UvA-DARE (Digital Academic Repository)

From self-fulfilling mistakes to behavioral learning equilibria

Hommes, C.

DOI

[10.1007/978-3-319-44076-7_5](https://doi.org/10.1007/978-3-319-44076-7_5)

Publication date

2017

Document Version

Author accepted manuscript

Published in

Sunspots and nonlinear dynamics

[Link to publication](#)

Citation for published version (APA):

Hommes, C. (2017). From self-fulfilling mistakes to behavioral learning equilibria. In K. Nishimura, A. Venditti, & N. C. Yannelis (Eds.), *Sunspots and nonlinear dynamics: Essays in honor of Jean-Michel Grandmont* (pp. 97-123). (Studies in Economic Theory; Vol. 31). Springer. https://doi.org/10.1007/978-3-319-44076-7_5

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

From Self-Fulfilling Mistakes to Behavioral Learning Equilibria *

Cars Hommes^a

^a*CeNDEF, University of Amsterdam and Tinbergen Institute*

January 30, 2017

Abstract

This essay links some of my own work on expectations, learning and bounded rationality to the inspiring ideas of Jean-Michel Grandmont. In particular, my work on consistent expectations and behavioral learning equilibria may be seen as formalizations of JMG's ideas of self-fulfilling mistakes. Some of our learning-to-forecast laboratory experiments with human subjects have also been strongly influenced by JMG's ideas. Key features of self-fulfilling mistakes are multiple equilibria, excess volatility and persistence amplification.

JEL Classification: D84, D83, E32, C92

Keywords: expectations, learning, chaos, almost self-fulfilling equilibria, laboratory experiments.

E-mail: C.H.Hommes@uva.nl

*In: Nishimura, K., Venditti, A. and Yannelis, N.C. (Eds.), Sunspots and Non-Linear Dynamics. Essays in Honor of Jean-Michel Grandmont. Studies in Economic Theory, Springer Verlag, Cham, 2017, pp. 97-123.

Contents

1	Introduction	3
2	Consistent Expectations Equilibrium	5
3	Learning to believe in (noisy) chaos	10
4	Stochastic Consistent Expectations Equilibrium (SCEE)	13
5	Behavioral Learning Equilibria	15
5.1	Asset pricing model	17
5.2	A New Keynesian Philips curve	21
6	Learning-to-forecast experiments	27
7	Concluding Remarks	32

1 Introduction

The ideas of Jean-Michel Grandmont have inspired the work of many young scholars in economics. During my own PhD thesis work on chaos in economic models (Hommes, 1991), I have for example been studying his seminal contribution on chaos in overlapping generations models (Grandmont, 1985). For many years thereafter, another seminal contribution Grandmont (1998)¹ on expectations formation and stability in large socio-economic systems has provided inspiration for my work on expectations, learning and bounded rationality in the last two decades (see e.g., Hommes, 2013).

Let me start off by quoting JMG at length (Grandmont, 1998, pp.776-777):

Complex “learning equilibria” may be at first sight good candidates to explain why agents keep making significant and recurrent mistakes when trying to predict the fate of socioeconomic systems in which they participate. To be acceptable, however, the observed patterns along such “learning equilibria” should display some reasonable degree of consistency with the agents’ beliefs. One might envision situations in which agents do believe (wrongly) that the world is relatively simple (e.g. linear) but subject to random shocks, and in which the corresponding (deterministic) “learning equilibria” are complex enough to make the agents’ forecasting mistakes “self-fulfilling” in a well defined sense. For instance, the agents might be assumed to have at their disposal a reasonably wide, but nevertheless limited, battery of statistical tests (“bounded rationality”) which would not allow them to reject the hypothesis that their recurrent forecasting mistakes are attributable to random disturbances ... It is not quite clear to me at this stage whether such a program can actually generate operational results or is even feasible (for a first step, see Sorger (1997), Hommes and Sorger (1997)²). Yet progress on this front, if possible, might provide an interesting alternative to our current paradigms, which rely very heavily on extreme, and often criticized, rationality axioms.

This essay summarizes some of my work emphasizing how it has been follow-

¹An essential part of this work was already presented at JMG’s Presidential address at the World Meetings of the Econometric Society, Barcelona, 1990.

²Grandmont cited the working paper versions 1997, published in 1998 in the *Journal of Economic Behavior & Organization* and *Macroeconomic Dynamics*.

ing these ideas. Section 2 starts off from the concept of a consistent expectations equilibrium (CEE), as introduced in Hommes and Sorger (1998), which may be seen as a formalization of Grandmont’s idea of a *self-fulfilling mistake*. Along a self-fulfilling mistake agents incorrectly believe that the economy follows a stochastic process, whereas the actual dynamics is generated by a deterministic chaotic process which is indistinguishable from the former (stochastic) process by linear statistical tests. The concept of CEE was motivated by the fact that piecewise linear asymmetric tent maps generate deterministic chaotic time series with exactly the same autocorrelations structure as a stochastic AR(1) process. Along a (chaotic) CEE agents use a simple linear, AR(1) forecasting rule and, given this belief, the economy follows a nonlinear chaotic asymmetric tent map dynamics with the same autocorrelation structure. Hommes and Sorger (1998) showed the existence of chaotic CEE in the cobweb “hog cycle” model with a backward bending supply curve. They also studied the stability of CEE under learning, introducing *sample autocorrelation* (SAC-) *learning*, where agents learn the two parameters of the AR(1) forecasting rule by the observed sample average and (first-order) sample autocorrelation coefficient.

Section 3 discusses an application of CEE in Hommes and Rosser (2001), in a fishery model with backward bending supply. They simulated stochastic nonlinear models where agents *learn to believe in chaos*, that is, the system converges to a noisy chaotic system, with SAC-learning parameters converging to sample average and sample autocorrelations. This situation qualifies as a *self-fulfilling mistake*: agents can not reject the hypothesis that the economy follows an AR(1) process, while the true law of motion of the economy follows a noisy chaotic process. Section 4 discusses more recent work of Hommes et al (2013) on *stochastic consistent expectations equilibria* (SCEE), generalizing the notion of CEE to a nonlinear stochastic framework. A SCEE is a self-fulfilling mistake where agents learn the correct AR(1) rule, in terms of sample average and sample autocorrelations, in a nonlinear stochastic environment.

A CEE may be viewed as an early example of a *Restricted Perceptions Equilibrium* (RPE), as in Evans and Honkapohja (2001), based on the idea that agents have misspecified beliefs, but within the context of their forecasting model they are unable to detect their misspecification³.

³In his survey Branch (2006) argues that the RPE is a natural alternative to rational expectation

In Section 5 we discuss recent work of Hommes and Zhu (2014), who apply the idea of SCEE in a stochastic *linear* modeling framework. The idea here is that agents use a simple (misspecified) univariate AR(1) forecasting rule in a higher dimensional linear framework. A *behavioral learning equilibrium* (BLE) or, more precisely, a first-order stochastic consistent expectations equilibrium (SCEE), arises when the sample average and the first-order autocorrelations of the AR(1) rule coincide with observed realizations. Hence, along a BLE the parameters of the AR(1) rule are not free, but pinned down by two simple observable statistics, the sample average and the first-order sample autocorrelation. Such a simple, parsimonious learning equilibrium may be a more plausible outcome of the coordination process of individual expectations in large complex socio-economic systems. An interesting feature of BLE is that *multiple equilibria* may arise in very simple settings. Section 6 discusses laboratory experiments on expectations, stressing the empirical relevance of coordination on almost self-fulfilling equilibria in positive feedback systems (Heemeijer et al., 2009) and recent experiments of Arifovic et al., 2015 in a complex overlapping generations framework a la Grandmont (1985). The final section concludes.

2 Consistent Expectations Equilibrium

Consider an expectations feedback system of the form

$$p_t = F(p_t^e), \tag{1}$$

where p_t is the state (or price) of the economy, p_t^e the forecast of the price in period t and F the *actual law of motion* of the economy. In general, the map F may be complex and nonlinear. A well known example of (1) is the classical cobweb “hog cycle” model, where $F = D^{-1}S$ is the composition of inverse demand and supply curves.

Throughout this paper, we assume that agents are boundedly rational and do *not* know the law of motion F of the economy. Rather agents form a belief about the price generating process. Assume that all agents believe that prices are generated by equilibrium (REE) because it is to some extent consistent with Muth’s original hypothesis of REE, while allowing for bounded rationality by restricting the class of the perceived law of motion.

a stochastic AR(1) process, that is, their perceived law of motion (PLM) is given by

$$p_t = \alpha + \beta(p_{t-1} - \alpha) + \delta_t, \quad (2)$$

where α and $\beta \in [-1, 1]$ represent the long run mean and the first-order autocorrelations coefficient of the PLM, and δ_t is an IID noise term. Given the PLM (2) and prices known up to p_{t-1} , the optimal forecast, that is, the prediction for p_t minimizing the mean squared prediction error, is⁴

$$p_t^e = \alpha + \beta(p_{t-1} - \alpha). \quad (3)$$

Given that agents use the linear forecast (3), the *implied actual law of motion* becomes

$$p_t = F_{\alpha,\beta}(p_{t-1}) := D^{-1}S(\alpha + \beta(p_{t-1} - \alpha)). \quad (4)$$

The (observable) *sample average* of a time series $(p_t)_{t=0}^\infty$ is

$$\bar{p} = \lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T p_t \quad (5)$$

and the (observable) *sample autocorrelation* coefficients are given by

$$\rho_j = \lim_{T \rightarrow \infty} \frac{c_{j,T}}{c_{0,T}}, \quad j \geq 1, \quad (6)$$

where

$$c_{j,T} = \frac{1}{T+1} \sum_{t=0}^{T-j} (p_t - \bar{p})(p_{t+j} - \bar{p}), \quad j \geq 0. \quad (7)$$

A consistent expectations equilibrium (CEE) is defined as (Hommes and Sorger, 1998)⁵

Definition A triple $\{(p_t)_{t=0}^\infty; \alpha, \beta\}$, where $(p_t)_{t=0}^\infty$ is a sequence of prices and α and β are real numbers, $\beta \in [-1, 1]$, is called a *consistent expectations equilibrium* (CEE) if

1. the sequence $(p_t)_{t=0}^\infty$ satisfies the implied actual law of motion (4) and is bounded,
2. the sample average $\bar{p}$ in (5) exists and is equal to α , and

⁴More generally, for a (nonlinear) stochastic process Y the optimal forecast conditional on X minimizing the mean squared error is the conditional expectation $E(Y|X)$; see e.g. Hamilton (1994) for a discussion and a proof.

⁵Extensions and applications of CEE include Sögner and Mitlöhner (2002), Tuinstra (2003), Branch and McGough (2005), Lansing (2009) and Bullard et al. (2008, 2010).

3. the sample autocorrelation coefficients ρ_j , $j \geq 1$, in (6) exist and one of the following is true:
 - a. if $(p_t)_{t=0}^\infty$ is a convergent sequence, then $\text{sgn}(\rho_j) = \text{sgn}(\beta^j)$, $j \geq 1$;
 - b. if $(p_t)_{t=0}^\infty$ is not convergent, then $\rho_j = \beta^j$, $j \geq 1$.

A CEE is a price sequence together with AR(1) belief parameters α and β , such that expectations are *self-fulfilling* in terms of the *observable* sample average and sample autocorrelations. The two parameters α and β of the AR(1) forecasting rule are not free, but pinned down by simple observable statistics. Along a CEE expectations are thus correct in a linear statistical sense. Hommes and Sorger (1998) showed that, given an AR(1) belief, there are (at least) three different types of CEE:

- a *steady state CEE* in which the price sequence $(p_t)_{t=0}^\infty$ converges to a steady state p^* , with $\alpha = p^*$ and $\beta = 0$;
- a *2-cycle CEE* in which the price sequence $(p_t)_{t=0}^\infty$ converges to a period two cycle $\{p_1^*, p_2^*\}$, $p_1^* \neq p_2^*$, with $\alpha = (p_1^* + p_2^*)/2$ and $\beta = -1$;
- a *chaotic CEE* in which the price sequence $(p_t)_{t=0}^\infty$ is chaotic, with sample average α and autocorrelations β^j .

A steady state CEE is a REE (at least in the long run) corresponding to some fixed point where demand D and supply S intersect. A 2-cycle CEE also is a REE, where the price jumps back and forth between two different intersection points of the demand and supply curves. A chaotic CEE is a non-rational equilibrium, where agents believe in a linear stochastic law of motion, while the true law of motion is nonlinear (e.g., a piecewise linear tent map) and chaotic. Which of these cases occurs in the cobweb model depends on the implied actual law of motion, i.e. upon the composite mapping $D^{-1}S$ in (4), determined by demand and supply curves. In general, different types of CEE may co-exist as will be discussed below.

Sample autocorrelation learning

The notion of CEE involves an AR(1) belief with fixed parameters α and β , which have been pinned down by two simple statistics, the sample average and the (first order)

sample autocorrelation. But how would agents learn these parameters? Assume that agents use *adaptive learning* to update their belief parameters α_t and β_t , as additional observations become available. There is a large literature on adaptive learning in macroeconomics, see e.g. Sargent (1993) and Evans and Honkapohja (2001) for extensive discussion and overviews. Many adaptive learning algorithms use standard econometrics/statistical tools such as (recursive) ordinary least squares.

Hommes and Sorger (1998) proposed another natural and simple learning scheme called *sample autocorrelation learning* (SAC-learning), with parameters based upon sample average and first order sample autocorrelation coefficient⁶:

$$\alpha_t = \frac{1}{t+1} \sum_{i=0}^t p_i, \quad t \geq 1, \quad (8)$$

$$\beta_t = \frac{\sum_{i=0}^{t-1} (p_i - \alpha_t)(p_{i+1} - \alpha_t)}{\sum_{i=0}^t (p_i - \alpha_t)^2}, \quad t \geq 1. \quad (9)$$

When, in each period, the belief parameters are updated according to (8) and (9) the (temporary) law of motion (4) becomes

$$p_{t+1} = F_{\alpha_t, \beta_t}(p_t) = D^{-1}S(\alpha_t + \beta_t(p_t - \alpha_t)), \quad t \geq 0. \quad (10)$$

One can also rewrite SAC-learning in recursive form. Define

$$R_t = \frac{1}{t+1} \sum_{i=0}^t (p_i - \alpha_t)^2,$$

then the SAC-learning is equivalent to the following recursive dynamical system (Hommes and Sorger, 1998).

$$\begin{cases} \alpha_t = \alpha_{t-1} + \frac{1}{t+1}(p_t - \alpha_{t-1}), \\ \beta_t = \beta_{t-1} + \frac{1}{t+1}R_t^{-1} \left[(p_t - \alpha_{t-1})(p_{t-1} + \frac{p_0}{t+1} - \frac{t^2 + 3t + 1}{(t+1)^2}\alpha_{t-1} - \frac{1}{(t+1)^2}p_t) \right. \\ \quad \left. - \frac{t}{t+1}\beta_{t-1}(p_t - \alpha_{t-1})^2 \right], \\ R_t = R_{t-1} + \frac{1}{t+1} \left[\frac{t}{t+1}(p_t - \alpha_{t-1})^2 - R_{t-1} \right]. \end{cases} \quad (11)$$

⁶Although not identical, SAC-learning is closely related to the ordinary least squares (OLS-learning) scheme; see the discussion in Hommes and Sorger (1998). A convenient feature of the SAC estimate β_t in (9) is that it always lies in the interval $[-1, 1]$, reflecting the fact that the first order autocorrelation coefficient is not explosive, while the OLS-estimate may be outside this interval.

An important feature of CEE and SAC-learning is that both have a simple, intuitive *behavioral interpretation*. In a CEE agents use a linear forecasting rule with two parameters, the mean α and the first-order autocorrelation β . Both can be observed from past observations by inferring the average price level and the (first-order) persistence of the time series. For example, $\beta = 0.5$ means that, on average, prices mean revert toward their long-run mean by 50 percent. The linear univariate AR(1) rule and the SAC-learning process are examples of simple forecasting heuristics that can be used without any knowledge of statistical techniques, simply by observing a time series and roughly "guestimating" its sample average and its first-order persistence⁷.

Which type of CEE exist in the nonlinear cobweb model and to which of them will the SAC-learning dynamics converge? Hommes and Sorger (1998) show that in the simplest case, when demand is decreasing and supply is increasing, the *only* CEE is the REE steady state price p^* . This means that, even when the underlying market equilibrium equations are *not* known, agents will be able to learn and coordinate on the REE price if they learn the correct sample average and sample autocorrelations. Hence, in a nonlinear cobweb economy with monotonic demand and supply, boundedly rational agents should, at least in theory, be able to learn the unique REE from time series observations⁸.

Hommes and Sorger (1998) study CEE in a cobweb model with linear demand and a non-monotonic, piecewise linear backward bending supply curve. They present examples of 2-cycle and chaotic CEE, where, given an AR(1) perceived law of motion, the implied actual law of motion is a chaotic piecewise linear tentmap. Different types of CEE, steady state, 2-cycle and chaotic, may co-exist and the SAC-learning dynamics exhibits *path-dependence*, with the long run CEE depending upon initial states. In the long run however, the SAC-learning always settles down to one of the CEE, where agents have learned the correct sample average and sample autocorrelation.

⁷In learning-to-forecast laboratory experiments for many subjects forecasting behavior is well described by simple rules, such as a simple AR(1) rule; see Section 6.

⁸For the cobweb model, Bray and Savin (1986) show that OLS learning also converges to the REE steady state. In the laboratory experiments of Hommes et al. (2007) however, prices do not always converge to the REE steady state but exhibit excess volatility when the cobweb model is strongly unstable.

3 Learning to believe in (noisy) chaos

Hommes and Rosser (2001) consider another example of a cobweb model with a backward bending supply curve having its origin in a fishery model. This example serves to illustrate how agents may “*learn to believe in chaos*” in a stylized nonlinear, stochastic environment. That is, in an unknown nonlinear environment agents learn the parameters of a simple, linear AR(1) forecasting rule, while the law of motion of the economy is nonlinear and chaotic. In the long run, the sample mean and first order autocorrelation coefficient of the AR(1) rule converge to the observed sample means and first order autocorrelation of the unknown nonlinear chaotic process. Under SAC-learning, the two parameters of the AR(1) rule thus converge, while the implied actual law of motion of the economy converges to a chaotic map. Moreover, agents can not reject the null hypothesis of a stochastic AR(1) process through statistical hypothesis testing and therefore in a linear statistical sense beliefs and realizations coincide⁹.

SAC-learning of a chaotic CEE may be seen as an example of an *approximate rational expectations equilibrium* (Sargent, 1998) or a Restricted Perceptions Equilibrium (RPE) (Evans and Honkapohja, 2001, Branch, 2006). Agents have misspecified beliefs, but within the context of their forecasting model they are unable to detect their misspecification, and they learn the optimal misspecified forecasts. We also would like to stress the *behavioural rationality* interpretation of CEE and SAC-learning, because the simple AR(1) rule is intuitively plausible and SAC-learning may be seen as a learning heuristic through guestimating the sample average and first order sample autocorrelation.

⁹The notion *learning to believe in chaos* has been introduced in Hommes (1998, p.360), and the first examples have been given by Sorger (1998) and Hommes and Sorger (1998). For related work on the instability of OLS learning, see e.g. Bullard (1994) and Grandmont (1998). Schönhofer (1999) has used the notion of learning to believe in chaos in a somewhat different context, namely when the entire OLS-learning process fluctuates chaotically. In Schönhofers’ examples belief parameters of the OLS-learning scheme do *not* converge, but keep fluctuating chaotically and at the same time, due to inflation, prices diverge to infinity, so that agents are in fact running an OLS-regression on a non-stationary time series. Tuinstra and Wagener (2007) consider the same model with heterogeneous expectations, with agents switching between different OLS-estimation methods.

We generalize the law of motion to a nonlinear stochastic system. SAC-learning is given by (8), (9), as before, but we add a noise term to the implied actual law of motion, i.e.

$$p_{t+1} = F_{\alpha_t, \beta_t}(p_t) + \epsilon_t = D^{-1}S(\alpha_t + \beta_t(p_t - \alpha_t)) + \epsilon_t, \quad t \geq 0, \quad (12)$$

where ϵ_t is an independently identically distributed (IID) random process.

Figure 1 illustrates an example of *learning to believe in noisy chaos*. Under SAC-learning the belief parameters $\alpha_t \rightarrow \alpha^* \approx 5400$ and $\beta_t \rightarrow \beta^* \approx -0.87$ converge to constants, while the underlying law of motion F_{α^*, β^*} converges to a chaotic map. Prices then keep fluctuating chaotically with noise. Recall that our boundedly rational agents have no knowledge about underlying market equilibrium equations, and therefore do not know the implied actual law of motion. They only observe time series and update their forecasting parameters based upon simple statistics, the sample average and the first order sample autocorrelation coefficient. Would agents in the long run be satisfied with their linear forecasting rules and stick to their AR(1) belief?

Figure 1e shows that the forecasting errors under SAC-learning are uncorrelated. Agents therefore do not make systematic mistakes, or at least there is no linear structure in their forecasting errors. As a next step, one could do statistical hypothesis testing of the linear forecasting rule. Would boundedly rational agents be able to reject their stochastic AR(1) belief or perceived law of motion by linear statistical hypothesis testing? Hommes and Rosser (2001) show that in this example the null hypothesis that prices follow a stochastic AR(1) process can not be rejected at the 10% level. Agents thus *learn to believe in noisy chaos*.

These equilibria are persistent with respect to dynamic noise. In fact, the presence of noise may increase the probability of convergence to such learning equilibria. Agents are using a simple, but misspecified model to forecast an unknown, possibly complicated actual law of motion. Without noise, boundedly rational agents using time series analysis might be able to detect the misspecification and improve their forecast model. In the presence of dynamic noise however, misspecification becomes harder to detect and boundedly rational agents using linear statistical techniques can do no better than stick to their optimal, simple linear model of the world. This

Figure 1: Learning to believe in noisy chaos. In the presence of noise, SAC learning converges to a (noisy) chaotic CEE, with chaotic prices fluctuations (top left) and at the same time convergence of the belief parameters $\alpha_t \rightarrow \alpha^* \approx 5400$ (mid left) and $\beta_t \rightarrow \beta^* \approx -0.87$ (mid right). Forecasting errors (top right) are (noisy) chaotic and seemingly unpredictable. The ACF of forecasting errors (bottom plot)) shows that errors are uncorrelated.

clearly satisfies Grandmont’s earlier quote of a self-fulfilling mistake (Grandmont, 1998, pp.776-777):

“One might envision situations in which agents do believe (wrongly) that the world is relatively simple (e.g. linear) but subject to random shocks, and in which the corresponding (deterministic) “learning equilibria” are complex enough to make the agents’ forecasting mistakes “self-fulfilling” in a well defined sense. For instance, the agents might be assumed to have at their disposal a reasonably wide, but nevertheless limited, battery of statistical tests (“bounded rationality”) which would not allow them to reject the hypothesis that their recurrent forecasting mistakes are attributable to random disturbances ...”

4 Stochastic Consistent Expectations Equilibrium (SCEE)

Hommes et al. (2013) have generalized the notion of consistent expectations equilibrium to a stochastic setting. Let the law of motion of an economic system be given by the stochastic system

$$x_t = f(x_{t+1}^e, u_t), \quad (13)$$

where x_t is the state of the system at date t , x_{t+1}^e is the expected value of x at date $t+1$, $\{u_t\}$ is an IID noise process with mean zero and f is a continuous (nonlinear) function. Note that the timing is different and (13) has the form of a *temporary equilibrium* map, with the state x_t depending on the expected future state x_{t+1}^e . As before, agents are boundedly rational and do not know the exact form of the (nonlinear) law of motion (13), but rather agents’ perceived law of motion is a stochastic AR(1) process. Given this perceived law of motion, the 2-period ahead forecast x_{t+1}^e that minimizes the mean-squared forecasting error is

$$x_{t+1}^e = \alpha + \beta^2(x_{t-1} - \alpha). \quad (14)$$

Here we use the convention of the learning literature that x_t in (13) is not yet observable when the forecast x_{t+1}^e is made. Combining the forecast (14) and the law of

motion of the economy (13), we obtain the implied actual law of motion (ALM)

$$x_t = f(\alpha + \beta^2(x_{t-1} - \alpha), u_t). \quad (15)$$

Hommes et al. (2013) define a *first-order stochastic consistent expectations equilibrium* (SCEE) as follows.

Definition 4.1 *A triple (μ, α, β) , where μ is a probability measure and α and β are real numbers with $\beta \in (-1, 1)$, is called a first-order stochastic consistent expectations equilibrium (SCEE) if the following three conditions are satisfied:*

- S1 The probability measure μ is a nondegenerate invariant measure for the stochastic difference equation (15);*
- S2 The stationary stochastic process defined by (15) with the invariant measure μ has unconditional mean α , that is, $E_\mu(x) = \int x d\mu(x) = \alpha$;*
- S3 The stationary stochastic process defined by (15) with the invariant measure μ has unconditional first-order autocorrelation coefficient β .*

A first-order SCEE is thus characterized by two consistency requirements: the unconditional mean and the unconditional first-order autocorrelation coefficient generated by the actual (unknown) stochastic process (15) coincide with the corresponding statistics of the perceived linear AR(1) process. Along a SCEE the two parameters α and β of the AR(1) forecasting rule are thus not free, but pinned down by two simple observable statistics. This means that along a first-order SCEE agents correctly perceive the mean and the first-order autocorrelation (i.e., the persistence) of the stochastic state of the economy, without fully understanding its (nonlinear) structure.

Under SAC-learning the actual law of motion becomes

$$x_t = f(\alpha_{t-1} + \beta_{t-1}^2(x_{t-1} - \alpha_{t-1}), u_t), \quad (16)$$

with time-varying parameters α_t, β_t as before in (8-9).

Hommes et al. (2013) study SAC-learning of SCEE in the highly nonlinear, chaotic overlapping generations model of Grandmont (1985) of the form

$$p_t = g(p_{t+1}^e) + \epsilon_t, \quad (17)$$

where g is a non-monotonic map with infinitely many periodic and chaotic equilibria. An interesting finding is that SAC-learning always converges to a simple equilibrium, either a steady state or a 2-cycle, as illustrated in Figure 4. In such a complex OLG-economy, SAC-learning of an AR(1) rule thus leads to learning-to-believe- in a steady state or learning-to-believe in a two-cycle.

The nonlinear framework for SCEE is very general. A drawback of the nonlinear framework however is that computation of first-order autocorrelations is typically not analytical tractable. The next section presents a simpler *linear* framework for SCEE, where agents use a simple, but *misspecified* univariate AR(1) rule in a higher dimensional linear framework.

5 Behavioral Learning Equilibria

Hommes and Zhu (2014) apply the first order SCEE to a *linear* framework, in which the univariate AR(1) forecasting rule is *misspecified*. The simplest class of models arises when the actual law of motion of the economy is a one-dimensional linear stochastic process x_t , driven by an exogenous AR(1) process y_t . More precisely, the actual law of motion of the economy is given by

$$x_t = f(x_{t+1}^e, y_t, u_t) = b_0 + b_1 x_{t+1}^e + b_2 y_t + u_t, \quad (18)$$

$$y_t = a + \rho y_{t-1} + \varepsilon_t, \quad (19)$$

with parameters $b_0 > 0$, b_1 in the interval $(-1, 1)$ ¹⁰, $b_2 > 0$, $a > 0$ and $0 < \rho < 1$; u_t are IID shocks.

The rational expectations equilibrium x_t^* of (18-19) is a linear function of the driving variable y_t , and is given by

$$x_t^* = \frac{b_0}{1 - b_1} + \frac{ab_1b_2}{(1 - b_1\rho)(1 - b_1)} + \frac{b_2}{1 - b_1\rho}y_t + u_t. \quad (20)$$

Its unconditional mean and first-order autocorrelation are given by (Hommes and

¹⁰This assumption is made to ensure stationarity; for $|b_1| > 1$ the dynamics under learning easily becomes explosive.

Figure 2: Convergence of SAC-learning in the OLG-model of Grandmont (1985). The model has infinitely many periodic and chaotic equilibria, but SAC-learning always selects a simple equilibrium outcome, either a steady state (top panels) or a 2-cycle (bottom panels). The plots show the price (P), expected price (EXPP), forecast error and the time-varying parameters α_t and β_t .

Zhu, 2014):

$$\overline{x^*} := E(x_t^*) = \frac{b_0(1 - \rho) + ab_2}{(1 - b_1)(1 - \rho)}, \quad (21)$$

$$Corr(x_t^*, x_{t-1}^*) = \frac{\rho b_2^2}{b_2^2 + (1 - b_1 \rho)^2 (1 - \rho^2) \frac{\sigma_u^2}{\sigma_\varepsilon^2}}. \quad (22)$$

Note that in the special case $\sigma_u = 0$, the above expression reduces to $Corr(x_t^*, x_{t-1}^*) = \rho$, that is, when there is no exogenous noise u_t in (18), the persistence of the REE coincides exactly with the persistence of the exogenous driving force y_t .

Hommes and Zhu (2014) show that for the linear system (18-19) at least one nonzero first-order SCEE (α^*, β^*) exists, with $\alpha^* = x^*$ and $0 < \beta < 1$. They call this equilibrium a *behavioral learning equilibrium (BLE)*, since it provides a simple, parsimonious forecasting rule, with the parameters pinned down by the simple statistics sample average and sample autocorrelation, on which a population of agents in large socio-economic systems may coordinate. Two important applications of this general framework are an asset pricing model driven by AR(1) dividends and a New Keynesian Phillips Curve (NKPC) where inflation is driven by an AR(1) process for marginal costs.

5.1 Asset pricing model

Consider an asset pricing model with a risky asset that pays stochastic dividends y_t following an AR(1) process. The equilibrium price of the risky asset p_t is given by

$$p_t = \frac{1}{R} [p_{t+1}^e + a + \rho y_t], \quad (23)$$

where $R > 1$ is the gross risk free rate of return. Compared to the general framework (18), we have $b_0 = \frac{a}{R}$, $b_1 = \frac{1}{R}$, $b_2 = \frac{\rho}{R}$ and $\sigma_u = 0$.

Using (20), the rational expectations equilibrium p_t^* becomes

$$p_t^* = \frac{aR}{(R-1)(R-\rho)} + \frac{\rho}{R-\rho} y_t. \quad (24)$$

In particular, if $\{y_t\}$ is IID, i.e., $a = \bar{y}$ and $\rho = 0$, then $p_t^* \equiv \frac{a}{R-1} = \frac{\bar{y}}{R-1}$ is constant. The corresponding mean, variance and first-order autocorrelation coefficient of the

rational expectation price p_t^* are given by, respectively,

$$\bar{p}^* := E(p_t^*) = \frac{a}{(R-1)(1-\rho)} = \frac{\bar{y}}{R-1}, \quad (25)$$

$$Var(p_t^*) = E(p_t^* - \bar{p}^*)^2 = \frac{\rho^2 \sigma_\varepsilon^2}{(R-\rho)^2(1-\rho^2)} \quad (26)$$

$$Corr(p_t^*, p_{t-1}^*) = \rho. \quad (27)$$

Under the assumption that agents are boundedly rational and believe that the price p_t follows a univariate AR(1) process, the implied actual law of motion for prices is

$$\begin{cases} p_t = \frac{1}{R} [\alpha + \beta^2(p_{t-1} - \alpha) + a + \rho y_t], \\ y_t = a + \rho y_{t-1} + \varepsilon_t. \end{cases} \quad (28)$$

A straightforward computation shows that the corresponding first-order autocorrelation coefficient $F(\beta)$ of the ALM (28) is given by

$$F(\beta) = \frac{\beta^2 + R\rho}{\rho\beta^2 + R}. \quad (29)$$

Hommes and Zhu (2014) show that in the asset pricing model (28), the BLE (α^*, β^*) is unique, $\alpha^* = \frac{\bar{y}}{R-1} = \bar{p}^*$ and $\beta^* > \rho$. This means that along the BLE the forecast is on average *unbiased* and prices exhibit *persistence amplification*, that is, the persistence β^* is larger than the persistence ρ under RE. Furthermore, the BLE is stable under SAC-learning.

Figure 3a illustrates the existence of a unique BLE $(\alpha^*, \beta^*) = (1, 0.997)$. The time series of fundamental prices and market prices along the BLE $(\alpha^*, \beta^*) = (1, 0.997)$ are shown in Figure 3b, illustrating that the BLE exhibits *excess volatility* compared to the RE solution. Furthermore, along the BLE the first-order autocorrelation coefficient β^* of market prices is larger than that of the fundamental prices ρ , implying that the market price exhibit *persistence amplification*. The autocorrelation functions of the market prices and the fundamental prices are shown in Figure 3c. Persistence amplification leads to much slower decay of the ACF, and the autocorrelation coefficients of prices along a BLE are substantially higher than those of the RE fundamental price.

Figure 4 illustrates how the persistence amplification and excess volatility depend on the autocorrelation coefficient ρ of dividends, which is also the autocorrelation

Figure 3: (a) BLE $\beta^*(= 0.997)$ is the intersection point of the first-order autocorrelation coefficient $F(\beta) = \frac{\beta^2 + R\rho}{\rho\beta^2 + R}$ (bold curve) with the perceived first-order autocorrelation β (dotted line); (b) RE fundamental prices (dotted curve) and market prices (bold curve) along the BLE; (c) Autocorrelation Functions (ACF) of RE fundamental prices (lower dots) and market prices (higher stars) along the BLE. Parameters: $R = 1.05, \rho = 0.9, a = 0.005, \varepsilon_t \sim \text{IID } U(-0.01, 0.01)$ (i.e. uniform distribution on $[-0.01, 0.01]$).

coefficient of the fundamental price. The first-order autocorrelation β^* of market prices is significantly higher than that of fundamental prices, especially for $\rho > 0.4$ (Figure 4a). For $\rho \geq 0.5$ we have $\beta^* > 0.9$, implying that asset prices are close to a random walk and therefore quite unpredictable. Based on empirical findings, e.g. Timmermann (1996) and Branch and Evans (2010), the autoregressive coefficient of dividends ρ is about 0.9, where the corresponding $\beta^* \approx 0.997$, very close to a random walk. In the case $\rho > 0.4$, the corresponding unconditional variance of market prices is larger than that of fundamental prices. As illustrated in Figure 4b, the ratio of the variance of market prices and the variance of fundamental prices is greater than 1 for $0.4 < \rho < 1$, with a peak around 3.5 for $\rho = 0.7$. For $\rho = 0.9$, $\frac{\sigma_p^2}{\sigma_{p^*}^2} \approx 2.5$, that is, excess volatility by a factor of more than two for empirically relevant parameter values.

Figure 5 illustrates that the unique BLE (α^*, β^*) is stable under SAC-learning. Figure 5a shows that the sample mean of the market prices under SAC-learning, α_t , tends to the mean $\alpha^* = 1$, while Figure 5b shows that the first-order sample autocorrelation coefficient of the market prices under SAC-learning, β_t , tends to the

(a) first-order autocorrelation in BLE

(b) Ratio unconditional variances in BLE

Figure 4: (a) first-order BLE β^* with respect to ρ ; (b) ratio of unconditional variances of market prices and fundamental prices with respect to ρ , where $R = 1.05$.

(a)

(b)

(c)

Figure 5: (a) Time series $\alpha_t \rightarrow \alpha^*(1.0)$; (b) time series $\beta_t \rightarrow \beta^*(0.997)$; (c) time series of market prices under SAC-learning and fundamental prices.

first-order autocorrelation coefficient $\beta^* = 0.997$. Figure 5c shows the asset price under SAC-learning, using the same sample path of noise, as the time series of the BLE in Figure 3c. Since the times series are almost the same, SAC-learning converges to the BLE rather quickly.

In summary, the BLE and SAC-learning offer an explanation of high persistence, excess volatility and bubbles and crashes in asset prices within a stationary time series framework.

5.2 A New Keynesian Philips curve

A second application of BLE and SAC-learning uses the New Keynesian macro model (Woodford, 2003). In the New Keynesian Philips curve (NKPC) with inflation driven by an exogenous AR(1) process y_t for the firm's real marginal cost or the output gap, inflation and the real marginal cost (output gap) evolve according to

$$\begin{cases} \pi_t = \delta \pi_{t+1}^e + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t, \end{cases} \quad (30)$$

where π_t is the inflation at time t , π_{t+1}^e is the subjective expected inflation at date $t+1$, y_t is the output gap or real marginal cost, $\delta \in [0, 1)$ is the representative agent's subjective time discount factor, $\gamma > 0$ is related to the degree of price stickiness in the economy and $\rho \in [0, 1)$ describes the persistence of the AR(1) driving process. u_t and ε_t are IID stochastic disturbances with zero mean and finite absolute moments with variances σ_u^2 and σ_ε^2 , respectively. The most important difference with the asset pricing model is that (30) includes two stochastic disturbances, namely the shock ε_t of the AR(1) driving variable and an additional noise term u_t in the New Keynesian Philips curve. We refer to u_t as a supply shock (or markup shock), and to ε_t as a demand shock, that is uncorrelated with the supply shock. We will see that this extra shock allows for the possibility of *multiple equilibria*. Compared with our general framework (18), the corresponding parameters are $b_0 = 0$, $b_1 = \delta$ and $b_2 = \gamma$.

Under the assumption that agents are boundedly rational and believe that inflation

π_t follows a univariate AR(1) process, the implied actual law of motion becomes¹¹

$$\begin{cases} \pi_t = \delta[\alpha + \beta^2(\pi_{t-1} - \alpha)] + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t. \end{cases} \quad (31)$$

The corresponding first-order autocorrelation coefficient $F(\beta)$ of the implied ALM (31) is computed as

$$F(\beta) = \delta\beta^2 + \frac{\gamma^2\rho(1 - \delta^2\beta^4)}{\gamma^2(\delta\beta^2\rho + 1) + (1 - \rho^2)(1 - \delta\beta^2\rho) \cdot \frac{\sigma_u^2}{\sigma_\varepsilon^2}}. \quad (32)$$

Hommes and Zhu (2014) show that for $0 < \rho < 1$ and $0 \leq \delta < 1$, there exists at least one nonzero BLE (α^*, β^*) for the New Keynesian Philips curve (31) with $\alpha^* = \frac{\gamma a}{(1-\delta)(1-\rho)} = \bar{\pi}^*$. Moreover, a BLE is stable under SAC-learning if $F'(\beta^*) < 1$.

For the New Keynesian Philips curve (31), multiple BLE may coexist. In the simulations below, we fix the parameters $\delta = 0.99$, $\gamma = 0.075$, $a = 0.0004$, $\rho = 0.9$, $\sigma_\varepsilon = 0.01$ [$\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$], and $\sigma_u = 0.003162$ [$u_t \sim N(0, \sigma_u^2)$], so that $\frac{\sigma_u^2}{\sigma_\varepsilon^2} = 0.1$. Figure 6a illustrates an example where $F(\beta)$ has three fixed points $\beta_1^* \approx 0.3066$, $\beta_2^* \approx 0.7417$ and $\beta_3^* \approx 0.9961$. Hence, we have coexistence of three first-order BLE (α^*, β_j^*) , $j = 1, 2, 3$. Figures 6b and 6c illustrate the time series of inflation along the coexisting BLE. Inflation has low persistence along the BLE (α^*, β_1^*) , but very high persistence along the BLE (α^*, β_3^*) . The time series of inflation along the high persistence BLE in Figure 6c has in fact similar persistence characteristics and amplitude of fluctuation as in empirical inflation data, e.g., in Tallman (2003). Furthermore, Figure 6c illustrates that inflation in the high persistence BLE has much stronger persistence than REE inflation, where the first-order autocorrelation coefficient of REE inflation is 0.865, significantly less than $\beta_3^* = 0.9961$.

If multiple BLE coexist, the convergence under SAC-learning depends on the initial state of the system, as illustrated in Figure 7. Since $0 < F'(\beta_j^*) < 1$, for $j = 1$ and $j = 3$, while $F'(\beta_2^*) > 1$, (see Figure 6a), the first-order BLE (α^*, β_1^*) and (α^*, β_3^*) are (locally) stable under SAC-learning, while (α^*, β_2^*) is unstable. For initial state $(\pi_0, y_0) = (0.028, 0.01)$ (Figures 7a and 7b), the SAC-learning dynamics (α_t, β_t)

¹¹As in the asset pricing model, we assume that boundedly rational agents do not recognize or do not believe that inflation is driven by output or marginal costs, but simply forecast inflation by an univariate AR(1) rule.

Figure 6: (a) The first-order autocorrelation β^* of the BLE correspond to the three intersection points of $F(\beta)$ in (32) (bold curve) with the perceived first-order autocorrelation β (dotted line); (b) time series of inflation in low-persistence BLE $(\alpha^*, \beta_1^*) = (0.03, 0.3066)$; (c) times series of inflation in high-persistence BLE $(\alpha^*, \beta_3^*) = (0.03, 0.9961)$ (bold curve) and time series of REE inflation (dotted curve).

Figure 7: Time series of α_t and β_t under SAC-learning for different initial values. (a-b) For $(\pi_0, y_0) = (0.028, 0.01)$ SAC-learning converges to the low persistence BLE $(\alpha^*, \beta_1^*) = (0.03, 0.3066)$; (c-d) For $(\pi_0, y_0) = (0.1, 0.15)$ SAC-learning converges to the high persistence BLE $(\alpha^*, \beta_1^*) = (0.03, 0.9961)$.

converges to the stable low-persistence BLE $(\alpha^*, \beta_1^*) = (0.03, 0.3066)$. Figure 7b also illustrates that the convergence of the first-order autocorrelation coefficient β_t to the low-persistence first-order autocorrelation coefficient $\beta_1^* = 0.3066$ is very slow. For a different initial state, $(\pi_0, y_0) = (0.1, 0.15)$, our numerical simulation shows that the sample mean α_t still tends to $\alpha^* = 0.03$, but only slowly (see Figure 7c), while β_t tends to the high persistence BLE $\beta_3^* \approx 0.9961$ ¹² (see Figure 7d).

Numerous simulations (not shown) show that for initial values π_0 of inflation higher than the mean $\alpha^* = 0.03$, the SAC-learning β_t generally enters the high-persistence region. In particular, a large shock to inflation may easily cause a jump of the SAC-learning process into the high-persistence region.¹³ In the following we further indicate how high and low persistence BLE depend on different parameters.

Multiple equilibria and parameter dependence

Figure 8 illustrates how the number of BLE depends on the parameter γ . For sufficiently small $\gamma (< 0.05)$, there exists only one, low persistence BLE β^* (Figure 8a). Moreover, since

$$\frac{\partial F}{\partial \gamma} = \frac{2\rho(1 - \delta^2\beta^4)(1 - \rho^2)(1 - \delta\beta^2\rho)\frac{\sigma_u^2}{\sigma_\varepsilon^2}}{\gamma^3[(\delta\beta^2\rho + 1) + (1 - \rho^2)(1 - \delta\beta^2\rho)\frac{1}{\gamma^2}\frac{\sigma_u^2}{\sigma_\varepsilon^2}]^2} > 0,$$

the graph of $F(\beta)$ in (32) shifts upward as γ increases. At some critical γ -value, a tangent bifurcation occurs. Immediately thereafter, there exist three BLE, β_1^* , β_2^* and β_3^* (see Figure 8b). The low persistence BLE β_1^* and the high persistence BLE β_3^* are stable under SAC-learning, since $0 < F'(\beta_j^*) < 1$, $j = 1$ and $j = 3$, separated by an unstable BLE β_2^* , with $F'(\beta_2^*) > 1$. As γ further increases, another tangent bifurcation occurs and the low persistence BLE disappears. A unique high persistence BLE then remains, which is stable under SAC-learning (Figure 8c).

The dependence of the number of BLE and their persistence upon the parameter γ are quite intuitive. Recall that γ in (30) measures the relative strength of the

¹²As shown in Figure 6a, $F'(\beta_3^*)$ is close to 1 and, hence, the convergence of SAC-learning is slow.

¹³Hommes and Zhu (2014) also simulate the NKPC under SAC-learning with a constant gain parameter (see the online Supplementary Material) and, similar to Branch and Evans (2010), obtained irregular regime switching between phases of very low persistence and phases of high persistence with near unit root behavior.

Figure 8: The figure illustrates how the (co-)existence of low and high persistence BLE β^* depends upon the parameter γ , measuring the relative strength of inflation upon the driving variable, the output gap. (a) $\gamma = 0.01$; (b) $\gamma = 0.075$; (c) $\gamma = 0.1$. Other parameters: $\frac{\sigma_u^2}{\sigma_\varepsilon^2} = 0.1$, $\rho = 0.9$ and $\delta = 0.99$.

driving variable, the output gap or marginal costs, to inflation. When the driving force is relatively weak, a unique, stable low persistence BLE prevails, with much weaker autocorrelation than in the driving variable. At the other extreme, when the driving force is sufficiently strong, a unique, stable high persistence BLE prevails, with significantly stronger autocorrelation and higher persistence than in the driving variable. In the intermediate case, multiple BLE coexist and the system exhibits path dependence, where, depending on initial conditions, inflation converges to a low or a high persistence BLE.

In a similar way, the dependence of the BLE upon the noise ratio $\frac{\sigma_u^2}{\sigma_\varepsilon^2}$ can be analyzed. $F(\beta)$ in (32) can be rewritten as

$$F(\beta) = \delta\beta^2 + \frac{\rho(1 - \delta^2\beta^4)}{(\delta\beta^2\rho + 1) + (1 - \rho^2)(1 - \delta\beta^2\rho) \cdot \frac{\sigma_u^2}{\sigma_\varepsilon^2} \frac{1}{\gamma^2}}.$$

Consequently, the effect of the noise ratio $\frac{\sigma_u^2}{\sigma_\varepsilon^2}$ is inversely related to the effect of γ . Hence, when the ratio $\frac{\sigma_u^2}{\sigma_\varepsilon^2}$ is high, that is, when the markup shocks to inflation are high compared to the noise of the driving variable, a unique, stable low persistence BLE prevails. If on the other hand, the markup shocks to inflation are low compared to the noise of the driving variable, a unique, stable high persistence BLE prevails.

Furthermore, Figure 9 illustrates how the BLE β^* , together with the first-order autocorrelation coefficient of REE inflation, depends upon the parameter ρ , measuring

Figure 9: First-order autocorrelation coefficient of REE inflation (dotted red curve), stable BLE β^* with respect to ρ (bold curves), unstable BLE β^* (dotted curve), where $\gamma = 0.075, \sigma_u = 0.003162, \sigma_\varepsilon = 0.01, \delta = 0.99$.

the persistence in the driving variable. For intermediate values of ρ ($\in [0.84, 0.918]$), two stable BLE β^* coexist separated by an unstable BLE. In the high persistence BLE, β^* is larger than the first-order autocorrelation coefficient of REE inflation, while in the low persistence BLE β^* is smaller than the first-order autocorrelation coefficient of REE inflation. For small values of ρ , $\rho < 0.84$, a unique, stable low persistence BLE prevails, while for large values of ρ , $\rho > 0.918$, a unique, stable high persistence BLE prevails.

Simulations show that, for plausible values of ρ around 0.9, for a large range of initial values of inflation, the SAC-learning converges to the stable, high persistence BLE β^* with very strong persistence in inflation (see e.g. Figure 7d). This result is consistent with the empirical finding in Adam (2007) that the Restricted Reception Equilibrium (RPE) describes subjects' inflation expectations surprisingly well and provides a better explanation for the observed persistence of inflation than REE.

In summary, the dependence of the number of equilibria and whether the persistence is high or low are quite intuitive. This intuition essentially follows from the signs of the partial derivatives of the first-order autocorrelation coefficient $F(\beta)$ in (32) of the implied ALM (31) satisfying:

$$\frac{\partial F}{\partial \gamma} > 0 \quad \frac{\partial F}{\partial (\frac{\sigma_u^2}{\sigma_\varepsilon^2})} < 0 \quad \frac{\partial F}{\partial \rho} > 0 \quad \frac{\partial F}{\partial \delta} > 0. \quad (33)$$

Hence, as in Figure 8, the graph of $F(\beta)$ shifts upwards when γ increases, $\frac{\sigma_u^2}{\sigma_\varepsilon^2}$ decreases, ρ increases or δ increases, and consequently, the equilibria shift from low persistence to high persistence equilibria. When the nonlinearity is strong and F is S-shaped, e.g., as in Figure 6 for empirically relevant parameter values in the NKPC, both the persistence and the number of equilibria shift, and a transition from a unique stable low persistence BLE, through coexisting stable low and high persistence equilibria, to a unique stable high persistence equilibrium occurs. Such a transition from a unique low persistence BLE, through coexisting low and high persistence BLE, toward a unique high persistence BLE occurs when the strength of the AR(1) driving force (the parameter γ) increases, when the ratio of the model noise compared to the noise of the driving force (i.e. $\frac{\sigma_u^2}{\sigma_\varepsilon^2}$) decreases, when the autocorrelation (i.e., the parameter ρ) in the driving force increases, and when the strength of the expectations feedback (i.e., the parameter δ) increases.

6 Learning-to-forecast experiments

In order to study the empirical relevance of different theories of expectations and learning I have been involved in many so-called *learning-to-forecast experiments* (LtFE) in controlled laboratory settings with human subjects. A LtFE, introduced by Marimon and Sunder (1993) and Marimon, Spear and Sunder (1993), is an experiment where a group of subjects repeatedly forecast the price of a good, whose value is endogenously determined by the average group forecast. A LtFE may be seen as an empirical test of the expectations hypothesis of a dynamic economic expectations feedback system, where all other decisions in the economy –consumption, production, trading, etc.– are computerized, consistent with the underlying model assumptions. A LtFE thus becomes an empirical test of the expectations hypothesis, with all other assumptions under the control of the experimenter. See Assenza et al. (2014) for a recent review of LtFEs and Duffy (2014) for a recent collection of state of the art work in experimental macroeconomics.

To my best knowledge, Jean-Michel Grandmont has never conducted experiments with human subjects himself, but his ideas and suggestions have certainly influenced my own experiments and in particular triggered the positive/negative feedback exper-

imental design in Heemeijer et al. (2009). In a workshop in December 2002, in honor of Volker Böhm, I made the claim that lab experiments in positive feedback environments are *more unstable* than those under negative feedback, based on a comparison between asset pricing experiments (Hommes et al., 2005) and cobweb experiments (Hommes et al., 2007). Jean-Michel –who was in the audience– correctly pointed out that this claim was not warranted, as these experiments differed in multiple dimensions. These suggestions led to the design of a new lab experiment comparing positive versus negative feedback systems in Heemeijer et al. (2009).

Positive expectations feedback is characteristic of speculative asset markets, where an increase of the average price forecast of investors causes the realized market price to rise through higher speculative demand. Negative feedback is more important in producer driven markets of perishable consumption goods, where more optimistic expectations about the price of the good lead to higher production and therefore to lower realized market prices. Heemeijer et al. (2009) investigate how the expectations feedback structure affects individual forecasting behaviour and aggregate market outcomes by considering market environments that *only* differ in the sign of the expectations feedback, but are equivalent along all other dimensions. The realized price is a linear map of the average of the individual price forecasts $p_{i,t}^e$ of six subjects. The (unknown) price generating rules in the *negative* and *positive* feedback systems were respectively:

$$p_t = 60 - \frac{20}{21} \left[\left(\sum_{h=1}^6 \frac{1}{6} p_{ht}^e \right) - 60 \right] + \epsilon_t, \quad \text{negative feedback} \quad (34)$$

$$p_t = 60 + \frac{20}{21} \left[\left(\sum_{h=1}^6 \frac{1}{6} p_{ht}^e \right) - 60 \right] + \epsilon_t, \quad \text{positive feedback} \quad (35)$$

where ϵ_t is a (small) exogenous random shock to the pricing rule. The positive and negative feedback systems (34) and (35) have the same unique RE equilibrium steady state $p^* = 60$ and *only* differ in the sign of the expectations feedback map. Both are linear near-unit-root maps, with slopes $20/21 \approx -0.95$ resp. $+20/21$ ¹⁴. Fig. 10 (top panels) illustrates the difference in the negative and positive expectations feedback maps. Both have the same unique RE fixed point 60. A striking feature

¹⁴In both treatments, the absolute value of the slopes is 0.95, implying in both cases that the feedback system is stable under naive expectations.

Figure 10: Laboratory experiments of negative (left panels) versus positive (right panels) feedback systems. Linear feedback maps (top panels), realized market prices (middle panels), six individual predictions (bottom panels) and individual errors (inside bottom panels). In the negative expectations feedback market (left panels) the realized price quickly converges to the RE benchmark 60. In all positive feedback markets (right panels) individuals coordinate on the "wrong" price forecast and as a result the realized market price persistently deviates from the RE benchmark 60.

of the near-unit-root positive feedback map, is that each point is in fact an *almost self-fulfilling equilibrium*. In near unit root positive feedback systems, agents only make small mistakes and these mistakes are almost self-fulfilling. Such near unit root behavior is typical in many macro and financial asset pricing models, where near unit roots arise e.g. due to discount factors close to 1. Will subjects in LtFEs be able to coordinate on the unique RE fundamental price, the only equilibrium that is perfectly self-fulfilling?

Figure 10 (bottom panels) shows realized market prices as well as six individual predictions in two typical groups. Aggregate price behavior is very different under positive than under negative feedback. In the negative feedback case, the price settles down to the RE steady state price 60 relatively quickly (within 10 periods), but in the positive feedback treatment the market price does not converge but rather oscillates around its fundamental value. Individual forecasting behavior is also very different: in the case of positive feedback, coordination of individual forecasts occurs extremely fast, within 2-3 periods. The coordination however is on a “wrong”, i.e., a non-RE-price. Individual errors are small, but strongly coordinated, leading to large aggregate deviations from the rational fundamental price. In contrast, in the negative feedback case coordination of individual forecasts is slower and takes about 10 periods. More heterogeneity of individual forecasts however ensures that, the realized price quickly converges to the RE benchmark of 60 (within 5-6 periods), after which individual predictions coordinate on the correct RE price.

OLG experiments

Arifovic et al. (2016) recently performed LtFEs in the overlapping generations modeling framework of Grandmont (1985)¹⁵. In this highly nonlinear OLG economy, infinitely many different periodic and chaotic perfect foresight equilibria may exist. Grandmont (1985) also showed that each periodic or chaotic perfect foresight equilibrium is stable under some suitable adaptive learning algorithm. The purpose of the LtFE is to test empirically on which of these infinitely many different equilibria a group of subjects may coordinate their individual expectations. The expectations

¹⁵See also Heemeijer et al. (2012) for related individual LtFEs in a different OLG economy.

feedback system is of the form

$$p_t = G_{\rho_2}(p_{t+1}^e), \quad (36)$$

where p_t is the price of the consumption good, p_{t+1}^e is the average price forecast of young consumers, G is a non-monotonic map and ρ_2 is a parameter (the degree of relative risk aversion of the old consumers). For small values of ρ_2 the dynamics of the map is simple, and the system has a stable steady state. For increasing values of ρ_2 the steady state becomes locally unstable and the dynamics exhibits a period-doubling bifurcation route to chaos and generates infinitely many periodic and chaotic perfect foresight solutions.

Figure 11 shows four typical groups of the LtFE of Arifovic et al. (2016). For $\rho_2 = 5$, the map G_5 has an unstable steady state and a stable 2-cycle. In the LtFEs for $\rho_2 = 5$, one group (top left panel) coordinates on a (noisy) 2-cycle, while another group coordinates on a steady state (top right panel). For $\rho_2 = 12$, in the chaotic region of the map, two typical groups are shown, both coordinating on a 2-cycle, one after a long transient and to a somewhat noisy 2-cycle (bottom left panel), and another one converging relatively fast to an almost perfect 2-cycle (bottom right panel). In this LtFE of the highly nonlinear OLG economy, subjects thus *learn-to-believe in a steady state* or *learn-to-believe in a 2-cycle* in a complex chaotic environment.

Figure 12 shows that SAC-learning provides a good fit on these laboratory data, and explains convergence to a steady state (top right panel) and to 2-cycles (top left and both bottom panels). These results are consistent with Hommes et al. (2013) (see Section 4), who showed that SAC-learning in Grandmont's OLG model framework either selects a steady state or a 2-cycle. In Figure 12a constant gain version of the SAC-learning is used, with constant gain parameter $\kappa = 0.2$. The SAC-learning converges rather quickly to the steady state (top right panel) and somewhat slower to the 2-cycle (top left and both bottom panels), but matches both long run outcomes quite nicely. SAC-learning thus explains coordination on simple– steady state and two-cycle– equilibria in a complex environment. If agents (wrongly) believe that prices follow a stochastic AR(1) process, this belief becomes self-fulfilling and selects simple equilibria in the complex, nonlinear OLG-economy as the only long run outcomes.

Figure 11: Lab experiments for the OLG-economy of Grandmont (1985). For $\rho_2 = 5$ the map has a stable 2-cycle and the LtFEs converge either to a 2-cycle (top left) or to a steady state (top right). For $\rho_2 = 12$ the map is chaotic, but both groups in the LtFEs converge to a (noisy) 2-cycle (bottom panels).

7 Concluding Remarks

I am grateful to Jean-Michel Grandmont for much of his inspiration over the years. My work discussed here may be seen as formalizations of his ideas. In particular, the idea of self-fulfilling mistakes (Grandmont, 1998), where agents (wrongly) believe that prices follow a (simple) stochastic process, while the true (unknown) law of motion is a nonlinear complex system, and agents are unable to statistically reject these beliefs, has inspired the research program on bounded rationality and learning. Coordination on almost self-fulfilling equilibria in positive feedback systems in laboratory experiments show the empirical relevance of these ideas. Much work on bounded rationality and learning in large socio-economic systems remains to be done, particularly to study how policy should manage self-fulfilling mistakes.

Figure 12: Two-step ahead simulations under SAC learning with gain coefficient $\kappa = 0.2$. SAC-learning explains coordination on a stable steady state or a stable 2-cycle.

References

- [1] Arifovic, J., Hommes, C.H. and Salle, I. (2016), Learning to believe in simple equilibria in a complex OLG economy - evidence from the lab, Working Paper, University of Amsterdam.
- [2] Assenza, T., Bao, T., Hommes, C.H. and Massaro, D. (2014), Experiments on expectations in macroeconomics and finance, In: Duffy, J. (Ed.), Experiments in Macroeconomics. Research in Experimental Economics, Volume 17, Emerald.
- [3] Branch, W.A., 2006. Restricted perceptions equilibria and learning in macroeconomics, in: Colander, D. (Ed.), Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model. Cambridge University Press, New York, pp. 135-160.
- [4] Branch, W., and B. McGough (2005), Consistent Expectations and Misspecification in Stochastic Non-linear Economies, *Journal of Economic Dynamics and Control*, 29, 659–676.

- [5] Branch, W.A., Evans, G.W., 2010. Asset return dynamics and learning. *Review of Financial Studies* 23(4), 1651-1680.
- [6] Bray, M. and Savin, N. (1986), Rational expectations equilibria, learning and model specification, *Econometrica* 54, 1129-1160.
- [7] Brock, W.A. and Hommes, C.H. (1997), A Rational Route To Randomness, *Econometrica* 65, 1059-1095.
- [8] Brock, W.A., and Hommes, C.H., (1998), Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control* 22, 1235-1274.
- [9] Bullard, J., (1994). Learning Equilibria. *Journal of Economic Theory* 64, 468-485.
- [10] Bullard, J., Evans, G.W., Honkapohja, S., 2008. Monetary policy, judgment and near-rational exuberance. *American Economic Review* 98, 1163-1177.
- [11] Bullard, J., Evans, G.W., Honkapohja, S., 2010. A model of near-rational exuberance. *Macroeconomic Dynamics* 14, 166-188.
- [12] Duffy, J. (Ed.), *Experiments in Macroeconomics. Research in Experimental Economics*, Volume 17, Emerald.
- [13] Evans, G.W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- [14] Grandmont, J.-M., (1985), On endogenous competitive business cycles, *Econometrica* 53, 995-1045.
- [15] Grandmont, J.-M., (1998), Expectation formation and stability in large socio-economic systems, *Econometrica* 66, 741-781.
- [16] Hamilton, J.D. (1994), *Time series analysis*, Princeton University Press, Princeton.

- [17] Heemeijer, P., Hommes, C.H., Sonnemans, J., Tuinstra, J., (2009). Price Stability and Volatility In Markets With Positive And Negative Expectations Feedback. *Journal of Economic Dynamics & Control* 33, 1052-1072.
- [18] Heemeijer, P., Hommes, C.H., Sonnemans, J., Tuinstra, J., (2012). An Experimental Study on Expectations and Learning in Overlapping Generations Models, *Studies in Nonlinear Dynamics & Econometrics* 16(4), article 1.
- [19] Hommes, C.H. (1991), Chaotic dynamics in economic models. Some simple case studies, Groningen theses in Economics, Management & Organization, Wolters-Noordhoff, Groningen.
- [20] Hommes, C.H. (1998), On the consistency of backward-looking expectations: The case of the cobweb. *Journal of Economic Behavior and Organization* 33, 333–362.
- [21] Hommes, C.H., (2013), Reflexivity, Expectations Feedback and Almost Self-Fulfilling Equilibria: Economic Theory, Empirical Evidence and Laboratory Experiments, *Journal of Economic Methodology* 20, 406-419.
- [22] Hommes, C.H., (2013), Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems, Cambridge University Press, Cambridge.
- [23] Hommes, C.H., Rosser, J.B., 2001. Consistent expectations equilibria and complex dynamics in renewable resource markets. *Macroeconomic Dynamics* 5, 180-203.
- [24] Hommes, C.H., Sonnemans, J.H., Tuinstra, J. and van de Velden, H, (2005). “Coordination of Expectations in Asset Pricing Experiments”. *Review of Financial Studies*, 18 (3), 955-980.
- [25] Hommes, C., J. Sonnemans, J. Tuinstra and H. van de Velden (2007), Learning in cobweb experiments, *Macroeconomic Dynamics* 11 (S1), 8-33.
- [26] Hommes, C.H., Sorger, G., 1998. Consistent expectations equilibria. *Macroeconomic Dynamics* 2, 287-321.

- [27] Hommes, C.H., Sorger, G., Wagener, F., 2013. Consistency of linear forecasts in a nonlinear stochastic economy. In: Bischi, G.I., Chiarella, C. and Sushko, I. (Eds.), *Global Analysis of Dynamic Models in Economics and Finance*, Springer-Verlag Berlin Heidelberg, pp. 229-287.
- [28] Hommes, C.H. and Zhu, M. (2014), Behavioral Learning Equilibria, *Journal of Economic Theory* 150, 778–814.
- [29] Lansing, K.J., 2009. Time-varying U.S. inflation dynamics and the new Keynesian Phillips curve. *Review of Economic Dynamics* 12, 304-326.
- [30] Lansing, K.J., 2010. Rational and near-rational bubbles without drift. *Economic Journal* 120, 1149-1174.
- [31] Marimon, R., S.E. Spear and S. Sunder, (1993), "Expectationally Driven Market Volatility: an Experimental Study," *Journal of Economic Theory*, 61, 74–103.
- [32] Marimon, R. and S. Sunder, "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence", *Econometrica*, 61(5), 1993, 1073-1107.
- [33] Schönhofer, M. (1999). Chaotic Learning Equilibria. *Journal of Economic Theory* 89, 1-20.
- [34] Shiller, R.J., (2000), *Irrational Exuberance*, Princeton University Press, Princeton.
- [35] Sögner, L., Mitlöhner, H., 2002. Consistent expectations equilibria and learning in a stock market. *Journal of Economic Dynamics & Control* 26, 171-185.
- [36] Sorger, G. (1998), Imperfect foresight and chaos: An example of a self-fulfilling mistake. *Journal of Economic Behavior and Organization* 33, 363-383.
- [37] Tuinstra, J., 2003. Beliefs equilibria in an overlapping generations model. *Journal of Economic Behavior & Organization* 50, 145-164.
- [38] Tuinstra, J. and Wagener, F.O.O. (2007), On learning equilibria, *Economic Theory* 30, 493-513.

- [39] Woodford, M., (1990). Learning to Believe in Sunspots. *Econometrica* 58, 277-307.