

2.1 Introduction

Effective policymaking requires that the monetary authority commit to a systematic policy approach. As long as price setting depends on expectations, a credible central bank may benefit from a better short-run trade-off between inflation and output and accordingly it may reduce inflation at less cost. For policymakers to succeed in steering expectations and reaping the benefits of commitment, however, agents must be able to fully anticipate the future impact of monetary decisions on the economy, which is feasible only if the economic environment is stationary and expectations are rational. When instead the structure of the economy is subject to shifts and knowledge is imperfect, agents must find other methods for anticipating future events, which may alter the policy trade-offs dramatically.

A recent stream of literature, most notably Orphanides and Williams (2002), has shown that imperfect knowledge makes stabilisation policies more difficult. This happens when central banks ascribe too much importance to output stabilisation, because overly activist policies are prone to create situations in which the public's inflation expectations become uncoupled from the policy

objective. An additional complication arises from the fact that if knowledge is incomplete, then committing to a systematic policy becomes problematic because the private sector cannot easily verify whether the central bank is delivering on its promises, while the benefits of such commitment are severely reduced, because expectations, being backward- rather than forward-looking, cannot be manipulated to increase policy effectiveness.

This chapter examines how relaxing the rational expectations hypothesis changes the way monetary policy is set, taking a principal-agent approach: in order to improve the discretionary equilibrium, society (the principal) assigns a loss function to the central bank (the agent) that may differ from society's preferences. The assumption underlying this approach is that it is possible to commit the monetary authority to a particular loss function but that its minimisation is discretionary. Society has standard quadratic preferences on output and inflation and the central banker may have either quadratic or a lexicographic preference ordering. The literature on monetary policy almost always posits quadratic preferences, but this type of loss function does not reflect the task that most developed countries assign to the monetary authority and cannot account for the role of factors that are vital in actual policymaking.¹ Lexicographic preferences overcome some of these shortcomings: they capture the hierarchical ordering of alternatives that is typical of the mandate of virtually every inflation-targeting central bank² and allow the policymaker to focus on different policy objectives under different circumstances.

The main findings are the following. First, it is confirmed that when agents do not possess complete knowledge on the structure of the economy and rely

¹To mention just a few: (i) uncertainty (see Al-Nowaihi and Stracca (2002)); (ii) imperfect observability of the state variables (see Svensson and Woodford (2000)); (iii) path dependence and differential valuation of deviations from the inflation and output targets (see Orphanides and Wilcox (1996)).

²Buiter (2006) presents a list of the central banks whose mandate reflects a lexicographic preference ordering, with price stability ranking first and all other desiderata coming after.

on an adaptive learning technology, a bias towards conservatism arises, suggesting that social welfare is improved by designating a policymaker who is more inflation-averse than society itself. Even if the economy has no intrinsic dynamics, the attempts of agents and policymakers to learn adaptively introduce inertia, which makes it costly for the central bank not to respond promptly and forcefully to shocks. Second, what matters for social welfare is that the monetary authority's inflation aversion be strong enough to prevent expectations fluctuating too much. The specific form the loss function of the central bank takes is of second-order importance. Third, the bias against stabilisation policies and towards conservatism and the relative efficiency of alternative monetary strategies do not depend on whether the memory of the learning process is finite or infinite.

This chapter relates to the literature in several ways. It parallels, under more general conditions, the work of Terlizzese (1999) and Driffill and Rondoni (2003) in deriving the properties of a monetary strategy whose prime objective is price stability, with an inflation cap; it models two-sided learning in the vein of Evans and Honkapohja (2002); it uses the study of Orphanides and Williams (2002) as a benchmark for the quantitative analysis. The most closely related of these contributions is Orphanides and Williams: model simulations are designed so as to replicate their experiments, and the objective here is largely the same as theirs, namely to understand how the economy responds to alternative monetary strategies when agents have bounded rationality and imperfect knowledge.

My original contribution consists in extending the findings of Orphanides and Williams (2002) in three different ways. First, the work assumes that both private agents and the monetary policymaker have imperfect knowledge; under this framework, policy effectiveness ultimately depends both on inflation and on output variability, so that the bias towards conservatism, if confirmed, cannot be attributed to the limited role of output volatility in the reference

model. Second, it tests whether society can increase welfare by appointing a policymaker whose preference ordering is lexicographic. Third, it tests the claim that imperfect knowledge affects economic stabilisation adversely under a set of alternative learning mechanisms.

The chapter is organised as follows. Section 2 outlines the model and contrasts the implications of assuming quadratic and lexicographic preferences. Section 3 introduces econometric learning and studies how different policies affect the speed at which learning algorithms converge on the rational expectations equilibrium. Section 4 presents some evidence, produced by simulations, on the distortions of monetary policymaking that are generated by postulating agents' bounded rationality; the issues are whether adaptive learning induces a bias toward conservatism and whether appointing a central banker with lexicographic preferences is welfare-improving. Section 5 concludes.

2.2 The model

The model has two basic components: (i) the non-observability of the supply shock; (ii) an unknown and time-varying output-inflation trade-off. I solve it first under rational expectations, then under adaptive learning.

2.2.1 The structure of the economy and the delegation problem

The economy is characterised by an expectations-augmented Phillips curve relationship, linking inflation surprises $\pi - \pi^e$ to (detrended) output y .

$$y = \alpha (\pi - \pi^e) + \varepsilon \tag{2.1}$$

Inflation is the policy instrument and is controlled without error by the monetary authority; the natural level of output is normalised to zero. Output also responds to a zero-mean supply shock ε , unobservable to the central

their inflation expectations for the period. The information advantage of the central bank creates a role for stabilisation policies and takes into account the fact that policy decisions can be made more frequently than most wage and price decisions.

Society is assumed to have quadratic utility and to dislike both inflation and output variability. Its welfare function is

$$W^S = -E \left[(y - k)^2 + \beta^S \pi^2 \right] \quad (2.2)$$

where k is the target level of output; the expectation operator is due to the non-observability of the output-inflation trade-off α and output shock ε . The assumption that $k > 0$ is usually justified on the grounds that the presence of labour and goods market distortions leads to an inefficiently low level of output in equilibrium; alternatively, $k > 0$ is interpreted as arising from political pressures on the central bank. In the principal-agent approach, society, whose preferences are quadratic, can designate a central banker whose loss function may differ from its own. It is assumed that the choice is restricted to policymakers with quadratic or lexicographic preferences; in the case of quadratic preferences, the inflation aversion parameter β can be different from β^S .

2.2.2 The central bank's loss function

Though in general a lexicographic preference ordering cannot be represented by a function, in the simplified case in which the monetary authority has only two objectives, such an ordering can be described by a loss function involving only the secondary objective, subject to a constraint involving the primary target. I accordingly assume that the central bank seeks to stabilise output at a non-zero level, provided that inflation is kept below a known upper bound.⁵

⁵The existence of a lower bound on inflation is not considered. In a model where the output shock is observable and the trade-off between output and inflation is time-invariant, Terlizzese (1999) shows that

Though there is some dispute on the correctness of formulations depicting central bankers as affected by an inflation bias, the assumption is retained because otherwise the only rational expectation for inflation would be the zero target itself and the inflation constraint would never be binding.

In formal terms, the problem solved by the central bank is

$$\begin{aligned} \min_{\pi} \quad & \frac{1}{2} E (y - k)^2 \\ \text{s.t.} \quad & \begin{cases} \pi \leq \bar{\pi} \\ y = \alpha (\pi - \pi^e) + \varepsilon \end{cases} \end{aligned} \quad (2.3)$$

Notice that k cannot exceed μ , the upper bound of the output shock: in what follows, it will be assumed that k is not too high. Under the standard hypothesis of time-separability of preferences, the problem is static and involves no trade-off between current and future utility, so that the optimal policy does not have to rely on the strategic interactions described by Ellison and Valla (2001), Bertocchi and Spagat (1993) and Wieland (2003).⁶

To highlight the implications of endowing the monetary authority with lexicographic preferences, the policy problem is also analysed under the standard assumption that the loss function is quadratic. In this case, the problem solved by the central bank can be formulated as

$$\begin{aligned} \min_{\pi} \quad & \frac{1}{2} E \left[(y - k)^2 + \beta \pi^2 \right] \\ \text{s.t.} \quad & y = \alpha (\pi - \pi^e) + \varepsilon \end{aligned} \quad (2.4)$$

where β measures the weight attached to the inflation objective relative to output stabilisation.

the main features of the monetary policy problem are largely unaffected by the inclusion of a lower bound on inflation. Intuitively, what explains this result is the asymmetric nature of the inflation bias that is assumed to characterise the monetary authority's preferences: if the central bank tries to push output above the natural level, it will tend to inflate, so that while the upper bound will often be binding, the lower one will not.

⁶Buiter (2006) provides an alternative, more restrictive, representation of a lexicographic preference ordering.

Regardless of the specific form of the loss function, be it (2.3) or (2.4), the model features an inflation bias, arising from the policymaker's incentive to create surprise inflation in order to keep output above the natural level. Many economists think that this feature makes the model irrelevant for monetary policy analysis, as central bankers are not in the business of fooling people.⁷ Jensen (2003) however shows that a variant of the same model can generate an inflation bias even without attributing any role to inflation surprises.

2.2.3 Signal extraction and the rational expectations equilibrium

Given the structure of the problem, the issue of estimating the unobserved output shock and that of setting the optimal inflation rate can be kept separate and solved sequentially. Before deciding the optimal policy, the central bank has to solve a signal extraction problem. The first step is therefore to derive the probability distribution of $z = \varepsilon + \xi$ and the conditional mean $E(\varepsilon|z)$. In Lemma 1 the density function of the signal z is derived, while in Proposition 1 the first moment of the distribution of the output shock ε conditional on z is calculated.

Lemma 1 If $z = \varepsilon + \xi$ and ε and ξ are independent uniform random variables, with support on the interval $[-\mu, \mu]$, then the density function of z is equal to $f(z) = \frac{1}{2\mu} + \frac{1}{4\mu^2} [\min(z, 0) - \max(0, z)] = \frac{1}{2\mu} - \frac{1}{4\mu^2}|z|$.

Proof See the appendix.

Proposition 1 If ε and ξ are uniform random variables, defined on the same close interval $[-\mu, \mu]$, and $z = \varepsilon + \xi$, then the optimal estimate of ε conditional on z is $E(\varepsilon|z) = \frac{z}{2}$

Proof See the appendix.

⁷See, for instance, the quotations from Blinder, Vickers and Issing listed in Jensen (2003).

Given the assumption that governs the flow of information, the central bank sets the inflation rate on the basis of the observed signal and the private sector's inflation expectations. Under lexicographic preferences, it will choose the inflation rate that solves the first order condition $E[\alpha(\alpha(\pi - \pi^e) + \varepsilon - k) | z] = 0$ - provided that the inflation constraint is satisfied - and will choose $\pi = \bar{\pi}$ otherwise, i.e.

$$\pi = \begin{cases} \pi^e - \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2} \left(\frac{z}{2} - k \right) = \pi^e - \frac{z}{\phi} + \frac{2k}{\phi} & \text{if } z \geq 2k + \phi(\pi^e - \bar{\pi}) = \Lambda \\ \bar{\pi} & \text{otherwise} \end{cases} \quad (2.5)$$

where $\phi^{-1} \equiv \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2} \frac{1}{2}$ and $\Lambda = 2k + \phi(\pi^e - \bar{\pi})$. The optimal policy can be written also in a more compact - but somewhat less transparent - form, namely

$$\pi = \pi^e - \frac{1}{\phi} [\max(z, 2k + \phi(\pi^e - \bar{\pi})) - 2k] = \pi^e - \frac{1}{\phi} [\max(z, \Lambda) - 2k]$$

The optimal policy depends, in a non-linear way, on the value of z : when output shocks are strongly negative and the primary objective is at risk, the central bank acts like an inflation nutter; when the signal indicates more favourable disturbances, it displays greater activism, favouring output stabilisation. Notice that the optimal policy depends on the parameters of the distribution of the output-inflation trade-off α . Two cases are considered, one corresponding to the rational expectations equilibrium (REE), and the other assuming bounded rationality and least-square learning. In the first case, it is assumed that α is not observed but $\bar{\alpha}$ and σ_α^2 are known by both the central bank and the private sector;⁸ in the second case, $\bar{\alpha}$ and σ_α^2 are unknown and must be estimated.

It is worthwhile noting a few points that illustrate the main properties of the optimal policy. First, the optimal policy is not altered by the non-

⁸This assumption could be justified if ε , though unobserved at the time when expectations and the inflation rate were set, was observed with a one-period lag.

observability of the output shock, except that the policymaker responds to an efficient estimate rather than the actual value of the state variable. It is well known that a linear model with a quadratic loss function and a partially observable state of the economy is characterised by certainty-equivalence; since the assumed preference ordering is not quadratic, the result applies only when inflation is within the admissible range. Second, uncertainty about the multiplier of the policy instrument makes it optimal to react less than fully to the output shock. There is nothing to be gained by more activist policy in order to learn from experimentation, since the model is static and the current welfare loss incurred by overreacting will not be offset by future gains. The reduction in policy activism caused by parameter uncertainty, originally shown by Brainard (1967), reflects the direct impact of the policy instrument on the variability of the target variable.

For the equilibrium to be fully characterised, one must provide the solution for expected inflation. Under rational expectations, agents understand the incentives driving the actions of the central bank and on average their expectations coincide with realisations. Accordingly $\pi^e = \int \pi dF(z)$, where $F(z)$ is the distribution function of the signal z . Proposition 2 gives the full characterisation of the rational expectations equilibrium and provides the closed-form solution for expected inflation under the simplifying assumption that $k = \frac{\mu}{6}$, the value used henceforth.⁹

Proposition 2 If the central bank has lexicographic preferences and output is determined as in (2.1), there is a unique RE equilibrium, where $\pi = \min \left[\pi^e - \frac{1}{\phi} (z - 2k), \bar{\pi} \right]$, or, more compactly, $\pi = \pi^e - \frac{1}{\phi} (\max(z, \Lambda) - 2k)$, where $\Lambda = 2k + \phi(\pi^e - \bar{\pi})$. Expected inflation is found by solving the equilibrium condition $2k = E[\max(z, \Lambda)]$. In the simplifying case where $k = \frac{\mu}{6}$, the closed-form solution for expected inflation is known

⁹Setting $k = \frac{\mu}{6}$ amounts to assuming that $2k + \phi(\pi^e - \bar{\pi}) = 0$ and implies that the central bank chooses its best static response when $z > 0$, while it has no discretion for negative values of the signal.

and is equal to $\pi^e = \bar{\pi} - \frac{2k}{\phi} = \bar{\pi} - \frac{\mu}{3\phi}$.

Proof See the appendix.

Equilibrium is non-cooperative Nash: the central bank and the private sector try to maximise their respective objective functions taking the other player's actions as given. The assumption of rational expectations implicitly defines the loss function of the private sector as $E(\pi - \pi^e)^2$: given the public's understanding of the central bank's decision problem, its choice of π^e is the one that minimises disutility.

From the expression for π^e , it is apparent that the existence of an upper bound on inflation helps to stabilise expected inflation: for any value of k , the lower $\bar{\pi}$, the lower expected inflation. Another feature of the policy is that the larger the support of the output shock, the closer to zero π^e . The reason for this result is straightforward: positive (and larger than k) output shocks trigger a reaction from the central bank, which creates negative inflation surprises to stabilise output; large negative disturbances, on the other hand, cannot be neutralised, because excessively high inflation rates are not admissible. Widening the support of ε has an asymmetric effect on the actions of the monetary authority: it extends the range of cases in which the central bank finds it optimal to deflate but does not affect its incentives to inflate.

Two features of the optimal monetary policy are worth stressing. First, for values of the signal in the non-empty interval $[-2\mu, 0)$, inflation is constant and equal to $\bar{\pi}$, which is higher than π^e : the central bank keeps price dynamics above inflation expectations and thus sustains output, though it cannot cushion it against shocks. Second, even in the face of favourable output shocks, the policymaker cannot fully stabilise output at the desired level. Two factors contribute to attenuate the policymaker's response: uncertainty about the output-inflation trade-off and non-observability of ε ; the former

reduces the response of the central bank by a factor of $\frac{2\alpha}{\phi}$, while the latter leaves part of the output shock, namely $\varepsilon - \frac{z}{2}$, unchecked. Since $\frac{z}{2}$ is an unbiased estimate of ε , unobservability of the output shock increases volatility, but does not affect the degree of activism of the policy response; on the contrary, unobservability of the output-inflation trade-off has a bearing on the policy strategy, since it favours more cautious policies. It turns out that output volatility is therefore smaller than $\frac{\mu^2}{3}$, the variance of the output shock, implying some degree of stabilisation on the side of monetary policy.

To assess the distinguishing traits of a lexicographic-preference policy, let us contrast it with the optimal policy arising under the standard assumption of quadratic loss function. If (2.4) describes the central bank's problem, the policy instrument is set according to the rule

$$\begin{aligned}\pi &= \frac{\bar{\alpha}^2 + \sigma_\alpha^2}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta} \pi^e + \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta} \left(k - \frac{z}{2}\right) = \rho \pi^e + \frac{\rho}{\phi} (2k - z) \\ &= \pi^e - \frac{\rho}{\phi} z\end{aligned}\tag{2.6}$$

where $\rho \equiv \frac{\bar{\alpha}^2 + \sigma_\alpha^2}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta}$, with $0 < \rho \leq 1$, and $\phi^{-1} \equiv \frac{1}{2} \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2}$, as in (2.5). Expected inflation is equal to $\frac{\bar{\alpha}}{\beta} k$.

Several differences are apparent. When output shocks are not too strongly negative, rule (2.5) ensures greater output stabilisation, as the increase in the inflation rate required to counteract the supply disturbance does not have a negative impact on welfare and hence does not bring a trade-off between the output and the inflation objectives. The contrary holds when ε is large and negative, because in that case output stabilisation is sacrificed to the primary objective of price stability. More activist policies are possible at the cost of greater inflation variability: in general¹⁰ there exists a value $\bar{\beta}$ such that, for $\beta \in (\bar{\beta}, \infty)$, output variability under lexicographic preferences (strategy LEX) is lower than under quadratic utility (strategy QUA), and there exists a value $\underline{\beta}$ such that, for $\beta \in [0, \underline{\beta})$, $E(\pi - \pi^e)^2$ is smaller under strategy

¹⁰For a formal proof, see Locarno (2006), proposition 4.

LEX. Since $\underline{\beta} < \overline{\beta}$, strategy LEX apparently cannot outperform strategy QUA in terms of both objectives, but in reality this is not necessarily the case, since what is important for social welfare is not $E(\pi - \pi^e)^2$ but $E\pi^2$. Under strategy LEX the central bank can use an additional instrument - the upper bound on inflation $\overline{\pi}$ - to seek both lower output volatility and lower mean square inflation: the reason why strategy LEX is appealing compared with the Rogoff's solution¹¹ is that the reduction of the inflation bias does not come at the cost of the output stabilisation objective;¹² the drawback is that it tends to stabilise output too much when $\pi < \overline{\pi}$.

2.3 Adaptive learning and monetary policy regimes

In the real world, where shifts in policies and in the economic structure are by no means rare, people must often determine how the environment has changed and the least costly way of adapting decision rules to suit the new framework. In such a context, strict application of the rational expectations hypothesis (REH) is not a convincing theoretical solution. Alternatives have long been suggested. Herbert Simon (1957), for instance, supported some kind of bounded rationality and proposed to create a theory with behavioural foundations where agents learn in the same way as econometricians. An increasingly important literature building on the pioneering work of Bray (1982) and Marcet and Sargent (1989), recently revived in particular by Evans and Honkapohja, has introduced a specific form of bounded rationality – adaptive learning – in which agents adjust their forecast rule as new data becomes available. This approach provides an asymptotic justification for the REH and allows us to ignore non-learnable solutions in models with

¹¹See Rogoff (1985).

¹²This is easily seen by noting that the choice of $\overline{\pi}$ has no effect on the set of values of z corresponding to inactive policies. This implies that it is always optimal for the policymaker to set $\overline{\pi}$ so that π^e is equal to zero.

multiple equilibria. Convergence of the adaptive learning process may be studied by checking the E-stability conditions, but the speed of convergence is in general unknown. Benveniste, Metiver and Priouret (1990) show that root- t convergence¹³ of the learning process holds when the real part of the largest eigenvalue of the Jacobian of the ordinary differential equation (ODE) is less than $-\frac{1}{2}$. When this condition on the eigenvalues is not met, no analytic results on the asymptotic distribution are known, since the influence of the initial conditions fails to die out quickly enough. Marcet and Sargent (1992) suggest a numerical procedure to obtain an estimate of the rate of convergence. The starting point is the assumption that there is a δ for which

$$t^\delta (\theta_t - \theta) \xrightarrow{D} F \quad (2.7)$$

where θ_t is the vector of parameters of the perceived law of motion (PLM), θ is its asymptotic limit and F is some non-degenerate, well-defined mean-zero distribution. Marcet and Sargent show that, for large t , a good approximation of the rate of convergence δ is given by the expression

$$\delta = \frac{1}{\log l} \log \sqrt{\frac{E(\theta_t - \theta)^2}{E(\theta_{tl} - \theta)^2}} \quad (2.8)$$

The expectations can be approximated by simulating a large number of independent realisations of length t and tl (i.e. l times t), and calculating the mean square across realisations.

In this section, adaptive learning is introduced to analyse the implications of imperfect knowledge on policy outcomes. The question is how the interaction between learning and central bank preferences affects aggregate welfare. For the sake of clarity, first I consider the case in which only the private sector learns; the model is then expanded to incorporate central bank learning.

¹³Root- t convergence means convergence at a rate of the same order as the root square of the sample size.

2.3.1 Private sector learning

Suppose that private agents have non-rational expectations, which they try to correct through adaptive learning. Assume further that the policymaker does not explicitly take agents' learning into account and continues to set policy according to either (2.5) or (2.6). The evolution of output and inflation is therefore described by the system

$$\begin{aligned} y &= \alpha \left(\pi - \widehat{E}^P \pi \right) + \varepsilon \\ \pi &= \begin{cases} \min \left[\widehat{E}^{CB} \pi - \frac{z}{\phi} + \frac{2k}{\phi}, \bar{\pi} \right] \\ \rho \widehat{E}^{CB} \pi - \frac{\rho}{\phi} z + \frac{\rho}{\phi} 2k \end{cases} \end{aligned} \quad (2.9)$$

where the inflation rate depends on the monetary authority's preferences. $\widehat{E}^P \pi$ represents the current estimate of the inflation rate of the private sector, while $\widehat{E}^{CB} \pi$ is the value of inflation expectations used in the central bank's control rule. It is assumed that private agents run regressions to set $\widehat{E}^P \pi$, while the monetary authority, which observes $\widehat{E}^P \pi$ before moving, has rational expectations and therefore sets $\widehat{E}^{CB} \pi = \widehat{E}^P \pi$.

In each period t , private agents have a PLM for inflation, which they use to make forecasts; it takes the form $\widehat{E}^P \pi_t = \pi_{P_t}^e$,¹⁴ where

$$\pi_{P_t}^e = \pi_{P_{t-1}}^e + \frac{1}{t} (\pi_{t-1} - \pi_{P_{t-1}}^e) \quad (2.10)$$

As inflation fluctuates around its mean depending on the value of the signal z , which is observed only by the central bank, private agents form expectations about inflation assuming that it is a constant and equals the arithmetic mean of past inflation rates.¹⁵

The estimate $\pi_{P_t}^e$ is updated over time using least squares; $\frac{1}{t}$ represents the gain parameter, which is a decreasing function of the sample size. Equation

¹⁴A time index is used only when strictly necessary, for instance when tracking the evolution over time of least squares learning.

¹⁵If agents knew $\bar{\pi}$ and $\text{prob}(\pi = \bar{\pi})$, they could also use the following PLM: $\text{prob}(\pi < \bar{\pi}) \pi_{P_t}^e + \text{prob}(\pi = \bar{\pi}) \bar{\pi}$, where $\pi_{P_t}^e$ is defined as in (2.10) and $\text{prob}(\pi < \bar{\pi}) + \text{prob}(\pi = \bar{\pi}) = 1$. This choice however would not affect the E-stability conditions.

(2.10), in line with the literature, is in recursive form, uses data up to period $t - 1$, and requires a starting value at time $t = 0$. The PLM has the same form as the RE solution for expected inflation: private agents estimate the parameter of the reduced form and set $\widehat{E}^P \pi_t = \pi_{Pt}^e$.

Consider first the case in which the central bank acts as if it had a lexicographic ordering of preferences. The actual law of motion (ALM) turns out to be

$$\pi_t = \begin{cases} \pi_{Pt}^e - \frac{z_t}{\phi} + \frac{2k}{\phi} & \text{if } z_t \geq 0 \\ \bar{\pi} & \text{otherwise} \end{cases} \quad (2.11)$$

The mapping between the PLM and the ALM generates the stochastic recursive algorithm

$$\pi_{Pt}^e = \begin{cases} \pi_{Pt-1}^e + \frac{1}{t} \left(-\frac{z_{t-1}}{\phi} + \frac{2k}{\phi} \right) & \text{if } z_{t-1} \geq 0 \\ \pi_{Pt-1}^e + \frac{1}{t} (\bar{\pi} - \pi_{Pt-1}^e) & \text{otherwise} \end{cases} \quad (2.12)$$

which is approximated by the following ODE

$$\frac{d}{d\tau} \pi_P^e = h(\pi_P^e) = \lim_{t \rightarrow \infty} E(\pi_{t-1} - \pi_P^e) \quad (2.13)$$

where

$$\begin{aligned} \lim_{t \rightarrow \infty} E(\pi_{t-1} - \pi_P^e) &= (\bar{\pi} - \pi_P^e) \int_{-2\mu}^0 \left(\frac{1}{2\mu} + \frac{z}{4\mu^2} \right) dz + \int_0^{2\mu} \left(-\frac{z}{\phi} + \frac{2k}{\phi} \right) \left(\frac{1}{2\mu} - \frac{z}{4\mu^2} \right) dz \\ &= \frac{1}{2} (\bar{\pi} - \pi_P^e) + \frac{k - \mu/3}{\phi} \\ &= \frac{1}{2} (\bar{\pi} - \pi_P^e) - \frac{\mu}{6\phi} \end{aligned}$$

Notice that the fixed point of the ODE, namely $\pi_P^e = h^{-1}(0) = \bar{\pi} - \frac{2k}{\phi} = \bar{\pi} - \frac{\mu}{3\phi}$, coincides with the unique RE equilibrium for expected inflation. The theorems on the convergence of stochastic recursive algorithms can be applied so that convergence is governed by the stability of the associated ODE.¹⁶ Since $\frac{d}{d\pi_P^e} h(\pi_P^e) = -\frac{1}{2} < 0$, the ODE is (globally) stable, and hence adaptive learning converges asymptotically to the RE equilibrium. Notice

¹⁶Chapter 6 of Evans and Honkapohja (2001) studies the conditions under which the convergence of

that the eigenvalue of $h(\pi_P^e)$ depends on the share of the support of z that corresponds to an active policy: in general, root- t convergence does not hold. Consider now the situation in which the central bank has quadratic preferences. The optimal policy for the monetary authority is to set $\pi_t = \rho\pi_{Pt}^e - \frac{\rho}{\phi}z_t + \frac{\rho}{\phi}2k$. Differently from the RE case, here the central bank does not completely offset inflation expectations and the parameter k explicitly enters the control rule: both features disappear asymptotically, provided that $\pi_P^e \rightarrow \pi^e = \frac{\bar{\alpha}}{\beta}k$.

If agents use recursive least squares, then expectations evolve according to the equation

$$\pi_{Pt}^e = \pi_{Pt-1}^e + \frac{1}{t} (\pi_{t-1} - \pi_{Pt-1}^e) = \pi_{Pt-1}^e + \frac{1}{t} \left[(\rho - 1) \pi_{Pt-1}^e - \frac{\rho}{\phi} z_{t-1} + \frac{\rho}{\phi} 2k \right] \quad (2.14)$$

and

$$h(\pi_P^e) = \lim_{t \rightarrow \infty} E \left[(\rho - 1) \pi_{Pt}^e - \frac{\rho}{\phi} z_{t-1} + \frac{\rho}{\phi} 2k \right] = (\rho - 1) \pi_P^e + \frac{\rho}{\phi} 2k \quad (2.15)$$

Also in this case, the fixed point of the ODE, namely $\pi_P^e = h^{-1}(0) = \frac{\rho}{(1-\rho)\phi} 2k = \frac{\bar{\alpha}}{\beta} k$, coincides with the unique RE equilibrium and the system is (globally) stable. In fact, $\frac{d}{d\pi_P^e} h(\pi_P^e) = \rho - 1 < 0$, since ρ is positive and smaller than one. Whether or not root- t convergence holds depends on the size of β : the greater the weight the central bank attaches to the inflation objective, the faster agents learn. The explanation of this result is easily grasped. In order to offset output shocks the central bank must generate inflation surprises, i.e. move the inflation rate away from expectations, so that

stochastic recursive algorithms is guaranteed. In the case of interest, this is equivalent to a demonstration that the process π , which is a linear function of z , is bounded and stationary and that the function driving the updating of the projection parameter, namely $\pi_{t-1} - \pi_{Pt-1}^e$, is bounded and is twice countinuously differentiable (with respect to both π_{t-1} and π_{Pt-1}^e), with bounded second derivatives. Whether stability holds locally or globally depends on whether the regularity conditions hold on an open set around the equilibrium or for all admissible values of π_P^e . These regularity conditions, in the present case, are clearly met.

in every period agents have to revise their estimate with values of π that may differ substantially from the unconditional mean: the larger inflation surprises (i.e. the lower β), the slower the learning process.

With lexicographic preferences, a similar result obtains: by reducing the support of the signal corresponding to an active policy (i.e. to a policy that aims at avoiding excessive output fluctuations), expectations adjust more quickly to the long-term equilibrium. Notice that for reasonable parameterization of the model the value of β must be high in order for $\frac{d}{d\pi_P^e} h(\pi_P^e)$ to be less than $-\frac{1}{2}$, meaning that only in the case of a highly inflation-averse central bank will root- t convergence hold.

The previous result is interestingly similar to that of Orphanides and Williams (2002), who use a dynamic model based on aggregate supply and demand equations. They find that, with imperfect knowledge, the ability of private agents to forecast inflation depends on the monetary policy in place, with forecast errors smaller on average when the central bank responds more aggressively to inflationary pressures. Significantly improved economic performance can be achieved by stronger emphasis on controlling inflation: more aggressive policies reduce the persistence of inflation and facilitate the formation of expectations, which in turn enhances economic stability and mitigates the effect of imperfect knowledge. Their conclusion turns out to be quite similar to the Rogoff solution to the central bank's credibility problem under discretion: to improve welfare, the responsibility of the conduct of monetary policy must be delegated to a policymaker who is more inflation-averse than the society.

2.3.2 Private sector and central bank learning

I now consider the case in which $\bar{\alpha}$ and σ_α^2 (or, alternatively, ϕ) are not known to the policymaker. The central bank needs to estimate them, since both parameters affect the policy rule via the degree of responsiveness to

the signal z . As usual, I assume that they are gauged by least squares and that the estimate is updated every time new realisations of y and π become available. This form of bounded rationality corresponds to the case in which ε is never observed, so that $\bar{\alpha}$ and σ_α^2 cannot be estimated directly on the basis of past realisations of the output shock.

To account for central bank learning, the previous model must be augmented with a new set of recursive equations, which are the same irrespective of the monetary authority's preferences, as learning involves parameters rather than variables so that the values to be estimated are not related to agents' behaviour.

The system of recursive least squares equations is now the following

$$\begin{aligned}\pi_{Pt}^e &= \pi_{Pt-1}^e + \frac{1}{t} (\pi_{t-1} - \pi_{Pt-1}^e) \\ \widehat{\alpha}_t &= \widehat{\alpha}_{t-1} + \frac{1}{t} R_{\pi,t-1}^{-1} (\pi_{t-1} - \pi_{Pt-1}^e) \left[(y_{t-1} - \frac{z_{t-1}}{2}) - \widehat{\alpha}_{t-1} (\pi_{t-1} - \pi_{Pt-1}^e) \right] \\ R_{y,t} &= R_{y,t-1} + \frac{1}{t} \left[(y_{t-1} - \frac{z_{t-1}}{2})^2 - R_{y,t-1} \right] \\ R_{\pi,t} &= R_{\pi,t-1} + \frac{1}{t} \left[(\pi_{t-1} - \pi_{Pt-1}^e)^2 - R_{\pi,t-1} \right]\end{aligned}\tag{2.16}$$

or, more compactly,

$$\theta_t = \theta_{t-1} + \frac{1}{t} Q(\theta_{t-1}, X_t)$$

where $\theta_t = (\pi_{Pt}^e, \widehat{\alpha}_t, R_{y,t}, R_{\pi,t})'$ and $X_t = (1, \alpha_t, z_t, \varepsilon_t)'$. The first equation is the same as in the previous section and captures private sector learning, while the others refer to the central bank's inference problem: $\widehat{\alpha}_t$ is an estimate of the mean value of the output-inflation trade-off; $R_{y,t}$ measures the sample variance of $y - \frac{z}{2}$, the policy-driven component of the output gap;¹⁷ $R_{\pi,t}$ is the second moment of the inflation surprise. As shown below, the central bank calculates the statistics $R_{y,t}$ and $R_{\pi,t}$ as an intermediate step in estimating the optimal response coefficient to the signal z in the policy rule.

¹⁷Since $E(\varepsilon|z) = \frac{z}{2}$, the difference $y - \frac{z}{2}$ represents an unbiased estimate of the share of the output gap that depends on the inflation surprise only.

While the recursion for $R_{\pi,t}$ is obvious, as it is simply the estimate of the variance of the inflation surprise, the other two equations require some explanation. To understand the recursion for $\widehat{\alpha}_t$, notice that the output equation can be rearranged as

$$y - \frac{z}{2} = \bar{\alpha}(\pi - \pi_P^e) + [\varepsilon - \frac{z}{2} + \tilde{\alpha}(\pi - \pi_P^e)] \quad (2.17)$$

The central bank observes the signal z and can efficiently estimate $\bar{\alpha}$ by regressing $y - \frac{z}{2}$ on the inflation surprise $(\pi - \pi_P^e)$. Using $y - \frac{z}{2}$ as the regressand allows for the elimination of the simultaneity bias: the inflation surprise, being a linear function of the signal z only, is uncorrelated with $\varepsilon - \frac{z}{2}$ – the residual of the regression of ε on z – and with $\tilde{\alpha}$, by assumption orthogonal to all other shocks in the model.

The justification for the recursion for $R_{y,t}$ is somewhat more complicated. A biased estimator of $E(\alpha^2)$ can be obtained from the sample average of the square of the policy-driven component of the output gap, scaled by the second moment of the inflation surprise

$$\frac{E(y - \frac{z}{2})^2}{E(\pi - \pi_P^e)^2} = \frac{E(\alpha^2) E(\pi - \pi_P^e)^2 + E(\varepsilon - \frac{z}{2})^2}{E(\pi - \pi_P^e)^2} = \bar{\alpha}^2 + \sigma_\alpha^2 + \frac{\frac{1}{2} \frac{\mu^2}{3}}{E(\pi - \pi_P^e)^2}$$

The bias is easily calculated, since it depends on $E(\pi - \pi_P^e)^2$ and on known parameters. The sample estimate of $\bar{\alpha}^2 + \sigma_\alpha^2$ is therefore obtained by using the expression $\psi_t \equiv \frac{R_{y,t} - \frac{1}{2} \frac{\mu^2}{3}}{R_{\pi,t}}$.¹⁸

Whether the stochastic recursive algorithm converges or not depends on the associated ODE, i.e. on the Jacobian of the matrix $h(\theta) = \lim_{t \rightarrow \infty} EQ(\theta, X_t)$.

¹⁸An alternative could have been to regress $(y - \frac{z}{2})^2$ on $(\pi - \pi_P^e)^2$. The drawback to this approach is that the bias is a more convoluted function of the model parameters than in the case considered here.

In the case of lexicographic preferences, it can be shown that the ODE is

$$\begin{bmatrix} \frac{d}{d\tau} \pi_P^e \\ \frac{d}{d\tau} \widehat{\alpha} \\ \frac{d}{d\tau} R_y \\ \frac{d}{d\tau} R_\pi \end{bmatrix} = h(\theta) = \begin{bmatrix} \frac{1}{2} (\bar{\pi} - \pi_P^e) - \frac{\widehat{\alpha}}{2 \frac{R_y - \frac{1}{2} \frac{\mu^2}{3}}{R_\pi}} \frac{\mu}{6} \\ R_\pi^{-1} E (\pi - \pi_P^e)^2 (\bar{\alpha} - \widehat{\alpha}) \\ E (y - \frac{z}{2})^2 - R_y \\ E (\pi - \pi_P^e)^2 - R_\pi \end{bmatrix} \quad (2.18)$$

while for the standard quadratic case it is equal to

$$\begin{bmatrix} \frac{d}{d\tau} \pi_P^e \\ \frac{d}{d\tau} \widehat{\alpha} \\ \frac{d}{d\tau} R_y \\ \frac{d}{d\tau} R_\pi \end{bmatrix} = h(\theta) = \begin{bmatrix} -\frac{\beta}{\frac{R_y - \frac{1}{2} \frac{\mu^2}{3}}{R_\pi} + \beta} \pi_P^e + \frac{\widehat{\alpha}}{\frac{R_y - \frac{1}{2} \frac{\mu^2}{3}}{R_\pi} + \beta} \frac{\mu}{6} \\ R_\pi^{-1} E (\pi - \pi_P^e)^2 (\bar{\alpha} - \widehat{\alpha}) \\ E (y - \frac{z}{2})^2 - R_y \\ E (\pi - \pi_P^e)^2 - R_\pi \end{bmatrix} \quad (2.19)$$

It is apparent that while the specific form of the loss function does not affect the inference problem of the central bank, it does have a bearing on private sector learning.

Both systems are recursive. $R_\pi \rightarrow E (\pi - \pi_P^e)^2$ from any starting point, which implies that $R_\pi^{-1} E (\pi - \pi_P^e)^2 \rightarrow I$, provided that R_π is invertible along the path. The same goes for R_y . Hence, the stability of the differential equation for $\widehat{\alpha}$ may be assessed regardless of the rest of the system. Conditional on $\widehat{\alpha}$, R_y and R_π approaching the true parameter values, convergence to the REE of private sector expectations is determined on the basis of the eigenvalues of the ODE for π_P^e . Note that the probability limit of the latter does not depend on the information set of the central bank and is the same whether or not the monetary authority knows the full structure of the economy. The conditions for learnability of the REE under both lexicographic and quadratic preferences are stated in the next proposition.

Proposition 3 Assume that the economy has agents that rely on adaptive learning to form expectations; further assume that the central bank has

only incomplete information about the structure of the economy and uses recursive least squares (RLS) to estimate the unknown parameters. Then the asymptotic behaviour of the system is described by (2.18) and (2.19) and, regardless of whether the policymaker has quadratic or lexicographic preferences, the discretionary REE is unique and E-stable: the estimates $(\hat{\alpha}_t, \psi_t)$ converge locally to $(\bar{\alpha}, \bar{\alpha}^2 + \sigma_\alpha^2)$ and the expectations of private agents tend in the limit to the RE values.

Proof See the appendix.

As in the case in which only the private sector learns, the effect of preferences on the speed of convergence is not clear. For small values of β , a central bank setting policy so as to minimise a quadratic loss function seems to be less effective in driving the economy towards the REE; when β is high, the opposite is true. However, positing imperfect knowledge for the central bank adds a layer of interaction between monetary policy and economic outcomes, and the dynamics of the model cannot be properly analysed by focusing only on asymptotic distributions. In particular, when the learning process is disturbed by shocks from several sources, the ODE becomes an acceptable approximation to the stochastic recursive algorithm only for large values of t , and the asymptotic distribution is of little help in understanding the properties of the system. The problem is even more serious in models where there are multiple equilibria, since in such cases, in early time periods, when estimates are based on very few degrees of freedom, large shocks can displace θ_t outside the domain of attraction of the ODE and the system can therefore converge to any of the equilibrium points.¹⁹ It follows then that

¹⁹When there is a unique equilibrium and the ODE is stable, it can be shown that $\theta_t \rightarrow \theta^*$ with probability 1 from any starting point. When there are multiple equilibria, however, such a strong result does not apply, unless one artificially constrains θ_t to an appropriate neighbourhood of the locally stable equilibrium θ^* . In the earlier literature, local convergence was obtained by making an additional assumption about the algorithm, known as the *projection facility*. As a reference, see Evans and Honkapohja (2001), section 6.4.

when the agents' information set is severely constrained, both the asymptotic and the finite sample behaviour of the system are relevant. Theoretical results are therefore no longer sufficient, necessitating simulation experiments and numerical results.

2.4 Imperfect knowledge and policy effectiveness

Model simulations are used to show how learning affects the dynamic properties of inflation, inflation expectations and output. First the performance of the forecasting rules is assessed; then the issue of the relative speed of convergence is considered. The output-inflation variability trade-off under alternative monetary regimes is also assessed: society, whose preferences are quadratic, can appoint a central banker whose loss function is lexicographic (strategy LEX) or quadratic (strategy QUA), depending on which monetary regime is more welfare-enhancing. To account for the finding of Orphanides and Williams (2002), that excessive stress on output stabilisation can produce episodes in which the public's inflation expectations are uncoupled from the policy objective, additional simulations mimic the impact of a string of negative supply shocks on the economy. Finally, as a further check on the extent to which the results depend on the learning mechanism selected, the assumption of infinite memory is dropped and perpetual learning is considered.²⁰

Each experiment runs 500 replications; all simulations cover an interval of 2000 periods; subsamples of 500 observations are also considered in order to estimate the speed of convergence, as measured by the parameter δ , defined in (2.8). Initial conditions for the lagged variables in the RLS algorithm are

²⁰Perpetual learning is sometimes used as a synonym of constant-gain learning. A constant gain algorithm is preferable when the agents believe that the economic environment is subject to frequent structural changes. In such cases, observations from the distant past are no longer informative and can be a source of distortion.

drawn randomly from the distribution corresponding to the RE equilibrium. The results reported in the tables are calculated excluding the first 150 periods, so as to minimise the impact of initial observations, which could be too far away from the equilibrium solution. The model is calibrated according to the estimates in Ellison and Valla (2001); the parameter values selected are reported in Table 1. Concerning β , the relative weight of the inflation objective in the loss function, I consider three values, namely $\beta = \{.1765, 1, 5.666\}$. Under lexicographic preferences, it is assumed that $\bar{\pi}$ is chosen so as to drive inflation expectations to zero.

Table1: Baseline calibrated parameters

Parameter	Value
$\bar{\alpha}$	1.75
σ_{α}	0.5
μ	0.0175
k	0.0029

Tables 2a and 3a report the simulation results. Table 2a describes the “plain” RLS learning rule (unconstrained estimator, UE for short), while Table 3a shows the results for constrained estimation (CE for short). The latter is tantamount to imposing some minimal restriction on the admissible regions of the estimates of $\bar{\alpha}$ and $\bar{\alpha}^2 + \sigma_{\alpha}^2$: in the first case it is assumed that only positive values of $\hat{\alpha}_t$ are admissible, since positive inflation surprises cannot have a negative impact on output; in the second, it is assumed that only values of ψ_t greater than $\hat{\alpha}_t^2$ are sensible, since variances cannot be negative. The constraints act as substitutes for a *projection facility*:²¹ though they

²¹Convergence of the learning process to the REE almost surely holds when there is a unique solution and the ODE is globally stable; in general, convergence with probability 1 is not guaranteed, since the ODE is not a reliable approximation of the stochastic recursive algorithm for small values of t . Almost sure convergence holds only when the algorithm is supplemented with a *projection facility*, i.e. when θ_t is artificially constrained to remain in an appropriate neighbourhood of θ (see section 6.4 and Corollary 6.8 in Evans and Honkapohja (2001)). The hypothesis of a projection facility however is often

cannot guarantee almost sure convergence of the learning algorithm, they can in principle reduce the number of non-convergent replications.

Expected inflation, the estimate of the policy rule coefficient, the standard deviation of output and inflation and the performance index (PI) are reported in the tables, as is the speed of convergence for π_P^e , $\widehat{\alpha}_t$ and ψ_t . The performance index is defined as

$$\frac{Ey_{LEX}^2 + \beta^S E\pi_{LEX}^2}{Ey_{QUA(\beta^S)}^2 + \beta^S E\pi_{QUA(\beta^S)}^2} \quad \text{or alternatively} \quad \frac{Ey_{QUA(\beta)}^2 + \beta^S E\pi_{QUA(\beta)}^2}{Ey_{QUA(\beta^S)}^2 + \beta^S E\pi_{QUA(\beta^S)}^2}$$

depending on which strategy is evaluated. Ey_{LEX}^2 is the second moment of output under strategy LEX and the other terms have a similar meaning; β (i.e. the degree of inflation aversion in the central banker's loss function) is in general different from β^S (i.e. the degree of inflation aversion for society), though the case $\beta = \beta^S$ is also considered. The index is equal to the ratio between the social loss function under the delegated policymaker and the loss that would be obtained by appointing a central banker with the same preferences as society's: when $PI < 1$ it is efficient to delegate to a policymaker with a utility function different from society's. For ease of comparison, the values of the performance index are also shown for the REE.

Regardless of central bank type, the estimates of expected inflation are precise but biased downward: the higher the values of β , the greater the accuracy of the estimate of π_P^e . Imposing constraints on the support of $\bar{\alpha}$ and $\bar{\alpha}^2 + \sigma_\alpha^2$ diminishes the bias in the case of lexicographic preferences, but increases it in the case of quadratic utility. One of the most striking findings is the modest effect on the variability of output and inflation of assuming imperfect central bank knowledge: the increase in volatility with respect to REE is surprisingly small, in most cases just a few percentage points. This is remarkable, since the estimation problem faced by the monetary authority is quite convoluted,

criticized, however, because it is not easily justified on economic grounds and is clearly inappropriate for decentralised markets.

involving non-linearities and the computing of higher order moments. In the vast majority of cases, the increase in the volatility of output and inflation is well below 10% regardless of the preferences of the central bank, suggesting that the cost for the policymaker of having only partial knowledge of the working of the economy is not disproportionately large. It is worth remembering, though, that the model lacks intrinsic dynamics, which explains why deviations from the REE tend to be short-lived.

The analysis of the performance index provides several insights into the relative efficiency of the two strategies. The first finding is that with decreasing gain learning strategy LEX performs well regardless of the exact value of β^S : under UE it is at least as good as strategy QUA unless the degree of inflation aversion is extremely great; under CE it is uniformly better. The second finding is that the shorter the memory of the learning process, the poorer the performance of strategy LEX, possibly because of the non-linearity of the policy rule: when the policymaker has no discretion, the inflation surprises do not yield information about the output-inflation trade-off and the recursive estimates are less accurate, in particular when the sample is small. The third finding is that the relative ranking of the two strategies is the same regardless of the way expectations are formed, which suggests that the learning process converges quite quickly to the REE. Closer inspection of the simulation evidence shows that for low values of β , a central bank with lexicographic preferences is more effective in keeping inflation expectations under control and is also successful in stabilising output fluctuations, even when bounds are not imposed on the RLS algorithm. When β is high the situation is reversed. A downward bias is evident in the estimate of the parameter that measures the response to the signal z , but it mostly disappears in the CE case; the imprecision in guessing the value of $\frac{1}{\phi}$ is responsible for some undesired fluctuations in output, while the excessive volatility of inflation is not attributable to the surprise component, but rather to movements in

private sector inflation expectations, which under adaptive learning are not constant as in the REE. Since the bounds imposed on the RLS algorithm mimic the working of a projection facility, the rejection rate in the CE case turns out to be substantially lower.²² Indeed, because of the complexity of the filtering problem that the monetary authority faces, in a large number of replications shocks displace the recursive algorithm outside of the domain of attraction of the ODE and the estimate of the optimal response coefficient in the policy rule remains far off the true value. If the estimate of ϕ is very large at time t , the monetary authority has no incentive to respond aggressively to the signal z and changes in y mostly reflect output shocks ε : the data become uninformative about the output-inflation trade-off and the estimate of ϕ gets larger and larger. Expectations become self-fulfilling and the economy gets stuck indefinitely on a suboptimal path, characterised by excessively passive monetary policy, as if the policymaker's degree of inflation aversion were enormously higher than society's.

Except for high β s, strategy QUA is outperformed by strategy LEX, but it is more effective in enhancing agents' learning process, as we can see from the more precise estimate of the policy rule coefficient, which guarantees that the equilibrium under imperfect knowledge matches the REE very closely. There is no clear evidence that excessively low or high inflation aversion has repercussions on the accuracy of the estimate of $\frac{\rho}{\phi}$. As to the speed of con-

²²The rejection rate is calculated on the basis of the estimated value of the second moment of the output-inflation trade-off. Replications are considered as diverging if the estimate of $\bar{\alpha}^2 + \sigma_\alpha^2$ is at least three times the true value. The first 150 observations are not used. In the case of lexicographic preferences, absent constraints, the rejection rate turns out to be quite high (some 20%); it falls by a factor of 4 if the RLS algorithm is augmented with lower bounds. In the case of quadratic preferences, the number of diverging replications is on average much smaller and so is the gain obtained by imposing constraints on the learning process; for high β s, however, the rejection rate rises towards that observed under lexicographic preferences. The estimated number of diverging replications decreases significantly if less restrictive criteria are used. Notice that divergence pertains to central bank learning and is defined in terms of the estimates of the policy parameters $\frac{1}{\phi}$ and $\frac{\rho}{\phi}$, which become very close to zero: neither the output gap nor the inflation rate actually deviates boundlessly from equilibrium.

vergence, the two rules are roughly equivalent, although under CE strategy LEX seems to be preferable; δ is very close to one half, so that convergence to a Gaussian distribution of both sequences $\{\theta_t\}^{LEX}$ and $\{\theta_t\}^{QUA}$ cannot be ruled out. As expected, strategy QUA reduces output variability more than strategy LEX when β is low, while the opposite holds when inflation aversion is very high. The simulation results confirm that when learning is posited in lieu of rational expectations, a bias against activist policies arises. Hawkish policies are welfare-enhancing because they offset the negative impact of two distortions: (i) the policymaker's desire to push output above the natural level and (ii) the uncoupling of expectations from policy objectives. Whereas the latter distortion is considered in Orphanides and Williams (2002), the former is not.

Tables 2b and 3b present evidence for the case of perpetual learning; the statistics for speed of convergence are not shown, of course, since under constant-gain learning θ_t may converge at most to a probability distribution, not to a non-stochastic point. No meaningful differences from the previous case are apparent. Given the structure of the model, there is no benefit in discarding observations, so it is no surprise that in most cases RLS estimates are less accurate and policies are less successful in stabilising both output and inflation; the deterioration in policy effectiveness seems to be relatively greater for strategy LEX.

The evidence suggests that a benevolent government may be better off appointing a central banker whose preferences are lexicographic if the degree of inflation aversion is not known with certainty or if it changes over time, since strategy LEX can very nearly maximise welfare for a large set of values of β^S . Strategy LEX can be implemented by giving the central bank a mandate specifying an upper (and possibly a lower) bound on inflation and not requiring the government to find the perfect policymaker with the right

preferences.²³

An additional set of simulations were run to analyse the dynamic response of output and inflation to a sequence of unanticipated shocks. The experiment postulates that the economy is hit by a string of negative output shocks that decline gradually in magnitude and vanish after 12 periods. With rational expectations, the impact of the shocks is short-lived and causes only a temporary fall in output and a rise in inflation, while under imperfect knowledge the response of the economy is prolonged and amplified by agents' learning. The experiment tests whether the finding of Orphanides and Williams (2002) that activist policies ultimately cause the perceived process for inflation to be uncoupled from the policymaker's objectives, is generally applicable and extends also to the theoretical framework adopted here.

Tables 4a and 4b report the outcome of the experiment. Regardless of the central bank's preferences, activist policies would not appear to pay off: the lower β , the more volatile inflation and output, especially the latter. The simulation supports the Rogoff's thesis that it is welfare-improving to appoint a central banker who attaches greater relative importance to the inflation objective than society does. No policy can effectively offset the impact on economic activity of a sequence of negative output shocks: it still pays off to be hawkish, but the attempts to reduce inflation volatility translate into output fluctuations that are much wider than under rational expectations. What worsens the performance of monetary policy is the uncoupling between expectations and policy targets: since expectation depends on past values of inflation, they cannot be easily anchored, unless the policymaker acts like an inflation nutter. According to the performance index, strategy LEX is highly

²³This feature also characterises inflation zone targeting, see Mishkin and Westelius (2006). It is easy to see that when the cost of overshooting the upper or undershooting the lower bound is extremely large ($C \rightarrow \infty$ in Mishkin's and Westelius' notation) and the policymaker does not care greatly about inflation variability inside the range ($\omega_\pi \rightarrow 0$), strategy LEX and inflation zone targeting become more and more alike.

successful in promoting social welfare. In the UE case, except for very high values of β , it is more effective than strategy QUA, and in the CE case, it is uniformly better. A central banker endowed with lexicographic preferences ensures both robustness and effectiveness of monetary policy.

2.5 Conclusions

The essay presented in this chapter focuses on the implications for the effectiveness of monetary policymaking of discarding the assumption of rational expectations and applies a principal-agent approach to deal with the time-inconsistency problem that arises when the central bank cannot commit. It is assumed that society can delegate monetary policy to a central banker with either quadratic or lexicographic preferences. Special attention is paid to the latter case, which seems a better fit with the objectives of inflation-targeting central banks. The main focus is on validating the hypothesis that policies designed to be efficient under rational expectations can perform very poorly when knowledge is incomplete and agents learn adaptively.

The evidence produced confirms that when agents do not possess complete knowledge on the structure of the economy and rely instead on an adaptive learning technology, a bias toward conservatism arises, suggesting that society is better off by appointing a policymaker whose degree of inflation aversion is higher than its own. The explanation is that agents' and policymakers' attempts to learn adaptively introduce inertia into the system and induce prolonged deviations of output and inflation from target, thereby raising the costs for the central bank of not responding promptly and forcefully to shocks. The chapter also shows that the strategy of implementing a lexicographic preference ordering performs very well, on average. It comes close to maximising social welfare for a wide range of values of β^S and outperforms the quadratic-preference strategies unless society is extremely inflation

averse.

These findings tally closely with those of Orphanides and Williams (2002), which is surprising given the differences in the theoretical framework. First, the model adopted here has no intrinsic dynamics, and the only source of persistence is the assumption that agents learn adaptively: the uncoupling between actual and perceived inflation is much less probable with such a simple dynamic structure, though presumably the lack of dynamics in the economy is compensated for by the inertia induced by the efforts of the central bank to estimate the mean and variance of the output-inflation trade-off. Second, though only inflation expectations have a direct impact on the equilibrium outcome, output gap uncertainty does affect the central bank's estimates of the moments of α and hence the policy setting: it is by no means obvious that a strategy that penalises output variability will be conducive to higher welfare. The justification for the bias in favour of hawkish policies lies in the role of central bank learning: excessively activist policies reduce the information content of the output gap and make estimates of the coefficients of the policy rule too volatile and unreliable.

2.6 Appendix

Proof of Lemma 1. Consider first two random variables, u and v , defined on the unit segment $[0, 1]$. Their sum, $w = u + v$, is defined on the close interval $[0, 2]$. The distribution function of w , for $0 \leq w \leq 1$, is given by $H(w) = \int_0^w du \int_0^{w-u} dv = \frac{w^2}{2}$, while, for value of w comprised in the interval $(1, 2]$, it is equal to $H(w) = \int_0^w du \int_0^{w-u} dv - \int_1^w du \int_0^{w-u} dv - \int_1^w dv \int_0^{w-v} du = 1 - \frac{(2-w)^2}{2}$. The corresponding density function is $h(w) = w$ for $0 \leq w \leq 1$ and $h(w) = 2 - w$ for $1 < w \leq 2$, or, more compactly,

$$h(w) = \min[w, 1] - \max[0, w - 1].$$

Consider now the case in which the two random variables, rather than having support on the unit interval, are both defined on $[-\mu, \mu]$. One can write $\varepsilon = -\mu + 2\mu u$ and $\xi = -\mu + 2\mu v$; their sum, $z = \varepsilon + \xi = -2\mu + 2\mu(u + v) = -2\mu + 2\mu w$, is a linear transformation of the random variable w , i.e. $z = g(w)$. One can use the change-of-variable technique to compute the density function $f(\cdot)$ of the variable z , i.e. $f(z) = \left| \frac{d}{dz} g^{-1}(z) \right| h(g^{-1}(z))$. Since $g^{-1}(z) = 1 + \frac{z}{2\mu}$ and $\left| \frac{d}{dz} g^{-1}(z) \right| = \frac{1}{2\mu}$, one has that

$$\begin{aligned} \text{for } -2\mu \leq z \leq 0 \quad f(z) &= \left(1 + \frac{z}{2\mu}\right) \left| \frac{d}{dz} \left(1 + \frac{z}{2\mu}\right) \right| = \frac{1}{2\mu} + \frac{z}{4\mu^2} \\ \text{for } 0 \leq z \leq 2\mu \quad f(z) &= \left[2 - \left(1 + \frac{z}{2\mu}\right)\right] \left| \frac{d}{dz} \left(1 + \frac{z}{2\mu}\right) \right| = \frac{1}{2\mu} - \frac{z}{4\mu^2} \end{aligned}$$

which can be written in a more compact way as $f(z) = \frac{1}{2\mu} + \frac{1}{4\mu^2} [\min(z, 0) - \max(0, z)]$. The corresponding distribution function is

$$\begin{aligned} \text{for } -2\mu \leq z \leq 0 \quad F(z) &= \frac{1}{8\mu^2} (2\mu + z)^2 \\ \text{for } 0 \leq z \leq 2\mu \quad F(z) &= 1 - \frac{1}{8\mu^2} (2\mu - z)^2 \end{aligned}$$

Proof of Proposition 1. By definition, $E(\varepsilon|z) = \int_D \varepsilon f(\varepsilon|z) d\varepsilon = \int_D \varepsilon \frac{f_1(z|\varepsilon)f_2(\varepsilon)}{f_3(z)} d\varepsilon$,

where $f_i(\cdot)$, $i = 1, 2, 3$, denotes a (conditional or marginal) density function, and D is the domain of $\varepsilon|z$, which clearly depends on the current realisation of the signal z . The density function of z conditional on ε is the same as the density function of ξ , which is $\frac{1}{2\mu}$, while the probability law of z , as shown in Lemma 1, is $f_3(z) = \frac{1}{2\mu} + \frac{1}{4\mu^2} [\min(z, 0) - \max(0, z)]$. It follows that

$$\begin{aligned} E(\varepsilon|z) &= \int_D \varepsilon f(\varepsilon|z) d\varepsilon = \int_D \varepsilon \frac{\frac{1}{2\mu} \cdot \frac{1}{2\mu}}{\frac{1}{2\mu} + \frac{1}{4\mu^2} [\min(z, 0) - \max(0, z)]} d\varepsilon \\ &= \int_D \frac{\varepsilon}{2\mu + \min(z, 0) - \max(0, z)} d\varepsilon \end{aligned}$$

The remaining problem is to find D , the support of $\varepsilon|z$. Since $\varepsilon = z - \xi$, a given value \bar{z} of the signal shifts the support of the random variable ε , which

becomes $D = [-\mu, \mu] \cap [-\mu + \bar{z}, \mu + \bar{z}]$. Two cases are possible, depending on whether \bar{z} is positive or negative. If $z = -\bar{z}_1 < 0$, then $\varepsilon \in [-\mu, -\bar{z}_1 + \mu]$ and

$$E(\varepsilon|z) = \int_{-\mu}^{-\bar{z}_1+\mu} \frac{\varepsilon}{2\mu-\bar{z}_1} d\varepsilon = \frac{1}{2\mu-\bar{z}_1} \frac{\varepsilon^2}{2} \Big|_{-\mu}^{-\bar{z}_1+\mu} = \frac{1}{2\mu-\bar{z}_1} \frac{(-\bar{z}_1+\mu)^2 - \mu^2}{2} =$$

$$\frac{1}{2\mu-\bar{z}_1} \frac{\bar{z}_1(\bar{z}_1-2\mu)}{2} = -\frac{\bar{z}_1}{2} = \frac{z}{2}$$

If, instead, $z = \bar{z}_2 > 0$, then $\varepsilon \in [-\mu + \bar{z}_2, \mu]$ and

$$E(\varepsilon|z) = \int_{-\mu+\bar{z}_2}^{\mu} \frac{\varepsilon}{2\mu-\bar{z}_2} d\varepsilon = \frac{1}{2\mu-\bar{z}_2} \frac{\varepsilon^2}{2} \Big|_{-\mu+\bar{z}_2}^{\mu} = \frac{1}{2\mu-\bar{z}_2} \frac{\mu^2 - (\bar{z}_2-\mu)^2}{2} =$$

$$\frac{1}{2\mu-\bar{z}_2} \frac{\bar{z}_2(2\mu-\bar{z}_2)}{2} = \frac{\bar{z}_2}{2} = \frac{z}{2}$$

Notwithstanding the distribution of the variables is not normal, the optimal estimate for the unobserved shock ε is the same as in the (standard) case of a Gaussian variable.

Proof of Proposition 2. If the output shock is not too unfavourable, the central bank's problem has an internal solution, obtained from the first order condition

$$E[\alpha(\alpha(\pi - \pi^e) + \varepsilon - k)] = E(\alpha^2)(\pi - \pi^e) + E(\alpha)\left(\frac{z}{2} - k\right) = 0$$

The optimal inflation rate is therefore

$$\pi = \pi^e - \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2} \left(\frac{z}{2} - k\right) = \pi^e - \frac{1}{\phi} (z - 2k)$$

where $\phi^{-1} \equiv \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2} \frac{1}{2}$. If instead z signals a value of ε close to $-\mu$, the value of π minimising the loss function is not admissible and the central bank will choose $\pi = \bar{\pi}$. The optimal strategy for the monetary authority is therefore the one described by (2.5), i.e. $\pi = \pi^e - \frac{1}{\phi} (\max(z, \Lambda) - 2k)$, where $\Lambda = 2k + \phi(\pi^e - \bar{\pi})$.

The upper bound on inflation implies that $\pi = \pi^e - \frac{1}{\phi}(z - 2k) \leq \bar{\pi}$, which holds for $z \in [2k + \phi(\pi^e - \bar{\pi}), 2\mu]$: the higher k , the higher expected inflation and the smaller the probability that the central bank succeeds in stabilising output. Since the density function of the signal has a kink at zero, it matters for the computation of expected inflation whether $2k + \phi(\pi^e - \bar{\pi})$ is positive or negative, for its value determines how the support of the signal is split.

In a RE equilibrium, beliefs are on average correct, implying that $\pi^e = E\pi$, which holds when $2k = E[\max(z, \Lambda)]$. When the previous condition is satisfied, Λ is determined and so is π^e . The proof of the existence and uniqueness of the RE equilibrium runs as follows. First, let us rewrite the left-hand side of the previous equality as $2k = \Lambda - \phi(\pi^e - \bar{\pi}) \geq \Lambda$ (as $\bar{\pi} > \pi^e$ for $k < \mu$ and $\bar{\pi} = \pi^e$ otherwise). Then, let us divide everything by 2μ , so that the equilibrium condition becomes:

$$\lambda - \frac{\phi(\pi^e - \bar{\pi})}{2\mu} = E(\max[x, \lambda]) \quad (\text{A1})$$

where $x = z/(2\mu)$ is a random variable with support on the interval $[-1, 1]$ and $\lambda = \Lambda/(2\mu)$.

A straightforward application of the change-of-variable formula gives the density function of x , which is $f(x) = 1 + [\min(x, 0) - \max(0, x)] = 1 - |x|$. In order to compute $E(\max[x, \lambda])$ it is necessary to study how the function $\max[x, \lambda]$ changes as λ increases from $-\infty$ to $+\infty$. Three cases are to be distinguished: (1) if $\lambda \leq -1$, $\max[x, \lambda] = x$ for all x in $[-1, 1]$; (2) if $\lambda \in (-1, 1)$, $\max[x, \lambda] = \lambda$ for $-1 < x < \lambda$ and $\max[x, \lambda] = x$ for

$\lambda \leq x < 1$; (3) if $\lambda \geq 1$, $\max[x, \lambda] = \lambda$. Accordingly,

$$E(\max[x, \lambda]) = \begin{cases} I(\lambda < -1) \cdot \int_{-1}^1 x(1 - |x|) dx \\ + I(-1 \leq \lambda < 0) \cdot \left\{ \int_{-1}^{\lambda} \lambda(1 + x) dx + \int_{\lambda}^1 x(1 - |x|) dx \right\} \\ + I(0 \leq \lambda \leq 1) \cdot \left\{ \int_{-1}^{\lambda} \lambda(1 - |x|) dx + \int_{\lambda}^1 x(1 - x) dx \right\} \\ + I(\lambda > 1) \cdot \lambda \end{cases}$$

where $I(\cdot)$ is the indicator function. After computing the integrals, the previous expression for $E(\max[x, \lambda])$ simplifies to

$$\begin{aligned} E(\max[x, \lambda]) &= 0 \cdot I(\lambda < -1) + \frac{1}{6}(1 + \lambda)^3 \cdot I(-1 \leq \lambda < 0) \\ &+ \frac{1}{6}(1 + 3\lambda + 3\lambda^2 - \lambda^3) \cdot I(0 \leq \lambda \leq 1) + \lambda \cdot I(\lambda > 1) \end{aligned}$$

The intersection between $E(\max[x, \lambda])$ and the straight line $\lambda - \frac{\phi(\pi^e - \bar{\pi})}{2\mu}$ detects a value of λ – which depends on k – that pinpoints the RE equilibrium: π^e is obtained as $\pi^e = \bar{\pi} + \frac{2\mu}{\phi} \left(\lambda(k) - \frac{2k}{2\mu} \right)$.

The equilibrium exists because the left-hand side of equation (A1) is below the right-hand side for low values of λ , while it is above it for high values of λ : as both functions are continuous, they must cross at least once. The equilibrium is unique because the slope of $E[\max(x, \lambda)]$ is never greater than that of $\lambda - \frac{\phi(\pi^e - \bar{\pi})}{2\mu}$ and accordingly there cannot be more than one intersection.²⁴

In the general case, to determine π^e it is necessary to solve a third-order polynomial, but the analysis is greatly simplified if one sets $2k + \phi(\pi^e - \bar{\pi}) = 0$ equal to zero. Under this assumption, expected inflation is the solution to the following equation

²⁴For $\lambda \leq -1$ the slope of $E[\max(x, \lambda)]$ is 0; for $\lambda \in (-1, 1)$ it is positive but below 1; for $\lambda \geq 1$ it is equal to 1. The slope of $\lambda - \frac{\phi(\pi^e - \bar{\pi})}{2\mu}$ on the other hand is always equal to 1. The only case when there are multiple solutions (actually infinite ones) is when the functions $E[\max(x, \lambda)]$ and $\lambda - \frac{\phi(\pi^e - \bar{\pi})}{2\mu}$ overlap for $\lambda \geq 1$, which happens when $k \geq 2\mu$, implying $\pi = \pi^e = \bar{\pi}$ for all values of the supply shock ε .

$$\begin{aligned}
\pi^e &= \int_{-2\mu}^0 \bar{\pi} \left(\frac{1}{2\mu} + \frac{z}{4\mu^2} \right) dz + \int_0^{2\mu} \left(\pi^e - \frac{z-2k}{\phi} \right) \left(\frac{1}{2\mu} - \frac{z}{4\mu^2} \right) dz \\
&= \frac{\bar{\pi}}{2} + \frac{1}{2} \left(\pi^e + \frac{2k}{\phi} \right) - \int_0^{2\mu} \frac{z}{\phi} \left(\frac{1}{2\mu} - \frac{z}{4\mu^2} \right) dz \\
&= \frac{\bar{\pi}}{2} + \frac{1}{2} \left(\pi^e + \frac{2k}{\phi} \right) - \frac{\mu}{3\phi}
\end{aligned}$$

Expected inflation is therefore $\pi^e = \bar{\pi} + \frac{2k}{\phi} - \frac{2\mu}{3\phi}$, which simplifies to

$$\pi^e = \bar{\pi} - \frac{\mu}{3\phi}$$

once k is substituted out using the restriction $2k + \phi(\pi^e - \bar{\pi}) = 0$. The implied value of the target level of output is $k = \frac{\phi(\bar{\pi} - \pi^e)}{2} = \frac{\mu}{6}$.

Proof of Proposition 3. Regardless of the preferences of the monetary authority, the recursive system representing the learning process is of the form $\theta_t = \theta_{t-1} + \frac{1}{t}Q(\theta_{t-1}, X_t)$, where $\theta_t = \left(\pi_{Pt}^e, \widehat{\alpha}_t, R_{y,t}, R_{\pi,t} \right)'$ and $X_t = (1, \alpha_t, z_t, \varepsilon_t)$. To show the asymptotic stability of the REE under learning, the procedure is the following: first, it must be verified that there exists a non-trivial open domain that contains the equilibrium point and in which the learning algorithm satisfies a few regularity conditions concerning the updating function $Q(\theta_{t-1}, X_t)$ and the stochastic process driving the state variables $X_t(\theta_{t-1})$; second, the local (or global) stability of the ODE associated to the stochastic recursive system must be established.

Consider first the case of lexicographic preferences. The system (2.16) has a unique equilibrium point θ^* , where $\pi_P^e = \bar{\pi} - \frac{\mu/3}{\phi}$, $\widehat{\alpha} = \bar{\alpha}$, $R_\pi = \frac{2}{3\phi^2} \left(\frac{\mu^2}{3} \right)$ and $R_y = (\bar{\alpha}^2 + \sigma_\alpha^2) \left[\frac{2}{3\phi^2} \left(\frac{\mu^2}{3} \right) \right] + \frac{1}{2} \frac{\mu^2}{3}$. It can be easily seen that θ^* is the REE. The stochastic process $X_t(\theta_{t-1})$ is white noise, with finite absolute moments, so that regularity conditions (B.1) and (B.2) in Evans and Honkapohja (2001) are satisfied.²⁵ In addition, the gain sequence approaches zero asymptotically and is not summable. Finally, provided that R_π and R_y are non

²⁵Chapter 6 in Evans and Honkapohja (2001) lists the regularity conditions required for the assessment

zero along the learning path, $Q(\theta_{t-1}, X_t)$ satisfies a Lipschitz condition²⁶ on a compact set containing the equilibrium point θ^* , which ensures that regularity conditions (A.1)-(A.3) in Evans and Honkapohja (2001) are also met. Convergence of the learning process to the REE hinges therefore on the stability of the associated ODE (2.18). The system is recursive and the asymptotic behaviour of the subsystem describing central bank learning can be assessed independently of the expectations formation mechanism of the private agents. Indeed, provided that R_π and R_y are invertible along the convergence path, $R_y \rightarrow E(y - \frac{z}{2})^2$ and $R_\pi \rightarrow E(\pi - \pi_P^e)^2$ from any starting point. As $R_\pi^{-1} E(\pi - \pi_P^e)^2 \rightarrow I$, it is easily seen that $\hat{\alpha}_t \rightarrow \bar{\alpha}$, since the eigenvalue of the Jacobian of the corresponding differential equation has a negative real part. Conditional on $\hat{\alpha}_t \rightarrow \bar{\alpha}$, convergence of private sector inflation forecasts follows, since the associated ODE is stable.

A more formal proof of the convergence of the learning process to the REE requires proving that the Jacobian of the ODE, evaluated at the REE θ^* , has eigenvalues whose real parts are negative. In order to show that this is indeed the case, first note that $R_\pi = E(\pi - \pi_P^e)^2$ at θ^* , implying that R_π does not appear in the first three equations of the ODE evaluated at θ^* . A similar result holds for R_y . In the last two equations, the derivatives of R_π and R_y (and, accordingly, of $E(\pi - \pi_P^e)^2$ and $E(y - \frac{z}{2})^2$) cancel out, so that the Jacobian has the following upper triangular, block-recursive structure:

$$Dh(\theta^*) = \begin{bmatrix} -\frac{1}{2} & -\frac{1}{\bar{\alpha}^2 + \sigma_\alpha^2} \frac{\mu}{12} & \frac{3}{2\mu} \frac{1}{\bar{\alpha}} & -\frac{3}{2\mu} \frac{\bar{\alpha}^2 + \sigma_\alpha^2}{\bar{\alpha}} \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

It is easily checked that its eigenvalues are $(-\frac{1}{2}, -1, -1, -1)$. They are all

of the asymptotic behaviour of the stochastic recursive algorithm. Local stability is treated in section 6.2, while global convergence is analysed in section 6.7.

²⁶ $Q(\theta_{t-1}, X_t)$ satisfies a Lipschitz condition if it is bounded and twice continuously differentiable, with bounded second derivatives.

negative and the system is therefore E-stable.

Consider now the case of quadratic preferences. The unique equilibrium point θ^* of the system (2.16) is now $\pi_P^e = \frac{\bar{\alpha}}{\beta}k$, $\hat{\alpha} = \bar{\alpha}$, $R_\pi = E(\pi - \pi_P^e)^2 = 2\left(\frac{\rho}{\phi}\right)^2\left(\frac{\mu^2}{3}\right)$ and $R_y = E(y - \frac{z}{2})^2 = (\bar{\alpha}^2 + \sigma_\alpha^2)\left[2\left(\frac{\rho}{\phi}\right)^2\left(\frac{\mu^2}{3}\right)\right] + \frac{1}{2}\frac{\mu^2}{3}$. The stochastic process $X_t(\theta_{t-1})$ is independent of θ and is the same as in the previous case, so that regularity conditions (B.1) and (B.2) in Evans and Honkapohja (2001) are satisfied. The same holds for the assumptions (A.1)-(A.3) on the gain sequence and the updating function $Q(\theta_{t-1}, X_t)$. The stability of the associated ODE (2.19) can be proved in the same way as for the system (2.18). Provided that R_π and R_y are invertible along the convergence path, $R_\pi \rightarrow E(\pi - \pi_P^e)^2$ from any starting point; central bank estimates converge to the true parameter values $\bar{\alpha}$, since the eigenvalue of the Jacobian of the corresponding differential equation has a negative real part, and $\pi_P^e \rightarrow \frac{\bar{\alpha}}{\beta}k$, since the associated ODE is stable.

As in the previous case, the structure of the Jacobian justifies the sequential solution of the system. At θ^* , the derivative matrix of the ODE (2.19) is equal to

$$Dh(\theta^*) = \begin{bmatrix} -\frac{\beta}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta} & \frac{1}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta} \frac{\mu}{6} & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

The lower block for R_π and R_y can be solved first; then, triangularity of the upper block ensures that convergence for $\hat{\alpha}_t$ does not depend on the asymptotic behaviour of π_{Pt}^e . The eigenvalues of the Jacobian are $\left(-\frac{\beta}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta}, -1, -1, -1\right)$ and the system is therefore E-stable.

Table 2a - Least squares learning and the volatility of output and inflation*(unconstrained estimator - decreasing gain sequence)*

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation, the performance index and the rate at which estimates of π^e_p , α and ψ converge to the REE. The performance index is the ratio between the social loss function under the delegated policymaker and that that would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to have infinite memory, implying a decreasing gain sequence. Recursive least squares estimates are unconstrained (UE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic preferences		Quadratic preferences	
	RE	T=2000	RE	T=2000
$\beta = 0.176$				
mean π^e_p	0.0000	-0.0005	0.0289	0.0283
policy rule coefficient	0.2642	0.1807	0.2508	0.2308
output variability	0.0084	0.0088	0.0074	0.0076
inflation variability	0.0022	0.0023	0.0291	0.0298
Performance index: $\beta^S=0.176$	0.3494	0.3628	1.0000	1.0000
$\beta^S=1.0$	0.8203	0.8453	9.8428	9.7605
$\beta^S=5.667$	1.0925	1.1462	54.5380	55.0000
convergence speed	π^e_p : $\delta = 0.3351$		π^e_p : $\delta = 0.3570$	
	α : $\delta = 0.4951$		α : $\delta = 0.2435$	
	ψ : $\delta = 0.4609$		ψ : $\delta = 0.4884$	
$\beta = 1.0$				
mean π^e_p	0.0000	-0.0005	0.0051	0.0048
policy rule coefficient	0.2642	0.1807	0.2029	0.1844
output variability	0.0084	0.0088	0.0076	0.0078
inflation variability	0.0022	0.0023	0.0059	0.0059
Performance index: $\beta^S=0.176$	0.3494	0.3628	0.3105	0.3168
$\beta^S=1.0$	0.8203	0.8453	1.0000	1.0000
$\beta^S=5.667$	1.0925	1.1462	2.8307	2.8287
convergence speed	π^e_p : $\delta = 0.3351$		π^e_p : $\delta = 0.2185$	
	α : $\delta = 0.4951$		α : $\delta = 0.2728$	
	ψ : $\delta = 0.4609$		ψ : $\delta = 0.4892$	
$\beta = 5.667$				
mean π^e_p	0.0000	-0.0005	0.0009	0.0008
policy rule coefficient	0.2642	0.1807	0.0975	0.0775
output variability	0.0084	0.0088	0.0086	0.0088
inflation variability	0.0022	0.0023	0.0017	0.0016
Performance index: $\beta^S=0.176$	0.3494	0.3628	0.3627	0.3680
$\beta^S=1.0$	0.8203	0.8453	0.8319	0.8340
$\beta^S=5.667$	1.0925	1.1462	1.0000	1.0000
convergence speed	π^e_p : $\delta = 0.3351$		π^e_p : $\delta = 0.4080$	
	α : $\delta = 0.4951$		α : $\delta = 0.3642$	
	ψ : $\delta = 0.4609$		ψ : $\delta = 0.4891$	

Table 2b - Least squares learning and the volatility of output and inflation*(unconstrained estimator - constant gain sequence)*

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation and the performance index. The performance index is the ratio between the social loss function under the delegated policymaker and that that would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to use a finite number of observations in computing RLS estimates, implying a constant gain sequence. Recursive least squares estimates are unconstrained (UE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic preferences		Quadratic preferences	
	RE	T=2000	RE	T=2000
$\beta = 0.176$				
mean π^e_p	0.0000	-0.0005	0.0289	0.0288
policy rule coefficient	0.2642	0.2084	0.2508	0.2001
output variability	0.0084	0.0090	0.0074	0.0076
inflation variability	0.0022	0.0026	0.0291	0.0291
Performance index: $\beta^S=0.176$	0.3494	0.3974	1.0000	1.0000
$\beta^S=1.0$	0.8203	0.9326	9.8428	1.2276
$\beta^S=5.667$	1.0925	1.2008	54.5380	48.8786
$\beta = 1.0$				
mean π^e_p	0.0000	-0.0005	0.0051	0.0051
policy rule coefficient	0.2642	0.2084	0.2029	0.4747
output variability	0.0084	0.0090	0.0076	0.0077
inflation variability	0.0022	0.0026	0.0059	0.0059
Performance index: $\beta^S=0.176$	0.3494	0.3974	0.3105	0.3163
$\beta^S=1.0$	0.8203	0.9326	1.0000	1.0000
$\beta^S=5.667$	1.0925	1.2008	2.8307	2.5821
$\beta = 5.667$				
mean π^e_p	0.0000	-0.0005	0.0009	0.0009
policy rule coefficient	0.2642	0.2084	0.0975	0.0503
output variability	0.0084	0.0090	0.0086	0.0090
inflation variability	0.0022	0.0026	0.0017	0.0018
Performance index: $\beta^S=0.176$	0.3494	0.3974	0.3627	0.3944
$\beta^S=1.0$	0.8203	0.9326	0.8319	0.8952
$\beta^S=5.667$	1.0925	1.2008	1.0000	1.0000

Table 3a - Least squares learning and the volatility of output and inflation
(constrained estimator - decreasing gain sequence)

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation, the performance index and the rate at which estimates of π^e_p , α and ψ converge to the REE. The performance index is the ratio between the social loss function under the delegated policymaker and that that would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to have infinite memory, implying a decreasing gain sequence. Recursive least squares estimates are constrained to belong to a subset of the parameter space (CE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic preferences		Quadratic preferences	
	RE	T=2000	RE	T=2000
$\beta = 0.176$				
mean π^e_p	0.0000	-0.0003	0.0289	0.0267
policy rule coefficient	0.2642	0.2365	0.2508	0.2411
output variability	0.0084	0.0087	0.0074	0.0075
inflation variability	0.0022	0.0025	0.0291	0.0271
Performance index: $\beta^S=0.176$	0.3494	0.4153	1.0000	1.0000
$\beta^S=1.0$	0.8203	0.9092	9.8428	8.7535
$\beta^S=5.667$	1.0925	0.7491	54.5380	28.4679
convergence speed	π^e_p : $\delta = 0.4193$ α : $\delta = 0.4670$ ψ : $\delta = 0.4765$		π^e_p : $\delta = 0.2972$ α : $\delta = 0.4557$ ψ : $\delta = 0.4892$	
$\beta = 1.0$				
mean π^e_p	0.0000	-0.0003	0.0051	0.0047
policy rule coefficient	0.2642	0.2365	0.2029	0.1913
output variability	0.0084	0.0087	0.0076	0.0077
inflation variability	0.0022	0.0025	0.0059	0.0056
Performance index: $\beta^S=0.176$	0.3494	0.4153	0.3105	0.3475
$\beta^S=1.0$	0.8203	0.9092	1.0000	1.0000
$\beta^S=5.667$	1.0925	0.7491	2.8307	1.5986
convergence speed	π^e_p : $\delta = 0.4193$ α : $\delta = 0.4670$ ψ : $\delta = 0.4765$		π^e_p : $\delta = 0.1616$ α : $\delta = 0.4525$ ψ : $\delta = 0.4918$	
$\beta = 5.667$				
mean π^e_p	0.0000	-0.0003	0.0009	0.0014
policy rule coefficient	0.2642	0.2365	0.0975	0.0805
output variability	0.0084	0.0087	0.0086	0.0105
inflation variability	0.0022	0.0025	0.0017	0.0026
Performance index: $\beta^S=0.176$	0.3494	0.4153	0.3627	0.6033
$\beta^S=1.0$	0.8203	0.9092	0.8319	1.2990
$\beta^S=5.667$	1.0925	0.7491	1.0000	1.0000
convergence speed	π^e_p : $\delta = 0.4193$ α : $\delta = 0.4670$ ψ : $\delta = 0.4765$		π^e_p : $\delta = 0.3331$ α : $\delta = 0.4544$ ψ : $\delta = 0.4870$	

Table 3b - Least squares learning and the volatility of output and inflation
(constrained estimator - constant gain sequence)

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation and the performance index. The performance index is the ratio between the social loss function under the delegated policymaker and that that would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to use a finite number of observations in computing RLS estimates, implying a constant gain sequence. Recursive least squares estimates are constrained to belong to a subset of the parameter space (CE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic preferences		Quadratic preferences	
	RE	T=2000	RE	T=2000
$\beta = 0.176$				
mean π^e_p	0.0000	-0.0003	0.0289	0.0288
policy rule coefficient	0.2642	0.1917	0.2508	0.2070
output variability	0.0084	0.0101	0.0074	0.0075
inflation variability	0.0022	0.0023	0.0291	0.0290
Performance index: $\beta^S=0.176$	0.3494	0.5040	1.0000	1.0000
$\beta^S=1.0$	0.8203	1.1546	9.8428	9.6551
$\beta^S=5.667$	1.0925	1.4084	54.5380	51.4546
$\beta = 1.0$				
mean π^e_p	0.0000	-0.0003	0.0051	0.0051
policy rule coefficient	0.2642	0.1917	0.2029	0.1603
output variability	0.0084	0.0101	0.0076	0.0077
inflation variability	0.0022	0.0023	0.0059	0.0058
Performance index: $\beta^S=0.176$	0.3494	0.5040	0.3105	0.3192
$\beta^S=1.0$	0.8203	1.1546	1.0000	1.0000
$\beta^S=5.667$	1.0925	1.4084	2.8307	2.6668
$\beta = 5.667$				
mean π^e_p	0.0000	-0.0003	0.0009	0.0009
policy rule coefficient	0.2642	0.1917	0.0975	0.0588
output variability	0.0084	0.0101	0.0086	0.0089
inflation variability	0.0022	0.0023	0.0017	0.0016
Performance index: $\beta^S=0.176$	0.3494	0.5040	0.3627	0.3900
$\beta^S=1.0$	0.8203	1.1546	0.8319	0.8799
$\beta^S=5.667$	1.0925	1.4084	1.0000	1.0000

Table 4a - Dynamic response to contractionary shocks
(decreasing gain sequence)

The table reports a few statistics measuring how the equilibrium outcome under learning differs from the perfect knowledge - i.e. rational expectations - benchmark. Results are presented for both the "plain" RLS algorithm (UE) and the constrained version (CE); agents are assumed to have infinite memory, implying a decreasing gain sequence. The first two columns refer to lexicographic preferences, while the next two to quadratic (dis)utility. To describe the dynamic response of output and inflation and to compare the outcomes under adaptive learning and rational expectations, three measures are computed: (1) the ratio of the volatility of the target variables under adaptive learning and under rational expectations; (2) the trough and (3) the peak of the responses of output and inflation. In the last three lines of each section of the table the value of the performance index is presented. All statistics are computed on the first 50 observations.

	Lexicographic preferences		Quadratic preferences	
	UE	CE	UE	CE
$\beta = 0.176$				
$\sigma_{y}^{AL}/\sigma_{y}^{RE}$	6.6936	6.0532	1.9096	2.1572
min y	-0.0177	-0.0149	-0.0132	-0.0196
max y	0.0025	0.0025	0.0067	0.0067
$\sigma_{\pi}^{AL}/\sigma_{\pi}^{RE}$	2.8320	1.0090	1.0370	0.9868
min π	-0.0028	-0.0012	0.0286	0.0261
max π	0.0013	0.0013	0.0330	0.0330
Performance index: $\beta^S=0.176$	0.1634	0.1409	1.0000	1.0000
$\beta^S=1.0$	0.5891	0.4687	17.9880	16.7520
$\beta^S=5.667$	1.4124	0.8177	190.9600	155.5900
$\beta = 1.0$				
$\sigma_{y}^{AL}/\sigma_{y}^{RE}$	6.6936	6.0532	2.6290	2.6582
min y	-0.0177	-0.0149	-0.0138	-0.0144
max y	0.0025	0.0025	0.0054	0.0054
$\sigma_{\pi}^{AL}/\sigma_{\pi}^{RE}$	2.8320	1.0090	1.0060	0.9812
min π	-0.0028	-0.0012	0.0048	0.0048
max π	0.0013	0.0013	0.0084	0.0084
Performance index: $\beta^S=0.176$	0.1634	0.1409	0.1460	0.1564
$\beta^S=1.0$	0.5891	0.4687	1.0000	1.0000
$\beta^S=5.667$	1.4124	0.8177	7.2675	6.2611
$\beta = 5.667$				
$\sigma_{y}^{AL}/\sigma_{y}^{RE}$	6.6936	6.0532	5.7385	5.5524
min y	-0.0177	-0.0149	-0.0162	-0.0149
max y	0.0025	0.0025	0.0028	0.0028
$\sigma_{\pi}^{AL}/\sigma_{\pi}^{RE}$	2.8320	1.0090	0.6207	0.9806
min π	-0.0028	-0.0012	-0.0003	0.0008
max π	0.0013	0.0013	0.0025	0.0025
Performance index: $\beta^S=0.176$	0.1634	0.1409	0.1361	0.1365
$\beta^S=1.0$	0.5891	0.4687	0.4769	0.4726
$\beta^S=5.667$	1.4124	0.8177	1.0000	1.0000

Table 4b - Dynamic response to contractionary shocks
(constant gain sequence)

The table reports a few statistics measuring how the equilibrium outcome under learning differs from the perfect knowledge - i.e. rational expectations - benchmark. Results are presented for both the "plain" RLS algorithm (UE) and the constrained version (CE); agents are assumed to use a finite number of observations in computing RLS estimates, implying a constant gain sequence. The first two columns refer to lexicographic preferences, while the next two to quadratic (dis)utility. To describe the dynamic response of output and inflation and to compare the outcomes under adaptive learning and rational expectations, three measures are computed: (1) the ratio of the volatility of the target variables under adaptive learning and under rational expectations; (2) the trough and (3) the peak of the responses of output and inflation. In the last three lines of each section of the table the value of the performance index is presented. All statistics are computed on the first 50 observations.

	Lexicographic preferences		Quadratic preferences	
	UE	CE	UE	CE
$\beta = 0.176$				
$\sigma_y^{AL}/\sigma_y^{RE}$	7.0633	5.5241	2.0835	1.6911
min y	-0.0170	-0.0138	-0.0156	-0.0111
max y	0.0025	0.0025	0.0067	0.0067
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	1.3607	0.7174	0.9863	0.9938
min π	-0.0017	-0.0002	0.0283	0.0286
max π	0.0025	0.0013	0.0330	0.0330
Performance index: $\beta^S=0.176$	0.1938	0.1217	1.0000	1.0000
$\beta^S=1.0$	0.6838	0.4498	17.8460	19.4540
$\beta^S=5.667$	1.2303	0.7748	168.8800	182.7800
$\beta = 1.0$				
$\sigma_y^{AL}/\sigma_y^{RE}$	7.0633	5.5241	2.6372	2.3021
min y	-0.0170	-0.0138	-0.0148	-0.0122
max y	0.0025	0.0025	0.0054	0.0054
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	1.3607	0.7174	0.9321	0.9540
min π	-0.0017	-0.0002	0.0046	0.0048
max π	0.0025	0.0013	0.0084	0.0084
Performance index: $\beta^S=0.176$	0.1938	0.1217	0.1528	0.1281
$\beta^S=1.0$	0.6838	0.4498	1.0000	1.0000
$\beta^S=5.667$	1.2303	0.7748	6.2019	6.7111
$\beta = 5.667$				
$\sigma_y^{AL}/\sigma_y^{RE}$	7.0633	5.5241	5.5729	5.3428
min y	-0.0170	-0.0138	-0.0149	-0.0143
max y	0.0025	0.0025	0.0028	0.0028
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	1.3607	0.7174	0.7987	0.8049
min π	-0.0017	-0.0002	0.0005	0.0007
max π	0.0025	0.0013	0.0025	0.0025
Performance index: $\beta^S=0.176$	0.1938	0.1217	0.1383	0.1308
$\beta^S=1.0$	0.6838	0.4498	0.4987	0.4982
$\beta^S=5.667$	1.2303	0.7748	1.0000	1.0000

Chapter 3

Monetary policy with misspecified, heterogeneous and ever-changing expectations

3.1 Introduction and motivation

The vast literature on adaptive learning focuses overwhelmingly on small linear models. Issues like the stability of the equilibrium, the speed of convergence and the dynamics of the learning process are dealt with only for models limited to a handful of equations. And the implications for monetary policymaking are analysed in this very restricted setting, sharply narrowing the range of possible uses. This neglect stems chiefly from the complications of studying stochastic recursive algorithms in large, non-linear systems, but this is unfortunate, because several issues that are relevant only in the context of large-scale models are not paid due attention.

In most of the literature on adaptive learning, it is assumed that the perceived law of motion (PLM) coincides with the minimum state variable (MSV) solution of the corresponding rational expectations equilibrium (REE). This is a convenient simplification that avoids the complexities of dealing with a potential multitude of alternative PLMs and allows straightforward analysis of

the asymptotic properties of the learning algorithm. With non-linear models, however, this is no longer possible, since a closed-form MSV solution does not generally exist; still, if the model is medium-sized or large, no unique and commonly accepted linear approximation will be available either, given the large number of state variables that could be included in the forecasting equation.¹

Absent an MSV solution acting as focal point, agents have to pick out a PLM from a profusion of alternatives, deciding on the basis of some predetermined criterion and taking into account costs of information-gathering and data-processing: different agents end up selecting different forecasting equations and no one sticks to the same PLM indefinitely, preferring to switch based on observed forecasting performances. Evolutionary game theory, which studies the behaviour of large populations who repeatedly engage in strategic interactions, provides the tools for modelling how agents choose among predictors. Because of degrees-of-freedom constraints, each PLM chosen includes only a handful of explanatory variables and accordingly represents just a projection on a small-dimensional space of the actual law of motion, and the ensuing solution is a restricted-perceptions equilibrium. Misspecified expectations also have non-trivial implications for policymaking. Under least squares learning, beliefs become fully rational only if the learning process is E-stable, which depends on the properties of the function $h(\theta) = \lim_{t \rightarrow \infty} EQ(\theta, X_t)$, where $Q(\theta, X_t)$ describes how the estimates of the vector θ of coefficients of the PLM is updated every period. When only a subset of the state vector X_t enters the PLM, the asymptotic limit of $EQ(\theta, X_t)$ depends on the covariances

¹Assume that the model describing the economy contains l free endogenous variables and n predetermined (endogenous and exogenous) variables; for each endogenous jump variable there are $M(n) = \sum_{j=0}^n \binom{n}{j} = 2^n$ alternative linear approximations to the RE solution, an overabundance of options even for small models. If expectations are multi-step ahead, the curse of dimensionality becomes even more uncontrollable, as the right-hand side variables entering the PLM must themselves be forecast: the alternative forecasting models become $M(n) \in O(2^{\psi n})$, with $\psi \gg 1$.

between the variables entering the PLM and those characterising the ALM. The main implication is that even asymptotically the equilibrium solution depends on the specific form of the PLM, as does the learning process. Accordingly, if policymakers can affect the shape of the function $Q(\theta, X_t)$, by guiding the choice of agents' PLMs, they can influence economic outcomes. Central banks, for instance, can improve the ability of financial markets to price long-term assets by providing credible information on how monetary policy rates are set, i.e. by choosing the right degree of transparency.

As is apparent from the foregoing, introducing learning in a high-dimensional non-linear model entails many complexities and makes it difficult to generalise the findings of the recent literature on adaptive learning. There are problems of model underparameterisation, heterogeneous beliefs, ever-changing expectations models, and non-ergodicity in expectations formation. The main implication is that the long-run properties of the learning algorithm change; in particular, under suitable but not too restrictive conditions, the asymptotic equilibrium no longer coincides with the REE, but becomes indeterminate, depending on the specific form of the expectations equations.

Not only analytic issues, but also policy prescriptions depend on the structure of the model. For monetary policymaking, Orphanides and Williams (2007) have shown that when agents learn adaptively, the incentives and constraints facing monetary authorities change substantially: compared with the rational expectations case, imperfect knowledge² (i) reduces the scope for stabilisation of the real economy; (ii) requires more strongly inflation-averse policies and (iii) increases the inertia in interest rate setting.

Besides the degree of activism, departing from the RE paradigm clearly changes the way transparency affects monetary policy effectiveness. Is it still the case that central banks enhance welfare by providing information

²Since imperfect knowledge is a precondition for bounded rationality and learning, that expression is used here and henceforth as a synonym of learning and as an antonym of rational expectations, following Orphanides and Williams (2002).

to households and firms, or should they rather exploit private information to generate inflation surprises? In the case of the standard New Keynesian model, Berardi and Duffy (2006) show that when the central bank operates under commitment, the effects of transparent policies are unambiguously positive, while under discretion there are cases when opaqueness may ensure better outcomes. Eusepi (2005) shows that a sufficient degree of transparency helps make the monetary policy rule robust to expectational errors. These findings are of limited generality, however, since in both papers uncertainty is restricted to the inflation objective and the functional form of the policy rule.³ In a large model, where non-linearities abound, the flow of information from the monetary authority to the private sector is potentially much richer and the role of communication more important. To a considerable extent the monetary authority can decide on the amount of information to provide to the public so as to influence the equilibrium outcomes.

This chapter uses a medium-size model to analyse expectations formation under adaptive learning, heterogeneous beliefs and ever-changing forecasting equations. Empirical rather than analytical results are presented. Two monetary policy issues are used as case studies, the first involving the optimal degree of activism - as in Orphanides and Williams - and the second concerning the benefits associated with transparent policies.

This work makes a number of original contributions.

First, it assumes that in order to anticipate the future path of economic variables, agents can choose among a set of alternative forecasting equations, picking the one with the best track record. Second, it allows agents to have heterogeneous expectations: the share of people selecting a given forecasting model follows a law of motion that is a discrete-time version of the replicator dynamics, implying a gradual movement from worse to better models,

³In Berardi and Duffy (2006), uncertainty about the monetary policy strategy means that agents do not know whether the lagged output gap is part of the reaction function of the central bank (i.e. whether policy is conducted under discretion or under commitment).

unlike another important class of dynamic processes, namely best response dynamics, which involves instantaneous movement to the best-performing strategies. Third, it analyses learning in an economy where expectations have a pervasive role, which is unmatched in the literature: the overwhelming majority of the very few papers studying bounded rationality in large non-linear models introduce learning only in the exchange rate equation.

The chapter is organised as follows. The next section presents a survey of the literature on adaptive learning and presents the replicator dynamics developed in evolutionary game theory to model predictor selection; section 3 outlines the model used in the simulations and introduces stochastic gradient learning. Section 4 presents some evidence, obtained by means of simulation, on the impact on monetary policymaking of departing from the assumption that agents are fully rational. Sensitivity analysis is presented in the following part. Section 6 concludes.

3.2 The literature

There are very few papers on learning in large non-linear models and they deal mainly with the asymptotic convergence of the learning algorithm, disregarding all the monetary policy implications. Garratt and Hall (1997) use the LBS macromodel, adjusted to include adaptive learning schemes to form expectations on the exchange rate, to study whether the choice of the PLM affects the uniqueness and the stability of the equilibrium and whether the volatility of the transition path depends on how agents learn. The issue is whether adaptive learning is E-stable even when the PLM is overparameterised.⁴ Absent analytical results due to the size of the LBS macromodel, they assume that E-stability is achieved when the parameters of the expec-

⁴A learning process that converges to the REE even when the PLM is overparameterised is said to be strongly E-stable.

tations rule cease changing. Garratt and Hall, who use the Kalman filter for the updating of the learning parameters, find that the choice of the PLM modifies the volatility and the speed of convergence of the learning process but obtain less clear-cut evidence on strong E-stability. The end-point for output seems to be the same regardless of the specific form of the PLM, but that for inflation does not. The authors also find that the dynamics and end-value responses of output and inflation are weakly affected by the choice of the expectations rule, but are sensitive to the hyperparameters of the model, i.e. the values of the covariance matrices of the transition and observation equations. The paper is interesting and innovative but has several shortcomings: (i) the forecasting model is the same for all agents; (ii) the learning process relies on hyperparameters that are calibrated rather than estimated; (iii) only exchange rate expectations play a role; (iv) policy issues are entirely neglected; (v) the empirical analysis is based on very short time horizons (less than 10 years).

Beeby, Hall and Henry (2001) go one step further and propose three methods to select a “sensible” PLM when an obvious choice is not available. The first option estimates the effects of a shock to each of the variables on the exchange rate and selects the variables that have a large impact; the second method prescribes computing the rolling correlation (on a 4-quarter window) between the exchange rate and each potential regressor, and ranking the correlations by standard deviations, and choosing the series with the less volatile correlations; the last procedure selects the variables that move most closely with the first few principal components. Beeby, Hall and Henry find also that, regardless of the method used, learning algorithms are quite effective in extracting information from any series, so that the exact form of the rule is unimportant, but they all differ substantially from the RE solution, suggesting that even small deviations from the benchmark of full information and full rationality may have a strong impact on model properties. An obvious

weakness of the paper is that the choice of the best-fitting model is made at the outset once and for all and no heterogeneity in expectations is allowed. Dieppe *et al.* (2011) represents an original attempt to incorporate learning in a large non-linear model, namely the multi-country model of the European Central Bank: it assumes that agents adopt as PLM the reduced form of the equation whose future value they want to anticipate, disregarding all other information. The coefficients of the forecasting equation are updated by means of the Kalman filter, whose hyperparameters are calibrated. The paper has strengths and weaknesses: it sheds light on the impact of learning in a model where beliefs are among the main drivers of the equilibrium outcomes but allows no heterogeneity in expectations formation and relies on convoluted and *ad-hoc* assumptions to specify the PLMs.

The existence of differences in expectations, invariably observed in the real world, is the subject of a vast literature. Evans and Ramey (1992) explicitly introduce the costs of calculation into the process of forming expectations: in any period, agents can revise expectations on the basis of a correct model of the economy if they are willing to pay a price, or can keep their expectations unchanged, incurring no cost. Full convergence to rational expectations happens only when the calculation algorithm becomes infinitely fast and resource costs approach zero. Sethi and Franke (1995) show that persistent heterogeneity can be derived on the basis of evolutionary dynamics in the presence of optimisation costs: the use of sophisticated methods is favoured when optimisation costs are low or when the environment has a high degree of exogenous variability. Deterministic and dynamically stable environments favour the use of simpler and cheaper forecasting methods. Dynamic predictor selection is considered by Brock and Hommes (1997) in a model where agents adapt their beliefs over time by choosing from a finite set of different expectations functions on the basis of costs and of a measure of fit, which is publicly available. Brock and Hommes find that a large response to goodness

of fit can lead to high-order cycles and chaotic dynamics. The rationale is simple: when agents use the cheaper and less accurate predictors, the steady-state equilibrium is unstable, whereas the costly, sophisticated models are stabilising; near the steady-state it pays to use the cheap predictors, which moves the economy away from the steady-state. For a large enough response, this tension leads to local instability and complex global dynamics. Branch and Evans (2006), working on a similar model, find different results. They assume that agents choose on the basis of the predictive performance among a list of costless, misspecified econometric models and obtain conditions under which there is an equilibrium with agents heterogeneously split between the misspecified models even as the intensity of choice becomes arbitrarily large. Branch and McGough (2008) introduce the replicator dynamics into a model with rationally heterogeneous expectations and show that (i) it is possible to generalise the results of Sethi and Franke to a model with an arbitrarily large number of predictors and (ii) complicated dynamics can arise also in setups that are more general than those of Brock and Hommes. Finally, Parke and Waters (2006) study the conditions under which initially heterogeneous beliefs eventually converge to a single forecasting procedure, based on fundamentals and resembling rational expectations.

A common feature of the papers focusing on heterogeneity in expectations formation is that they work with highly simplified models that can be solved analytically and are therefore unsuitable for studying policy issues.

The impact on monetary policymaking of assuming boundedly rational agents is the subject of the paper by Orphanides and Williams (2007). The authors examine the performance and robustness of alternative monetary policy rules by estimating a macroeconomic model in which private agents and the central bank possess imperfect knowledge about the true structure of the economy. They find that policies that appear to be optimal under perfect knowledge can perform very poorly when knowledge is incomplete, partly as

a result of the persistent policy errors due to misperceptions of the natural rates and partly as a result of the learning process that agents use to form expectations. Efficient policies that take account of private learning and of non-observability of natural rates have two features: first, they call for more aggressive responses to inflation; second, they exhibit a high degree of inertia in the setting of the monetary policy rate. Indeed, difference rules (i.e. rules having on the right-hand-side the lagged interest rate with a coefficient equal to 1), which circumvent the need to rely on uncertain estimates of the natural rates, appear to be robust to potential misspecifications of private sector learning and to the magnitude of variation in natural rates.

The value of communication in monetary policy under imperfect knowledge is studied in several papers, including Ferrero and Secchi (2010), Berardi and Duffy (2006) and Eusepi and Preston (2007). Ferrero and Secchi find mixed results on the impact of transparency on the effectiveness of monetary policymaking: when the central bank reveals information about its own expected interest path, conditions for stability under learning become more stringent and the speed of convergence slows down; on the contrary, the announcement of expected inflation and output gap enlarges the set of policy rules which are consistent with stability and a fast process of convergence. Berardi and Duffy link monetary policy transparency to the specification of the forecast rule adopted by the private sector, unlike the traditional view that equates transparency with more or better information. They adopt the standard cashless, three-equation, New Keynesian model and find that under commitment central bank communication is unequivocally welfare-enhancing, while under discretion the relative value of transparency is ambiguous and depends on target values. Eusepi and Preston find that in a dynamic stochastic general equilibrium model with imperfect knowledge, under no communication the policy rule fails to stabilise macroeconomic dynamics, fostering expectations-driven fluctuations. However, by announcing the details of the

policy process, stability is restored: communication permits households and firms to construct more accurate forecasts of future macroeconomic conditions. They further find that if the central bank only announces the desired inflation target, economies with persistent shocks will frequently be prone to self-fulfilling expectations.

While offering significant insights into monetary policymaking, these papers also share two weaknesses, namely the inability to deal with heterogeneous and ever-changing expectations and an overly simplified description of the monetary policy transmission channels.

3.3 The model

The model used here is a reduced-scale version (a so-called *maquette*), reproducing the basic features of the Bank of Italy Quarterly Model.⁵ The sample that has been used to estimate the model covers a 30-year horizon, from the early 1970s to the late 1990s, before Italy joined the European Monetary Union.

The behavioural equations are consistent with maximising agents, but the

⁵A detailed description of the theoretical underpinnings of the Bank of Italy Quarterly Model is in Terlizzese (1994) and Busetti *et al.* (2005). The model is Keynesian in the short run, with the level of economic activity primarily determined by aggregate demand, and neo-classical in the long run. Along a steady-state growth path, the dynamics of the model stem solely from capital accumulation, productivity growth, foreign inflation and demographics; in the short run, a number of additional features matter, namely (i) the stickiness of prices and wages, (ii) the putty-clay nature of capital and (iii) expectational errors. Agents are not fully rational and form expectations by projecting the variables of interest on a subset of predetermined variables; however, unlike adaptive learning, the coefficients of the PLM are not updated whenever new observations are available. In equilibrium - i.e. when no shocks affect the model, expectations are fulfilled and all adjustment processes are completed - the model describes a full-employment economy, in which output, employment and the capital stock are consistent with an aggregate production function, relative prices are constant and inflation equals the exogenous rate of growth of foreign prices. Money is neutral, though not super-neutral. In the taxonomy proposed by Fukač and Pagan (2009), the Bank of Italy Quarterly Model belongs to the 2nd Generation, but shares some of the features of 3rd (i.e. stock-flow consistency and prominence of steady-state properties).

model is not *strictu sensu* microfounded,⁶ since it does not contain all the cross equations restrictions that hold when agents are fully rational, and its structural equations are not tied down exclusively by taste and technology parameters.⁷ Identification is achieved by imposing that the responses of model variables to exogenous shocks are consistent with stylised facts and theoretical presumptions. Like the model in Orphanides and Williams (2007), its main merit is to fit the sample data reasonably well. The learning framework accommodates the Lucas critique, in the sense that expectations formation is endogenous and adjusts to changes in policy or in the structure of the economy,⁸ and accordingly it is legitimate to measure the welfare implications of competing interest-rate rules.

The maquette has some 90 endogenous and 70 exogenous variables. Taking into account expectations formation, the model is described by the following set of vector equations:

$$\begin{aligned}\tilde{y}_t &= f_1 \left(\hat{E}_{t-1}\tilde{y}_t, \hat{E}_{t-1}\tilde{y}_{t+1}, \dots, \hat{E}_{t-1}\tilde{y}_{t+h}, \hat{y}_t, x_t; \Psi_1 \right) + u_{1,t} \\ \hat{y}_t &= f_2 (x_t; \Psi_2) + u_{2,t} \\ w_t &= f_3 (w_{t-1}, w_{t-2}, \dots, w_{t-q}; \Psi_3) + u_{3,t}\end{aligned}\tag{3.1}$$

where \tilde{y}_t and \hat{y}_t are, respectively, the vectors of free and predetermined endogenous variables; w_t indicates the set of exogenous variables, including the intercept; $u_{j,t}$ ($j = 1, 2, 3$) are innovations; x_t is the vector that assembles w_t and all the lags of \tilde{y}_t , \hat{y}_t and w_t ; the matrices Ψ_j ($j = 1, 2, 3$) are collections of parameters; \hat{E} is the (nonrational) expectations operator; h represents the

⁶Many of the econometric models used for forecasting purposes by central banks are not microfounded. In the euro area, most central banks except the Finnish use semi-structural models like the Bank of Italy's; in the United States, the FRB/US, FRB/MCM and FRB/World, which are not truly structural, are still nevertheless the prime large-scale macro models currently in use at the Fed. See Fagan and Morgan (2005) for the euro area and Pescatori and Zaman (2011) for the United States.

⁷Incidentally, one could convincingly object that a model that assumes imperfect knowledge should **not** feature structural equations that are consistent with full rationality.

⁸Orphanides and Williams (2004) use the expression “noisy rational expectations” as a synonym of adaptive learning.

maximum lead with which free variables enter the system; q is the order of the possibly non-linear autoregression for the vector w_t .

Beliefs, which have a direct impact on \tilde{y}_t but affect \hat{y}_t only indirectly (through lags of \tilde{y}_t), enter the model in several ways: ex-ante real interest rates affect the demand for both consumption and capital goods; next-period expected inflation drives current-period wage claims; beliefs about future price developments affect the policy interest rates⁹ and the term structure, which is modelled according to the expectations hypothesis;¹⁰ anticipated changes in the nominal exchange rate bear upon competitiveness and the terms of trade. Moreover, beliefs play a direct role in shaping policy decisions, since natural rates are non-observable and the central bank has to estimate them, before deciding on the proper monetary stance.

The monetary policy transmission mechanism, which is described in detail, works in three phases: first, a change in the policy interest rate spills over to other segments of the capital market, affecting financial asset returns (namely yields on long-term bonds and exchange rates); next, the movements in financial prices interact with the spending behaviour of households and firms; and finally, the change in output and unemployment gaps, driven by the response of consumption and investment, induces wages and prices to adjust to restore the equilibrium. The adjustment process induces modifications in the composition of private and public sector balance sheets, which in turn exert second-round effects on interest rates, thus setting the stage for the response of aggregate demand and supply: the interaction between the real and the financial side of the economy continues until a new equilibrium is reached.

Interest rates affect output through five transmission channels: (i) the cost-of-

⁹The short-term (policy) interest rate depends on the current unemployment gap and on next-period inflation, the latter variable expressed in terms of deviations from target inflation. Some inertia in the policy instruments is allowed by including the lagged interest rate among the arguments of the policy rule.

¹⁰Long-term interest rates are assumed to be a weighted average of current and future short-term rates, with the term spread constant.

capital channel, which works through changes in the optimal capital-output ratio; (ii) the substitution-effect-in-consumption channel, involving the response to financing costs of the relative price of present as opposed to future consumption; (iii) the income and cash-flow channel, based on how capital income flows affect disposable income, whose effects depend on the financial structure of the economy and on borrowers' and lenders' relative propensity to spend; (iv) the wealth channel, that takes into account how fluctuations in borrowing conditions affect the discounted value of future expected payoffs of physical and financial assets; and (v) the exchange rate channel, which measures how fluctuations in exchange rates – triggered by the uncovered interest-rate parity condition – affect competitiveness, the price of imported goods, aggregate demand and inflation.

3.3.1 The learning mechanism

Bounded rationality may be modelled by using recursive least squares (RLS) learning. A convenient alternative to RLS is the stochastic gradient (SG) algorithm, whose main advantage is that it does not rely on information on the second moments of the variables in the forecasting equation. SG learning, which under standard conditions is consistent but not efficient, has been found to work well in complex environments, suggesting that it has robustness properties that RLS lacks. The main drawbacks are: (i) it is not invariant with respect to changes in the units of measurement of the variables in the PLM and (ii) E-stability does not always imply convergence of SG learning.

Recently, Evans et al. (2006) have proposed a generalisation of the SG algorithm, called Generalised Stochastic Gradient (GSG) learning, which solves the invariance problem. They also show that the GSG algorithm has other important justifications: first, it approximates a Bayesian estimator in models where parameters drift; second, it is a maximally robust optimal pre-

diction rule when there is parameter uncertainty; third, though conditions for the stability of generalised stochastic gradient learning differ in general from those governing stability under least squares learning, E-stability in most cases remains a necessary condition for asymptotic convergence of GSG learning.

In all the experiments described in this paper, expectations are modelled by means of a GSG algorithm, as described in (3.2).

$$\begin{aligned}
\widehat{E}_{t-1}^j \widetilde{y}_{i,t} &= \varphi_{j,t-1}^{iT} D_i^j x_t & j \in \{0, \dots, k_i\} \\
\varphi_{j,t}^i &= \varphi_{j,t-1}^i + \gamma_t \Gamma D_i^j x_t \left(y_t - \varphi_{j,t-1}^{iT} D_i^j x_t \right) \\
x_{j,t}^i &= x_{j,t-1}^i g \left(\widetilde{y}_{i,t} - \widehat{E}_{t-1}^j \widetilde{y}_{i,t}, \widetilde{y}_{i,t} - \widehat{E}_{t-1} \widetilde{y}_{i,t}, \right) & g_1 < 0, g_2 > 0 \quad (3.2) \\
\widehat{E}_{t-1} \widetilde{y}_{i,t} &= \sum_{j=1}^{k_i} x_{j,t-1}^i \widehat{E}_{t-1}^j \widetilde{y}_{i,t}
\end{aligned}$$

$\widetilde{y}_{i,t}$ indicates the i^{th} free variable, γ_t is the gain sequence and $\varphi_{j,t}^i$ represents the vector of coefficients of the PLM estimated as of time $t - 1$. The first equation represents the j^{th} PLM for variable $\widetilde{y}_{i,t}$: as agents have imperfect knowledge, they use only a subset of the state vector x_t , i.e. $D_i^j x_t$, to make predictions;¹¹ moreover, being uncertain about the data generating process, they use k_i forecasting equations to predict $\widetilde{y}_{i,t}$, switching from one to another depending on some measure of fit. The share of individuals choosing model $j \in \{1, 2, \dots, k_i\}$ is equal to $x_{j,t}^i$, which is a function of its forecasting accuracy $\left(\widetilde{y}_{i,t} - \widehat{E}_{t-1}^j \widetilde{y}_{i,t} \right)$ relative to the average performance of all forecasting models $\left(\widetilde{y}_{i,t} - \widehat{E}_{t-1} \widetilde{y}_{i,t} \right)$. The last equation in (3.2) states that the expected value of $\widetilde{y}_{i,t}$ is the weighted average of the forecasts of the k_i PLMs.

Though the structural relationships among variables are in general non-linear, the PLMs are assumed to be linear. The gist of the GSG algorithm is captured by the (fixed) matrix Γ , which does not coincide with the inverse of the

¹¹ D_i^j is a selector matrix, i.e. a matrix whose rows have all zeros and a single 1.

second-moment matrix of the regressors (as in recursive least squares) and is not equal to the identity matrix (as in the gradient learning algorithm).

Stability is governed by the differential equation

$$\frac{d\varphi}{d\tau} = \Gamma M_x (T(\varphi) - \varphi) \quad (3.3)$$

where $T(\varphi)$ is the (expected value of the) ALM, $M_x = \lim_{t \rightarrow \infty} E D_i x_t x_t^T D_i^T$ and τ is notional time. When both Γ and M_x are positive definite, the fixed point of (3.3) is the REE $\bar{\varphi}$ and (local) stability is achieved when the eigenvalues of the linearisation of the above matrix differential equation¹² have negative real parts. With RLS learning, the term ΓM_x cancels out and (local) stability depends only on the eigenvalues of the Jacobian of $T(\varphi) - \varphi$, i.e. $DT - I$. In general, the stability conditions for the RLS and GSG algorithms do not coincide and neither implies the other; they become equivalent when the matrix $DT - I$ is H-stable or, alternatively, when Γ is such that $\Gamma M_x = I$.¹³

In all the simulations described here, the scaling matrix Γ is set equal to the inverse of the sample covariance matrix of the regressors estimated on historical data.

3.3.2 Heterogeneous expectations and predictor selection

If agents can pick one out of a large number of forecasting equations, none of which is clearly superior, some problems arise: first, expectations can be heterogenous, since there is no guarantee that everyone will choose the same PLM (or the same sample period); second, agents may elect to change their forecasting equation if they perceive its accuracy as poor; third, several PLMs

¹²The linearisation of the RHS of equation (3.3) is $(\Gamma M_x \otimes I) (DT' - I)$, where DT is the Jacobian of the vectorised mapping $T(\varphi)$.

¹³A matrix C , whose eigenvalues have negative real parts, is said to be H-stable if the eigenvalues of HC have negative real parts whenever the matrix H is positive definite. See Evans *et al.* (2010), in particular Proposition 3 and 4.

can coexist asymptotically, though enough observations are available to tell which performs best.

Evolutionary game theory provides the tool for constructing an explicit model of the process by which agents select the strategy to play in a repeated game.¹⁴ In the typical evolutionary game-theoretic model, there is a large population of agents whose payoff is a function not only of their own strategy but also of other players' behaviour: if an agent can maximise and knows other players' actions, then he can choose the best response; if he does not, he can learn from the observed history of play, which conveys information about how the opponents are likely to play and suggests which strategies are most successful. Agents gradually learn to play an equilibrium if they play the same game (or similar games) repeatedly: once all players have learned how their opponents are playing, and if all are maximising, then they converge to a Nash equilibrium. But how do they reach such an equilibrium? The simplest evolutionary model one could use is the replicator dynamics, which specifies that agents tend to select strategies that do better than the population average and discard those that do worse.

Evolutionary models exhibit learning as a primary ingredient, but are not structural models of learning or bounded rationality: individuals are not explicitly modelled and are treated as naive learners, who do not understand that their behaviour can affect the future play of their opponents and do not take into account that their competitors behave just like them. Agents do not look for patterns in historical data but behave as if the world is stationary, presuming that other players' experience is relevant for them, which justifies imitation.

Here I use the discrete-time version of the replicator dynamics. The economy is populated by a large but finite number of individuals, who play strategy

¹⁴Mailath (1998) and Samuelson (2002) are short but very good surveys of evolutionary game theory; Weibull (1995) is a comprehensive and detailed reference.

$i \in \{1, 2, \dots, K\}$ in a symmetric two-player game with mixed-strategy simplex $\Delta \subset \mathbb{R}^{K-1}$. Let $p_t^i \geq 0$ be the number of individuals who currently select pure strategy i (i.e. who choose model i as a predictor for variable y_t) and let $p_t = \sum_{i=1}^K p_t^i > 0$ be the total population; the share of agents adopting strategy i is accordingly defined as $x_t^i \equiv \frac{p_t^i}{p_t}$ and the vector of predictor proportions (also referred as population state) is $\mathbf{x}_t = \begin{bmatrix} x_t^1 & x_t^2 & \dots & x_t^K \end{bmatrix}^T \in \Delta$, showing that a population state is formally identical with a mixed strategy. The payoff to any pure strategy i at a random match when the population is in state $\mathbf{x}_t \in \Delta$ is $u(e_t^i, \mathbf{x}_t)$, where e_t^i is a vector with 1 in the i^{th} position and 0 elsewhere, representing a pure strategy (i.e. a vertex of the simplex Δ); the associated average payoff is $\sum_{i=1}^K x_t^i u(e_t^i, \mathbf{x}_t)$. The p^i s evolve according to the following laws of motion:

$$p_t^i = (g + u(e_{t-1}^i, \mathbf{x}_{t-1})) p_{t-1}^i, \forall i$$

where g represents the (steady-state) growth rate of the population (the so-called background net birthrate), which implies that

$$p_t = \sum_{i=1}^K x_{t-1}^i (g + u(e_{t-1}^i, \mathbf{x}_{t-1})) p_{t-1} = \left(g + \sum_{i=1}^K x_{t-1}^i u(e_{t-1}^i, \mathbf{x}_{t-1}) \right) p_{t-1}$$

The discrete-time replicator dynamics is accordingly:

$$x_t^i = \frac{g + u(e_{t-1}^i, \mathbf{x}_{t-1})}{g + \sum_{i=1}^K x_{t-1}^i u(e_{t-1}^i, \mathbf{x}_{t-1})} x_{t-1}^i$$

For the vector \mathbf{x}_t to be a proper population state, it must belong to the unit simplex in \mathbb{R}^{K-1} and each of its elements must satisfy the constraint that $0 \leq x_t^i \leq 1$: both conditions are clearly satisfied in the standard case when $u(\cdot, \mathbf{x})$ and g are positive.¹⁵

¹⁵Branch and McGough (2008) use a rule for updating predictor proportions that is state-contingent. They distinguish the strategies $j \in B(x_{t-1})$ that perform worse than average from the strategies $i \in G(x_{t-1})$ that perform better. To impose that $\sum_i x_t^i = 1$, they compute $\sum_{j \in B(x_{t-1})} |\Delta x_t^j|$ and distribute that amount to the strategies $i \in G(x_{t-1})$ in proportion to their payoffs.

In the empirical section an exponential transformation of the mean-square error is used as the payoff function, namely $u(e_t^i, \mathbf{x}_t) = \exp[-\lambda(MSE_t^i + C_i)]$,¹⁶ where C_i is the cost of using model i and $MSE_t^i = (1 - \omega_t)MSE_{t-1}^i + \omega_t(y_t - \hat{E}_{t-1}^i y_t)^2$, with $\omega_t \in \{\frac{1}{t}, \bar{\omega}\}$ and \hat{E}_{t-1}^i being the (conditional) expectations operator based on model i . For simulation purposes, the parameters of the replicator dynamics have been given the following values: $g = .02$; $\lambda = 1000$; $C_i = 0, \forall i$.

3.3.3 The role of expectations

Expectations play a pervasive role in the model: they enter the price- and wage-setting equations, affect monetary policy decisions and drive prices in asset markets. Both the central bank and private agents are assumed to be boundedly rational: the monetary authority learns about inflation and the natural rates of interest and unemployment; households and firms learn about the policy rate, inflation and the exchange rate. It is assumed that the central bank does not consciously attempt to influence the speed of learning by adjusting the degree of activism in policymaking.¹⁷

Unlike the central bank, which uses a single forecasting model for each variable of interest, the private sector employs several predictors jointly. There is no way to avoid unwarranted assumptions in specifying the multiple and mutually coexisting forecasting models that boundedly rational agents use. The problem is how to constrain the information set in an intelligent way, choosing among the innumerable possible ways of doing so. The solution adopted here is to consider only specifications that are sensible economically and that generate predictions that track realisations reasonably well: good-

¹⁶A convex mapping like the exponential function does not reorder the ranking of the payoffs, but alters players' reaction to small and large forecast errors.

¹⁷Ellison and Valla (2001) show that strategic interactions create a connection between the activism of the central bank and the volatility of inflation expectations: the latter reacts to the former because an activist policy produces more information, helping the learning process.

ness of fit is assessed using the criteria proposed in Beeby, Hall and Henry (2001).

The central bank is assumed to set the policy instrument¹⁸ i_t according to the following reaction function:

$$i_t = \rho i_{t-1} + (1 - \rho) \left[r^* + \bar{\pi} + \alpha_\pi \left(\hat{E}_{t-1}^{CB} \pi_{t+1} - \bar{\pi} \right) - \alpha_u (u_t - u^*) \right] \quad (3.4)$$

where \hat{E}^{CB} indicates central bank expectations and r^* and u^* are, respectively, the non-observable natural real interest rate and unemployment rate, which the policymaker seeks to estimate by computing the sample average of the corresponding observables. The central bank's PLMs for π , r^* and u^* are:

$$\begin{cases} \hat{E}_{t-1}^{CB} \pi_t &= \pi_{t-1} + \alpha_{1,t-1} \Delta i_{t-1} + \alpha_{2,t-1} \Delta \pi_{t-1} \\ \hat{E}_{t-1}^{CB} r_t^* &= r_{t-1}^* + \gamma_t (i_{t-1} - \pi_{t-1} - r_{t-1}^*) \\ \hat{E}_{t-1}^{CB} u_t^* &= u_{t-1}^* + \gamma_t (u_{t-1} - u_{t-1}^*) \end{cases} \quad (3.5)$$

where γ_t is the gain sequence.¹⁹ The PLM for inflation is admittedly simple, but it captures the idea that inflation is sticky and depends on the monetary policy stance. The specification was chosen because it minimises the standard error of the regression in a two-variable equation and exhibits a high and stable correlation with survey measures of inflation expectations.²⁰

¹⁸The Bank of Italy Quarterly Model includes several interest rates. To keep the size of the maquette small, all money market rates were reduced to one - the monetary policy instrument - defined as the weighted average of the yields of 3, 6, and 12-month Treasury bills.

¹⁹In the case of decreasing gain $\gamma_t = \frac{1}{t}$, while for perpetual learning $\gamma_t = \bar{\gamma}$.

²⁰Besides the short-term interest rate and lagged inflation, the following variables were considered as eligible regressors: (1) the output gap; (2) the growth rate of GDP; (3) the unemployment rate; (4) the oil price; (5) the nominal effective exchange rate. Absent a unique procedure for selecting the regressors, an evaluation was made on the basis of four criteria: (1) the standard error of the regression; (2) the correlation between $\hat{E}_{t-1}^{CB} \pi_t$ and survey measures of inflation expectations; (3) the rolling correlation (with a 4-year window) with actual inflation; (4) the co-movement with the 1st and 2nd principal components. The last two criteria are suggested by Beeby *et al.* (2001) on the grounds that one picks the variables whose correlation with inflation is high and stable and the other helps select regressors that do not overlap in the amount of predictive information. In principle, the maximisation of the correlation between $\hat{E}_{t-1}^{CB} \pi_t$ and survey-based inflation expectations is what one should be concerned with in choosing the specification of the PLM; in practice, survey data are not a fully satisfactory proxy of households' and firms' anticipations

The specification is in first differences, so that it is consistent with a time-varying inflation objective, reflecting the historical experience of monetary policymaking in Italy in the 1970s and 1980s.

According to the PLM chosen, the central bank's expectations for next-period inflation are equal to:

$$\hat{E}_{t-1}^{CB} \pi_{t+1} = A_{1,t-1} \pi_{t-1} + A_{2,t-1} \pi_{t-2} + A_{3,t-1} \Delta i_t + A_{4,t-1} \Delta i_{t-1} \quad (3.6)$$

where $A_{1,t-1} = \frac{1-\alpha_{2,t-1}^3}{1-\alpha_{2,t-1}}$, $A_{2,t-1} = 1-A_{1,t-1}$, $A_{3,t-1} = \alpha_{1,t-1}$ and $A_{4,t-1} = (1 + \alpha_{2,t-1}) \alpha_{1,t-1}$.

Individuals neither observe the central bank's inflation expectations nor compute model-based estimates. Rather, they pick a forecast out of a limited number of alternatives, sold for a fee C_k by professional forecasters, who use their own model to estimate the future value of economic variables. Absent any empirical evidence for estimating the cost parameters C_k , it is assumed that $C_k = 0$, $\forall k$, so that the only factor affecting the choice of a given forecasting model is accuracy. The relative performance of each model, measured by its mean square error, is common knowledge and in each period agents buy the forecast with the best track record.

Concerning inflation predictions, it is assumed that agents choose among the following 5 predictors:

$$\begin{aligned} \hat{E}_{t-1}^{\pi 1} \pi_t &= \vartheta_{2,t-1}^1 \Delta y_{t-1} + \vartheta_{3,t-1}^1 \pi_{t-1} + \vartheta_{4,t-1}^1 \pi_{t-2} + \vartheta_{5,t-1}^1 \Delta i_{t-1} \\ \hat{E}_{t-1}^{\pi 2} \pi_t &= \vartheta_{0,t-1}^2 + \vartheta_{1,t-1}^2 u_{t-1} \\ \hat{E}_{t-1}^{\pi 3} \pi_t &= \vartheta_{0,t-1}^3 + \vartheta_{2,t-1}^3 \Delta y_{t-1} + \vartheta_{6,t-1}^3 \Delta e_{t-1} \\ \hat{E}_{t-1}^{\pi 4} \pi_t &= \vartheta_{0,t-1}^4 + \vartheta_{7,t-1}^4 \Delta ulc_{t-1} + \vartheta_{8,t-1}^4 \Delta p_{t-1}^M \\ \hat{E}_{t-1}^{\pi 5} \pi_t &= \vartheta_{0,t-1}^5 + \vartheta_{3,t-1}^5 \pi_{t-1} \end{aligned}$$

of future price dynamics. Principal component analysis suggests that two factors explain most of the sample variance and hence two-regressor models are considered. Among the specifications featuring only two regressors, that with lagged inflation and the policy interest rate (i) minimises the standard error of the regression; (ii) exhibits the second-highest correlation with survey-based expected inflation; (iii) has the highest and most stable correlation with actual inflation and (iv) presents regressors moving closely with the first principal component.

where $\widehat{E}_{t-1}^{\pi^k}$ is the expectations operator referring to the k^{th} inflation predictor; y_t is output; u_t the unemployment rate; e_t is the exchange rate; ulc_t unit labour costs; p_t^M the import deflator.²¹ The first equation captures the idea that inflation is sticky and responds to changes in the monetary policy stance and in output growth; the second is a simplified Phillips curve; the third and fourth equations model consumer price dynamics as the sum of domestic costs, proxied either by output growth or by changes in unit labour costs, and foreign inflation, measured by the exchange rate or, alternatively, the import deflator; the last equation models inflation as an AR(1) process. Private sector inflation expectations are equal to $\widehat{E}_{t-1}^{\pi}\pi_t = \sum_{j=1}^5 x_{t-1}^j \widehat{E}_{t-1}^{\pi^j}\pi_t$, i.e.

$$\begin{aligned}\widehat{E}_{t-1}^{\pi}\pi_t &= \Theta_{0,t-1} + \Theta_{1,t-1}u_{t-1} + \Theta_{2,t-1}\Delta y_{t-1} + \Theta_{3,t-1}\pi_{t-1} + \Theta_{4,t-1}\pi_{t-2} \\ &+ \Theta_{5,t-1}\Delta i_{t-1} + \Theta_{6,t-1}\Delta e_{t-1} + \Theta_{7,t-1}\Delta ulc_{t-1} + \Theta_{8,t-1}\Delta p_{t-1}^M\end{aligned}$$

where $\Theta_{0,t-1} = \sum_{j=2}^5 x_{t-1}^j \vartheta_{0,t-1}^j$, $\Theta_{1,t-1} = x_{t-1}^2 \vartheta_{1,t-1}^2$, $\Theta_{2,t-1} = x_{t-1}^1 \vartheta_{2,t-1}^1 + x_{t-1}^3 \vartheta_{2,t-1}^3$, $\Theta_{3,t-1} = x_{t-1}^1 \vartheta_{3,t-1}^1 + x_{t-1}^5 \vartheta_{3,t-1}^5$, $\Theta_{4,t-1} = x_{t-1}^1 \vartheta_{4,t-1}^1$, $\Theta_{5,t-1} = x_{t-1}^1 \vartheta_{5,t-1}^1$, $\Theta_{6,t-1} = x_{t-1}^3 \vartheta_{6,t-1}^3$, $\Theta_{7,t-1} = x_{t-1}^4 \vartheta_{7,t-1}^4$ and $\Theta_{8,t-1} = x_{t-1}^4 \vartheta_{8,t-1}^4$. Individuals rely upon professional forecasters for (short-term) interest rate expectations as well; they are aware which rate is the central bank's instrument, but they do not know either the precise form of the interest-rate rule or how natural rates are estimated. It is assumed that agents can choose

²¹All variables but u_t are log transformations.

among the following 7 predictors:

$$\begin{aligned}
\widehat{E}_{t-1}^{i1} i_t &= \theta_{0,t-1}^1 + \theta_{1,t-1}^1 u_{t-1} + \theta_{3,t-1}^1 \pi_{t-1} \\
\widehat{E}_{t-1}^{i2} i_t &= \theta_{0,t-1}^2 + \theta_{1,t-1}^2 u_{t-1} + \theta_{3,t-1}^2 \pi_{t-1} + \theta_{4,t-1}^2 i_{t-1} \\
\widehat{E}_{t-1}^{i3} i_t &= \theta_{0,t-1}^3 + \theta_{1,t-1}^3 u_{t-1} + \theta_{3,t-1}^3 \widehat{E}_{t-1}^{i3} \pi_{t+1} + \theta_{4,t-1}^3 i_{t-1} \\
\widehat{E}_{t-1}^{i4} i_t &= \theta_{0,t-1}^4 + \theta_{2,t-1}^4 \widehat{E}_{t-1}^{i4} \Delta y_t + \theta_{3,t-1}^4 \widehat{E}_{t-1}^{i4} \pi_{t+1} + \theta_{4,t-1}^4 i_{t-1} \\
\widehat{E}_{t-1}^{i5} i_t &= \theta_{0,t-1}^5 \\
\widehat{E}_{t-1}^{i6} i_t &= \theta_{0,t-1}^6 + \theta_{2,t-1}^6 \widehat{E}_{t-1}^{i6} \Delta y_t + \theta_{3,t-1}^6 \widehat{E}_{t-1}^{i6} \pi_{t+1} + \theta_{4,t-1}^6 i_{t-1} + \theta_{5,t-1}^6 \widehat{E}_{t-1}^{i6} \Delta e_t \\
\widehat{E}_{t-1}^{i7} i_t &= \theta_{0,t-1}^7 + \theta_{4,t-1}^7 i_{t-1}
\end{aligned} \tag{3.7}$$

where \widehat{E}_{t-1}^{ij} is the expectations operator referring to the j^{th} interest-rate predictor. Models 1 to 4 reflect the main finding of the model comparison project conducted by Bryant, Hooper and Mann (1993), namely that effective interest-rate rules react to both inflation and economic slackness, the latter measured in terms of the unemployment rate or, alternatively, the GDP growth rate. The four specifications differ also with regard to policy inertia and the timing of the arguments of the interest-rate rule. Models 5 and 7 capture the naive belief that the central bank seeks to keep the nominal interest rate constant, allowing at most temporary deviations from the target level. Equation 6 includes the exchange rate among the variables affecting the monetary policy stance, which is not uncommon for small open economies.²² In some of the above forecasting models, predictions of future variables appear among the regressors, which in principle would require specifying additional (and possibly multiple) PLMs for each of them. To simplify matters, the following solution has been adopted: expectations of the right-hand-side variables are obtained under the assumption that they evolve according to simple $AR(1)$ processes, namely $\widehat{E}_{t-1} z_t = \psi_{0,t-1}^z + \psi_{1,t-1}^z z_{t-1}$, where z_t is, alternatively, π_t , u_t , Δy_t or Δe_t .²³

²²Bryant, Hooper and Mann (1993) find that interest-rate rules that react to the exchange rate perform worse on average than those that neglect it. Taylor and Williams (2009) make a similar claim.

²³It is implicitly assumed that the professional forecasters predicting the short-term interest rate are

The average expected short-term interest rate at time t is therefore equal to

$$\begin{aligned}
\widehat{E}_{t-1}^i i_t &= \sum_{k=1}^7 x_{t-1}^k \widehat{E}_{t-1}^{ik} i_t \\
&= \Omega_{0,t-1} + \Omega_{1,t-1} u_{t-1} + \Omega_{2,t-1} \Delta y_{t-1} \\
&\quad + \Omega_{3,t-1} \pi_{t-1} + \Omega_{4,t-1} i_{t-1} + \Omega_{5,t-1} e_{t-1}
\end{aligned} \tag{3.8}$$

where

$$\begin{aligned}
\Omega_{0,t-1} &= \sum_{k=1}^7 x_{t-1}^k \theta_{0,t-1}^k + \psi_{0,t-1}^{\Delta y} \sum_{k \in \{4,6\}} x_{t-1}^k \theta_{2,t-1}^k \\
&\quad + (1 + \psi_{1,t-1}^{\pi}) \psi_{0,t-1}^{\pi} \sum_{k \in \{3,4,6\}} x_{t-1}^k \theta_{3,t-1}^k + \psi_{0,t-1}^{\Delta e} x_{t-1}^6 \theta_{5,t-1}^6 \\
\Omega_{1,t-1} &= \sum_{k=1}^3 x_{t-1}^k \theta_{1,t-1}^k \\
\Omega_{2,t-1} &= \psi_{1,t-1}^{\Delta y} \sum_{k \in \{4,6\}} x_{t-1}^k \theta_{2,t-1}^k \\
\Omega_{3,t-1} &= (\psi_{1,t-1}^{\pi})^2 \sum_{k \in \{3,4,6\}} x_{t-1}^k \theta_{3,t-1}^k \\
\Omega_{4,t-1} &= \sum_{k \in \{2,3,4,6\}} x_{t-1}^k \theta_{4,t-1}^k \\
\Omega_{5,t-1} &= \psi_{1,t-1}^{\Delta e} x_{t-1}^6 \theta_{5,t-1}^6
\end{aligned}$$

Expectations of short-term interest rates form part of the equation of the yield curve. According to the expectations hypothesis, k -year bond yields are equal to the k -year moving average of current and the future short-term interest rates plus a constant term premium that agents estimate using the historical mean. To prevent forecast errors from accumulating when computing multi-step ahead interest-rate expectations, the term premium is corrected for the mean difference between expected and actual past policy rates:

$$\widehat{E}_{t-1} term_t = term_{t-1} + \gamma_t \left[(i_{t-1}^L - i_{t-1}) + \frac{1}{6} \sum_{j=1}^6 \xi_{t-j} - term_{t-1} \right] \tag{3.9}$$

not the same as those forecasting inflation.

where $\xi_{t-j} \equiv i_{t-1} - \widehat{E}_{t-1-j} i_{t-1}$ measures the surprise on the policy interest rate.

By taking the j^{th} lead of equation (3.8), for $1 \leq j \leq 5$, and replacing all non-predetermined variables, one obtains the following expression for $\widehat{E}_{t-1} i_{t+j}$:

$$\begin{aligned} \widehat{E}_{t-1}^i i_{t+j} &= \Omega_{0,t-1}^j + \Omega_{1,t-1}^j u_{t-1} + \Omega_{2,t-1}^j \Delta y_{t-1} \\ &+ \Omega_{3,t-1}^j \pi_{t-1} + \Omega_{4,t-1}^j i_{t-1} + \Omega_{5,t-1}^j \Delta e_{t-1} \end{aligned} \quad (3.10)$$

where

$$\begin{aligned} \Omega_{0,t-1}^j &= \Omega_{0,t-1} + \Omega_{1,t-1} \frac{1 - (\psi_{0,t-1}^u)^{j+1}}{1 - \psi_{0,t-1}^u} + \Omega_{2,t-1} \frac{1 - (\psi_{0,t-1}^{\Delta y})^{j+1}}{1 - \psi_{0,t-1}^{\Delta y}} \\ &+ \Omega_{3,t-1} \frac{1 - (\psi_{0,t-1}^\pi)^{j+1}}{1 - \psi_{0,t-1}^\pi} + \Omega_{4,t-1} \Omega_{0,t-1}^{j-1} + \Omega_{5,t-1} \frac{1 - (\psi_{0,t-1}^{\Delta e})^{j+1}}{1 - \psi_{0,t-1}^{\Delta e}} \\ \Omega_{1,t-1}^j &= \Omega_{1,t-1} (\psi_{1,t-1}^u)^{j+1} + \Omega_{4,t-1} \Omega_{1,t-1}^{j-1} \\ \Omega_{2,t-1}^j &= \Omega_{2,t-1} (\psi_{1,t-1}^{\Delta y})^{j+1} + \Omega_{4,t-1} \Omega_{2,t-1}^{j-1} \\ \Omega_{3,t-1}^j &= \Omega_{3,t-1} (\psi_{1,t-1}^\pi)^{j+1} + \Omega_{4,t-1} \Omega_{3,t-1}^{j-1} \\ \Omega_{4,t-1}^j &= \Omega_{4,t-1} \Omega_{4,t-1}^{j-1} \\ \Omega_{5,t-1}^j &= \Omega_{5,t-1} (\psi_{1,t-1}^{\Delta e})^{j+1} + \Omega_{4,t-1} \Omega_{5,t-1}^{j-1} \end{aligned}$$

with $\Omega_{k,t-1}^0 = \Omega_{k,t-1}$, $k = 1, 2, \dots, 5$. It is clear from the above expression that the long-term interest rate depends on expectations of several variables, so that policies that focus on a single objective at the cost of others are unlikely to be welfare-enhancing.

Private sector expectations also affect the value of the domestic currency. As there is no well-established specification for the exchange rate equation, only two competing models are considered: the first relates exchange rate dynamics to the ratio of net foreign assets to nominal GDP; the second captures the belief that the value of the domestic currency is a random walk.

$$\begin{aligned} \widehat{E}_{t-1}^{e1} e_t &= e_{t-1} + \beta_{1,t-1} \Delta \frac{FA_{t-1}}{Y_{t-1}} + \beta_{2,t-1} \Delta e_{t-1} \\ \widehat{E}_{t-1}^{e2} e_t &= e_{t-1} \end{aligned} \quad (3.11)$$

Among the set of two-regressor specifications, the model selected (i) minimises the standard error of the regression; (ii) exhibits the second highest

correlation with survey measures of exchange rate changes; (iii) has explanatory variables that move closely in line with the first two principal components; (iv) presents the second largest and most stable correlation with the change in the exchange rate. Along with the UIP, equation (3.11) determines e_t as a function of its own lags, the interest rate differential and foreign indebtedness.

3.4 Simulation results

Monetary policy rules are ranked on the basis of their impact on social welfare. Society dislikes both price and output variability, defined as the unconditional variances of inflation and GDP growth. The target value of both variables is the steady-state equilibrium value and the two objectives have the same weight in the welfare function, which is equal to:

$$W = - \left[E (\pi_t - \bar{\pi})^2 + E (\Delta y_t - \overline{\Delta y})^2 \right] \quad (3.12)$$

Using unconditional variances rather than discounted future losses implicitly favours policies that minimise the overall impact of shocks, penalising those that trade smaller fluctuations today for larger ones tomorrow. Unlike Orphanides and Williams (2007), here the welfare function factors in output rather than unemployment but the change is inconsequential, since in all the experiments the ranking of the policy rules is the same regardless of the argument variable. Interest rate volatility is not included, but it affects social welfare indirectly, since the term structure exerts a powerful influence on GDP.

Model simulations are used to illustrate how the interaction between the expectations formation mechanism and the monetary policy rule affects the equilibrium outcomes. The optimal policy is selected via a grid search on the

parameters $\{\rho, \alpha_\pi, \alpha_u\}$ of the Taylor-type reaction function: in order to save on computation time, the step-length of the grid search is initially quite large (.1 for ρ ; .5 for α_π and α_u), but gradually diminishes once the region containing the welfare-maximising triplet is located. Each experiment consists of 500 replications and all simulations cover an interval of 490 years (from year 2011 to year 2500). In the first 90 periods, the main stochastic equations²⁴ are shocked to test how effectively the monetary policy rule stabilises the economy;²⁵ in the subsequent 400 years, all shocks are reset to zero and the model settles down on the steady-state equilibrium growth path, which makes it possible to assess the convergence properties of the learning algorithm. The GSG algorithm is initialised using OLS estimates on historical data.

3.4.1 Optimal monetary policy under rational expectations

Under rational expectations the optimal monetary policy has a small degree of interest rate smoothing, a strong response to inflation and a non-negligible concern for changes in the unemployment rate (see Fig.1): the welfare-maximising coefficients are $\rho = 0.4$, $\alpha_\pi = 2$ and $\alpha_u = 1.5$.

Comparing the performance of alternative rules provides some notable insights. First, the degree of inertia does not matter greatly: the welfare function is quite flat for positive values of ρ up to 0.7. For higher values, both output and inflation variability increase, suggesting that too smooth a path of the policy interest rate fails to stabilise the economy; for $\rho \geq 0.9$ the system no longer converges, showing that difference rules are not a viable alternative. Second, the equilibrium outcomes are not overly sensitive

²⁴The equations are: (i) household consumption, (ii) exports, (iii) the private sector value added deflator and (iv) the consumption deflator. The white-noise shocks may be interpreted as referring to domestic and foreign household preferences and domestic and foreign mark-ups.

²⁵To ensure a fair comparison across policy rules, the same sequences of random draws are used for each triplet $\{\rho, \alpha_\pi, \alpha_u\}$.

to the value of α_u , possibly because of the role of fiscal policy in stabilising the economy. Close to the local optimum, the welfare function exhibits a hump-shaped response to α_u ; away from it, no well-defined relationship is apparent. Third, the policymaker's response to deviations from the inflation target ought to be quite strong: the optimum is achieved when $\alpha_\pi = 2$, while for $\alpha_\pi \leq 1$ the model is not stable, suggesting that the Taylor principle holds. This finding is not trivial, since unlike small closed-economy models with no government, the maquette of the Bank of Italy Quarterly Model provides for channels other than monetary policy that help to tame inflationary pressures.²⁶ Social welfare turns out to be very sensitive to changes in α_π , contrary to what happens with α_u or ρ : other things equal, it falls by nearly one sixth when $\alpha_\pi = 3$ and by one third when $\alpha_\pi = 4$. Fourth, mild changes in the weighting of the objectives of the loss function are inconsequential: the optimal policy stays the same when the weight of output stabilisation is halved and remains close to optimal when it is doubled.

3.4.2 Optimal monetary policy under learning

The foregoing results are based on four partly interrelated hypotheses: (i) the economic environment is *stationary*, since equations do not change over time; (ii) agents know the structure of the economy; (iii) expectations are *rational* and (iv) the central bank is *credibly committed* to an unchanging policy rule. Each assumption has a strong impact on the properties of the system and on the policymaker's incentives and constraints. Uncertainty about the structure of the economy forces policymakers to rely on estimates of the unobserved natural rates; imperfect knowledge on how the economy works alters the way monetary policy and private sector expectations inter-

²⁶An increase in inflation worsens price competitiveness and reduces the real value of non-indexed financial wealth; the resulting decline in exports and private-sector spending translates into less employment and decelerating costs. Besides, there is usually fiscal drag weakening aggregate demand. Since these channels are at work in the model, the Taylor principle may not be a necessary condition for determinacy.

act; learning makes an otherwise stationary environment non-stationary; and the imperfect credibility of the central bank reduces the authority’s ability to steer market expectations.

In order to assess the impact of these assumptions on the central bank’s strategy, I run three sets of simulations. In the benchmark case, labelled “no transparency”, I assume that agents do not know the current value of the policy interest rate when they take their decisions but observe it with a one-period delay (i.e. the private sector forms expectations before the monetary policy rate for the current period is set). In the second experiment, I assume that the central bank pre-announces the current-period monetary policy stance, so that $\hat{E}_{t-1}^P i_t = i_t$; this case is dubbed “partial transparency”, because the authority communicates neither its own estimates of the natural rates ($\hat{E}_{t-1}^{CB} r_t^*$ and $\hat{E}_{t-1}^{CB} u_t^*$) nor the coefficients of the reaction function (ρ , α_π and α_u). The final set of simulations posits a “fully transparent” central bank that provides private agents with all the information it processes in making policy decisions.²⁷ Comparing the first experiment with the rational expectations equilibrium, one can assess the welfare losses and the changes of the optimal monetary policy rule due to imperfect knowledge; comparing the other two experiments with the benchmark “no transparency” hypothesis, one can gauge the gains from an effective communication strategy.

The size and complexity of the model make it hard to disentangle the channels through which monetary policy decisions affect the economy and to assess how the parameters of the interest-rate rule bear on output and inflation volatility. In order to determine which factors affect welfare most strongly, social welfare and its drivers have been regressed on the standard deviations of the main macroeconomic variables. Table 1 reports the t -statistics of these regressions: a negative correlation between welfare and the standard

²⁷The expression “full transparency” is not perfectly appropriate here, as the central bank does not communicate everything to the public, e.g. its PLM for inflation and the variables entering the interest-rate rule.

deviation of a variable implies that the lower the volatility of that variable, the greater the increase in social welfare. The entries in the table suggest that the anchoring of inflation expectations and the stabilisation of wages and prices are the primary sources of welfare movements: when transparency is not complete, lower variability of surprise and expected inflation results in higher welfare; when transparency is full, wage fluctuations are the main factor in economic instability. The econometric evidence suggests that by controlling nominal variables the monetary policymaker succeeds in keeping the real variables in check as well. The exchange rate and the term spread do not appear to be significant drivers of welfare.

The benchmark case

Table 2a presents summary statistics describing how alternative policy rules work. Monetary policies are appraised according to two indices: the level of welfare and the rejection rate (i.e. the percentage of non-converging replications). Results on the optimal rule are presented in the first row; the other policies are arranged so that only one parameter at a time changes, making it easier to see how sensitive the rule's performance is to changes in each element of the triplet $(\rho, \alpha_\pi, \alpha_u)$. For every combination of parameters, the table shows the first and second moment from steady-state of (i) output, (ii) inflation, (iii) inflation surprises, (iv) private-sector nominal wage growth, (v) the (mean) intercept of the inflation PLMs, (vi) the difference between central bank and private-sector inflation expectations, (vii) the optimal capital-output ratio, (viii) the short and (ix) the long-term interest rate. In the last two columns, the table shows the entropy of inflation and of the policy rate, measuring the uncertainty agents face in choosing the forecasting model.

The entropy H of a discrete random variable X , whose values are $\{x_1, x_2, \dots, x_K\}$,

is defined as $H(X) = - \sum_{k=1}^K p(x_k) \log_b p(x_k)$, where $p(x_k)$ is the probability of drawing x_k and b is the base of the logarithm. $H(X)$ reaches its maximum when, for each x_k , $p(x_k) = 1/K$ and its minimum when $p(x_k) = 1$ for one x_k and zero otherwise. Common values for b are 2, e or 10; alternatively, one can choose $b = K$ so that $H(X) \in [0, 1]$. In this paper, x_k represents the k^{th} forecasting model and $p(x_k)$ is the share of agents buying its predictions; $H(X) = 1$ means that the data do not help to discriminate among models, while $H(X) = 0$ indicates that one predictor dominates and precludes the others.

Comparing the equilibrium outcomes under rational expectations and learning, it seems that neither the uncertainty about the natural rate nor the expectations formation mechanism entail substantial welfare losses: the optimal policy under learning achieves nearly the same welfare level as the optimal policy under rational expectations, and even suboptimal ones perform quite well in most cases. What does change is the shape of the optimal policy rule: under adaptive learning, the optimal policy requires a weaker response to deviations of inflation from target ($\alpha_\pi = 0.4$ rather than $\alpha_\pi = 2.0$) and a stronger concern for output stabilisation ($\alpha_u = 3.5$ rather than $\alpha_u = 1.5$). This outcome depends mostly on the exchange rate: under adaptive learning, exchange rate expectations are stickier and the value of the currency - and hence inflation - is less volatile, which induces the monetary policymaker to pay more attention to output stabilisation. The degree of inertia is roughly the same ($\rho = 0.3$ rather than $\rho = 0.4$), but it does not seem to play a substantial role: social welfare is to a large extent unaffected by the coefficient of the lagged interest rate and does not change significantly for values of ρ in the range $[0.3, 0.5]$; as ρ increases, output fluctuates less and inflation more. Concerning the performance of alternative monetary policy rules under learning, table 2a offers several insights. In particular, it shows that inflation

volatility is much more responsive than output volatility to changes in ρ , α_π and α_u : the range of variation of the standard deviation of inflation is about seven times that of output. Not surprisingly, the best-performing rules are those that anchor prices better, even if this comes at the cost of wider fluctuations in the level of economic activity. If the inertia of the policy rule decreases from 0.8 to 0.3, inflation volatility decreases by one-fourth, while that of output increases by less than one-fiftieth; if α_π falls from 2.5 to 0.5, the second moment of both of the central bank's objectives diminishes; if α_u rises from 1.0 to 2.5, inflation counter-intuitively becomes much less erratic and output fluctuates more. More effective rules have a low or even zero rejection rate; moreover, the share of non-converging replications seems to be proportional to the degree of inertia and inversely related to responsiveness to the unemployment gap.

The best-performing policies are therefore those that (1) have low inertia; (2) do not overreact to changes in inflation; and (3) lean strongly against aggregate demand shocks. These rules succeed in keeping price fluctuations under control mostly through expectations management. Applying the law of total variance, the second moment of inflation (i.e. $Var(\pi_t)$) can be decomposed into the sum of the expected value of its conditional variance ($E(Var(\pi_t|I_{t-1}))$) and the variance of its conditional expectation ($Var(E(\pi_t|I_{t-1}))$): the first term can be proxied by the variance of time- t inflation surprises averaged across time, the second by the variance of expected inflation, where expected inflation in each period is computed averaging across replications. The recipe for effectiveness is therefore to make inflation predictable (which reduces the first term) and to prevent expectations from decoupling from targets (which minimises the second).²⁸ If these two requirements are not met, wages - which depend on inflation expecta-

²⁸The alignment between target and expected inflation is also measured by the first and second moments of the intercept of the forecasting models used for predicting inflation. Both statistics are shown in table 1.

tions - become excessively erratic and nominal instability is transmitted to households' and firms' spending decisions. Table 2a confirms that a high level of welfare is in general associated with predictable inflation and stable expectations: predictability is inversely related to the volatility of inflation surprises.

Expectations are one of the key elements in understanding how the economy responds to monetary policy actions. If beliefs are not homogeneous, a natural question is whether model heterogeneity disappears as data accumulates: the answer is a resounding no, regardless of the central bank's communication strategy.

With regard to short-term interest-rate expectations, the two PLMs that include output growth and next-period inflation (the 4th and 6th models) outperform the others, despite the fact that the central bank policy rule uses the unemployment rate as a proxy for slackness in economic activity. The population state gradually converges towards a situation where more than 80% of the agents use one of these two models. Similar results are found for expected inflation. The best-fitting PLM has unit labour costs and the import deflator as regressors, while the second-best has the unemployment rate as the sole explanatory variable; taken together, they account for nearly 85% of agents' picks. The ranking of the forecasting models is more or less the same across replications and end-of-sample proportions cluster together quite neatly. For the exchange rate the picture is different: neither of the two forecasting models clearly stands out and the relative accuracy of the competing PLMs seems to be driven by a combination of shocks to the economy. Unlike the other variables, exchange rate expectations are highly dispersed across replications.

One finding is common to all three cases and to all transparency regimes: highly inaccurate forecasting models tend to be discarded, but no PLM succeeds in ruling out all the others, possibly because there is not enough in-

formation in the data. Heterogeneity in expectations formation seems to be an intrinsic feature of the model: even for the policy interest rate - which depends on a small set of variables and is more accurately tracked by agents' expectations - two PLMs coexist. Does this finding depend on the value of the sensitivity parameter λ of the payoff function? Not really. Fig. 2 shows that even for very high values of the responsiveness of the payoff function to forecast errors, the heterogeneity in expectations formation does not disappear. On the contrary, excessively high values of λ tend to reduce the share of agents choosing the two best forecasting models, especially for inflation: heterogeneity in expectations stops decreasing, and the selection of forecasting models becomes more and more erratic. Indeed, when λ exceeds a certain threshold (i.e. when $\lambda \geq 2500$) social welfare deteriorates: the volatilities of output growth and inflation increase, as agents tend to switch too frequently from one forecasting model to another, making predictions inaccurate and disanchoring expectations.

These results clash with those of Orphanides and Williams (2007), who find that when private agents have imperfect knowledge, the central bank benefits from more strongly inflation-averse policies, which help prevent expectations from decoupling from target inflation. This contrasting evidence is explained by differing monetary policy transmission mechanisms. The Orphanides-Williams model is a plain-vanilla three-equation New-Keynesian model: the policy instrument affects aggregate demand directly and inflation indirectly (through the output gap); as long as interest rate changes offset inflationary pressures, they stabilise the economy and have negligible spillovers on social welfare.²⁹ The model used in this chapter has a much richer transmission mechanism, where expectations not only drive monetary policy choices but also affect wage setting, competitiveness and asset prices: a strong interest-rate response to price shocks makes actual inflation more erratic and less

²⁹Provided of course that interest rate volatility does not have a large weight in the loss function.

predictable. What happens is that by overreacting to inflationary pressures, the policymaker induces greater fluctuations in consumption and investment, putting additional pressures on prices. The net effect is to amplify rather than attenuate the initial shock.

An additional channel affecting the transmission of monetary impulses works through asset prices. A tightening of the policy stance results in an appreciation of the currency, which keeps price dynamics in check, both directly (through a lower import deflator) and indirectly (through the impact of a deterioration in price competitiveness on economic activity). But the simulation results suggest that this channel plays only a minor role in shaping the response of the economy to monetary impulses. Policy stimuli bear upon the yields of long-term bonds and the slope of the term structure of interest rates also. There is no easily discernible relationship among the coefficients of the policy rule, the volatility of the term structure and social welfare. Considering the volatility of inflation-adjusted yields, it clearly has a positive effects on the standard deviation of output growth, but the impact on welfare is distorted by the response of inflation, whose fluctuations seem to be dampened by more volatile real interest rates.³⁰

Two other findings are worth mentioning. First, the Taylor principle does not apply: though the welfare-maximising value of α_π is 0.4, the model is stable and learnable, and the rejection rate is zero.³¹ Second, the optimal rule has the lowest entropy associated with the predictor proportions of PLMs for the short-term interest rate, suggesting that one ingredient in a successful policy

³⁰The link between learning and interest rates is not a novel feature of this paper. Dewachter and Lyrio (2006) present a macroeconomic model in which agents learn about the central bank's inflation target and the real interest rate to explain the joint dynamics of output, inflation and the term structure of interest rates. Learning generates endogenous stochastic endpoints that act as level factors for the yield curve. They find that their model has a better fit than those based on rational expectations and generates sufficiently volatile endpoints to match the variation in long-maturity yields and in surveys of inflation expectations.

³¹Svensson (2000) explains why the Taylor principle does not hold in open economies.

is enabling agents to discriminate between good and bad forecasting models.

The case of partial transparency

Transparency of monetary policy refers to the absence of information asymmetries between policymakers and the private sector. Perfect transparency, in the setup used here, implies that the central bank discloses to the general public both its estimates of the natural rates and the precise form of the policy rule; incomplete transparency is defined as a situation where the policymaker communicates in advance only the monetary stance (i.e. the value of i_t). In this case, expectations about future policy rates, which are needed to price long-term securities, are formed with a PLM that differs from the true interest-rate rule, namely:

$$\begin{aligned}\widehat{E}_{t-1}^i i_{t+j} = & \Omega_{0,t-1}^j + \Omega_{1,t-1}^j u_{t-1} + \Omega_{2,t-1}^j \Delta y_{t-1} \\ & + \Omega_{3,t-1}^j \pi_{t-1} + \Omega_{4,t-1}^{j-1} i_t + \Omega_{5,t-1}^j \Delta e_{t-1}\end{aligned}$$

Table 2b shows the results of the simulations under partial transparency. There seems to be only a modest gain from being transparent: the optimal policy achieves a level of welfare that is just slightly better than the best outcome under opaqueness. Some benefits are discernible in lower rejection rates and in the way the central bank manages to steer agents' behaviour: when agents know in advance what the central bank is going to do, they behave in a way that is consistent with the monetary stance, fostering the achievement of the objectives with smaller changes in the policy instrument. Less volatile short-term interest rates promote a somewhat flatter term structure and are conducive to a more precise appraisal of the unobserved natural rates, though the evidence is not unambiguous.

The case of full transparency

The equilibrium outcomes change substantially when the central bank is fully transparent and discloses all the information that it uses in choosing its monetary stance. Full transparency holds when no information asymmetry between the central bank and the general public exists. Since the central bank informs market participants of the coefficients of the policy rule, the inflation objective, and its own estimates of the natural rates, expectations about future policy rates are set according to the following equation:

$$\begin{cases} \widehat{E}_{t-1}^i i_{t+j} &= \rho \widehat{E}_{t-1}^i i_{t-1+j} + (1 - \rho) i_{t+j}^* \\ i_{t+j}^* &= \widehat{E}_{t-1}^{CB} r_t^* + \bar{\pi} + \alpha_\pi \left(\widehat{E}_{t-1}^i \pi_{t+1+j} - \bar{\pi} \right) - \alpha_u \left(\widehat{E}_{t-1}^i u_{t+j} - \widehat{E}_{t-1}^{CB} u_t^* \right) \end{cases}$$

for $j > 0$. The only remaining information asymmetry is the one about the PLMs for inflation and the unemployment rate, which are not the same for the central bank and the private sector. As shown in Table 2c, the best performing rule features a much higher degree of inertia, a stronger inflation aversion and a lower concern for output fluctuations. The sensitivity to changes in the value of ρ is high: for $\rho = .3$, welfare is nearly 20 p.p. lower than at the optimum. A low degree of inertia tends to destabilise the exchange rate and raises substantially the cost of financing, which justifies the deterioration of the policy performance.³² Welfare is also sensitive to the value of α_u , since too weak a response to unemployment gaps injects variability in inflation. It is worth stressing that though the optimal strategy exhibits a larger α_π and a smaller α_u than in the partial and no-transparency cases, output fluctuates less and inflation more; moreover, for most combinations of $\{\rho, \alpha_\pi, \alpha_u\}$, the policy interest rate tends to be less volatile, but the long-term yield exhibits much larger fluctuations. A possible explanation is that when the natural rates and the policy parameters are not estimated, but

³²As shown in Table 2c, the bias of the long-term interest rate - i.e. the difference between the mean value across time and replications and the steady-state value - is always larger than 200 basis points.

provided by the central bank, there is no automatic error-correction mechanism working through recursive learning, so that expected future policy rates become extremely erratic and the term structure biased and volatile.³³

One noteworthy feature is that in all cases, even when the overall performance of the policy rule is poor, the standard deviation of output is smaller than under partial transparency or opaqueness.³⁴ Higher inflation volatility is traded for lower output volatility, as witnessed also by the standard deviation of the capital-output ratio, which is substantially smaller than in the other two cases. Notwithstanding the relatively large variability of inflation, the mean intercept of the inflation PLMs turns out to be much less biased and unstable.

While the optimal rule does not improve significantly upon the partial and no-transparency case, suboptimal strategies seem on average to perform better, suggesting that transparency may be conducive to robustness. All in all, it seems that central bank talk has a beneficial but very modest impact on agents' expectations and behaviour. The explanation of this finding echoes the warning of Amato, Morris and Shin (2002), who note that central bank communication has a dual function: on the one hand, it provides signals about the policymaker's private information; on the other hand, it serves as a coordination device for the beliefs of private agents and may at times induce agents to do away with their own private information. The first effect is welfare-enhancing; the second may be welfare-reducing. Which effect prevails cannot be said in general: in the case considered, it seems that the

³³This guess is confirmed by the value - not reported in the table - of the 1st and 2nd moments of the variable measuring interest-rate missperceptions, i.e. $\frac{1}{6} \sum_{j=1}^6 (i_{t-1} - \hat{E}_{t-1-j} i_{t-1})$, which are much larger than in the previous cases.

³⁴It is not certain however that this outcome is to be attributed to monetary policy. An alternative possibility is that this result is due to fiscal policy: at the optimum point, the standard deviation of the tax rate on disposable income (which is the fiscal policy instrument used to keep the debt-to-GDP ratio close to its target of 0.6) is more volatile and much higher than in the steady-state equilibrium; in the partial and no-transparency cases, the opposite happens.

benefits of adopting a completely transparent policy are largely offset by its shortcomings.

3.4.3 Perpetual learning

The canonical justification for adopting gain sequences that remain bounded above zero is that the economy is subject to structural shifts and, accordingly, past observations should be given less weight than recent data in the learning algorithm. There is actually a second rationale for using constant-gain estimators that fits the model in this paper perfectly: the possibility of nonconvergence to the REE. If convergence to the perfect information equilibrium is for whatsoever reason unlikely, then the actual stochastic process followed by the economy may best be modelled - given the PLMs employed by agents - as undergoing structural change over time. The main implication of constant-gain learning is that agents' estimates are always subject to sampling variation and never converge to fixed values; for this reason, some authors name this adaptive scheme "perpetual learning".

Table 3a to 3c report the simulation results under the three alternative communication strategies in the case of perpetual learning. Under no transparency, there is hardly any difference between the decreasing and constant gain cases. The best policy is essentially the same, just a bit more inertial ($\rho = .31$ rather than $\rho = .3$) and slightly less reactive to fluctuations in real activity ($\alpha_u = 3.2$ rather than $\alpha_u = 3.5$). Welfare is apparently not affected by the memory of the learning algorithm: it is either the same or slightly lower, suggesting that observations far away in the past are indeed barely informative. The ranking of suboptimal policies is not altered either: the worst outcomes are achieved when either α_π is too high or α_u is too low.

Similar results are obtained when the central bank discloses the information it uses in making policy decisions. Under partial transparency, the welfare-maximising policy features a slightly milder response to the unemployment

gap ($\alpha_u = 3$ vs. $\alpha_u = 3.2$). Under full transparency the opposite happens: the optimum is achieved with a somewhat higher value of α_u and a somewhat lower value of α_π . In general, the simulations confirm that when the monetary policymaker reacts too aggressively to price shocks or too meekly to demand fluctuations, the economy becomes unstable and social welfare plunges.

3.5 Sensitivity analysis

The results just described are based on several ad-hoc assumptions. On some of them - the number of replications in each experiment or the initial conditions of the learning process - a thorough sensitivity analysis can be conducted; on others - the choice of the PLMs - no fully-satisfactory testing procedure is available: with hundreds of variables, there are too many PLMs that can be chosen, most of them indistinguishable in terms of parsimony or fitting.

To test the generality of the findings described in the previous section, four sensitivity analysis exercises are conducted: in the first, the model is simulated with 10,000 replications and the results compared with those obtained in the baseline experiment, to test whether the latter are distorted by the small number of replications; in the second, the initial conditions of the learning algorithm are changed, by increasing/decreasing the (fixed) covariance matrix of the regressors, that drives the size of the Kalman gain and accordingly the extent of the revisions in expectations once new data becomes available; in the third, the sensitivity of the optimal monetary policy rule to changes in the welfare function is assessed; in the final experiment, initial conditions for predictor proportions x_t^i are set randomly, by drawing from a uniform distribution with support in $[0, 1]$, instead of imposing that they are equal to the reciprocal of the number of forecasting models and constant across replications.

3.5.1 Experiment #1: the number of replications

For each experiment, the number of replications has been chosen so as to guarantee reliable results while keeping the time needed for a full search of the optimal policy at an acceptable level. The model, augmented with the learning recursions, contains nearly 300 equations: when all 500 replications converge, it takes roughly two minutes to complete them; when some of them diverge, it can require two hours of computer time. Since the search for the optimum policy calls for the evaluation of more than 300 combinations of the Taylor-rule coefficients, 500 replications has been viewed as an acceptable compromise.

To assess whether the results shown in Tables 2a to 2c are affected by small sample bias, the equilibrium outcomes of the three communication regimes at the optimum have been compared with the results obtained by running 10,000 replications. Table 4 presents a summary of the findings. Only three variables are compared: social welfare, output growth and inflation; for the latter two, both the first (bias) and the second moment (volatility) from the steady-state equilibrium are considered. Each entry is the ratio between the value computed in 10,000 replications and that obtained in 500 ones; for all ratios, the mean, the median, the maximum and the minimum across replications are shown.

According to the evidence presented in Table 4, the size of the small sample bias is negligible: regardless of the transparency regime, the difference in welfare does not reach 2 percentage points and the discrepancy is even smaller for the volatility of output growth and inflation. The estimates of the biases are less alike and sometimes even change sign, but this is no evidence of the existence of a significant small-sample bias: both the numerator and the denominator of the ratios are close to zero, so that even small differences can lead to high jumps in the ratio. The precision of the estimates based on few

replications is confirmed by looking at the ratios between the maxima and minima, which are surprisingly low.

3.5.2 Experiment #2: the size of Γ

A second type of sensitivity analysis has been conducted on initial conditions of the learning algorithm. A critical parameter is Γ , the moment matrix of the regressors entering the (generalised) stochastic gradient learning recursive equations: unlike the coefficients of the PLM, the matrix Γ is not updated, but held fixed at some assigned level. To assess the influence of the value of Γ on the ranking of the policy rules, other simulations have been run, using $k\Gamma$ as the moment matrix of the regressors. Six cases have been considered, corresponding to $k = \{\frac{9}{10}, \frac{11}{10}, \frac{3}{4}, \frac{5}{4}, \frac{1}{2}, \frac{3}{2}\}$. Table 5 shows the results for the three monetary regimes and the 6 values of k ; the entries in the table indicate the rank of each policy rule in terms of social welfare. In the final two rows, the Spearman ρ and the Kendall τ rank correlation coefficients are presented. The results are reassuring. In the full transparency case, there is no uncertainty about which is the welfare-maximising policy rule: all values of $k\Gamma$ point to the same rule. Something similar happens in the partial transparency case, where the optimal policy is identified for all values of k except $\frac{3}{2}$, while the ranking in the no-transparency regime seems to be somewhat more dependent on choice of Γ . The sample values of the rank correlation coefficients - surprisingly high in nearly all cases - confirm that the actual value of Γ is quite irrelevant not only in detecting the optimal policy, but also in ordering suboptimal ones.

3.5.3 Experiment #3: specification of the welfare function

The welfare function (3.12) used in the paper does not penalise interest rate instability and attributes the same importance to the volatility of inflation

and that of output growth. The first feature is justified on the grounds that in a sufficiently large model excess volatility of the monetary policy instrument trasmits to other asset prices, affecting private-sector spending decisions, so that it is implicitly incorporated in the volatility of output and inflation; the second feature reflects the desire to treat evenly fluctuations in nominal and real variables, as in Orphanides and Williams (2007).

As it is unclear which are the appropriate weights of the different arguments of the welfare functions, it is advisable to test how sensitive is the choice of the best-performing monetary policy rule to the specification of social preferences. A more general specification of the welfare function is

$$W = - \left[\zeta E (\pi_t - \bar{\pi})^2 + (1 - \zeta) E (\Delta y_t - \overline{\Delta y})^2 + \omega E (i_t - \bar{i})^2 \right] \quad (3.13)$$

The parameter ζ measures the degree of inflation aversion, while non-zero values of ω signal that society dislikes interest rate volatility as well. For $\zeta = .5$ and $\omega = 0$, (3.13) coincides with (3.12); $\zeta = .5$ and $\omega = 0.125$ are instead the values used in Orphanides and Williams (2007).

Tables 6a to 6c show how different combinations of the parameters (ζ, ω) affect the ranking of monetary policy rules. Each row of the table corresponds to a policy rule, while the columns refer to alternative values of the weights of the interest rate and inflation objectives relative to that of output growth. In the last two rows of the table, the Spearman's and Kendall's rank correlation coefficients are shown.

Three findings are worth stressing: (1) save the case when the degree of inflation aversion of the monetary policymaker is very low, changes in the weight of the inflation objective have no impact on the choice of the best-performing rule: the rank correlation coefficients is in all but one case not just high, but equal to 1; (2) adding interest rate volatility to the welfare function does not influence the ordering of the policy rules, unless its weight

is unreasonably high. Using the same specification as in Orphanides and Williams (2007), does not alter the results shown in table 2a to 2c; (3) in the full transparency case, the ranking of the policy rules turns out to be much more sensitive to the inclusion of interest rate volatility in the welfare function, though the main features of the best-performing policy changes only marginally. In general, more inertial and less activist policies seem to become more effective. The main rationale of this outcome is that under full transparency the short-term interest rate is much more volatile than inflation and output growth, so that even for low values of ω the shape of the welfare function changes in a non-negligible way.

3.5.4 Experiment #4: stochastic initial conditions for model proportions x_t^i

In all the experiments described so far, initial conditions for predictor proportions x_t^i are set equal to the reciprocal of the number of models used to forecast a given variable and are kept constant across PLMs and replications, under the presumption that this creates a level playing field for all competing forecasting models. To assess whether this assumption does indeed leave the model selection process unaffected, additional simulations are run, this time drawing initial conditions from a uniform distribution with support in (a finite subset of) \mathbb{R}^+ ; the constraints on model proportions are enforced by rescaling each draw so as to ensure a unit sum. Table 7 shows predictor proportions at the end of the simulation horizon under fixed and random initial conditions, together with two other statistics: the standardised difference between average predictor proportions and the correlation between initial conditions and limit values of model shares.

For all the variables - short-term interest rate, inflation and exchange rate - the table clearly shows that initial conditions do not matter much: the ranking of the models is the same regardless of the way initial conditions are set

and the standardised difference between predictor proportions is always well below one and quite close to zero. The correlation between initial and final values of the predictor proportions is in general non-negligible, suggesting that initial conditions are not irrelevant, though not important enough to change the long-run behaviour of the system.

3.6 Conclusions

This paper has analysed the properties of a large non-linear model populated by boundedly rational and incompletely informed agents. When the economy is sufficiently complex, individuals do not know the "true" law of motion of the variables they need to predict and are confronted with a host of equally plausible forecasting models. If agents can pick one out of a large number of predictors, none of which clearly superior, there is no guarantee that everyone selects the same one; in addition, they may choose to change to a different forecasting model if the predictive accuracy of the one that they are using deteriorates. Expectations therefore end up being misspecified, heterogeneous and ever-changing, even asymptotically, when enough observations are available to detect which forecasting model exhibits the best predictive performance. As the equilibrium to which the economy asymptotically converges differs from the REE and depends on the specific form of the expectations equations, central bank communication may be beneficial if it helps private agents to coordinate their beliefs. The paper is an attempt to assess whether in such a model economy the implications for monetary policymaking are similar to those found in the literature for small, linear systems and whether higher degrees of transparency are welfare-enhancing. The main findings are the following. First, expectations heterogeneity is an intrinsic feature of the economy: regardless of the monetary policy in place, no PLM succeeds in ruling out all the other forecasting models, though the

most inaccurate ones are eventually dismissed. Second, the monetary policymaker has much weaker incentives (than, e.g., in the paper by Orphanides and Williams) to adopt more inflation-averse policies, since too strong a reaction to price shocks increases both inflation and output volatility and tends to make the model unstable and non-learnable. At first sight, this outcome seems quite counterintuitive: a central bank that is committed to tame inflationary pressures is presumably more credible and more effective in anchoring long-run inflation expectations and bond yields. This connection however is not present in the model and credibility depends on outcomes, not intentions: agents learn from the data and what matters is whether monetary policy makes the economy more stable. Third, more transparent policies are in some cases mildly welfare-enhancing, but they never warrant sizeable improvements; the degree of transparency alters the form of the optimal policy rule also, as it increases inflation aversion. Disclosing more information is however not always beneficial.

Fig.1 – Social welfare under rational expectations





Fig.2 - Predictor proportions and payoff function



Note: The values reported on the x-axis are scaled by 1000, implying that the value $\lambda=1000$ used in the simulations corresponds to the value of 1 on the x-axis. The blue and green lines represent the sum of the predictor proportions of the two most successful forecasting models for the short-term interest rate and for inflation; the red line shows the share of agents adopting the best predictive model for the exchange rate.

Table 1 - Impact on welfare of the volatility of the main macroeconomic variables

The table reports the t -statistic of the simple regression of welfare and the standard deviation of the arguments of the welfare function on the volatility of a subset of the main macroeconomic variables included in the model. The first column lists the set of regressors; the subsequent ones, coming in groups of three (one group for each transparency regime), show the t -statistic of the regression of, respectively, welfare, the volatility of output growth ($\sigma_{\Delta y}$) and the volatility of inflation (σ_{π}) on the variable indicated in the first column. The notation is used as follows: π^e is expected inflation; $\pi - \pi^e$ is surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; k^* is the optimal capital-output ratio; $\Delta \pi^e$ is the difference between central bank and private-sector inflation expectations; i and i^L are the short (policy) and long-term interest rate; $i - \pi$ and $i^L - \pi$ are the corresponding real rates; $i - i^L$ is the term spread; e is the exchange rate.

	No Transparency			Partial Transparency			Full Transparency		
volatility of:	<i>Welfare</i>	$\sigma_{\Delta y}$	σ_{π}	<i>Welfare</i>	$\sigma_{\Delta y}$	σ_{π}	<i>Welfare</i>	$\sigma_{\Delta y}$	σ_{π}
π^e	-34.19	-4.52	71.86	-37.97	-5.55	68.02	-5.53	-8.05	7.38
$\pi - \pi^e$	-35.57	-4.53	193.68	-42.04	-5.54	180.24	-35.59	-5.15	42.75
Δw	-21.87	-5.12	37.32	-21.89	-6.23	33.43	-41.33	-4.23	19.24
Θ_0	-17.28	-5.40	26.67	-21.65	-6.24	31.79	0.09	1.24	-0.44
k^*	-1.03	-0.97	1.12	-0.12	-0.27	0.15	-4.19	-1.28	3.32
$\Delta \pi^e$	1.72	3.30	-1.88	1.50	2.52	-1.64	-6.39	-9.53	8.38
i	-0.03	0.63	-0.15	0.07	0.57	-0.17	-0.67	-1.89	1.04
i^L	0.20	0.76	-0.39	0.23	0.60	-0.32	0.19	-1.21	0.18
$i^L - \pi$	3.67	8.45	-4.38	3.37	5.21	-3.75	0.27	-1.15	0.11
$i - \pi$	1.29	1.93	-1.53	1.45	1.96	-1.59	-0.33	-1.75	0.72
$i^L - i$	0.10	0.84	-0.30	0.33	0.90	-0.45	0.47	-0.81	-0.11
e	-0.85	-1.00	0.95	-0.11	-0.31	0.14	0.53	-0.94	-0.14

Table 2a - Monetary policy effectiveness under no transparency
(decreasing gain learning)

The table reports a summary of the results (500 replications) of the model simulations (for the initial 140 periods). For each set of parameters of the central bank's interest-rate rule, mean (*bias*) and standard deviation (*volatility*) with respect to the steady-state values are reported (in p.p.). W stands for welfare (as a ratio to the optimum, reported in row 1) and RR for the rejection rate; Δy is the growth rate of GDP, π and $\pi - \pi^e$ are actual and surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; $\Delta \pi$ is the difference between central bank and private-sector inflation expectations; k^* is the optimal capital-output ratio; i and i^e are the actual and expected short-term (policy) interest rate; i^L is the yield on Treasury bonds; $H(\pi^e)$ and $H(i^e)$ are the entropy associated, respectively, with the choice of the forecasting model for π^e and i^e .

ρ	α_π	α_u		W	RR	Δy	π	$\pi - \pi^e$	Δw	Θ_0	$\Delta \pi^e$	k^*	i	i^L	$H(\pi^e)$	$H(i^e)$
$\rho=0.3$	$\alpha_\pi=0.4$	$\alpha_u=3.5$	vol.	-0.026	0.0	1.115	1.137	1.076	1.15	1.90	0.15	3.22	1.905	0.543	0.061	0.024
			bias			0.005	0.057	-0.003	0.06	1.89	-0.13	3.22	0.589	0.012	0.786	0.641
$\rho=0.8$			vol.	0.711	22.0	1.043	1.546	1.29	2.15	3.11	0.14	3.23	1.834	0.518	0.08	0.021
			bias			0.004	0.064	-0.01	0.07	2.96	-0.11	3.23	0.509	-0.053	0.71	0.744
$\rho=0.6$			vol.	0.852	0.8	1.052	1.340	1.18	1.68	2.53	0.14	3.22	1.763	0.492	0.07	0.022
			bias			0.003	0.058	-0.01	0.06	2.45	-0.11	3.22	0.652	0.026	0.75	0.720
$\rho=0.5$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.890	0.4	1.057	1.297	1.16	1.58	2.40	0.14	3.22	1.730	0.478	0.07	0.020
			bias			0.003	0.056	-0.01	0.06	2.34	-0.12	3.22	0.673	0.045	0.76	0.714
$\rho=0.4$			vol.	0.912	0.4	1.060	1.273	1.15	1.52	2.32	0.14	3.22	1.712	0.470	0.07	0.021
			bias			0.003	0.056	-0.01	0.06	2.27	-0.12	3.22	0.688	0.058	0.76	0.709
$\rho=0.3$			vol.	0.922	0.4	1.063	1.260	1.14	1.48	2.27	0.14	3.22	1.704	0.466	0.07	0.022
			bias			0.003	0.055	-0.01	0.06	2.23	-0.12	3.21	0.699	0.069	0.76	0.706
	$\alpha_\pi=0.5$		vol.	0.814	4.4	1.046	1.394	1.21	1.80	2.75	0.14	3.24	1.623	0.481	0.08	0.021
			bias			0.004	0.073	-0.01	0.08	2.65	-0.11	3.23	0.400	-0.080	0.74	0.713
	$\alpha_\pi=1.0$		vol.	0.804	4.0	1.047	1.405	1.22	1.83	2.78	0.14	3.23	1.795	0.505	0.08	0.023
			bias			0.003	0.060	-0.01	0.06	2.68	-0.11	3.22	0.608	-0.004	0.73	0.729
$\rho=0.7$	$\alpha_\pi=1.5$	$\alpha_u=2.0$	vol.	0.787	3.2	1.049	1.424	1.23	1.88	2.80	0.13	3.22	1.988	0.551	0.08	0.027
			bias			0.003	0.050	-0.01	0.05	2.70	-0.11	3.21	0.796	0.062	0.73	0.743
	$\alpha_\pi=2.0$		vol.	0.750	1.6	1.052	1.460	1.25	1.97	2.82	0.13	3.21	2.215	0.617	0.08	0.034
			bias			0.002	0.044	-0.01	0.05	2.71	-0.10	3.21	0.971	0.116	0.73	0.752
	$\alpha_\pi=2.5$		vol.	0.764	1.6	1.055	1.453	1.24	1.97	2.82	0.13	3.20	2.379	0.662	0.08	0.039
			bias			0.002	0.037	-0.01	0.04	2.71	-0.10	3.20	1.114	0.166	0.72	0.754
		$\alpha_u=1.0$	vol.	0.684	55.6	1.040	1.607	1.32	2.24	3.18	0.14	3.23	1.436	0.367	0.08	0.020
			bias			-0.002	0.082	-0.03	0.09	3.02	-0.11	3.23	0.554	0.009	0.70	0.749
	$\alpha_\pi=1.0$	$\alpha_u=1.5$	vol.	0.711	16.0	1.041	1.546	1.29	2.11	3.06	0.14	3.23	1.712	0.472	0.08	0.021
			bias			0.002	0.072	-0.02	0.08	2.92	-0.11	3.23	0.561	-0.011	0.71	0.739
$\rho=0.7$		$\alpha_u=2.0$	vol.	0.804	4.0	1.047	1.405	1.22	1.83	2.78	0.14	3.23	1.795	0.505	0.08	0.023
			bias			0.003	0.060	-0.01	0.06	2.68	-0.11	3.22	0.608	-0.004	0.73	0.729
		$\alpha_u=2.5$	vol.	0.863	0.0	1.057	1.320	1.17	1.66	2.49	0.14	3.22	1.887	0.536	0.07	0.024
			bias			0.004	0.056	-0.01	0.06	2.42	-0.11	3.22	0.657	0.006	0.76	0.723

Table 2b - Monetary policy effectiveness with transparency
(decreasing gain learning)

The table reports a summary of the results (500 replications) of the model simulations (for the initial 140 periods). For each set of parameters of the central bank's interest-rate rule, mean (*bias*) and standard deviation (*volatility*) with respect to the steady-state values are reported (in p.p.). W stands for welfare (as a ratio to the optimum, reported in row 1) and RR for the rejection rate; Δy is the growth rate of GDP, π and $\pi - \pi^e$ are actual and surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; $\Delta \pi$ is the difference between central bank and private-sector inflation expectations; k^* is the optimal capital-output ratio; i and i^e are the actual and expected short-term (policy) interest rate; i^L is the yield on Treasury bonds; $H(\pi^e)$ and $H(i^e)$ are the entropy associated, respectively, with the choice of the forecasting model for π^e and i^e .

				W	RR	Δy	π	$\pi-\pi^e$	Δw	Θ_0	$\Delta\pi^e$	k^*	i	i^L	$H(\pi^e)$	$H(i^e)$
$\rho=0.26$	$\alpha_\pi=0.7$	$\alpha_u=3.2$	vol.	-0.025	0.0	1.106	1.118	1.064	1.129	1.910	0.148	3.219	1.781	0.543	0.057	0.025
			bias			0.005	0.057	-0.005	0.057	1.900	-0.126	3.217	0.616	0.045	0.797	0.649
$\rho=0.8$			vol.	0.730	17.2	1.041	1.490	1.260	2.041	2.950	0.137	3.225	1.737	0.525	0.075	0.023
			bias			0.003	0.060	-0.015	0.064	2.820	-0.108	3.222	0.489	0.023	0.726	0.757
$\rho=0.6$			vol.	0.878	0.4	1.051	1.288	1.154	1.580	2.410	0.137	3.217	1.656	0.507	0.067	0.021
			bias			0.003	0.057	-0.011	0.059	2.350	-0.115	3.214	0.621	0.076	0.766	0.729
$\rho=0.5$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.902	0.0	1.056	1.257	1.137	1.498	2.310	0.138	3.215	1.638	0.502	0.064	0.022
			bias			0.003	0.058	-0.012	0.061	2.250	-0.117	3.213	0.647	0.092	0.773	0.720
$\rho=0.4$			vol.	0.923	0.0	1.059	1.235	1.125	1.442	2.240	0.139	3.213	1.624	0.499	0.062	0.024
			bias			0.003	0.057	-0.012	0.059	2.200	-0.119	3.211	0.661	0.105	0.777	0.714
$\rho=0.3$			vol.	0.933	0.0	1.062	1.224	1.119	1.408	2.200	0.139	3.211	1.618	0.500	0.061	0.026
			bias			0.003	0.056	-0.013	0.058	2.160	-0.120	3.209	0.672	0.119	0.778	0.711
	$\alpha_\pi=0.5$		vol.	0.824	2.4	1.047	1.350	1.185	1.714	2.610	0.140	3.233	1.539	0.496	0.071	0.023
			bias			0.004	0.074	-0.013	0.079	2.520	-0.114	3.230	0.387	-0.042	0.752	0.721
	$\alpha_\pi=1.0$		vol.	0.829	2.8	1.046	1.351	1.186	1.729	2.630	0.137	3.220	1.690	0.516	0.073	0.022
			bias			0.003	0.058	-0.012	0.061	2.540	-0.112	3.217	0.581	0.056	0.748	0.741
$\rho=0.7$	$\alpha_\pi=1.5$	$\alpha_u=2.0$	vol.	0.802	2.0	1.047	1.376	1.201	1.791	2.660	0.134	3.208	1.881	0.572	0.075	0.021
			bias			0.002	0.045	-0.013	0.047	2.560	-0.109	3.206	0.751	0.147	0.744	0.758
	$\alpha_\pi=2.0$		vol.	0.784	1.2	1.049	1.397	1.213	1.846	2.680	0.132	3.198	2.071	0.641	0.077	0.023
			bias			0.001	0.034	-0.015	0.035	2.580	-0.106	3.195	0.897	0.235	0.740	0.766
	$\alpha_\pi=2.5$		vol.	0.784	1.2	1.051	1.403	1.217	1.874	2.690	0.130	3.188	2.233	0.709	0.078	0.032
			bias			0.001	0.025	-0.015	0.025	2.590	-0.104	3.185	1.017	0.318	0.735	0.759
	$\alpha_u=1.0$		vol.	0.681	48.4	1.032	1.587	1.313	2.198	3.120	0.140	3.221	1.410	0.411	0.082	0.020
			bias			-0.002	0.080	-0.033	0.089	2.960	-0.112	3.218	0.545	0.095	0.708	0.764
	$\alpha_u=1.5$		vol.	0.728	11.6	1.035	1.494	1.261	2.013	2.960	0.138	3.222	1.633	0.493	0.075	0.022
			bias			0.001	0.066	-0.021	0.069	2.820	-0.110	3.218	0.546	0.063	0.726	0.752
$\rho=0.7$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.829	2.8	1.046	1.351	1.186	1.729	2.630	0.137	3.220	1.690	0.516	0.073	0.022
			bias			0.003	0.058	-0.012	0.061	2.540	-0.112	3.217	0.581	0.056	0.748	0.741
	$\alpha_u=2.5$		vol.	0.895	0.0	1.056	1.266	1.143	1.555	2.380	0.137	3.219	1.762	0.540	0.067	0.021
			bias			0.004	0.056	-0.007	0.058	2.320	-0.114	3.217	0.609	0.051	0.770	0.733

Table 2c - Monetary policy effectiveness with full transparency
(decreasing gain learning)

The table reports a summary of the results (500 replications) of the model simulations (for the initial 140 periods). For each set of parameters of the central bank's interest-rate rule, mean (*bias*) and standard deviation (*volatility*) with respect to the steady-state values are reported (in p.p.). *W* stands for welfare (as a ratio to the optimum, reported in row 1) and *RR* for the rejection rate; Δy is the growth rate of GDP, π and $\pi - \pi^e$ are actual and surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; $\Delta \pi^e$ is the difference between central bank and private-sector inflation expectations; k^* is the optimal capital-output ratio; i and i^e are the actual and expected short-term (policy) interest rate; i^L is the yield on Treasury bonds; $H(\pi^e)$ and $H(i^e)$ are the entropy associated, respectively, with the choice of the forecasting model for π^e and i^e .

		<i>W</i>	<i>RR</i>	Δy	π	$\pi - \pi^e$	Δw	Θ_0	$\Delta \pi^e$	k^*	i	i^L	$H(\pi^e)$	$H(i^e)$
$\rho=0.72$	$\alpha_\pi=1.4$	<i>vol.</i>	<i>bias</i>	1.053	1.202	1.096	1.443	0.350	0.124	2.967	1.557	4.169	0.036	0.022
	$\alpha_u=2.3$			-0.014	-0.163	-0.010	-0.128	0.130	-0.088	2.965	0.595	3.897	0.214	0.780
$\rho=0.8$		<i>vol.</i>	<i>bias</i>	1.044	1.260	1.134	1.582	0.360	0.126	2.998	1.316	3.098	0.038	0.027
				-0.016	-0.145	-0.001	-0.143	0.130	-0.087	2.997	0.532	2.816	0.220	0.811
$\rho=0.6$		<i>vol.</i>	<i>bias</i>	1.031	1.252	1.122	1.507	0.350	0.128	2.965	1.654	4.474	0.034	0.023
				-0.012	-0.155	-0.022	-0.105	0.120	-0.089	2.962	0.719	4.227	0.211	0.768
$\rho=0.5$	$\alpha_\pi=1.0$	<i>vol.</i>	<i>bias</i>	1.022	1.303	1.149	1.591	0.340	0.131	2.969	1.835	4.924	0.033	0.022
	$\alpha_u=2.0$			-0.008	-0.137	-0.037	-0.059	0.120	-0.090	2.966	0.802	4.667	0.207	0.753
$\rho=0.4$		<i>vol.</i>	<i>bias</i>	1.013	1.363	1.182	1.695	0.340	0.134	2.977	2.006	5.264	0.033	0.021
				-0.004	-0.113	-0.052	-0.008	0.120	-0.090	2.974	0.880	4.990	0.204	0.742
$\rho=0.3$		<i>vol.</i>	<i>bias</i>	1.006	1.426	1.218	1.806	0.340	0.137	2.987	2.162	5.522	0.032	0.021
				0.001	-0.087	-0.067	0.044	0.120	-0.091	2.983	0.953	5.230	0.202	0.733
	$\alpha_\pi=0.5$	<i>vol.</i>	<i>bias</i>	1.033	1.233	1.112	1.481	0.350	0.128	2.970	1.507	3.975	0.036	0.027
				-0.015	-0.161	-0.012	-0.134	0.130	-0.089	2.968	0.747	3.731	0.216	0.783
	$\alpha_\pi=1.0$	<i>vol.</i>	<i>bias</i>	1.039	1.223	1.106	1.473	0.350	0.126	2.971	1.473	3.878	0.036	0.025
				-0.015	-0.161	-0.010	-0.138	0.130	-0.089	2.969	0.633	3.631	0.215	0.786
$\rho=0.7$	$\alpha_\pi=1.5$	<i>vol.</i>	<i>bias</i>	1.046	1.217	1.106	1.473	0.350	0.123	2.973	1.465	3.795	0.036	0.024
	$\alpha_u=2.0$			-0.015	-0.160	-0.008	-0.140	0.130	-0.088	2.971	0.530	3.542	0.215	0.788
	$\alpha_\pi=2.0$	<i>vol.</i>	<i>bias</i>	1.053	1.214	1.106	1.481	0.350	0.122	2.975	1.479	3.725	0.036	0.024
				-0.016	-0.158	-0.006	-0.141	0.130	-0.087	2.973	0.440	3.464	0.214	0.789
	$\alpha_\pi=2.5$	<i>vol.</i>	<i>bias</i>	1.060	1.214	1.107	1.494	0.350	0.120	2.977	1.511	3.669	0.036	0.024
				-0.016	-0.157	-0.005	-0.141	0.130	-0.086	2.975	0.362	3.396	0.213	0.789
	$\alpha_u=1.0$	<i>vol.</i>	<i>bias</i>	1.020	1.451	1.243	1.927	0.360	0.137	3.041	1.213	2.282	0.039	0.027
				-0.013	-0.106	-0.010	-0.115	0.140	-0.088	3.038	0.527	2.044	0.220	0.844
	$\alpha_u=1.5$	<i>vol.</i>	<i>bias</i>	1.028	1.283	1.146	1.600	0.360	0.128	2.998	1.297	3.056	0.037	0.028
	$\alpha_\pi=1.0$			-0.016	-0.141	-0.006	-0.141	0.130	-0.089	2.996	0.568	2.830	0.218	0.814
$\rho=0.7$	$\alpha_u=2.0$	<i>vol.</i>	<i>bias</i>	1.039	1.223	1.106	1.473	0.350	0.126	2.971	1.473	3.878	0.036	0.025
				-0.015	-0.161	-0.006	-0.138	0.130	-0.089	2.969	0.633	3.631	0.215	0.786
	$\alpha_u=2.5$	<i>vol.</i>	<i>bias</i>	1.049	1.210	1.098	1.445	0.350	0.127	2.969	1.765	4.802	0.035	0.020
				-0.009	-0.152	-0.023	-0.083	0.120	-0.090	2.966	0.802	4.514	0.211	0.765

Table 3a - Monetary policy effectiveness under no transparency
(constant gain learning)

The table reports a summary of the results (500 replications) of the model simulations (for the initial 140 periods). For each set of parameters of the central bank's interest-rate rule, mean (*bias*) and standard deviation (*volatility*) with respect to the steady-state values are reported (in p.p.). W stands for welfare (as a ratio to the optimum, reported in row 1) and RR for the rejection rate; Δy is the growth rate of GDP, π and $\pi - \pi^e$ are actual and surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; $\Delta \pi$ is the difference between central bank and private-sector inflation expectations; k^* is the optimal capital-output ratio; i and i^e are the actual and expected short-term (policy) interest rate; i^L is the yield on Treasury bonds; $H(\pi^e)$ and $H(i^e)$ are the entropy associated, respectively, with the choice of the forecasting model for π^e and i^e .

ρ	α_π	α_u		W	RR	Δy	π	$\pi - \pi^e$	Δw	Θ_0	$\Delta \pi^e$	k^*	i	i^L	$H(\pi^e)$	$H(i^e)$
$\rho=0.31$	$\alpha_\pi=0.3$	$\alpha_u=3.2$	vol.	-0.026	0.0	1.112	1.136	1.074	1.16	1.93	0.15	3.22	1.801	0.575	0.064	0.018
			bias			0.004	0.047	-0.018	0.05	1.92	-0.12	3.22	0.884	0.024	0.783	0.621
$\rho=0.8$			vol.	0.747	26.4	1.038	1.497	1.25	2.05	3.23	0.13	3.22	1.776	0.571	0.07	0.023
			bias			0.002	0.052	-0.03	0.05	3.03	-0.10	3.22	0.589	0.014	0.71	0.722
$\rho=0.6$			vol.	0.872	1.2	1.052	1.314	1.16	1.63	2.58	0.13	3.21	1.737	0.552	0.07	0.019
			bias			0.002	0.044	-0.02	0.04	2.49	-0.11	3.21	0.715	0.080	0.75	0.700
$\rho=0.5$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.915	0.8	1.058	1.269	1.14	1.52	2.44	0.14	3.21	1.700	0.537	0.07	0.019
			bias			0.002	0.042	-0.02	0.04	2.37	-0.11	3.21	0.731	0.096	0.76	0.693
$\rho=0.4$			vol.	0.929	0.4	1.062	1.250	1.13	1.47	2.36	0.14	3.21	1.690	0.532	0.07	0.020
			bias			0.002	0.041	-0.02	0.04	2.30	-0.11	3.21	0.744	0.107	0.76	0.688
$\rho=0.3$			vol.	0.941	0.4	1.066	1.236	1.12	1.43	2.30	0.14	3.21	1.679	0.527	0.07	0.022
			bias			0.002	0.041	-0.02	0.04	2.25	-0.11	3.21	0.752	0.115	0.76	0.685
	$\alpha_\pi=0.5$		vol.	0.850	5.2	1.044	1.351	1.18	1.71	2.78	0.14	3.22	1.582	0.522	0.07	0.020
			bias			0.002	0.059	-0.03	0.06	2.66	-0.11	3.22	0.501	-0.014	0.74	0.689
	$\alpha_\pi=1.0$		vol.	0.824	4.8	1.045	1.380	1.20	1.78	2.84	0.13	3.22	1.770	0.567	0.07	0.021
			bias			0.002	0.046	-0.03	0.05	2.71	-0.10	3.21	0.678	0.055	0.74	0.710
$\rho=0.7$	$\alpha_\pi=1.5$	$\alpha_u=2.0$	vol.	0.793	4.0	1.047	1.413	1.21	1.86	2.89	0.13	3.21	1.980	0.627	0.07	0.023
			bias			0.002	0.037	-0.03	0.04	2.76	-0.10	3.21	0.843	0.111	0.73	0.732
	$\alpha_\pi=2.0$		vol.	0.775	3.6	1.051	1.435	1.23	1.92	2.92	0.13	3.20	2.182	0.687	0.07	0.026
			bias			0.001	0.031	-0.03	0.03	2.78	-0.09	3.20	0.995	0.158	0.73	0.751
	$\alpha_\pi=2.5$		vol.	0.741	2.0	1.056	1.474	1.25	2.01	2.93	0.13	3.20	2.422	0.762	0.07	0.031
			bias			0.001	0.027	-0.03	0.03	2.78	-0.09	3.20	1.147	0.196	0.72	0.760
		$\alpha_u=1.0$	vol.	0.790	67.2	1.025	1.460	1.24	1.96	3.15	0.14	3.22	1.295	0.393	0.07	0.020
			bias			-0.003	0.062	-0.04	0.06	2.96	-0.11	3.21	0.627	0.072	0.72	0.718
		$\alpha_u=1.5$	vol.	0.770	22.4	1.035	1.468	1.24	1.96	3.15	0.13	3.22	1.627	0.517	0.07	0.021
			bias			0.000	0.058	-0.04	0.06	2.96	-0.10	3.22	0.621	0.049	0.72	0.715
$\rho=0.7$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.824	4.8	1.045	1.380	1.20	1.78	2.84	0.13	3.22	1.770	0.567	0.07	0.021
			bias			0.002	0.046	-0.03	0.05	2.71	-0.10	3.21	0.678	0.055	0.74	0.710
		$\alpha_u=2.5$	vol.	0.889	0.8	1.057	1.293	1.15	1.61	2.53	0.13	3.21	1.857	0.592	0.07	0.021
			bias			0.002	0.040	-0.02	0.04	2.44	-0.10	3.21	0.730	0.066	0.75	0.707

Table 3b - Monetary policy effectiveness with full transparency
(constant gain learning)

The table reports a summary of the results (500 replications) of the model simulations (for the initial 140 periods). For each set of parameters of the central bank's interest-rate rule, mean (*bias*) and standard deviation (*volatility*) with respect to the steady-state values are reported (in p.p.). W stands for welfare (as a ratio to the optimum, reported in row 1) and RR for the rejection rate; Δy is the growth rate of GDP, π and $\pi - \pi^e$ are actual and surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; $\Delta \pi$ is the difference between central bank and private-sector inflation expectations; k^* is the optimal capital-output ratio; i and i^e are the actual and expected short-term (policy) interest rate; i^L is the yield on Treasury bonds; $H(\pi^e)$ and $H(i^e)$ are the entropy associated, respectively, with the choice of the forecasting model for π^e and i^e .

				W	RR	Δy	π	$\pi-\pi^e$	Δw	Θ_0	$\Delta\pi^e$	k^*	i	i^L	$H(\pi^e)$	$H(i^e)$
$\rho=0.26$	$\alpha_\pi=0.7$	$\alpha_u=3.0$	vol.	-0.025	0.0	1.106	1.118	1.064	1.137	1.950	0.147	3.213	1.731	0.577	0.058	0.027
			bias			0.003	0.045	-0.017	0.043	1.930	-0.121	3.211	0.661	0.073	0.793	0.642
$\rho=0.8$			vol.	0.794	23.6	1.033	1.416	1.213	1.896	3.050	0.133	3.215	1.645	0.561	0.067	0.025
			bias			0.002	0.049	-0.028	0.050	2.880	-0.101	3.212	0.559	0.073	0.728	0.743
$\rho=0.6$			vol.	0.900	0.8	1.052	1.266	1.140	1.536	2.490	0.135	3.208	1.634	0.559	0.065	0.022
			bias			0.002	0.045	-0.024	0.045	2.410	-0.109	3.206	0.675	0.119	0.763	0.718
$\rho=0.5$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.931	0.4	1.057	1.231	1.122	1.449	2.370	0.135	3.207	1.611	0.551	0.063	0.023
			bias			0.002	0.044	-0.023	0.043	2.310	-0.111	3.205	0.694	0.133	0.769	0.709
$\rho=0.4$			vol.	0.949	0.4	1.061	1.211	1.111	1.394	2.300	0.136	3.205	1.597	0.547	0.063	0.025
			bias			0.002	0.043	-0.022	0.043	2.250	-0.113	3.203	0.705	0.144	0.771	0.704
$\rho=0.3$			vol.	0.945	0.0	1.064	1.208	1.109	1.375	2.260	0.137	3.204	1.605	0.552	0.063	0.027
			bias			0.002	0.043	-0.023	0.044	2.200	-0.114	3.202	0.718	0.158	0.773	0.703
	$\alpha_\pi=0.5$		vol.	0.879	4.0	1.045	1.298	1.153	1.611	2.660	0.136	3.223	1.482	0.526	0.067	0.023
			bias			0.003	0.060	-0.028	0.064	2.560	-0.108	3.220	0.464	0.007	0.752	0.706
	$\alpha_\pi=1.0$		vol.	0.842	3.2	1.045	1.334	1.173	1.693	2.730	0.134	3.211	1.672	0.573	0.068	0.024
			bias			0.002	0.046	-0.026	0.045	2.620	-0.105	3.209	0.642	0.101	0.748	0.731
$\rho=0.7$	$\alpha_\pi=1.5$	$\alpha_u=2.0$	vol.	0.817	2.8	1.046	1.362	1.189	1.762	2.780	0.131	3.200	1.867	0.637	0.069	0.023
			bias			0.001	0.034	-0.025	0.032	2.650	-0.101	3.198	0.795	0.191	0.743	0.753
	$\alpha_\pi=2.0$		vol.	0.795	2.4	1.048	1.388	1.203	1.830	2.820	0.129	3.191	2.063	0.713	0.071	0.021
			bias			0.000	0.023	-0.024	0.021	2.680	-0.098	3.188	0.929	0.275	0.737	0.762
	$\alpha_\pi=2.5$		vol.	0.771	1.6	1.051	1.417	1.219	1.902	2.830	0.127	3.182	2.268	0.796	0.073	0.028
			bias			-0.001	0.015	-0.024	0.013	2.700	-0.094	3.180	1.050	0.354	0.732	0.754
	$\alpha_u=1.0$		vol.	0.779	60.8	1.023	1.453	1.233	1.949	3.140	0.135	3.210	1.283	0.435	0.072	0.020
			bias			-0.002	0.061	-0.039	0.065	2.950	-0.106	3.207	0.610	0.147	0.722	0.748
	$\alpha_u=1.5$		vol.	0.772	16.8	1.034	1.438	1.224	1.907	3.090	0.135	3.213	1.573	0.539	0.069	0.022
			bias			0.000	0.056	-0.035	0.058	2.910	-0.103	3.210	0.602	0.110	0.727	0.739
$\rho=0.7$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.842	3.2	1.045	1.334	1.173	1.693	2.730	0.134	3.211	1.672	0.573	0.068	0.024
			bias			0.002	0.046	-0.026	0.045	2.620	-0.105	3.209	0.642	0.101	0.748	0.731
	$\alpha_u=2.5$		vol.	0.901	0.0	1.057	1.256	1.135	1.532	2.440	0.134	3.210	1.755	0.598	0.066	0.025
			bias			0.002	0.042	-0.021	0.042	2.370	-0.107	3.208	0.678	0.098	0.768	0.724

Table 3c - Monetary policy effectiveness with full transparency
(constant gain learning)

The table reports a summary of the results (500 replications) of the model simulations (for the initial 140 periods). For each set of parameters of the central bank's interest-rate rule, mean (*bias*) and standard deviation (*volatility*) with respect to the steady-state values are reported (in p.p.). W stands for welfare (as a ratio to the optimum, reported in row 1) and RR for the rejection rate; Δy is the growth rate of GDP, π and $\pi - \pi^e$ are actual and surprise inflation; Δw is wage growth; Θ_0 is the mean intercept of the forecasting equations for inflation; $\Delta \pi$ is the difference between central bank and private-sector inflation expectations; k^* is the optimal capital-output ratio; i and i^e are the actual and expected short-term (policy) interest rate; i^L is the yield on Treasury bonds; $H(\pi^e)$ and $H(i^e)$ are the entropy associated, respectively, with the choice of the forecasting model for π^e and i^e .

				W	RR	Δy	π	$\pi\text{-}\pi^e$	Δw	Θ_0	$\Delta\pi^e$	k^*	i	i^L	$H(\pi^e)$	$H(i^e)$
$\rho=0.72$	$\alpha_\pi=1.1$	$\alpha_u=2.5$	vol.	-0.026	0.0	1.061	1.191	1.089	1.411	0.350	0.127	2.968	1.619	4.359	0.036	0.031
			bias			-0.013	-0.160	0.026	-0.117	0.130	-0.087	2.966	0.666	4.061	0.214	0.740
$\rho=0.8$			vol.	0.947	0.0	1.045	1.253	1.129	1.565	0.360	0.126	3.005	1.295	2.999	0.038	0.034
			bias			-0.016	-0.140	0.024	-0.141	0.130	-0.082	3.003	0.528	2.712	0.221	0.787
$\rho=0.6$			vol.	0.971	0.0	1.039	1.239	1.114	1.481	0.350	0.129	2.967	1.608	4.274	0.034	0.030
			bias			-0.014	-0.156	0.022	-0.115	0.130	-0.088	2.965	0.681	4.024	0.212	0.733
$\rho=0.5$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.934	0.0	1.031	1.286	1.137	1.555	0.350	0.132	2.966	1.775	4.689	0.033	0.028
			bias			-0.011	-0.144	0.013	-0.081	0.120	-0.089	2.964	0.746	4.427	0.208	0.715
$\rho=0.4$			vol.	0.889	0.0	1.024	1.341	1.166	1.651	0.340	0.136	2.970	1.933	5.001	0.033	0.026
			bias			-0.008	-0.127	0.004	-0.043	0.120	-0.089	2.967	0.806	4.723	0.205	0.700
$\rho=0.3$			vol.	0.841	0.0	1.017	1.400	1.196	1.754	0.340	0.139	2.975	2.080	5.239	0.032	0.026
			bias			-0.005	-0.109	-0.006	-0.006	0.120	-0.089	2.971	0.861	4.941	0.202	0.690
	$\alpha_\pi=0.5$		vol.	0.988	0.0	1.038	1.221	1.105	1.456	0.350	0.130	2.976	1.455	3.798	0.036	0.032
			bias			-0.016	-0.157	0.025	-0.136	0.130	-0.087	2.974	0.697	3.549	0.217	0.751
	$\alpha_\pi=1.0$		vol.	0.991	0.0	1.044	1.213	1.103	1.454	0.350	0.127	2.977	1.440	3.723	0.036	0.032
			bias			-0.016	-0.157	0.026	-0.139	0.130	-0.086	2.975	0.612	3.472	0.216	0.757
$\rho=0.7$	$\alpha_\pi=1.5$	$\alpha_u=2.0$	vol.	0.989	0.0	1.050	1.210	1.103	1.459	0.350	0.124	2.979	1.447	3.660	0.036	0.032
			bias			-0.016	-0.157	0.027	-0.141	0.130	-0.085	2.977	0.536	3.402	0.216	0.763
	$\alpha_\pi=2.0$		vol.	0.985	0.0	1.056	1.209	1.104	1.470	0.350	0.122	2.980	1.472	3.608	0.036	0.031
			bias			-0.016	-0.156	0.028	-0.142	0.130	-0.084	2.979	0.468	3.341	0.215	0.769
	$\alpha_\pi=2.5$		vol.	0.977	0.0	1.063	1.211	1.106	1.487	0.350	0.120	2.982	1.514	3.568	0.036	0.031
			bias			-0.016	-0.156	0.029	-0.143	0.130	-0.083	2.980	0.410	3.289	0.214	0.774
	$\alpha_u=1.0$		vol.	0.793	4.0	1.017	1.447	1.233	1.919	0.360	0.137	3.047	1.225	2.235	0.039	0.030
			bias			-0.014	-0.101	0.008	-0.113	0.140	-0.081	3.044	0.547	1.991	0.220	0.826
	$\alpha_u=1.5$		vol.	0.936	0.0	1.030	1.277	1.142	1.588	0.360	0.128	3.004	1.291	2.964	0.037	0.034
			bias			-0.016	-0.136	0.021	-0.137	0.130	-0.085	3.003	0.576	2.731	0.219	0.790
$\rho=0.7$	$\alpha_\pi=1.0$	$\alpha_u=2.0$	vol.	0.991	0.0	1.044	1.213	1.103	1.454	0.350	0.127	2.977	1.440	3.723	0.036	0.032
			bias			-0.016	-0.157	0.026	-0.139	0.130	-0.086	2.975	0.612	3.472	0.216	0.757
	$\alpha_u=2.5$		vol.	0.995	0.0	1.058	1.199	1.092	1.419	0.350	0.128	2.967	1.683	4.551	0.035	0.029
			bias			-0.012	-0.157	0.024	-0.104	0.130	-0.088	2.965	0.716	4.260	0.212	0.733

Table 4 - Sensitivity analysis: number of replications*(decreasing gain learning)*

Each entry in the table is the ratio between the value of the first or second moment of welfare, output growth and inflation computed on 10,000 and 500 replications. For output growth and inflation not only the mean, but also the maximum, minimum and the median of each set of replications are presented.

NO TRANSPARENCY				
Welfare ratio = 0.991				
<i>volatility ratios</i>				
	max	min	mean	median
Δy	1.036	0.845	1.010	1.006
π	1.172	0.924	1.000	1.003
<i>bias ratios</i>				
	max	min	mean	median
Δy	1.110	1.016	1.041	1.125
π	1.096	0.360	1.005	1.010
PARTIAL TRANSPARENCY				
Welfare ratio = 0.990				
<i>volatility ratios</i>				
	max	min	mean	median
Δy	1.046	0.848	1.011	1.008
π	1.167	0.898	1.000	1.002
<i>bias ratios</i>				
	max	min	mean	median
Δy	1.113	1.000	1.043	1.143
π	1.073	0.457	1.005	1.020
FULL TRANSPARENCY				
Welfare ratio = 0.989				
<i>volatility ratios</i>				
	max	min	mean	median
Δy	1.066	0.889	1.012	1.020
π	1.078	0.894	1.001	1.008
<i>bias ratios</i>				
	max	min	mean	median
Δy	1.352	1.010	0.993	0.987
π	0.990	1.324	1.002	0.997

Table 5 - Sensitivity analysis: initial conditions
(constant gain learning)

The table reports the ranking in terms of welfare of the competing policy rules for a set of values of the Γ matrix (the normalising factor of the generalised stochastic gradient algorithm). The values of Γ considered are (1) the one used in the baseline simulations; (2) Γ scaled up and down by 10 p.p.; (3) Γ multiplied by 1.25 and 0.75; (4) Γ set equal to 1.5 and 0.5 times the benchmark value. As in the previous tables, only the initial 140 observations are used in computing the welfare ranking. For each policy rule, the model is simulated 500 times. Each row of the table refers to a policy rule, while the columns are divided into three subgroups, corresponding to the alternative monetary regimes (i.e. no transparency, partial transparency and full transparency). In the last two rows of the table, the Spearman's and Kendall's rank correlation coefficients are shown.

	<i>no transparency</i>								<i>partial transparency</i>								<i>full transparency</i>							
<i>k</i> Γ where <i>k</i> is:	1	0.9	1.1	3/4	5/4	1/2	3/2	1	0.9	1.1	3/4	5/4	1/2	3/2	1	0.9	1.1	3/4	5/4	1/2	3/2			
<i>ρ</i> =0.8	11	12	11	12	11	11	12	11	11	12	12	12	12	12	8	8	9	8	8	8	8			
<i>ρ</i> =0.6	5	5	5	5	5	5	5	5	5	5	5	5	5	5	7	7	7	7	7	7	7			
<i>ρ</i> =0.5 <i>α_π</i> =1.0 <i>α_u</i> =2.0	3	3	3	3	3	3	3	3	3	4	3	3	3	3	10	10	8	10	10	10	10			
<i>ρ</i> =0.4	2	2	2	2	2	2	2	2	2	2	2	2	2	2	11	11	11	11	11	11	11			
<i>ρ</i> =0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	12	13	12	12	12	12	12			
<i>α_π</i> =0.5	6	6	6	6	6	6	6	7	6	6	6	6	6	6	5	6	3	6	4	5	5			
<i>α_π</i> =1.0	7	7	7	7	7	7	7	6	7	7	7	7	7	7	3	4	2	3	2	3	2			
<i>ρ</i> =0.7 <i>α_π</i> =1.5 <i>α_u</i> =2.0	8	8	8	8	8	8	8	8	8	8	8	8	8	8	2	2	4	2	3	2	3			
<i>α_π</i> =2.0	10	9	9	9	10	10	10	9	9	9	9	9	9	9	4	3	5	4	5	4	4			
<i>α_π</i> =2.5	9	10	10	10	9	9	9	10	10	10	10	10	10	10	6	5	6	5	6	6	6			
<i>α_u</i> =1.0	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	12	13	13	13	13	13			
<i>ρ</i> =0.7 <i>α_π</i> =1.0 <i>α_u</i> =1.5	12	11	12	11	12	12	11	12	12	11	11	11	11	11	9	9	10	9	9	9	9			
<i>α_u</i> =2.5	4	4	4	4	4	4	4	4	4	3	4	4	4	4	1	1	1	1	1	1	1			
Spearman <i>ρ</i> (%)	0.99	0.99	0.99	1.00	1.00	1.00	0.99	0.99	0.98	0.99	0.99	0.99	0.99	0.99	0.98	0.96	0.99	0.99	1.00	0.99				
Kendall <i>τ</i> (%)	0.95	0.97	0.95	1.00	1.00	0.97		0.97	0.92	0.95	0.95	0.95	0.95		0.92	0.87	0.97	0.95	1.00	0.97				

Table 6a - Sensitivity analysis: specification of the welfare function
(no transparency)

The table reports the rankings in terms of welfare of the competing policy rules for alternative specifications of the welfare function. The parameter ω is the weight attached to the unconditional variance of the interest rate (σ^2 while the parameter ζ and $(1-\zeta)$ measure, respectively, the relevance of inflation (σ^2_{π}) and output growth ($\sigma^2_{\Delta y}$) volatility. Each row of the table corresponds to a policy rule, while the columns refer to alternative values of the weights of the interest rate and inflation objectives relative to that of output growth. In the last two rows of the table, the Spearman's and Kendall's rank correlation coefficients are shown.

$$\text{Welfare function: } -[\zeta\sigma^2_{\pi} + (1-\zeta)\sigma^2_{\Delta y} + \omega\sigma^2_i]$$

			$\omega/1-\zeta$ (for $\zeta=0.5$):						$\zeta/1-\zeta$ (for $\omega=0$):						
ρ	α_{π}	α_u	0	1/10	1/4	1/2	3/4	9/10	1/4	1/2	3/4	5/4	6/4	7/4	2
$\rho=0.3$	$\alpha_{\pi}=0.4$	$\alpha_u=3.5$	1	1	1	4	7	7	6	1	1	1	1	1	1
$\rho=0.8$			13	13	12	12	11	11	13	13	13	13	13	13	13
$\rho=0.6$			6	5	5	6	6	6	5	6	6	6	6	6	6
$\rho=0.5$	$\alpha_{\pi}=1.0$	$\alpha_u=2.0$	4	4	4	3	4	4	3	4	4	4	4	4	4
$\rho=0.4$			3	3	3	2	2	2	2	3	3	3	3	3	3
$\rho=0.3$			2	2	2	1	1	1	1	2	2	2	2	2	2
	$\alpha_{\pi}=0.5$		7	7	6	5	3	3	7	7	7	7	7	7	7
	$\alpha_{\pi}=1.0$		8	8	8	8	8	8	8	8	8	8	8	8	8
$\rho=0.7$	$\alpha_{\pi}=1.5$	$\alpha_u=2.0$	9	9	9	11	12	12	9	9	9	9	9	9	9
	$\alpha_{\pi}=2.0$		11	10	13	13	13	13	10	11	11	11	11	11	11
	$\alpha_{\pi}=2.5$		10	12	14	14	14	14	11	10	10	10	10	10	10
		$\alpha_u=1.0$	14	14	10	9	5	5	14	14	14	14	14	14	14
		$\alpha_u=1.5$	12	11	11	10	10	10	12	12	12	12	12	12	12
$\rho=0.7$	$\alpha_{\pi}=1.0$	$\alpha_u=2.0$	9	9	9	11	12	12	9	9	9	9	9	9	9
		$\alpha_u=2.5$	5	6	7	7	9	9	4	5	5	5	5	5	5
Spearman ρ			0.98	0.90	0.84	0.59	0.59	0.59	0.93	1.00	1.00	1.00	1.00	1.00	1.00
Kendall τ			0.93	0.76	0.63	0.43	0.43	0.43	0.87	1.00	1.00	1.00	1.00	1.00	1.00

Table 6b - Sensitivity analysis: specification of the welfare function
(partial transparency)

The table reports the rankings in terms of welfare of the competing policy rules for alternative specifications of the welfare function. The parameter ω is the weight attached to the unconditional variance of the interest rate (σ^2 while the parameter ζ and $(1-\zeta)$ measure, respectively, the relevance of inflation (σ^2_{π}) and output growth ($\sigma^2_{\Delta y}$) volatility. Each row of the table corresponds to a policy rule, while the columns refer to alternative values of the weights of the interest rate and inflation objectives relative to that of output growth. In the last two rows of the table, the Spearman's and Kendall's rank correlation coefficients are shown.

$$\text{Welfare function: } -[\zeta\sigma^2_{\pi} + (1-\zeta)\sigma^2_{\Delta y} + \omega\sigma^2_i]$$

ρ α_{π} α_u			$\omega/1-\zeta$ (for $\zeta=0.5$):						$\zeta/1-\zeta$ (for $\omega=0$):						
			0	1/10	1/4	1/2	3/4	9/10	1/4	1/2	3/4	5/4	6/4	7/4	2
$\rho=0.26$	$\alpha_{\pi}=0.7$	$\alpha_u=3.2$	1	1	1	4	6	6	6	1	1	1	1	1	1
$\rho=0.8$			12	13	11	12	11	11	13	13	13	12	12	12	12
$\rho=0.6$			6	6	5	6	5	5	5	6	6	6	6	6	6
$\rho=0.5$	$\alpha_{\pi}=1.0$	$\alpha_u=2.0$	4	4	4	3	4	4	3	4	4	4	4	4	4
$\rho=0.4$			3	3	3	2	2	2	2	3	3	3	3	3	3
$\rho=0.3$			2	2	2	1	1	1	1	2	2	2	2	2	2
	$\alpha_{\pi}=0.5$		8	7	7	5	3	3	8	8	8	8	8	8	8
	$\alpha_{\pi}=1.0$		7	8	8	8	8	8	7	7	7	7	7	7	7
$\rho=0.7$	$\alpha_{\pi}=1.5$	$\alpha_u=2.0$	9	9	9	11	12	12	9	9	9	9	9	9	9
	$\alpha_{\pi}=2.0$		10	10	13	13	13	13	10	10	10	10	10	10	10
	$\alpha_{\pi}=2.5$		11	12	14	14	14	14	11	11	11	11	11	11	11
		$\alpha_u=1.0$	14	14	12	9	9	7	14	14	14	14	14	14	14
		$\alpha_u=1.5$	13	11	10	10	10	10	12	12	12	13	13	13	13
$\rho=0.7$	$\alpha_{\pi}=1.0$	$\alpha_u=2.0$	7	8	8	8	8	8	7	7	7	7	7	7	7
		$\alpha_u=2.5$	5	5	6	7	7	9	4	5	5	5	5	5	5
Spearman ρ			0.98	0.92	0.82	0.74	0.66	0.66	0.93	1.00	1.00	1.00	1.00	1.00	1.00
Kendall τ			0.93	0.80	0.60	0.52	0.45	0.45	0.87	0.98	0.98	1.00	1.00	1.00	1.00

Table 6c - Sensitivity analysis: specification of the welfare function
(full transparency)

The table reports the rankings in terms of welfare of the competing policy rules for alternative specifications of the welfare function. The parameter ω is the weight attached to the unconditional variance of the interest rate (σ^2 while the parameter ζ and $(1-\zeta)$ measure, respectively, the relevance of inflation (σ^2_{π}) and output growth ($\sigma^2_{\Delta y}$) volatility. Each row of the table corresponds to a policy rule, while the columns refer to alternative values of the weights of the interest rate and inflation objectives relative to that of output growth. In the last two rows of the table, the Spearman's and Kendall's rank correlation coefficients are shown.

$$\text{Welfare function: } -[\zeta\sigma^2_{\pi} + (1-\zeta)\sigma^2_{\Delta y} + \omega\sigma^2_i]$$

			$\omega/1-\zeta$ (for $\zeta=0.5$):						$\zeta/1-\zeta$ (for $\omega=0$):						
ρ	α_{π}	α_u	0	1/10	1/4	1/2	3/4	9/10	1/4	1/2	3/4	5/4	6/4	7/4	2
$\rho=0.72$	$\alpha_{\pi}=1.4$	$\alpha_u=2.3$	1	3	7	8	9	9	8	3	1	1	1	1	1
$\rho=0.8$			9	7	2	2	2	2	10	10	9	9	9	9	9
$\rho=0.6$			8	10	9	10	10	10	3	7	8	8	8	8	8
$\rho=0.5$	$\alpha_{\pi}=1.0$	$\alpha_u=2.0$	11	11	12	12	12	12	7	11	11	11	11	11	11
$\rho=0.4$			12	12	13	13	13	13	11	12	12	12	12	12	12
$\rho=0.3$			13	14	14	14	14	14	13	13	13	13	13	13	13
	$\alpha_{\pi}=0.5$		6	5	6	6	7	7	1	1	5	7	7	7	7
	$\alpha_{\pi}=1.0$		4	2	3	4	4	5	2	2	3	4	5	5	5
$\rho=0.7$	$\alpha_{\pi}=1.5$	$\alpha_u=2.0$	3	1	1	3	3	4	4	5	4	3	3	3	3
	$\alpha_{\pi}=2.0$		5	4	5	5	5	6	9	6	6	5	4	4	4
	$\alpha_{\pi}=2.5$		7	6	8	7	8	8	12	8	7	6	6	6	6
		$\alpha_u=1.0$	14	13	11	9	6	3	14	14	14	14	14	14	14
		$\alpha_u=1.5$	10	8	4	1	1	1	6	9	10	10	10	10	10
$\rho=0.7$	$\alpha_{\pi}=1.0$	$\alpha_u=2.0$	4	2	3	4	4	5	2	2	3	4	5	5	5
		$\alpha_u=2.5$	2	9	10	11	11	11	5	4	2	2	2	2	2
Spearman ρ			0.83	0.55	0.36	0.24	0.10	0.10	0.59	0.90	0.99	1.00	0.99	0.99	0.99
Kendall τ			0.74	0.43	0.27	0.21	0.14	0.14	0.43	0.78	0.96	0.98	0.96	0.96	0.96

Table 7 - Sensitivity analysis: impact of initial conditions on predictor proportions

The table shows the limiting behaviour of predictor proportions under fixed and random initial conditions. The first column lists the forecasting models used for forming expectations (7 for the short-term interest rate; 5 for inflation; 2 for the exchange rate). The next columns show - for each of the 3 variables - 4 statistics: the average across replications of the share of population selecting model i under random (\mathbf{x}^r) and fixed (\mathbf{x}^f) initial conditions; the standardised difference (\mathbf{z}) between \mathbf{x}^r and \mathbf{x}^f ; the correlation (ρ) between random initial conditions and limit values of predictor proportions. The denominator of \mathbf{z} is the simple average of the standard deviations of the limit values of predictor proportions under random and fixed initial conditions.

	Short-term Interest Rate				Inflation				Exchange Rate			
	\mathbf{x}^r	\mathbf{x}^f	\mathbf{z}	ρ	\mathbf{x}^r	\mathbf{x}^f	\mathbf{z}	ρ	\mathbf{x}^r	\mathbf{x}^f	\mathbf{z}	ρ
PLM #1	0.000	0.000	0.000	-0.034	0.000	0.000	0.000	0.029	0.687	0.634	0.206	0.555
PLM #2	0.006	0.006	-0.003	0.755	0.319	0.265	0.479	0.795	0.313	0.366	-0.206	0.555
PLM #3	0.149	0.158	-0.320	0.872	0.050	0.053	-0.021	0.671	—	—	—	—
PLM #4	0.449	0.464	-0.539	0.837	0.517	0.574	-0.502	0.735	—	—	—	—
PLM #5	0.000	0.000	0.000	0.022	0.113	0.108	0.045	0.628	—	—	—	—
PLM #6	0.396	0.372	0.862	0.556	—	—	—	—	—	—	—	—
PLM #7	0.000	0.000	0.000	0.024	—	—	—	—	—	—	—	—

Chapter 4

Monetary policy uncertainty and the stock market

4.1 Introduction

The stance of monetary policy is a crucial variable for investors who have a short/medium-term perspective. In fact, the effect of inflation on future cash flows largely depend on the way monetary authorities respond. In the case of a non-accommodating policy, higher inflation induces a monetary tightening, which adversely affects future growth, producing lower cash flows and stock prices. In the case of an accommodating policy however, inflation simply indicates an accelerating cycle and signals higher cash flows and stock prices. However transparent monetary policy is, investors cannot anticipate the central bank's moves with certainty. This is due to the genuine uncertainty surrounding policy decisions¹ and to asymmetries in data availability and information processing that differentiate private investors and the central bank.

¹As Christiano *et al.* (1999) note, there is a discretionary component of monetary policymaking, unrelated to economic developments, that is driven by “exogenous shocks to the preference of the monetary authority, perhaps due to stochastic shifts in the relative weight given to unemployment and inflation. These shifts could reflect shocks to the preferences of the members of the Federal Open Market Committee, or to the weights by which their views are aggregated. A change in weights may reflect shifts in the political power of individual committee members or in the factions they represent.”.

Incomplete knowledge about the central bank's actions generates uncertainty and so affects investors' decisions. To reduce this uncertainty, capital market participants exploit any available signal to learn something about the stance of monetary policy. By giving the investors an opportunity to study the central banker's reaction, inflation is helpful in this regard. Inflation may either reinforce or disconfirm beliefs in a given monetary policy stance. If investors' expectations are confirmed, uncertainty diminishes and so do the risk premium and the required rate of return. If this reduction in stock returns is large enough, it may induce an unexpected negative correlation between inflation and returns, which may help to explain the so-called Fisher "puzzle", i.e. the empirical observation that expected inflation and stock returns fail to move together but actually move in opposite directions.

Here we explicitly link stock prices to monetary policy uncertainty, making a few original contributions. We show that uncertainty over monetary policy carries a price, which helps to explain the Fisher puzzle. We start by laying out a theoretical model that directly relates the effects of inflation on stock returns to the way in which investors learn about the stance of monetary policy. We assume monetary policy to be an exogenous process and study how investors react to the uncertainty generated by the non-observability of the policymaker's preferences. Neglecting the systematic (i.e. state-contingent) component of monetary policy is admittedly a strong simplification, but it helps us to focus on the discretionary elements that obscure the public's understanding of the central bank's strategies. We then show how this uncertainty affects stock prices. Inflation is one of the signals on which investors rely to learn the stance of monetary policy. We prove that, depending on the equilibrium beliefs of the investors on the monetary policy stance, a change in consumer prices has different effects on the risk premium: a change in inflation that confirms investors' beliefs leads to a reduction in the risk premium; one that contradicts them increases the premium.

We argue that this can explain the Fisher puzzle. That is, an increase in inflation may reduce uncertainty and therefore risk premia, thus producing a negative correlation between inflation and returns, if the signal embedded in inflation is consistent with investors' beliefs. The opposite happens if there is a deceleration in prices when investors are anticipating an accommodative monetary stance; symmetrically, a decrease in inflation can induce either a fall or a rise in excess returns depending on whether it corroborates or invalidates investors' beliefs. Realistically calibrating the model with US data, we show that the reduction in risk premia during the periods when investors beliefs are confirmed by the change in consumer prices is sufficient to generate the negative correlation.

We then test the restrictions of the model, using US data for the period 1965-1998. We construct a market-based proxy of monetary policy uncertainty and show that it is priced. Moreover, conditioning on the fundamental and monetary policy uncertainty makes the Fisher puzzle disappear. The empirical results largely support our hypothesis, showing that uncertainty about monetary policy is indeed priced.

The chapter is organized as follows. The next section briefly surveys the literature. Section 3 sets out a simple model describing the interaction between monetary policy uncertainty and stock prices. Section 4 reports the testable restrictions that the model implies. Section 5 describes the construction of a proxy for monetary policy uncertainty and Section 6 provides the main empirical findings. A brief conclusion follows.

4.2 The literature

The relationship between monetary policy uncertainty, inflation and stock returns has scarcely been studied directly. Monetary economists have analyzed monetary policy uncertainty extensively but have usually neglected to

consider its influence on asset prices; financial economists have studied the impact of expected and unexpected inflation on stock returns thoroughly but without attributing any significant role to uncertainty about the monetary stance.

With this caveat in mind, we can identify three relevant strands of literature dealing with the Fisher effect or monetary policy uncertainty or both: (1) essays on the connections between transparency and credibility on the one hand and monetary policy effectiveness on the other; (2) work testing whether investing in real assets provides protection against inflation shocks and, if not, why the Fisher effect does not hold; (3) assessments of the impact of changes in monetary policy regimes on asset prices, in particular on the correlation between inflation and stock returns.

The first strand is mostly concerned with gauging the role of transparency and credibility in the conduct of monetary policy and in determining its effectiveness (Barro and Gordon (1983), Cuckierman and Meltzer (1986), Goodfriend (1986), Kydland and Prescott (1977), Rogoff (1989), Stein (1989), Svensson (1999)). Credibility is either assumed constant (Barro and Gordon (1983), Kydland and Prescott (1977)) or modelled as a sequence of idiosyncratic unobservable shifts (Cuckierman and Meltzer (1986)). In the latter case, transparency, credibility and reputation are derived in models where the characteristics of the monetary authority are non-observable to the private sector and inferred from the policy outcome. In this context, the uncertainty over the type of monetary policy (i.e. accommodating or non-accommodating) is relevant for its impact on the effectiveness of the conduct of monetary policy itself. For example, Faust and Svensson (1997) show how a low-credibility central bank may optimally conduct a more inflationary policy than a high-credibility one. However, no direct link is explicitly formulated between monetary policy uncertainty and asset prices.

Financial economists, on the other hand, have mostly focused on dividend

uncertainty and touched upon inflation uncertainty; they have very rarely investigated the direct impact of monetary policy uncertainty on asset prices. More strikingly, even models that properly account for the relationship between inflation and stock returns are rare. In a general equilibrium context, Stulz (1986) shows that an increase in expected inflation, by reducing real wealth, induces investors to choose a portfolio of risky assets with a lower risk-return profile, which generates a negative correlation between stock returns and inflation. Thorbecke (1997) examines the effects of monetary policy innovations on stock returns and finds strong evidence that monetary policy is a common risk factor and that assets must pay positive premia to compensate for exposure to it.

The most commonly studied feature of the relationship between inflation and asset prices is the so-called "Fisher puzzle", i.e. the fact that, contrary to the conventional wisdom, common stocks do not fully insure against inflation and, actually, move in the "wrong" direction when prices accelerate (or decelerate). In particular, Fama and Schwert (1977) show a negative correlation between inflation and stock returns, real and nominal, as well as excess returns. The fact that their specification is neither structural nor reduced-form implies that their results may be due to spurious correlation, caused by the omission of relevant variables. Fama (1981) himself suggests that the anomalous return-expected inflation relation may be due to errors in the specification: his hypothesis is that anticipated changes in economic activity affect inflation and stock returns in opposite directions, so that the failure to control for expectations of future output growth generates a spurious negative correlation.² Boudoukh and Richardson (1993) challenge the

²In particular, when a proxy for future real activity is included as a regressor, expected inflation loses most of its explanatory power and becomes insignificant. A similar objection has been raised more recently by Groenewold *et al.* (1997), who show that when stock returns and macroeconomic variables are jointly modelled, the reduced-form equation for common stock returns features as regressors all the exogenous variables (government consumption, tax rates and foreign variables). When the proper theoretical restrictions are imposed and the entire structure of the model is taken into account, the sign

view that stock returns and expected inflation are negatively correlated, with evidence that over long holding periods the "Fisher puzzle" disappears.

A modified version of Fama's argument is offered by Geske and Roll (1983), who argue that the causality runs from stock returns to inflation expectations, not the reverse. Given that the government derives most of its revenues from income taxes, when stock prices go down in response to an anticipated deterioration of business conditions, so do government revenues and unless public expenditure is adjusted accordingly, fluctuations in revenues will be reflected in larger deficits, which will be financed partly by printing money. As investors foresee a future monetization of the debt, inflation expectations rise. Thus movements in stock prices caused by changes in anticipated economic conditions will be negatively correlated with changes in both expected and unexpected inflation.

Boudoukh, Richardson and Whitelaw (1994) investigate the cross-sectional implications of the Fisher relation and find that the stock returns of non-cyclical industries tend to covary positively with expected inflation, while for cyclical industries the contrary holds.

A third strand in the literature seeks to explain the relationship between stock returns and inflation by looking directly at the role of switches in monetary policy regimes. Evans and Lewis (1995), for example, find that during the post-war period there were significant shifts in inflation; although anticipated, infrequent regime changes in the inflation process induce a significant small-sample bias in the estimate of the long-run Fisher relationship, which creates the appearance of a decoupling between ex ante nominal returns and inflation.

Kaul (1990) maintains that counter-cyclical money supply is responsible for the negative correlation between inflation and asset returns; a different regime

of the effect of expected inflation is not univocally determined and depends on the specific value of the coefficients.

- pro-cyclical response of the money supply to output fluctuations - would lead to insignificant or even positive correlations. Along the same lines, Söderlind (1999) shows that if the central bank attempts to stabilize output fluctuations, nominal returns and inflation should move in parallel since the monetary authority keeps real interest rates fixed in order to cushion shocks to economic activity; conversely, if the central bank is mostly concerned with inflation, movements in nominal rates are, to a large extent, likely to reflect changes in real rates.³ However, neither of these models considers investors' efforts to discover the type of monetary policy and the uncertainty generated by the learning process.

In general, the three strands of the literature have not investigated the way investors react to the informational uncertainty concerning the stance of monetary policy and are therefore unable to account for the way financial markets use inflation as a signal to price the uncertainty resulting from the type of monetary regime. Our contribution is to provide a framework that models and links uncertainty over fundamentals with that over monetary policy: unlike Fama (1981), Boudoukh, Richardson and Whitelaw (1994) or, less markedly, Kaul (1990) and Söderlind (1999), we do not view the omission of future growth prospects as the main cause of the Fisher puzzle; unlike Evans and Lewis (1995) and Geske and Roll (1983), we do not attribute the negative correlation between asset returns and inflation to the use of unsophisticated econometric techniques; we disagree with Boudoukh and Richardson (1993) that confining the Fisher puzzle to short-run holding periods is an acceptable solution. Our main objective is to show that both in theory and in practice uncertainty about monetary policy is the missing risk factor that explains why stocks apparently fail to insure against inflation.

³Thorbecke (1997) seeks to measure the influence of monetary policy on stock returns, finding that though such influence is quite sizeable, it is not enough to explain the predictive power of other financial variables (e.g. dividend yield) or to account for the predictable volatility in excess returns.

4.3 A simple model

The literature in general agrees that, while the actions of the monetary policymaker are rather explicit, their interpretation is not immediate. Indeed, any action has to be considered in terms of macroeconomic developments and depends on the overall monetary objectives. Investors do not have access to the same information set as the monetary authority and they are uncertain about the objectives of the central bank. We focus on the process through which investors seek to judge the monetary policy stance from inflation outcomes.

We do not try to describe the behavior of the monetary authority. Instead we assume that monetary policy is an exogenous process and examine how asset prices are affected by the reaction of investors to the uncertainty generated by the non-observability of the central bank's preferences. The model, which builds on David (1987), Veronesi (1999) and David and Veronesi (2000), posits an exchange economy with a single consumption good and a single risky asset.

The model and the empirical estimation are based on a Markov switching framework. As observed by Sims (1982) and Cooley et al. (1984), it is doubtful at best to characterize changes in the policy framework as permanent changes in the parameters of a reaction function. In fact, genuine changes of regime are rare indeed, since agents, knowing the menu of choices available to the policymakers, form their expectations on the basis of past experience, accounting for all possible outcomes: they have a probability distribution ranging over all possible policy rules and use it to forecast the behavior of the policymakers. From this perspective, "regime" changes are neither rare events nor abnormal policy shifts but should be seen as variations in central bank policy that cannot be accounted for as a reaction to the state of the economy. These changes might reflect *exogenous shocks to the policy-*

maker's preferences, due perhaps to stochastic shifts in the relative weights assigned to the various policy objectives, or may be induced by changes in the composition of the monetary policy committee. Alternatively, they might be caused by incomplete non-observability of the state of the economy at the time decisions are to be taken (Christiano *et al.*,1999).

In this framework, a Markov switching model is versatile enough to cover both once-and-for-all structural changes and policies that are set period-by-period. Any intermediate case can be obtained by an appropriate choice of the parameters of the transition matrix.

4.3.1 Monetary policy regimes

Monetary policy is represented by the exogenous Poisson process θ_t , which can take two values: a and b . When $\theta_t = a > 0$, monetary policy is accommodating; when $\theta_t = b < 0$, it is restrictive. The matrix of transition probabilities between time t and time $t + dt$ is:

$$\begin{array}{c|cc}
 & \mathbf{a} & \mathbf{b} \\
 \hline
 \mathbf{a} & 1 - \lambda dt & \lambda dt \\
 \mathbf{b} & \lambda dt & 1 - \lambda dt
 \end{array} \tag{4.1}$$

For simplicity, we assume that $a = 1$ and $b = -1$. This framework is meant to capture the fact that while the actions of the monetary authority are rather explicit, their interpretation is not immediate or direct. In fact, investors do not have access to the information set of the central bank and do not know the monetary objectives in detail, so they have to infer them from the information and signals available (i.e., consumer and stock prices).

Given that our focus is the reaction of investors to monetary policy uncertainty, the description of the behavior of the central bank is necessarily highly stylized. One possible objection to our approach is that there is really nothing in the model that differentiates shifts in the monetary stance from

movements in payoff-relevant variables. This issue is addressed below.

4.3.2 Investors' behavior and stock returns

The model features a representative investor, endowed with a power utility function $u(t, C_t) = e^{-\phi t} \frac{C_t^{1-\rho} - 1}{1-\rho}$, where C_t is the consumption level, ρ is the degree of risk aversion and ϕ is the discount rate; the investor can invest in a risky asset (stock) and in a safe bond, which delivers a nominal riskless return equal to r_t^n . Uncertainty is represented by the vector of Wiener processes $\mathbf{z}_t = (z_{D,t}, z_{p,t}, z_{e,t})'$ that collects the shocks driving the time path of real dividends, inflation and the endowment process. By assumption, the elements of the vector \mathbf{z}_t are pairwise uncorrelated. The laws of motion of real dividends (D_t), the price level (p_t) and consumption (C_t) are given by the following set of Brownian motions:

$$\begin{aligned} \frac{dD_t}{D_t} &= [\mu_D + \beta\theta_t] dt + \mathbf{b}_D d\mathbf{z}_t \\ \frac{dp_t}{p_t} &= [\mu_p + \delta\theta_t] dt + \mathbf{b}_p d\mathbf{z}_t \\ \frac{dC_t}{C_t} &= [\mu_c + \gamma\theta_t] dt + \mathbf{b}_c d\mathbf{z}_t \end{aligned} \tag{4.2}$$

The systematic component of the growth rate of dividends is equal to the sum of the drift term μ_D , which captures productivity developments, and β , which is positive by assumption and measures the sensitivity of earnings to monetary policy; stochastic fluctuations around the mean growth rate are driven by the term $\mathbf{b}_D d\mathbf{z}_t$, where $\mathbf{b}_D = (\sigma_D, 0, 0)$. Accordingly, in the short run a restrictive monetary stance reduces the real growth of the economy ($\theta_t < 0$), while an expansionary stance ($\theta_t > 0$) raises it; in the long run, monetary policy is neutral, as $E\theta_t = 0$.

The specification of the law of motion of the price level of the single consumption good in the economy (p_t) follows Stulz (1986). μ_p is the long-run inflation level and δ , a positive parameter, measures the sensitivity of the price level to the stance of monetary policy. We assume it to be positive. β

and δ characterize the type of economy and define how costly it is for the central bank to tame inflation: if β is high and δ is low, the deflationary impact of curbing inflation is substantial, while if β is low and δ is high, the central bank can reduce inflation without major output losses. As for the stochastic component, we assume that $\mathbf{b}_p = (0, \sigma_p, 0)$.

Our modelling of the transmission of monetary policy impulses is consistent with the prototype New Keynesian model (Clarida, Gali and Gertler, 1999). Also, notice that although monetary policy affects dividend growth and inflation, the model does not imply that it is the sole cause of booms and recessions.

The investor has two sources of income: dividends and endowment (Berk, Green and Naik, 1999, Cecchetti, Lam and Mark, 1993, Campbell and Cochrane, 1999, and Barberis, Huang and Santos, 2001). We define their aggregate value as total real consumption (David and Veronesi, 2000). The law of motion of consumption is represented by the third equation in (4.2), where μ_c is long-term mean consumption growth, γ is the impact of monetary policy on it, and $\mathbf{b}_c = (\sigma_{cD}, 0, \sigma_e)$. We also assume that the impact of monetary policy on consumption differs from its impact on dividends and that $\gamma \leq \beta$. Three signals are available to the investors (D_t , p_t and C_t) and there are four sources of uncertainty, which can be used to learn about the value of θ_t .⁴

Theorem 4.3.1. *The expected value of the nominal excess stock return over the riskless asset (R_t) is:*

$$E\left[\frac{dS_t^n}{S_t^n} + \frac{D_t^n}{S_t^n} - r_t^n dt\right]/dt = \Gamma_t + \Theta_t + \Phi_t = F_t + \Phi_t \quad (4.3)$$

where S_t^n is the stock price, $\Gamma_t = \mu_D - \beta(1 - 2\pi_{a,t}) - \phi - \rho\mu_c + \frac{1}{2}\rho(\rho+1)\mathbf{b}_c\mathbf{b}_c'$, $\Theta_t = (1 - 2\pi_{a,t}) \left[\frac{2(\beta-\rho\gamma)\lambda}{\Psi_t} + \gamma\rho \right]$, $\Phi_t = \pi_{a,t}(1 - \pi_{a,t}) \frac{4(\beta-\rho\gamma)}{\Psi_t} (\beta\Omega_t + \delta - \gamma)$, $\pi_{a,t}$

⁴Given the assumption of a single representative agent, the observation of dividends provides the same information as the observation of asset prices or returns (i.e., stock and bond rates).

is the probability that the monetary stance is accommodating and the values of Ω_t and Ψ_t are defined in the Appendix (proof in the Appendix).

Γ_t consists of the conditional expected real dividend return ($E_t[\frac{dD_t}{D_t}]/dt = \mu_D - \beta(1 - 2\pi_{a,t})$) and some adjustment for consumption growth and risk. Θ_t represents the risk of inflation and in particular the risk that in order to curb inflation the central bank will curb output growth. These first two terms represent the part of the risk due to both the real sources of income (dividend and endowment or consumption) and inflation. We aggregate them and define $F_t = \Gamma_t + \Theta_t$ as the “fundamental” component of the risk premium. Finally, the term Φ_t captures monetary policy uncertainty: it is highest at $\pi_{a,t} = 0.5$ and lowest at $\pi_{a,t} = 0$ or $\pi_{a,t} = 1$.

Equation (4.3) provides two key insights: one on pricing and the other on the Fisher puzzle. It also gives us a way of directly quantifying the degree of monetary policy uncertainty. In the next section we examine these implications separately. First, however, we define expected inflation:

$$E \left[\frac{dp_t}{p_t} \right] / dt = [\mu_p + (2\pi_{a,t} - 1)\delta] \quad (4.4)$$

This equation shows that there is a direct mapping between expected inflation ($E_t[dp_t/p_t]$) and the beliefs on the stance of monetary policy ($\pi_{a,t}$). This implies that we can use one as a direct proxy for the other.⁵

4.4 The role of monetary policy uncertainty

4.4.1 Pricing implications

Equation (4.3) shows that the risk premium can be broken down into a component due to “fundamental uncertainty” (F_t) and one due to “monetary

⁵In particular, from a simple application of the Implicit Function Theorem, we can see that there is a direct relationship between the impact of expected inflation on stock returns and the impact of a change in the beliefs about the type of monetary regime (π_a) on stock returns. That is, $E_t \left\{ \frac{dR}{dE[\frac{dp}{p}]} = \frac{1}{2\delta} \frac{dR}{d\pi_a} \right\}$.

policy uncertainty” (Φ_t). The former is a function of the uncertainty over the real fundamentals (dividend and endowment) and inflation. The latter (Φ_t) is a function of investors’ beliefs about the monetary policy stance ($\pi_{a,t}$) or, according to equation (4.4), about inflation.

A graph of the relationship between risk premia and investors’ beliefs is given in Figures 1 and 2. We graph the risk premium for different values of β , δ , $\pi_{a,t}$ and the percentage share of dividends out of aggregate consumption. The graphs are based on parameters calibrated to the data for the period January 1965 - December 1998. Details on this calibration are provided in the Appendix. The risk premium has a hump-shaped relationship with respect to beliefs: when $\pi_{a,t} < 0.5$, i.e. when the market is confident that the monetary policy is restrictive, an increase in inflation disconfirms expectations and raises uncertainty and the risk premium; conversely, when $\pi_{a,t} > 0.5$, i.e. when the market is confident that the stance is accommodating, an increase in inflation reduces uncertainty and hence the risk premium. An acceleration in price dynamics confirms the priors of an investor who believes the central bank is accommodative and therefore reduces the uncertainty about central bank’s actions.

Exhibit 1: Effects of expected inflation on risk premia

	Perceived Monetary Policy Regime	
	Non-accommodating	Accommodating
	$(\pi_a < 0.5)$	$(\pi_a > 0.5)$
Expected inflation increases	Positive	Negative
Expected inflation decreases	Negative	Positive

Therefore, an increase in inflation when investors believe that the monetary policy is tight (first quadrant in Exhibit 1) or a reduction when they believe it to be expansionary (fourth quadrant) will increase the risk premium. On the contrary, an increase of inflation when investors believe monetary policy is accommodating (second quadrant) or a reduction when they believe it to be

restrictive (third quadrant) will reduce the risk premium. The points on the diagonal represent outcomes where a change in expected prices ($Infl_t^e$) moves against investors' prior beliefs; conversely, the off-diagonal points represent outcomes where a change in expected prices is aligned with the priors.

The intuition is that a change in consumer prices gives investors an opportunity to study how the central bank reacts and infer the monetary policy stance. Low inflation can reinforce beliefs in a non-accommodating policy where they exist, as well as disprove beliefs in an accommodating stance. If the signal reinforces investors' beliefs, it reduces the risk premium. In particular, both an increase in inflation when investors believe that monetary policy is tight and a reduction when they believe it is expansionary will increase the risk premium. Vice-versa, an increase in inflation when investors believe that the central banker is an inflation dove and a reduction when the central banker is believed to be a hawk will reduce the risk premium.

Equation (4.3) provides a directly testable restriction. If we use the pricing kernel representation, excess stock returns can be expressed as:

$$E[R_{jt}|\Omega_{t-1}] = \sum_{i=1}^L \lambda_{i,t-1} cov[R_{jt}, \Upsilon_{it}|\Omega_{t-1}] \quad (4.5)$$

where R_{jt} is the excess rate of return on the j^{th} stock at time t , λ_i is the price of risk, Υ_{it} is the return on the portfolio that proxies for the i^{th} factor of the economy, for $i = 1, \dots, L$, and Ω_{t-1} is investors' information set. We consider two cases. In the unrestricted case, there are 4 factors (i.e., $L = 4$): the 3 Fama and French factors (FF henceforth), which proxy for the fundamentals (i.e., F_t), and the factor proxying for monetary policy uncertainty (i.e., Φ_t). In the restricted case, we consider only the 3 FF factors.

If we define M_t as the marginal rate of substitution between returns at t and at $t - 1$, the first order conditions of the portfolio choice problem can be

expressed as:

$$E[M_t(1 + r_t^n)|\Omega_{t-1}] = 1 \text{ and } E[M_t R_{jt}|\Omega_{t-1}] = 0 \quad (4.6)$$

where r_t^n is the conditionally riskless rate of interest at time $t-1$. By substituting equation (4.5) into equation (4.6), we have:

$$M_t = [1 - \lambda_{o,t-1} - \sum_{i=1}^L \lambda_{it-1} \Upsilon_{it}]/(1 + r_t^n). \quad (4.7)$$

Equation (4.6) allows us to proceed to define the orthogonality conditions of the econometric specification. We use a set of instrumental variables \mathbf{Z}_{t-1} , in order to proxy for the information set Ω_{t-1} .⁶ We assume that the price of risk (λ) is linearly related to such state variables, according to:

$$\lambda_{o,t-1} = -\mathbf{Z}_{t-1}\delta \text{ and } \lambda_{it,t-1} = \mathbf{Z}_{t-1}\phi_i, \quad (4.8)$$

where δ and $\phi_1, \phi_2, \dots, \phi_L$ are time-invariant vectors of weights. If we define the innovation u_t in the marginal rate of substitution as:

$$u_t = 1 - M_t(1 + r_t^n) \quad (4.9)$$

or, using equation (4.7):

$$u_t = -\mathbf{Z}_{t-1}\delta + \sum_{i=1}^L \mathbf{Z}_{t-1}\phi_i \Upsilon_{it} \quad (4.10)$$

and we define $h_{jt} = R_{jt} - R_{jt}u_t$, we can rewrite the pricing conditions of equation (4.6) in terms of the orthogonality conditions:

$$E[u_t|\Omega_{t-1}] = 0 \quad (4.11)$$

$$E[\mathbf{h}_t|\Omega_{t-1}] = 0 \quad (4.12)$$

where \mathbf{h}_t is the vector that contains the h_{jt} stacked for all the considered portfolios. In particular, if we define the vector of residuals $\varepsilon_t = (u_t, \mathbf{h}_t)$, we

⁶The vector \mathbf{Z}_{t-1} contains the state variables on which investors condition their portfolio decision.

can rewrite the system of equations (4.11) and (4.12) as:

$$E[\varepsilon_t | \mathbf{Z}_{t-1}] = 0 \quad (4.13)$$

or, considering its sample version:

$$\mathbf{Z}'\varepsilon = 0 \quad (4.14)$$

where Z is the $T \times Q$ matrix containing the sample values of the instruments, ε the $T \times (1 + m)$ matrix containing the residuals, T the sample size, Q the number of instruments and m the number of portfolios used to test the model (25 in the book-to-market and size specification and 17 in the industry specification).

Equation (4.14) allows us to test whether the model is correct and whether the additional factor we are considering is relevant and priced. If the factors are priced, equations (4.11) and (4.12) must hold and the quadratic form of equation (4.14) must be asymptotically distributed as a χ^2 .

Also, from equations (4.7) and (4.8), we can derive the pricing power of the monetary policy factor by looking directly at the significance of the vector of coefficients ϕ_{Φ_t} , where $\lambda_{\Phi_t, t-1} = \mathbf{Z}_{t-1}\phi_{\Phi_t}$, for monetary policy uncertainty: if monetary policy uncertainty is priced, at least some of the coefficients of ϕ_{Φ_t} should be significant. When the risk factors are correlated, testing for the significance of $\lambda_{\Phi_t, t-1}$ (ϕ_{Φ_t}) is the appropriate strategy. In fact, the λ s are the multiple regression coefficients of M_t on $[F_t, \Phi_t]$ and capture whether one factor is marginally useful in pricing assets, given the presence of the other factors (Cochrane, 2000).

4.4.2 New Perspective on the Fisher Puzzle

Equation (4.3) has direct implications for the Fisher puzzle. While the Fisher relationship requires that:

$$Corr [R_{jt}, Infl_t^e] = 0, \quad \forall j \quad (4.15)$$

the “Fisher puzzle” is the empirical finding that:

$$\text{Corr} [R_{jt}, \text{Infl}_t^e] < 0 \quad (4.16)$$

On the basis of our working hypothesis, this negative correlation can now be explained in terms of the monetary policy risk premium. *If a change in consumer price dynamics reduces the portion of the risk premium due to monetary policy uncertainty substantially enough, the net effect can be a negative correlation between inflation and stocks returns.*

Unconditionally, the sign of the correlation between expected inflation (Infl_t^e) and risk premium depends on the relative frequency of periods when monetary policy is perceived as non-accommodating and on the size of the component of the risk premium due to policy uncertainty relative to the component due to fundamental uncertainty. The correlation between an increase in consumer prices and risk premia is positive in periods of non-accommodating monetary policy, negative during accommodating periods. The fact that monetary policy has been perceived as accommodating for most of our sample period (Figure 3) may justify the negative relationship found in the literature. This is also consistent with the fact that on a longer sample the Fisher relationship appears to hold (Boudoukh and Richardson (1993)).

Conditionally, if we control properly for fundamental uncertainty (F_t) and monetary policy uncertainty (Φ_t), we should be able to get back to the expected zero correlation between inflation and excess stock returns. A directly testable restriction is therefore:

$$\text{Corr} [R_{jt}, \text{Infl}_t^e \mid F_t, \Phi_t] = 0 \quad (4.17)$$

As a strategy, we accordingly verify whether unconditionally we find the Fisher puzzle in our sample and whether, after conditioning on fundamentals and policy uncertainty, the puzzle disappears.

4.5 A proxy for monetary policy uncertainty

Equation (4.3) also suggests a way to construct a measure of monetary policy uncertainty: $I = \pi_a(1 - \pi_a)$. This indicator peaks ($I_t = 0.25$) when investors are highly uncertain about the monetary policy regime (i.e. $\pi_{a,t} = 0.5$), and is lowest ($I_t = 0$) when investors have a well-defined belief (i.e., $\pi_{a,t} = 0$ or $\pi_{a,t} = 1$). To apply this indicator, we need to find suitable proxies for π_a , e.g. using survey measures of inflation expectations or exploiting the information content of financial data. The former approach is limited by the lack of a good proxy. One possibility is the dispersion of the forecasts contained in the surveys of professional forecasters and economists. For example, the Livingston Survey, the ASA-NBER Survey of Professional Forecasters and the University of Michigan Inflation Expectations survey contain information about expected inflation. The former two also have some information on the dispersion of forecasts among analysts. However, the Livingston Survey is relatively infrequent (semi-annual), while the ASA-NBER Survey has a variable number of responses over time and covers quite a short period.⁷ Furthermore, in both cases the latent variable about which investors are uncertain is not the monetary policy regime but inflation itself.

A direct quantification of the monetary policy stance is provided by the indexes of Boschen-Mills (1995) and Bernanke-Mihov (1998). These have been constructed using standard monetary variables and provide the best identification of the monetary policy regime. Unfortunately, they are based on data not immediately available to the market. Nor do they offer any way of quantifying the degree of uncertainty. We therefore need to construct a measure that allows us to quantify both the market's monetary policy expectations and the uncertainty over them.

Accordingly, we go for the direct approach. We first estimate investors' beliefs

⁷See in particular, Dean Croushore, Introducing: the Survey of Professional Forecasters, *Business Review*, Nov./Dec. 1993.

about the type of monetary regime and their volatility and then construct the index of monetary policy uncertainty. We assume that policy changes represent a latent stochastic process which can be estimated jointly with the parameters of the model. This accounts properly for the instability of the regimes, avoids arbitrary assumptions and provides a direct measure of the regimes *as they are perceived by the market*.

4.5.1 Monetary policy regimes

To identify the monetary policy regimes, we use a two-state Markov switching VAR. The VAR includes five variables: the excess return on the market portfolio, the corresponding dividend yield, the risk-free (real) rate, consumer price inflation (CPI) and the real GDP growth rate. The non-financial data is provided by the Federal Reserve. The model is estimated by maximum likelihood (ML), using quarterly data from 1965:3 to 1998:4. An EM algorithm is used to estimate the vector of parameters and the hidden Markov process (see Appendix E). Unlike the theoretical model, the unconditional probabilities of the two monetary policy regimes are not restricted to be the same.

The probability of the non-accommodating regime ($\pi_{b,t}$) is shown in Figure 3, which also presents a proxy of monetary policy uncertainty. To identify the regimes, we look at the correlations with standard indexes of monetary policy. The correlation is very high for both the Boschen-Mills index (38.3%) and the Bernanke-Mihov measure (48.1%), which is strong evidence of the quality of our estimates. As further evidence, we consider the correlation with the Romer and Romer (1989) index of monetary policy, based on the “minutes” of the FOMC meetings. For the period for which both indicators are available, the four episodes of restrictive policy (December 1968, April 1974, August 1978 and October 1979) identified by $\pi_{b,t}$ coincide with those identified by the Romer and Romer index. This strongly supports our identification, as

the unconditional probability of the non-accommodating regime is just 0.31 and the likelihood that this result is due to chance is well below 1%.

It is nevertheless possible that these correlations may be due to spurious effects if, for instance, the Markov process simply captures cyclical movements in the economy. To address this issue, we regress $\pi_{b,t}$ on alternative measures of the monetary policy stance and on a proxy for business cycle fluctuations. We estimate the model:

$$\pi_{b,t} = \alpha + \beta M_t + \gamma BC_t + \varepsilon_t \quad (4.18)$$

where M_t and BC_t are, respectively, the index of monetary policy stance (either Bernanke-Mihov or Boschen-Mills one) and a proxy for business cycle fluctuations. Regarding the latter, indicators mostly based on the NBER Business Cycle Reference Dates are used: (i) the series for cyclical peaks, (ii) the series for cyclical troughs; (iii) the difference between peaks and troughs; (iv) the sequence of NBER turning points linearly interpolated; (v) James Stock's coincident, leading and recession indexes (and some transformations of them).⁸ If the correlation between our indicator and the exogenous measures of the monetary policy stance were due to spurious correlation, we would expect it to disappear when we control for business cycles.

In Table 1 we report the results for the case when the Bernanke-Mihov index is used.⁹ The results are quite clear-cut: the β coefficient exhibits a high t -statistic and has the expected sign in all the specifications. Furthermore, the statistics for γ are in general less significant than the corresponding statistics for β .¹⁰

Finally, direct inspection of the index (Figure 3) suggests that the regimes

⁸A detailed description of these indicators is reported in the Appendix. We thank J. Eberly and R. Jaganathan for suggesting some of them.

⁹The analysis based on the Boschen-Mill Index provides analogous results. However, these results are less reliable, as the Boschen-Mill Index is a discrete variable with limited range of variation.

¹⁰The residual correlation with the cyclical indicators are due to the inherent link between monetary policy and business cycle.

are identified correctly. The accommodating regime is more frequent from 1965 to the 1970s, with the exception of 1966-1967. This is consistent with the standard explanation that “money growth accelerated ... and persisted through the 1970s. US inflation began to accelerate in 1964, with a pause in 1966-1967, and was not curbed until 1980” (Bordo and Schwartz, 1999). It was only in the second part of the 1970s that monetary policy gradually became tighter with a change in operating procedures.¹¹ Tightening can be clearly identified during the Volcker era. The accommodating regime seems to prevail also at the end of the sample, when stable and low inflation allowed the Fed to follow a less tight policy and to accommodate the “irrational exuberance” of the stock market.

4.5.2 A market-based measure of monetary policy uncertainty

Once the perceived regimes of monetary policy have been estimated, the construction of a measure of policy uncertainty is immediate. The obvious choice is to use the index $I_t = \pi_{a,t}(1 - \pi_{a,t})$, where $\pi_{a,t}$ is the probability that investors attach to monetary policy being accommodating, as estimated from the Markov-switching VAR model. From Figure 3 it is easy to see that higher volatility and lower persistence are more common in expansionary regimes. The next step is to construct a tracking portfolio that mimics I_t , which provides a direct measure of the possibility for investors to hedge the risk due to monetary policy uncertainty. The construction of the tracking portfolio shows whether this particular source of uncertainty is actually hedgeable (Campbell, Lo and MacKinlay, 1997, Cochrane, 2000) and if hedgeability

¹¹In the early 1970s, the Fed gradually abandoned indirect targeting in favor of direct targeting of the federal funds rate, allowing movements only within a narrow band (usually 25 basis point), specified by the FOMC each time it met. In 1975, the Fed started to adopt and announce one-year money growth targets, in application of Congressional Resolution 133. In 1979, targeting of non-borrowed reserves in place of direct federal funds rate targeting was adopted. Finally, in October 1982 the FOMC decided to abandon non-borrowed reserves targeting in favor of managing borrowed reserves.

changes over time. However, it is worth stressing that *the results do not depend upon it*. Indeed, replicating the analysis on the index of policy uncertainty without the tracking portfolio yields essentially the same results.

We use the technique developed by Lamont (2000) and Vassallou (2001), whose underlying intuition is that innovations in returns reflect changes in expectations of future cash flows and discount rates. Provided that market expectations are properly accounted for, portfolios whose innovations have a strong correlation with revisions in expectations about fundamentals can be used to explain the cross-sections of asset returns.

We therefore regress I_t , viewed as a measure of information uncertainty, on a set of portfolios, the so-called base assets, and on instrumental variables, which summarize the information available to market participants and proxy for their expectations. The regression model is:

$$I_{t+k} = \alpha + \theta \mathbf{B}_t + \zeta \mathbf{Z}_{t-1} + \varepsilon_{t+k} \quad (4.19)$$

where \mathbf{B}_t represents the set of base assets,¹² \mathbf{Z}_t is the set of instruments and I_{t+k} is the realized future value of uncertainty.¹³ Given that the tracking portfolios are generated regressors, we also compute the White-corrected t -statistics. Table 2 reports the estimates of equation (4.19). Note that the R^2 , while lower than in Lamont and Vassallou, is still comfortably high, especially considering that the factors we are trying to mimic are not so persistent or

¹²The base assets are *term*, *junk* and *ME1/ME5*. “Term” stands for the spread between the yield on 10-year Treasury bonds and 3-month Treasury bills; “junk” is the difference between Moody’s Baa and Aaa corporate yields; “ME1/ME5” is the return on an arbitrage portfolio that is long on stocks of small firms (first NYSE market equity quintile) and short on stocks of large firms (fifth NYSE market equity quintile). Their returns are in excess of the riskless rate.

The instrumental variables are expected inflation, actual inflation, producer price inflation, the excess return on the market portfolio, and the share of durables in total consumption.

¹³The tracking portfolios are selected so as to maximize the fit of equation (4.19). In order to assess whether the results are robust, a few other tracking portfolios, which combine portfolios constructed on the basis of book-to-market categorization and portfolios constructed on the basis of industrial categorization (2-digit SIC codes), have been estimated and used as an alternative. No notable differences were found, apart from a worse fit of the regression model (4.19).

predictable.

The tracking portfolio $\theta \mathbf{B}_t$ is our measure of monetary policy uncertainty Φ_t . To proxy for fundamental uncertainty we rely on recent studies (Liew and Vassallou, 2001 and Vassallou, 2000) that show that the FF factors are good proxies of news on future GDP growth and are correlated with innovations on other macroeconomic fundamentals. We consider both a CAPM model and a specification based on the three FF factors. In the former case we proxy fundamental uncertainty with the excess return on the market portfolio, in the latter we use the three FF factors, i.e., Market, SMB and HML.¹⁴

4.6 Evidence on the market price of policy uncertainty

We proceed as follows. First, we estimate whether the Fisher puzzle shows up also in our sample (restriction (4.16)). Then we focus on the pricing relationship (4.3) and assess whether monetary uncertainty affects stock returns and is therefore priced. Finally, relying on this pricing relationship, we go back to the Fisher puzzle, compare restrictions (4.15) and (4.17) and see whether, by properly conditioning on the uncertainty factors, the puzzle is resolved.

4.6.1 Inflation and stock returns: the Fisher puzzle

We start by testing whether the Fisher puzzle actually exists in our sample using the 25 size and book-to-market portfolios and 17 industry portfolios. Descriptive statistics of the sample data are provided in Table 3, Panel A

¹⁴“Market” is the excess return of the aggregate market portfolio over the riskless rate, book-to-market (“HML”) is the difference between the average returns on the two portfolios with high book-to-market ratios and the average returns on the two portfolios with low book-to-market ratios. Size (“SMB”) is the difference between the average returns on three small stock portfolios and the average returns on the three big stock portfolios.

and B. We estimate the following equation (4.16):

$$R_{i,t} = \alpha + \beta_i Infl_t^e + \varepsilon_{i,t} \quad (4.20)$$

where $Infl_t^e$ is inflation expected on the basis of information up to time $t - 1$ and $R_{i,t}$ represents the excess returns of the i -th portfolio. Expected inflation is constructed by using the one-period-ahead forecast based on the Markov-switching VAR specification allowing for regime shifts, i.e.: $Infl_t^e = \pi_{a,t}Infl_{a,t}^e + \pi_{b,t}Infl_{b,t}^e$, where $\pi_{a,t}$ and $\pi_{b,t} = 1 - \pi_{a,t}$ are the probabilities (conditional upon information as of time $t - 1$) of the two monetary policy regimes and $Infl_{a,t}^e$ and $Infl_{b,t}^e$ are expectations as of time $t - 1$ of time- t inflation in the first and second regime respectively.¹⁵

Table 4 reports the results for the size and book-to-market portfolios (Panel A) and for the industry portfolios (Panel B). In all cases and for all portfolios there is a highly significant negative correlation between expected inflation and excess returns, confirming that the “Fisher puzzle” holds regardless of sample period or the criterion for forming portfolios. Expected inflation affects risk premia and the relationship is negative. Let us now explain why.

4.6.2 Evidence of Pricing

To assess whether monetary policy uncertainty is priced, we estimate equation (4.14). The additional factor (Φ_t) is measured by the tracking portfolio described above. In the unrestricted case, the vector of factors is $\tilde{\mathbf{F}}_t = [\mathbf{F}_t, \Phi_t]'$. Data on the FF factors is derived from Kenneth French’s website. The information variables are those used in the literature (Ferson and Harvey, 1991, 1993 and 1999, Dumas and Solnik, 1995): a constant, a January dummy, the one-month T-Bill yield (T -bill), the dividend yield of the S&P 500 index (div), the term premium, i.e. the spread between a 10-year and 1-

¹⁵The correlation between $Infl_t^e$ - as estimated by our Markov-switching VAR - and University of Michigan Inflation Expectations is 0.93.

year Treasury bond yield (*term*), the junk premium, i.e. the spread between Moody's Baa and Aaa corporate bond yields (*junk*), the spread between the one-month return on a three-month T-bill and the return on a one-month T-bill (*hb3*).

The tests are based on the GMM estimation of the $Q \times (1 + m)$ orthogonality conditions described in equation (4.14). We consider both industry portfolios and size and book-to-market portfolios. In addition to the specification using \mathbf{F}_t and Φ_t as risk factors, we also estimate two alternative models: one in which the factor that proxies for policy uncertainty is orthogonalized with respect to the FF factors (Panels C and D) and one in which the FF factors are orthogonalized with respect to Φ_t (Panels E and F).¹⁶ Though the FF factors are expected to proxy for fundamental uncertainty, they may also be related to monetary policy and could capture part of the effects of monetary policy uncertainty. Estimating the two alternative models allows us to better isolate the component of the risk premium that depends on monetary policy uncertainty, removing from the FF factors the part that is not due to fundamentals.

We report the values of the estimated coefficients ($\phi_f s$) in Table 5: Panels A, C and E for the size portfolios and Panels B, D and F for the industry portfolios. The last row of each panel reports the χ^2 of the model and the associated *p-value*. The results show that we cannot reject the null at any confidence level, as the *p-value* is in general close to 0.99. This is evidence for our working hypothesis that monetary policy uncertainty is priced, since the analysis shows that most of the coefficients contained in ϕ_{Φ_t} are highly significant. In particular, the price of risk of monetary policy uncertainty is related to *div*, *junk*, *term*, *hb3*, the January dummy and the constant in the case of the book-to-market and size portfolios and to *div*, *junk*, *term*, the

¹⁶The orthogonalization of the three risk factors is obtained by taking the residuals of the regression of each of them on the tracking portfolio for I_t .

January Dummy and the constant in the case of the book-to-market and size portfolios. The ϕ_{Φ_t} are different from zero at any significance level.

4.6.3 Inflation and stock returns: the Fisher relationship revived

We can now verify how uncertainty about monetary policy regimes affects the relationship between stock returns and inflation. In particular, we want to check whether the joint use of fundamental uncertainty (F_t) and monetary policy uncertainty (Φ_t), as specified in equation (4.3), can explain the Fisher puzzle (restriction (4.17)). Accordingly, we estimate the following equation:

$$R_{i,t} = \alpha + \delta'_i \mathbf{F}_t + \gamma_i \Phi_t + \beta_i \text{Infl}_t^e + \varepsilon_{i,t} \quad (4.21)$$

Restrictions (4.15) and (4.17) require that $\gamma_i > 0$ and $\beta_i = 0$. That is, if information uncertainty is priced, the relationship between expected inflation and returns should be non-negative. For the more uncertain investors are over future monetary policy actions, the higher the risk premium should be.

Table 6 reports the results for the three-factor specification ¹⁷. The results strongly support our working hypothesis. The coefficient on monetary policy uncertainty is positive and highly significant, both in the case of book-to-market and size portfolios and in the case of industry portfolios (i.e., $\gamma_i > 0$). The estimated coefficient remains highly significant even after the application of the White correction to control for the problem of generated regressors. In particular, the tracking portfolios are positive and strongly significant in 24 out of 25 cases for the size and book-to-market portfolios and in 13 out of 17 cases for the industry portfolios.

The impact of policy uncertainty would appear to be negatively related to the size of the company: the larger the company, the less the impact. This effect is approximately monotonic and regards both the value of the coefficients and their statistical significance: the value of γ increases from a minimum of

¹⁷The one-factor specification (CAPM) gives the same results.

-0.06 for the largest companies to a maximum of 0.82 for the smallest, with the respective *t-statistics* rising from -1.09 to 21.30. This seems to suggest that the largest companies are more capable of hedging uncertainty, either because they operate in many sectors (industrial diversification) or because they are financially more sophisticated (financial diversification).

In the case of industry portfolios, those that do not appear to be affected by monetary policy uncertainty are mostly those concentrating in such sectors as Utilities, Oil, Consumer Products and Financial Services. In the case of Utilities and Oil stocks, this suggests a sort of “built in hedge” against monetary policy shocks, in that the periods when monetary policy is tighter are presumably those with higher inflation and often coincide with periods when oil- and energy-related stocks fare better. In the case of financial services, the lack of correlation between monetary uncertainty and the financial sector confirms that the returns on financial stocks, banks stocks in particular, are positively affected by tight monetary policy.¹⁸

In order to assess the robustness of the results and in line with the findings of Brennan, Chordia and Subrahmanyam (1998), we also estimate equation (4.21) for portfolios grouped according to three alternative criteria: (i) the ratio between cash flow and market price, (ii) the price-earning ratio and (iii) the dividend yield. The results, reported in Table 6, Panel C, agree with our previous findings. In particular, monetary policy uncertainty is strongly correlated with stock returns in 8 out of 10 portfolios in the cases of both

¹⁸In general, the literature has identified channels through which inflation and an accommodating monetary policy adversely affect banks’ cash flows. In particular, Kessel (1956) and Alchian and Kessel (1960) argue that banks’ shareholders would suffer from inflation because banks are net holders of financial assets whose contractual characteristics are fixed in nominal terms. More recently, Dermine (1985 and 1987) considers the tax burden. He argues that as taxes are calculated on nominal profits, “the increase in after-tax earnings fueled by inflation is not sufficient to finance a constant level of real dividends and the retained earnings that are required to satisfy an exogenous capital adequacy ratio”. Therefore, exogenously imposed capital adequacy ratios together with taxation of nominal returns, he argues, means that inflation reduces banks’ cash flows. Empirical evidence seems to confirm this hypothesis (Dermine, 1999).

cash-flow-to-price and earnings-to-price portfolios and in 9 out of 10 portfolios in the case of dividend-yield portfolios. Moreover, in most instances β_i is not statistically different from zero. In particular, expected inflation is significant in only 6 out of 25 cases for the size and book-to-market portfolios and 4 out of 17 for the industry portfolios. Furthermore, for both cash-flow-to-price and earnings-to-price portfolios, expected inflation is not significant in 9 out of 10 cases and is never significant for dividend-yield portfolios. This supports the main theoretical findings.

As an additional robustness check, we have also investigated the relative importance of fundamental uncertainty and monetary policy uncertainty, by re-estimating equation (4.21) without the proxy for monetary policy uncertainty. The results (not reported) show that the model's fit, as measured by the Adjusted R^2 , deteriorates sharply, which suggests that a significant fraction of the explanatory power of model (4.21) is captured by the tracking portfolio. When Φ_t is not included in the regression, the average Adjusted R^2 falls from 0.9 to 0.8 on average. Moreover, the coefficient of expected inflation turns statistically significant.¹⁹ This suggests that fundamental uncertainty by itself is not enough to resolve the Fisher puzzle; that is, that the joint presence of both fundamental and policy uncertainty is required, as the model indicates.

4.7 Conclusion

We have studied the relationship between inflation and stock returns from an original perspective. If investors do not know the monetary policy stance and use inflation as a signal to it, the learning process generates uncertainty that increases the risk premium. A change in consumer prices has effects on

¹⁹The sign of the coefficient is now positive. This contrasts with the strict definition of the Fisher relationship, which would impose a zero correlation between excess returns and expected inflation (restriction 4.17)

the equity risk premium that vary with investors' beliefs. We show that a change in consumer prices that confirms investors' beliefs leads to a reduction in risk premia, while a change that contradicts them has the opposite effect. We then construct a market-based proxy of monetary policy uncertainty, showing that it is priced and demonstrating that this helps to explain the Fisher puzzle. We further show that, by conditioning on it, the Fisher puzzle disappears.

Our results provide a link between asset pricing and monetary economics. They suggest a new channel through which the central bank affects financial markets, one that has not been properly explored to date. Moreover, we shed new light on the “rules versus discretion” debate, quantifying the cost - in terms of higher risk premium - of discretion. It would be interesting to extend this analysis to other countries to see whether different degrees of disclosure of the central bank's targets are related to differing impact of monetary policy uncertainty on stock returns.

4.8 Appendix

Proof of Theorem 1. We first solve investor's learning problem and then we define the equilibrium stock price. Investors observe D_t , C_t and p_t (signals) and try to infer the value of θ_t .²⁰ The unobservable component (θ_t) can take values a and b (that is $E = [a, b]$). From Liptser and Shirayayev (pag. 333) we know that the posterior probability of a is:

$$d\pi_{a,t} = (1 - 2\pi_{a,t})\lambda dt + \pi_{a,t}(\mu_{i,t} - \bar{\mu}_t)'(\Sigma\Sigma')^{-\frac{1}{2}}d\mathbf{v}_t \quad (4.22)$$

where $\mu_{i,t} = (\mu_D + \beta\theta_t, \mu_p + \delta\theta_t, \mu_c)'$, $\bar{\mu}_t = \sum_{j \in \{a,b\}} \mu_{j,t}\pi_{j,t}$, $\Sigma = (\mathbf{b}_D, \mathbf{b}_p, \mathbf{b}_c)'$ and $\mathbf{b}_D = (\sigma_D, 0, 0)'$, $\mathbf{b}_p = (0, \sigma_p, 0)'$ and $\mathbf{b}_c = (\sigma_{cD}, 0, \sigma_e)'$. The vector

²⁰Investors observe nominal Dividends (D_t^n). However, given that they also observe consumption prices (p_t), the informational content of the dividends coincides with their real component (D_t).

$\nu = (\nu_D, \nu_p, \nu_e)$, defined on the agents' new filtration, follows:

$$d\mathbf{v}_t = (\Sigma \Sigma')^{-\frac{1}{2}} (d\mathbf{W}_t - \bar{\mu}_t dt) \quad (4.23)$$

We can also rewrite equation (4.22) as:

$$\begin{aligned} d\pi_{a,t} &= (1 - 2\pi_{a,t})\lambda dt + \pi_{a,t}(1 - \pi_{a,t})(a - b)(\beta\Omega_t d\nu_{D,t} + \delta d\nu_{p,t} - \gamma d\nu_{e,t}) = \\ &= \mu_{\pi_{a,t}} dt + \sigma_{\nu,t} d\nu_t \end{aligned} \quad (4.24)$$

where: $\mu_{\pi_{a,t}} = (1 - 2\pi_{a,t})\lambda$, $\Omega_t = \frac{\sigma_{cD} - \sigma_e}{\sigma_e}$, $\sigma_{\nu,t} = (\sigma_{\nu_D,t}, \sigma_{\nu_p,t}, \sigma_{\nu_e,t})'$, where $\sigma_{\nu_D,t} = \pi_{a,t}(1 - \pi_{a,t})(a - b)\beta\Omega_t$, $\sigma_{\nu_p,t} = \pi_{a,t}(1 - \pi_{a,t})(a - b)\delta$ and $\sigma_{\nu_e,t} = -\pi_{a,t}(1 - \pi_{a,t})(a - b)\gamma$.

We can now define the equilibrium stock price. The price of the stock is the present discounted value of its future dividends, discounted by the stochastic discount factor (Campbell and Kyle, 1993, Wang, 1993), that is:

$$S_t^n = E_t \left[\int_t^\infty \frac{n_s D_s^n ds}{n_t} \right] \quad (4.25)$$

where S_t^n is the nominal value of the stock, D_t^n is the nominal dividend and n_t is the nominal stochastic discount factor. We define $S(s)$ as the real price of the stock and r^n the nominal interest rate. Given the law of motion of the real consumption process described in equation (4.2), in equilibrium the *real* pricing kernel is: $m_t = u_c(t, C_t) = e^{-\phi t} C^{-\rho}$. This implies that the law of motion of the real stochastic discount factor is:

$$dm_t = k_t m_t dt - \rho m_t \sigma_D dz_{D,t} \quad (4.26)$$

where $k_t = -\phi - \rho(\mu_c + \gamma\theta_t) + \frac{1}{2}\rho(\rho + 1)\mathbf{b}_c \mathbf{b}_c'$.

Given that the nominal stochastic discount factor (n_t) is a function of both real stochastic factor (m_t) and the price level (p_t), it can be defined as: $n_t = \frac{m_t}{p_t}$. The nominal dividend can be therefore expressed in terms of the real dividend (D_t) and of the price level (p_t) as: $D_t^n = D_t p_t$. Its law of motion

follows:

$$\begin{aligned} dD_t^n &= p_t dD_t + D_t dp_t = D_t^n [\mu_D + \beta\theta_t + \mu_p + \delta\theta_t] dt + D_t^n [\sigma_D dz_{D,t} + \sigma_p dz_{p,t}] = \\ &= D_t^n \mu_{D^n} dt + D_t^n [\sigma_D dz_{D,t} + \sigma_p dz_{p,t}] \end{aligned} \quad (4.27)$$

To determine the price of the stock we use Proposition 1 from Veronesi (1999) which relates the value of a stock to the present discounted value of future dividends. For the general case where the underlying regimes $(\theta_{i,t})$ can take $i = 1 \dots N$ values, the nominal price of the stock (S_t^n) can be defined as a function of nominal dividends and nominal discount factors, such as:

$$\begin{aligned} S_t^n n_t &= E_t \left[\int_t^\infty n_s D_s^n ds \right] = D_t^n E_t \left[\int_t^\infty \frac{n_s D_s^n}{n_t D_t^n} ds \right] = \\ &= D_t^n E_t \left[\int_t^\infty \frac{m_s D_s}{m_t D_t} ds \right] = D_t^n \sum_{i=1}^N E_t \left[\int_t^\infty \frac{\Xi_s}{\Xi_t} ds | \theta_t = \theta_{i,t} \right] \pi_{i,t} = \\ &= D_t^n \sum_{i=1}^N C_{i,t} \pi_{i,t} \end{aligned} \quad (4.28)$$

where $\Xi_t = m_t D_t$ and $C_{i,t} = E_t \left[\int_t^\infty \frac{\Xi_s}{\Xi_t} ds | \theta_t = \theta_{i,t} \right]$. Ξ_t follows:

$$d\Xi_t = [k + \mu_D + \beta\theta_t - \rho\sigma_D^2] dt + [1 - \rho] \sigma_D dz_{D,t} = \mu_{\Xi,t} dt + [1 - \rho] \sigma_D dz_{D,t} \quad (4.29)$$

The solution requires us to determine the value of $C_{i,t}$. Let's define $\mathbf{A}_t = -\mathbf{\Lambda} - \text{diag}(\mu_{\Xi_t})$. Given that we consider only two regimes, (θ_t can only take values a and b), we have that:

$$\mathbf{A} = - \begin{bmatrix} -\lambda & \lambda \\ \lambda & -\lambda \end{bmatrix} + \begin{bmatrix} -\mu_{\Xi_t|\theta_t=a} & 0 \\ 0 & -\mu_{\Xi_t|\theta_t=b} \end{bmatrix} = \begin{bmatrix} -\mu_{\Xi_t|\theta_t=b} + \lambda & -\lambda \\ -\lambda & -\mu_{\Xi_t|\theta_t=a} + \lambda \end{bmatrix} \quad (4.30)$$

$C_{i,t}$ can be computed using the relationship $\mathbf{C}_t = \mathbf{A}_t^{-1} \mathbf{1}_m$, where $\mathbf{1}_m$ is a unity vector. That is, we have that: $C_{a,t} = \frac{2\lambda - \mu_{\Xi_t|\theta_t=b}}{(-\mu_{\Xi_t|\theta_t=a} + \lambda)(-\mu_{\Xi_t|\theta_t=b} + \lambda) - \lambda^2}$ and

$C_{b,t} = \frac{2\lambda - \mu_{\Xi_t|\theta_t=a}}{(-\mu_{\Xi_t|\theta_t=a} + \lambda)(-\mu_{\Xi_t|\theta_t=b} + \lambda) - \lambda^2}$. The nominal price is:

$$S_t^n = D_t^n \sum_{i \in \{a,b\}} \pi_{i,t} C_{i,t} \quad (4.31)$$

The law of motion of the nominal price of the asset is determined by totally differentiating equation (4.31). That is:

$$\begin{aligned} \frac{dS_t^n}{S_t^n} = & \left[\mu_{D,t}^n dt + D_t^n [\sigma_D dz_{D,t} + \sigma_p dz_{p,t}] \right] + \frac{\sum_{i \in \{a,b\}} C_{i,t} \{ \mu_{\pi,t} dt + \sigma_{\nu,t} d\nu_t \}}{\sum_{i \in \{a,b\}} C_{i,t} \pi_{i,t}} + \\ & \frac{\sum_{i \in \{a,b\}} C_{i,t} \pi_{i,t} \{ D_t^n \sigma_{\nu,t} d\nu_t (\sigma_{D,t} dz_{D,t} + \sigma_{p,t} dz_{p,t}) \}}{\sum_{i \in \{a,b\}} C_{i,t} \pi_{i,t}} \end{aligned} \quad (4.32)$$

Let's specify that $a = 1$ and $b = -1$. Using equation (4.23), we have that: $dz_{D,t} = \frac{1}{\sigma_D} \{ d\nu_{D,t} - \beta [\theta_t + (1 - 2\pi_{a,t})] dt \}$ and $dz_{p,t} = \frac{1}{\sigma_p} \{ d\nu_{p,t} - \delta [\theta_t + (1 - 2\pi_{a,t})] dt \}$. After substituting out for $C_{a,t}$ and $C_{b,t}$, we can write:

$$\begin{aligned} \frac{dS_t^n}{S_t^n} = & \{ \mu_D + \mu_p - (1 - 2\pi_{a,t})(\beta + \delta) + (1 - 2\pi_{a,t}) \frac{2(\beta - \rho\gamma)\lambda}{\Psi_t} + \\ & \pi_{a,t}(1 - \pi_{a,t}) \frac{4(\beta - \rho\gamma)}{\Psi_t} (\beta\Omega_t + \delta - \gamma) \} dt + \sigma_{s,t} d\nu_t \end{aligned} \quad (4.33)$$

where $\Psi_t = \sum_{i \in \{a,b\}} C_{i,t} \pi_{i,t} = C_{a,t} \pi_{a,t} + C_{b,t} \pi_{b,t} = -\mu_D - k + 2\lambda + (2\pi_{a,t} - 1)(\beta - \gamma\rho) + \rho \mathbf{b}_D \mathbf{b}_D'$, $\Omega_t = \frac{\sigma_{cD} - \sigma_e}{\sigma_e}$, $\sigma_{s,t} = (1 + \frac{(\mu_{\Xi_t|\theta_t=a} - \mu_{\Xi_t|\theta_t=b})}{\Psi_t} \sigma_{\nu_{D,t}}, 1 + \frac{(\mu_{\Xi_t|\theta_t=a} - \mu_{\Xi_t|\theta_t=b})}{\Psi_t} \sigma_{\nu_{p,t}}, \frac{(\mu_{\Xi_t|\theta_t=a} - \mu_{\Xi_t|\theta_t=b})}{\Psi_t} \sigma_{\nu_{e,t}})$ and $k = -\phi - \rho\mu_c + \frac{1}{2}\rho(\rho + 1)b_c b_c'$.

Therefore, we can define the expected value of the nominal stock return as:

$$E_t \left(\frac{dS_t^n}{S_t^n} \right) / dt = \mu_{D,t} + \mu_{p,t} - (1 - 2\pi_{a,t})(\beta + \delta) + (1 - 2\pi_{a,t}) \frac{2(\beta - \rho\gamma)\lambda}{\Psi_t} + \Phi_t \quad (4.34)$$

where $\Phi_t = \pi_{a,t}(1 - \pi_{a,t}) \frac{4(\beta - \rho\gamma)}{\Psi_t} (\beta\Omega_t + \delta - \gamma)$. We can also derive the equilibrium riskless real rate of return, as:

$$E[r_t] = \phi + \rho\mu_c - \frac{1}{2}\rho(\rho + 1)\mathbf{b}_c \mathbf{b}_c' + \rho\gamma\theta_t \quad (4.35)$$

and expected inflation as:

$$E \left[\frac{dp_t}{p_t} \right] / dt = [\mu_p + (2\pi_{a,t} - 1)\delta] \quad (4.36)$$

The expected nominal riskless rate is:

$$\begin{aligned} E[r_t^n] &= E[r_t] + \mu_p + \delta(2\pi_{a,t} - 1) = \\ &= \phi + \rho\mu_c - \frac{1}{2}\rho(\rho + 1)\mathbf{b}_c \mathbf{b}_c' + \mu_p + (2\pi_{a,t} - 1)(\delta + \gamma\rho) \end{aligned} \quad (4.37)$$

Using equation (4.37) and (4.34), we find that the expected excess rate of return on the stock (risk premium) is:

$$\begin{aligned}
E[R_t] &= E\left[\frac{dS_t^n}{S_t^n} + \frac{D_t^n}{S_t^n} - r_t^n dt\right]/dt = \\
&= \mu_{D,t} + \mu_{p,t} - (1 - 2\pi_{a,t})(\beta + \delta) + (1 - 2\pi_{a,t})\frac{2(\beta - \rho\gamma)\lambda}{\Psi_t} + \\
&\quad \frac{D_t^n}{S_t^n} - \phi - \rho\mu_c + \frac{1}{2}\rho(\rho + 1)\mathbf{b}_c\mathbf{b}_c' - \mu_p - (2\pi_{a,t} - 1)(\delta + \gamma\rho) + \Phi_t \\
&= F_t + \Phi_t
\end{aligned} \tag{4.38}$$

Calibration of the model. The model is calibrated using data provided by the Federal Reserve Bank of St. Louis ("FRED") for the period January 1965-December 1998. Given that the unit period is a quarter, all the parameters are defined on a quarterly basis. The following values are used: $\mu_p = 0.0112$, $\sigma_p = 0.0062$, $\mu_{D^n} = 0.0220$, $\sigma_{D^n} = 0.0222$, $\mu_D = 0.0108$ and $\sigma_D = 0.023$. The values of the parameters for real consumption are $\mu_c = 0.0077$ and $\sigma_c = 0.009$. The decomposition of the volatility of consumption is based on the assumption that dividends represent roughly 5% of overall consumption as reported in the literature (see Berk, Green and Naik, 1999, Cecchetti, Lam and Mark, 1993, Campbell and Cochrane, 1999, Barberis, Huang and Santos, 2001). This implies that $\sigma_e = \sqrt{\sigma_c^2 - 0.05^2\sigma_{cD}^2/0.95^2}$, where $\sigma_{cD} = \sigma_D$. The impact of monetary policy on consumption is equal to 5% of that on dividends ($\gamma = 0.05\beta$); we consider different values of β in the interval [0.001-0.5]: the case of $\beta = 0.5$ corresponds to an elasticity of consumption to (and GDP) to θ_t equal to 2.5%. The impact of monetary policy on prices is assumed to be equal to 2% yearly, which corresponds also to the sample value of price volatility.

In order to determine the value of λ we consider the transition probability $P_t(a, b)$, that is the probability of moving from one regime to the other. Using quarterly data the Markov-switching model delivers an estimate of

$P_t(a, b) = 0.1640$. The degree of risk aversion (ρ) is assumed to be equal to 2. The results are robust to changes in ρ .

In order to determine λ we apply Karlin and Taylor (pag. 151-152). We have that:

$$P_s = \frac{1}{2\lambda} \begin{pmatrix} \lambda(1 + e^{-2\lambda s}) & \lambda(1 - e^{-2\lambda s}) \\ \lambda(1 - e^{-2\lambda s}) & \lambda(1 + e^{-2\lambda s}) \end{pmatrix} \quad (4.39)$$

and λ can be obtained as the solution to $0.5(1 - e^{-2\lambda s}) = 0.1640$, where $s = 1$ for a unit period equal to a quarter. The solution is $\lambda = 0.1987$.

The data. The data and the procedure for constructing portfolios come from K. French's web page. Industry portfolios are constructed by first assigning each NYSE, AMEX, and NASDAQ stock to an industry portfolio at the end of June of year t based on its four-digit SIC code at that time and then computing the returns from July of t to June of $t + 1$. The size and book-to-market portfolios are constructed at the end of June of each year as the intersections of 5 portfolios formed on (i) size (market equity, ME) and (ii) the ratio of book equity to market equity (BE/ME). For size the breakpoints for year t are the NYSE market equity quintiles at the end of June of each year; for BE/ME they depends on book equity for the last fiscal year divided by ME in December of year $t - 1$. The BE/ME breakpoints are NYSE quintiles. We considered the period July 1965-December 1998.²¹

Cash-flow-to-price portfolios are constructed by grouping stocks into 10 deciles. Portfolios are formed on the basis of the ratio of cash flow (CF) to price (P) computed at the end of June of year t , using NYSE breakpoints. The CF used in June of year t is total earnings before extraordinary items, plus equity share of depreciation, plus deferred taxes (if available) for the last fiscal year end in $t - 1$. P is price times the number of shares outstanding at the end of December of $t - 1$. Earnings-to-price portfolios and dividend-yield-to-price

²¹Fama and French provide 38 portfolios, but given that 5 of them contain a large number of missing values in the sample period, we use only 33 portfolios.

portfolios are constructed analogously. Portfolios are formed on the basis of the value of earnings, dividend yields and prices computed at the end of June; earnings used in June of year t are total earnings before extraordinary items for the last fiscal year end in $t - 1$. The dividend yield used to form portfolios in June of year t is the total dividend paid from July of year $t - 1$ to June of year t per dollar of equity.

Indexes of Business Cycle. We use two sets of indexes: NBER Business Cycle Reference Dates (peaks, troughs and peaks-troughs) and James Stock's coincident and leading indexes. In particular, NBERPEAK is the series showing the Business Cycle Reference Dates of the peaks: the index takes value 1 at peaks and zero elsewhere; NBERTROU contains the Business Cycle Reference Dates of the troughs: the index takes value 1 at the trough and zero elsewhere; NBERDATE is constructed as the difference between the previous two indexes; NBERCYCL is a linear interpolation of NBERPEAK and NBERTROU: it attaches a value of 1 to NBER peaks and a value of -1 to troughs and then connects peaks and troughs by means of linear segments. XLI is Stock's NBER Experimental Leading Index: it is the forecast of the growth of the Experimental Consumer Index in the following 6 months; XRI is Stock's NBER Experimental Recession Index, measuring the probability that the economy is in recession in the next 6 months; XLI_2 is Stock's NBER Alternative Non Financial Experimental Leading Index, and XRI_2 is Stocks' NBER Alternative Non Financial Experimental Recession Index. All these indexes are constructed as forecasts six months ahead. Therefore, to account for possible delays in investors' reactions or for misalignments with financial market variables, alternative measures constructed as the 6th order lags of Stock's indexes (respectively XLIL6, XLI_2L6, XRIL6 and XRI_2L6) are also used. We also consider Stock's Experimental Coincident Recession Index: given the non-stationarity of this index, we use two transformations, that is (i) the logarithm of its first differences (XCI_1) and (ii) the detrended

logarithm (XCI_2).

The Markov-Switching VAR. To identify monetary policy regimes, a two-state Markov-switching VAR model is used. The state-space representation of the model is the following:

$$\begin{aligned}
y_t &= c_{s_t} + A_{1,s_t}y_{t-1} + \dots + A_{p,s_t}y_{t-p} + \varepsilon_t = \\
&= \left(\xi_t' \otimes I \right) \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} + (\xi_t' \otimes I) \begin{bmatrix} A_{11} \\ A_{12} \end{bmatrix} y_{t-1} + \dots + \left(\xi_t' \otimes I \right) \begin{bmatrix} A_{p1} \\ A_{p2} \end{bmatrix} y_{t-p} + \varepsilon_t \\
\xi_t &= F' \xi_{t-1} + \eta_t
\end{aligned}$$

where: $\xi_t = [1, 0]'$ if $s_t = 1$ and $\xi_t = [0, 1]'$ if $s_t = 2$; s_t is an unobserved random variable that takes the values 1 or 2 according to which regime the process is in at time t . $F = \{p_{ij}\}_{i,j=1,2}$ is the transition matrix and p_{ij} is the probability that $s_t = j$ given that $s_{t-1} = i$.

The hypotheses underlying the statistical model are standard: the error term in the observation equation, ε_t , is assumed to be i.i.d. normal, with covariance matrix Σ_{s_t} ; η_t is a martingale difference sequence, independent of ε_t and of all available information, past values of s_t included. The VAR is stable in both states. The vector y_t contains five variables: the excess return on the market portfolio, the corresponding dividend yield, the risk-free real rate, CPI inflation and the growth rate of real GDP; the sample period starts in 1965:3 and ends in 1998:4.²² The model is estimated by maximum likelihood, using the EM algorithm. Given an ML-estimate of the vector of parameters (i.e. $\psi = (c_1', c_2', \text{vec}(A_{11})', \text{vec}(A_{12})', \dots, \text{vec}(A_{p2})', \text{vec}(\Sigma_1)', \text{vec}(\Sigma_2)', \text{vec}(F)')$), the hidden Markov process is estimated by iterating on the following set of equa-

²²The same model, estimated on monthly data, is much more noisy. The reason is that at higher frequencies it takes too high order a VAR to provide an adequate account of the correlation structure of the data, which inevitably reduces the efficiency of the estimates.

tions:

$$\hat{\xi}_{t|t} = \frac{\hat{\xi}_{t|t-1} \odot \eta_t}{1' (\hat{\xi}_{t|t-1} \odot \eta_t)}, \quad \hat{\xi}_{t+1|t} = F' \hat{\xi}_{t|t} \text{ and } \hat{\xi}_{t|T} = \hat{\xi}_{t|t} \odot \left\{ F \left[\hat{\xi}_{t+1|T} \div \hat{\xi}_{t+1|t} \right] \right\} \quad (4.40)$$

where $\hat{\xi}_{t|t-k} = E(\xi_t | Y_{t-k})$, $Y_t = \{y_1, y_2, \dots, y_t\}$, η_t represents the 2×1 vector whose i -th element is the conditional density $f(y_t | s_t = i, Y_{t-1}; \psi)$, 1 is a 2×1 unit vector and the symbols \odot and \div denote element-by-element multiplication and division.

While the index of switches in monetary policy regimes is quarterly, the data on returns are monthly. In order not to waste degrees of freedom, the index is therefore disaggregated to a monthly frequency by applying the method suggested by Chow and Lin (1971), with the inflation rate used as the indicator variable. The method consists of estimating the model using quarterly data, under the assumption that the error term is first-order autocorrelated, and then using the GLS coefficients to estimate the endogenous variable at missing points.

Figure 1 - Excess Returns and the value of the monetary policy parameter δ






Figure 2 - Excess Returns and the share of dividends over consumption






Figure 3 - Perceived regimes of monetary policy and uncertainty

The figure reports the 6th-order centred moving average of the probability of a non-accommodative monetary policy as perceived by the investors and estimated by means of a Markov-Switching VAR. It also reports the uncertainty about the monetary policy, calculated as the product between the probability of tight monetary policy and the complement to one.

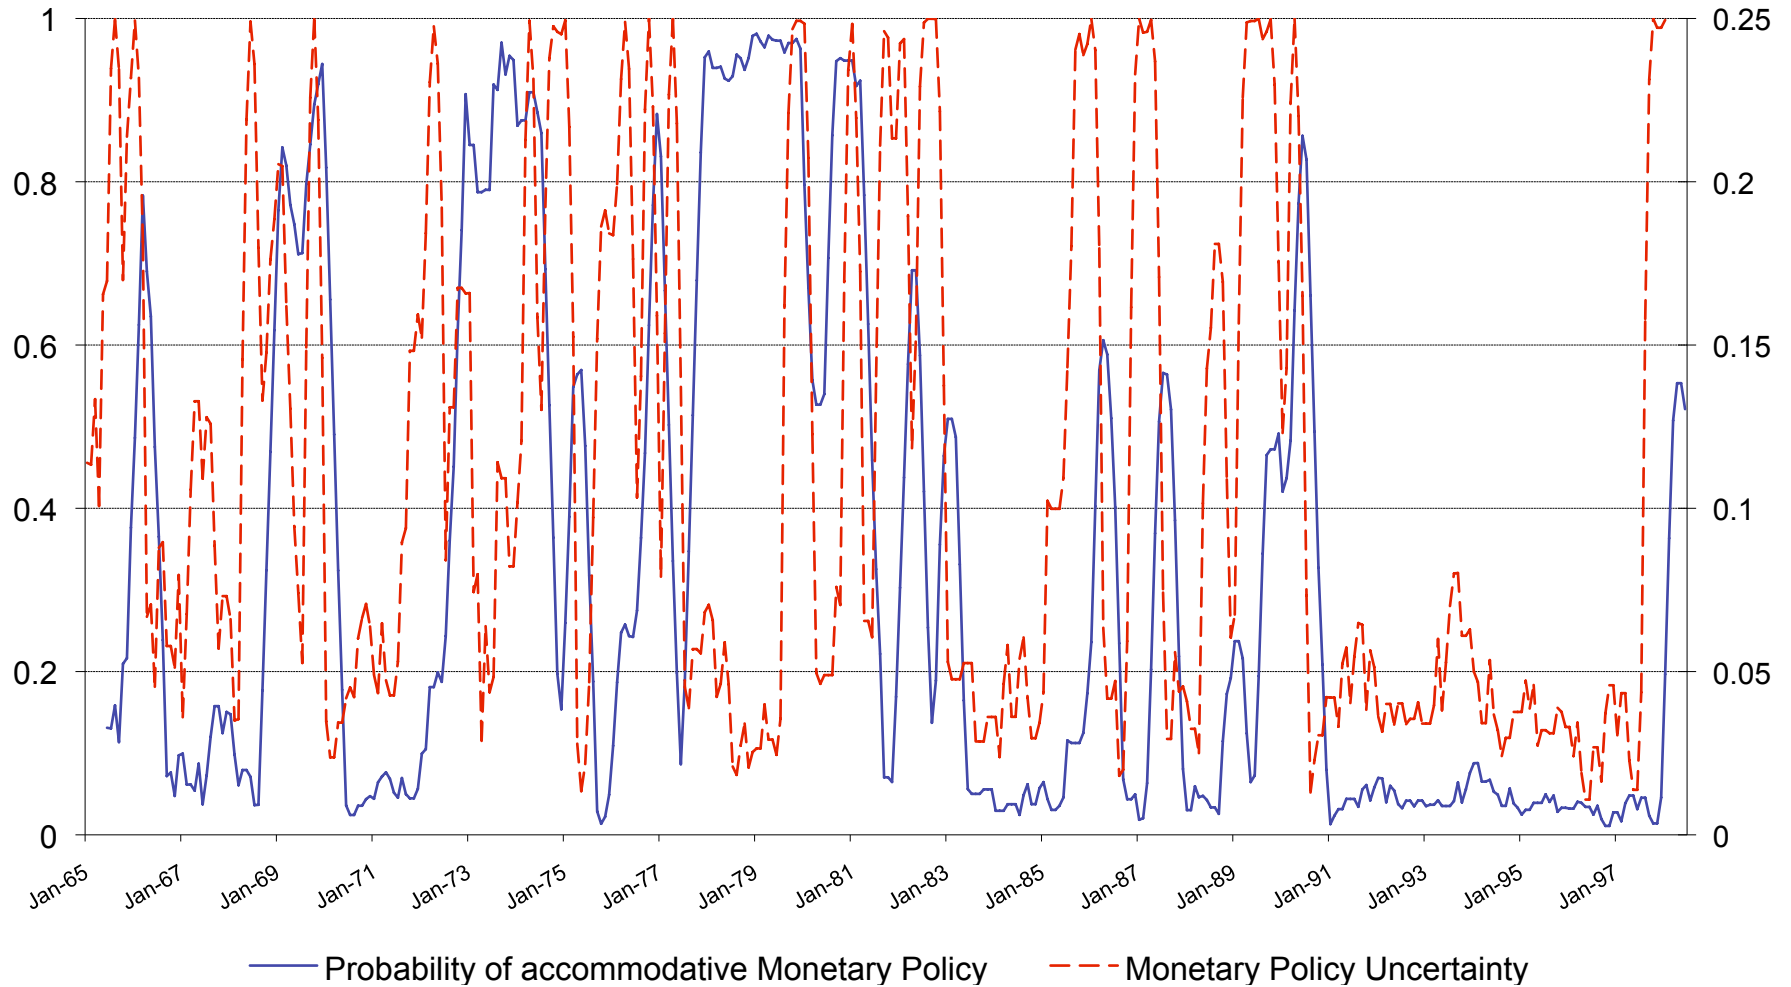


Table 1:
Probability of the Monetary Policy Regime and Business Cycle

The table reports summary statistics about the regression: $\pi_t = \alpha + \beta\Phi_t + \gamma BC_t + \varepsilon_t$, where π_t is the probability of accommodative monetary policy regime derived from the estimated Markow-switching model; Φ_t is the index of monetary policy stance computed by Bernanke-Mihov and BC is either a business cycle index derived from NBER dating or one of James Stock's indices. The index of Bernanke-Mihov increases as the monetary policy becomes more accommodative. The first column indicates which business cycle index has been used in the regression; the second column shows the adjusted R^2 of the regression; the third and fourth ones present, respectively, the estimated coefficient and the corresponding t -statistics of the monetary policy index; finally, the last two columns report the point estimate and the t -statistic of the coefficient of the business cycle index. The sample spans the period from January 1961 to December 1996.

<i>Business Cycle index</i>	<i>Adjusted R^2</i>	<i>β</i>	<i>t</i>	<i>γ</i>	<i>t</i>
NBERCYCL	0.200	2.97	5.14	-0.10	-2.07
NBERDATE	0.202	3.68	9.06	-0.26	-2.20
NBERPEAK	0.200	3.73	9.21	-0.33	-2.00
NBERTROU	0.194	3.78	9.31	0.18	1.11
XLI	0.229	2.87	6.34	0.03	4.26
XLI_2	0.221	3.37	8.14	0.03	3.80
XRI	0.309	2.13	4.95	-0.62	-7.94
XRI_2	0.249	3.27	8.11	-0.86	-5.36
XRI_C	0.219	3.81	9.59	-0.24	-3.66
XCI_1	0.200	3.70	9.10	7.30	2.07
XCI_2	0.262	1.96	3.94	-3.28	-5.94
XLIL6	0.196	3.82	9.48	0.01	1.46
XLI_2L6	0.192	3.86	9.50	0.01	0.77
XRIL6	0.224	3.71	9.36	-0.28	-3.94
XRI_2L6	0.197	3.90	9.62	-0.26	-1.61

Table 2:
Economic tracking portfolio for I_t : regression coefficients and portfolio weights

This table reports results from the following regression: $I_{t+12} = \alpha + \beta \mathbf{B}_t + \chi \mathbf{Z}_{t-1} + u_t$, where \mathbf{B}_t denotes the vector of base assets and \mathbf{Z}_{t-1} the set of control variables. The base assets are two bond portfolios and one equity portfolio. *Term* stands for the spread between the yield on 10-year Treasury bonds and 3-month Treasury bills; *junk* is the difference between Moody's Baa and Aaa corporate yields; *ME1/ME5* is the return on an arbitrage portfolio which is long on stocks of small firms (first NYSE market equity quintile) and short on stocks of big firms (fifth NYSE market equity quintile). Their returns are in excess of the riskless rate. The vector of control variables includes expected inflation (inf^e), actual CPI inflation ($infl$), inflation measured by the year-on-year rate of change of the index of producer prices for finished goods ($infl^{ppi}$) and the share of household expenditure on durable goods out of total expenditure ($shrdc$). T-values in the last two columns are computed by using the OLS estimator of the covariance matrix of the estimated coefficients in the first case and by using the White estimator in the second case, so as to correct for heteroscedasticity in the residuals.

<i>Regressors</i>	<i>Coefficients</i>	<i>t-statistics (OLS)</i>	<i>t-statistics (White)</i>
Base assets			
<i>term</i>	-0.397	-2.120	-2.264
<i>junk</i>	32.493	2.642	2.563
<i>ME1/ME5</i>	0.049	2.589	2.510
Control variables			
<i>Inf^e</i>	-12.781	-3.845	-3.670
<i>Infl</i>	3.031	1.530	1.514
<i>Inf^{ppi}</i>	0.397	2.187	2.077
<i>exmkt</i>	-0.306	-3.332	-3.188
<i>shrdc</i>	1.082	3.132	3.322
<i>R²</i>	0.099		
<i>Adjusted R²</i>	0.080		
<i>Standard error</i>	0.072		
<i>Residual Autocorrelation</i>	$F(12,368)=1.551$		<i>p-value</i> =0.104
<i>Heteroschedasticity (linear)</i>	$\chi(8)=27.907$		<i>p-value</i> =0.001
<i>Heteroschedasticity (exp.)</i>	$\chi(8)=8.361$		<i>p-value</i> =0.399
<i>Current sample</i>	1965.8-1997.12		

Table 3: Summary Financial Statistics

Excess returns on 25 book/market and 17 industry portfolios. *MARKET* is the market portfolio and *EXMKT* is the excess return of the market portfolio. The one-month Treasury bill proxies for the riskless rate. The sample covers the period July 1965 - December 1998. Both the sample means and their standard deviations are annualized. ρ_j , $j=1,2,3,4,12,24$, is the sample autocorrelation of order j .

<i>Portfolio</i>	<i>Mean</i>	<i>Std. Dev.</i>	ρ_1	ρ_2	ρ_3	ρ_4	ρ_{12}	ρ_{24}
Panel A: Size and book-to-market portfolios								
<i>S1/B1</i>	1.888	26.484	0.234	0.018	-0.011	0.002	0.054	-0.023
<i>S1/B2</i>	8.703	23.119	0.226	-0.002	-0.025	-0.002	0.065	-0.016
<i>S1/B3</i>	8.887	20.943	0.221	0.005	-0.012	-0.028	0.090	0.013
<i>S1/B4</i>	11.551	19.771	0.224	-0.013	-0.015	-0.020	0.124	-0.011
<i>S1/B5</i>	12.795	20.857	0.248	-0.006	-0.024	-0.034	0.178	0.053
<i>S2/B1</i>	4.907	25.228	0.168	-0.036	-0.053	-0.039	-0.016	-0.036
<i>S2/B2</i>	7.725	21.157	0.183	-0.031	-0.036	-0.042	0.041	0.004
<i>S2/B3</i>	10.546	19.080	0.184	-0.030	-0.029	-0.021	0.048	-0.035
<i>S2/B4</i>	11.385	18.000	0.171	-0.041	-0.029	-0.009	0.092	0.009
<i>S2/B5</i>	12.280	19.906	0.160	-0.067	-0.063	-0.041	0.135	0.026
<i>S3/B1</i>	5.449	23.173	0.144	-0.026	-0.036	-0.064	-0.008	-0.044
<i>S3/B2</i>	8.589	19.240	0.174	-0.012	-0.001	-0.060	0.014	-0.003
<i>S3/B3</i>	8.421	17.540	0.155	-0.039	-0.039	-0.045	0.012	-0.018
<i>S3/B4</i>	10.334	16.498	0.161	-0.023	-0.004	-0.038	0.051	0.032
<i>S3/B5</i>	11.585	18.850	0.153	-0.075	-0.067	-0.029	0.095	-0.007
<i>S4/B1</i>	6.157	20.235	0.107	-0.027	-0.026	-0.057	-0.021	-0.036
<i>S4/B2</i>	5.421	18.488	0.128	-0.028	-0.032	-0.026	-0.028	-0.011
<i>S4/B3</i>	8.436	17.104	0.080	-0.026	-0.014	-0.072	0.003	-0.008
<i>S4/B4</i>	9.736	16.184	0.082	0.000	0.001	-0.062	0.055	0.013
<i>S4/B5</i>	11.127	18.887	0.044	-0.035	-0.019	-0.020	0.035	0.002
<i>S5/B1</i>	6.248	16.771	0.055	-0.002	0.005	-0.018	0.056	-0.012
<i>S5/B2</i>	5.965	16.111	0.036	-0.060	-0.002	0.007	-0.002	-0.020
<i>S5/B3</i>	6.067	15.162	-0.034	-0.052	0.007	-0.035	-0.012	0.019
<i>S5/B4</i>	7.839	14.822	-0.055	0.008	0.064	-0.084	0.035	0.022
<i>S5/B5</i>	8.601	16.184	0.021	-0.007	-0.039	0.003	0.046	0.012
<i>MARKET</i>	12.607	15.475	0.054	-0.039	-0.011	-0.033	0.019	-0.012
<i>EXMKT</i>	6.259	15.547	0.060	-0.034	-0.009	-0.031	0.016	-0.015
Panel B: Industry Portfolios								
<i>Food</i>	8.559	15.970	0.067	-0.060	0.001	-0.023	0.070	-0.049
<i>Mines</i>	4.457	22.527	0.047	-0.011	-0.017	-0.057	-0.028	0.060
<i>Oil</i>	6.854	18.305	-0.013	-0.040	0.022	0.020	0.004	-0.043
<i>Clths</i>	5.493	21.358	0.236	0.045	-0.032	-0.059	0.068	-0.046
<i>Durbl</i>	6.992	18.855	0.107	0.046	0.002	-0.027	0.033	-0.008
<i>Chems</i>	5.319	18.533	0.013	-0.052	0.044	-0.013	-0.038	0.024
<i>Cnsum</i>	9.718	16.946	0.019	-0.005	-0.054	0.004	0.100	-0.002
<i>Cnstr</i>	7.136	20.499	0.109	-0.042	-0.025	-0.073	0.027	-0.003
<i>Steel</i>	2.576	21.235	-0.004	-0.055	-0.081	-0.014	-0.100	0.087
<i>FabPr</i>	5.879	18.165	0.153	-0.061	-0.038	-0.052	0.023	-0.006
<i>Machn</i>	6.314	19.937	0.097	0.008	-0.013	-0.051	0.015	0.053
<i>Cars</i>	5.529	20.317	0.132	-0.017	-0.018	-0.050	0.026	-0.031
<i>Trans</i>	6.411	20.909	0.124	-0.006	-0.087	0.044	-0.003	0.029
<i>Utils</i>	4.215	13.699	0.017	-0.108	0.011	0.029	0.044	0.020
<i>Rtail</i>	7.702	19.827	0.170	0.006	-0.073	-0.054	0.035	-0.061
<i>Finan</i>	8.099	17.913	0.126	-0.042	-0.037	-0.017	0.037	-0.051
<i>Other</i>	6.554	16.370	0.049	-0.038	-0.008	-0.075	-0.013	0.005
<i>MARKET</i>	12.607	15.475	0.054	-0.039	-0.011	-0.033	0.019	-0.012
<i>EXMKT</i>	6.259	15.547	0.060	-0.034	-0.009	-0.031	0.016	-0.015

Table 4: Expected Inflation and Excess Returns

Size and book-to-market portfolio excess returns are regressed on a constant and the proxy for expected inflation derived from the estimated Markov-switching model. The first column shows the adjusted R^2 of the model while the second one reports value of the regression coefficient of expected inflation. The last column reports the t -statistic for testing whether the expected inflation parameter is statistically significant. To correct for the bias due to the presence of a generated regressor, White standard errors have been used. The sample period is July 1965-December 1998. Returns are in excess of the 30-day Treasury bill.

Portfolio	Adjusted R ²	Expected Inflation	
		coefficient	t-stat
Panel A: Size and book-to-market portfolios			
S1/B1	0.011	-3.761	-1.996
S1/B2	0.020	-4.201	-2.596
S1/B3	0.019	-3.728	-2.521
S1/B4	0.024	-3.909	-2.905
S1/B5	0.023	-4.032	-2.693
S2/B1	0.013	-3.849	-2.171
S2/B2	0.017	-3.583	-2.386
S2/B3	0.027	-3.984	-2.850
S2/B4	0.025	-3.648	-2.816
S2/B5	0.018	-3.479	-2.437
S3/B1	0.018	-4.069	-2.546
S3/B2	0.022	-3.630	-2.631
S3/B3	0.027	-3.670	-2.816
S3/B4	0.031	-3.668	-3.029
S3/B5	0.021	-3.549	-2.410
S4/B1	0.021	-3.790	-2.617
S4/B2	0.024	-3.657	-2.802
S4/B3	0.019	-3.028	-2.380
S4/B4	0.032	-3.658	-3.142
S4/B5	0.022	-3.633	-2.601
S5/B1	0.051	-4.730	-4.230
S5/B2	0.035	-3.819	-3.456
S5/B3	0.033	-3.498	-3.330
S5/B4	0.029	-3.195	-3.044
S5/B5	0.043	-4.195	-3.878
Panel B: Industry market portfolios			
Food	0.039	-3.899	-4.110
Mines	0.001	-1.534	-1.122
Oil	0.014	-2.818	-2.532
Clths	0.020	-3.865	-2.992
Durbl	0.055	-5.462	-4.855
Chems	0.023	-3.615	-3.200
Cnsum	0.035	-3.971	-3.872
Cnstr	0.017	-3.493	-2.802
Steel	0.008	-2.636	-2.034
FabPr	0.020	-3.228	-3.009
Machn	0.036	-4.663	-3.928
Cars	0.040	-5.030	-4.151
Trans	0.023	-4.051	-3.196
Utils	0.019	-2.438	-2.907
Rtail	0.029	-4.263	-3.567
Finan	0.034	-4.079	-3.806
Other	0.030	-3.511	-3.623

TABLE 5: Evidence of pricing

This table reports the Generalized Method of Moments tests of the moment conditions of equations 17-20 in the text. We consider the standard Fama and French factors (Panels A and B) and the “orthogonalized” ϕ_t factor (Panels C and D), where the monetary policy uncertainty factor (ϕ_t) has been previously orthogonalized by regressing it on the Fama and French factors. Panels A and C report the estimates for 25 book-to-market and size portfolios, while Panels B and D report the estimates for 17 industry portfolios. The vectors, δ , ϕ_{MKT} , ϕ_{HML} , ϕ_{SMB} , ϕ_{ϕ_t} , contain the coefficients of the linear relationship between λ , λ_{MKT} , λ_{HML} , λ_{SMB} , λ_{ϕ_t} and the vector of instruments, \mathbf{Z} . The instrumental variables are a constant, one month T-bill yield (*T-bill*), dividend yield of the S&P 500 index (*Div*), term premium – spread between a 10 years and 1year Treasury bond yield (*Term*), junk premium – spread between Moody’s Baa and Aaa corporate bond yields (*Junk*), difference between the one month returns of a three month and one month T-bill (*Hb3*) and a January dummy that takes value 1 for January and 0 otherwise. We report the estimated coefficients as well as the *t-stat*. Last rows of each panel reports the test for overidentifying restrictions and the Wald test for the significance of the ϕ_{ϕ_t} coefficients. The value of the χ^2 statistic, the degrees of freedom and the *p-value* are reported. The coefficients for the dividend yield (*Div*), the difference between the one month returns of a three month and one month T-bill (*Hb3*), the Treasury Bill (*T-bill*) and the junk premium (*Junk*) have been divided by 100.

Non-orthogonalized factors

Panel A: Book-to-market and size portfolios

	δ		ϕ_{MKT}		ϕ_{HML}		ϕ_{SMB}		ϕ_{ϕ}	
	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>
<i>Constant</i>	0.84	1.73	1.33	0.27	51.72	6.27	49.94	4.77	-57.92	-2.97
<i>Div</i>	17.46	6.38	-120.72	-4.70	-174.03	-4.18	392.40	7.68	-709.50	-6.49
<i>Junk</i>	-55.87	-7.68	155.19	3.09	35.65	0.41	-717.43	-9.18	1992.41	10.03
<i>Term</i>	50.35	7.17	-649.59	-11.84	499.78	5.20	698.94	5.91	-2223.51	-9.41
<i>Hb3</i>	-1.18	-2.90	17.70	5.35	51.18	8.09	12.45	2.19	3.05	0.26
<i>T-bill</i>	-4.45	-4.62	52.98	6.00	17.32	1.44	-214.24	-9.97	204.66	6.40
<i>Dummy_{JANUARY}</i>	-7.67	-11.95	-22.11	-5.13	87.53	10.66	110.82	11.08	-61.28	-2.96

Overidentifying restrictions test: $\chi^2 = 61.65$; degrees of freedom: 147; *p-value*: 0.99.

Wald test: $\chi^2 = 323.11$; degrees of freedom: 7; *p-value*: 0.00.

Panel B: Industry portfolios

	δ		ϕ_{MKT}		ϕ_{HML}		ϕ_{SMB}		ϕ_{ϕ}	
	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>	<i>Mean</i>	<i>t-test</i>
<i>Constant</i>	5.33	7.33	5.48	0.90	-1.79	-0.17	11.27	0.67	-128.60	-4.93
<i>Div</i>	-33.01	-7.48	-25.60	-0.98	193.09	2.97	248.96	4.02	1215.85	8.11
<i>Junk</i>	12.65	1.18	228.21	3.95	425.35	3.06	-665.45	-4.73	-1240.16	-4.07
<i>Term</i>	-15.87	-1.95	72.43	0.88	613.22	5.21	-199.71	-1.03	686.83	2.30
<i>Hb3</i>	4.03	6.03	-0.74	-0.12	4.18	0.31	66.77	6.15	-133.77	-6.98
<i>T-bill</i>	-0.10	-0.08	-11.27	-0.99	-201.36	-9.34	-171.29	-7.00	-34.42	-0.83
<i>Dummy_{JANUARY}</i>	0.33	0.50	-12.77	-2.37	29.82	2.61	-12.81	-1.02	-46.37	-2.09

Overidentifying restrictions test: $\chi^2 = 49.09$; degrees of freedom: 91; *p-value*: 0.99.

Wald test: $\chi^2 = 237.81$; degrees of freedom: 7; *p-value*: 0.00.

Orthogonalized Φ_t factor

Panel C: Book-to-market and size portfolios

	δ		Φ_{MKT}		Φ_{HML}		Φ_{SMB}		Φ_{ϕ}	
	Mean	t-test	Mean	t-test	Mean	t-test	Mean	t-test	Mean	t-test
<i>Constant</i>	0.84	1.73	-1.64	-0.31	38.35	3.60	23.16	3.06	-57.92	-2.97
<i>Div</i>	17.46	6.38	-157.21	-5.32	-337.88	-6.58	64.34	2.48	-709.51	-6.49
<i>Junk</i>	-55.86	-7.68	257.62	4.71	495.73	5.18	203.76	2.73	1992.41	10.03
<i>Term</i>	50.35	7.17	-763.90	-12.45	-13.76	-0.12	-329.19	-3.35	-2223.41	-9.41
<i>Hb3</i>	-1.18	-2.90	17.86	5.20	51.88	6.46	13.86	2.76	3.05	0.26
<i>T-bill</i>	-4.45	-4.61	63.50	6.49	64.58	4.97	-119.61	-9.71	204.64	6.40
<i>Dummy_{JANUARY}</i>	-7.67	-11.95	-25.26	-5.68	73.39	8.16	82.49	11.49	-61.28	-2.96

Overidentifying restrictions test: $\chi^2 = 61.35$; degrees of freedom: 147; *p-value*: 1.

Wald test: $\chi^2 = 320.04$; degrees of freedom: 7; *p-value*: 0.00.

Panel D: Industry portfolios

	δ		Φ_{MKT}		Φ_{HML}		Φ_{SMB}		Φ_{ϕ}	
	Mean	t-test	Mean	t-test	Mean	t-test	Mean	t-test	Mean	t-test
<i>Constant</i>	5.33	7.33	-1.13	-0.17	-31.49	-2.23	-48.18	-3.42	-128.59	-4.93
<i>Div</i>	-33.01	-7.48	36.90	1.31	473.86	5.94	811.12	12.37	1215.82	8.11
<i>Junk</i>	12.65	1.18	164.42	2.77	138.94	0.90	-1238.86	-8.63	-1240.11	-4.07
<i>Term</i>	-15.87	-1.95	107.75	1.26	771.83	5.65	117.86	0.83	686.86	2.30
<i>Hb3</i>	4.03	6.03	-7.62	-1.22	-26.70	-1.70	4.92	0.50	-133.77	-6.98
<i>T-bill</i>	-0.10	-0.08	-13.04	-1.12	-209.31	-8.51	-187.21	-6.98	-34.42	-0.83
<i>Dummy_{JANUARY}</i>	0.33	0.50	-15.16	-2.68	19.11	1.46	-34.24	-4.40	-46.36	-2.09

Overidentifying restrictions test: $\chi^2 = 49.09$; degrees of freedom: 91; *p-value*: .99.

Wald test: $\chi^2 = 237.81$; degrees of freedom: 7; *p-value*: 0.00.

Table 6: Expected Inflation and monetary uncertainty

Different groupings of portfolio excess returns are regressed on a constant, the three Fama-French risk factors (F_t), the tracking portfolio mimicking monetary policy uncertainty (Φ_t) and the proxy for expected inflation derived from the estimated Markov-switching model (Inf_t^e). We consider the 25 size and book-to-market portfolios (Panel A), the 17 industry portfolios (Panel B), cash flow-to-price portfolios (Panel C), earnings-to-price portfolios (Panel D) and dividend-to-price portfolios (Panel E). In the case of Cash flows to price portfolios (CF/P), cash flow is the cash flow at the last fiscal year end of the prior calendar year, while price is represented by the market capitalization (ME) at the end of December of the prior year. In the case of earnings-to-price, we consider the excess returns on portfolios formed on deciles of the distribution of E/P, where E/P are the earnings before extraordinary at the last fiscal year end of the prior calendar year divided by ME at the end of December of the prior year. In the case of dividends-to-price portfolios, we consider the excess returns on portfolios formed on deciles of the distribution of D/P, where D/P is the dividend yield. In all the cases stocks are grouped into 10 deciles.

The first column shows the adjusted R^2 of the model. The subsequent two columns present, respectively, the point vale and the t -statistic for the coefficient on the tracking portfolio. The last two columns report the corresponding evidence for the expected inflation parameter. Due to the presence of a generated regressor, the error term in the equation in not homoskedastic. To correct for the bias in the OLS estimate of the covariance matrix of the estimated coefficients, White standard errors have been used. The sample spans the period from July 1965 to December 1998. Returns are measured in excess of the return on a 30-day Treasury bill.

The first column shows the adjusted R^2 of the model. The subsequent two columns present, respectively, the point vale and the t -statistic for the coefficient on the tracking portfolio. The last two columns report the corresponding evidence for the expected inflation parameter. Due to the presence of a generated regressor, the error term in the equation in not homoskedatic. To correct for the bias in the OLS estimate of the covariance matrix of the estimated coefficients, White standard errors have been used. The sample spans the period from July 1965 to December 1998. Returns are measured in excess of the return on a 30-day Treasury bill.

$$\text{Regression Model: } R_{i,t} = \alpha_i + \delta_i' F_t + \gamma_i(\Phi_t) + \beta_i \text{Infl}_t^e + \varepsilon_{i,t}$$

Portfolio	Adjusted R ²	Tracking Portfolio		Expected Inflation	
		coefficient	t-stat	coefficient	t-stat
Panel A: Size and book-to-market portfolios					
S1/B1	0.931	0.799	9.867	0.436	1.001
S1/B2	0.954	0.802	16.182	-0.579	-1.957
S1/B3	0.962	0.821	21.300	-0.472	-1.808
S1/B4	0.959	0.753	18.547	-0.842	-3.217
S1/B5	0.957	0.805	18.594	-0.797	-2.707
S2/B1	0.955	0.497	9.275	0.899	2.540
S2/B2	0.959	0.603	12.977	0.292	0.973
S2/B3	0.954	0.503	10.990	-0.424	-1.576
S2/B4	0.946	0.446	10.119	-0.068	-0.229
S2/B5	0.953	0.546	11.521	0.428	1.473
S3/B1	0.949	0.419	7.843	0.591	1.755
S3/B2	0.938	0.483	9.499	0.114	0.356
S3/B3	0.925	0.353	7.354	0.018	0.054
S3/B4	0.926	0.367	7.701	-0.230	-0.666
S3/B5	0.918	0.407	7.042	0.422	1.110
S4/B1	0.943	0.228	4.692	0.600	2.090
S4/B2	0.915	0.316	5.894	0.499	1.356
S4/B3	0.915	0.236	4.342	1.005	2.853
S4/B4	0.897	0.270	4.530	0.097	0.230
S4/B5	0.877	0.473	6.902	0.423	1.066
S5/B1	0.932	0.126	2.612	-0.836	-2.400
S5/B2	0.918	0.118	2.341	0.162	0.525
S5/B3	0.859	-0.064	-1.087	0.439	1.067
S5/B4	0.891	0.151	2.696	0.635	1.864
S5/B5	0.809	0.267	3.696	-0.657	-1.321
Panel B: Industry portfolios					
Food	0.746	0.219	2.680	-0.671	-1.295
Mines	0.550	0.530	3.448	1.971	2.023
Oil	0.541	0.189	1.481	0.504	0.626
Clths	0.796	0.822	8.325	-0.459	-0.735
Durbl	0.798	0.358	4.101	-1.776	-3.214
Chems	0.768	0.377	4.079	0.238	0.407
Cnsum	0.754	-0.046	-0.529	-0.411	-0.749
Cnstr	0.888	0.272	3.867	1.292	2.897
Steel	0.650	0.443	3.438	1.391	1.703
FabPr	0.811	0.294	3.717	0.525	1.049
Machn	0.809	0.411	4.633	-0.991	-1.765
Cars	0.623	0.630	4.948	-1.750	-2.170
Trans	0.814	0.328	3.531	0.415	0.706
Utils	0.576	-0.098	-1.061	0.593	1.013
Rtail	0.743	0.565	5.482	-0.779	-1.193
Finan	0.866	-0.091	-1.363	0.633	1.492
Other	0.938	0.188	4.570	0.318	1.223

$$\text{Regression Model: } R_{i,t} = \alpha_i + \delta'_i F_t + \gamma_i(\Phi_t) + \beta_i \text{Inff}_t^e + \varepsilon_{i,t}$$

Portfolio	Adjusted R^2	Tracking portfolio		Expected Inflation	
		coefficient	t-test	coefficient	t-test
Panel C: Cash flows-to-price portfolios (CF/P)					
Decile 1	0.924	0.170	2.900	-0.530	-1.400
Decile 2	0.926	0.174	3.561	-0.523	-1.446
Decile 3	0.913	0.139	2.549	0.207	0.577
Decile 4	0.918	0.093	1.775	0.532	1.535
Decile 5	0.890	0.205	3.577	1.005	2.732
Decile 6	0.898	0.068	1.255	0.228	0.621
Decile 7	0.876	0.154	2.760	0.422	1.135
Decile 8	0.868	0.211	3.583	-0.009	-0.020
Decile 9	0.867	0.299	4.091	-0.237	-0.563
Decile 10	0.889	0.325	4.976	-0.029	-0.064
Panel D: Earnings-to-price portfolios (E/P)					
Decile 1	0.905	0.277	4.314	-0.860	-2.376
Decile 2	0.938	0.100	2.176	0.368	1.345
Decile 3	0.915	0.208	3.962	0.320	0.887
Decile 4	0.894	0.114	1.929	0.396	1.060
Decile 5	0.911	0.157	3.336	0.600	1.841
Decile 6	0.894	0.118	1.774	0.667	1.798
Decile 7	0.889	0.148	2.909	0.305	0.934
Decile 8	0.875	0.163	2.743	-0.175	-0.372
Decile 9	0.886	0.285	4.361	-0.475	-1.307
Decile 10	0.894	0.264	4.347	-0.224	-0.483
Panel E: Dividends-to-price portfolios (D/P)					
Decile 1	0.918	0.179	2.885	0.458	1.162
Decile 2	0.922	0.198	3.955	0.331	0.951
Decile 3	0.929	0.147	2.870	-0.124	-0.374
Decile 4	0.924	0.114	2.053	0.242	0.690
Decile 5	0.908	0.170	3.232	0.435	1.147
Decile 6	0.896	0.160	2.964	0.550	1.605
Decile 7	0.898	0.090	1.717	0.596	1.848
Decile 8	0.873	0.196	3.265	0.521	1.531
Decile 9	0.827	0.167	2.495	-0.149	-0.368
Decile 10	0.683	0.240	2.493	-1.092	-1.731

Bibliography

- [1] Al-Nowaihi, A. and L. Stracca (2002), Non-standard central bank loss functions, skewed risks and certainty equivalence, Working paper no. 129, European Central Bank.
- [2] Amato, J., S. Morris and H.S. Shin (2002), Communication and Monetary Policy, *Oxford Review of Economic Policy*, **18**(4), 495-503.
- [3] Amihud, Y. (1996), Unexpected Inflation and Stock Returns Revisited. Evidence from Israel, *Journal of Money, Credit and Banking*, **28**(1), 22-33.
- [4] Ando, A. and F. Modigliani (1975), Some reflections on describing structures of financial sectors, in Fromm, G. and L.R. Klein (eds), *The Brookings Model: Perspective and Recent Developments*, North-Holland Publishing Co.
- [5] Balduzzi, P. (1995), Stock returns, inflation and the ‘proxy hypothesis’: a new look at the data, *Economic Letters*, **18**, 47-53.
- [6] Barberis, N., M. Huang and J. Santos (2001), Prospect Theory and Asset Pricing, *Quarterly Journal of Economics*, **116**, 1-53.
- [7] Barro, R. and D. Gordon (1983), A Positive Theory of Monetary Policy in a Natural rate Model, *Journal of Political Economy*, **91**(4), 589-610.

- [8] Battenberg, D., J. Enzler and A. Havenner (1975), MINNIE: a small version of the MIT-PENN-SSRC econometric model, Special Studies Papers no.63, Board of Governors of the Federal Reserve System.
- [9] Beeby, M., S.G. Hall and S.B. Henry (2001), Rational Expectations and Near Rational Alternatives: How Best to Form Expectations, ECB Working Paper no.86.
- [10] Benveniste, A., M. Metivier and P. Priouret (1990), *Adaptive Algorithms and Stochastic Approximations*, Springer-Verlag, Berlin Heidelberg.
- [11] Berardi, M. (2007), Heterogeneity and misspecifications in learning, *Journal of Economic Dynamics and Control*, **31**(10), 3203-3227.
- [12] Berardi, M. and J. Duffy (2006), The Value of Central Bank Transparency When Agents are Learning, mimeo.
- [13] Berk, J. , R. Green and V. Naik (1999), Optimal Investment, Growth Options and Security Returns, *The Journal of Finance*, **54**(5), 1533-1608.
- [14] Bernanke, B. S. and I. Mihov (1998), Measuring Monetary Policy, *Quarterly Journal of Economics* **120**(1), 387-422.
- [15] Bertocchi, G. and M. Spagat (1993), Learning, experimentation and monetary policy, *Journal of Monetary Economics*, **32**(1), 169-183.
- [16] Blinder, A.S., M. Ehrmann, M. Fratzscher, J. De Haan and D. Jansen (2008), Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence, NBER Working Paper no. 13932.
- [17] Bordo, M.D and A.J. Schwartz (1999), Monetary Policy Regimes and Economic Performance: the Historical Record, in Taylor J.B. and M.

- Woodford (eds), *Handbook of Macroeconomics*, vol.1A, Elsevier, Amsterdam.
- [18] Boschen, J. F. and L.O. Mills (1995), The Relation between Narrative and Money Market Indicators of Monetary Policy, *Economic Inquiry*, **33**(1), 24-44.
 - [19] Boudoukh, J., and M. Richardson (1993), Stock Returns and Inflation: A Long-Horizon Perspective, *The American Economic Review*, **83**(5), 1346-1355.
 - [20] Boudoukh, J., M. Richardson and R.F. Whitelaw (1994), Industry returns and the Fisher effect, *the Journal of Finance*, **49**(5), 1595-1616.
 - [21] Brainard, W. (1967), Uncertainty and the Effectiveness of Policy, *American Economic Review*, **57**(2), 411-425.
 - [22] Branch, W.A. and G.W. Evans (2006), Intrinsic heterogeneity in expectation formation, *Journal of Economic Theory*, **127**(1), 264-295.
 - [23] Branch, W.A. and B. McGough (2008), Replicator Dynamics in a Cobweb Model with Rationally Heterogeneous Expectations, *Journal of Economic Behavior and Organization*, **65**(2), 224-244.
 - [24] Bray, M.M. (1982), Learning, Estimation, and the Stability of Rational Expectations, *Journal of Economic Theory*, **26**(2), 318-339.
 - [25] Bray, M.M. and N.E. Savin (1986), Rational Expectations Equilibria, Learning and Model Specification, *Econometrica*, **54**(5), 1129-1160.
 - [26] Brennan, M. J., T. Chordia and A. Subrahmanyam (1998), Alternative factor specifications, security characteristics, and the cross-section of expected stock returns, *Journal of Financial Economics*, **49**, 345-373.

- [27] Brock, W.A. and C.H. Hommes (1997), A Rational Route to Randomness, *Econometrica*, **65**(5), 1059-1095.
- [28] Bryant, R., P. Hooper and C. Mann (1993), Evaluating Policy Regimes: New Empirical Research in Empirical Macroeconomics, Brookings Institution, Washington, D.C.
- [29] Buiter, W.H. (2006), How Robust is the New Conventional Wisdom? The Surprising Fragility of the Theoretical Foundations of Inflation Targeting and Central Bank Independence, CEPR Discussion Paper no. 5772.
- [30] Bullard, J. (1997), Learnability Dynamics: An Essay, mimeo.
- [31] Bullard, J. (2006), The Learnability Criterion and Monetary Policy, Federal Reserve Bank of Saint Louis Review, May/June.
- [32] Bullard, J. and I. Cho (2005), Escapist Policy Rules, *Journal of Economic Dynamics and Control*, **29**(11), 1841-1865.
- [33] Bullard, J. and S. Eusepi (2008), When Does Determinacy Imply Expectational Stability?, Federal Reserve Bank of Saint Louis Working Paper 2008-007A.
- [34] Bullard, J. and K. Mitra (2001), Determinacy, Learnability and Monetary Policy Inertia, mimeo.
- [35] Campbell, J.Y., A.W. Lo and A.C. MacKinlay, (1997), *The Econometrics of Financial Markets*, Princeton University Press.
- [36] Cecchetti, S.G., P. Lam and N.C. Mark (1993), The Equity Premium and the Risk Free Rate: Matching the Moments, *Journal of Monetary Economics*, **31**, 21-45.

- [37] Chen, Nai-Fu, R. Roll and S. A. Ross (1986), Economic Forces and the Stock Market, *Journal of Business*, **59**(3), 83-103.
- [38] Chow, G. and A.L. Lin (1971), Best Linear Unbiased Interpolation, Distribution and Extrapolation of Time Series by Related Series, *Review of Economics and Statistics*, **53**(4), 372-375.
- [39] Christiano, L.J., M. Eichenbaum and C.L. Evans (1999), *Monetary Policy Shocks: What Have we Learned and to What End?*, in Taylor J.B. and M. Woodford (eds), *Handbook of Macroeconomics*, vol.1A, Elsevier, Amsterdam.
- [40] Clarida, R., J. Galí and M. Gertler (2000), Monetary policy rules and macroeconomic stability: evidence and some theory, *Quarterly Journal of Economics*, **115**(1), 1661-1707.
- [41] Cochrane, J.H. (2000), *Asset Pricing*, Princeton University Press, Princeton and Oxford.
- [42] Cook, T. and T. Hahn (1989), The effect of changes in the federal funds rate target on market interest rates in the 1970s, *Journal of Monetary Economics*, **24**, 331-351.
- [43] Cuckierman, A. and A.H. Meltzer (1986), A Theory of Ambiguity, Credibility and Inflation under Discretion and Asymmetric Information, *Econometrica*, **54**(5), 1099-1128
- [44] David, A. (1997), Fluctuating confidence in stock markets: implications for returns and volatility, *Journal of Financial and Quantitative Analysis*, **32**(4), 427-462.
- [45] David, A. and P. Veronesi, (1999), Option prices with uncertain fundamentals, mimeo.

- [46] DeCanio, S. (1979), Rational Expectations and Learning from Experience, *Quarterly Journal of Economics*, **92**(4), 47-58.
- [47] Dermine, J. (1985), Inflation, Taxes and Banks' Market Value, *Journal of Business, Finance and Accounting*, **12**(1), 65-73.
- [48] Dermine, J. and F. Lajeri (1999), Unexpected Inflation and Bank Stock Returns: the Case of France 1977-1991, *Journal of Banking and Finance*, **23**(6), 939-953.
- [49] Dewachter, H. and M. Lyrio (2006), Learning, Macroeconomic Dynamics and the Term Structure of Interest Rates, mimeo.
- [50] Dieppe, A., A. Gonzalez Pandiella, S. Hall and A. Willman (2011), The ECB's New Multi-Country Model for the euro area, NMCM - with boundedly rational learning expectations, ECB Working Paper no. 1316.
- [51] Driffill, J. and Z. Rotondi (2003), Monetary Policy and Lexicographic Preference Ordering, mimeo.
- [52] Dumas, B. and B. Solnik (1995), The world price of foreign exchange risk, *Journal of Finance*, **50**(2), 445-479.
- [53] Ellison, M. and N. Valla (2001), Learning, uncertainty and central bank activism in an economy with strategic interactions, *Journal of Monetary Economics*, **48**, 153-171.
- [54] Eusepi, S. (2005), Central Bank Transparency under Model Uncertainty, Federal Reserve Bank of New York, Staff Reports no. 199.
- [55] Eusepi, S. and B. Preston (2007), Central Bank Communication and Expectations Stabilization, NBER working paper no.13259.

- [56] Evans, M.D.D. and K. Lewis (1995), Do expected Shifts in Inflation Affect Estimates of the Long-Run Fisher Relation?, *The Journal of Finance*, **50**(1), 225-253.
- [57] Evans, G.W. (1985), Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models, *Quarterly Journal of Economics*, **100**(4), 1217-1233.
- [58] Evans, G.W. and S. Honkapohja (2002), Expectations and the Stability Problem for Optimal Monetary Policy, mimeo.
- [59] Evans, G.W. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press, Princeton and Oxford.
- [60] Evans, G.W. and S. Honkapohja (2008), Expectations, Learning and Monetary Policy: An Overview of Recent Research, Centre for Dynamic Macroeconomic Analysis, Working Paper Series no. 08/02.
- [61] Evans, G.W., S. Honkapohja and N. Williams (2006), Generalized Stochastic Gradient Learning, mimeo.
- [62] Evans, G.W. and B. McGough (2005), Monetary policy, indeterminacy, and learning, *Journal of Economic Dynamics and Control*, **29**(11), 1809-1840.
- [63] Evans, G.W. and B. McGough (2008), Representations and Sunspot Stability, mimeo.
- [64] Evans, G.W. and G. Ramey (1992), Expectation Calculation and Macroeconomic Dynamics, *American Economic Review*, **82**(1), 207-224.
- [65] Fagan, G. and J. Morgan (2005), *Econometric Models of the Euro-area Central Banks*, Edward Elgar, Cheltenham.

- [66] Fama, E.F. (1981), Stock Returns, Real Activity, Inflation and Money, *The American Economic Review*, **71**(4), 545-565.
- [67] Fama, E.F. and G.W. Schwert (1977), Asset Returns and Inflation, *Journal of Financial Economics*, **5**(1), 115-146.
- [68] Faust, J. and L.E.O. Svensson (1997), Transparency and Credibility: Monetary policy with Unobservable Goals, mimeo
- [69] Faust, J. and L.E.O. Svensson (1999), The Equilibrium Degree of Transparency and Control in Monetary Policy, mimeo
- [70] Ferrero, G. (2007), Monetary policy, learning and the speed of convergence, *Journal of Economic Dynamics and Control*, **31**(9), 3006-3041.
- [71] Ferson, W.E. (1990), Are the Latent Variables in Time-Varying Expected Returns Compensation for Consumption Risk?, *Journal of Finance*, **45**(2), 397-429.
- [72] Ferson, W.E. and C.R. Harvey (1991), The Variation of Economic Risk Premiums, *Journal of Political Economy*, **99**(2), 385-415.
- [73] Ferson, W. E. and C.R. Harvey (1993), An Exploratory Investigation of the Fundamental Determinants of National Equity Market Returns, NBER Working Paper no. 4595.
- [74] Ferson, W. E. and C.R. Harvey (1993), The Risk and Predictability of International Equity Returns, *The Review of Financial Studies*, **6**(3), 527-566.
- [75] Ferson, W. E. and C.R. Harvey (1995), Predictability and Time-Varying Risk in World Equity Markets, in A.H. Chen (ed) *Research in finance*, JAI Press, Elsevier: Amsterdam.

- [76] Ferson, W. E. and R. Schadt (1996), Measuring Fund Strategy and Performance in Changing Economic Conditions, *Journal of Finance*, **51**(2), 425-461.
- [77] Ferson, W.E. and C.R. Harvey (1999), Conditioning Variables and the Cross Section of Stock Returns, *Journal of Finance*, **54**(4), 1325-1360.
- [78] Ferson, W. E., S. Kandel and R. Stambaugh (1987), Tests of Asset Pricing with Time-Varying Expected Risk Premiums and Market Betas, *Journal of Finance*, **42**(2), 201-220.
- [79] Friedman, M. and A. Schwarz (1976), *From Gibson to Fisher*, Explorations in Economic Research, vol. 3, no. 2.
- [80] Frydman, R. (1982), Towards an Understanding of Market Processes: Individual Expectations, Learning, and Convergence to Rational Expectations Equilibrium, *American Economic Review*, **72**(4), 652-668.
- [81] Fukac, M. and A. Pagan (2006), Issues in Adopting DSGE Models for Use in the Policy Process, Czech National Bank Working Paper Series no. 6.
- [82] Garratt, A. and S.G. Hall (1997), E-equilibria and adaptive expectations: Output and inflation in the LBS model, *Journal of Economic Dynamics and Control*, **21**(7), 1149-1171.
- [83] Geske, R. and R. Roll (1983), The Fiscal and Monetary Linkage between Stock Returns and Inflation, *The Journal of Finance*, **38**(1), 1-33.
- [84] Giannitsarou, C. (2003), Heterogeneous Learning, *Review of Economic Dynamics*, **6**(4), 885-890.
- [85] Goodfriend, M. (1986), Monetary Mystique: Secrecy and Central Banking, *Journal of Monetary Economics*, **17**(1), 63-92.

- [86] Goodfriend, M. (1997), Monetary Policy Comes of Age: A 20th Century Odyssey, Federal Reserve Bank of Richmond Economic Quarterly, vol. 83, no. 1.
- [87] Goodfriend, M. (1998), Using the Term Structure of Interest Rates for Monetary Policy, Federal Reserve Bank of Richmond Economic Quarterly, vol. 84, no. 3.
- [88] Groenewold, N., G. O'Rourke and S. Thomas (1997), Stock Returns and Inflation: a Macro Analysis, *Applied Financial Economics*, **7**(2), 127-136.
- [89] Guse, E. (2008), Learning in a Misspecified Multivariate Self-Referential Linear Stochastic Model, *Journal of Economic Dynamics and Control*, **32**(5), 1517-1542.
- [90] Haas, R. and P. Masson (1986), MINIMOD: Specification and Simulation Results, IMF Staff Papers, no. 4.
- [91] Hommes, C.H. and G. Sorger (1998), Consistent Expectations Equilibria, *Macroeconomic Dynamics*, **2**(3), 287-321.
- [92] Honkapohja, S. and K. Mitra (2006), Learning Stability in Economies with Heterogeneous Agents, *Review of Economic Dynamics*, **9**(2), 284-309.
- [93] Jensen, H. (2003), Explaining an Inflation Bias without Using the Word "Surprise", mimeo.
- [94] Karlin, S. and H.M. Taylor (1996), *A First Course in Stochastic Processes*, Academic Press, New York.
- [95] Kaul, G. (1987), Stock Returns and Inflation: the Role of the Monetary Sector, *Journal of Financial Economics*, **18**, 253-276.

- [96] Kaul, G. (1990), Monetary Regimes and the Relation between Stock Returns and Inflationary Expectations, *Journal of Financial and Quantitative Analysis*, **25**(3), 307-321.
- [97] Kent, D.D. and D.A. Marshall (1998), Consumption-Based Modelling of Long-Horizon Returns, Federal Reserve Bank of Chicago, Working Paper 98-18.
- [98] Kessel, R. A. (1956), Inflation-caused wealth redistribution: A Test of a Hypothesis, *American Economic Review*, **46**(1), 128-141.
- [99] Kessel, R.A., and A.A.Alchian (1960), The Meaning and Validity of the Inflation Induced Lags of Wages behind Prices, *American Economic Review*, **50**(1), 43-66.
- [100] Kydland, F. and E. Prescott (1977), Rules rather than Discretion: the Inconsistency of Optimal Rules, *Journal of Political Economy*, **85**(3), 473-492.
- [101] Lamont, O.A. (2001), Economic tracking portfolios, *Journal of Econometrics*, **105**(1), 161-184.
- [102] Lee, Bong-Soo (1992), Causal Relations among Stock Returns, Interest Rates, Real Activity and Inflation, *The Journal of Finance*, **47**(4), 1591-1603.
- [103] Liew, J. and M. Vassallou (2000), Can book-to-market, size and momentum be risk factors that predict economic growth?, *Journal of Financial Economics*, **57**, 221-245.
- [104] Liptser, R.A. and A.N. Shiriyayev (1977), *Statistics of Random Processes*, Springer-Verlag, New York.

- [105] Mailath, G.J. (1998), Do People Play Nash Equilibrium? Lessons From Evolutionary Game Theory, *Journal of Economic Literature*, **36**(3), 1347-1374.
- [106] Marcet, A. and T.J. Sargent (1989), Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information, *Journal of Political Economy*, **97**(6), 1306-1322.
- [107] Marcet, A. and T.J. Sargent (1989), Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models, *Journal of Economic Theory*, **48**(2), 337-368.
- [108] Marcet, A. and T.J. Sargent (1995), Speed of convergence of recursive least squares: learning with autoregressive moving-average perceptions, in Kirman, A. and M. Salmon (eds), *Learning and Rationality in Economics*, Blackwell, Oxford.
- [109] Marshall, D.A. (1992), Inflation and Asset Returns in a Monetary Economy, *The Journal of Finance*, **47**(4), 1315-1342.
- [110] Masson, P. (1989), International Dimensions of Monetary Policy: Coordination Versus Autonomy, Monetary Policy Issues in the 1990s, proceedings of a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming.
- [111] McCallum, B.T. (2007), E-stability vis-à-vis Determinacy Results for a Broad Class of Linear Rational Expectations Models, *Journal of Economic Dynamics and Control*, **31**(4), 1376-1391.
- [112] McCallum, B.T. (2008), Determinacy, Learnability and Plausibility in Monetary Policy Analysis: Additional Results, mimeo.
- [113] McGough, B. (2006), Shocking Escapes, *Economic Journal*, **116**(2), 507-528.

- [114] Mishkin, F.S. (1992), Is the Fisher effect for real?, *Journal of Monetary Economics*, **30**(2), 195-215.
- [115] Mishkin, S. and N.J. Westelius (2006), Inflation Band Targeting and Optimal Inflation Contracts, NBER Working Paper no. 12384.
- [116] Orphanides, A. and D.W. Wilcox (1993), The Opportunistic Approach to Disinflation, Federal Reserve Board, Finance Econ. Discussion Paper no. 24.
- [117] Orphanides, A. and J.C. Williams (2002), Imperfect Knowledge, Inflation Expectations and Monetary Policy, Board of Governors of the Federal Reserve System, Working paper no. 27.
- [118] Orphanides, A. and J.C. Williams (2004), Imperfect Knowledge, Inflation Expectations, and Monetary Policy, in *The Inflation-Targeting Debate*, in B.S. Bernanke and Woodford, M. (eds), University of Chicago Press, Chicago.
- [119] Orphanides, A. and J.C. Williams (2007), Robust Monetary Policy with Imperfect Knowledge, *Journal of Monetary Policy*, **54**(5), 1406-1435.
- [120] Parke, W.R. and G.A. Waters (2006), An Evolutionary Route to Rational Expectations, mimeo.
- [121] Parigi, G. and S. Siviero (2001), An investment-function-based measure of capacity utilization. Potential output and utilised capacity in the Bank of Italy's quarterly model, *Economic Modelling*, **18**(4), 525-550.
- [122] Patelis, A.D. (1997), Stock Return Predictability and the Role of Monetary Policy, *The Journal of Finance*, **52**(5), 1951-1972.
- [123] Pescatori, A. and S. Zaman (2011), Macroeconomic Models, Forecasting, and Policymaking, Economic Commentary 2011-19, Federal Reserve Bank of Cleveland.

- [124] Poole, W. (2001), Expectations, Federal Reserve Bank of St. Louis Review, March/April.
- [125] Ram, R. and D.E. Spencer (1983), Stock returns, real activity, inflation and money: comment, *American Economic Review*, **73**(3), 463-470.
- [126] Rogoff, K. (1985), The Optimal Commitment to an Intermediate Monetary Target, *Quarterly Journal of Economics*, **100**(4), 1169-1189.
- [127] Rogoff, K. (1989), Reputation, Coordination and Policy, in Barro, R.J., *Modern Business Cycle Theory*, Harvard University Press, Harvard.
- [128] Romer, C.D. and D.H. Romer (1989), Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz, in Blanchard, O.J. and S. Fischer (eds), *NBER Macroeconomics Annual*, MIT Press.
- [129] Samuelson, L. (2002), Evolution and Game Theory, *Journal of Economic Perspectives*, **16**(2), 47-66.
- [130] Sargent, T.J. (1993), *Bounded Rationality in Macroeconomics*, Clarendon Press, Oxford.
- [131] Sargent, T.J. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton.
- [132] Sethi, R. and R. Franke (1995), Behavioural heterogeneity under evolutionary pressure: macroeconomic implications of costly optimisation, *Economic Journal*, **105**(May), 583-600.
- [133] Simon, H.A. (1957), *Models of Man: Social and Rational*, Wiley, New York.
- [134] Sims, C.A. (1982), Policy analysis with econometric models, *Brookings Papers on Economic Activity* **1982**(1), 107-164.

- [135] Söderlind, P. (2001), Monetary Policy and the Fisher Effect, *Journal of Policy Modelling*, **23**(5), 491-495.
- [136] Stein, J.C. (1989), Cheap Talk and the Fed: a Theory of Imprecise Policy Announcements, *American Economic Review*, **79**(1), 32-42.
- [137] Stulz, R.M. (1986), Asset Pricing and Expected Inflation, *The Journal of Finance*, **41**(1), 209-223.
- [138] Svensson, L.E.O (1999), Inflation Targeting as a Monetary Rule, *Journal of Monetary Economics*, **43**(3), 607-654.
- [139] Svensson, L.E.O. (2000), Open-Economy Inflation Targeting, *Journal of International Economics*, **50**(1), 155-183.
- [140] Svensson, L.E.O. and M. Woodford (2000), Indicator Variables for Optimal Policy, NBER Working Paper no. 7953.
- [141] Taylor, J.B. and J.C. Williams (2009), Simple and Robust Rules for Monetary Policy, forthcoming in Friedman, B. and M. Woodford (eds), *the Handbook of Monetary Economics*, vol.III, Elsevier, Amsterdam.
- [142] Terlizzese, D. (1994), Il modello econometrico della Banca d'Italia: una versione in scala 1:15, *Ricerche quantitative per la politica economica 1993*, Banca d'Italia.
- [143] Terlizzese, D. (1999), A note on lexicographic ordering and monetary policy, mimeo.
- [144] Thorbecke, W. (1997), On Stock Market Returns and Monetary Policy, *Journal of Finance*, **52**(2), 635-654.
- [145] Vassallou, M. (2001), News related to GDP growth as a risk factor in equity returns, mimeo.

- [146] Veronesi, P. (2000), How does information quality affect stock returns?, *Journal of Finance*, **55**(2), 807-837.
- [147] Veronesi, P. (1999), Stock market overreaction to bad news in good times: a rational expectations equilibrium model, *The Review of Financial Studies*, **12**(5), 975-1007.
- [148] Veronesi, P. and A. David (2000), *Option Prices with Uncertain Fundamentals: Theory and Evidence on the Dynamics of Implied Volatilities*, Chicago University Working Paper.
- [149] Waters, G.A. (2007), Chaos in the Cobweb Model with a New Learning Dynamic, mimeo.
- [150] Weibull, J.W. (1995), *Evolutionary Game Theory*, The MIT Press.
- [151] Wieland, V. (2003), Monetary Policy and Uncertainty about the Natural Unemployment Rate, CEPR Discussion Paper no. 3811.
- [152] Woodford, M. (1990), Learning to Believe in Sunspots, *Econometrica*, **58**(2), 277-307.
- [153] Yared, F. and P. Veronesi (2000), Short and long horizon term and inflation risk premia in the US term structure, mimeo.