Optimal monetary policy in a New Keynesian model with heterogeneous expectations^{*}

Giovanni Di Bartolomeo Sapienza University of Rome giovanni.dibartolomeo@uniroma1.it

Marco Di Pietro Sapienza University of Rome marco.dipietro@uniroma1.it

Bianca Giannini Sapienza University of Rome bianca.giannini@uniroma1.it

2015

Abstract

In a world where expectations are heterogeneous, what is the design of the optimal policy? Are canonical policies robust when heterogeneous expectations are considered or would they be associated with large welfare losses? We aim to answer these questions in a stylized simple New Keynesian model where agents' beliefs are not homogeneous. Assuming that a fraction of agents can form their expectations by some adaptive or extrapolative schemes, we focus on an optimal monetary policy by second-order approximation of the policy objective from

^{*}The authors are grateful to James Bullard, John Driffill, Elton Beqiraj, Jordi Galì, Luisa Lambertini, Leonardo Melosi, Salvatore Nisticò, Carolina Serpieri, Patrizio Tirelli and Andrea Zaghini for useful comments on earlier drafts. We have benefited from comments on various workshops and conferences including the ICMAIF (University of Crete), MMF Research Group Annual Conference (Cardiff University), Dynare conference (National Bank of Belgium, Brussels). The authors also acknowledge financial support by Sapienza University of Rome.

the consumers' utility function. We find that the introduction of bounded rationality in the New Keynesian framework matters. The presence of heterogeneous agents adds a new dimension to the central bank's optimization problem—consumption inequality. Optimal policies must be designed to stabilize the cross-variability of heterogeneous expectations. In fact, as long as different individual consumption plans depend on different expectation paths, a central bank aiming to reduce consumption inequality should minimize the cross-sectional variability of expectations. Moreover, the traditional trade-off between the price dispersion and aggregate consumption variability is also quantitatively affected by heterogeneity.

Jel codes: E52, E58, J51, E24.

Keywords: monetary policy, bounded rationality, heterogeneous expectations.

1 Introduction

The New Keynesian approach has undoubtedly become the workhouse for academic and practical discussions about monetary policy. Optimal monetary policies are usually designed on the rational expectations paradigm, although heterogeneity in the expectations formation mechanism is well documented in both survey data and laboratory experiments.¹ Our paper aims to investigate the impact on optimal monetary policy in a New Keynesian framework where not all agents are rational in forming their expectations. We are interested in determining the degree to which the presence of heterogeneous expectations affects the canonical prescriptions with respect to the conduct of monetary policy and, therefore, their degree of robustness. In our investigation, we focus on optimal monetary policies aiming at maximizing a welfare-based criterion derived from the agent's utility function.

Several approaches have been proposed that include, in a standard model, some modifications that take into account empirical evidence showing that inflation and output forecasts are not rational, at least for some agents. For instance, Krusell and Smith (1996), Mankiw (2000), Amato and Laubach (2003) and Galí *et al.* (2004, 2007) introduce in macroeconomic models a fraction of agents who are not completely rational, as they make their decisions according to some "rule of thumb," i.e., consuming all their income.² Mankiw and Reis (2002, 2007) propose an approach based on sticky information, while Evans and Honkapohja (2001, 2003) focus on learning models.

¹An alternative approach is that of agent-based models (see Cirillo *et al.*, 2011).

 $^{^{2}}$ An alternative interpretation of the assumption is to see it as a short-cut to model limited asset market participation (see, e.g., Bilbiie, 2008).

Along the above lines, a notable number of studies explicitly consider non-homogeneous expectations and address the issue of how this heterogeneity potentially affects aggregate economic dynamics. These studies include Brock and Hommes (1997), Preston (2006), Branch and Evans (2006), Branch and McGough (2009), and Massaro (2013). In particular, by using two alternative approaches to model the heterogeneity of expectations, Branch and McGough (2009) and Massaro (2013) develop parsimonious micro-founded sticky price models that are consistent with boundedly rational individuals.

In Branch and McGough (2009) and Massaro (2013), expectations operators may differ across groups of agents, who are of two kinds—a fraction of them are rational and the remaining are not—and agents' optimal choices are modeled to be consistent with their specific forecasts. By aggregating individual decision rules, Branch and McGough (2009) and Massaro (2013) derive aggregate demand and supply equations of New Keynesian kind embedding bounded rationality.

Formally, and consistent with Preston (2006), Massaro (2013) focuses on long-horizon forecasts and assumes that agents with subjective expectations choose optimal plans while considering forecasts of macroeconomic conditions over an infinite horizon. By contrast, Branch and McGough (2009) assume that individuals with subjective beliefs choose optimal plans that satisfy their individual Euler equations (short-horizon forecasts). As a result, in Massaro (2013), the predicted aggregate dynamics hinge on long-horizon forecasts, while in Branch and McGough (2009), the aggregate dynamics depend on one-period-ahead subjective heterogeneous forecasts.

The primary interest of Branch and McGough (2009) and Massaro (2013) is in the impact of heterogeneity on the existence of sunspot equilibria. At least implicitly, they use the concept of Heterogeneous Expectations Equilibrium (HEE) as "a non-explosive solution to the reduced form expectational difference equation given that expectations are formed heterogeneously" (Branch and McGough, 2004: 1).

Both Branch and McGough (2009) and Massaro (2013) show that heterogeneous expectations can undermine some standard results. However, they do not consider optimal policies; rather, they assume that the central bank sets the interest rate according to a Taylor-type rule. In other words, they investigate the effects of policies based on simple interest rules and show that specifications that are determined under rationality may exhibit explosive or multiple equilibria in the case of bounded rationality.

Gasteiger (2014) extends Branch and McGough (2009) by considering optimal monetary policy. In his context, the optimal path of inflation and the output gap is obtained by minimizing an *ad hoc* loss function and is implemented by a central bank's reaction function. Focusing on determinacy, he shows that when the central bank uses an expectational-based reaction function, it is able to control inflation and the output gap. By contrast, fundamental reaction functions lead to indeterminacy.

Our paper extends Branch and McGough (2009) and Gasteiger (2014) in several ways.

- 1. In a heterogeneous expectations model with short-horizon forecasters, we compute the optimal monetary policy by minimizing a loss function that, differently from Gasteiger (2014), is obtained as a second-order approximation of the utility function of the individuals who populate our economy. Specifically, we derive a second-order approximation of the policy objective from the consumers, who are assumed to have the same specification for the utility functions, while different agents only differ in the way they form their beliefs.
- 2. We show that under bounded rationality, the presence of a fraction of heterogeneous agents adds a new dimension to the central bank's optimization problem, that of consumption inequality. Indeed, the central bank problem now consists of three dimensions: i) minimization of the variability of aggregate consumption, ii) minimization of the cost associated with price dispersion and iii) minimization of the cross-sectional variance of consumption. The latter derives from the fact that rational and boundedly rational agents have different levels of consumption, and it requires the stabilization of the size of the group of boundedly rational agents, but as long as it continues to grow, the emphasis on stabilizing inflation versus output declines. This occurs because the dynamics of price dispersion are more complex than those in the standard case, where the dynamics rely only on inflation. Here, price dispersion also depends on output stabilization.
- 3. We study the optimal stabilization path of the endogenous variables of the model assuming different policy regimes: commitment and discretion. In the former, the monetary authority affects the private sector's expectations by committing to a policy plan. In the latter, in each period, the central bank minimizes the welfare loss taking private agents' forecasts as given. We find that the introduction of bounded rationality in the New Keynesian framework matters, as the existence of a group of non-rational agents, who form their forecasts by adaptive or extrapolative mechanisms, implies that optimal policies must be designed to stabilize the cross-variability of heterogeneous expectations. The rationale of our result is based on the fact that under heterogeneous agents beliefs, the central bank's optimization problem should also consider the minimization of the consumption inequality.

As long as different individual consumption plans depend on different expectations paths, a central bank aiming to reduce the variability in individual consumption should minimize the cross-sectional variability of expectations.

- 4. We study how commitment marginal gains over discretion are affected by the degree of bounded rationality and to what extent simple interest rate rules (i.e., Taylor rules) become suboptimal according to a change in the fraction of non-rational agents who populate our economy.
- 5. We investigate the potential costs associated to the presence of model uncertainty. In detail, we explore the potential dangers of ignoring expectations heterogeneity under a welfare-maximizing perspective. We find that the optimal robust policy requires taking into account expectations heterogeneity, as assuming the presence of bounded rationality is less harmful than ignoring it, i.e., additional welfare losses are smaller.

The rest of the paper is organized as follows. Section 2 presents a general heterogeneous expectations model with n-types of short-run forecasters based on Branch and McGough's (2009) axioms that encompasses the representative rational agent benchmark as a special case. Section 3 derives the welfare criterion consistent with our setup as a second-order approximation of the policy objective, assuming that the steady state is not distorted. Then, by modeling expectations heterogeneity similar to Branch and McGough (2009) and Gasteiger (2014), Section 4 illustrates the properties of optimal policies under bounded rationality, comparing these to the canonical policies and Taylor rules. This section then discusses the implications of agents' heterogeneity for commitment and discretion. Section 5 discusses the dangers of ignoring bounded rationality. Section 6 concludes the paper.

2 The economy

We consider a simple yeoman-farmer economy. The economy is populated by a continuum of mass one of infinitely lived households who produce and consume. Each household produces a differentiated good by using its own labor, and it consumes a "bundle," a composite good composed of all products. We assume price stickiness by assuming that in every period each yeoman-farmer faces an exogenous constant probability of being able to reset its price.

Population is partitioned into n groups of agents. Agents belonging to the same group use the same forecasting rule. We index the size of each group by a_{τ} with $\tau \in \{1, 2..n\}$ and $\sum_{\tau=1}^{n} a_{\tau} = 1$.

Moreover, we set $b_{\tau} = \sum_{k=1}^{\tau} a_k$ and $b_0 = 0$, so that $\int_{b_{\tau-1}}^{b_{\tau}} C_t(i) di = a_{\tau} C_t^{\tau}$ is the total consumption of agents of kind τ and C_t^{τ} is the per-capita consumption. We assume that agents do not learn parameters of their forecasting rule and cannot change their type.

We share the assumption with Preston (2006), Branch and McGough (2009), Massaro (2013), and Gasteiger (2014) that expectational schemes are fixed and that forecasters cannot change their predictor over time. This assumption can be removed and learning introduced. The introduction of a learning process would transform the log-linear description of the economy into a non-linear system (see, e.g., Branch and McGough, 2010). However, we keep the economy that constrains the policymaker linear, as our aim is to generalize the Woodford's LQ approach to optimal policy to the case of heterogeneous agents endowed with different expectation predictors.

Our framework can be interpreted as a study on the optimal prescriptions for monetary policy in an HEE resulting from the convergence of different learning processes based on different specifications of the forecasting model. Berardi (2007), e.g., shows how an HEE a la Branch and McGough (2004) can emerge as a learnable equilibrium when agents underparameterize their model with respect to the common factor representation. Clearly, the learning process can itself be influenced by the conduct of the central bank. Therefore, a limitation of our approach is that it ignores an additional channel for policy to impact predictor selection.

The model, borrowed from Branch and McGough (2009), is described in the following subsections. The axiomatic approach used to model heterogeneous expectations is presented in the next sub-section, while the remaining sub-sections describe the model equations, namely, the private sector's first-order conditions and the log-linearized economy.

2.1 The axiomatic approach to heterogeneous expectations

Heterogeneous expectations are introduced by following the axiomatic approach developed by Branch and McGough (2009). Formally, denoting by \mathcal{E}^{τ} a generic (subjective) expectations operator (i.e., $\mathcal{E}_t^{\tau} x_{t+1}$ is the time t expectation on the value assumed by variable x at t+1 formed by an agent of type τ), we impose the following assumptions: i) each expectation operator, \mathcal{E}^i , fixes observables; ii) all agents' beliefs coincide in the steady state; iii) \mathcal{E}^{τ} is a linear operator; iv) if for all $k \ge 0$, x_{t+k} and $\sum_k \beta^{t+k} x_{t+k}$ are forecasted by agents, then $\mathcal{E}_t^{\tau} \sum_k \beta^{t+k} x_{t+k} =$ $\sum_k \beta^{t+k} \mathcal{E}_t^{\tau} x_{t+k}; v$) \mathcal{E}^{τ} satisfies the law of iterated expectations; vi) if x is a variable forecasted by agents at time t and time t + k such that $\mathcal{E}_t^{\tau} \mathcal{E}_{t+k}^{\tau'} x_{t+k} = \mathcal{E}_t^{\tau} x_{t+k}$, then $\tau \neq \tau'$; and vii) all agents have common expectations on expected differences in limiting wealth.

As discussed by Branch and McGough (2009), assumptions i to v are consistent with rea-

sonable specifications of agent behavior. Assumptions vi) to vii) are necessary for aggregation. The former implies that agents' forecasts satisfy the law of iterated expectations at an aggregate level (i.e., it imposes a particular structure on higher-order beliefs). The axiom that agents agree on limiting wealth distributions does not allow for wealth distribution dynamics that otherwise affect the formulation of forecasts based on expectation type, thus causing a problem for aggregation. This allows us to remain close in form to the homogeneous case.³ By using these assumptions, in fact, we can define the aggregate expectations as a weighted average of group expectations, i.e., $\mathcal{E}_t x_t = \int_0^1 \mathcal{E}_t^i x_t di = \sum_{\tau=1}^n a_\tau \mathcal{E}_t^\tau x_t$, where \mathcal{E}^i is the expectation of agent *i*.

2.2 The model

Each household i produces, as a monopolist, its own differentiated product and directly purchases a composite good seeking to maximize the expected value of the following utility function:

$$\mathcal{E}_0^{\tau} \sum_{t=0}^{\infty} \beta^t \left[u\left(C_t^i \right) - \nu(Y_t(i)) \right] \tag{1}$$

where \mathcal{E}^{τ} denotes the generic (subjective) expectation operator of agent *i*, who belongs to group τ , and $\beta \in (0, 1)$ is the discount factor. The terms $u(C_t^i)$ and $\nu(Y_t(i))$ indicate the utility from consuming the composite good (C_t^i) and the disutility from producing the differentiated product $(Y_t(i))$, respectively.⁴

In each period, a number of randomly selected agents are allowed to change their prices. Each firm may reset its price only with the probability $1 - \xi_p$ in any given period and independent of the time elapsed since the last adjustment occurred.

Because optimal prices depend on the expectations regarding future marginal costs, they differ between agents forming expectations in different ways. Moreover, the Calvo lottery implies heterogeneity within each type of agent, as only a fraction of them reset their prices. Acting as price setters, individual agents face the risk associated with the Calvo lottery. Following a common procedure among heterogeneous agent models to somehow limit heterogeneity between types, it is assumed that agents are engaged in a form of risk sharing to protect themselves from the risk of Calvo price setting. A benevolent financial regulator collects all income and then

 $^{^3\}mathrm{For}$ further details, see Branch and McGough (2009).

⁴There exists a continuum of goods represented by an interval [0,1]; $C_t^i \equiv \left(\int_0^1 C_t^i(j)^{\frac{\varepsilon-1}{\varepsilon}} dj\right)^{\frac{\varepsilon}{\varepsilon-1}}$ is then a Dixit-Stiglitz consumption basket, and $C_t^i(j)$ is the quantity of good j consumed by household i in period t. The consumer price index is defined as $P_t \equiv \left[\int_0^1 P_t(i)^{1-\varepsilon} di\right]^{\frac{1}{1-\varepsilon}}$.

redistributes to each type of household the average income of that agent's type.⁵

It follows that individuals are fully insured against the risk associated with the possibility that they will not be able to adjust prices. Insurances are within agents belonging to the same kind. It follows that since agents belonging to different groups have different expectations, the insurance contracts imply different amounts of expected real income. Formally, the real income of an agent of the kind τ is $\Omega_t^{\tau} = \frac{1}{a_{\tau}} \int_{b_{\tau-1}}^{b_{\tau}} \frac{P_t(i)Y_t(i)}{P_t} di$. Aggregating by kind, we obtain the real output of the economy: $Y_t = \sum_{\tau=1}^n a_{\tau} \Omega_t^{\tau}$.

Because of the insurance mechanism, the real budget constraint of household i of kind τ is:

$$C_t^i + B_t^i = \frac{1 + i_{t-1}}{1 + \pi_t} B_{t-1}^i + \Omega_t^\tau$$
(2)

where $B_t^i P_t$ is the quantity of one period nominal riskless discount bonds purchased in period tand maturing in t + 1 held by agent i. Each bond pays one unit of money at maturity and its price is $Q_t = (1 + i_t)^{-1}$. The term $1 + i_t$ indicates the gross nominal interest rate on a riskless one period bond purchase in period t, and $1 + \pi_t$ defines the gross inflation rate.

Finally, the existence of fully enforceable contracts requires the agents to behave as if they will receive their full marginal revenue when producing more by assuming that agents should choose price and output as if they faced their perceived trade-off. In fact, each agent's income is independent of his efforts because of the presence of the insurance company. Thus, without enforceable contracts, any agent would choose an effort equal to zero due to free-riding behavior.

Each household must decide its optimal consumption (saving) plan, i.e., allocating its consumption expenditures among the different goods and choosing the optimal price (i.e., effort) if selected in a Calvo lottery. The optimal consumption plan of household i is obtained by the maximization of (1) subject to (2) and a solvency constraint. The optimal plan is the same among agents belonging to the same type, and it can be obtained by a simple variational argument. As a result, agents of each type τ make choices about consumption with respect to the intertemporal Euler equation:

$$\frac{1}{1+i_t} = \beta \mathcal{E}_t^{\tau} \left[\frac{P_t}{P_{t+1}} \frac{u_C \left(C_{t+1}^{\tau} \right)}{u_C \left(C_t^{\tau} \right)} \right]$$
(3)

The optimal allocation of household consumption expenditures among different goods, $C_t(j)$,

⁵Among others, the same risk-sharing mechanism is used by Kocherlakota (1996), Shi (1999) and Mankiw and Reis (2007). It is entirely standard in models with heterogeneous agents. Alternatives are also discussed by Branch and McGough (2009).

requires that the consumption index C_t is maximized for any level of expenditure $\int_0^1 P_t(i) C_t(j) dj$.

Each producer j belonging to type τ chooses the price (P_t^j) solving:

$$\max_{P_{t}^{j}} \mathcal{E}_{0}^{\tau} \sum_{i=0}^{\infty} \left(\beta \xi_{p}\right)^{i} \left[(1-T) \lambda_{t+i}(P_{t}^{j}) Y_{t+i}(j) - \nu \left(Y_{t+i}(j)\right) \right]$$
(4)

subject to the demand for good j.⁶ The first term of the sum (4) is the marginal utility of additional nominal income, which can be interpreted as the contribution to utility derived by sales revenues. The second term is the production cost in terms of effort. A production subsidy T is introduced to eliminate distortions in the steady state.

By substituting the demand function into the household's objective function and calculating with respect to P_t^j , we obtain:

$$\mathcal{E}_{0}^{\tau} \sum_{i=0}^{\infty} \left(\beta \xi_{p}\right)^{i} \left[\frac{\partial u\left(C_{t+i}^{\tau}\right)}{\partial C_{t+i}^{\tau}} \left(\frac{P_{t}^{j}}{P_{t+i}}\right)^{-\varepsilon} Y_{t+i} - \left(\frac{P_{t}^{j}}{P_{t+i}}\right)^{-\varepsilon-1} Y_{t+i} \frac{\partial \nu\left(Y_{t+i}\left(j\right)\right)}{\partial Y_{t+i}\left(j\right)} \right] = 0 \quad (5)$$

where we used the fact that $T = -(\varepsilon - 1)^{-1}$ to implement the optimal steady state and $\lambda_t^j P_t = \partial u(C_t^j) / \partial C_t^j$.

2.3 The log-linear economy

By log-linearization of (3) and after some substitutions, we obtain

$$c_t^{\tau} = \mathcal{E}_t^{\tau} c_{t+1}^{\tau} - \sigma (i_t - \mathcal{E}_t^{\tau} \pi_{t+1}) \tag{6}$$

where $\sigma^{-1} \equiv -\bar{C} \ u_{CC}/u_C > 0$ is the inverse of the intertemporal elasticity of substitution of consumption, i.e., the coefficient of relative risk aversion. We define the log deviations of the variables from their steady state values with lower case letters, i.e., $c_t^{\tau} = \log(C_t^{\tau}/\bar{C})$.

By using the budget constraint, $c_t^{\tau} = \omega_t^{\tau} \equiv \left(\beta^{-1}\log(B_{t-1}^{\tau}) - \log(B_t^{\tau})\right)/\bar{Y} + \log(\Omega_t^{\tau}/\bar{\Omega})$,⁷ the log-linearized consumer Euler equation (6) is then rewritten in terms of wealth, provided that the axioms *i*) to *vii*) are satisfied by the expectational operator \mathcal{E}_t^{τ} . Using the bond market clearing condition, $\sum_{\tau=1}^n a_\tau \log(B_t^{\tau}) = 0$ and aggregating (6), it is possible to derive an IS curve that is similar to the relation derived in the standard New Keynesian framework with the exception

⁶By solving the intratemporal goods allocation problem for each household type τ and by aggregating and using the bond market clearing condition, a standard demand for good j can be derived: $Y_t(j) = \left(P_t^j/P_t\right)^{-\varepsilon} Y_t$.

⁷The relationship is obtained by exploiting the fact that $B^{\mathcal{R}} = B^{\mathcal{B}} = 0$. See Branch and McGough (2009: 1040).

of the conditional expectation operator, which is substituted by a convex combination of the heterogeneous expectation operators of the two types of agents.

$$y_t = \mathcal{E}_t y_{t+1} - \sigma \left(i_t - \mathcal{E}_t \pi_{t+1} \right) \tag{7}$$

where $\mathcal{E}_t y_{t+1} = \sum_{\tau=1}^n a_\tau \mathcal{E}_t^\tau y_t$ and $\mathcal{E}_t \pi_{t+1} = \sum_{\tau=1}^n a_\tau \mathcal{E}_t^\tau \pi_t$. Equation (7) is similar to the relation derived in the standard New Keynesian framework with the exception of the conditional expectation operator, which is substituted by a convex combination of the heterogeneous expectation operators of the two types of agents.

Regarding the supply side, by making use of assumptions *iii*) to *vi*), the log-linear version of the optimal price equation of agent j belonging to type τ (5) can be rewritten as:

$$\log\left(\frac{P_t^j}{P_t}\right) = \xi_p \beta \mathcal{E}_t^{\tau} \pi_{t+1} + \left(1 - \beta \xi_p\right) \left[\frac{\sigma^{-1}}{1 + \eta \varepsilon} \omega_t^{\tau} + \frac{\eta}{1 + \eta \varepsilon} y_t\right] + \beta \xi_p \mathcal{E}_t^{\tau} \log\left(\frac{P_{t+1}^j}{P_{t+1}}\right) \tag{8}$$

where $\eta = \nu_{YY} \bar{Y} / \nu_Y$. It is worth noting that each type of producer sets the optimal price based on its specific kind of expectations.

Aggregating (8) for different types of agents and combining the definition of the aggregate price dynamics yields an AS curve similar to the relation derived in the standard New Keynesian framework with the same expectation operator used for (7):

$$\pi_t = \beta \mathcal{E}_t \pi_{t+1} + \frac{\left(1 - \xi_p\right) \left(1 - \beta \xi_p\right) \left(\eta + \sigma^{-1}\right)}{\xi_p \left(1 + \varepsilon \eta\right)} y_t + e_t \tag{9}$$

Note that (9) has been augmented by a supply disturbance, e_t .

3 Welfare criterion

Our aim is to study the optimal HEE. Like Gasteiger (2014: 1539-1541), by optimal, we mean the central bank is committed to the economic well-being of the individuals. The "paternalistic" central bank maximizes the utility of individuals subject to the correct model (understanding how agents form their expectations). As said, we differ from Gasteiger (2014) in one major respect: we consider that the central bank minimizes a loss derived by a second-order approximation of the utility function of the individuals instead of an *ad hoc* loss.

It is worth noting that there is an extensive debate in heterogeneous beliefs models about whether a social planner should be paternalistic or respect the beliefs of the agent, i.e., whether welfare should be measured as ex ante or ex post welfare. Our formulation would correspond to the latter, as in Gasteiger (2014).

Formally, we compute a quadratic Taylor series approximation of utility for household i. The first term of (1) is approximated as

$$\tilde{u}(C_t^i) = \bar{C}u_C\left(c_t(i) + \frac{1 - \sigma^{-1}}{2}c_t^2(i)\right) + t.i.p. + \mathcal{O}\left(\left\|\xi^3\right\|\right)$$
(10)

Note that in the steady state, $\bar{C}^i = \bar{C}$ for any *i* by assumption; the term $\mathcal{O}(||\xi^3||)$ indicates the terms of order greater than two, and *t.i.p.* collects the terms independent of policy. Integrating (10) over *i*, we obtain

$$\int_{0}^{1} \tilde{u}\left(C_{t}^{i}\right) di = \bar{C}u_{C}\left\{c_{t} - \frac{1}{2}var_{i}c_{t}(i) + \frac{1 - \sigma^{-1}}{2}\left[c_{t}^{2} + var_{i}\left(c_{t}(i)\right)\right]\right\} + t.i.p. + \mathcal{O}\left(\left\|\xi^{3}\right\|\right) \quad (11)$$

where we use the relation $var_i(c_t(i)) = E_i c_t (i)^2 - (E_i c_t (i))^2$ and the fact that up to a secondorder $c_t = E_i c_t (i) + \frac{1}{2} var_i c_t (i)$.

Regarding effort, it is noted that each agent potentially supplies a different quantity of output. This depends on the agent's type and on whether he is or is not extracted in a Calvo lottery. A second-order approximation of the second term of (1) leads to

$$\tilde{\nu}(Y_t(i))di = u_N \bar{N}\left(y_t(i) + \frac{1+\eta}{2}y_t^2(i)\right) + t.i.p. + \mathcal{O}\left(\left\|\xi^3\right\|\right)$$
(12)

and, after integration, we obtain

$$\int_{0}^{1} \tilde{\nu}(Y_{t}(i)) di = u_{N} \bar{N} \left(y_{t} - var_{i}(y_{t}(i)) + \frac{1+\eta}{2} \left([E_{i}y_{t}(i)]^{2} + var_{i}(y_{t}(i)) \right) \right) + t.i.p. + \mathcal{O} \left(\left\| \xi^{3} \right\| \right)$$
(13)

where $\int_0^1 \bar{N}(i) = \bar{N}$, as the price dispersion is zero in the steady state. Then, by considering $Y(i) = (P_t(i)/P_t)^{-\varepsilon} Y_t$, the cross-sectional variance of $y_t(i)$ can be expressed as $var_i(y_t(i)) = \varepsilon^2 var_i(p_t(i))$. Thus, in the non-distorted steady state, where the equality $\bar{C}u_C/u_N = \bar{N}$ holds, the above expression is rewritten as

$$\int_0^1 \tilde{\nu}(Y_t(i)) di = \bar{C}u_C \left(y_t + \frac{\varepsilon^2 \eta}{2} var_i(p_t(i)) + \frac{1+\eta}{2} y_t^2 \right) + t.i.p. + \mathcal{O}\left(\left\| \xi^3 \right\| \right)$$
(14)

Combining (11) and (14), the approximated intertemporal utility can be expressed as

$$\sum_{t=0}^{\infty} \beta^{t} \frac{\mathcal{U}_{t}}{\bar{C}u_{C}} = -\sum_{t=0}^{\infty} \beta^{t} \left[L_{t} + t.i.p. + \mathcal{O}\left(\left\| \xi^{3} \right\| \right) \right]$$
(15)

where the instantaneous loss is

$$L_t = \frac{1}{2} \left[\left(\eta + \frac{1}{\sigma} \right) y_t^2 + \left(\varepsilon^2 \eta \right) var_i(p_t(i)) + \frac{1}{\sigma} var_i(c_t(i)) \right]$$
(16)

The welfare-based loss (16), as in the textbook case, depends on two components related to the costs associated with aggregate-consumption variability (y_t^2) and price dispersion $(var_i(p_t(i)))$. Moreover, differently from the standard case, an additional term $(var_i(c_t(i)))$ in (16) captures the cost linked to inequality in the consumption across types. This cost is generated by the different forecast paths, and it vanishes if expectations are not heterogeneous. Note that consumption variability can be expressed as $var_i(c_t(i)) = \sum_{\tau=1}^n a_\tau \left[(c_t^\tau)^2 - c_t^2 \right].$

It is worth noting that price dispersion, Δ_t , evolves according to

$$\Delta_t = \xi_p \Delta_{t-1} + \frac{1}{1 - \xi_p} \sum_{\tau=1}^n a_\tau \left(\log p_t^\tau - P_{t-1} \right)^2 - \pi_t^2 \tag{17}$$

The dispersion of prices (17) does not depend solely on current inflation, unless agents are homogeneous (n = 1).⁸ It has a complex structure that hinges on the assumptions about the price setting of the different agents' groups, which, in turn, depends on how they have formed their expectations in the past.

To stress the impact of beliefs' heterogeneity on welfare compared to the standard textbook case, equation (16) can be written as the sum of three components:

$$L_{t} = L_{0} + \underbrace{\frac{\varepsilon^{2} \eta}{2} \left[var_{i}(p_{t}(i)) - var_{i}^{0}(p_{t}(i)) \right]}_{L_{1}} + \underbrace{\frac{1}{2\sigma} var_{i}(c_{t}(i))}_{L_{2}}$$
(18)

where $L_0 = \frac{1}{2} \left[\left(\eta + \frac{1}{\sigma} \right) y_t^2 + \varepsilon^2 \eta var_i^0(p_t(i)) \right]$ is the loss the policymaker faces when all the agents are rational (the term $var_i^0(p_t(i))$ denotes the price dispersion in this homogenous case); L_1 is the additional cost implied by the larger price dispersion associated with the presence of heterogeneity;⁹ and L_2 is the costs stemming from consumption inequality, which is absent in

⁸ In the case of rational agents, the price dispersion is $\frac{\xi_p}{(1-\beta\xi_p)(1-\xi_p)}\pi_t^2$. ⁹Note that $\left[var_i(p_t(i)) - var_i^0(p_t(i))\right] > 0$, i.e., the price dispersion is larger when agents' beliefs are heterogeneous.

the homogeneous cases.

4 Optimal policies under bounded rationality

4.1 Functional forms and calibration

Following Branch and McGough (2009) and Gasteiger (2014), we assume that our economy is only populated by two kind of agents, who differ in the way they form their expectations (i.e., n = 2). Specifically, a fraction α of them (rational households) have rational expectations; the remaining fraction (non-rational households) form expectations according to a mechanism of bounded rationality. All non-rational households use the same predictor. The two kinds of households are indexed by \mathcal{R} and \mathcal{B} , which refer to rational and boundedly rational households, respectively. Their shares are captured by α , so that $a_{\mathcal{R}} = \alpha$ and $a_{\mathcal{B}} = 1 - \alpha$. Like Branch and McGough (2009) and Gasteiger (2014), we assume that the share α is fixed.

Accordingly, we assume that $\mathcal{E}_t^{\mathcal{R}} x_{t+1} = E_t x_{t+1}$, i.e., rational agents have one-step-ahead perfect foresight on economic variables. Non-rational individuals, instead, form their beliefs on the basis of a simple perceived linear law of motion, i.e., $x_t = \theta x_{t-1}$, where θ is defined as the adaption operator. It follows that $\mathcal{E}_t^{\mathcal{B}} x_t = \theta x_{t-1}$, and applying the law of iterated expectations, we obtain $\mathcal{E}_t^{\mathcal{B}} x_{t+1} = \theta^2 x_{t-1}$. The operator $\mathcal{E}^{\mathcal{B}}$ is a form of adaptive ($\theta < 1$) or extrapolative ($\theta > 1$) expectations. We refer to $\theta = 1$ as the case of naive expectations.¹⁰

The RE model associated with the HEE can then be written as

$$y_t = \alpha E_t y_{t+1} + (1-\alpha)\beta \theta^2 y_{t-1} - \sigma \left[i_t - \alpha E_t \pi_{t+1} - (1-\alpha)\theta^2 \pi_{t-1} \right]$$
(19)

$$\pi_t = \alpha \beta E_t \pi_{t+1} + (1-\alpha) \beta \theta^2 \pi_{t-1} + \kappa y_t + e_t$$
(20)

$$c_t^{\tau} = \mathcal{E}_t^{\tau} c_{t+1}^{\tau} - \sigma(i_t - \mathcal{E}_t^{\tau} \pi_{t+1}) \quad \tau \in \{\mathcal{R}, \mathcal{B}\}$$
(21)

which is the framework used by Branch and McGough (2009) and Gasteiger (2014).

The variances in expression (16) consistent with (19)-(20) are

$$var_i(\log p_t(i)) = \delta \pi_t^2 + \frac{\delta \xi_p \left(1 - \alpha\right)}{\alpha} \left[\pi_t - \beta \theta^2 \pi_{t-1} - \kappa \left(\frac{c_t^{\mathcal{B}} + \eta \sigma y_t}{1 + \eta \sigma}\right)\right]^2$$
(22)

$$var_i(c_t(i)) = \alpha \left(1 - \alpha\right) \left(c_t^{\mathcal{R}} - c_t^{\mathcal{B}}\right)^2$$
(23)

¹⁰See Branch and McGough (2009) for details. See also Pesaran (1987), Brock and Hommes (1997, 1998), Branch and McGough (2005) for empirical support and further considerations.

where $\delta = \frac{\xi_p}{(1-\beta\xi_p)(1-\xi_p)}$ and $\kappa = \frac{(1-\xi_p)(1-\beta\xi_p)(\eta+\sigma^{-1})}{\xi_p(1+\varepsilon\eta)}$.

As discussed before, under heterogeneity of beliefs, the dispersion of prices (22) has a complex structure. It positively depends on the proportion of boundedly rational agents. For a given level of y_t , it increases in the fraction of agents who form their expectations in a non-rational way. Cross-variability of consumption is captured by an index of consumption inequality. This cost is not linear in the degree of bounded rationality. Rather, it has a peak when $\alpha = 0.5$, because in this case, *ceteris paribus*, the distribution of agents exhibits the highest dispersion (see (23)). It is noted that inequality is increasing in the variability of expectations across types as the different beliefs drive different choices.

We calibrate the model to the U.S. economy. The time unit is one quarter. The calibrations of the structural parameters are chosen to equal those estimated or calibrated by Rotemberg and Woodford (1997) by using the structural vector auto-regression (SVAR) methodology and microeconomic evidence. We assume that the subjective discount rate β is 0.99, such that $(\beta^{-1} - 1)$ equals the long-run average real interest rate. In the goods market, the intratemporal elasticity of substitution between the differentiated goods (price elasticity of demand) is set equal to 7.84, thus implying a markup of 15%. The parameter ξ_p , which represents the frequency of price adjustment, is set at 0.66, and thus, prices are fixed, on average, for three quarters. Finally, the elasticity of the marginal disutility of producing output with respect to an increase in output, η , is set to 0.47 based on data regarding labor costs.¹¹ Our calibration is summarized in Table 1.

Table 1 – Baseline calibration $\beta = 0.99$ $\sigma^{-1} = 0.16$ $\varepsilon = 7.84$ $\eta = 0.47$ $\xi_p = 0.66$

Regarding the parameters characterizing heterogeneity (α, θ) , we consider different calibrations. We assume a baseline value where α is equal to 0.7, thus implying that 30% of households form their expectations using a mechanism of bounded rationality, whereas 70% of the households are rational. We explore the effects of α for a range $\alpha \in [0.5, 1]$.¹² The model for $\alpha = 1$ clearly represents the homogeneous-rational agents standard New Keynesian model.

The baseline value for the adaption parameter θ is set equal to one, which implies that boundedly rational households recognize, with a one-period lag, changes in inflation and the output gap (naive expectations). Again, we consider different specifications to test the robustness

¹¹See Rotemberg and Woodford (1997) for details.

 $^{^{12}}$ We explored the model properties and robustness results for the whole existence field of α . Results are available upon request.

of our results, and we report numerical simulations for a range between 0.9 and 1.1. It is noted that different values of θ imply different "rules of thumb" for non-rational households when setting their expectations. Lower θ values involve mean-reverting strategies according to which deviations from the average are expected to revert it. For example, if inflation is below the average, increases in inflation are expected. Higher θ values entail trend-following behaviors, and thus, deviations from the average are expected to be confirmed, such that when the current inflation trend is upward (above the average), the expectation is that the inflation will continue to follow that trend.

4.2 Optimal monetary policy and expectations stabilization

Optimal policy here consists in minimizing the quadratic loss function (16) derived from a second order approximation of the agents' preferences, constrained by the linear system (19)-(21). In (16), $var_i(p_t(i))$ and $var_i(c_t(i))$ are consistently defined by (23) and (22).¹³ In solving this LQ minimization problem, we consider two possible policy regimes: discretion and timeless commitment.

The two policy regimes differ in their ability to internalize the private sector expectations.

- 1. In discretion, the central bank minimizes the loss period-by-period, considering private expectations as already formed, since any promises made in the past by the policy-maker do not constrain current decisions. In our setup, a closed form for the solution of this problem cannot be derived. We solve it by using the codes developed by Soderlind (1999).
- 2. In timeless commitment, the government minimizes the loss at some initial point and commits to a time-contingent strategy that influences expectations. The "optimal" plan is derived by ignoring the initial optimality conditions (when expectations were already formed) and considering that optimal policy solves an optimization problem "from some date forward as being optimal from a timeless perspective, rather than from the perspective of the particular time at which the policy is actually chosen" (Woodford, 2011: 744).

Like, e.g., Benigno and Woodford (2004), Steinsson (2003), and Ravenna and Walsh (2011), we define the optimal policy as the first-order conditions that solve the problem discussed above in the different regimes. We do not discuss the implementation of optimal policy by an interest rate rule. In an HEE with *ad hoc* welfare functions, Gasteiger (2014) shows that fundamentals-based

¹³The problem is formally illustrated in the Appendix.

interest rate rules designed to implement optimal policy always lead to indeterminacy, whereas expectations-based reaction functions do not. He thus confirms Evans and Honkapohja's (2003) results in a homogenous agents setting. We implicitly assume that expectations-based reaction functions are used.¹⁴

We restrict our attention to the comparison between discretion and a timeless perspective. However, it should be noted that the timeless-perspective policy is not necessarily the timeinvariant policy that minimizes the policymaker's objective. On average, the timeless-perspective policy performs worse than the time-invariant policy that minimizes the unconditional expected value of the policy objective (optimal unconditional continuation policy). In general, averaging across all initial conditions can be quite important (see, e.g., Blake, 2001; Jensen and McCallum, 2002, 2010).

The introduction of heterogeneity in expectation formation affects the optimal stabilization policies. Differently from the standard case (homogeneous rational agents), optimal policies now tend to exhibit fluctuations before returning to the steady state. These oscillations are amplified by increases in the adaption parameter.¹⁵ Figure 1 plots the optimal path of the nominal interest rate after a cost-push shock when $\alpha = 0.7$. Optimal policies are computed assuming timeless perspective commitment.

In our baseline ($\alpha = 0.7$ and $\theta = 1$), the path of the interest rate is characterized by oscillations that are amplified for larger values of θ . The rationale of the policy is that the policymaker attempts to reduce the forecast error of boundedly rational agents and, therefore, the crosssectional dispersion in consumption that affects welfare.¹⁶

To understand the intuition, assume that at time t = 1, a shock hits the economy. The policymaker then raises the interest rate to compress inflationary pressure, thereby creating a small deflation. Thus, the fraction of boundedly rational agents form their expectations, assuming that inflation is equal to its previous value. However, their belief is biased. Consequently, they overestimate the inflation, as it was brought down by the monetary contraction pursued by the central banker. In the next period, aggregate inflation exhibits a slight recovery, but again, the inflation forecasts of the non-rational agents are wrong, as they are now underestimating the inflation level. The policymaker must now adjust the nominal interest rate up and down to

 $^{^{14}}$ It is worth noting that a discussion about implementation in our context is more complex, since the first-order conditions simultaneously determine the path for the output gap, inflation and interest rate, as welfare also depends on the consumption distribution that is determined by the path of the interest rate.

¹⁵Note that θ magnifies the effects of the behavior of boundedly rational agents. By contrast, for very low adaption parameter values, the model converges to a purely forward-looking framework.

 $^{^{16}}$ Recall that cross-sectional dispersion in consumption depends on the difference in expectations between rational and boundedly rational agents.

stabilize inflation expectations.



Figure 1 - Nominal interest rate IRF to a cost-push shock for several values of θ (commitment regime).

Figure 2 illustrates the IRFs for inflation and for the output gap when the economy is hit by a cost-push shock. We compare the dynamics obtained from our baseline framework (solid lines) to those stemming from the standard model where all the agents are rational, i.e., $\alpha = 1$ (dashed lines). The figure reports the dynamics under both commitment and discretion in the left and right panels, respectively.



Figure 2 - Optimal stabilization policy for both regimes (cost-push shock).

The figure shows how bounded rationality affects the central bank's optimal choice facing the inflation-output trade-off implied by the supply shock. Optimal policies exhibit some inertia when heterogeneity is introduced, as bounded agents react to the policy with some lags. The results in Figure 2 are robust for different calibrations regarding the share of boundedly rational agents. Within a reasonable range, we obtain IRFs that are qualitatively similar to those represented. Regarding the adaption parameter, similar results are also obtained for different calibrations. However, for large values of θ , the oscillations in the interest rate are very large and translate to inflation and the output gap, which are also characterized by an oscillating dynamic.¹⁷

 $^{^{17}\}mathrm{Results}$ are available upon request.

4.3 Gains from commitment and bounded rationality

We now examine how heterogeneous beliefs affect the relative gains of commitment over discretion. Our results are summarized in Table 2. We consider several combinations of α and θ .¹⁸ Different shares of α are reported in the columns, whereas different adaption parameters θ are displayed in the panels (labeled (a), (b), and (c)). The rows indicate the welfare losses under discretion and commitment and the percentage loss of the former over the latter. The table illustrates several results.

	Share	nare of rational agents (α)			
(a) Adaptive expectations		0.9	0.7	0.5	
(1) Commitment		140.1	223.8	351.6	
(2) Discretion		177.2	275.3	404.9	
(3) Difference $(\%)$		26.4	23.0	15.2	
	Share	Share of rational agents (α)			
(b) Naive expectations		0.9	0.7	0.5	
(1) Commitment		145.4	248.4	409.3	
(2) Discretion		185.5	314.8	483.4	
(3) Difference $(\%)$		27.5	26.72	18.09	
	Share	are of rational agents (α)			
(c) Extrapolative expectations		0.9	0.7	0.5	
(1) Commitment		152.5	281.6	492.9	
(2) Discretion		196.1	373.5	606.4	

Table 2 – Welfare loss and commitment gains

First, a larger share of non-rational agents always involves higher welfare losses in each policy regime. Everything equal, losses increase as α becomes smaller, since the effects of the shocks become more persistent.

28.5

32.6

23.0

Second, parameter θ also affects welfare. Lower levels of θ lead to smaller losses. This result hinges on the fact that by considering adaptive expectations schemes, one removes unit root beliefs from the system, implying lower volatility.

(3) Difference (%)

 $^{^{18}{\}rm Other}$ parameters are calibrated as indicated in Table 1.

Finally, let us now compare the outcomes of the two policy regimes. Commitment always leads to welfare gains compared to discretion (row (1) vs. row (2) in each panel). The best absolute performance of commitment is expected.

By contrast, the relative gains of commitment (row (3)) are affected by the fraction of boundedly rational agents $(1 - \alpha)$ and the way they form their expectations (θ).

- 1. If the proportion of non-rational agents is larger than 30% (i.e., $\alpha \leq 0.7$), a decrease of α always reduces the gains of commitment. For instance, the relative gain is 23% when the share is 30%; it reduces to 15.2% when the share is 50%.¹⁹
- 2. If the share of boundedly rational agents is smaller than 30% (i.e., α ∈ [0.7, 1]), commitment gains depend on the adaption parameter. Two cases are possible. *i*) When the expectations are either adaptive or naive, an increase in the fraction of non-rational agents reduces the gains of commitment as before.²⁰ *ii*) When expectations are extrapolative, an increase in this share instead increases the relative gains of commitment.²¹

The rationale of our last result can be found by examining the effects of bounded rationality on inflation persistence and price dispersion. These are, in fact, the main factors that explain the marginal gains of commitment over discretion. In general, we can distinguish two effects:

- 1. *Persistence effect.* Differently from discretion, rational expectations under commitment are formed after the policy is set. By influencing these expectations, the policymaker is able to eliminate the so-called stabilization bias. Thus, the benefit of commitment tends to vanish as the backward-looking component becomes predominant with respect to the forward one (e.g., Steinsson, 2003).
- 2. Dispersion effect. Compared to discretion, commitment requires stabilizing inflation more than output variability; thus, it performs relatively better when price dispersion is high.

In our framework, a higher share of non-rational agents entails more backward-looking behavior. Thus, it tends to reduce the gains of commitment via the persistence effect (as shown in panels (a) and (b)). By contrast, extrapolative expectations (panel (c)) increase price dispersion. Therefore, high θ values increase the relative gains of commitment via the dispersion effect. When the share of non-rational agents is not too large ($\alpha \in [0.7, 1]$), a higher share

 $^{^{19}}$ Cf. the two last two columns of row (3) in each panel.

 $^{^{20}}$ Cf. the first two columns of row (3) in panels (a) and (b).

 $^{^{21}}$ Cf. the first two columns of row (3) in panel (c).

of non-rational agents increases the gains of commitment, since the dispersion dominates the persistence effect.

4.4 Optimal policies vs. Taylor rules

It is often argued that Taylor rules well describe the conduct of monetary policies. Moreover, under certain circumstances, they also mimic optimal discretionary policies. In the canonical model, a simple Taylor rule, responding to inflation and the output gap according to the Taylor principle, may lead to similar dynamics compared to optimal discretionary policies. We check the robustness of this result by examining how the gains from the discretion over the Taylor rule are affected by the degree of bounded rationality.

We consider two model specifications that differ only in how the interest rate is set. In one case, the policy maker acts under discretion, whereas in the other, the central bank adjusts the nominal interest rate according to a Taylor rule responding to both inflation and the output gap:

$$i_t = \phi_\pi \pi_t + \phi_y y_t \tag{24}$$

where $\phi_{\pi} = 1.5$ and $\phi_{y} = 0.125$.

In the table below, we provide the results of our numerical simulations for several values of α and θ . The remaining parameters are calibrated as reported in Table 1.

	St	nare of r	rational	agents	(α)
(a) Adaptive expectations	1	0.9	0.8	0.7	0.6
Discretion	140.9	177.2	221.9	275.3	336.6
Taylor rule	141.2	180.1	230.4	294.4	372.6
Difference $(\%)$	0.2	1.6	3.8	6.9	10.7
	Sh	are of r	rational	agents	(α)
(b) Naive expectations	1	0.9	0.8	0.7	0.6
Discretion	140.9	185.5	243.2	314.8	396.8
Taylor rule	141.2	189.8	258.4	356.7	495.8
Difference $(\%)$	0.2	2.3	6.2	13.3	24.9
	Share of rational agents (α)				(α)
(c) Extrapolative expectations	1	0.9	0.8	0.7	0.6
Discretion	140.9	196.1	272.4	373.5	490.8
Taylor rule	141.2	202.6	301.2	485.4	1019.4
Difference $(\%)$	0.2	3.3	10.5	29.9	107.6

Table 3 - Optimal policy vs. Taylor rule: Welfare losses

In the standard framework, where all the agents are rational (i.e., $\alpha = 1$), the Taylor rule is suboptimal while discretion gains are small (just 0.2%). Note that these gains are independent of θ . As we introduce heterogeneous beliefs, we observe that the welfare losses associated with a Taylor rule progressively grow and that the relative gains of discretion become large. The costs associated with Taylor rule-based policies are between 1.6% and 107%, depending on the values assigned to α and θ .²²

Our results indicate that although the costs of a Taylor rule compared to optimal inflation targeting are small when the standard case is considered, they may become consistent once some degree of bounded rationality is introduced. The costs of the Taylor rules increase because they cannot stabilize the variability in individual beliefs, and thus, they are associated with increasing costs in terms of inequality. Moreover, considering the canonical trade-off between price and aggregate consumption stabilization, as explained, the weight in the Taylor rule should decrease in α to mimic optimal policies.

²²Similar results are found for a Taylor rule responding only to inflation rate, i.e., a policy that somehow replicates an inflation targeting. Results are available upon request.

In Figure 3, we plot the marginal gains of discretion, expressed in percentage terms, over a Taylor rule. According to Table 3, they are very close to zero when all the agents are rational; however, as the share of rational agents decreases, acting following a simple Taylor rule could induce exceptionally high marginal losses.



Figure 3 – Welfare gains of discretion over a Taylor rule

5 The dangers of ignoring expectation heterogeneity

We investigate the potential costs associated with the presence of model uncertainty. Gasteiger (2014) also tests the robustness to misspecification in Branch and McGough (2009). However, by using an *ad hoc* loss function, he only focuses on determinacy.²³ Because we have derived a welfare measure consistent with the distribution of the agents between types, we instead explore the

²³Specifically, he derives an expectational-based reaction function for monetary policy assuming that all the agents are rational, and he tests the determinacy of the equilibrium when the economy obeys to a Phillips curve obtained by assuming that part of the agents have a nonrational forecast. He finds that the central banker always controls inflation and the output gap as long as $\theta \leq 1$ (i.e., adaptive or naive expectations). Indeterminacy may only arise when nonrational agents form their expectations by extrapolative schemes (i.e., $\theta > 1$).

potential dangers of ignoring expectation heterogeneity under a welfare-maximizing perspective.

Specifically, inspired by Brock *et al.* (2007) and Sbordone (2007), we consider an environment characterized by model uncertainty, assuming that the central bank is unable to recognize how agents form their beliefs. We assume that the central bank knows the general form of the welfarebased loss (16), but it ignores the true value of α , and we explore the optimal robust policy that consists of minimizing the maximum regret of choosing the wrong α (Savage, 1951).

Savage's criterion can be defined by introducing the concept of regret. Consider a policymaker facing model uncertainty and some policy options. Possible models and policies are defined in the sets M and P, respectively. The regret is the difference between the loss incurred by implementing a certain policy in a certain given model specification and the loss incurred by using the optimal policy in that given model. Formally, the regret of policy $p_i \in P$ in model $m_j \in M$ is $R(p_i, m_j) = L(p_i, m_j) - \min_{p \in P} L(p, m_j)$, where $L(p_i, m_j)$ is the loss. Note that $\min_{p \in P} L(p, m_j)$ is the optimal policy in m_j to which a regret equal to zero is associated. Savage's robust policy p^* is then $p^* = \min_{p \in P} \max_{m \in M} R(p, m)$. Operatively, we first compute the maximum regrets associated with different policies, and then we choose their minimum.

Let us use the criterion. We first restrict our attention to a dichotomous scenario where α can only assume two given values (1 or 0.7). Basically, the policymaker can make two kinds of mistakes: assume nonrational expectations in a world where all the agents are rational or neglect heterogeneous beliefs in a context where a share of agents is not fully rational. What we want to investigate is which error is less costly in terms of welfare and, accordingly, which policy should be chosen to minimize the potential additional welfare loss arising in the presence of uncertainty.

We denote by $\tilde{\alpha} \in \{1, 0.7\}$ the true value of α ; $\alpha^{CB} \in \{1, 0.7\}$ is instead the value of α that the central bank assume to be true. The central bank uses α^{CB} in its loss, but $\tilde{\alpha}$ characterizes its constraints.²⁴ For all possible combinations, losses are computed, followed by regrets. The results are described in Table 4, which reports welfare losses (between brackets) and regrets under model uncertainty. For instance, 117.19 is the welfare loss arising when $\tilde{\alpha} = 1$, but the central bank sets the optimal policy assuming that $\alpha^{CB} = 0.7$; 6.34 is the policy regret of choosing $\alpha^{CB} = 0.7$ when $\tilde{\alpha} = 1$, i.e., the difference between 117.19 and the loss associated with the optimal policy under certainty, 110.85.

Making a mistake always involves additional welfare costs. Therefore, the (maximum) regret is always that associated with choice of the wrong α . In the case of timeless commitment (panel (a)), the maximum regret of choosing $\alpha^{CB} = 0.7$ is 6.34, while the maximum regret of choosing

 $^{^{24}\}mathrm{Details}$ are provided in the Appendix.

 $\alpha^{CB} = 1$ is 75.9.

According to the *minimax* regret criterion, then, the optimal robust policy requires taking into account expectations heterogeneity, as this policy selects the minimum among the maximum regret. Thus, ignoring expectation heterogeneity is more harmful, in welfare terms, than considering it. It is easy to test that the same result holds for discretion (see panel (b)).

	(a) timeless commitment		(b) discretion			
True α is $(\tilde{\alpha})$		α is $(\tilde{\alpha})$	True α is $(\tilde{\alpha})$			
Policy (α^{CB}) is	0.7	1.0	0.7	1.0		
0.7	0.00	6.34*	0.00	4.96*		
	(248.40)	(117.19)	(314.84)	(145.92)		
1.0	75.9	0.00	17.99	0.00		
	(324.26)	(110.85)	(332.83)	(140.96)		
Legend:	The table repo	rts policy regrets	and, in the brack	kets, welfare losses.		
	(*) optimal robust policy (<i>minimax</i> regret).					

Table 4 – Welfare losses and regrets under model uncertainty

The intuition behind our result is as follows. When the central bank faces a world characterized by nonrational forecasting, the policymaker's problem has a further dimension represented by the need to minimize also the cross-sectional variance of consumption among agents. If the monetary authority ignores this, the optimal policy does not consider the minimization of consumption among types of agents, and aggregate inflation is not fully stabilized (see (17)), inducing additional welfare losses. By contrast, when the central bank erroneously assumes the presence of bounded rationality, it includes consumption inequality in its welfare loss, but the additional losses are smaller than in the previous case. This result comes from the fact that in a fully rational world all the agents consume the same quantity of goods and thus there are no costs associated with inequality in consumption.

The above result holds for several calibrations, as we tested a share of nonrational agents ranging from 50 to 90% vs. the canonical model where $\alpha = 1$. Our test is based on pair-wise comparisons, i.e., we assume that the cardinality of both M and P is two. We can refine our investigation considering a larger spectrum of uncertainty. Specifically, we consider that α can assume 0.5, 0.6, 0.7, 0.8, 0.9, and 1.²⁵ Optimal robust policy requires setting $\alpha^{CB} = 0.7$ and $\alpha^{CB} = 0.6$ with timeless commitment and discretion, respectively.

²⁵Details are available upon request.

6 Conclusions

Starting from well-documented observations that people form their expectations according to different mechanisms, this paper studied the impact of heterogeneous expectations on optimal monetary policies in a New Keynesian framework. The point of departure of our work is Branch and McGough (2009), who introduce bounded rationality in a small-scale New Keynesian DSGE model and provide an equilibrium determinacy analysis when the policymaker sets monetary policy according to a Taylor rule. We use their setup to investigate the implications of the heterogeneity of optimal monetary policy design. We also test the robustness of canonical inflation-targeting policies under both discretion and commitment regimes and the equilibrium determinacy when some agents behave according to bounded rationality.

Our simulations show that optimal policies and welfare losses quantitatively and qualitatively depend on the share of boundedly rational agents. As the share of boundedly rational agents increases, welfare deteriorates because the economy becomes more persistent, and simultaneously, price dispersion increases, thereby inducing further welfare costs. In this context, optimal policies depart from those optimal in the canonical framework, which can sometimes be represented by appropriated Taylor rules. The assumption that some agents form their expectations using an adaptive or extrapolative mechanism in fact introduces a new dimension to the standard policy problem faced by the central bank, that of consumption inequality. Moreover, we find that optimal policies are always associated with equilibrium determinacy under both discretion and commitment, which differs from recent literature based on Taylor rules in the same context.

The assumption that some agents form their expectations using an adaptive or extrapolative mechanism implies that the central bank also needs to minimize the economic inequality (cross-sectional variance of consumption) in addition to the usual two dimensions of the central bank's policy problems (i.e., minimizing the variability of aggregate consumption and the costs associated with price dispersion.) In other words, it now faces a three-dimensional problem.

To minimize the economic inequality, the central bank should stabilize expectation variability. As long as different individual consumption plans depend on different expectation paths, the variability in individual consumption is reduced by minimizing the cross-sectional variability of expectations. The usual trade-off between the stabilization of aggregate consumption and price dispersion is also affected by heterogeneity. The cost of price dispersion increases with the size of the group of boundedly rational agents, but the emphasis of stabilizing inflation versus output declines as long as heterogeneity increases, due to the more complex dynamics of price dispersion that, in this case, also depend on output stabilization.

Comparing welfare performances under different policy regimes, we find, as expected, that commitment always guarantees the lowest welfare losses. Then, following Steinsson (2003), we investigate the relative gains of commitment over discretion by considering different degrees of heterogeneity. In general, the relative gains of commitment depend on two different effects, which are higher when i) the forward-looking component of the Phillips curve progressively enhances compared to the backward one (persistence effect) and ii) price dispersion is more costly (dispersion effect). We find that an increase in the share of boundedly rational subjects has an ambiguous effect on the relative gains of commitment because, on the one hand, it reduces the lead component of the Phillips curve, but, on the other hand, it raises the price dispersion. In our setup, commitment gains are more likely to be observed when agents form their expectations using an extrapolative mechanism and the fraction of non-rational agents is small (dispersion effect prevails over persistence effect). By contrast, commitment gains decrease whenever expectations are adaptive or naive (the persistence effect dominates the dispersion effect).

We highlight the importance of pursuing an optimal stabilization policy rather than following a simple exogenous interest rule. We show that in a world characterized by a fraction of non-rational agents, the costs of neglecting an optimal policy rule barely increase, leading to significant welfare losses. The rationale is that a Taylor rule is unable to mimic the optimal policy design because it overreacts to inflation and does not stabilize individual expectations. The former involves suboptimal choices in the trade-off between aggregate consumption and price dispersion stabilization, while the latter induces costs in terms of consumption inequality between individuals.

Similarly, under model uncertainty, we highlight the dangers of ignoring heterogeneity. Because our main lesson is that heterogeneity calls for taking account of consumption inequality across types, the cost of ignoring it when it is present is greater than the cost of taking account of it when it is not. In the former, in fact, even if the policymaker assigns a weight to the consumption inequality, this does not affect his decisions, since there is no inequality across agents. By contrast, in the latter, if the weight is not assigned, the cost of inequality is not minimized, causing welfare costs. Clearly, when the policymaker misunderstands the existence of heterogeneity, the different "wrong" weight on inflation will instead affect welfare in both cases.

Appendix

Optimal policy can be obtained by solving a LQ problem constrained by the REE associated to the HEE. By using $y_t = \alpha c_t^{\mathcal{R}} + (1 - \alpha) c_t^{\mathcal{B}}$, the problem can be conveniently written as

$$W = \min \sum_{t=0}^{\infty} \beta^{t} \left[\frac{\sigma \eta + 1}{\sigma} y_{t}^{2} + \frac{\alpha \left(y_{t} - c_{t}^{\mathcal{R}} \right)^{2}}{(1 - \alpha) \sigma} + \varepsilon^{2} \eta \delta \left\{ \pi_{t}^{2} + \frac{\xi_{p} \left(1 - \alpha \right)}{\alpha} \left[\pi_{t} - \beta \theta^{2} \pi_{t-1} - \kappa y_{t} - \frac{\alpha \kappa \left(y_{t} - c_{t}^{\mathcal{R}} \right)}{(1 + \eta \sigma) \left(1 - \alpha \right)} \right]^{2} \right\} \right]$$
(25)

$$y_t = \alpha E_t y_{t+1} + (1-\alpha)\beta \theta^2 y_{t-1} - \sigma \left[i_t - \alpha E_t \pi_{t+1} + (1-\alpha)\theta^2 \pi_{t-1} \right]$$
(26)

$$\pi_t = \alpha \beta E_t \pi_{t+1} + (1 - \alpha) \beta \theta^2 \pi_{t-1} + \kappa y_t + e_t \tag{27}$$

$$c_t^{\mathcal{R}} = E_t c_{t+1}^{\mathcal{R}} - \sigma(i_t - E_t \pi_{t+1})$$
(28)

Defining $X_t = [y_t, y_{t-1}, \pi_t, \pi_{t-1}]'$, (25) can be written as

$$W = \sum_{t=0}^{\infty} \beta^t X_t' Q X_t \tag{29}$$

where Q is an appropriate matrix of weights; (26)-(28) can be put in a state-space form:

$$X_{t+1} = AX_t + Bi_t + Ce_t \tag{30}$$

where A and B are appropriate matrices.

s.t.

By the above representation, standard solution algorithms can then be used to minimize discounted sum of period losses and at the same time solving for rational expectations equilibrium. We use the algorithms developed by Söderlind (1999).

An algorithm solves for optimal discretionary rules using dynamic programming principles. The policymaker reoptimizes every period by taking the process by which private agents form their expectations as given, but where the expectations are consistent with actual policy. Since the model is linear-quadratic, the solution in any time gives a value function which is quadratic in the state variables. Timeless commitment is based on Lagrangian methods. Assuming $\alpha = \alpha^{CB}$ (i.e., policy parameter) in (29) and $\alpha = \tilde{\alpha}$ (i.e., share of rational households) in (30), we can obtain the outcomes discussed in the paper. In particular, for $\alpha = \alpha^{CB} = \tilde{\alpha}$, we obtain our main results. Instead, by considering possible differences between α^{CB} and $\tilde{\alpha}$, we can discuss optimal robust policies.

References

- Amato, J.D. and T. Laubach (2003), "Rule-of-thumb behaviour and monetary policy," European Economic Review, 47(5): 791-831.
- Benigno, P. and M. Woodford (2004), "Optimal monetary and fiscal policy: A linear-quadratic approach," NBER Chapters, in NBER Macroeconomics Annual 2003, Volume 18: 271-364.
- Berardi M. (2007), "Heterogeneity and misspecifications in learning," Journal of Economic Dynamics and Control, 31(10): 3203-3227.
- Blake, A.P. (2001), "A 'timeless perspective' on optimality in forward-looking rational expectations models," NIESR Discussion Papers 188, National Institute of Economic and Social Research.
- Bilbiie, F.O. (2008), "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of Economic Theory*, 140(1): 162-196.
- Blanchard, O.J. and C.M. Kahn (1980), "The solution of linear difference models under rational expectations," *Econometrica*, 48(5): 1305-1311.
- Branch, W.A. and G.W. Evans (2006), "Intrinsic heterogeneity in expectation formation," Journal of Economic Theory, 127(1): 264-295.
- Branch, W.A. and B. McGough (2004), "Multiple equilibria in heterogeneous expectations models," *The B.E. Journal of Macroeconomics*, 4(1): 1-16.
- Branch, W.A. and B. McGough (2005), "Consistent expectations and misspecification in stochastic non-linear economies," *Journal of Economic Dynamics and Control*, 29(4): 659-676.
- Branch, W.A. and B. McGough (2009), "A New Keynesian model with heterogeneous expectations," Journal of Economic Dynamics and Control, 33(5): 1036-1051.
- Brock, W.A. and C.H. Hommes (1997), "A rational route to randomness," *Econometrica*, 65(5): 1059-1096.
- Brock, W.A. and C.H. Hommes (1998), "Heterogeneous beliefs and routes to chaos in a simple asset pricing model," *Journal of Economic Dynamics and Control*, 22(8-9): 1235-1274.
- Brock, W.A., S.N. Durlauf, J.M. Nason and G. Rondina (2007), "Simple versus optimal rules as guides to policy," *Journal of Monetary Economics*, 54(5): 1372-1396.
- Cirillo, P., D. Delli Gatti, S. Desiderio, E. Gaffeo and M. Gallegati (2011), "Macroeconomics from the bottom-up," *Milan: Springer*.

- Calvo, G.A. (1983), "Staggered prices in a utility-maximizing framework," Journal of Monetary Economics, 12(3): 383-398.
- Clarida, R., J. Galí and M. Gertler (1999), "The science of monetary policy: A New Keynesian perspective," *Journal of Economic Literature*, 37(4): 1661-1707.
- Duffy J. and W. Xiao (2007), "The value of interest rate stabilization policies when agents are learning," *Journal of Money, Credit and Banking*, 39(8): 2041-2056.
- Evans, G.W. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press.
- Evans, G.W. and S. Honkapohja (2003), "Adaptive learning and monetary policy design," Proceedings, Federal Reserve Bank of Cleveland: 1045-1084.
- Gaffeo, E., I. Petrella, D. Pfajfar and E. Santoro (2014), "Loss aversion and the asymmetric transmission of monetary policy," *Journal of Monetary Economics*, 68(3): 19-36.
- Galì, J., J.D. López-Salido and J. Vallés (2004), "Rule of thumb consumers and the design of interest rate rules," *Journal of Money, Credit and Banking*, 36(4): 739-763.
- Galì, J., J.D. López-Salido and J. Vallés (2007), "Understanding the effects of government spending on consumption," *Journal of the European Economic Association*, 5(1): 227-270.
- Gasteiger, E. (2014), "Heterogeneous expectations, optimal monetary policy, and the merit of policy inertia," Journal of Monetary, Credit and Banking, 46(7): 1533-1554.
- Kocherlakota, N. (1996), "Implications of efficient risk sharing without commitment," Review of Economic Studies, 63(4): 595-609.
- Krusell, P. and A.A. Smith (1996), "Rule of thumb in macroeconomic equilibrium: A quantitative analysis," *Journal of Economic Dynamics and Control*, 20(4): 527-558.
- Jensen, C. and B.T. McCallum, (2002), "The non-optimality of proposed monetary policy rules under timeless perspective commitment," *Economics Letters*, 77(2): 163-168.
- Jensen, C. and B.T. McCallum, (2010), "Optimal continuation versus the timeless perspective in monetary policy," *Journal of Money, Credit and Banking*, 42(6): 1093-1107.
- Mankiw, N.G. (2000), "The saver-spenders theory of fiscal policy," American Economic Review, 90(2): 120-125.
- Mankiw, N.G. and R. Reis (2002), "Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 117(4): 1295-1328.
- Mankiw, N.G. and R. Reis (2007), "Sticky information in general equilibrium," Journal of the European Economic Association, 5(2-3): 603-613.
- Massaro, D. (2013), "Heterogeneous expectations in monetary DSGE models," Journal of Economic Dynamics and Control, 37(3): 680-692.
- Pesaran, M.H. (1987), "The limits to rational expectations," Oxford: Basil.
- Preston, B. (2006), "Adaptive learning, forecast-based instrument rules and monetary policy," Journal of Monetary Economics, 53(3): 507-535.

- Ravenna, F. and C.E. Walsh, (2011), "Welfare-Based Optimal Monetary Policy with Unemployment and Sticky Prices: A Linear-Quadratic Framework," *American Economic Journal: Macroeconomics*, 3(2): 130-62.
- Rotemberg, J.J. and M. Woodford (1997), "An optimization-based econometric framework for the evaluation of monetary policy," *NBER Macroeconomics Annual 1997*, 12: 297-361.
- Savage, L.J. (1951), "The theory of statistical decision," Journal of the American Statistical Association, 46(253): 55-67.
- Shi, S. (1999), "Search, inflation and capital accumulation," Journal of Monetary Economics, 44(1): 81-103.
- Sbordone, A. (2007), "Inflation persistence: Alternative interpretations and policy implications," Journal of Monetary Economics, 54(5): 1311-1339.
- Steinsson, J. (2003), "Optimal monetary policy in an economy with inflation persistence," Journal of Monetary Economics, 50(7): 1425-1456.