Journal of Economic Dynamics & Control 45 (2014) 353–365

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

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

Solving the income fluctuation problem with unbounded rewards

Huiyu Li [a,b,*], John Stachurski [a,b]

[a] Department of Economics, Stanford University, United States
[b] Research School of Economics, Australian National University, Australia

ARTICLE INFO

Article history:
Received 18 February 2013
Received in revised form
13 November 2013
Accepted 5 June 2014
Available online 12 June 2014

JEL classification:
C63
E21

Keywords:
Coleman operator
Policy iteration
Time iteration
Global convergence

ABSTRACT

This paper studies the income fluctuation problem without imposing bounds on utility, assets, income or consumption. We prove that the Coleman operator is a contraction mapping over the natural class of candidate consumption policies when endowed with a metric that evaluates consumption differences in terms of marginal utility. We show that this metric is complete, and that the fixed point of the operator coincides with the unique optimal policy. As a consequence, even in this unbounded setting, policy function iteration always converges to the optimal policy at a geometric rate.

1. Introduction

The income fluctuation problem refers to a classic decision problem that lies at the heart of modern macroeconomic theory. In the problem, agents choose a state-contingent path for savings and consumption in order to maximize expected lifetime utility, taking as given the rate of return on assets, an exogenous stream of non-capital income, and, in many cases, a borrowing constraint. The model has been used to analyze household behavior in many fundamental economic and financial applications. The literature is too large to enumerate, but some broadly representative examples include Schechtman and Escudero (1977), Deaton (1991), Huggett (1993), Aiyagari (1994), Krusell and Smith (1998), Deaton and Laroque (1992, 1996) and Angeletos (2007).

Early work on consumption behavior focused on highly simplified problems with closed-form solutions. It turned out that these models have only limited ability to fit consumption data (see, e.g., Carroll, 2001). Adding more realistic features has led to better models, but in these settings computation cannot be avoided. The computational problem remains a nontrivial one because in most modeling exercises the consumer problem is embedded in a larger equilibrium or estimation problem, and needs to be solved quickly, accurately and reliably for many different parameter values.

A variety of solution techniques have been proposed for the income fluctuation problem specifically or for optimization problems that subsume the income fluctuation problem. The literature now contains many numerical studies presenting

* Corresponding author.
E-mail addresses: huiyuli@stanford.edu (H. Li), john.stachurski@anu.edu.au (J. Stachurski).

http://dx.doi.org/10.1016/j.jedc.2014.06.003

simulation results for particular solution methods according to particular criteria in particular applications and at particular sets of parameter values. While such studies can certainly complement theoretical analysis, they cannot substitute for it, and there remains a lack of clear analytical results proving convergence at given rates for a given method over a continuum of standard applications, parameter values and initial conditions.

In this paper, we provide analytical results on convergence for the common solution method known as policy function iteration (or, in some circles, time iteration). The basic ideas behind policy function iteration were illuminated by Coleman (1990). As is well-known, when the utility function is bounded, policy iteration is globally convergent. The reason is that the operator that implements policy iteration—the Coleman operator—is essentially conjugate to the Bellman operator (Coleman, 1990). When rewards are bounded, global geometric stability of the Bellman operator is guaranteed by classical dynamic programming theory. By applying this conjugacy between the two operators, one can then show that the Coleman operator has all of the same properties. Rendahl (2013) makes use of these ideas to provide a detailed treatment of policy iteration in the bounded reward case, working with an abstract optimization problem that permits occasionally binding constraints.

For standard income fluctuation models, however, utility is unbounded, and consumption can become either arbitrarily small or arbitrarily large. In these settings, the Bellman operator is not a contraction mapping in the usual metric, and we cannot claim that iterates of the Bellman operator converge uniformly to the value function. In fact the uniform deviation is typically infinite, regardless of how many iterations are performed. Thus the standard dynamic programming theory does not apply.

In response to these issues, the present paper develops an alternative approach to the income fluctuation optimization problem that delivers sharper results than previously obtained—even when rewards are bounded. Our focus is directly on the Coleman operator, rather than drawing connections to the Bellman operator. Because we work with the Coleman operator and policy function iteration, our main results are formulated in policy function space rather than value function space, and unbounded rewards cause no difficulties for the analysis.

As our most significant theoretical result, we show that a version of the Coleman operator adapted to the income fluctuation problem is in fact a contraction mapping in a complete metric space of candidate consumption policies, even when rewards are unbounded. We also prove that the asset-consumption path associated with the fixed point of Coleman's operator satisfies the sequential Euler equation and transversality conditions, and that the Euler equation and transversality conditions are sufficient for optimality. Putting these facts together, we show that a unique optimal consumption policy exists, and that, for any well-behaved initial condition, policy function iteration converges to this optimal policy at a geometric rate. In particular, we prove that the pointwise deviation between the n-th iterate and the optimal policy converges to zero at a geometric rate, and the same is true for the uniform deviation over any bounded set. (As will be discussed later, this is in a sense the best possible result for policy function iteration in the unbounded setting.) Moreover, we give a computable upper bound on the deviation in terms of observable quantities.

All of these results are obtained in a setting that can accommodate a broad range of standard applications. In particular, no specific structure is imposed on utility beyond differentiability, concavity and the usual slope conditions. Utility can be unbounded both above and below. In addition, non-capital income and the asset space are allowed to be unbounded. The income process is permitted to be nonstationary, as is required in certain applications.[1]

In terms of connections to the existing literature, perhaps the most closely related results are those found in a recent paper on heterogeneous agent incomplete market economies by Kuhn (2013). Like us, Kuhn permits unbounded rewards and unbounded asset and shock spaces. As one component of his investigation into decentralized equilibria, he studies the same consumer problem considered in this paper. By applying an order-theoretic approach to the analysis of the Coleman operator, he establishes existence of a fixed point, which corresponds to an optimal consumption policy, and provides some convergence results for policy function iteration. On one hand, the present paper is much narrower than Kuhn's paper, in the sense that we concern ourselves only with the consumer's problem. On the other hand, our results on the consumer problem's are considerably sharper. We obtain not only the existence of a fixed point but also uniqueness, as well as geometric rates of convergence of policy function iteration.

Regarding earlier literature, the Coleman operator was originally introduced as a constructive iterative method for solving stochastic optimal growth models (Coleman, 1990). It has often been used to establish existence of equilibrium in economies with distortions, notably by Coleman (1991), Greenwood and Huffman (1995), Datta et al. (2002), Morand and Reffett (2003), Datta et al. (2005) and Mirman et al. (2008). In these papers, fixed points of the Coleman operator were analyzed using a variety of methods related to order preserving structures, continuity, compactness and concavity. The last four papers derive fixed point results in very general settings, but always with either bounded utility, compact state spaces or both.

There are other approaches to the optimization problem treated in this paper besides analysis of the Coleman operator, even in the unbounded setting. One such alternative is value function iteration paired with weighted supremum norms rather that standard supremum norms. While the weighted supremum norm strategy is well suited to convergence analysis,

[1] Predictions of this class of problems can be highly sensitive to the persistence and stationarity of the shock process—hence the need to include the possibility of nonstationary income dynamics. Recent papers addressing this point include Kaplan and Violante (2010), Blundell et al. (2008), Moll (2012) and Kuhn (2013).