*Macroeconomic Dynamics*, **18**, 2014, 727–752. Printed in the United States of America.
doi:10.1017/S1365100512000582

ARTICLES

# LEARNING THE OPTIMAL BUFFER-STOCK CONSUMPTION RULE OF CARROLL

MURAT YILDIZOĞLU, MARC-ALEXANDRE SÉNÉGAS, ISABELLE SALLE, AND MARTIN ZUMPE
*GREThA*
*UMR CNRS 5113*
*and*
*University of Bordeaux*

This article questions the rather pessimistic conclusions of Allen and Carroll [*Macroeconomic Dynamics* 5 (2001), 255–271] about the ability of consumers to learn the optimal buffer-stock-based consumption rule. To this end, we develop an agent-based model in which alternative learning schemes can be compared in terms of the consumption behavior that they yield. We show that neither purely adaptive learning nor social learning based on imitation can ensure satisfactory consumption behavior. In contrast, if the agents can form adaptive expectations, based on an evolving individual *mental model*, their behavior becomes much more interesting in terms of its regularity and its ability to improve performance (which is a clear manifestation of learning). Our results indicate that assumptions on bounded rationality and on adaptive expectations are perfectly compatible with sound and realistic economic behavior, which, in some cases, can even converge to the optimal solution. This framework may therefore be used to develop macroeconomic models with adaptive dynamics.

**Keywords:** Consumption Decisions, Learning, Expectations, Adaptive Behavior, Computational Economics

## 1. INTRODUCTION

Recent developments of the standard approach to individual and aggregate consumption behavior in the last two decades[1] have been driven mainly by the quest for a better matching with the stylized facts observed in this field. In an extensive set of influential studies provided in this respect, Carroll (1992, 1997, 2001) shows that an amended version of the life cycle/permanent income hypothesis model is able to deliver outcomes that are, in a broad sense, consistent with the main features of the related empirical evidence and, in any case, far more consistent than those

We are very grateful to the anonymous referees, and the participants at 2011, San Francisco SCE Conference, as well as to a seminar at BETA Strasbourg, for their criticisms and comments. Address correspondence to: Murat Yıldızoğlu, GREThA (UMR CNRS 5113), Bordeaux University, Avenue Léon Duguit, F-33608 PESSAC Cedex, France; e-mail: yildi@u-bordeaux4.fr.

 1365-1005/13 

stemming from the traditional modeling frameworks that have been called upon before, such as the perfect certainty model with constant relative risk aversion utility or the certainty equivalent model. In his version of the consumption model, Carroll shows that, under quite mild conditions—regarding consumer behavior under uncertainty—the solution to the optimal consumption problem does exhibit the properties of a "buffer stock" rule, according to which the individual consumer behaves as if she had a target level for a stock of financial assets in mind, and used it to smooth her consumption in the face of an uncertain, periodic income stream.[2] With this consumption rule at hand, the model is able to explain at least three empirical puzzles that cannot be solved under the aforementioned alternative, settings: the "consumption/income parallel," the "consumption/income" divergence, and the stability of the "household age/wealth profile."[3] Moreover, according to Carroll, the buffer-stock model provides a reliable framework for formalizing the Friedmannian conception of the permanent income hypothesis, by explicitly acknowledging the importance of precautionary saving induced by uncertainty about future labor income.

One of the main purposes of Carroll's investigations is to try to reconcile the predictions of one model of individual consumption behavior based on intertemporal optimization and rational expectations with what we do observe in terms of actual consumption and savings patterns. Whether such a framework may plausibly be assumed to underlie the consumption behavior of real-life consumers remains an open question, however. As Carroll himself recognizes, "the sophisticated mathematical apparatus [that is] required to solve [numerically] the optimal consumption problem" [Carroll (2001, p. 41)] seems to act as a sufficient impediment for considering that consumers could be endowed with such numerical capabilities in reality. Indeed, despite its intuitive interpretation and heuristic simplicity, the exact solution to the optimization problem takes the form of a complex nonlinear consumption rule without any explicit analytical formula. Yet, as Allen and Carroll show [Allen and Carroll (2001)], this optimal strategy may be approximated by a linear rule whose adoption generates utility streams that are only slightly lower than those associated with the exact and fully nonlinear solution. This rule recasts the nonlinear optimization problem into a two-dimensional framework that has an intuitive interpretation: the intercept of the rule formula determines the target wealth and the slope, the speed with which the consumer tries to get back to the latter when away from it. As such, this rule may in turn provide a plausible candidate for learning and a relevant basis for testing whether consumers are able to adopt a nearly optimal (intertemporal) consumption behavior in real life. Allen and Carroll (2001) address this issue by considering a set of consumers who engage in a process of trial and error regarding alternative linear consumption rules, and select them according to their welfare properties. On the basis of their simulations, they observe that "the simplified linear [nearly optimal] consumption function is enormously difficult to find by trial and error . . . it takes about a million "years" of model time to find a reasonably good consumption rule by trial and error." Hence, their conclusion: although the "empirical evidence suggests that typical households engage in buffer-stock saving behavior," the "question remains of how consumers come by their consumption rules."

In this paper, we reassess the case for learning the linear buffer-stock rule of Allen and Carroll by considering alternative assumptions about the learning process followed by consumers. By doing so, we aim to investigate which features of learning may be key in this context for pushing the consumption behavior close to the optimal solution. The conclusion of Allen and Carroll seems to indicate that no matter the length of the trial and the number of repetitions (at least for plausible values of their combination), one simple (but systematic) exploration of the strategy space (which obviously includes the linear approximation of the fully optimal solution) proves to be a rather inefficient process for selecting the relevant rule. Other forms of learning processes which, in contrast, do embed feedback from experience to the (dynamic) choice of strategies by the individual consumer may, however, lead to more efficient results. In the following we will analyze three of them, which are usually considered in the learning literature: purely individual adaptive learning based on a combination of discovered strategies and random experimenting [Arifovic (1994); Vriend (2000); Yildizoglu (2002); Vallée and Yıldızoğlu (2009)]; social learning based on imitation [Arifovic (1994); Vriend (2000)]; adaptive individual learning where the strategies are chosen on the basis of adaptive expectations formed by the agents, as a consequence of their experience in the economy [Yildizoglu (2001)].

The first mechanism relies on the adaptation of agents' behavior through both random experimenting and combining already discovered strategies. Arifovic (1994) provides one of the first analyses of this approach in an economic context. What is modeled here is the capacity of the agents to *refine* a population of strategies as a consequence of the performance they obtain with these strategies in their environment, as well as their capacity to adapt their strategies to the evolution of this environment, in a dynamic context. The formalization of this approach usually corresponds to a particular adaptation of the genetic algorithms (GA) introduced by Holland (1992). Many applications of this approach in economics take the form of a social learning process, combining random experimentation by individual agents with imitation of strategies between agents. The originality of our approach is the adoption of a framework that includes purely individual learning in the first place. We nevertheless also analyze the potential role of the social dimension of learning, by introducing an imitation process between agents. This social dimension corresponds to the second mechanism we analyze. Hence, we echo the suggestion of Allen and Carroll according to which "there may be more hope of consumers finding reasonably good rules in a 'social learning' context in which one can benefit from the experience of others" [Allen and Carroll (2001)]. Moreover, Vriend (2000) indicates that social and individual learning can yield very contrasted results [see also Vallée and Yıldızoğlu (2009)].

With the third learning scheme, we introduce a richer framework that aims to overcome the main shortcoming of the preceding schemes: the absence of forward-looking behavior by the agents. Looking forward is important when the

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

agents compare different strategies, before choosing one of them for the current period. If they do not form any expectations, they can only base that decision on the performance that they have observed in the past. And, to this end, they can only compare the strategies they have actually used in the past (and, moreover, in specific economic contexts). To assess the potential performance of these strategies in the current context (which can be completely new to them), or to assess how the strategies they have recently discovered would perform, even if they have not yet used them, they must be able to generalize from past observations. Such generalization requires that the agents develop a representation of their environment and the connection between their decisions and performance. With rational expectations, agents are assumed to know and use the real model of the economy, whereas in a purely adaptive context this assumption is not relevant. In this respect, the approach we introduce is original, because it adopts a framework in which the agents are able to build a representation of their environment [a *mental model*, Holland et al. (1989)], but only on the basis of their past experience. Moreover, this representation evolves as a consequence of this experience.

To analyze the ability of consumers to learn through these mechanisms, we develop a simple computational agent-based model (ABM) directly inspired by the original setup of Allen and Carroll (2001). First, we introduce in this model adaptive learning without expectations, including also a social component that we modulate through a dedicated parameter. In a second stage, learning with adaptive expectations is introduced by endowing each consumer with an artificial neural network (ANN) that captures her mental model of the economy. To our knowledge, this is the first article that considers such a learning process in this setup.[4]

Two articles tackling the same question as ours may be contrasted with the approach adopted here. They are both based on a specific scheme of learning: reinforcement learning. Reinforcement learning corresponds to the selection of an action rule in a set of rules, with a probability that increases with the relative success observed in the past for each rule [Sutton and Barto (1998)]. Howitt and Özak (2009) consider such a reinforcement learning process and show that consumers can discover optimal consumption strategies. But to obtain this result, they need to enhance reinforcement learning with a very sophisticated adjustment mechanism. This latter feature is rather difficult to accept under bounded rationality assumptions, even if the complete learning process is parsimonious in terms of information used by the consumer. Lettau and Uhlig (1999) introduce a much simpler learning framework: a classifier system reduced to its reinforcement learning component. In this setting, agents choose, in each period, the consumption strategy that has obtained the highest average performance in the past. They observe that this mechanism introduces a bias in favor of strategies that yield high performance in periods with high incomes. These strategies are adopted instead of the optimal one, which is introduced into the population from the start. However, the authors disregard the most interesting dimension of classifier systems [Holland and Miller (1991); Holland et al. (1989)], their ability to generalize using a flexible correspondence between the states of the environment and the chosen strategies. It ensues that Lettau and Uhlig use a reinforcement mechanism that is exclusively dependent on the past performance of the strategies. In contrast, we aim to build a framework that is perfectly compatible with bounded rationality, and in which agents can form adaptive expectations by generalizing from their past experience.

In what follows, we proceed through numerical simulations, and analyze our results through standard statistical and econometric methods. Another innovation of this article is the methodology used to conduct the sensitivity analysis in these simulations. Instead of the commonly used Monte Carlo approach, we adopt a design of experiments (DOE) method based on nearly orthogonal Latin hypercubes (NOLH). This method is very promising in simulation studies, because it allows the exploration of the parameter space in a very parsimonious way. The structure of the ABM and our methodology are presented in dedicated sections.

Two main insights may be drawn from our analysis. First, the social dimension of learning does not appear to significantly improve the ability of consumers to discover (and adopt) nearly optimal consumption behavior. Endowing the consumer with the capacity to imitate the best strategy that has been used in the previous period adds only marginally to the performance associated with a purely individual learning scheme. This result suggests that sharing rules corresponding to different contexts (in terms of income and wealth) does not yield more efficient learning when consumers face heterogeneous income shocks. Therefore, contrary to what Allen and Carroll (2001) expect,[5] the social learning process may not lead, in such an environment, to quicker convergence onto the optimal strategy.

What seems crucial (and this is the second insight) for learning better consumption rules is the ability of consumers to build, with the help of their experience, structured representations of their environment. Assuming the existence of such mental representations ensures a much better outcome in terms of convergence toward the optimal rule than what would obtain when the strategy space was explored in an unstructured manner, through random experimenting and some simple combination of already discovered strategies. Moreover, giving to the consumer the ability to look forward over several periods using this representations (i.e., to form expectations about the intertemporal consequences of her current decisions), enhances the convergence process.

Finally, our results show that the genuinely adaptive learning process that we have considered yields rather realistic behavior for the agents (stability of behavior and increasing performance over time). In the simple setup of Allen and Carroll (2001), we further observe that such a process may even converge toward the optimal solution, without assuming that the consumers possess rational expectations beforehand. This feature looks promising for studying adaptive macroeconomic dynamics with learning agents.

We proceed as follows. The next section introduces the original setting of Allen and Carroll (2001) and the buffer-stock rule for consumption, as well as the numerical experiments carried out by these authors. The learning mechanisms explored in our article are presented in the third section. We first quickly present learning without expectations and follow with a more detailed presentation of learning with adaptive expectations. Our simulation protocol and methods of analysis are introduced in the fourth section. Our results are discussed in the fifth section. We first show that purely adaptive individual and social learning do not yield satisfactory consumption behaviours. Only a learning process directed by adaptive expectations gives rise to economically sound consumption behavior. The last section concludes and discusses our results.

## 2. THE ORIGINAL PROBLEM

Following Allen and Carroll (2001), let us consider the intertemporal consumption problem of an individual agent. The consumer aims to maximize discounted utility flows from consumption over the remainder of a (possibly infinite) lifetime,

$$\max_{\{C_s\}_t^{\infty}} E_t \left[ \sum_{s=t}^{\infty} \beta^{s-t} u(C_s) \right], \tag{1}$$

in a setting characterized by the equations

$$A_s = X_s - C_s, \tag{2}$$

$$X_{s+1} = R_{s+1} \cdot A_s + Y_{s+1}, \tag{3}$$

$$C_s \leq X_s \ \forall s, \tag{4}$$

and where the variables are $\beta$, time discount factor; $X_s$, resources available for consumption ("cash on hand"); $A_s$, assets after all actions have been taken in period $s$; $C_s$, consumption in period $s$; $R_s$, interest factor $(1 + r)$ from period $s - 1$ to $s$; $u(C)$, utility derived from consumption; $Y_s$, noncapital income in period $s$.

Allen and Carroll (2001) adopt some more specific assumptions with respect to this general setting. First, they specify the utility function as $u(C) \equiv C^{1-\rho}/(1-\rho)$, with $\rho = 3$, implying that

$$u\,(C) = -\frac{1}{2C^2} < 0, \quad C \neq 0. \tag{5}$$

Furthermore, they set $R = 1$ and $\beta = 0.95$. Finally, they consider a three-point distribution for income:

| $Y$ | 0.7 | 1 | 1.3 |
|---|---|---|---|
| Probability | 0.2 | 0.6 | 0.2 |

with $E[Y] = 1$.

In this framework, Carroll (2004) shows that $C^*(X_t)$ may be rewritten[6] as $C^*(X_t) = 1 + f(X_t - \overline{X}^*)$ for some functional form $f(.)$ with specific properties (but no analytical expression), and with $\overline{X}^*$ a target level for cash on hand (which is assumed to be larger than 1 for the latter expression to be valid).[7] Then a linear (Taylor) expansion of $C^*(X_t)$ can be obtained around the point $X_t = \overline{X}^*$, and is written as[8]

$$C^*\,(X_t) \simeq 1 + \gamma^* \cdot (X_t - \overline{X}^*).$$

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

This expression gives the linear "optimal" buffer-stock rule. Allen and Carroll then consider the family of functions $C^{\theta}(X_t)$ that are indexed by $\theta \equiv \{\gamma_\theta, \overline{X}_\theta\}$ and written as

$$C^{\theta}(X_t) = \begin{cases} 1 + \gamma_\theta(X - \overline{X}_\theta) & \text{if } \gamma_\theta(X - \overline{X}_\theta) \leq \overline{X}_\theta \\ X & \text{if } \gamma_\theta(X - \overline{X}_\theta) > \overline{X}_\theta. \end{cases} \tag{6}$$

Each consumption strategy of the agent can then be represented by a vector: $\theta = (\gamma_\theta, \overline{X}_\theta)$. By construction, when $\theta = \theta^* \equiv (\gamma^*, \overline{X}^*)$, $C^{\theta^*}(X_t)$ corresponds to the Taylor approximation of $C^*(X_t)$ around $X_t = \overline{X}^*$.

In their numerical analysis, Allen and Carroll (2001) adopt the following search space of consumption strategies:

$$\gamma \in [0.05, 1]\,, \Delta\gamma = 0.05,$$

$$\overline{X} \in [1, 2.9]\,, \Delta\overline{X} = 0.1.$$

This setup corresponds to 20 steps for each component, generating a complete strategy space of 400 combinations to explore. Let $\Theta$ be the complete set of these strategies.

Given the steps used for constructing the search space, the element of $\Theta$ that is the closest to the optimal strategy is

$$\theta^* = (\gamma^*, \overline{X}^*) = (0.25, 1.2), \tag{7}$$

$$\Rightarrow C^*(X) = 1 + 0.25\,(X - 1.2)\,. \tag{8}$$

Allen and Carroll (2001) test whether consumers can discover this optimal strategy through a systematic exploration of the strategy space and an estimation of their infinite-horizon utility flow. Each consumer tests all the possible 400 strategies by using each of them to decide on her consumption during $n$ periods, starting with a given initial cash on hand $S_0$. In order to estimate the expected utility flow over all possible random income flows, this $n$-period consumption process is repeated $m$ times. The strategy that gives the highest estimate for the utility flow is then selected by each consumer. Allen and Carroll use the numerical approximation of the value function to evaluate the distance to the optimal value flow observed with this best strategy, and they call this distance the sacrifice value. They consider this process for 1,000 consumers for each combination $(S_0, n, m)$ and compute the average sacrifice value over this population to assess how close this process can get to the optimal utility for the corresponding combination. They show that a sufficiently small sacrifice can only be obtained for a very large number of consumption decisions (in the most extreme case, $n = 50$, $m = 200$, each consumer taking 10,000 effective consumption decisions with each strategy).

Their results indicate that it is not easy for individual consumers to get sufficiently close to the infinite horizon optimum without explicitly solving the full optimization problem, even if one assumes that they use the more parsimonious buffer-stock rule:

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

> ⋯ even when the goal is to learn only this simple approximation, pure trial-and-error learning requires an enormous amount of experience to allow consumers to distinguish good rules from bad ones—far more experience than any one consumer would have over the course of a single lifetime. [Allen and Carroll (2001, p. 268)]

The aim of the following sections is to show that alternative learning schemes could yield more interesting outcomes.

## 3. THREE LEARNING SCHEMES

We now present three different learning schemes that we analyze in the context of the buffer-stock model. The first learning scheme we study is a simple one, based on random experimenting by the agents and the combination of already discovered strategies. We also allow for a possibility of imitation of the strategies between consumers (see the next paragraph).

### 3.1. Purely Adaptive Learning without Expectations

The economy is composed of $n$ consumers, each using an evolving population $\Theta_i \subset \Theta$ of $m$ strategies of type $\theta = (\gamma, \overline{X})$. In the initial period, these strategies are randomly drawn from $\Theta$, each with a random fitness $f \in [0, 1]$.

In each period, each consumer either imitates the behavior of another consumer or uses the strategy for which the highest performance (*fitness*) has been observed in the past (this maximal fitness is just random in the initial period). This performance is computed using the utility obtained with this strategy, the last time the consumer used it:

$$f(\theta) = \exp(u(C(\theta))). \tag{9}$$

When the consumer uses a strategy in a period, she updates its fitness using the utility obtained with that strategy.

Moreover, every *GArate* period, each consumer revises her strategy population through the following three steps:

(1) Reconducting the strategies for the next experimentation period, through a roulette wheel based on the relative performance of the strategies: this selection operator creates a new population of strategies, where the probability of each strategy to be reproduced is proportional to its relative performance $(f_l / \sum_j f_j)$.
(2) Combining the already discovered strategies (crossover): with a probability $p^{\mathrm{C}}$, each strategy in the population can have the chance of being combined with another, randomly drawn, strategy. If strategies $\theta_i$ and $\theta_j$ are chosen, they are replaced with two new strategies: $\theta_k = (\gamma_i, \overline{X}_j)$ and $\theta_l = (\gamma_j, \overline{X}_i)$.
(3) Random experimenting (mutation): independently from the crossover, with a probability $p^m$, each strategy can see one of its components ($\overline{X}_j$ or $\gamma_j$) modified by drawing a new value from the corresponding strategy space.

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

1. (B) create $n$ consumers, each consumer $i$ using a population $\Theta_i \subset \Theta$ of $m$ strategies of type $\theta_{il} = (\gamma_{il}, \overline{X}_{il})_{l=1...m}$
2. (B) randomly draw the initial strategy population of each consumer and the corresponding fitness values
3. (B) draw randomly an element of $\Theta_i$ as the initial strategy of the consumer
4. (B) compute the consumption of each consumer with this strategy given the common initial resources $X_0 \in \{0, 1, 2, 3\}$

$$C_0 = \min\left\{1 + \gamma_0\left(X_0 - \overline{X}_0\right), X_0\right\}$$

5. (B) compute the initial utility of each consumer $u_0 = U(C_0)$
6. (S) compute the consumption level and the utility performance that the consumer would have attained using the optimal strategy $\theta^*$

$$C_0^* = \min\left\{1 + \gamma^*\left(X_0 - \overline{X}^*\right), X_0\right\}, \; u_0^* = U\left(C^*\right)$$

7. (S) compute the distances to these optimal levels with the following indicators:

$$\Delta Z \equiv Z^* - Z, \; Z = \overline{X}, \gamma, C, X, u$$

8. (B) compute the new cash on hand of the consumer $X_1 = X_0 - C_0$ and (S) that which would have resulted from the use of the optimal strategy $X_1^* = X_0 - C_0^*$
9. (S) compute other individual and global indicators
10. for $t \leq T$ (T is the length of each run),
    (a) (B) draw a new income for each consumer and compute the new cash in hand $X_t = X_{t-1} - C_{t-1} + Y_t$ and (S) the corresponding optimal flow $X_t^* = X_{t-1} - C_{t-1}^* + Y_t$
    (b) (B) select a strategy for each consumer:
        i. with a probability $p^{\mathrm{I}}$ the consumer imitates the best strategy observed in $t-1$, in the population of agents, and the imitated strategy replaces the strategy with the lowest fitness
        ii. with a probability $(1 - p^{\mathrm{I}})$ the consumer uses the best strategy in $\Theta_i$
    (c) execute steps (4) – (9) using the selected strategy and $X_t$
    (d) (B) if $t$ mod $GARate = 0$, the strategy population is updated using selection, crossover, mutation operators

**FIGURE 1.** Pseudo code of the learning-without-expectations model. We distinguish computations related to the behavior of agents (B) from the computation of statistics or indicators for analysis (S).

In each period, we measure the distance between the observed consumption and corresponding utility levels on the one hand, and the optimal values we would observe with the behavior given by equation (7), on the other hand (see Section 4.2 for more details on these indicators).

The complete structure of the model (its pseudo code) is given in Figure 1.

### 3.2. Social Dimension of Learning: Imitating Successful Consumers

With a probability $p^{\mathrm{I}}$ each consumer can imitate the strategy used in the preceding period by the consumer who has obtained the highest utility. Allen and Carroll evoke a potentially positive role for social learning in the search for the optimal consumption strategy. Imitation is indeed the simplest way of introducing the diffusion of good strategies within the population of consumers.

We will analyze together the outcomes related to those two learning processes.

### 3.3. Learning with Adaptive Expectations

The previous learning schemes are based on the use of the discovered best strategy by the consumer. She chooses a particular strategy on the basis of the performance observed the last time she used this strategy, even if this performance was obtained under specific circumstances (resulting mainly from past income shocks and consumption decisions). But this is not necessarily a very relevant basis for assessing the performance of a strategy under current conditions. In other words, these learning schemes are purely adaptive, and the decisions are not based on the projection of past performances on the future circumstances. Such a projection would require a capacity of the agents to *generalize* or form *expectations*. This generalization would in turn entail that the agent develops a representation of her environment [*mental model*, Holland et al. (1989)]. We consider now consumers who are able to develop such a representation as a result of their past experience.

The mental model of each agent summarizes the state of the agent's knowledge and evolves as a consequence of evolution of this knowledge. It guides the decision process because it enables the agent to test the connections between the alternative choices and their consequences. The presence of such an internal model can reflect the intentionality of decisions. Obviously, in this context, the concept of "model" must be understood in a very loose sense. More than a mathematical construction, it consists in a representation of the agent's perception of the environment: "In . . . situations [that are not sufficiently simple as to be transparent to the human mind], we must expect that the mind will use such imperfect information as it has, will simplify and represent the situation as it can, and make such calculations as are within its powers" [Simon (1976, p. 144)]. These calculations are "as if" experiments that enable agents to evaluate the possible consequences of their decisions. In other words, before making a decision, the agent simulates the potential outcomes of different decisions by using her internal model. The output of these simulations provides the expectations of the agent. The agent makes a decision on the basis of these expectations. This decision yields an effective outcome, which can be compared with the expected one resulting from the simulations. Discrepancies between those outcomes may lead to an update of the mental model. Hence, we have a dynamic structure that evolves as depicted in Figure 2 [Yildizoglu (2001)].

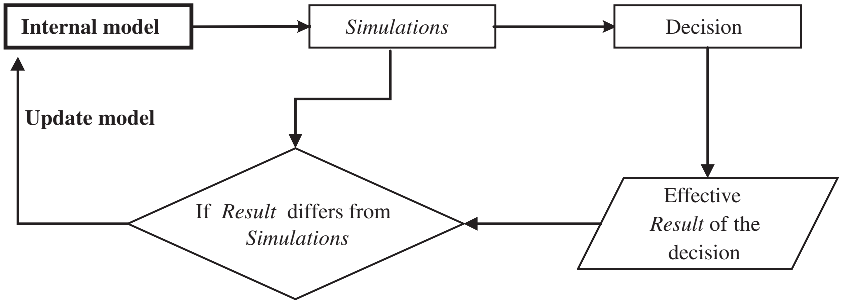


**FIGURE 2.** Dynamics of the mental model of the agents.

Although this line of thought is quite obvious, its integration into economic models is problematic. This is why purely adaptive models (see the preceding section) generally neglect the dynamic process of expectation formation. This representation of learning, as the product of an evolutionary algorithm, does enable the elaboration of better decision rules, but only through trial and error. In this case, the agent can only judge decisions that have been used before. In contrast, the vision based on the dynamics of the internal model admits that agents can have a relatively precise (if not perfect) perception of the value of their decisions, even if they have never been made before. This is made possible by means of simulations using the internal model.

The standard way of formalizing such a model is to rely upon the subjective probabilities approach of Savage. In this case, the internal model of the agent corresponds to a set of conditional probability distributions. The update of this model can be imagined through successive least-squares estimates or applications of Bayes's rule. The Bayesian approach has the advantage of not assuming any particular structure for the internal model. But it is very demanding in terms of agents' rationality. Moreover, "there is substantial evidence that Bayes' theorem lacks empirical relevance and hence its procedural justification is weak" [Salmon (1995, p. 245)].

Recursive least-squares estimates have been used, in this perspective, albeit at the aggregate level, by the recent macroeconomic learning literature [Evans and Honkapohja (2001)]. However, this method relies upon a specific functional structure for the internal model. We adopt, here, a more flexible tool. Our approach is independent of the structure and the parameterization of the internal model, incorporating only its most primitive dimensions: its existence and its influence on the decisions of agents. In this respect, an artificial neural network (ANN) is a good candidate for representing the role of the internal model, and its adaptive nature. With only minimal structural assumptions, namely the list of endogenous and explicative variables, and the structure of the hidden layer, it can represent the fact that the agent adjusts the internal model to the flow of experience. For many practical problems, even a very simple feedforward ANN with one hidden layer of

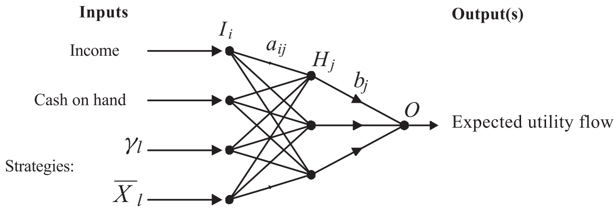


**FIGURE 3.** A feedforward ANN with one hidden layer.

a few hidden nodes gives quite robust results [see Masters (1993) for a discussion of properties of ANNs].

Another potentially interesting modeling approach is learning classifier systems (LCSs). A complete LCS, combining a generalization capability with a reinforcement learning mechanism [such as the XCS developed by Wilson (1995)], could model context-dependent choice of strategies by agents. One of the authors has already tested this approach in modeling industrial dynamics. The main limit of this approach, in the context of our discussion, is the fact that expectations included in the rules of the XCS are necessarily implicit, and it is impossible to separate them from the actions. The use of a mental model represented by an ANN allows such a separation. Moreover, the behavior modeled using this representation seems more realistic (exhibiting some inertia, but also, performance increasing in time).[9]

More particularly, an ANN provides a time-varying flexible functional form that delivers an approximation of the connections between the inputs and the output of the internal model. This approximation is obtained by the calibration of the parameters of the ANN ($a_{ij}$ and $b_j$ in Figure 3) according to the series of input and output data submitted to the ANN in successive training periods. To train the ANN, the complete past history of inputs and outputs can be used, or only observations for a given number of past periods (*windowSize*). In each training period (an *epoch*), a number of passes (*numEpoch*) through the ANN are executed in order to correct the error observed between the observed outputs and the predicted ones. Each pass adjusts the parameters $a_{ij}, b_j$ in order to correct a fraction *learnRate* of the residual error. This repetitive adjustment process aims to minimize the prediction errors of the ANN, indicating better adaptation of the ANN to the environment.

Parameters $a_{ij}, b_j$ reflect the intensity of the connections in the network. A better approximation can be achieved through the introduction of hidden nodes in the network, that is, nodes that represent unobserved state variables or, more particularly, the unobserved variables of the internal model of the agent. ANN thus covers a wide range of models from the simplest linear one where there is no hidden layers, to the increasingly sophisticated ones where the number of hidden nodes (*numHidden*) increases. This number can even be used to represent the complexity of the agent's internal model.

In our case consumers are placed in a very simple context. They can observe two contextual variables: their *cash on hand* and income. Their strategies have two components: $\theta = (\gamma, \overline{X})$. These four elements naturally constitute the inputs of their model. The strategy component of the inputs is used by them to compare different potential strategies on the basis of the expected utility flows they can yield (hence the unique output of the mental model). This comparison serves as a basis for selecting the consumption strategy that will be used in the current period, after observation of the corresponding income.

More particularly, at each period $t$, each consumer uses the ANN as follows. At the beginning of the period, she compares strategies on the basis of the expected utility flow resulting from them. She *feeds* the ANN with the state of the environment and the components of each strategy, and observes the utility flow predicted by the ANN:

$$U_t^e = \sum_{\tau=0}^{\tau=forwardLook} \beta^\tau u_{t+\tau}^e. \tag{10}$$

This utility flow depends on the horizon that is considered by the consumer. This horizon is characterized by the parameter *forwardLook*. If *forwardLook* = 0, the consumer is only interested in her immediate utility (she is myopic); otherwise she tries to take into account the future utility impact of her current consumption strategy. She adopts the strategy that yields the highest expected utility flow.

At the end of period $t$, she acquires a new observation point $(X_t, Y_t, \gamma_t, \overline{X}_t; u_t)$, and she can adjust her mental model by training it, using data for the last period for which she now has a complete set of observations. If the consumer is only interested in the expectation of the current utility $(u_t^e)$, each period's observations can be used to train the ANN before its use in the following period. If the consumer is less myopic, *forwardLook* observations of the output are necessary to train the ANN in each period. To this aim, in period $t$, the consumer can compute the difference (error) between, on the one hand, the expectations formed and used in period $t_0(= t - forwardLook)$, and on the other hand, the *forwardLook* observations of utility between $t_0$ and $t$ (because $u_t$ is necessary to compute the complete utility flow that has resulted from the strategy used in period $t_0$). Then she trains the ANN using the following supplementary inputs and output:

$$\begin{gathered}\left(X_{t-forwardLook}, Y_{t-forwardLook}, \gamma_{t-forwardLook}, \overline{X}_{t-forwardLook}\right) \\ \rightarrow U_{t-forwardLook} = \sum_{\tau=0}^{\tau=forwardLook} \beta^\tau u_{t-forwardLook+\tau}.\end{gathered} \tag{11}$$

As a consequence, in our model, the role of *forwardLook* is twofold: on one hand, a longer horizon yields less myopic decisions; on the other, it imposes on the agent the use of a more out-of-date mental model for forming expectations.

The pseudo code of the model is summarized in Figure 4.

1. (B) create $n$ consumers, each consumer $i$ using a population $\Theta_i \subset \Theta$ of $m$ strategies of type $\theta_{il} = (\gamma_{il}, \overline{X}_{il})_{l=1...m}$;
2. (B) $\gamma$ and $\overline{X}$ belong to the original strategy space of Allen and Carroll (2001)
3. (B) randomly draw the initial strategy population of each consumer
4. (B) randomly initialize the ANN of each consumer
5. (B) draw randomly an element of $\Theta_i$ as the initial strategy of the consumer
6. (B) compute the consumption of each consumer with this strategy given the common initial resources $X_0 \in \{0, 1, 2, 3\}$

$$C_0 = \min\left\{1 + \gamma_0\left(X_0 - \overline{X}_0\right), X_0\right\}$$

7. (B) compute the initial utility of each consumer $u_0 = U(C_0)$
8. (S) compute the behavior and performance that the consumer would have using the optimal strategy $\theta^*$

$$C_0^* = \min\left\{1 + \gamma^*\left(X_0 - \overline{X}^*\right), X_0\right\},\ u_0^* = U\left(C^*\right)$$

9. (S) compute the distances to the optimal behavior and results

$$\Delta Z \equiv Z^* - Z,\ Z = \overline{X}, \gamma, C, X, u$$

10. (S) compute other individual and global indicators
11. (B) compute the new cash on hand of the consumer $X_1 = X_0 - C_0$ and (S) the one she would have using the optimal strategy $X_1^* = X_0 - C_0^*$
12. for $t \leq T$ (T is the length of each run),
    (a) (B) draw a new income for each consumer and compute the new cash in hand $X_t = X_{t-1} - C_{t-1} + Y_t$ and (S) the corresponding optimal flow $X_t^* = X_{t-1} - C_{t-1}^* + Y_t$
    (b) (B) selection of the strategy:
        i. with a probability $p^I$ the consumer imitates the best strategy observed in $t - 1$, in the population of agents, and the imitated strategy replaces the strategy with the lowest fitness
        ii. with a probability $(1 - p^I)$ and if $t > forwardLook$ the consumer chooses a new strategy from her strategy population using her expectations given by her mental model; if $t \leq forwardLook$, consumers chooses randomly a strategy $\Theta_{it}$
    (c) execute steps $(6) - (10)$ using the selected strategy and $X_t$
    (d) (B) if $t > forwardLook$, train the ANN with the observation corresponding to the period $t - forwardLook$, using *windowSize* past observations
    (e) (B) if $t$ mod $GARate = 0$, the strategy population is updated using selection, crossover, mutation operators and the expected fitness of the elements of the new population is updated using the actual state of the ANN
    (f) (S) compute individual and global indicators

**FIGURE 4.** Pseudo code of the learning with expectations model. We distinguish computations related to the behavior of agents (B) from computations of statistics or indicators for analysis (S)

In each period, the training (step 12.d) is done using only *windowSize* past observations of the input (including effectively used pairs $(X, \gamma)$) and output (utility flow) vector. The most recent observations come from the period $(t - forwardLook)$. When forming expectations to select strategies (steps 12.b and 12.e), the agent only updates the expected performance of consumption strategies that are currently available in his or her strategy population of size $m = 40$. These are the only strategies that are *visible* to the agent at this period.

As we have noted, several parameters condition the learning capacity of the ANN: the number of hidden nodes (*numHidden*); the data window used for the training (*windowSize*); the error correction rate in each epoch of training (*learnRate*); the number of passes used in each training epoch (*numEpoch*). Consequently, we use the standard back–propagation of errors to train the ANN. The names and explored values of these parameters are given in the Appendix. Except in extreme cases (when *windowSize* is low, for example: 50), their values do not play a major role in our results.

## 4. SIMULATION PROTOCOL AND METHODS OF ANALYSIS

### 4.1. Experimentation Protocol

Large sampling methods such as Monte Carlo simulations come at a computational cost if there are numerous parameters with large experiment domains.

We would indeed need to implement a very large number of simulations to obtain a representative sample of all parameter configurations. In this context, a design of experiments (DOE) approach[10] allows us to minimize the sample size under a constraint of representativeness. This method provides a sample, namely a design, of the whole set of parameter (or factor) values. The chosen configurations are called *design points*. Some properties of the design are useful. Uniform designs [see, for example, Fang et al. (2000)], such as Latin hypercubes, typically have good space-filling properties; i.e. they correctly cover the whole parameter space.[11] Moreover, Latin hypercubes ensure that linear effects of the factors are noncorrelated, and they are widely used in computer simulations [Ye (1998); Butler (2001)]. Nevertheless, this orthogonality comes at the cost of deteriorated space-filling properties. Accordingly, Cioppa (2002) proposes a nearly orthogonal (NOLH) design that offers an efficient trade-off between orthogonality and space-filling properties [see also Kleijnen et al. (2005); Cioppa and Lucas (2007)].

For each version of the model, we use the same NOLH design to sample the parameter space using Sanchez (2005). Up to 11 factors, the resulting NOLH design provides 33 design points [see Sanchez (2005) for further details, and Table A.1 in the Appendix for the values of the parameters used in the experiments]. We launch 20 replications of each experiment, with a duration of $T = 1{,}250$ period, in order to take into account the diversity of the random draws. This setup corresponds to 660 runs in total and we sample the results every 50 periods during each run. We have $n = 20$ consumers and each consumer is given a strategy population of size $m = 20$.

### 4.2. Indicators and Analysis of Results

The main indicators that we use in the analysis are dedicated to measuring the distance to optimal behavior and performance, as indicated in step 9 of Figure 4:

$$\begin{aligned} \text{sumDistCons} &= \textstyle\sum_{i=1}^{i=n} |C_i^* - C_i|, \\ \text{sumDistUtility} &= \textstyle\sum_{i=1}^{i=n} |U_i^* - U_i|, \\ \text{sumDistGamma} &= \textstyle\sum_{i=1}^{i=n} |\gamma^* - \gamma_i|, \\ \text{sumDistX} &= \textstyle\sum_{i=1}^{i=n} |X^* - X_i|. \end{aligned} \tag{12}$$

Using absolute values gives a full assessment of the distance, because we eliminate all possible compensation between the distances of the consumers. We also consider the variances of these absolute distances, in order to check if individual consumers' behavior converges. We use simple time plots and boxplots to study the evolution of these distances in time and their distributions between different configurations. Boxplots give the four quartiles of the distribution, and the median corresponds to the middle bar. We use R-project [R Development Core Team (2003)] and the `ggplot2` library [Wickham (2009)] to conduct this analysis.

## 5. RESULTS

### 5.1. Individual Learning

Figure 5 shows that the agents are able to learn the optimal consumption levels somewhat, and we observe that the variance of the consumption levels is also decreasing in time. But their performance in terms of utility is not satisfactory at all. Even in the latest periods, they remain collectively far from the optimum and a very high discrepancy exists between their individual performances. Figure 6 shows that even if they are able to converge toward $\overline{X}^*$, the distance to $\gamma^*$ increases and remains high until the end of the simulations. Basing the selection of the consumption strategy to be used only on past individual performance is not able to really structure the learning of the agents. This type of learning is not able to discover strategies specifically adapted to the current consumption context of the agent.

### 5.2. Social Dimension of Learning

Different social learning profiles are pooled together in the preceding results. If we distinguish configurations where imitation is frequent from the ones where it is rarer, we can observe the role played by social learning. Figure 7 distinguishes results in different configurations according to the corresponding intervals

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

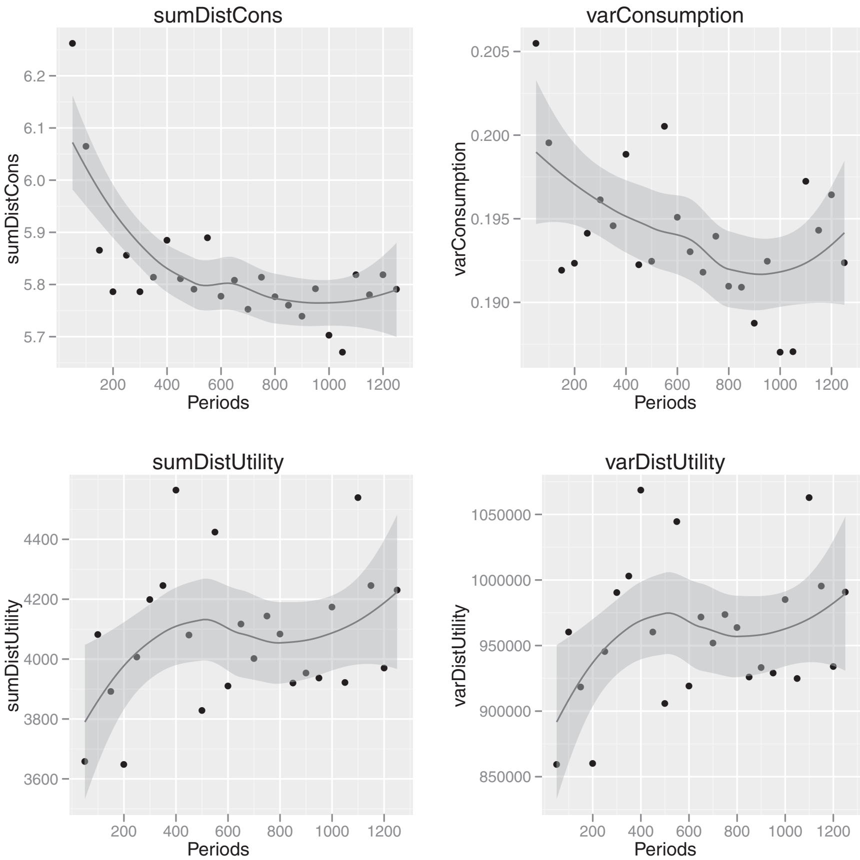


**FIGURE 5.** Individual learning and convergence to the optimal strategy (average of each indicator, in each period, over all experiments and all runs).

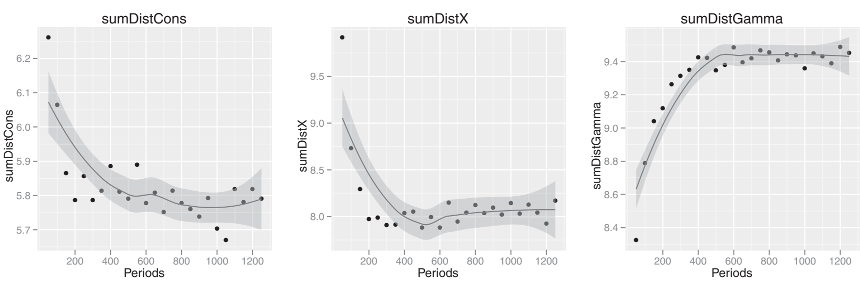


**FIGURE 6.** Learning without expectations: convergence in time on optimal consumption, but not really on its components (average of each indicator, in each period, over all experiments and all runs).

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

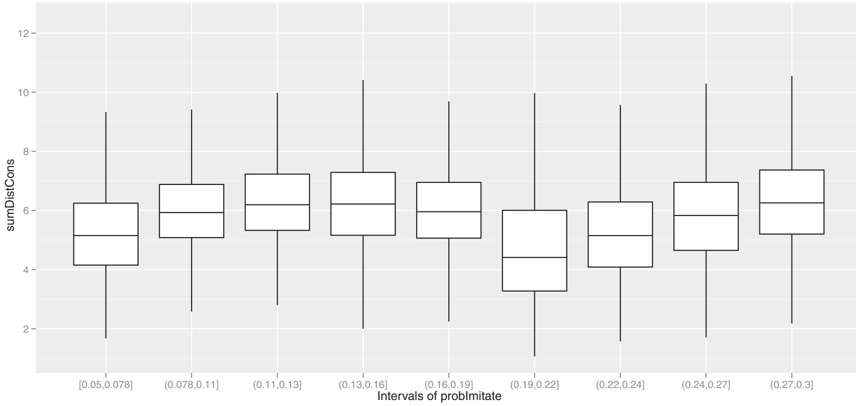


**FIGURE 7.** Individual learning and the role of imitation (distribution of sumDistCons over the corresponding experiments and all runs, for $t > T/2$).

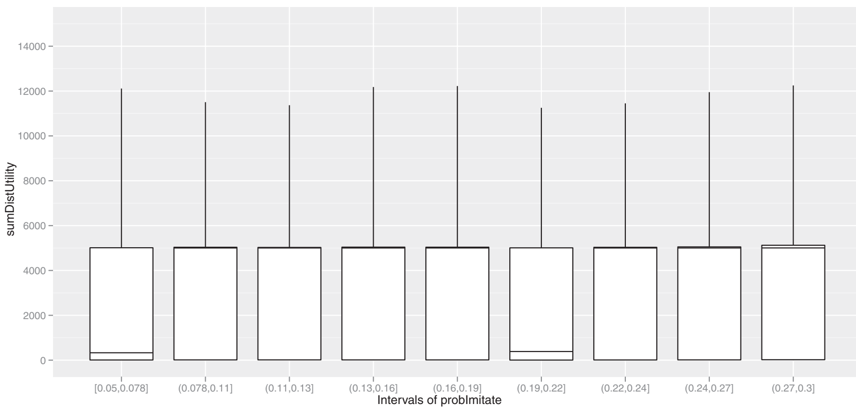


**FIGURE 8.** Role of imitation regarding utility sacrifice (distribution of sumDistUtility over the corresponding experiments and all runs, for $t > T/2$).

of $p^{\mathrm{I}}$ values. It exhibits an intermediate range of imitation probability that minimizes the total distance to optimal consumption levels. Figure 8 confirms this result from the point of view of the total utility sacrifice: it is minimal for the same configurations, $p^{\mathrm{I}} \in (0.19, 0.22]$, but still remains high. As with the selection of strategies on the basis of past individual performance, guiding this choice by the past performance of other individuals does not correctly take into account the current context, and yields relatively mediocre performance. Sharing these past experiences is not really able to correctly guide the learning process either.

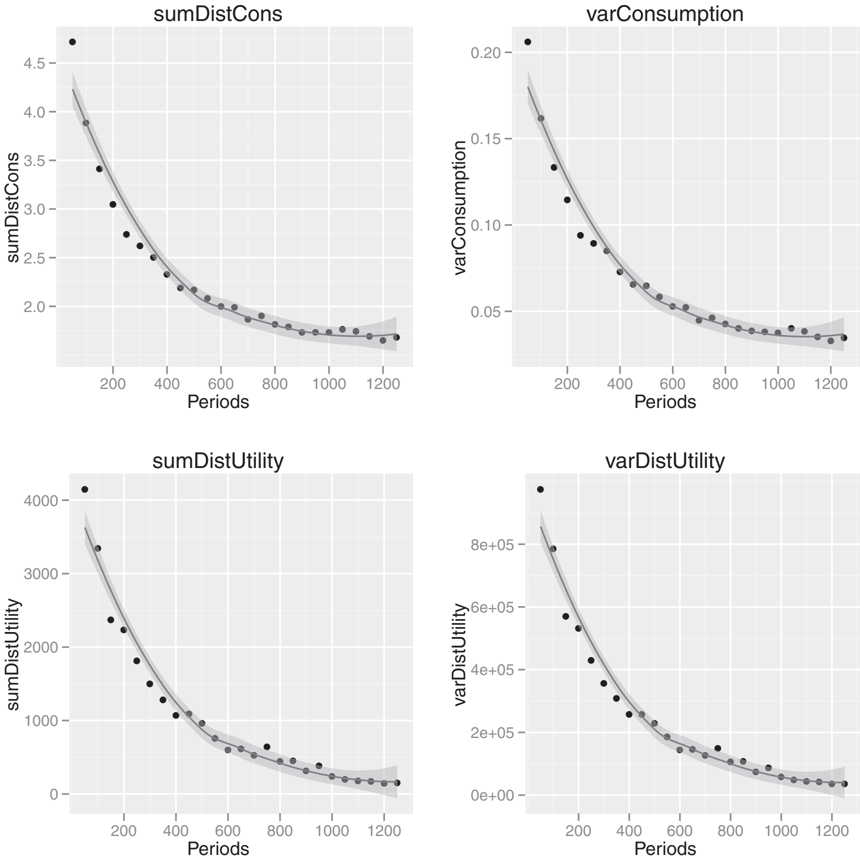


**FIGURE 9.** Learning with expectations (average of each indicator, in each period, over all experiments and all runs).

## 5.3. Individual Learning with Expectations

In contrast with the preceding outcomes, learning with expectations corresponds to a continuous improvement in the performance of the consumers. Figure 9 shows that the total distance to optimal consumption and to optimal level of utility decreases in time, as well as the distance between consumers. These results can clearly be distinguished from the ones obtained earlier. Forming adaptive expectations allows consumers to better discover consumption strategies that improve their utility.

We should also remark that a total distance of 2 corresponds to an average individual distance of 0.1 from the optimal consumption level for each consumer. This is a remarkable performance if we consider that these consumers are not supposed to solve an infinite-horizon optimization problem.

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

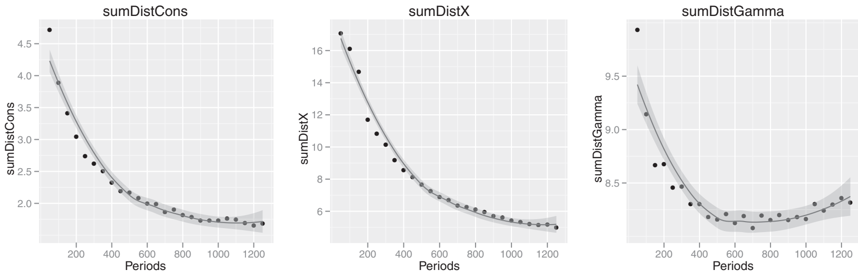


**FIGURE 10.** Learning with expectations: convergence in time on the optimal consumption strategy and its components (average of each indicator, in each period, over all experiments and all runs).

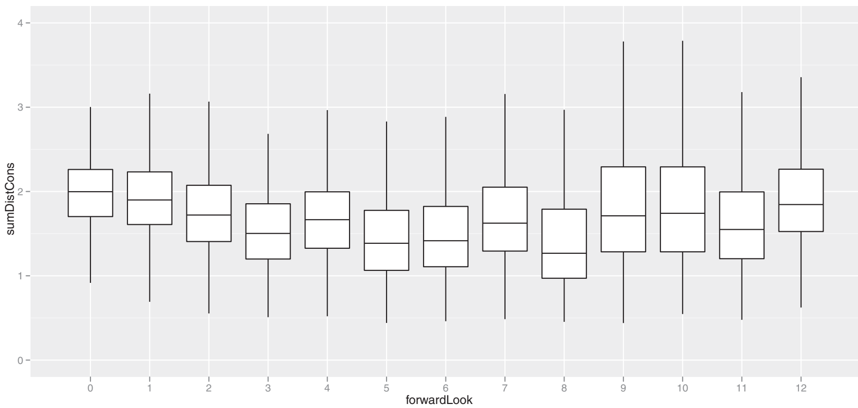


**FIGURE 11.** Learning with expectations and looking forward (distribution of sumDistCons over the corresponding experiments and all runs, for $t > T/2$).

Moreover, Figure 10 shows that they can now better converge toward the optimal consumption strategy $\theta^*$, even if, again, discovering $\gamma^*$ is more difficult for them.

The role of forward looking can also be analyzed from the same point of view. First, Figure 11 shows that, even with myopic forward looking ($forwardLook = 0$), the total distance to optimal consumption is significantly lower than the one observed with the previous learning schemes. Second, giving the consumer the ability to look forward over several periods ($forwardLook > 0$) enhances the convergence process. We indeed observe in the graphic an intermediate zone where the distance is minimal, but, from $forwardLook = 8$ on, it begins to increase again. With a long horizon, the agent uses a more out-of-date mental model to form his or her expectations. As a consequence, looking very far is not necessarily preferable with this adaptive behavior. Figure 12 confirms these results in terms of utility sacrifice.

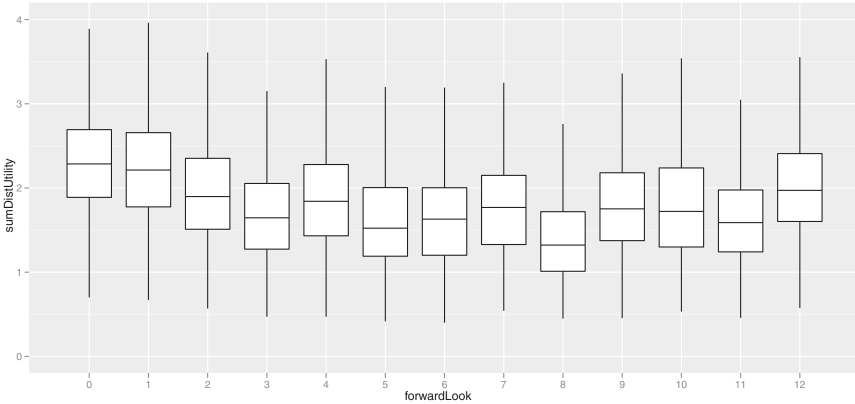


**FIGURE 12.** Learning with expectations and looking forward: utility sacrifice (distribution of sumDistUtility over the corresponding experiments and all runs, for $t > T/2$).

## 6. CONCLUSION

In this article, we developed a computational agent-based model (ABM) to reassess the case for learning regarding the linear buffer-stock rule of Allen and Carroll (2001), by considering alternative assumptions about the learning process of consumers. By doing so, we try to investigate which features of learning may be key in this context for pushing the consumption behavior close to the optimal solution obtained in a rational expectations intertemporal setting.

In this ABM, we consider three learning mechanisms: purely adaptive learning based on random experiment and combinations of already discovered consumption strategies; social learning based on imitation of strategies between consumers; adaptive learning guided by adaptive expectations. The first two mechanisms are modeled using a framework similar to genetic algorithms. The last mechanism combines this kind of learning with adaptive expectations formed by the agents on the basis of their *mental models* of the economy. These mental models are represented as a personal artificial neural network used by each consumer to build a representation of the economy from her experience in this economy. We show that only the last approach yields economically sound consumption behavior. Consumers develop consumption behavior preserved from unrealistic erratic fluctuations (a common shortcoming of purely adaptive learning schemes), while attaining performances that increase in time. This corresponds to the emergence of effective learning on their side. The ability to look forward helps them in this process and an intermediate expectation horizon yields the best results. We should nevertheless notice that such performance is only obtained after 1,000 periods. Even if the total number of experiments used in our case is smaller than the one adopted by Allen and Carroll, it remains quite significant. In-line learning by an ANN appears to be quite demanding in terms of experimentation by the consumer. However, our results show that the use of a mental model to represent the forward-looking dimension of agents' learning is probably the correct way of modeling adaptive behavior, even if we need yet to invent more frugal, and hence more realistic, ways to represent the mental model and its adaptation.

Overall, these results look promising in the perspective of building macroeconomic models based on adaptive learning dynamics. Agent-based modeling would be a natural framework for such investigations, as it would make possible understanding the aggregate outcomes resulting from coordination problems between agents endowed with bounded rationality. For example, the authors are developing an ABM inspired by the canonical NK model, in order to analyze the effects of different monetary rules à la Taylor with learning agents.

## NOTES

1. See Deaton (1991, 1992) for an overview of the state of the art at the beginning of the nineties.
2. See Carroll (1997) for a thorough examination of those properties and Carroll (2001) for a didactic presentation and comparative analysis.
3. See Carroll (1997) for a detailed documentation of those puzzles.
4. See Yildizoglu (2001) for an example of this approach in industrial economics.
5. Carroll is, however, skeptical about the added value of considering social learning with respect to the problem at hand: "even the social learning model will probably take considerable time to converge on optimal behavior, so this model provides no reason to suppose that consumers will react optimally in the short or medium run to the introduction of new elements into their environment" [Carroll (2001, p. 42)].
6. This equivalence is only valid under the *impatience* condition $R\beta^{1/\rho} < G$, with $G$ the income growth factor. In this case we assume $G = 1$, and the condition is satisfied.
7. This target level is a key element of the buffer-stock savings model of Carroll. The proof of its existence is set up in Carroll (1997).
8. By construction, $\gamma^* \equiv f'(0)$.
9. See http://yildizoglu.info/essid/learnapplet/index.html for a Java applet that can be used to simulate firm behavior resulting from different learning mechanisms.
10. See for example Goupy and Creighton (2007) for a pedagogical statement. This method is widely used in areas such as industry, chemistry, computer science, and biology. To our knowledge, Happe (2005) and Oeffner (2008) are the only applications to an economic-agent-based model.
11. They also respect the noncollapsing criterion, which ensures that each point is uniquely tested.

## REFERENCES

Allen, Todd W. and Christopher D. Carroll (2001) Individual learning about consumption. *Macroeconomic Dynamics* 5(2), 255–271.

Arifovic, J. (1994) Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control* 18, 3–28.

Butler, N.A. (2001) Optimal and orthogonal latin hypercube designs for computer experiments. *Biometrika* 88(3), 847–857.

Carroll, Christopher D. (1992) The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity* (2), 61–156.

Carroll, Christopher D. (1997) Buffer stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112(1), 1–56.

Carroll, Christopher D. (2001) A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives* 15(3), 23–46.

Carroll, Christopher D. (2004) Theoretical Foundations of Buffer Stock Saving. Working paper 10867 NBER.

Cioppa, T.M. (2002) *Efficient Nearly Orthogonal And Space-Filling Experimental Designs For High-Dimensional Complex Models*. Ph.D. dissertation, in philosophy in operations research, Naval Postgraduate School.

Cioppa, Thomas M. and Thomas W. Lucas (2007) Efficient nearly orthogonal and space-filling latin hypercubes. *Technometrics* 49(1), 45–55.

Deaton, Angus (1991) Saving and liquidity constraints. *Econometrica* 59(5), 1221–1248.

Deaton, Angus (1992) *Understanding Consumption*. New York: Oxford University Press.

Evans, G.W. and S. Honkapohja (2001) *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.

Fang, K.T., D.K.J. Lin, P. Winker, and Y. Zhang (2000) Uniform design: Theory and application. *Technometrics* 42(3), 237–248.

Goupy, J. and L. Creighton (2007) *Introduction to Design of Experiments with JMP Examples*, 3rd ed. Cary, NC: SAS Institute.

Happe, K. (2005) Agent-Based Modelling and Sensitivity Analysis by Experimental Design and Metamodelling: An Application to Modelling Regional Structural Change. Paper prepared for the XIth International Congress of the European Association of Agricultural Economists.

Holland, John H. (1992) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Cambridge, MA: MIT Press.

Holland, J. and J. H. Miller (1991) Artificial adaptive agents in economic theory. *American Economic Review Papers and Procedings* 81(2), 363–370.

Holland, J.H., K.J. Holyoak, R.E. Nisbett, and P.R. Thagard (1989) *Induction. Processes of Inference, Learning and Discoverey*. Cambridge, MA: MIT Press.

Howitt, Peter and Ömer Özak (2009) Adaptive Consumption Behavior. Working paper 15427, NBER.

Kleijnen, Jack P.C., Susan M. Sanchez, Thomas W. Lucas, and Thomas M. Cioppa (2005) A user's guide to the brave new world of designing simulation experiments. *INFORMS Journal on Computing* 17(3), 263–289.

Lettau, Martin and Harald Uhlig (1999) Rules of thumb versus dynamic programming. *American Economic Review* 89(1), 148–174.

Masters, Timothy (1993) *Practical Neural Network Recipes in C++*. New York: Academic Press.

Oeffner, M. (2008) *Agent-Based Keynesian Macroeconomics—An Evolutionary Model Embedded in an Agent-Based Computer Simulation*. Doctoral dissertation, Bayerische Julius—Maximilians Universitat.

R Development Core Team (2003) *R: A Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing. Available at http://www.r-project.org.

Salmon, Mark (1995) Bounded rationality and learning: Procedural learning. In Alan Kirman and Mark Salmon (eds.), *Learning and Rationality in Economics*, pp. 236–275. Oxford, UK: Blackwell.

Sanchez, S.M. (2005) Nolhdesigns Spreadsheet. Available at http://diana.cs.nps.navy.mil/SeedLab/.

Simon, Herbert A. (1976) From substantial to procedural rationality. In S.J. Latsis (ed.), *Method and Appraisal in Economics*, pp. 129–148. Cambridge, UK: Cambridge University Press.

Sutton, Richard S. and Andrew G. Barto (1998) *Reinforcement Learning: an Introduction*. Cambridge, MA: MIT Press.

Vallée, Thomas and Murat Yıldızoğlu (2009) Convergence in the finite cournot oligopoly with social and individual learning. *Journal of Economic Behaviour and Organization* (72), 670–690.

Vriend, Nicolaas (2000) An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control* 24(1), 1–19.

Wickham, Hardley (2009) *ggplot2 Elegant Graphics for Data Analysis*. Dordrecht: Springer.

Wilson, Stewart W. (1995) Classifier fitness based on accuracy. *Evolutionary Computation* 3(2), 149–175.

Ye, K.Q. (1998) Orthogonal column latin hypercubes and their application in computer experiments. *Journal of the American Statistical Association* 93(444), 1430–1439.

Yildizoglu, Murat (2001) Connecting adaptive behaviour and expectations in models of innovation: The potential role of artificial neural networks. *European Journal of Economics and Social Systems* 15(3), 203–220.

Yildizoglu, Murat (2002) Competing R&D strategies in an evolutionary industry model. *Computational Economics* 19(1), 52–65.

# APPENDIX: MODEL PARAMETERS AND SIMULATION EXPERIMENTS

Table A.1 gives the values of the parameters explored in the simulations. These values have been generated using Sanchez (2005). For other parameters, we have adopted the following assumptions:

$n = 20$: number of consumers;
$m = 40$: number of elements in the strategy population of each agent;
$T = 1250$ : number of simulation periods in each run;
$\beta = 0.95$;
$\rho = 3$;
$windowSize = 150$: the training of the ANN uses observations from the last 150 periods;
$u\,(c \leq 0.01) \equiv -5000$: truncation of utility computation, in order to avoid buffer overflow problems resulting from the utility function adopted by Allen and Carroll (2001).

https://doi.org/10.1017/S1365100512000582 Published online by Cambridge University Press

**TABLE A.1.** Experiments

| | Parameter | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | *initialWealth* | *probCrossOver* | *probMutate* | *probImitate* | *forwardLook* | *gaRate* | *numeEpoch* | *learnRate* | *nbHidden* |
| Min | 0 | 0.05 | 0.05 | 0.05 | 0 | 1 | 20 | 0.01 | 2 |
| Max | 3 | 0.4 | 0.4 | 0.3 | 12 | 10 | 50 | 0.1 | 6 |
| Experiment | | | | | | | | | |
| 0 | 3 | 0.08 | 0.2 | 0.1 | 11 | 7 | 41 | 0.05 | 6 |
| 1 | 3 | 0.4 | 0.09 | 0.14 | 6 | 3 | 43 | 0.04 | 6 |
| 2 | 3 | 0.2 | 0.37 | 0.09 | 0 | 6 | 42 | 0.01 | 3 |
| 3 | 2 | 0.36 | 0.4 | 0.15 | 11 | 2 | 44 | 0.02 | 4 |
| 4 | 3 | 0.06 | 0.21 | 0.1 | 8 | 7 | 32 | 0.06 | 2 |
| 5 | 3 | 0.38 | 0.16 | 0.12 | 5 | 3 | 25 | 0.09 | 2 |
| 6 | 2 | 0.21 | 0.39 | 0.11 | 0 | 7 | 31 | 0.09 | 6 |
| 7 | 2 | 0.29 | 0.38 | 0.14 | 11 | 3 | 27 | 0.1 | 4 |
| 8 | 2 | 0.14 | 0.13 | 0.18 | 9 | 4 | 20 | 0.03 | 4 |
| 9 | 2 | 0.28 | 0.15 | 0.22 | 3 | 6 | 23 | 0.04 | 5 |
| 10 | 2 | 0.13 | 0.31 | 0.29 | 4 | 2 | 24 | 0.02 | 4 |
| 11 | 2 | 0.3 | 0.28 | 0.28 | 9 | 10 | 34 | 0.05 | 3 |
| 12 | 2 | 0.1 | 0.12 | 0.19 | 7 | 2 | 49 | 0.08 | 3 |
| 13 | 3 | 0.26 | 0.18 | 0.27 | 2 | 6 | 48 | 0.07 | 3 |
| 14 | 2 | 0.12 | 0.35 | 0.28 | 5 | 1 | 40 | 0.08 | 5 |
| 15 | 2 | 0.27 | 0.26 | 0.3 | 10 | 9 | 37 | 0.07 | 5 |
| 16 | 2 | 0.23 | 0.23 | 0.18 | 6 | 6 | 35 | 0.06 | 4 |
| 17 | 0 | 0.37 | 0.25 | 0.25 | 2 | 4 | 29 | 0.06 | 2 |
| 18 | 0 | 0.05 | 0.36 | 0.21 | 6 | 8 | 28 | 0.07 | 2 |
| 19 | 0 | 0.25 | 0.08 | 0.26 | 12 | 5 | 28 | 0.1 | 5 |
| 20 | 1 | 0.09 | 0.05 | 0.2 | 1 | 9 | 26 | 0.09 | 4 |

**Table A.1.** (Continued.)

| | Parameter | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | *initialWealth* | *probCrossOver* | *probMutate* | *probImitate* | *forwardLook* | *gaRate* | *numeEpoch* | *learnRate* | *nbHidden* |
| 21 | 0 | 0.39 | 0.24 | 0.25 | 4 | 4 | 38 | 0.05 | 6 |
| 22 | 0 | 0.07 | 0.29 | 0.23 | 7 | 8 | 45 | 0.02 | 6 |
| 23 | 1 | 0.24 | 0.06 | 0.24 | 12 | 4 | 39 | 0.02 | 3 |
| 24 | 1 | 0.16 | 0.07 | 0.21 | 1 | 8 | 43 | 0.01 | 4 |
| 25 | 1 | 0.31 | 0.32 | 0.17 | 3 | 7 | 50 | 0.08 | 4 |
| 26 | 1 | 0.17 | 0.3 | 0.13 | 9 | 5 | 47 | 0.07 | 3 |
| 27 | 1 | 0.32 | 0.14 | 0.06 | 8 | 9 | 46 | 0.09 | 5 |
| 28 | 1 | 0.15 | 0.17 | 0.07 | 3 | 1 | 36 | 0.06 | 5 |
| 29 | 1 | 0.35 | 0.33 | 0.16 | 5 | 9 | 21 | 0.03 | 5 |
| 30 | 0 | 0.19 | 0.27 | 0.08 | 10 | 5 | 22 | 0.04 | 5 |
| 31 | 1 | 0.33 | 0.1 | 0.07 | 8 | 10 | 30 | 0.03 | 3 |
| 32 | 1 | 0.18 | 0.19 | 0.05 | 2 | 2 | 33 | 0.04 | 3 |