



Discrete beliefs space and equilibrium: a cautionary note

Michele Berardi¹ 

Published online: 29 June 2020
© The Author(s) 2020

Abstract

Bounded rationality requires assumptions about ways in which rationality is constrained and agents form their expectations. Evolutionary schemes have been used to model beliefs dynamics, with agents choosing endogenously among a limited number of beliefs heuristics according to their relative performance. This work shows that arbitrarily constraining the beliefs space to a finite (small) set of possibilities can generate artificial equilibria that can be stable under evolutionary dynamics. Only when “enough” heuristics are available are beliefs in equilibrium not artificially constrained. I discuss these findings in light of an alternative approach to modelling beliefs dynamics, namely, adaptive learning.

Keywords Expectations · Evolutionary dynamics · Learning · Equilibrium

JEL Classification C62 · D83 · D84 · E32

1 Introduction and related literature

The aim of this work is to shed some light on some results derived in macroeconomics using models with bounded rationality. In particular, I look at the practice of modelling expectations using evolutionary schemes to allow agents to choose among a limited set of alternative rules or heuristics. I will highlight possible perils of such practice and then draw a comparison with another common way of modelling agents' expectations under bounded rationality, namely, adaptive learning.

✉ Michele Berardi
michele.berardi@manchester.ac.uk

¹ The University of Manchester, Manchester, UK

Rational expectations (RE) represents a powerful way of closing and solving an economic model but they impose strong requirements on agents in terms of information acquisition and computational capabilities. In recent years, a growing number of studies have replaced RE with more realistic ways of modelling agents' expectations and their dynamics.

One popular method to model bounded rationality is adaptive learning, by which agents are treated as econometricians who repeatedly adjust the parameter values in their model economy based on previous forecast errors. Prominent examples are Sargent (1999), Evans et al. (2003, 2009) and, for an extensive treatise on methodology and applications, Evans and Honkapohja (2001).

A second class of methods used to model the way agents revise their expectations over time relies on evolutionary dynamics applied to a predetermined set of heuristics. Agents are allowed to select among the available expectation formation rules on the basis of their relative performance. While adaptive learning can be interpreted to represent learning at the individual level, evolutionary dynamics on expectation formation rules better capture the idea of social learning, where agents tend to imitate strategies that are more successful in the population. Under adaptive learning, in fact, agents adapt their own forecast strategy¹ based on previous performance, as measured against actual outcomes; under evolutionary models, instead, agents adopt strategies (again, expectations formation rules) that have been relatively more successful in the population, i.e., the performance of which has been assessed and evaluated positively against other strategies. Under evolutionary dynamics, thus, the performance of a strategy is not evaluated against actual outcomes but only against the performance of other available strategies. A strategy can thus perform quite badly in absolute terms, but still outperform all other available strategies and prevail in a market.

The two approaches, adaptive learning and evolutionary dynamics, have also been studied together in economic models: for example, Berardi (2015) combines evolutionary selection among heterogeneous classes of models with adaptive learning on the parameters of each model and finds that heterogeneous equilibria are possible but fragile; Branch and Evans (2006) propose the concept of misspecification equilibrium, where different underparameterized predictors are selected even in the limit under least-squares learning and dynamic predictor selection based on average profits; Guse (2010) considers a setting where agents can choose between a minimum state variable and a sunspot forecasting model and finds that, with an ad hoc cost to using the sunspot predictor, heterogeneity cannot be sustained under the combined evolutionary-adaptive learning dynamics.

The relation between learning and evolutionary dynamics has also been studied extensively in game theory. Fudenberg and Levine (1998, ch. 3) derive replicator dynamics from models of individual learning while Borgers and Sarin (1997) show that, in a game theoretical framework, a learning model based on the reinforcement of strategies with higher payoff converges, in a continuous time limit, to the replicator

¹Throughout this work, I will use the term strategy to mean a model/rule/heuristic for forming expectations. I will also call it “perceived law of motion” (or PLM), in line with the adaptive learning literature.

dynamics of evolutionary game theory. A key element to note is that, in these works, learning entails competition of strategies against each other, which differentiates crucially such learning mechanisms from adaptive learning.

More generally, the relation between learning and evolution has been debated in biology for a long time: Sznajder et al. (2012) recently reviewed some of the literature concerning the so called Baldwin effect Baldwin (1896), i.e., the hypothesis that learning, by improving fitness, accelerates evolution. Closer to our field of investigation, Marimon (1993) considers a game theoretical framework and compares key properties of adaptive learning (adaptation, experimentation and inertia) with their counterparts in evolutionary dynamics (reproduction, mutation and conservation).

An important difference between adaptive learning and evolutionary schemes as ways of modelling the evolution of expectations is that, in the first case, the researcher usually endows agents with a single forecasting model and then allows them to optimize the parameters in such a model through learning; in the second case, instead, the researcher selects a set of alternative forecasting models, and then allows agents to choose among this fixed set of rules based on their relative performance.² Anufriev and Hommes (2012), for example, consider a fixed set of rules, or heuristics, available to agents and investigate how the switching between these rules can explain results from experimental data. I will argue that the coarse discretization of the beliefs space under evolutionary schemes can lead to the creation of artificial equilibria and thus to misleading conclusions about possible outcomes in an economy. Such equilibria can be stable under evolutionary schemes, though they disappear when the beliefs space is sufficiently expanded.

From an empirical point of view, there seems to be quite a lot of dispersion in forecasts available through survey data (as seen, for example, from the Survey of Professional Forecasters). This means that households at any point in time have access, at least in principle, to a fairly large number of different predictors, represented by the different professional forecasters. Whether they actually take advantage of them, and whether any sort of convergence emerges among such predictors over time, are empirical questions that are beyond the scope of this paper. The aim of this paper is simply to point out the perils of artificially restricting, a priori, the set of available predictors to agents in a theoretical framework.

Another important question that is outside the scope of this paper pertains to the best way to model expectations under bounded rationality, whether through adaptive learning, evolutionary dynamics or possibly other methods. The answer, I speculate, would probably be case dependent, as different schemes could provide a better approximation of how agents actually behave under different circumstances. One, for example, could imagine that, in competitive environments, where the pressure for improvement comes from the need to outperform competitors, evolutionary dynamics could do a better job at approximating the way beliefs are formed and updated over time by agents; in problems of individual decision making, instead, adaptive learning might be a superior modelling strategy, as agents are not competing against each other but are trying to do the best they can for themselves in “absolute” terms.

²From an evolutionary perspective, only selection is allowed, but not mutation or crossover that could generate new alternatives.

Ultimately, this again seems to be an empirical question that can be settled only by empirical evidence. What I argue here is that it is important to understand the impact of using one or the other approach on the predictions that a model can deliver.

1.1 Related literature

The concept of evolutionary dynamics has been extensively applied in economics. In a seminal work, Brock and Hommes (1997) define the concept of adaptively rational equilibrium (ARE), in which agents adapt their beliefs over time by choosing from a finite set of different predictors according to a discrete logit model. They show how, in a simple cobweb model where agents can choose between rational (at a cost) and naive (free) expectations, highly irregular equilibrium prices can emerge. Building on this idea, Brock et al. (2005) introduce the concept of large type limit (LTL), describing the dynamic behavior of heterogeneous markets with many trader types. Applications of the LTL have been considered by Anufriev et al. (2013) and by Agliari et al. (2017): in both cases, from multiple steady states, the system moves to one unique steady state in the large type limit. This work will shed light on these results. In fact, in different settings, other studies - Aoki (2002) and Evstigneev et al. (2001) and Hens and Schenk-Hoppé (2005) - find that, as the number of strategies available to agents increases, heterogeneity in trading strategies tends to disappear. In related work, Evstigneev et al. (2006) find that only a rational market in which assets are evaluated by expected relative dividends is evolutionary stable. Similar to the LTL approach, Diks and Van der (2002) and Diks and van der (2003) propose a continuous beliefs system (CBS), where a continuum of strategies for forming beliefs are available to agents: while they are able to derive the limiting distributions for a number of performance measures, such distributions in fact degenerate into a point mass (and prices converge to the fundamental value) if agents are all allowed to choose the best possible predictor available. Again, the aim of this work is to understand such results more broadly.

An application of the ARE concept to a New Keynesian macroeconomic model of inflation and output has been proposed by Agliari et al. (2017), showing how this way of modelling expectations can alter the conclusions in terms of stability and uniqueness of equilibrium compared to rational expectations. Branch and McGough (2016) find cycles and chaotic dynamics in a monetary model à la Lagos and Wright (2005), with traders switching between (costly) rational and (costless) adaptive predictors according to a discrete logit model. De Grauwe (2011) generates endogenous waves of optimism and pessimism (“animal spirits”) in a macroeconomic model with beliefs switching behavior implemented through a discrete logit mechanism where agents can choose among a limited set of simple but biased rules (heuristics) to forecast future output and inflation. The author, discussing the results, recognizes that “[...] the menu of heuristics which is extremely small in this paper, will have to be broadened so that the selection of the “fittest” rules can occur using a wider pool of possible rules.” This is the aspect upon which this paper focuses.

In order to keep the analysis as general as possible, the evolutionary concept I will mainly rely upon is the evolutionary stable strategy (ESS), proposed by Maynard Smith and Price (1973, 1974). I will also make use, at times, of two additional

evolutionary schemes: the first is the concept of replicator dynamics proposed by Taylor and Jonker (1978) and used in economics, for example, by Sethi and Franke (1995) and more recently by Branch and McGough (2008); the second is a discrete logit model of beliefs selection, used, for example, in Brock and Hommes (1997).³

The choice of focusing on the notion of ESS has its benefits, but it is not without costs. While in fact this concept allows for a sharper characterization of some results, it does not allow us to discuss properly beliefs dynamics in a way that schemes such as replicator dynamics or the discrete logit model would. In particular, the notion of ESS relates to local stability properties of the model, and it has nothing to say about the global properties of the system under consideration. These various concepts are, nevertheless, closely related, and Weibull (1995) shows that every ESS is (locally) asymptotically stable under replicator dynamics, though the reverse does not generally hold. In other words, focusing on the notion of ESS will lead to results that are conservative: the equilibria found using ESS are also (locally) stable equilibria under replicator dynamics, though there might be other stable outcomes that emerge under replicator dynamics that do not correspond to ESS.

2 A stylized model

In order to investigate how the restriction of the beliefs space affects the set of possible equilibria, I consider a simple linear stochastic model with feedback from expectations to actual outcomes. The model, though very simple, admits multiple equilibria under rational expectation (RE).

The linear, stochastic model is represented by

$$y_t = B y_{t+1}^e + v_t, \quad (1)$$

where y_{t+1}^e denotes expectations at time t of y_{t+1} and v_t is an i.i.d. random variable with zero mean and variance σ_v^2 . Under RE, the model has a unique, fundamental (or minimal state variable - MSV) noisy steady state, with $y_{t+1}^e = 0$ and $y_t = v_t$, for $B \neq 1$.⁴

There is also a set of AR(1) RE equilibria of the form

$$y_t = \frac{1}{B} y_{t-1} - \frac{1}{B} v_{t-1}, \quad (2)$$

indexed by the initial condition y_0 . Such equilibria are stationary if $|B| > 1$.

The previous model, though most simple, captures nevertheless important features of economic models and can represent stylized economic systems.

The New Keynesian model presented in Agliari et al. (2017), for example, fits in this form, with y a vector of variables and B a suitable matrix. The Cobweb model used in Brock and Hommes (1997) has a different timing in terms of expectations,

³The discrete logit model can be derived from a random utility model (see, e.g., Manki and McFadden (1982) in which agents observe the performance of each rule with some noise; it can also be derived from a rational inattention problem, as shown by Matejka and McKay (2015).

⁴For $B = 1$, any arbitrary $y_{t+1}^e = k$ gives rise to a noisy steady state of the form $y_t = k + v_t$.

but admits the same fundamental (noisy) RE steady state (though no AR(1) REE), and considerations about the consequences of the discretization of beliefs derived in Eq. 1 can be carried over to that model.

3 Modelling beliefs

Stepping back from RE, beliefs dynamics under evolutionary forces can be modelled in different ways. Common choices in this respect are replicator dynamics and a discrete logit model, which I will discuss below. In game theory, a common way to analyze evolutionary stability is through the notion of an ESS. I will use this last concept in conducting my analysis and discuss the relationship between ESS and stability under both replicator dynamics and the discrete logit model. As all these concepts require a measure of the relative performance of different predictors, I first introduce the measure I will be using throughout the paper, the mean squared error.

3.1 A measure of relative performance

Agents are characterized by the strategy they use to form their expectations. I define a strategy for an agent as a forecasting model, also called a perceived law of motion (*PLM*), used to form their beliefs about actual outcomes. The performance of a strategy (*PLM*) is then measured by its mean squared error (*MSE*),⁵ defined as

$$MSE = E(y_t - y_t^e)^2, \quad (3)$$

where y_t^e are the expectations of y_t (formed at time $t - 1$) based on the *PLM* used.⁶

Given two competing strategies in use among agents, say PLM_1 and PLM_2 , their performance will depend in general on the fraction of agents using each strategy (as this determines the actual outcome y_t), determined by μ .⁷ Appropriate use of notation would thus require one to denote the performance of each strategy as $MSE(\mu)_1$ and $MSE(\mu)_2$, respectively, to show their dependence on μ . One could then define a measure of relative performance as

$$\Delta(\mu) = MSE(\mu)_2 - MSE(\mu)_1. \quad (4)$$

⁵The choice of *MSE* as measure of performance for the evolutionary schemes is consistent with the use of recursive least squares under adaptive learning, the objective of which is indeed to choose parameter estimates that minimize the *MSE* of the forecasting rule.

⁶Using the mean squared error as a driver for evolutionary dynamics might seem to impose strong informational requirements on agents. An alternative would be to assume that agents use an adaptive process to estimate the *MSE* over time, such as

$$MSE_t = (1 - g_t)MSE_{t-1} + g_t(y_t - y_t^e)^2$$

with g_t a small fixed or decreasing gain. For simplicity and analytical tractability, I will assume in this work instead that agents have full knowledge of the relevant *MSEs* at each point in time.

⁷More precisely, μ is the fraction of agents using PLM_1 , with $1 - \mu$ the fraction of agents using PLM_2 .

For simplicity, in the rest of the paper I will drop the notation that shows dependence on μ and simply write Δ , MSE_1 , MSE_2 . Note that, in all the settings considered in this paper (and generally in all well behaved economic models), $\Delta(\mu)$ is continuous, differentiable and monotonic in μ .

3.2 ESS and evolutionary schemes

While I will mainly adopt the concept of ESS in my analysis, the literature often models the evolution of beliefs according to evolutionary schemes such as replicator dynamics and discrete logit. There is a deep connection between these concepts. In particular, if such schemes are driven by the difference in forecasting performance, as described above, results of this paper derived using the ESS notion can be applied to settings using either of the two evolutionary schemes, under certain conditions. I elaborate on this point below, describing in briefs such schemes and showing their connection with the ESS concept.

3.2.1 Nash equilibrium and evolutionary stable strategy

Generally speaking, a Nash equilibrium is characterized by the fact that no agent has an incentive to change strategy. A strategy profile is defined by the strategies used by agents and the relative fraction of the population using each strategy. The following definition adapts this concept to the present setting, with two competing predictor strategies available to agents:

Definition 1 Consider a setting with two competing belief strategies available to agents, defined by PLM_1 and PLM_2 , used respectively by a proportion μ and $(1 - \mu)$ of agents. A Nash equilibrium is defined by the triple (PLM_1, PLM_2, μ) and one of the following three conditions: i) a non-negative fraction of agents uses each PLM available (i.e., $\mu \in [0, 1]$) and $\Delta = 0$; ii) all agents use PLM_1 ($\mu = 1$) and $\Delta > 0$; iii) all agents use PLM_2 ($\mu = 0$) and $\Delta < 0$.

I also define an evolutionary stable strategy (ESS) in this context, both for homogeneous and heterogeneous equilibria. In a homogeneous equilibrium, all agents use the same strategy (i.e., form their beliefs using the same PLM), and the condition for such strategy to be an ESS is the following:

Definition 2 The triple $(PLM_1, PLM_2, \mu = 1)$ is an ESS if $\exists \bar{\mu} \in (0, 1)$, s.t. $\Delta|_{\mu > \bar{\mu}} > 0$. Analogously, the triple $(PLM_1, PLM_2, \mu = 0)$ is an ESS if $\exists \bar{\mu} \in (0, 1)$, s.t. $\Delta|_{\mu < \bar{\mu}} < 0$.

In a heterogeneous (Nash) equilibrium, instead, a positive fraction of agents use each of the two strategies available. A strategy profile is then a ESS if the following condition holds:

Definition 3 The triple $(PLM_1, PLM_2, \tilde{\mu} \in (0, 1))$ is an ESS if $\exists \varepsilon$, s.t. for $\mu = \tilde{\mu} + \varepsilon$, $\Delta < 0$ and for $\mu = \tilde{\mu} - \varepsilon$, $\Delta > 0$, with $\Delta = 0$ for $\mu = \tilde{\mu}$.

These definitions follow from the definition of ESS in the seminal works Maynard Smith and Price (1973) and Maynard Smith (1974). It is well known that every ESS is a Nash equilibrium, but the converse is not necessarily true. Note also that Taylor and Jonker (1978) define an ESS as an evolutionary stable state (rather than strategy), which seems more appropriate for this setting. A state might be determined by a homogeneous group of agents using the same predictor (strategy), or by different groups using different predictors (strategies): the combination of strategies and the relative fraction of agents using each strategy defines a strategy profile.

Note that it is not possible in this setting to have a homogeneous Nash equilibrium that is not an ESS, since there is continuity of payoffs (Δ) in μ , and equilibrium requires Δ to be negative (positive) for $\mu = 1$ ($\mu = 0$).⁸ There could, though, be a heterogeneous Nash equilibrium that is not an ESS: a heterogeneous equilibrium, in fact, requires $\Delta = 0$, but nothing guarantees that the conditions set out in Definition (3) are satisfied.⁹

Finally, a note on terminology: in this paper I will often refer to an equilibrium as being an ESS (or not). This represents a slight misuse of terminology and must be understood as meaning that the strategy profile associated with such equilibrium is an ESS (or not).

3.2.2 Replicator dynamics

While replicator dynamics have been defined both in continuous and discrete time in the literature, the continuous time version is more widely used than its discrete time counterpart. In addition, and crucially for this work, it allows for a direct representation of the dynamics in terms of Δ .

Under continuous time replicator dynamics, the fraction μ of agents using PLM_1 (with the remaining fraction $(1 - \mu)$ using PLM_2) evolves according to the ordinary differential equation (ODE)

$$\dot{\mu} = \mu(1 - \mu)\Delta(\mu) \quad (5)$$

where $\Delta(\mu)$ is defined by Eq. 4.

Clearly equilibrium points of these replicator dynamics are points where $\mu = 0$, $\mu = 1$ or $\Delta = 0$. Note that points where all agents use the same predictor ($\mu = 0$ and $\mu = 1$) are fixed points of ODE (5) but need not be Nash equilibria, according to the definition given above (e.g., $\mu = 1$, $\Delta < 0$ is an equilibrium under replicator dynamics, but it is not Nash, as agents would have an incentive to switch model).

⁸I abstract here from corner solutions where homogeneous Nash equilibria ($\mu \in \{0, 1\}$) have $\Delta = 0$: in those cases, an equilibrium could be Nash but not an ESS. For example, if $\mu = 0 \Rightarrow \Delta = 0$ but $\mu > 0 \Rightarrow \Delta > 0$. This reminds one of a self-confirming equilibrium: if nobody uses PLM_1 , no agent has any incentive to adopt it, but once someone starts using it, everyone will follow.

⁹This is the case, for example, in an economy where there is negative feedback from expectations to outcomes: as more agents adopt a strategy, its performance decreases.

Local stability of equilibrium points under replicator dynamics is governed by $\frac{\delta \dot{\mu}}{\delta \mu}$, with

$$\frac{\delta \dot{\mu}}{\delta \mu} = (1 - 2\mu) \Delta(\mu) + \mu(1 - \mu) \frac{\delta \Delta(\mu)}{\delta \mu}.$$

In particular, at $\mu = \{1, 0\}$, local asymptotic stability is governed only by the sign of $\Delta(\mu)$: $\mu = 0$ is stable for $\Delta(0) < 0$, while $\mu = 1$ is stable for $\Delta(1) > 0$.

Remark 4 An homogeneous ESS is always locally asymptotically stable under replicator dynamics. The converse is not necessarily true.

This remark is a well known result in the literature. See, e.g., Weibull (1995), Proposition 3.10, p. 100. It implies that the set of ESSes that will be found in my analysis is a subset of the possible equilibria under replicator dynamics. Since the local asymptotic stability conditions for homogeneous equilibria ($\mu = \{1, 0\}$) under replicator dynamics are the same as the conditions for Nash equilibrium in definition (1), the set of Nash equilibria in the model corresponds to the set of locally asymptotically stable homogeneous equilibria under replicator dynamics. Coupling this result with the fact that the set of ESSes is a (possibly improper) subset of the Nash equilibria, it leads to the above remark that every homogeneous ESS is locally asymptotically stable under replicator dynamics, but the converse is not necessarily true.

3.2.3 Discrete choice model

Brock and Hommes (1997) use a discrete logit model for modelling prediction choice dynamics.¹⁰ Defining as μ_t the fraction of agents using PLM_1 in period t , the predictor choice in the population evolves over time according to

$$\mu_{t+1} = \frac{\exp\{-\beta MSE_{1,t}\}}{\exp\{-\beta MSE_{1,t}\} + \exp\{-\beta MSE_{2,t}\}} \tag{6}$$

or equivalently

$$\mu_{t+1} = \frac{1}{2} \left(\tanh \left[\frac{\beta}{2} \Delta(\mu_t) \right] + 1 \right), \tag{7}$$

where β is the ‘‘intensity of choice’’ parameter, a measure of how fast agents switch predictors, and $MSE_{1,t}$, $MSE_{2,t}$ and Δ_t are the unconditional mean square errors and their difference computed with μ_t .

Note that, for $\beta = \infty$, the model reduces to a deterministic choice model where all agents choose the predictor with the smaller MSE : only the sign of Δ matters. This implies that, for $\Delta(\mu_{t-1}) > 0$, $\mu_t = 1$ and, for $\Delta(\mu_{t-1}) < 0$, $\mu_t = 0$. Stability of homogeneous equilibria then requires $\Delta(0) < 0$ for $\mu_t = 0$ and $\Delta(1) > 0$ for $\mu_t = 1$. These conditions once again correspond to the Nash equilibrium conditions in definition (1). This means that every homogeneous ESS is a stable equilibrium under the discrete logit model with $\beta = \infty$, but the converse does not necessarily hold. I summarize these observations in the following remark:

¹⁰This is equivalent to a so called softmax action selection based on a Gibbs or Boltzman distribution.

Remark 5 Every homogeneous ESS is asymptotically stable under the discrete logit model. More generally, every homogeneous Nash equilibrium defined in Definition (1) - sub-points ii) and iii) - is asymptotically stable under the discrete logit model with $\beta = \infty$.

For finite β , instead, homogeneous equilibria do not emerge for finite values of $\Delta(\mu_t)$, as, in this case, $\mu_t = 1$ requires $\Delta(\mu_{t-1}) = \infty$ and $\mu_t = 0$ requires $\Delta(\mu_{t-1}) = -\infty$.

The relationship between the full set of Nash equilibria and stationary states under the two different evolutionary schemes can lead to interesting results. In particular, I will show below that if a setting does not allow for a homogeneous Nash equilibrium but it includes an heterogeneous Nash equilibrium that is also an ESS, it is possible to generate simple cycles between the two homogeneous (non-Nash) states using a discrete logit model for predictor selection. The same, though, would not be possible under replicator dynamics, as homogeneous states ($\mu = 1$ or $\mu = 0$) are always fixed points in this case (though not necessarily stable, as discussed previously).

4 Steady state equilibria

In this section, I analyze (noisy) steady state equilibria of model (1) when belief dynamics can evolve according to evolutionary schemes.¹¹

Such steady states will be characterized by $y_t = y + v_t$, for some y , and by beliefs dynamics being in equilibrium, in the sense that no agent wishes to change the predictor he is using. In addition to being a Nash equilibrium in this sense, I will also verify whether such equilibrium is evolutionary stable, using the concept of ESS.

In a steady state, a *PLM* for an agent is simply represented by a number, the predicted value for y_t in steady state.

4.1 Optimistic/pessimistic versus fundamental agents

Consider a setting where two competing *PLMs* are available to agents in order to form expectations, defined as follow:

$$PLM_1 : y_t = 0 \tag{8}$$

$$PLM_2 : y_t = a, a \in \mathbb{R} - \{0\}. \tag{9}$$

While *PLM*₁ is consistent with the fundamental equilibrium, one could interpret *PLM*₂ as representing agents being optimistic (for $a > 0$) or pessimistic (for $a < 0$)

¹¹It is instructive to note that steady state equilibria of Eq. 1 correspond to steady state equilibria of the alternative model

$$y_t = B y_t^e + v_t$$

where y_t^e denotes expectations at time $t - 1$ of y_t . This follows from the fact that, in a steady state, $y_t^e = y_{t+1}^e$. The framework is then also equivalent to a Cobweb model.

about the state of the economy.¹² If μ agents use PLM_1 and $1 - \mu$ use PLM_2 , then for any arbitrary a , aggregate expectations are given by $(1 - \mu) a$ and the actual law of motion (ALM) for the economy is represented by

$$y_t = B (1 - \mu) a + v_t.$$

Mean squared errors for the two groups are

$$\begin{aligned} MSE_1 &= E (B (1 - \mu) a + v_t - 0)^2 \\ MSE_2 &= E (B (1 - \mu) a + v_t - a)^2 \end{aligned}$$

and therefore

$$\Delta = a^2 (1 - 2B (1 - \mu)).$$

The choice of predictors by agents depends on the sign of Δ , as seen above, and clearly

$$\Delta > 0 \Leftrightarrow B (1 - \mu) < \frac{1}{2}. \tag{10}$$

It follows that, for any B , it is possible to find a $\bar{\mu} \in (0, 1)$ such that $\Delta|_{\mu > \bar{\mu}} > 0$ (since, for $\mu = 1$, the condition is satisfied, and Δ is continuous in μ): the strategy profile ($PLM_1, PLM_2, \mu = 1$) is thus an ESS. One can also see that, if $\mu = 0$, $\Delta < 0$ requires $B > 1/2$, in which case it is always possible to find a $\bar{\mu} \in (0, 1)$: $\Delta|_{\mu < \bar{\mu}} < 0$ and thus the strategy profile ($PLM_1, PLM_2, \mu = 0$) is also an ESS if $B > \frac{1}{2}$.

Proposition 6 *From condition (10), it follows that the strategy profile ($PLM_1, PLM_2, \mu = 1$), with $PLMs$ defined by Eqs. 8-9, is an ESS $\forall a \neq 0$. Moreover, the strategy profile ($PLM_1, PLM_2, \mu = 0$) is also an ESS $\forall a \neq 0$ if $B > \frac{1}{2}$.*

The previous proposition shows that in a setting where agents can choose between beliefs consistent with the fundamental equilibrium (PLM_1) and biased beliefs (PLM_2), a situation in which all agents use such biased beliefs can represent a Nash equilibrium (and an ESS). For $B > 1/2$, there are thus two (noisy) steady states, one where $y_t = 0 + v_t$ (fundamental) and one where $y_t = a + v_t$.

As noted before, because in these equilibria the MSE of the PLM used is smaller than the MSE of the PLM not used (that is $\Delta < 0$ for $\mu = 0$ and $\Delta > 0$ for $\mu = 1$), these ESSes are stable both under replicator dynamics and under the discrete logit model with $\beta = \infty$. Analyzing this system with either of the two mechanisms for selecting beliefs, one would thus conclude that additional steady states exist (and are evolutionary stable) besides the fundamental equilibrium.

4.2 Heterogeneous equilibria

The above results have shown that additional homogeneous equilibria can emerge from artificially limiting the set of belief strategies available to agents. Heterogeneous equilibria, where different groups of agents maintain different beliefs (i.e., use

¹²The case with both pessimistic and optimistic agents simultaneously competing is analyzed in the [Appendix](#) and confirms results found here.

different *PLMs*), can also emerge from the same practice of limiting the set of available alternatives to agents. I will show how this is possible by considering the case with only two available strategies (but the case is easily generalizable to setting with more alternatives available).

Consider the setting where a fraction μ of agents use PLM_1 defined as

$$PLM_1 : y_t = a, \quad a \in \mathbb{R} - \{0\} \tag{11}$$

and the remaining fraction $1 - \mu$ use PLM_2

$$PLM_2 : y_t = \alpha a, \quad \alpha \in \mathbb{R}. \tag{12}$$

This framework captures the situation where agents can choose among two steady state predictors, one of which could be consistent with the fundamental steady state.

The respective *MSEs* are

$$\begin{aligned} MSE_1 &= E (Ba (\mu + (1 - \mu) \alpha) + v_t - a)^2 \\ MSE_2 &= E (Ba (\mu + (1 - \mu) \alpha) + v_t - \alpha a)^2 \end{aligned}$$

and therefore

$$\Delta = a^2 (\alpha - 1) (1 + \alpha - 2B (\mu + (1 - \mu) \alpha)). \tag{13}$$

It is clear that $\Delta < 0$ if: $\alpha > 1$ and $B (\mu + (1 - \mu) \alpha) > \frac{1+\alpha}{2}$; or if $\alpha < 1$ and $B (\mu + (1 - \mu) \alpha) < \frac{1+\alpha}{2}$; that is, if αa is closer than a to the actual outcome induced by the heterogeneous expectations.

An heterogeneous equilibrium instead exists when $\Delta = 0$ for $\mu \in (0, 1)$: this requires

$$\mu = \tilde{\mu} \equiv \frac{\alpha + 1 - 2B\alpha}{2B(1 - \alpha)}. \tag{14}$$

For example, if $\alpha = -1$, then $\mu = \frac{1}{2}$ implies $\Delta = 0$: a heterogenous equilibrium exists where half the population is optimistic and half is pessimistic, equally distant from the fundamental value. Note that, in general, restrictions on α and B must be placed to ensure that $\tilde{\mu} \in (0, 1)$.

Would this heterogeneous equilibrium (i.e., the triple $PLM_1, PLM_2, \tilde{\mu}$) be an ESS? The condition for this to happen is that agents don't have any incentive to deviate from the strategy played in equilibrium, i.e.,

$$\begin{aligned} \Delta &< 0 \text{ if } \mu > \tilde{\mu} \\ \Delta &> 0 \text{ if } \mu < \tilde{\mu}. \end{aligned}$$

This means that Δ has to be decreasing in μ at $\tilde{\mu}$, i.e., $\frac{d\Delta}{d\mu} |_{\mu=\tilde{\mu}} < 0$. It is easy to compute

$$\frac{d\Delta}{d\mu} = 2a^2 (\alpha - 1)^2 B$$

and thus

$$\frac{d\Delta}{d\mu} < 0 \Leftrightarrow B < 0.$$

The heterogeneous equilibrium corresponds to an ESS for $B < 0$. I summarize results of this section in the following Proposition:

Proposition 7 *In a setting where agents can select among the two PLMs defined by Eqs. 11 and 12, a heterogeneous expectations equilibrium exists with $\mu = \tilde{\mu}$ determined by Eq. 14. The corresponding strategy profile ($PLM_1, PLM_2, \mu = \tilde{\mu}$) is an ESS if $B < 0$.*

With an arbitrary restriction of the strategy space, not only one can have artificial homogenous equilibria, but it is also possible to construct heterogeneous equilibria, and such equilibria can also be ESSes.

4.3 Rational cycles: an example

In the analysis considered in Section 4.2, for $B < 0$, a stable heterogeneous equilibrium is coupled with two homogeneous states that are not Nash equilibria (in the sense that $\Delta \neq 0$, though they are by definition fixed points for the replicator dynamics), and thus not ESSes.¹³ In particular, for $\mu = 0, \Delta > 0$ and for $\mu = 1, \Delta < 0$: this means that under a discrete logit model, “rational” agents (i.e., with $\beta = \infty$) would keep switching between the two strategies and the system would cycle between two states.¹⁴

4.4 Refining the beliefs space

Having seen that arbitrary restrictions of the beliefs space can lead to artificial equilibria under evolutionary dynamics, this section will show that, as the set of possible strategies is enlarged, the only ESS that survives is the one where all agents use a PLM consistent with the fundamental equilibrium. In other words, any arbitrary equilibrium (i.e., not fundamental) that is an ESS when a finite number of belief strategies is available ceases to be an ESS once available strategies are not artificially constrained.

To model an unrestricted beliefs space, it is convenient to use the setting discussed in Section 4.2, modifying it to allow an infinite number of strategies to be available. To this end, consider the same model presented before, with two PLMs available to agents, $PLM_1 : y_t = a$, for some $a \in \mathbb{R} - \{0\}$, and $PLM_2 : y_t = \alpha a$, but where now the second PLM is parameterized by α (instead of being characterized by a given and fixed α). This last PLM thus represents in effect a class of PLMs: restrictions on α amount to restrictions to the beliefs space of agents, while an unrestricted space of belief strategies can be characterized by an unconstrained $\alpha \in \mathbb{R} - \{0\}$.

For PLM_1 to be an ESS against $PLM_2, \forall \alpha \in \mathbb{R} - \{0\}$, it must be that it is possible to find a $\tilde{\mu} \in (0, 1)$ such that $\Delta|_{\mu > \tilde{\mu}} > 0$. But the difference in performance, given by Eq. 13, for $\mu = 1$ reduces to

$$\Delta(\alpha) = a^2(\alpha - 1)(1 + \alpha - 2B), \tag{15}$$

¹³Cycles could not emerge between different ESSes (unless different costs are assumed), as ESSes are absorbing states for evolutionary dynamics.

¹⁴This sort of cycles would not emerge under continuous time replicator dynamics, as $\mu \in \{0, 1\}$ are fixed points of the relevant ODE.




Fig. 1 Shape of $\Delta(\alpha)$

where the dependence of Δ on α is made explicit. Clearly, for any $B \neq 1$,¹⁵ it is always possible to find an $\alpha \in \mathbb{R} - \{0\}$ s.t. $\Delta(\alpha) < 0$. In particular, since $sign(\Delta)$ depends on $(\alpha - 1)(1 + \alpha - 2B)$, for a given B , this requires

$$1 < \alpha < 2B - 1 \quad \text{if } B > 1 \quad (16)$$

$$2B - 1 < \alpha < 1 \quad \text{if } B < 1. \quad (17)$$

This means that, when agents are free to choose any strategy (that is, any α), not only can an arbitrary PLM_1 not be an ESS, but it does not even correspond to a Nash equilibrium.

To understand these restrictions, it is useful to look at Eq. 15, which shows that $\Delta(\alpha)$ takes a U shape, with $\Delta < 0$ iff α lies within the two roots of the quadratic equation, equal to $(2B - 1)$ and 1 , and a minimum at $\alpha = B$. For example, with $a = 2$ and $B = \frac{1}{2}$, it can easily be seen that $\Delta = 0$ at the two roots $\alpha = 1$ and $\alpha = 2B - 1 = 0$, with Δ reaching a minimum at $a = B = 1/2$: see Fig. 1.

Note that the middle point between the two roots is B , the actual effect of expectations on y_t . Intuitively, agents using PLM_1 were expecting $y_t = a$, while the actual outcome turns out to be aB , that is, aB is the individual best response to strategy a . It follows that, for generic $B \neq 1$, any arbitrary strategy $a \neq 0$ cannot be a Nash equilibrium of this system, and thus cannot be an ESS if the state space of beliefs is not restricted. Moreover, since $a = 0$ is the only Nash equilibrium, it is by necessity an ESS - see Weibull (1995).

Proposition 8 *When the space of beliefs is unconstrained, the only Nash steady state equilibrium (also corresponding to an ESS) in model (1) is the fundamental one where all agents use PLM_1 .*

¹⁵The non-generic case $B = 1$ is trivial, as it implies that any a is an equilibrium and an ESS.

Equilibria derived under evolutionary dynamics arise from comparing the performance of different available strategies against each other. Since the actual outcome is endogenous to the forecasting strategy used by agents, restrictions in this space mean that equilibria can emerge where it is not worth it for agents to switch to an alternative strategy, even though the one they are using is far from being optimal, i.e., it is not a fixed point in the map from beliefs to outcomes. For example, the second part of Proposition 6, which states that PLM_2 can be an ESS, depends crucially on comparing the two strategies PLM_1 and PLM_2 against each other. But any arbitrary strategy is not a best response against itself, and so it does not represent a Nash equilibrium when the set of possible strategies is expanded.

4.5 Stepping stones and individual best response

The previous section showed that when the space of beliefs is unconstrained, the fundamental equilibrium is the only Nash (and ESS) equilibrium. Though a continuum of strategies is a sufficient condition for this result, it is by no means necessary. What is necessary, instead, is to have just enough “stepping stones” to ensure that any arbitrary strategy cannot survive. Any arbitrary belief (commonly held) will generate an actual outcome: as long as an alternative belief exists, which is closer to the actual outcome than the belief that generated it, such arbitrary belief cannot be an equilibrium.

One possible way to generate these “stepping stones” is through an individual best response function, as shown below. This method, though, does not generate the minimum number of strategies needed to ensure that the situation where all agents use the fundamental PLM is the only Nash equilibrium: one could exploit the fact that, for any given strategy, there is an interval of alternative strategies that outperform it, as seen above, and select the one that is closer to the fundamental as a stepping stone; that is, at each iteration, instead of the best response (i.e., $\alpha = B$), one could select a “minimum strategy” that does at least as well as the current one (i.e., $\alpha = 2B - 1$) and use that as a stepping stone towards the fundamental equilibrium.

A fruitful way to analyze the set of equilibria of the system when there are no a priori constraints on the set of possible belief strategies available to agents is to look at the long run equilibria of an adaptive game where agents can replace strategies over time, moving from one to the next through an individual best response mechanism. In this perspective, a strategy profile represents a Nash equilibrium if it is a best response to itself.¹⁶

I thus define a best response function $f : a \rightarrow Ba$, which generates successive strategies for agents and determines the beliefs dynamics of the system.

Consider, for example the framework from the previous section: when all agents use a strategy defined by $PLM_1 : y_t = a_1$, the actual outcome is given by $y_t = Ba_1$. An individual agent, observing such outcome, could then use the new strategy $PLM_2 = a_2 = Ba_1$ as predictor for the steady state. This new strategy is derived as the best response to PLM_1 and represents the case where the new

¹⁶For an analysis of best response dynamics and Nash equilibria, see Gilboa and Matsui (1991).

belief is halfway in the range of possible values that would lead to a lower MSE . In the next round, then, the new best response strategy would be $Ba_2 = B^2a_1$. And so on.

This way of introducing new strategies through the best response function f reminds one of the naive expectations assumption common before the RE revolution, by which agents use the current observed value as a predictor for the future. It is clearly not optimal in general, though it can be in special cases. More importantly here, it provides a way to generate new strategies for evolutionary selection.

In the present case, such “adaptive” behavior would indeed be optimal for an atomistic agent: when all agents predict the outcome a_1 , the strategy $a_2 = Ba_1$ is indeed the best individual predictor, since it has minimum MSE . In other words, $a_2 = Ba_1$ is a best response strategy, in game theoretic terms, for an atomistic agent when everybody else is using a_1 . If all agents switch from one period to the other, only two long run outcomes are then possible: either $y_t \rightarrow 0$ (for $|B| < 1$) or $y_t \rightarrow \pm\infty$ (for $|B| > 1$), either monotonically or oscillating depending on the sign of B . The following proposition summarizes this result.

Proposition 9 *The only long run stationary (Nash) equilibrium in an adaptive game where new strategies are generated (synchronously) through an “individual best response” mechanism is the fundamental equilibrium with $y^e = 0$.*

5 Equilibria in autoregressive form

So far I have considered only steady state equilibria of Eq. 1, which have the special feature that a PLM (or forecasting strategy) is simply a constant. It is well known, though, that, under RE, in addition to the fundamental, steady state equilibrium, model (1) admits also equilibria in AR(1) form.¹⁷ It is thus interesting to analyze beliefs dynamics under evolutionary schemes that rely on AR(1) forecasting models.

Given the simple form of model (1), which does not include any lagged variable, in order for an AR(1) equilibrium to emerge it is necessary that (at least some) agents use an AR(1) model to form expectations.

5.1 Discrete beliefs space

I will start by showing how artificial AR(1) equilibria can also arise under evolutionary schemes when the beliefs space is constrained. In order to discipline the analysis, I will focus on two cases separately: i) the case where a steady state PLM and an AR(1) PLM are available to agents to form their expectations; ii) the case where two competing AR(1) $PLMs$ are available.

¹⁷Such equilibria require restrictions on B in order to be stationary.

5.1.1 AR(1) versus steady state models

Consider first the case where two *PLMs* are available to agents, one consistent with the economy being in a steady state and the other consistent with the economy being in an AR(1) equilibrium. In particular, the first model, denoted PLM_1 , is represented by

$$y_t = a, \quad a \in \mathbb{R} \tag{18}$$

while the second, PLM_2 , is consistent with the AR(1) equilibrium and represented by

$$y_t = by_{t-1}, \quad b \in \mathbb{R}. \tag{19}$$

Parameters a and b will be specified later on for different scenarios.

With a fraction μ of agents using the first model and the remaining fraction $(1 - \mu)$ using the second model, the ensuing aggregate expectations are

$$E_t y_{t+1} = \mu a + (1 - \mu) b^2 y_{t-1}$$

and the ALM for the system is given by

$$y_t = B\mu a + B(1 - \mu) b^2 y_{t-1} + v_t. \tag{20}$$

MSEs are given, respectively, by

$$\begin{aligned} MSE_1 &= E(y_t - a)^2 \\ MSE_2 &= E(y_t - by_{t-1})^2, \end{aligned}$$

and inserting the ALM into these equations leads to:

$$MSE_1 = E\left((B\mu - 1)a + B(1 - \mu)b^2 y_{t-1} + v_t\right)^2 \tag{21}$$

$$MSE_2 = E\left(B\mu a + (B(1 - \mu)b^2 - b)y_{t-1} + v_t\right)^2. \tag{22}$$

It follows that

$$\Delta = \frac{b^2(1 - 2bB(1 - \mu))(B^2\mu^2 a^2 + \sigma_v^2)}{1 - (B(1 - \mu)b^2)^2} - a^2(1 - 2B\mu) + \frac{2a^2 b B^2 \mu (b(1 - \mu) - \mu)}{1 - B(1 - \mu)b^2}. \tag{23}$$

For the general case $a, b \neq 0$, the analysis is complicated by the interactions between the unconditional mean of the model (which would be different from zero for $a \neq 0$) and the AR(1) component. Some special cases are of particular interest.

First, for $a \neq 0$ and $b = 0$, $\Delta > 0 \Leftrightarrow B\mu > \frac{1}{2}$: this condition ensures that the expectations derived with PLM_1 (equal to a) are closer (in expected value) to the actual realization ($B\mu a$) than those derived under PLM_2 (equal to 0). Thus $B > 1/2$ ensures that the strategy profile ($PLM_1, PLM_2, \mu = 1$) is an ESS¹⁸. For $\mu = 0$, instead, $\Delta < 0$, and thus the strategy profile ($PLM_1, PLM_2, \mu = 0$) is always an ESS (again, as already seen above).

More interesting is the case where $a = 0$ and $b \neq 0$, for which it is possible to derive precise results that shed light on the interactions between these components.

¹⁸This case, in fact, corresponds to one already seen above, since, from PLM_2 with $b = 0$, the expected value for y_t is always zero.

In particular, I will consider two sub-cases: i) a MSV *PLM* against an AR(1) REE *PLM*; ii) a MSV *PLM* against a “naive expectations” *PLM*. The intuition behind the results is always the same: the strategy that performs better is the one that delivers expectations that are closer, in MSE, to the actual realization, which is determined by the expectations of each group, weighted by their relative size, and mapped into actuals through the structural parameter *B*.

MSV versus AR(1) REE models Suppose, first, that the two groups of agents believe themselves to be in a RE equilibrium, but not in the same one: agents in the first group, of size μ and using *PLM*₁, believe themselves to be in the fundamental noisy steady state (thus $a = 0$) while agents in the other group, of size $1 - \mu$ and using *PLM*₂, believe themselves to be in the AR(1) RE equilibrium (thus $b = \frac{1}{B}$). The ensuing aggregate expectations are thus given by

$$E_t y_{t+1} = \frac{1 - \mu}{B^2} y_{t-1}$$

and the ALM for the economy is

$$y_t = \frac{1 - \mu}{B} y_{t-1} + v_t. \tag{24}$$

It is then straightforward to derive the MSEs for the two groups

$$MSE_1 = E \left(\frac{1 - \mu}{B} y_{t-1} + v_t \right)^2$$

$$MSE_2 = E \left(-\frac{\mu}{B} y_{t-1} + v_t \right)^2.$$

Using the fact that, from Eq. 24 and assuming stationarity of y_t ,¹⁹

$$E y_t = 0$$

$$E y_t^2 = \frac{B^2}{B^2 - (1 - \mu)^2} \sigma_v^2$$

it follows that

$$MSE_1 = \frac{B^2 \sigma_v^2}{B^2 - (1 - \mu)^2}$$

$$MSE_2 = \frac{(B^2 + 2\mu - 1) \sigma_v^2}{B^2 - (1 - \mu)^2}.$$

The stationarity condition implies that the denominator of these MSEs is positive. Clearly then, $MSE_2 > MSE_1 \Leftrightarrow \mu > \frac{1}{2}$: the model used by the majority of agents delivers better forecasts and, in the long run, prevails under evolutionary dynamics. This means that the two strategy profiles (*PLM*₁, *PLM*₂, $\mu = 1$) and (*PLM*₁, *PLM*₂, $\mu = 0$) represent ESSes: the first gives rise to the MSV REE, the second

¹⁹Note, from ALM (24), that stationarity requires $|1 - \mu| < |B|$. As $\mu \rightarrow 1$, this condition is satisfied $\forall B$. As $\mu \rightarrow 0$, instead, the requirement is the usual one for the existence of stationary AR(1) solutions seen above, i.e., $|B| > 1$.

to the AR(1) REE. This result also implies that, under replicator dynamics, both $\mu = 1$ and $\mu = 0$ are locally asymptotically stable equilibria, with basins of attraction separated by $\mu = \frac{1}{2}$, which represents an unstable heterogeneous equilibrium.

Proposition 10 *Both strategy profiles $(PLM_1, PLM_2, \mu = 1)$ and $(PLM_1, PLM_2, \mu = 0)$, with PLM_1 defined by Eq. 18 with $a = 0$ and PLM_2 defined by Eq. 19 with $b = 1/B$ represent ESSes. The first gives rise to the MSV REE, the second to the AR(1) REE.*

MSV versus naive expectations Another interesting case arises when $a = 0$ and $b = 1$: one group of agents thus uses a PLM consistent with the MSV fundamental equilibrium while the other uses a simple “naive expectations” model. It is straightforward to show, from condition (23), that: a) the strategy profile $(PLM_1, PLM_2, \mu = 1)$ is always an ESS; b) the strategy profile $(PLM_1, PLM_2, \mu = 0)$, where all agents use naive expectations, is an ESS if $B < -1$ or $\frac{1}{2} < B < 1$ (but, in the second case, the ensuing equilibrium is not stationary). Naive expectations can thus represent an ESS strategy against fundamental expectations in this setting for $B < -1$.

Proposition 11 *Consider the strategy profiles $(PLM_1, PLM_2, \mu = 1)$ and $(PLM_1, PLM_2, \mu = 0)$, with PLM_1 defined by Eq. 18 with $a = 0$ and PLM_2 defined by Eq. 19 with $b = 1$. The first represents an ESS for any value of B , while the second is an ESS if $B < -1$ or $\frac{1}{2} < B < 1$.*

5.1.2 Competing AR(1) models

The second case I consider is one where different group of agents use different AR(1) models as PLMs. Consider the two competing AR(1) PLM s, used, respectively, by a fraction μ and of $1 - \mu$ of agents:

$$PLM_1 : y_t = by_{t-1}, \quad b \in \mathbb{R} \tag{25}$$

$$PLM_2 : y_t = dy_{t-1}, \quad d \neq b \in \mathbb{R}. \tag{26}$$

Deriving the ensuing ALM for the system and computing the MSEs for each PLM , it is straightforward to find the following condition on $\Delta \equiv MSE_2 - MSE_1$:

$$\Delta > 0 \Leftrightarrow (d - b) \left(d + b - 2B \left(\mu b^2 + (1 - \mu) d^2 \right) \right) > 0 \tag{27}$$

or

$$d > b \Rightarrow \Delta > 0 \Leftrightarrow \frac{d + b}{2} > B \left(\mu b^2 + (1 - \mu) d^2 \right) \tag{28}$$

$$d < b \Rightarrow \Delta > 0 \Leftrightarrow \frac{d + b}{2} < B \left(\mu b^2 + (1 - \mu) d^2 \right). \tag{29}$$

Note that the r.h.s of these conditions represents the impact of y_{t-1} on y_t . These conditions, therefore, have a straightforward interpretation: the PLM that performs better is the one that predicts an impact of past y_{t-1} on current y_t closer to its actual value.

If one now sets $d = 1/B$, PLM_2 corresponds to the AR(1) REE strategy. One can thus check whether such strategy can dominate arbitrary AR(1) strategies (and vice-versa).

Setting $d = 1/B$ and $\mu = 0$ in condition (27) leads to

$$\Delta > 0 \Leftrightarrow -(1 - Bb)^2 > 0$$

and such condition is clearly never satisfied. Since $\Delta < 0$ for any possible value of b , the strategy $d = 1/B$ always outperforms any possible alternative strategy b when all agents use it (as it represents the best response to itself). Moreover, by continuity, is as also an ESS. Conversely, setting $d = 1/B$ and $\mu = 1$ in condition (27) leads to

$$\Delta > 0 \Leftrightarrow (1 - Bb) \left(1 + Bb - 2(Bb)^2 \right) > 0. \quad (30)$$

It follows that, $\forall B, \exists b$ such that condition (30) is satisfied and $\Delta > 0$:²⁰ the strategy defined by such b thus outperforms the strategy defined by $d = 1/B$ when all agents use it and it is an ESS.

More generally, it is possible to verify whether two arbitrary AR(1) models, defined by PLM_1 with arbitrary b and PLM_2 with arbitrary d , can represent ESSes. Clearly condition (27), computed with $\mu = 1$ and $\mu = 0$, respectively, shows that they can. Assuming, without loss of generality, that $d > b$, $\Delta > 0$ with $\mu = 1$ requires $d + b - 2Bb^2 > 0$ and $\Delta < 0$ with $\mu = 0$ requires $d + b - 2Bd^2 < 0$. Putting together the two conditions, it is thus possible to find, $\forall B$, a pair (b, d) that satisfies $\frac{d+b}{2b^2} > B$ & $\frac{d+b}{2d^2} < B$.

Proposition 12 *Consider the strategy profiles $(PLM_1, PLM_2, \mu = 1)$ and $(PLM_1, PLM_2, \mu = 0)$, with PLM_1 defined by Eq. 25 with arbitrary b and PLM_2 defined by Eq. 26 with arbitrary d , with $d > b$. Both strategy profiles can represent ESSes. The first requires $\frac{d+b}{2b^2} > B$ and the second $\frac{d+b}{2d^2} < B$. Selecting b and d appropriately, one can define a system where both strategy profiles are ESSes.*

With a limited number of strategies (two in this case), it is always possible to choose strategy profiles where arbitrary beliefs dominate the alternatives available and thus construct artificial equilibria (in the sense that they would be equilibria under beliefs dynamics driven by replicator dynamics or discrete logit model). But these equilibria are artificial, due to the arbitrary restriction of the beliefs space, as I show in the next section.

5.2 Refining the beliefs space

In order to model an unrestricted beliefs space, I consider as before a setting where one PLM is fixed while the other is parameterized to represent a continuum of $PLMs$, this time in AR(1) form. In particular, the two $PLMs$ are given by

²⁰This requires $-\frac{1}{2} < Bb < 1$. For example, one could simply choose $b = \frac{1}{2B}$.

$$PLM_1 : y_t = by_{t-1}, \quad b \in \mathbb{R} \tag{31}$$

$$PLM_2 : y_t = \alpha by_{t-1}, \tag{32}$$

where the second *PLM* is now parameterized by $\alpha \in \mathbb{R}$.

If μ agents use *PLM*₁ and $1 - \mu$ use *PLM*₂ to form expectations, the ensuing ALM is then

$$y_t = Bb^2 (\mu + (1 - \mu)\alpha^2) y_{t-1} + v_t, \tag{33}$$

and the two *MSEs* are

$$MSE_1 = E \left(\left(Bb^2 (\mu + (1 - \mu)\alpha^2) - b \right) y_{t-1} + v_t \right)^2$$

$$MSE_2 = E \left(\left(Bb^2 (\mu + (1 - \mu)\alpha^2) - \alpha b \right) y_{t-1} + v_t \right)^2.$$

Note that $Ey_t = 0$ and $Ey_t^2 = \frac{\sigma_v^2}{1 - (b^2 B (\mu + (1 - \mu)\alpha^2))^2}$, under the condition that Eq. 33 is stationary (which requires $|Bb^2 (\mu + (1 - \mu)\alpha^2)| < 1$). It is then possible to compute

$$\Delta(\alpha) = \sigma_v^2 \frac{b^2 (\alpha^2 - 1) + 2b(1 - \alpha)(b^2 B (\mu + (1 - \mu)\alpha^2))}{1 - (b^2 B (\mu + (1 - \mu)\alpha^2))^2}, \tag{34}$$

where again it is made clear that Δ is now a function of the free parameter α . The next step now is to check whether an arbitrary *PLM*₁ can be a Nash equilibrium against any other *PLM*₂, which is a necessary condition for it to be an ESS. To do so, I compute (34) with $\mu = 1$ and obtain

$$\Delta(\alpha)_{\mu=1} = \sigma_v^2 b^2 \frac{(\alpha^2 - 1) + 2Bb(1 - \alpha)}{1 - (b^2 B)^2}. \tag{35}$$

Note that stationarity of Eq. 33 with $\mu = 1$ requires $|Bb^2| < 1$ and thus the denominator is always positive. In order for the strategy profile (*PLM*₁, *PLM*₂, $\mu = 1$) to be a Nash equilibrium, it must be that $\Delta(\alpha) > 0$ everywhere. But one can see that it is always possible to find a value for α such that $\Delta(\alpha) < 0$. In particular, given B and b , this requires to choose α as follows:

$$\alpha > 1 \text{ if } Bb > 1$$

$$\alpha < 1 \text{ if } Bb < 1.$$

This shows that, unless $b = 1/B$, which has been shown before to represent a Nash equilibrium,²¹ it is always possible to find a value for α such that *PLM*₂ outperforms

²¹With $b = 1/B$, from Eq. 35 $\Delta = \frac{(\alpha-1)^2}{1-B^{-2}}$ and for $|B| > 1$ (required for stationarity), $\Delta > 0$ (in fact, $MSE_1 = \sigma_v^2$), $\forall \alpha \neq 1$: the strategy $b = 1/B$ is thus a Nash equilibrium. Moreover, because of continuity of Δ in μ , it is also an ESS.

PLM_1 for $\mu = 1$: any PLM defined by an arbitrary $AR(1)$ parameter b cannot, therefore, represent a Nash equilibrium (and thus cannot be an ESS).

Proposition 13 *When agents can choose among a set of unrestricted $AR(1)$ models to form their forecasts, the belief strategy defined by $b = 1/B$ is the only Nash equilibrium and ESS that gives rise to a stationary homogeneous $AR(1)$ equilibrium for $|B| > 1$.*

This result shows that any other $AR(1)$ strategy different from $b = 1/B$ can be outperformed and does not, therefore, represent an equilibrium in this setting. Such arbitrary strategies can only be found to be equilibria when the set of available predictors is arbitrarily restricted.

5.3 Stepping stones and individual best response

Again, it is insightful to look at the problem in terms of an adaptive game, where agents replace their strategies with better ones over time through an “individual best response” function.

Starting from the case where every agent uses the arbitrary $PLM_1 : y_t = by_{t-1}$, the individual best response to such strategy would be to use the new $PLM_2 : y_t = Bb^2y_{t-1}$. Repeating this process over time gives rise to the map

$$b \rightarrow Bb^2,$$

which defines a difference equation in b

$$b_{t+1} = Bb_t^2.$$

This non-linear difference equation admits two fixed points: $b = 0$ and $b = 1/B$. Such points correspond to the two homogeneous RE equilibria and their stability requires the derivative

$$\frac{\partial b_{t+1}}{\partial b_t} = 2Bb_t$$

to be within the unit circle. It is then immediate to see that $b = 0$ is locally stable while $b = 1/B$ is not under this “individual best response” dynamics.²² Thus, when agents are unconstrained in terms of their forecasting strategy, the only equilibrium that can emerge from the above adaptive process is the fundamental MSV equilibrium. Again, as before, it is not necessary to have a continuum of strategies available in the form of the support for b : only “enough” strategies to ensure that any arbitrary value of b that is not Nash can be selected away by a better strategy. Individual best

²²Remember that it was found instead that $b = 1/B$ can be a Nash equilibrium and an ESS, when played against the fundamental equilibrium. Again, this shows the central message of the paper: restricting the strategy space can generate artificial equilibria, which do not survive if the strategy space is sufficiently enlarged.

reply is a natural way to deploy such stepping stones strategies, though it is not the most parsimonious one.

6 Some discussion and a comparison

The central message of this paper is that an arbitrary limitation of beliefs under evolutionary schemes can give rise to artificial equilibria. Through the notion of “stepping stones” I have shown that, if enough strategies are available, only strategies that are Nash in an unrestricted beliefs space can emerge as equilibrium outcomes.

Researchers wishing to use evolutionary dynamics on a pre-defined set of heuristics, thus, should make sure that “enough” heuristics are made available to agents. It might well be that different strategies entail different costs (monetary, informational, or of other sort), but the choice should not be restricted a priori.

This paper helps explain some results in the existing literature. For example, Brock and Hommes (1997) show that, if the intensity of choice parameter $\beta \rightarrow \infty$, the equilibrium in their cobweb model with adaptively rational agents would be one where all agents use the naive predictor, as in equilibrium such a predictor is equivalent to the rational one, and it is cost free. Without costs on the rational predictor, all agents would instead use the RE predictor. In either case, the price (and the expected price) is equal to the fundamental value, which is the only Nash equilibrium and ESS. Note that, in that setting, while it would seem that only two predictors are available, the naive and the rational one, effectively the rational predictor is not defined as a fixed heuristic but instead as a belief equal to the actual realized price, which changes (among other things) with the fraction of agents using the naive predictor: it thus corresponds to an infinite set of heuristics in the terminology of this paper.

Similarly, Diks and van der (2003) find that, in a model similar to Eq. 1 with AR(1) strategies, a CBS leads to an equilibrium distribution that, as $\beta \rightarrow \infty$, degenerates into a mass point on the fundamental price. This is indeed the only stationary Nash and ESS for $B < 1$.

With a discrete logit model, when $\beta \rightarrow \infty$, heterogeneity by definition disappears. The question is, what strategy will prevail? This work has provided some insight into the answer to this question. It is important to note that results obtained are not due to specific features of the framework I’ve adopted, but instead come from the way evolutionary schemes select strategies, based on relative, rather than absolute, performance. A strategy can thus prevail even if it entails a large forecast error, as long as such error is smaller than the ones entailed by the available competing strategies. With a small number of competing strategies, such error can in fact be quite large. Note that, if all possible strategies were available, the only equilibrium points would be ones where the error is at its minimum, i.e., the error entailed by rational expectations. I am going to expand on this point in Section 6.1.

One important assumption held throughout this paper is that all agents switch beliefs at the same time (synchronous updating): it would be interesting to consider these issues in the case where agents are allowed to switch at different times. I leave this issue for future work.

6.1 A comparison with adaptive learning

In order to gain some more insights into the above results, it is useful to draw a comparison between evolutionary dynamics and adaptive learning as a means to model the evolution of beliefs.

Under adaptive learning, model (1) has a unique E-stable (that is, stable under adaptive learning)²³ equilibrium, the MSV, for $|B| < 1$ (E-stability requires $B < 1$; the other part of the constraint ensures a unique RE equilibrium). The AR(1) equilibrium instead is never E-stable, which means that it cannot be learned by agents acting as econometricians and recurrently estimating an AR(1) PLM.²⁴ This means that, if agents are given a steady state PLM, they can converge to the fundamental REE, while if they are given an AR(1) PLM, they cannot learn the AR(1) REE. In contrast, we have seen that, under evolutionary schemes with a limited number of beliefs heuristics available, non-fundamental equilibria can arise. With unrestricted beliefs, instead, the only ESS are the MSV equilibrium and the AR(1) REE (this last one requiring $|B| > 1$), though only the MSV equilibrium is obtainable through “best reply” dynamics.

Under adaptive learning, as long as beliefs are not consistent with actual outcomes (i.e., $a \neq Ba$ for the steady state case, or $b \neq Bb^2$ for the AR(1) case), beliefs are revised (i.e., new strategies, in the form of revised values for a or b , are generated through a form of directed search). Under evolutionary dynamics, instead, agents change beliefs (heuristics) only if better ones are available, i.e., they can only select a 's and b 's that are made available to them by the modeller.

Evolutionary dynamics and adaptive learning thus tend to deliver different results in terms of beliefs dynamics and equilibrium outcomes: while adaptive learning can only select (Nash) REE,²⁵ evolutionary dynamics allow for a richer set of outcomes to emerge when the set of available strategies is restricted. An important element in all these results is the self-referentiality of models such as Eq. 1, where actual outcomes depend on agents' beliefs. This means that the measure of fitness for beliefs, their forecasting accuracy, is endogenous to the choice of agents, and thus multiple equilibria can emerge.²⁶ While such self-referentiality is a feature of the model itself, irrespective of how expectations are modelled, its consequences are more dramatic under evolutionary dynamics because such dynamics are based on the relative performance of competing forecasting models (PLMs), rather than on the absolute performance of a given forecasting model against actual outcomes as is the case under adaptive learning. In a sense, beliefs under evolutionary dynamics are less constrained by outcomes, as they only need to prevail in relative terms, and this enhances the possibility of non-fundamental outcomes. This is particularly the case when a

²³For a discussion of E-stability, and in general for an exhaustive treatment of adaptive learning in macroeconomics, see Evans and Honkapohja (2001).

²⁴These are well known results in the adaptive learning literature. See, e.g., Evans and Honkapohja (2001).

²⁵This is true under the assumption that agents are given the correct PLM; otherwise, convergence could be to a different equilibrium, such as a restricted perceptions equilibrium. See Evans and Honkapohja (2001).

²⁶For example, with $B = 0$ it is easy to see from Proposition 14 that the arbitrary PLM_2 cannot be an ESS. The MSV model would always dominate under evolutionary dynamics.

small number of beliefs heuristics compete, as the “distance” in performance between different predictors can be quite large and allows for each set of beliefs to be an equilibrium (in the restricted sense of being Nash and an ESS against each other).

Evolutionary algorithms require two elements in order to be able to drive beliefs dynamics towards an optimum: a selection criteria (such as the forecasting performance in terms of MSE), and a mechanism to generate new strategies.²⁷ The two elements combine to produce dynamics that move from one strategy to the next, improving performance along the way. But if the mechanism that generates new strategies is muted, no evolution towards an optimum can take place. Under adaptive learning, instead, the two elements are joint: the dynamics from one strategy (represented by the current estimates for parameters in the PLM) to the next (updated estimates) are determined by the forecast error through some form of parameter estimation (e.g., recursive least square), that is, “mutation” in the strategy spaces takes place in the direction of the forecast error, and the magnitude is determined by the gain coefficient and the second moments of the regressors. This ensures that only fixed points from beliefs to outcomes (entailing a zero expected forecast error), i.e., Nash equilibria, can be equilibria of these learning schemes.

7 Conclusions

In this work I have shown how modelling beliefs using evolutionary schemes based on a limited number of predictors can lead to misleading conclusions about possible equilibrium outcomes in a model where current endogenous variables are affected by expectations.

After analyzing this issue in a simple, univariate forward-looking model, I have put results in context by drawing a comparison with adaptive learning and discussed the differences between the two approaches. The key difference lies in the fact that adaptive learning algorithms are driven by the difference between predictions and outcomes, while evolutionary schemes are driven by the difference in performance between different available predictors. Constraining the set of available strategies/predictors can thus allow for otherwise inferior beliefs to survive. The outcome is analogous to using genetic algorithms for optimization without allowing new population elements to be generated at each iteration: inferior solutions can prevail.

It must be noted, to conclude, that the aim of this work is not to argue against the use of evolutionary schemes to model expectations dynamics, but simply to warn of the possible consequences of such practice. Whether agents modify their forecasting models in the direction of actual outcomes, as implied by adaptive learning methods, or they select among a limited set of fixed heuristics, as assumed under evolutionary schemes, is ultimately an empirical question that remains to be settled and that is bound to be context specific. The aim of this paper is simply to clarify how different assumptions in this regard lead to different conclusions about equilibrium outcomes in self-referential economic models.

²⁷In genetic algorithms, this second step is usually achieved through the operations of crossover and mutation.

Compliance with Ethical Standards

Conflict of interests the author declares that he has no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix: Optimistic, pessimistic and fundamental agents

While the analysis in Section 4.1, and throughout the paper, is restricted to only two PLM_s available to agents, it can be extended to allow for more possibilities. Consider, for example, the case where three PLM_s are available to agents, defined as follows:

$$PLM_1 : y_t = 0 \quad (36)$$

$$PLM_2 : y_t = a, \quad a \in \mathbb{R} - \{0\} \quad (37)$$

$$PLM_3 : y_t = \alpha a, \quad \alpha \in \mathbb{R} - \{0\}. \quad (38)$$

A strategy profile is now defined by the 5-tuple $(PLM_1, PLM_2, PLM_3, \mu_1, \mu_2)$, where μ_1 and μ_2 represent the relative fraction of agents using, respectively, PLM_1 and PLM_2 (with $\mu_3 = 1 - \mu_1 - \mu_2$ representing the remaining fraction of agents using PLM_3). One can interpret PLM_1 as being a “fundamental” PLM , while PLM_2 and PLM_3 define different degrees of optimism/pessimism competing with each other.

Clearly, the strategy profile where all agents use the fundamental PLM_1 is always an ESS. It can also be shown that the strategy profile where all agents use the arbitrary PLM_2 is an ESS if

$$\alpha < 0 \ \& \ B > \frac{1}{2} \quad (39)$$

$$0 < \alpha < 1 \ \& \ B > \frac{1}{2} + \frac{\alpha}{2} \quad (40)$$

$$\alpha > 1 \ \& \ \frac{1}{2} < B < \frac{1}{2} + \frac{\alpha}{2}. \quad (41)$$

The conditions for the strategy profile where all agents use the arbitrary PLM_3 to be an ESS are then symmetric, with

$$\alpha < 0 \ \& \ B > \frac{1}{2} \quad (42)$$

$$0 < \alpha < 1 \ \& \ \frac{1}{2} < B < \frac{1}{2} + \frac{1}{2\alpha} \quad (43)$$

$$\alpha > 1 \ \& \ B > \frac{1}{2} + \frac{1}{2\alpha}. \quad (44)$$

A brief comment on these inequalities. Condition $B > \frac{1}{2}$ is required in order for the optimistic/pessimistic PLM to outperform the fundamental PLM_1 . The remaining part of the conditions ensures that PLM_2 outperforms PLM_3 when all agents use it (this is: $(B - 1)^2 < (B - \alpha)^2$), or PLM_3 outperforms PLM_2 when all agents use it (this is: $(B - 1)^2 < (B - \alpha^{-1})^2$): such conditions depend on whether $\alpha \leq 1$.

For example, selecting $\alpha = -1$, one obtains a setting with two non fundamental $PLMs$, one leading to optimistic and one to pessimistic forecasts, with optimism/pessimism symmetric around the fundamental steady state: in this case, the condition for the strategy profiles where all agents use one of the two non fundamental $PLMs$ to be ESSes reduces to $B > \frac{1}{2}$. Selecting instead an $\alpha \in -(0, \infty)$, the symmetry of optimism and pessimism would not necessarily hold anymore, but the condition for strategy profiles where either PLM_2 or PLM_3 are used by all agents to be ESSes would still be $B > \frac{1}{2}$.

Proposition 14 *Consider the three $PLMs$ defined in Eqs. 36-38. Each of the three 5-tuples $(PLM_1, PLM_2, PLM_3, \mu_1 = 1, \mu_2 = 0)$, $(PLM_1, PLM_2, PLM_3, \mu_1 = 0, \mu_2 = 1)$ and $(PLM_1, PLM_2, PLM_3, \mu_1 = 0, \mu_2 = 0)$ can be an ESS. In particular, the first one is always an ESS; the second one is an ESS if one of the three conditions (39)-(41) is satisfied; the third one is an ESS if one of the three conditions (42)-(44) is satisfied.*

The intuition for these results is straightforward: a PLM prevails on the others if its prediction is closer to the actual realization than each of the alternative's predictions, when all agents use it. Again, these ESSes correspond to equilibria that are stable both under replicator dynamics and under the discrete logit model with $\beta = \infty$.

References

- Agliari A, Massaro D, Pecora N, Spelta A (2017) Inflation targeting, recursive inattentiveness and heterogeneous beliefs. *J Money Cred Bank* 49:1587–1619
- Anufriev M, Hommes C (2012) Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *Amer Econ J Microecon* 4:35–64
- Anufriev M, Assenza T, Hommes C, Massaro D (2013) Interest rate rules and macroeconomic stability under heterogeneous expectations. *Macroecon Dyn* 17:1574–1604
- Aoki M (2002) Open models of share markets with two dominant types of participants. *J Econ Behav Organ* 49:199–216
- Baldwin JM (1896) A new factor in evolution. *Am Nat* 30:441–451
- Berardi M (2015) On the fragility of sunspot equilibria under learning and evolutionary dynamics. *J Econ Behav Organ* 112:251–265
- Borgers T, Sarin R (1997) Learning through reinforcement and replicator dynamics. *J Econ Theory* 77:1–14
- Branch WA, Evans GW (2006) Intrinsic heterogeneity in expectation formation. *J Econ Theory* 127:264–295
- Branch WA, McGough B (2008) Replicator dynamics in a cobweb model with rationally heterogeneous expectations. *J Econ Behav Organ* 65:224–244
- Branch WA, McGough B (2016) Heterogeneous beliefs and trading inefficiencies. *J Econ Theory* 163:786–818
- Brock WA, Hommes C (1997) A rational route to randomness. *Econometrica* 65:1059–1096

- Brock WA, Hommes C, Wagener FOO (2005) Evolutionary dynamics in markets with many traders types. *J Math Econ* 41:7–42
- De Grauwe P (2011) Animal spirits and monetary policy. *Econ Theory* 47:423–457
- Diks CGH, Van der WR (2002) Continuous Beliefs Dynamics ceNDEF Working Papers 02-11, Universiteit van Amsterdam, Center for Nonlinear Dynamics in Economics and Finance
- Diks CGH, van der WR (2003) Heterogeneity as a natural source of randomness. Tinbergen Institute Discussion Papers 03-073/1, Tinbergen Institute
- Evans GW, Honkapohja S (2001) Learning and Expectations in Macroeconomics. Princeton University Press
- Evans GW, Honkapohja S (2003) Expectations and the stability problem for optimal monetary policy. *Rev Econ Stud* 70:807–824
- Evans GW, Honkapohja S, Mitra K (2009) Anticipated fiscal policy and adaptive learning. *J Monet Econ* 56:930–953
- Evstigneev IV, Hens T, Schenk-Hoppé KR (2001) Market selection of financial trading strategies: global stability. *Math Financ* 12:329–339
- Evstigneev IV, Hens T, Schenk-Hoppé KR (2006) Evolutionary stable stock markets. *Econ Theory* 27:449–468
- Fudenberg D, Levine D (1998) The theory of learning in games. MIT Press, Cambridge
- Gilboa I, Matsui A (1991) Social stability and equilibrium. *Econometrica* 59:859–867
- Guse EA (2010) Heterogeneous expectations, adaptive learning, and evolutionary dynamics. *J Econ Behav Organ* 74:42–57
- Hens Th, Schenk-Hoppé R (2005) Evolutionary stability of portfolio rules in incomplete markets. *J Math Econ* 41:43–66
- Lagos R, Wright R (2005) A unified framework for monetary theory and policy analysis. *J Polit Econ* 113:463–484
- Manki C, McFadden D (1982) Structural analysis of discrete data with econometric applications. MIT Press, Cambridge
- Maynard Smith J (1974) The theory of games and the evolution of animal conflicts. *J Theor Biol* 47:209–221
- Maynard Smith J, Price GR (1973) The logic of animal conflict. *Nature* 246:15–18
- Marimon R (1993) Adaptive learning, evolutionary dynamics and equilibrium selection in games. *Eur Econ Rev* 37:603–611
- Matejka F, McKay A (2015) Rational Inattention to discrete choices: a new foundation for the multinomial logit model. *Am Econ Rev* 105:272–298
- Sargent TJ (1999) The Conquest of American Inflation. Princeton University Press
- Sethi R, Franke R (1995) Behavioral heterogeneity under evolutionary pressure: macroeconomic implications of costly optimisation. *Econ J* 105:583–600
- Sznajder B, Sabelis MW, Egas M (2012) How adaptive learning affects evolution: reviewing theory on the Baldwin effect. *Evol Biol* 39:301–310
- Taylor PD, Jonker LB (1978) Evolutionarily stable strategies and game dynamics. *Math Biosci* 40:145–156
- Weibull J (1995) Evolutionary game theory. MIT Press, Cambridge

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.