Heterogeneity in Inflation Expectations and Macroeconomic Dynamics under Evolutionarily Satisficing Learning

Gilberto Tadeu Lima
University of São Paulo
giltadeu@usp.br

Jaylson Jair da Silveira
Federal University of Santa Catarina
jaylson.silveira@ufsc.br

Abstract: Drawing on a considerable empirical literature that reveals persistent and endogenously time-varying heterogeneity in inflation expectations, this paper embeds two inflation forecasting strategies – one based on costly *ex ante* full rationality or perfect foresight, and the second based on costless *ex ante* bounded rationality or extrapolative trend-following – in a dynamic macroeconomic model. Drawing also on the significant empirical evidence that inflation forecast errors may have to exceed some threshold before agents abandon their previously selected inflation forecasting strategy, we describe agents as switching between inflation forecasting strategies according to evolutionarily satisficing learning dynamics. We find that convergence to a long-run equilibrium consistent with growth, unemployment and inflation at their natural levels may be achieved even when heterogeneity in inflation expectations (with predominance of the extrapolative trend-following foresight strategy) is an attractor of an evolutionarily satisficing learning dynamic perturbed by mutant agents. Therefore, in keeping with robust empirical evidence, heterogeneity in strategies to form inflation expectations (with prevalence of boundedly rational expectations) can be a stable long-run equilibrium.

Keywords: Heterogeneous inflation expectations; perfect foresight; extrapolative foresight; satisficing learning dynamics.

Resumo: Baseado em uma considerável literatura empírica que detecta que a heterogeneidade das expectativas de inflação é persistente e varia endogenamente no tempo, o artigo embute duas estratégias de previsão da inflação – uma, custosa *ex ante*, baseada em racionalidade plena ou previsão perfeita, a outra, sem custo *ex ante*, baseada em racionalidade limitada ou previsão extrapolativa – em um modelo macroeconômico dinâmico. Baseado também na evidência empírica de que os erros de previsão devem ultrapassar certo limiar antes que os agentes abandonem a estratégia de previsão da inflação previamente selecionada, supomos que os agentes mudam de estratégia de previsão da inflação conforme uma dinâmica de aprendizado evolucionariamente *satisficing*. Mostramos que a convergência para o equilíbrio de longo prazo em que as taxas de inflação, crescimento e desemprego estão em seus níveis naturais pode ser obtida mesmo quando a heterogeneidade de expectativas de inflação (com predominância da estratégia de previsão extrapolativa) é um atrator de uma dinâmica de aprendizado evolucionariamente *satisficing* perturbada por agentes mutantes. Logo, em conformidade com robusta evidência empírica, heterogeneidade nas estratégias de formação de expectativas de inflação (com predominância de expectativas limitadamente racionais) pode ser um equilíbrio de longo prazo estável.

Palavras-chave: expectativas de inflação heterogêneas; previsão perfeita; previsão extrapolativa; dinâmica de aprendizado de *satisficing*.

J.E.L. Classification Codes: E03; E31; C62; C73.

Classificação Anpec: Área 4 – Macroeconomia, Economia Monetária e Finanças.

* We are grateful to Larry Samuelson and William Brock for helpful comments and suggestions. Nonetheless, we take full responsibility for any remaining errors. We are also grateful to CNPq (Brazil) and FAPESP (Brazil) for research funding.

1. Introduction

There is compelling empirical evidence from both survey data and laboratory experiments that inflation expectations are persistently heterogeneous and formed mostly through boundedly rational, rule-of-thumb mechanisms or heuristics (see, e.g., Hommes, 2013, for a detailed overview). Meanwhile, there is also robust evidence from survey data that sometimes inflation forecast errors must pass some threshold before decision makers decide to abandon their previously selected inflation predictor (Branch, 2004, 2007).

Motivated by these several pieces of robust empirical evidence (which we report and discuss more thoroughly in the next section), we extend a dynamic macroeconomic model to explore implications of heterogeneous inflation expectations which vary over time according to *evolutionarily satisficing learning dynamics* (with and without mutation). From a methodological perspective, the intended contribution of this paper lies in particular in an alternative way of modeling of the dynamics of inflation predictor switching when inflation forecast errors have to cross some threshold in order to induce a switch. Our proposed modeling strategy is intended to incorporate this documented threshold or inertia effect in a more representative way than has been carried out in the existing literature (Branch, 2004, 2007; Lines and Westerhoff, 2010), in which such threshold is modeled as a component of the cost associated with the use of a predictor (and a cost which is usually taken as exogenously fixed and deterministic). It is true, as argued in Branch (2004, 2007), first, that if forecast errors have to pass some threshold in order to induce a switch, there is inertia in predictor switching; and second, that costs act as thresholds through which forecast errors must cross to induce switching between prediction methods. In our model, however, although the existence of a cost incurred by agents to use a predictor also imposes inertia, satisficing predictor choice in and of itself is already an inherent source of inertia in predictor switching. It is therefore analytically convenient, if not required, that these two distinct inherent sources of inertia in predictor switching are not bunched together in a single channel. For instance, in our model, as in Lines and Westerhoff (2010), there is an exogenously fixed cost associated with perfect inflation prediction. Yet, unlike in Lines and Westerhoff (2010), we make it explicit that this cost involves the fact that perfect inflation prediction requires perfectly knowing the frequency distribution of inflation predictors in the population of agents. As it was early pointed out by Brock & de Fontnouvelle (2000), the perfect foresight predictor involves very complex calculations given that it has to take into account the equilibrium effects of (all) the other predictor(s). Thus, it is analytically convenient that those two sources of inertia in predictor switching are incorporated to the model separately. Meanwhile, if the cost attached to the use of an inflation predictor, when broadly conceived, ensures that the empirical estimate of the proportion of agents using it most closely fit the data, as documented in Branch (2004, 2007), our suggested satisficing choice protocol can be seen as similarly empirically accommodating. Indeed, as acknowledged in the closing of Branch (2004), the data are supportive of the model of the paper, but not completely so, given that such model could not account for inertia-generating bias towards a predictor. The author then concludes that the empirical evidence found in the paper indicate that a theoretical account of these biases or predispositions is called for. In this context, the model set forth in this paper suggests an explanation of inflation predictor choice based on evolutionarily satisficing learning.

As further described in Section 3, our formal methodological approach follows the spirit of the suggestive contributions of Herbert Simon and conceives of choice as primarily intending to meet an acceptability threshold rather than to select the best of all possible alternatives, hence the notion of *satisficing behavior*. Consequently, while in both Branch (2004, 2007) and Lines and Westerhoff (2010) the inflation predictor selection follows the discrete choice model set forth in Brock and Hommes (1997), in this paper the inflation predictor choice follows an evolutionarily satisficing learning dynamic subject to a perturbation analogous to mutation in natural environments. In fact, by replacing the demand for choosing the best inflation predictor possible for the agent with a conceptually and formally precise notion of being acceptably or tolerably inaccurate an inflation prediction, our evolutionarily satisficing learning mechanism provides a fruitful alternative way of rationalizing the inertia finding in Branch (2004, 2007). In the latter, the finding that agents abandon a previously selected predictor only after forecast errors cross some threshold is seen as reflecting that agents have inherent predispositions to use one predictor over another. In the model herein, the same finding is seen as reflecting that agents have an inherent predisposition not to one predictor over another, but to satisficing choice behavior. Indeed, as our proposed satisficing behavior is subject to mutation, we are ascribing to agents another inherent predisposition, which is a predisposition to experiment once and a while (as explained below). According to Herbert Simon's original usage of the term "satisficing", choices are identified that are suitable enough by comparing their attributes to some standard. While in Simon (1955, 1956) choices are compared to the "aspiration level" of how good a choice might be made, in the model of this paper inflation predictor choices are compared analogously to the aspiration level of how acceptably or tolerably inaccurate a given inflation predictor choice might be made in light of the uncertainties surrounding the future.

As a result, while in Branch (2004) agents are admittedly perceived as inherently predisposed to using an inflation predictor in a way that does not depend on accuracy, our conception that agents are instead inherently predisposed to satisficing choice does not mean that accuracy is overlooked. What satisficing predictor choice means is that agents, instead of looking for the highest of all possible accuracies, in fact look for an acceptable level of accuracy or, to put it another way, are satisfied with (or tolerate) a positive level of inaccuracy in inflation prediction. Therefore, our notion of satisficing predictor choice is consistent with the viewpoint stated in Branch (2004) that "the agent will not switch his prediction method because the forecast error is 'reasonable'. Only forecast errors above a threshold will induce changes in predictor choice." (p. 613) Besides, we plausibly assume that the level of inaccuracy in inflation prediction that is deemed reasonable or acceptable or tolerable is agent-specific and randomly distributed acrosss agents. In fact, this heterogeneity across agents in acceptable prediction errors is in accordance with the empirical evidence found in Branch (2004).

Revealingly, the evolutionarily satisficing learning mechanism proposed herein, in addition to incorporating the threshold effect documented in Branch (2004, 2007) in a representative, conceptually sound and intuitive way, carries relevant implications for the long-run equilibrium configuration of a macroeconomic model as the one set forth in the related contribution by Lines and Westerhoff (2010). In fact, the evolutionarily satisficing dynamics across inflation predictor choice suggested in this paper is coupled to the equilibrium dynamics of the endogenous macroeconomic variables. Meanwhile, although our suggested evolutionarily satisficing choice mechanism is embedded in a specific macroeconomic model featuring inflation prediction, readers are nonetheless invited to contemplate the possibility that other instances of predictor (or, more broadly still, behavior) switching are also subject to satisficing evolutionary learning.

In keeping with the empirical evidence on inflation expectations and following Lines and Westerhoff (2010), the aggregate expected inflation is a weighted average of the inflation forecast of agents who play the costly *ex ante full rationality* or *perfect foresight* strategy and the inflation forecast of agents who play a costless *ex ante bounded rationality* or *extrapolative trend-following* foresight strategy. Differently from the discrete choice framework adopted in Lines and Westerhoff (2010), however, our evolutionarily satisficing learning mechanism, if not perturbed by a noise analogous to mutation in natural environments, necessarily generates pervasive extrapolative trend-following foresight (all agents eventually decide to use the bounded rationality inflation forecasting strategy) as a long-run equilibrium solution. Meanwhile, in the presence of mutation the evolutionarily satisficing learning mechanism necessarily culminates with heterogeneity in strategies to establish inflation expectations as a long-run equilibrium, which is in keeping with the empirical evidence. In either situation (with and without mutation) the long-run equilibrium is nonetheless characterized by the natural values of the output growth, unemployment and inflation being achieved. In fact, a long-run equilibrium configuration consistent with the achievement of the natural values of the endogenous macroeconomic variables does not emerge because all agents adopt the perfect foresight inflation forecasting strategy. Instead, it is either all agents (without mutation) or the greater proportion of them (with mutation) that eventually decide to follow the bounded rationality inflation forecasting strategy because the natural values of the output growth, unemployment and inflation are being achieved.

As the general case, the model herein features "mutation" as a perturbation in the satisficing evolutionary mechanism, leading some decision makers to select an inflation forecasting strategy at random. We should emphasize that "mutation" is a metaphor for all sorts of payoff-unrelated reasons, outside the more explicit part of the model, for why agents might switch inflation predictor, and we take the rate of mutation as being independent across agents and over time. Two rationales for the presence of mutation are that an agent exits the economy with some (fixed) probability and is replaced with a new agent who still does not know anything about the decision-making environment, or that each agent simply experiments occasionally with exogenously fixed probability. In any case, it follows that our model features three layers of randomness: first, the level of acceptable inflation forecast accuracy is randomly and independently determined across agents and over time; second, the specific agents who behave as mutants in each period, which measure is given by a mutation rate which is fixed and independent across agents and over time, are randomly drawn from the population; and third, agents who behave as mutants choose an inflation foresight strategy at random.

Meanwhile, a question that arises is whether the presence of mutation happens to prevent convergence towards a long-run equilibrium configuration consistent with the achievement of the natural values of the endogenous macroeconomic variables, and the answer is not necessarily. In fact, it turns out that the long-run equilibrium distribution of inflation forecasting strategies across agents and its local stability properties depend on whether the satisficing learning dynamics are perturbed by mutation: in the absence (presence) of mutant agents, the long-run equilibrium configuration, which may be a local attractor, is characterized by the natural values of the endogenous macroeconomic variables and the extinction of the perfect foresight strategy (survival of both inflation forecasting strategies). Thus, pervasive perfect foresight is neither a necessary condition for the achievement of the natural values of the rates of output growth, unemployment and inflation, nor an inevitable consequence of the achievement of these natural values. In fact, our analytical results show that heterogeneous inflation expectations on the part of agents need not necessarily thwart the successful handling of stabilization policy-making. This is actually the case even when inflation expectations are not only formed heterogeneously but also, as a direct consequence of the propensity of decision makers to switch between inflation forecasting strategies according to evolutionarily satisficing criteria, endogenously time-varying.

The remainder of the paper is organized as follows. The next section presents the motivating empirical evidence and briefly discusses (and broadly compares the present contribution to) a representative sample of the related literature, while Section 3 lays out the structure of the basic macroeconomic model on which our investigation is based. Section 4 explores the consequences for macroeconomic stability properties in the long-run equilibrium of heterogeneous inflation expectations that market participants switch between as described by evolutionarily satisficing learning dynamics with and without mutation. Section 5 then concludes.

2. Motivating evidence and related literature

The extensive empirical literature on inflation expectations formation can be divided into two groups depending on the type of evidence used in the investigation: studies based on survey data; and laboratory experiments with human subjects. In either case, a common finding is the strong empirical support for time-varying heterogeneity in inflation expectations with predominance of some form of bounded rationality (see, e.g., Duffy, 2008, and Hommes, 2011, 2013, for thorough overviews of experiments dealing with formation of expectations about macroeconomic variables). The principal reasons for this heterogeneity that have been suggested in the literature are that market participants rely on different models, have access to different information sets, or have different cognitive capacities for processing information.

In the experimental literature, Adam (2007) documents that subjects' inflation expectations are not captured by the rational expectations predictor, whereas predictors based on lagged inflation capture inflation expectations quite well. Pfajfar and Zakelj (2011) find evidence of persistent heterogeneity in the inflation forecasting process both across subjects and time. They find that subjects form expectations based on different models, and despite the most popular rule is trend extrapolation, a sizable proportion of subjects uses adaptive expectations, adaptive learning (as described in Evans and Honkapohja, 2001) or sticky information type models (as proposed in Mankiw and Reis, 2002). Pfajfar and Zakelj (2014) find evidence that switching between different rules to form inflation expectations describes the behavior of subjects better than a single rule. Experimental evidence for persistent heterogeneity in inflation expectations is found in Roos and Luhan (2013), who find that adaptive and static (or naïve) inflation expectations are often observed. Heterogeneous and time-varying inflation expectations have also been found in studies based on survey data. Mankiw, Reis and Wolfers (2004) and Pfajfar and Santoro (2010) use US survey data to find considerable heterogeneity in inflation expectations, and Blanchflower and MacCoille (2009) use UK survey data to find strong support for heterogeneous and backward-looking inflation expectations.

The literature also features macroeconomic models with heterogeneous inflation expectations from which predictions are sometimes tested empirically. Branch (2004) develops a model where agents form their forecasts of inflation by selecting a predictor from a set of costly alternatives whereby they may rationally choose a method other than the most accurate due to an inherent inclination to use one predictor over another. Agents are nonetheless seen as *rationally heterogeneous* in that each predictor choice is optimal for them; strictly speaking, agents' expectations are seen as boundedly rational and consistent with optimizing behavior. In our framework, agents' expectations are seen as boundedly rational precisely because they are instead consistent with satisficing behavior; in sum, agents are seen instead as *satisficingly heterogeneous*. In Brazier et al. (2008) and De Grauwe (2010), agents employ two heuristics to forecast inflation: one is based on one-period lagged inflation, the other on the official inflation target. While in both models heuristic switching is described by a discrete choice model à la Brock and Hommes (1997), in the model herein agents switch between forecasting strategies as described by evolutionarily satisficing learning dynamics with and without mutation. Branch and McGough (2010) embed predictor selection in a model with heterogeneity in inflation forecasts. They follow Brock and Hommes (1997) in setting the degree of heterogeneity to endogenously vary over time depending on past forecast errors (net of a fixed cost). Agents choose between a costly rational predictor and a costless adaptive predictor. In our model, in turn, the choice of inflation foresight strategy involves costly rational and costless trend-following foresight strategies, with agents switching between strategies as described by evolutionarily satisficing learning dynamics rather than the discrete choice model set forth in Brock and Hommes (1997).

Our contribution is mostly related to the model set forth in Lines and Westerhoff (2010), in which the interplay between heterogeneity in inflation expectations and the coupled dynamics of inflation, output growth and unemployment is explored. Agents switch between trend-following (which is costless *ex ante*) and perfect foresight (which is costly *ex ante*) strategies to predict future inflation, and the fractions of agents using one or the other inflation predictor is updated over time as described by the discrete choice model in Brock and Hommes (1997). The authors show analytically that the unique equilibrium of inflation, output growth and unemployment features the natural values of these macroeconomic variables, but this equilibrium outcome loses stability through a bifurcation at a critical value which depends on the cost associated with perfect foresight and the parameter measuring how sensitive agents are to selecting the most attractive predictor. The authors also do numerical simulations of the model to explore its global dynamic behavior and investigate conditions under which it leads to endogenous, complex fluctuations in macroeconomic variables and the proportions of agents using one or the other predictor.

Our contribution differs from the interesting one offered in Lines and Westerhoff (2010) in several broad respects (a comparison of specific results is effected in Section 4). First, although both contributions employ the same dynamic macroeconomic model (which features Okun's law, an expectations-augmented Phillips curve and an aggregate demand relation), the two corresponding models differ considerably in the manner they incorporate the empirical evidence documented in Branch (2004) that inflation forecast errors must exceed a threshold in order to induce a switch of inflation predictors. Lines and Westerhoff (2010) follow the suggestion made in Branch (2004, 2007) to treat the cost associated with the use of a predictor in the Brock and Hommes (1997) framework as also reflecting (along with information collection and processing costs) such inertial behavior (or predisposition effect) based on acceptable ranges of forecast errors. In our model, although we maintain the assumption made in Lines and Westerhoff (2010) that there is an *ex ante* cost associated with perfect inflation foresight, the inertial behavior found empirically in Branch (2004) is modeled more characteristically by appealing to Herbert Simon's intuitively suggestive notion of satisficing choice behavior. Indeed, this cogent notion conceives of choice as intending to meet an acceptability threshold rather than to select the best of all possible alternatives, which makes it a representative way to model inertial behavior arising from acceptable ranges of forecast errors. Also, we treat the threshold effect found empirically in Branch (2004) as randomly agent-specific. In Brock and Hommes (1997), agents use a discrete choice model to select a predictor where the deterministic part of the utility of the predictor is the performance measure, and the intensity of choice parameter specifies how good a measure the deterministic component is of the predictor utility. While the standard discrete choice model features deterministic and random individual-specific characteristics, in our model the payoff associated with an inflation foresight strategy features a deterministic component and a random agent-specific error-tolerance threshold. Yet in our model such random component is not random shocks faced by agents, but instead random agent-specific error-acceptance (or, equivalently, error-tolerance) thresholds; although in both models the random component affects the payoff of each of the possible choices. Incidentally, as it turns out, our satisficing evolutionary choice dynamics with agent-specific acceptable prediction errors that are randomly and independently distributed across agents and over time can be seen as in keeping with the following suggestion in Branch (2004): "It would be fruitful for future research to examine a 'richer' dynamical structure then the model at hand. One possible accounting for the inertia finding is that a predictor choice in period t is not independent from the choice in $t-1$. The serial correlation could be the result of time-varying random costs or random utility shocks." (p. 616)

Second, in Lines and Westerhoff (2010) the frequency distribution of inflation predictors in the population of agents evolves over time as described by the predictor switching mechanism set forth in Brock and Hommes (1997), while in our model this frequency distribution follows evolutionarily satisficing learning dynamics with and without mutation. As a result, in Lines and Westerhoff (2010), as in Brock and Hommes (1997), for a strictly positive and finite intensity of choice (or intensity with which firms react to increases in relative net benefit) and information cost, there is a strictly positive share of firms using the strictly dominated perfect foresight strategy even in the long-run equilibrium. In our model, in turn, with a strictly positive and finite information cost and the absence of mutation, a strictly dominated inflation foresight strategy inevitably vanishes asymptotically. In our model, therefore, the long-run evolutionarily satisficing equilibrium configuration is characterized either by all firms playing the trend-following extrapolative foresight strategy (in the absence of mutation) or by most firms playing such strategy (in the presence of mutation). In either situation, however, as in Lines and Westerhoff (2010), the long-run equilibrium is characterized by the natural values of the output growth, unemployment and inflation being achieved.

Third, the role played by mutation in our model is analytically analogous to the one played by a finite intensity of choice in the discrete choice framework followed in Lines and Westerhoff (2010). The analogy arises because in both cases the implication is that not all agents will select the best performing inflation predictor as reflected in the corresponding relative net attractiveness, but the underlying rationale is different in each model. Lines and Westerhoff (2010) define the intensity of choice as a measure of how sensitive agents are to selecting the most attractive predictor. In our satisficing choice framework, meanwhile, a mutant is either an agent who exits the economy with fixed probability and is replaced with a new agent who is unfamiliar with predictor choice, or an agent who experiments occasionally with fixed probability. In either case, a mutant agent chooses each available predictor with equal probability. Besides, recall that Brock and Hommes (1997) and Branch (2004) define the intensity of choice as a measure of how *fast* agents switch predictors for a given change in relative net benefit. Therefore, in the context of a discrete choice framework, one might suggest that it would be more representative to treat the evidence of inertia (or predisposition effect) found in Branch (2004) not as a component of the cost associated with the use of a predictor, but instead as a reflection of a finite intensity of choice. In fact, Branch (2004) concedes that the assumption of a finite intensity of choice is a simple way to incorporate inertia. Besides, although Branch (2004) declares to have found separate evidence of a finite intensity of choice and inertia stemming from the costs, he acknowledges that this inertia may be a reflection of agents' 'inattentiveness', as emphasized in Brock and Hommes (1997), and that his findings support a rich variety of behavior. In this context, our model suggests an alternative explanation of predictor choice behavior based on satisficing evolutionary learning which is consistent with these findings.[1]

Fourth, the model in Lines and Westerhoff (2010) can exhibit local instability of the long-run equilibrium and complicated global equilibrium dynamics, while in our model (with and without mutation) the long-run evolutionary equilibrium configuration is unique and possibly locally stable. In Lines and Westerhoff (2010), as in Brock and Hommes (1997), a rational choice between a cheap *ex ante* destabilizing (boundedly rational) predictor and a costly *ex ante* (perfect foresight) stabilizing predictor leads to the existence of complicated dynamics when the intensity of choice to switch predictors is high. In fact, the model in Lines and Westerhoff (2010) shows that with information costs it may be rational for firms to select methods other than perfect foresight, with the conflict between cheap free riding and costly sophisticated prediction being a potential source of instability and complicated global dynamics. In our model, meanwhile, the evolutionary competition between costless *ex ante* bounded rationality and costly *ex ante* sophisticated inflation prediction may be a source of local stability even if there is exogenous mutation. Thus, although instability is inherent in such circumstances when more sophisticated prediction methods are more expensive (Brock and Hommes, 1997), the evolutionarily satisficing learning suggested in this paper may take the strategy-revision process to a locally stable long-run equilibrium with either the majority or all agents playing a costless *ex ante* trend-following extrapolative foresight strategy.[2]

3. The model

3.1. The macroeconomic setting

As in Lines and Westerhoff (2010), for comparability of results, the macromodel on which the analysis in this paper is conducted can be stated as follows:

(1) $$u_t - u_{t-1} = -\beta(g_t - g_n)\,,$$

[1] Speaking of 'inattentiveness', Lima and Silveira (2015) provide evolutionary game-theoretic micro-foundations to the nominal adjustment of the price level in response to a monetary shock by introducing an analytical notion dubbed 'boundedly rational inattentiveness'. A firm can pay a cost (featuring a random component) to update its information set and establish the optimal price or use costless non-updated information to set a lagged optimal price. The distribution of information-updating strategies across firms and the extent of the nominal adjustment of the price level to a monetary shock co-evolve according to evolutionary dynamics. While in Lima and Silveira (2015) the population dynamics is driven by an asymmetric evolutionary protocol (only firms that do not pay the information-updating cost rely on random pairwise comparisons of losses), in the model herein the population dynamics is governed by a satisficing evolutionary protocol based on random agent-specific acceptable prediction errors.

[2] In most of the above respects, our model likewise differs from the one in Lines and Westerhoff (2012), which is closely related (and obtains qualitative results similar) to Lines and Westerhoff (2010). These two contributions employ the same macroeconomic model and specify inflation predictor switching following Brock and Hommes (1997). Yet in Lines and Westerhoff (2012) extrapolative expectations are paired with regressive expectations (which predict the return of inflation towards its normal value) as costly predictor instead of rational expectations. Even so, the associated cost is again supposed to reflect a predisposition effect in addition to information collection and processing.

(2) $$\pi_t = \pi_t^e - \alpha(u_t - u_n),$$

(3) $$g_t = m - \pi_t,$$

where u_t denotes the unemployment rate in period t, u_n represents the natural (or normal) unemployment rate, g_t is the output growth rate in period t, g_n is the natural (or normal) output growth rate, m stands for the nominal money growth rate, which is exogenously fixed, π_t is the actual rate of inflation in period t, π_t^e is the aggregate expected rate of inflation for period t, which is formed by agents in period $t-1$, and lower case Greek letters denote strictly positive parameters. Equation (1) is simply Okun's law, according to which changes in the rate of unemployment are negatively related to the rate of output growth, equation (2) is an expectations-augmented Phillips curve, and equation (3) is a simple linear aggregate demand relation.

Inserting (2) into (3) and the resulting expressions into (1), we get:

(4) $$u_t - u_n = \frac{1}{1+\alpha\beta}(u_{t-1} - u_n) + \frac{\beta}{1+\alpha\beta}(\pi_t^e - \pi_n),$$

where $\pi_n = m - g_n$ is the natural (or normal) inflation rate. Substituting (4) in (2), we obtain:

(5) $$\pi_t - \pi_n = -\frac{\alpha}{1+\alpha\beta}(u_{t-1} - u_n) + \frac{1}{1+\alpha\beta}(\pi_t^e - \pi_n).$$

Consequently, for a given vector of structural and policy parameters represented by $(\alpha, \beta, u_n, \pi_n = m - g_n)$, the state transition of the unemployment rate and the actual inflation rate depends not only on the current unemployment rate, but also on the aggregate expected inflation.

3.2. Inflation expectations formation

At a given period t, each agent either follows the costly *ex ante* perfect foresight strategy, so that her expected rate of inflation is given by:

(6) $$\pi_t^R = \pi_t,$$

or follows a costless *ex ante* extrapolative trend-following foresight strategy and form inflation expectations given by:

(7) $$\pi_t^E = \pi_{t-1} + \gamma(\pi_{t-1} - \pi_{t-2}),$$

where $\gamma \in \mathbb{R}_{++}$ is a parameter. Hence, again for comparability of results, the two available inflation foresight strategies are also the same as in Lines and Westerhoff (2010).

At a given period *t* there is a fraction $x_t \in [0,1] \subset \mathbb{R}$ of the population of agents, which may vary from one period to the next one, forming extrapolative trend-following inflation expectations, that is, they use the bounded rationality inflation predictor in (7). We refer to this type of agent as an *extrapolative forecaster*. The remaining fraction, $1-x_t$, is made up of fully informed agents that pay an exogenously fixed and strictly positive *ex ante* cost to play the perfect foresight strategy in (6). We refer to this type of agent as a *perfect forecaster*. Given this predictor distribution across the population of agents, the average aggregate expected rate of inflation rate can then be expressed as a weighted average of these two predictors as follows:

(8) $$\pi_t^e = x_{t-1}\pi_t^E + (1 - x_{t-1})\pi_t^R.$$

Therefore, the cost associated with perfect foresight is actually the cost associated with a potential heterogeneity in the choice of inflation forecasting strategy across agents. In fact, as shown in (8), playing the perfect inflation foresight strategy inevitably requires perfectly knowing the frequency distribution of inflation forecasting strategies across agents.

Using (6), (7) and (8), the gap between the average aggregate expected rate of inflation and the natural (or normal) rate of inflation can be expressed as follows:

(9) $$\pi_t^e - \pi_n = x_{t-1}[(1+\gamma)(\pi_{t-1} - \pi_n) - \gamma(\pi_{t-2} - \pi_n)] + (1 - x_{t-1})(\pi_t - \pi_n).$$

3.3 Average inflation expectations dynamics as evolutionarily satisficing learning

Although the frequency distribution of predictors is predetermined at a given period t, it evolves over time according to evolutionarily satisficing learning dynamics. Yet, as in Lines and Westerhoff (2010), we assume that agents strictly prefer inflation predictors with a high forecasting accuracy and, therefore, rely on squared prediction error as a (publicly observable) fitness measure. Following the nomenclature in Lines and Westerhoff (2010), the attractiveness of the extrapolative trend-following forecasting strategy can be defined as follows:

(10) $$a_t^E = -(\pi_t^E - \pi_t)^2.$$

Given (6), it follows that $\pi_t^R - \pi_t = 0$. Thus, the attractiveness of the perfect foresight strategy, as in Lines and Westerhoff (2010), is given by:

(11) $$a_t^R = -(\pi_t^R - \pi_t)^2 - \kappa = -\kappa,$$

where $\kappa \in \mathbb{R}_+$ stands for the exogenously fixed cost associated with playing the perfect foresight strategy, as this playing behavior does necessitate knowing the frequency distribution of inflation forecasting strategies across agents, which is not explicitly acknowledged in Lines and Westerhoff (2010). Moreover, in our model, more closely in accordance with the evidence in Branch (2004) that inflation prediction errors have to exceed some threshold to induce predictor switching, the attractiveness of a prediction strategy is compared with the attractiveness the agent considers acceptable.

Let us then describe the evolutionarily satisficing learning dynamics embedded in the law of motion of the proportion of extrapolative forecasters, x_t, which is in the spirit of the contributions of Herbert Simon.[3] As classically elaborated by Simon (1955, 1956), satisficing is a theory of choice centered on the process through which available alternatives are examined and evaluated. By conceiving of choice as intending to meet an *acceptability threshold* rather than to select the best of all possible alternatives, satisficing theory evidently contrasts with optimization theory. As Simon persuasively suggests, this contrast is analogous to 'looking for the sharpest needle in the haystack' (i.e., optimizing) versus 'looking for a needle sharp enough to sew with' (i.e., satisficing) (Simon, 1987, p. 244). Hence, as in Lines and Westerhoff (2010), in our model agents choose an inflation predictor from a set of alternative predictors whose *ex ante* cost ordering increases in a predictor's *ex ante* precision, whereas the proportion of agents in the population using a certain predictor varies positively with a formally precise measure of relative attractiveness. In our model, however, as it is quite plausibly assumed that there is heterogeneity in the choice idiosyncrasies of satisficing agents, an

[3] The formal derivation of this evolutionarily satisficing learning mechanism (without mutation) is based on Vega-Redondo (1996, p. 91).

acceptable range of prediction error is explicitly and separately modeled as randomly varying across agents.[4]

In our specification, an extrapolative forecaster i takes the current attractiveness given by (10) and compares it with the attractiveness she considers *acceptable*, which is denoted by a^i. Taking $t-1$ as the *current* period, if $|a^i| < |a^E_{t-1}|$, the extrapolative forecasters i considers changing her strategy for forming inflation expectations in t, that is, the extrapolative forecaster i becomes a strategy reviser. The attractiveness that is deemed acceptable by an agent depends, *inter alia*, on her idiosyncrasies. Reasonably, we assume that the attribute of acceptable attractiveness is randomly and independently determined across agents and over time. More precisely, we assume that $|a^i| = -a^i$ is a random variable with cumulative distribution function $F:\mathbb{R}_+ \to [0,1] \subset \mathbb{R}$ which is continuously differentiable. Consequently, the probability that a randomly drawn agent i in the subpopulation of extrapolative forecasters will consider that the currently observed attractiveness of her foresight strategy is unacceptable is simply:

(12) $$\Pr\left(|a^i| < |a^E_{t-1}|\right) = \Pr(-a^i < -a^E_{t-1}) = F(-a^E_{t-1}).$$

The measure of extrapolative forecasters who become perfect forecasters is then given by:

(13) $$x_{t-1} F(-a^E_{t-1}).$$

Analogously, the measure of perfect forecasters who becomes extrapolative forecasters is given by:

(14) $$(1-x_{t-1})\Pr\left(|a^i| < |a^R_{t-1}|\right) = (1-x_{t-1})\Pr(-a^i < -a^R_{t-1}) = (1-x_{t-1})F(-a^R_{t-1}).$$

Therefore, the subtraction of (13) from (14) and use of (7), (10) and (11) yields the following evolutionarily satisficing learning dynamic:

(15) $$x_t - x_{t-1} = (1-x_{t-1})F(\kappa) - x_{t-1}F\left(g(\pi_{t-1}, y_{t-1}, z_{t-1})\right),$$

where:

(16) $$g(\pi_{t-1}, y_{t-1}, z_{t-1}) \equiv -a^E_{t-1} = \left[(1+\gamma)(y_{t-1}-\pi_n) - \gamma(z_{t-1}-\pi_n) - (\pi_{t-1}-\pi_n)\right]^2,$$

which is specified by using (7), (10) and the following auxiliary variables:

(17) $$y_t = \pi_{t-1} \text{ and } z_t = y_{t-1}.$$

As it turns out, an increase in the relative attractiveness of the extrapolative trend-following foresight strategy in the current period (a decrease in $g(\cdot)$) leads to an increase in the proportion of agents playing this inflation foresight strategy in the next period, whereas the opposite occurs when the relative attractiveness of the perfect inflation foresight strategy rises (a decrease in κ). Thus, the evolutionarily satisficing learning dynamics in (15) reflects the operation of a selection mechanism according to which the proportion of agents playing a given inflation forecasting strategy at a given period in time varies positively with the relative fitness or attractiveness of such strategy.

[4] We should note that the model in Lines and Westerhoff (2010) also features "a permanent evolutionary competition between the prediction rules" (p. 247). The authors explicitly allude to Simon's suggestion that agents should be seen as boundedly rational when they recognize that because agents "lack the cognitive capabilities to derive fully optimal actions…they strive to do the right thing" (p. 249). However, although the gist of Simon's approach to choice behavior is that decision makers satisfice rather than optimize, the model in Lines and Westerhoff (2010) does not incorporate the evidence in Branch (2004) that inflation forecast errors have to cross some threshold in order to induce a switch by appealing to Simon's fitting notion of satisficing behavior.

In order to further enrich the predictor switching process and gain in analytical generality, we also consider the possibility that the evolutionarily satisficing learning dynamics in (15) operate in the presence of a perturbation mechanism, analogous to mutation in natural environments. In a biological setting, mutation is interpreted literally as consisting of random changes in genetic codes. In economic settings, as pointed out by Samuelson (1997, Ch. 7), mutation refers to a situation in which a player refrains from comparing payoffs and switches strategy at random. Thus, the present extension features mutation as an exogenous perturbation in the evolutionarily satisficing learning mechanism leading some decision makers to choose an inflation foresight strategy at random. Following Kandori, Mailath and Rob (1993), two rationales can be provided to this random choice. First, at every period a given measure of agents exit the economy with some (exogenously fixed) probability and are replaced one-to-one with new agents who happen to know nothing about (or are relatively inexperienced in) the decision-making process in question. Second, each agent simply "experiments" occasionally with some exogenously fixed probability. More generally, this noise component could also capture the effect of sporadic exogenous institutional factors such as changes of administration in the monetary authority or other changes in the policy-making framework. However, we would have to further assume that these changes do not affect the structure of the basic macroeconomic model in (1)-(3) (i.e., these institutional changes do not affect the structural stability of the model in (1)-(3)).

Drawing on Gale, Binmore and Samuelson (1995), mutation can be incorporated into the evolutionarily satisficing learning mechanism in (15) as follows. Let $\varepsilon \in (0,1) \subset \mathbb{R}$ be the measure of mutant agents that choose an inflation foresight strategy in a given revision period independently both of their acceptable forecast error and of the attractiveness of the inflation foresight strategy just played. Thus, there are εx_{t-1} extrapolative forecasters and $\varepsilon(1-x_{t-1})$ perfect forecasters behaving as mutants. As we assume that mutant agents choose either one or the other of the available inflation foresight strategies with equal probability, there are then $\varepsilon x_{t-1}/2$ extrapolative forecaster mutant agents and $\varepsilon(1-x_{t-1})/2$ perfect forecaster mutant agents changing foresight strategy. The net flow of mutant agents becoming extrapolative forecaster agents in a given revision period, which can be either positive or negative, is then the following:

(18) $$\varepsilon\left(1-x_{t-1}\right)\frac{1}{2}-\varepsilon x_{t-1}\frac{1}{2}=\varepsilon\left(\frac{1}{2}-x_{t-1}\right).$$

Following Gale, Binmore and Samuelson (1995), this exogenous perturbation can be linearly added to the evolutionary protocol in (15) to yield the following *noisy evolutionarily satisficing learning dynamic* of the frequency distribution of inflation foresight strategies across agents:

(19) $$x_t - x_{t-1} = (1-\varepsilon)\left[(1-x_{t-1})F(\kappa) - x_{t-1}F\left(g(y_{t-1}, z_{t-1}, \pi_{t-1})\right)\right] + \varepsilon\left(\frac{1}{2}-x_{t-1}\right).$$

4. Behavior of the model in the long-run evolutionary equilibrium

Note that the macrodynamics of the rates of output growth, unemployment and inflation is coupled with the microdynamics of the frequency distribution of inflation foresight strategies across firms. In fact, plugging (9) into (5) and considering the auxiliary variables defined in (17), we can establish that:

(20) $$\pi_t - \pi_n = \frac{1}{\alpha\beta + x_{t-1}}\left\{-\alpha(u_{t-1}-u_n) + x_{t-1}\left[(1+\gamma)(\pi_{t-1}-\pi_n) - \gamma(y_{t-1}-\pi_n)\right]\right\}.$$

Meanwhile, inserting (20) into (9) and the resulting expressions into (4), we get:

(21) $$u_t - u_n = \frac{x_{t-1}}{\alpha\beta + x_{t-1}}\left\{u_{t-1}-u_n + \beta\left[(1+\gamma)(\pi_{t-1}-\pi_n) - \gamma(y_{t-1}-\pi_n)\right]\right\}.$$

As a result, the state transition of the economy is determined by the first-order difference equation system given by (19)-(21) and the auxiliary variables defined in (17), which state space is given by $\Theta = \{(x_t, \pi_t, u_t, y_{t-1}, z_{t-1}) \in \mathbb{R}^5 : 0 \le x_t \le 1, 0 \le u_t \le 1\}$.

In the following proposition we establish the existence and uniqueness of the evolutionary long-run equilibrium of the dynamic system given by (19)-(21) and the auxiliary variables defined in (17).

Proposition 1: *For a given vector of parameters* $(\alpha, \beta, \gamma, u_n, \pi_n = m - g_n, \varepsilon)$, *the dynamic system given by* (19)-(21) *and the auxiliary variables in* (17) *has a unique long-run evolutionary equilibrium which is given by* $(x^*, \pi_n, u_n, \pi_n, \pi_n) \in \Theta$, *with* $x^* = \dfrac{(1-\varepsilon)F(\kappa) + \varepsilon/2}{(1-\varepsilon)F(\kappa) + \varepsilon} \in (1/2, 1] \subset \mathbb{R}$.

Proof: See Appendix 1.

In the long-run evolutionary equilibrium, therefore, the extrapolative forecasting strategy is either the only strategy to survive ($x^* = 1$), in the absence of mutant behavior ($\varepsilon = 0$), or the most played foresight strategy ($1/2 < x^* < 1$), in the presence of mutation $(0 < \varepsilon < 1)$. Nonetheless, note that the long-run equilibrium values of the endogenous macroeconomic variables (π_t, u_t, and g_t) are the same in either case. Although pervasive extrapolative inflation foresight (all agents eventually decide to follow the boundedly rational inflation forecasting strategy) is an evolutionarily satisficing long-run equilibrium outcome when mutation is absent, this equilibrium configuration is nonetheless characterized by the natural values of the rates of inflation, output growth and unemployment being achieved. The intuition is straightforward: when mutation is absent, once the economy converges to the long-run equilibrium, the extrapolative foresight strategy fares as well as the perfect foresight strategy, but it is less costly. In fact, given the nature of the information imperfection which imposes a cost on perfect foresight, all agents following the perfect foresight strategy ($x = 0$) is not a long-run equilibrium. The reason is that the substance of this result ($x = 0$) is that all agents are paying the cost to learn the measure of a behavioral heterogeneity that simply does not exist. Meanwhile, when mutation is absent, $x^* = 1$ is a long-run equilibrium, for the substance of this result is that no agent is paying the cost to find out the measure of a non-existing behavioral heterogeneity. Therefore, by the same logic, $1/2 < x^* < 1$ is a long-run equilibrium in the presence of mutant agents because in this situation there are agents paying a behavioral heterogeneity cost that actually has to be paid.

As in Lines and Westerhoff (2010), therefore, in our model the unique long-run equilibrium solution features the natural (or normal) values of the macroeconomic variables (rates of inflation, unemployment and output growth). Moreover, in each model the long-run equilibrium solution for the frequency distribution of inflation foresight strategies across agents depends on the respective key parameters: information cost associated with perfect foresight and intensity of choice (in Lines and Westerhoff, 2010) and information cost associated with perfect foresight and mutation rate (in our model). As elaborated in Section 2, a strictly positive mutation rate is nonetheless analytically analogous to a finite strictly positive intensity of choice, as they both imply that, due to some source of randomness, not all agents will select the best performing predictor as reflected in the respective relative net attractiveness. In fact, we would suggest that the rationales we have put forward for the presence of mutation may apply to (and add to the existing justifications for) a finite strictly positive intensity of choice.

While in our model the proportion of extrapolative trend-follower agents in the long-run equilibrium is given by the expression for x^* in Proposition 1, the respective value in Lines and Westerhoff (2010, p. 250) is given by:

(22) $$\mu = \frac{1}{1 + e^{-\lambda\kappa}} \in (1/2, 1] \subset \mathbb{R},$$

where $\lambda \in \mathbb{R}_+$ is the intensity of choice and $\kappa \in \mathbb{R}_+$ is also the exogenously fixed cost associated with playing the perfect foresight strategy. In the limit for $\lambda \to +\infty$, it follows that all agents choose the extrapolative trend-following strategy in the long-run equilibrium ($\mu \to 1^-$). Therefore, as when mutation is absent in our model ($\varepsilon = 0$), which is equivalent to the intensity of choice tending to positive infinity in the Lines and Westerhoff (2010) model, the perfect foresight strategy also ceases to be played in the long-run equilibrium ($x^* = 1$); and this is so no matter how low happens to be any strictly positive *ex ante* cost associated with playing the perfect foresight strategy. The intuition is that in the limit for either $\lambda \to +\infty$ or $\varepsilon \to 0^+$, the long-run equilibrium is characterized by both foresight strategies being equally successful in predicting inflation, but whereas the extrapolative trend-following strategy does so without any *ex ante* cost, the perfect foresight strategy does so at a strictly positive *ex ante* cost. Consequently, except in the limit for $\lambda \to +\infty$ (in the model in Lines and Westerhoff, 2010) or in the limit for $\varepsilon \to 0^+$ (in our model), in the unique long-run equilibrium, which is common to both models, although both strategies are effective in predicting inflation, there is nonetheless a tail of the population of agents that selects the least attractive inflation predictor due to randomness. And, per (10) and (11), this least attractive inflation predictor is the perfect foresight one. Meanwhile, as in the limit for $\varepsilon \to 1^-$ in our model, in the limit for $\lambda \to 0^+$ in the model in Lines and Westerhoff (2010), it follows that exactly half of the agents play each inflation foresight strategy in the long-run equilibrium. The intuition is straightforward: in either case, agents' choice of inflation predictor becomes completely random.

Additionally, in the limit for $\kappa \to 0^+$, the result obtains that $\mu = x^* = 1/2$, which substance is that the two available foresight strategies become equivalent in the long-run equilibrium due to their sharing of the same level of attractiveness. Meanwhile, in the limit for $\kappa \to +\infty$, it follows that $\mu \to 1^-$, whereas $x^* \to 1 - \varepsilon / 2$. Therefore, in the limit for the information cost associated with the perfect foresight strategy tending to positive infinity, while such strategy fails to survive in the long-run equilibrium in the Lines and Westehoff (2010) model, in our model the occurrence of mutation prevents the extinction of such strategy in the long-run equilibrium, albeit in a minority position for any $\varepsilon \in (0,1) \subset \mathbb{R}$. In effect, the randomness introduced by mutant behavior comes to rescue the perfect foresight strategy in such equilibrium configuration where the extrapolative trend-following strategy is permanently relatively more attractive. In this situation, intuitively, the higher the mutation rate, the lower the predominance of the extrapolative trend-following strategy in the long-run equilibrium.

However, a natural question that arises regards whether the dynamic described in (19)-(21) and (17) can take the economy to the long-run equilibrium configuration established in Proposition 1. For instance, in an economy with no mutant agents and where initially there is heterogeneity in inflation expectations, can all agents eventually learn to perfectly predict the rate of inflation, but without having to play the costly *ex ante* perfect foresight strategy? The answer is yes, as formally established in the following proposition.

Proposition 2: *For a given vector of parameters* $(\alpha, \beta, \gamma, u_n, \pi_n = m - g_n, \varepsilon)$, *the unique long-run evolutionary equilibrium* given by $(x^*, \pi_n, u_n, \pi_n, \pi_n) \in \Theta$ *of the dynamic system given by* (19)-(21) *and the auxiliary variables in* (17) *is locally asymptotically stable if* $\gamma < \dfrac{\alpha\beta + x^*}{2x^*} \equiv \gamma_c$.

<u>Proof</u>: See Appendix 2.

Consequently, the predictor switching game explored in this paper is subject to evolutionary learning dynamics that can take it to a long-run equilibrium configuration in which, albeit either the majority or even all agents choose to play the extrapolative foresight strategy to form inflation expectations, the rates of inflation, output growth and unemployment all achieve their natural values.

Thus, inflation expectations that are persistently heterogeneous need not prevent the successful conduct of macroeconomic stabilization policy through the setting of the growth rate of nominal money by the monetary authority. Besides, when there is mutation, so that the two available inflation forecasting strategies survive in the long-run equilibrium solution (yet the extrapolative forecasting strategy predominates), such long-run equilibrium level of heterogeneity varies expectedly, considering (A-5) in Appendix 1, with the *ex ante* cost associated with perfect foresight and the measure of mutant agents:

(23) $$\frac{\partial x^*}{\partial \kappa} = \frac{\varepsilon(1-\varepsilon)F'(\kappa)}{2[(1-\varepsilon)F(\kappa)+\varepsilon]^2} > 0,$$

(24) $$\frac{\partial x^*}{\partial \varepsilon} = \frac{-F(\kappa)}{2[(1-\varepsilon)F(\kappa)]+\varepsilon]^2} < 0.$$

Intuitively, in the long-run equilibrium configuration with coexistence of the two inflation forecasting strategies, the higher the *ex ante* cost associated with perfect foresight, the greater the predominance of agents playing the extrapolative foresight strategy (recall from Proposition 1 that, in the presence of mutation, the extrapolative forecasting strategy is the most played strategy in the long-run evolutionary equilibrium ($1/2 < x^* < 1$). Yet the predominance of agents who select the extrapolative foresight strategy varies negatively with the rate of mutation. The intuition is that, the greater the random noise in the selection process of inflation forecasting strategy, the greater has to be the proportion of perfect forecasters in the long-run equilibrium to ensure that the aggregate expected inflation is consistent with the natural rate of inflation.[5]

While in our model the unique long-run evolutionary equilibrium is locally asymptotically stable if the extrapolation parameter γ is lower than the critical value given by the expression for γ_c in Proposition 2, the respective critical value in the Lines and Westerhoff (2010, p. 250) model is given by $\gamma_c^{LW} \equiv \frac{\alpha\beta+\mu}{2\mu}$. Intuitively, in both models the corresponding critical value for the extrapolation parameter varies positively with the reaction parameters α and β, and negatively with the long-run equilibrium value of the proportion of extrapolative trend-following agents, which in turn varies positively with the information cost associated with perfect foresight. Consequently, in the limit for $\lambda \to +\infty$ (in the Lines and Westerhoff (2010) model) or in the limit for $\varepsilon \to 0^+$ (in our model), as all agents choose the extrapolative trend-following strategy in the common long-run equilibrium, the likewise common basin of attraction (or corridor of stability) given by $\gamma_c^{LW} = \gamma_c = 0.5(1+\alpha\beta)$ attains its smallest possible size. Meanwhile, as in the limit for $\varepsilon \to 1^-$ in our model, in the limit for $\lambda \to 0^+$ in the model in Lines and Westerhoff (2010) (or in the limit for $\kappa \to 0^+$ in both models), since exactly half of the agents choose each inflation foresight strategy in the long-run equilibrium, the resulting common basin of attraction given by $\gamma_c = \gamma_c^{LW} = 0.5 + \alpha\beta$ attains its largest possible size. Interestingly, therefore, in both models a larger extent of a specific kind of randomness makes for a locally more stability-prone long-run equilibrium. In fact, in both models the critical value for the extrapolation parameter varies negatively with the long-run equilibrium value of the proportion of extrapolative trend-following agents, which in turn varies negatively with the extent of randomness in inflation predictor switching as measured respectively by the mutation rate and intensity of choice (recall that our model features three layers of randomness, of which the distribution of tolerable prediction inaccuracy across agents is the only one unrelated to mutation).

[5] Moreover, since the stable eigenvalue corresponding to the evolutionarily satisficing learning dynamics is given by $\lambda_1 = (1-\varepsilon)[1-F(\kappa)]$ (see Appendix 2), the lower is the measure of mutant agents or the higher is the *ex ante* cost associated with perfect inflation foresight, the lower is the speed of convergence of the frequency distribution of inflation forecasting strategies to its long-run equilibrium configuration. The intuition is that perfect forecasters play a crucial role in guiding the evolutionarily satisficing learning dynamics.

Moreover, in our model, when mutant behavior is absent ($\varepsilon = 0$), the resulting critical value for the extrapolation parameter is given by $\gamma_c = 0.5(1+\alpha\beta)$, which is strictly lower than $\gamma_c^{LW} = 0.5\left[1+(1+e^{-\lambda\kappa})\alpha\beta\right]$ for any $0 \le \lambda < +\infty$ and $0 \le \kappa < +\infty$ in the Lines and Westerhoff (2010) model. In words, in the absence of mutation, the long-run equilibrium in our model is more locally instability-prone than the long-run equilibrium in the Lines and Westerhoff (2010) model with any strictly finite intensity of choice and information cost. In the latter model, however, either in the limit for $\lambda \to +\infty$ or in the limit for $\kappa \to +\infty$, the critical value for the extrapolation parameter is given by $\gamma_c^{LW} = 0.5(1+\alpha\beta)$, which is strictly lower than $\gamma_c = (\alpha\beta + x^*)/2x^*$ for any $0 < \varepsilon \le 1$ and $0 \le \kappa \le +\infty$ in our model, since in that case $0 < x^* < 1-\varepsilon/2 < 1$. In words, when the model in Lines and Westerhoff (2010) features either the intensity of choice or the information cost tending to positive infinity, its long-run equilibrium becomes more locally instability-prone than the long-run equilibrium in our model with any strictly positive mutation rate and any (even if extreme-value) information cost.[6]

Meanwhile, when the information cost attached to perfect foresight, the intensity of choice and the rate of mutation all assume non-extreme values (that is, $0 < \kappa < +\infty$, $0 < \lambda < +\infty$ and $0 < \varepsilon < 1$), so that $1/2 < \mu < 1$ and $1/2 < x^* < 1$, respectively, in both models the extent of the long-run predominance of extrapolative trend-followers varies positively (and hence the size of the basin of attraction of the respective unique long-run equilibrium varies negatively) with the information cost associated with perfect foresight. In fact, as discussed earlier, in the limit for $\kappa \to 0^+$, it follows that $\mu \to 1/2$ and $x^* \to 1/2$ in that limit. Meanwhile, in the limit for $\kappa \to +\infty$, it follows that $\mu \to 1$ and $x^* \to 1-\varepsilon/2$ in that limit. Moreover, in the limit for $\kappa \to +\infty$, it then follows that $(\mu - x^*) \to \varepsilon/2 > 0$ (which strictly positive value, in the presence of mutation, intuitively varies positively with the mutation rate) in that limit. Consequently, it follows that $1/2 < x^* < \mu < 1$ for a sufficiently high value of the information cost associated with perfect foresight, which implication that $\gamma_c > \gamma_c^{LW}$ means that the unique long-run equilibrium in our model is locally more stability-prone than the (also unique) long-run equilibrium in the Lines and Westerhoff (2010) model when the information cost associated with perfect foresight becomes relatively high.

Note, however, that it is not possible to unambiguously evaluate how the difference in terms of local stability properties as expressed by $\gamma_c - \gamma_c^{LW}$ varies with the information cost associated with perfect foresight. In fact, this difference is given by:

(25) $$\gamma_c - \gamma_c^{LW} = \frac{1}{2}\alpha\beta\left[\frac{\varepsilon}{\varepsilon - 2(\varepsilon-1)F(\kappa)} - e^{-\lambda\kappa}\right],$$

which varies with the information cost associated with perfect foresight according to:

(26) $$\frac{\partial(\gamma_c - \gamma_c^{LW})}{\partial\kappa} = \alpha\beta\left\{\frac{1}{2}\lambda e^{-\lambda\kappa} - \frac{\varepsilon(1-\varepsilon)F'(\kappa)}{[\varepsilon + 2(1-\varepsilon)F(\kappa)]^2}\right\},$$

which sign is ambiguous. Considering (A-5), (22) and (25), in the limit for $\kappa \to 0^+$, it follows that both $\mu \to 1/2$ and $x^* \to 1/2$ in that limit, so that $(\gamma_c - \gamma_c^{LW}) \to 0$ in that limit, which implies that in both

[6] Arguably, mutant agents, who are presently assumed to play either one or the other of the inflation foresight strategies with equal probability, might eventually switch to playing the extrapolative strategy with higher probability as the information cost associated with perfect foresight would become to be considered too high. In fact, an agent might even become reluctant to intentionally behave as mutant for experimentation reasons in such situation. However interesting, we abstract from these analytical possibilities in our paper and leave them for future research. Incidentally, the robust empirical evidence provided in Branch (2004, 2007) is that the intensity of choice and information costs are both finite.

models the common basin of attraction of the common unique long-run equilibrium attains its largest possible size in that limit. Meanwhile, in the limit for $\kappa \to +\infty$, it follows that $\mu \to 1$ and $x^* \to 1 - \varepsilon / 2$ in that limit, which then implies that $(\mu - x^*) \to \varepsilon / 2 > 0$ (which intuitively varies positively with the rate of mutation) in that limit. As a result, it follows that $x^* < \mu$ (and hence $\gamma_c > \gamma_c^{LW}$) for a sufficiently high value of the information cost associated with perfect foresight.

5. Conclusion

The empirical motivation for this paper comes from two series of evidence on the formation of inflation expectations. First, inflation expectations are endogenously time-varying, persistently heterogeneous across decision makers and formed mostly through boundedly rational procedures. Second, inflation forecast errors may have to pass some threshold before decision makers decide to abandon their previously selected inflation predictor. Against this backdrop, the intended theoretical contribution of this paper lies primarily in the exploration of macroeconomic implications of heterogeneous inflation expectations which are endogenously time-varying according to evolutionarily satisficing learning dynamics subject to a random noise analogous to mutation in natural environments. In the alternative modeling approach to inertia in inflation predictor switching devised in this paper, a key role is played by Herbert Simon's understanding of choice as intending to meet some acceptability threshold rather than to select the best of all possible alternatives, hence the core idea of satisficing instead of optimal choice.

In a specification of inflation predictor choice featuring perfect foresight in competition with extrapolative trend-following, we find that the frequency distribution of prediction strategies across agents in the long-run equilibrium depends on the relative importance of an objective measure of accuracy in inflation prediction as compared to random components operating at different layers of the predictor choice process. Also, there is a definite relationship between the frequency distribution of inflation prediction strategies and the local stability properties as two distinct features of the unique long-run equilibrium, which is characterized by the endogenous macroeconomic variables at their natural values and at least exactly one half of the population of agents playing the extrapolative strategy. In fact, the higher the proportion of extrapolative trend-following agents in the long-run equilibrium, the less locally stability-prone it is such evolutionary equilibrium. Meanwhile, the more inflation predictor choice is subject to random disturbance, the lower is the proportion of extrapolative trend-following agents in the long-run equilibrium outcome. Thus, randomness in inflation predictor choice plays a locally stabilizing role in our suggested evolutionarily satisficing learning dynamics.

Finally, we have suggested the notion of evolutionarily satisficing learning in the context of a given dynamic macroeconomic model and applied it to the exploration of the specific but relevant issue of whether endogenously time-varying (and eventually persistent) heterogeneity in strategies to predict inflation is necessarily detrimental to successful stabilization policy, which answer was found to be no. However, the likely prospect that other instances of choice behavior are also subject to satisficing learning deserves future investigation.

References

Adam, K. (2007) Experimental evidence on the persistence of output and inflation, *The Economic Journal*, 117(520), 603–636.

Blanchflower, D. G. and MacCoille, C. (2009) The formation of inflation expectations: an empirical analysis for the UK, *NBER Working Paper* 15388, September.

Branch, W. A. (2004) The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations, *The Economic Journal*, 114(497), pp. 592–621.

Branch, W. A. (2007) Sticky information and model uncertainty in survey data on inflation expectations, *Journal of Economic Dynamics and Control*, 31(1), pp. 245–276.

Branch, W. and McGough, B. (2010) Dynamic predictors selection in a New Keynesian model with heterogeneous expectations, *Journal of Economic Dynamics and Control*, 34(8), pp. 1492–1508.

Brazier, A., Harrison, R., King, M. and Yates, T. (2008) The danger of inflating expectations of macroeconomic stability: heuristic switching in an overlapping generations monetary model, *International Journal of Central Banking*, 4, pp. 219–254.

Brock, W. A. and de Fontnouvelle, P. (2000) Expectational diversity in monetary economies, *Journal of Economic Dynamics & Control*, 24(5-7), pp. 725-759.

Brock, W. A. and Hommes, C. H. (1997) A Rational route to randomness, *Econometrica*, 65, pp. 1059–1160.

De Grauwe, P. (2010) Animal spirits and monetary policy, *Economic Theory*, pp. 1-35.

Duffy, J. (2008) Experimental macroeconomics, in: Durlauf, S. and Blume, L. (eds.) *New Palgrave Dictionary of Economics*, New York: Palgrave Macmillan.

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

Farebrother, R. W. (1973) Simplified Samuelson conditions for cubic and quartic equations, *The Manchester School*, pp. 396-400.

Gale, J., Binmore, K. and Samuelson, L. (1995) Learning to be imperfect: the ultimatum game, *Games and Economic Behavior*, 8(1), pp. 56-90.

Hommes, C. H. (2011) The heterogeneous expectations hypothesis: some evidence from the lab, *Journal of Economic Dynamics and Control*, 35, pp. 1–24.

Hommes, C. H. (2013) *Behavioral rationality and heterogeneous expectations in complex economic systems*, Cambridge, MA: Cambridge University Press.

Kandori, M., Mailath, G. J. and Rob, R. (1993) Learning, mutation, and long run equilibria in games, *Econometrica*, 61(1), pp. 29-56.

Lima, G. T. and Silveira, J. J. (2015) Monetary neutrality under evolutionary dominance of bounded rationality, *Economic Inquiry*, 53(2), pp. 1108-1131.

Lines, M. and Westerhoff, F. (2010) Inflation expectations and macroeconomic dynamics: the case of rational versus extrapolative expectations, *Journal of Economic Dynamics and Control*, 34, pp. 246-257.

Mankiw, N. G. and Reis, R. (2002) Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve, *Quarterly Journal of Economics*, 117(4), pp. 1295-1328.

Mankiw, N. G., Reis, R. and Wolfers, J. (2004) Disagreement about inflation expectations, in: *NBER Macroeconomics Annual 2003*, v. 18, pp. 209-270.

Pfajfar, D. and Zakelj, B. (2011) Inflation expectations and monetary policy design: evidence from the laboratory, *Tilburg University CentER Discussion Papers*, 2011-2091.

Pfajfar, D. and Zakelj, B. (2014) Experimental evidence on inflation expectation formation, *Journal of Economic Dynamics and Control*, 44, pp. 147-168.

Roos, M. W. and Luhan, W. J. (2013) Information, learning and expectations in an experimental model economy, *Economica*, 80(319), pp. 513-31.

Samuelson, L. (1997) *Evolutionary Games and Equilibrium Selection*, Cambridge, MA: The MIT Press.

Simon H. A. (1955) A behavioral model of rational choice, *Quarterly Journal of Economics*, 69(1), pp. 99–118.

Simon H. A. (1956) Rational choice and the structure of the environment, *Psychological Review*, 63, pp. 129–138.

Simon H. A. (1987) Satisficing, in: Eatwell J., Milgate, M. and Newman, P. (eds.) *The New Palgrave: A Dictionary of Economics*, Vol. 4, New York: Stockton Press.

Vega-Redondo, F. (1996) *Evolution, Games and Economic Behaviour*, Oxford, UK: Oxford University Press.

Appendix 1: Proof of Proposition 1

An evolutionary long-run equilibrium of the dynamic system given by (19)-(21) and (17) should satisfy the conditions given by $x_t = x_{t-1} \equiv x^*$, $u_t = u_{t-1} \equiv u^*$, and $\pi_t = \pi_{t-1} = \pi_{t-2} = \pi_{t-3} \equiv \pi^*$ for any $t \in \{1, 2, 3, ...\} \subset \mathbb{N}$. Therefore, we can use (20) and (21) to establish that:

(A-1) $$\alpha\beta(\pi^* - \pi_n) + \alpha(u^* - u_n) = 0,$$

(A-2) $$\beta x^*(\pi^* - \pi_n) - \alpha\beta(u^* - u_n) = 0.$$

It is easy to see that for any given $x^* \in [0,1] \subset \mathbb{R}$ there is a unique solution for the homogeneous linear system (A-1)-(A-2), given by $\pi^* - \pi_n = u^* - u_n = 0$, such that:

(A-3) $$\pi^* = \pi_n \text{ and } u^* = u_n.$$

Considering (A-3), it follows that $\pi_t = \pi_{t-1} = \pi_{t-2} = \pi_{t-3} = \pi_n$. Therefore, based on (17), we have that $y_{t-1} = \pi_{t-2} = \pi_n$ and $z_{t-1} = y_{t-2} = \pi_{t-3} = \pi_n$ in the long-run evolutionary equilibrium. As a result, considering (16), it follows that $F\left(-a_{t-1}^E\right) = F\left(g(\pi_{t-1}, y_{t-1}, z_{t-1})\right) = F(0) = 0$ in the long-run equilibrium solution. From the noisy evolutionarily satisficing learning dynamic in (19), the following condition has then to be satisfied:

(A-4) $$(1-\varepsilon)(1-x^*)F(\kappa) + \varepsilon\left(\frac{1}{2} - x^*\right) = 0.$$

The solution for this condition is given by:

(A-5) $$x^* = \frac{(1-\varepsilon)F(\kappa) + \varepsilon/2}{(1-\varepsilon)F(\kappa) + \varepsilon}.$$

Let us prove that $x^* \in (1/2, 1] \subset \mathbb{R}$ for all $\varepsilon \in [0,1) \subset \mathbb{R}$. Since $\varepsilon \in [0,1) \subset \mathbb{R}$ and $F(\kappa) \subset (0,1) \subset \mathbb{R}$, it follows that $(1-\varepsilon)F(\kappa) > 0$. Therefore, we can use the latter inequality to establish that:

(A-6) $$2(1-\varepsilon)F(\kappa) + \varepsilon > (1-\varepsilon)F(\kappa) + \varepsilon \Rightarrow x^* = \frac{(1-\varepsilon)F(\kappa) + \varepsilon/2}{(1-\varepsilon)F(\kappa) + \varepsilon} > \frac{1}{2},$$

for all $\varepsilon \in [0,1) \subset \mathbb{R}$, and:

(A-7) $$(1-\varepsilon)F(\kappa) + \varepsilon/2 < (1-\varepsilon)F(\kappa) + \varepsilon \Rightarrow x^* = \frac{(1-\varepsilon)F(\kappa)] + \varepsilon/2}{(1-\varepsilon)F(\kappa)] + \varepsilon} < 1,$$

for all $\varepsilon \in (0,1) \subset \mathbb{R}$. Finally, it is easy to check that $x^* = 1$ for $\varepsilon = 0$. This completes the proof that there is one, and only one, equilibrium solution $(x^*, \pi_n, u_n, \pi_n, \pi_n) \in \Theta$ of the system represented by (19)-(21) and (17). □

Appendix 2: Proof of Proposition 2

Consider the Jacobian matrix of the linearization around the evolutionary long-run equilibrium of the system given by (19)-(21) and (17):

$$\text{(A-8)} \qquad J(x^*, \pi_n, u_n, \pi_n, \pi_n) = \begin{bmatrix} (1-\varepsilon)(1-F(\kappa)) & 0 & 0 & 0 & 0 \\ 0 & \dfrac{(1+\gamma)x^*}{\alpha\beta + x^*} & \dfrac{-\alpha}{\alpha\beta + x^*} & \dfrac{-\gamma x^*}{\alpha\beta + x^*} & 0 \\ 0 & \dfrac{\beta(1+\gamma)x^*}{\alpha\beta + x^*} & \dfrac{x^*}{\alpha\beta + x^*} & \dfrac{-\beta\gamma x^*}{\alpha\beta + x^*} & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

Let λ be an eigenvalue of the Jacobian matrix (A-8) and I the 5×5 identity matrix. We can then set the following characteristic equation of the linearization around the equilibrium:

$$\text{(A-9)} \qquad |J - \lambda I| = \begin{bmatrix} (1-\varepsilon)(1-F(\kappa)) - \lambda & 0 & 0 & 0 & 0 \\ 0 & \dfrac{(1+\gamma)x^*}{\alpha\beta + x^*} - \lambda & \dfrac{-\alpha}{\alpha\beta + x^*} & \dfrac{-\gamma x^*}{\alpha\beta + x^*} & 0 \\ 0 & \dfrac{\beta(1+\gamma)x^*}{\alpha\beta + x^*} & \dfrac{x^*}{\alpha\beta + x^*} - \lambda & \dfrac{-\beta\gamma x^*}{\alpha\beta + x^*} & 0 \\ 0 & 1 & 0 & -\lambda & 0 \\ 0 & 0 & 0 & 1 & -\lambda \end{bmatrix}.$$

We can use the Laplace expansion to express (A-9) as follows:

$$\text{(A-10)} \qquad |J - \lambda I| = \left[(1-\varepsilon)(1-F(\kappa)) - \lambda\right]\lambda \begin{vmatrix} \dfrac{(1+\gamma)x^*}{\alpha\beta + x^*} - \lambda & \dfrac{-\alpha}{\alpha\beta + x^*} & \dfrac{-\gamma x^*}{\alpha\beta + x^*} \\ \dfrac{\beta(1+\gamma)x^*}{\alpha\beta + x^*} & \dfrac{x^*}{\alpha\beta + x^*} - \lambda & \dfrac{-\beta\gamma x^*}{\alpha\beta + x^*} \\ 1 & 0 & -\lambda \end{vmatrix} = 0.$$

Consequently, two of the eigenvalues can already be explicitly determined, which are given by $\lambda_1 = (1-\varepsilon)[1-F(\kappa)]$ and $\lambda_2 = 0$. Given that $\varepsilon \in [0,1) \subset \mathbb{R}$ and $F(\kappa) \subset (0,1) \subset \mathbb{R}$, it follows that $0 < \lambda_1 < 1$.

Considering (A-10), the remaining eigenvalues have to satisfy:

(A-11) $$\begin{vmatrix} \frac{(1+\gamma)x^*}{\alpha\beta+x^*}-\lambda & \frac{-\alpha}{\alpha\beta+x^*} & \frac{-\gamma x^*}{\alpha\beta+x^*} \\ \frac{\beta(1+\gamma)x^*}{\alpha\beta+x^*} & \frac{x^*}{\alpha\beta+x^*}-\lambda & \frac{-\beta\gamma x^*}{\alpha\beta+x^*} \\ 1 & 0 & -\lambda \end{vmatrix} = 0,$$

which can be re-written as:

(A-11-a) $\lambda^3 - a\lambda^2 - b\lambda - c = 0$,

where:

(A-12) $a = \frac{(2+\gamma)x^*}{\alpha\beta+x^*}$, $b = \frac{(1+2\gamma)x^*}{\alpha\beta+x^*}$, and $c = \frac{\gamma x^*}{\alpha\beta+x^*}$.

Given (A-11-a), we can then make use of a set of stability conditions for a third order characteristic equation to determine if the remaining three eigenvalues are inside the unit circle. In fact, based on Farebrother (1973, p. 397, inequalities 3), we can then establish, as done in Lines and Westerhoff (2010, p. 256), the following set of necessary and sufficient for the cubic polynomial (A.11-a) to have roots that are strictly smaller than one in absolute value:

(A-13) $1-a-b-c>0$, $1+a-b+c>0$, $1+b+ac-c^2>0$, and $3-a+b+3c>0$.

Inserting (A-12) into (A-13), we obtain:

(A-14) $1-a-b-c = \frac{\alpha\beta}{\alpha\beta+x^*} > 0$,

(A-15) $1+a-b+c = \frac{\alpha\beta+4(1+\gamma)x^*}{\alpha\beta+x^*} > 0$,

(A-16) $1+b+ac-c^2 = \frac{\alpha\beta[\alpha\beta+x^*-2\gamma x^*]}{(\alpha\beta+x^*)^2} > 0$, if $\gamma < \frac{\alpha\beta+x^*}{2x^*} \equiv \gamma_c$,

(A-17) $3-a+b+3c = \frac{3\alpha\beta}{\alpha\beta+x^*} > 0$.

Therefore, if $\gamma < \frac{\alpha\beta+x^*}{2x^*} \equiv \gamma_c$, all the eigenvalues of (A-10) are strictly less than one in absolute value, so the equilibrium configuration given by $(x^*, \pi_n, u_n, \pi_n, \pi_n) \in \Theta$ of the system given by (19)-(21) and (17) is locally asymptotically stable. □