Journal of Economic Dynamics & Control 39 (2014) 237–254

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Optimal consumption under uncertainty, liquidity constraints, and bounded rationality

Ömer Özak

Department of Economics, Southern Methodist University, PO Box 0496, Dallas, TX 75275-0496, USA

ARTICLE INFO

Article history:
Received 5 February 2013
Received in revised form
5 December 2013
Accepted 9 December 2013
Available online 27 December 2013

JEL classification:
C6
D8
D9
E21

Keywords:
Adaptive learning models
Bounded rationality
Dynamic programming
Consumption function
Behavioral economics
Saving behavior

ABSTRACT

I study how boundedly rational agents can learn a "good" solution to an infinite horizon optimal consumption problem under uncertainty and liquidity constraints. Using an empirically plausible theory of learning I propose a class of adaptive learning algorithms that agents might use to choose a consumption rule. I show that the algorithm always has a globally asymptotically stable consumption rule, which is optimal. Additionally, I present extensions of the model to finite horizon settings, where agents have finite lives and life-cycle income patterns. This provides a simple and parsimonious model of consumption for large agent based models.

1. Introduction

Rationality is one of the main tenets of modern economics and though it has proven fruitful in all areas of economics, it has recently been subject to attacks both on theoretical and empirical grounds. In particular, the modern theory of consumption under liquidity constraints and uncertainty, which is one of the main building blocks of modern macroeconomics, has been criticized for its rationality requirements. For example, Carroll (2001) presents this theory and argues that "when there is uncertainty about the future level of labor income, it appears to be impossible under plausible assumptions about the utility function to derive an explicit solution for consumption as a direct (analytical) function of the model′s parameters". Similarly, Allen and Carroll (2001) admit that "finding the exact nonlinear consumption policy rule (as economists have done) is an extraordinarily difficult mathematical problem". In order to answer this line of critiques, economists have tried to provide bounded rationality foundations to optimal behavior, especially within game theory (Fudenberg and Levine, 1998) and macroeconomics (Evans and Honkapohja, 2001; Sargent, 1993). Still, the study of how agents learn the optimal policy to an infinite horizon dynamic programming problem under uncertainty, and the consumption function in particular, has largely been

E-mail address: ozak@smu.edu
URL: http://www.omerozak.com

0165-1889/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jedc.2013.12.007

ignored aside from a few exceptions (Allen and Carroll, 2001; Evans et al., 2009; Howitt and Özak, in press; Lettau and Uhlig, 1999).[1] While theoretical results have been mixed, empirical and experimental evidence suggests that agents *do learn to behave as if* they had solved the optimal consumption problem (see Brown et al., 2009, and references therein).

In this paper I study an infinite horizon optimal consumption problem under uncertainty, liquidity constraints, and bounded rationality. I follow the previous literature in assuming that boundedly rational agents use a consumption rule that is linear in wealth.[2] I endow agents with a learning algorithm, which I call the HO-algorithm, that is a generalization of the one studied in the numerical exercise of Howitt and Özak (in press). In particular, I assume that agents change their behavior in response to differences between their experienced marginal utility in one period and next period′s discounted marginal utility. Thus, agents react to ex-post mistakes in their Euler equation, and adjust their consumption rule if it failed to equalize the marginal utilities of consumption between yesterday and today.

There are various reasons why the HO-algorithm seems like a good candidate for behavior under bounded rationality. First, it does not require complex optimizing behavior by agents, which is a fundamental requirement of any theory of bounded rationality (Selten, 2001). Agents in this theory are only required to compare two marginal utilities in order to make changes. In this aspect it is similar to aspiration adaptation theory (Selten, 1998), reinforcement learning (Börgers and Sarin, 1997), and learning direction theory (Selten and Stoecker, 1986). This simplicity lowers the cognitive capabilities required of agents. Second, it has low informational and computational requirements, which are *independent* of agents' economic environment, including the number of states or the set of possible consumption rules. This is extremely important, since these factors put a heavy burden on agents' cognitive abilities. In particular, agents only need to remember a small amount of information and know basic algebraic operations. Third, in order to apply the algorithm, agents do neither need to fully understand the economic environment in which they are embedded nor the effects of changes in it. Thus, it gives them guidance even in unfamiliar situations. Fourth, unlike some models of bounded rationality, which are qualitative in nature (e.g., learning direction theory), the quantitative nature of the HO-algorithm facilitates its use in applied macroeconomic models. Fifth, its similarity to learning direction and regret theory gives it empirical relevance. Finally, as I will show below, various versions of the algorithm can easily accommodate different levels of rationality and social interaction.

My analysis shows that the behavior of agents using this learning scheme converges to a unique consumption rule that has good welfare and stability properties. I prove that for a general class of consumption problems under uncertainty and liquidity constraints, agents using variants of the HO-algorithm learn an "optimal" consumption rule. More specifically, I show that the HO-algorithm converges to the unique asymptotically stable point of a particular ordinary differential equation (ODE). This implies that applying the HO-algorithm to *any* initial linear consumption rule in an uncountable and compact set, for different assumptions about an agent′s rationality, her level of risk-aversion or impatience, or her income process, will with probability one converge to the globally asymptotically stable stationary point of this ODE. Additionally, I show that this stationary consumption rule maximizes her steady state expected life-time utility under the stationary distribution generated by this consumption rule. Thus, in a steady state agents would not have an incentive to choose a different linear consumption function, since their welfare under the current rule would be maximal. I also provide bounds for the rate of convergence to the optimal consumption rule. I show how this rate of convergence is affected by assumptions about agents' rationality and their social interactions. Although my main results are based on an infinite horizon analysis, I extend my analysis to cover the behavior of the HO-algorithm in finite time including the case of life-cycle consumers. The results suggest that the algorithm keeps its properties in these settings.

My approach differs from the one used by Lettau and Uhlig (1999) and Allen and Carroll (2001), who use the accumulated performance of a rule as measured by the discounted sum of utilities as a base for their learning mechanisms. In these papers, agents estimate the value function of their respective problem in order to select the best rule. The approach in these two papers has four main drawbacks. First, they require the set of rules and states to be finite. Second, the memory, processing, and rationality requirements increase in the number of rules and states. Third, they cannot determine the welfare properties of the rules that are learnt, especially when the rational rule is not available or if the rational rule is not equivalent to a mix of the available rules. And fourth, the algorithms put forward in these papers do not converge to the optimal rule or only converge very slowly. Thus they are "not an adequate description of the process by which consumers learn about consumer behavior" (Allen and Carroll, 2001, p. 268).

This paper is closely related to Evans et al. (2009), who use Euler equation and shadow price learning schemes. In their approach, agents are forward looking and forecast either shadow prices or the control variable and then choose the control variable optimally according to the first order conditions in their problem. In particular, they show that agents using Euler equation or shadow price learning can learn the optimal solution to a linear quadratic dynamic programming under the

[1] There is a large literature which studies dynamic programming problems in which agents do not hold Rational Expectations, but are otherwise fully rational. The objective of this literature is to understand the conditions under which the expectational mechanism held by agents converges to Rational Expectations (Sargent, 1993). This is not the problem I am alluding to here. In this setting agents are not able or willing to solve the optimal consumption problem, even if they had the correct expectational mechanism.

A related literature studies the problem of convergence of the computational methods applied to solve numerically dynamic programming problems, for example Puterman and Brumelle (1979) and Santos and Rust (2004). Although I am not studying this problem, one could apply the methods of this paper to find the optimal partition of the state space or to approximate the optimal solution.

[2] Gabaix (2011) has suggested that boundedly rational agents only use "sparse" rules of behavior. In this paper, the assumption is that agents focus only on wealth and disregard all other variables. As can be seen from the results and proofs below, they can be extended straightforwardly to include linearity in other variables, without affecting the results.