J. Math. Anal. Appl. • • • (• • • •) • • •–• • •

Contents lists available at ScienceDirect

Journal of Mathematical Analysis and Applications

www.elsevier.com/locate/jmaa

Discounted approximations to the risk-sensitive average cost in finite Markov chains ☆

Rolando Cavazos-Cadena [a,*], Daniel Cruz-Suárez [b]

[a] *Departamento de Estadística y Cálculo, Universidad Autónoma Agraria Antonio Narro, Boulevard Antonio Narro 1923, Buenavista, Saltillo, COAH 25315, Mexico*
[b] *División Académica de Ciencias Básicas, Universidad Juárez Autónoma de Tabasco, A.P. 24, Cunduacán, TAB 86690, Mexico*

ARTICLE INFO

Article history:
Received 22 September 2016
Available online xxxx
Submitted by A. Sulem

Keywords:
Exponential utility
Certainty equivalent
Vanishing discount method
Hölder's inequality
Convex function

ABSTRACT

This work concerns with Markov chains on a finite state space, which is endowed with a cost function. The evolution of the chain is observed by an agent with constant risk-sensitivity and, assuming that the state space is a communicating class, the relation between the risk-sensitive discounted and average performance criteria is studied. It is proved that, as the discount factor increases to 1, an appropriate normalization of the discounted value function converges to the average cost. Also, it is shown that if the classical normalization used in the risk-neutral case is applied in the risk-sensitive context, then the normalized discounted value function converges to an arithmetic mean of the average cost.

1. Introduction

This work concerns with (time-homogeneous) Markov cost chains (MCCs) on a finite state space. A state-dependent cost is incurred each time that a state is visited and, using the stream of costs, the overall performance of the system is measured by an agent with constant risk-sensitivity. Two performance criteria are considered, namely, the (long-run) average index, and the total discounted cost. The basic structural requirement on the model, stated as Assumption 2.1 below, is that the state space is a communicating class, a condition ensuring that the average cost does not depend on the initial state and is characterized via a single Poisson equation. In this context, the following problem is addressed:

- To obtain convergent approximations for the risk-sensitive average cost in terms of the family of risk-sensitive discounted value functions.

☆ This work was supported by the PSF Organization under Grant No. 0116-18.
* Corresponding author.
E-mail addresses: rcavazos@uaaan.mx (R. Cavazos-Cadena), daniel.cruz@ujat.mx (D. Cruz-Suárez).

http://dx.doi.org/10.1016/j.jmaa.2017.01.084
0022-247X/

Please cite this article in press as: R. Cavazos-Cadena, D. Cruz-Suárez, Discounted approximations to the risk-sensitive average cost in finite Markov chains, J. Math. Anal. Appl. (2017), http://dx.doi.org/10.1016/j.jmaa.2017.01.084

In the risk-neutral case, corresponding to a null risk-sensitivity coefficient, the solution to this problem is well-known; see for instance, [13,1,18], or [23]. For the case of a no-null risk-sensitivity coefficient, the above problem was studied in [6] and [11], and this note is motivated by an attempt to avoid some restrictive conditions in those works. In the first paper, MCCs on a denumerable state space were studied under a stability condition that is not generally satisfied in the present framework, whereas in the second paper, which concerns with *controlled* Markov chains, the above problem was analyzed under the assumption that the risk-sensitivity coefficient is small enough; see Remark 3.1 below.

The theory of discrete-time Markov models endowed with the risk-sensitive average criterion can be traced back, at least, to the seminal paper by Howard and Matheson [14] where, for *controlled* Markov chains with finite state and action spaces, the optimal average cost was characterized in terms of an optimality equation rendering an optimal stationary policy. More recently, the risk-sensitive average index has been intensively studied. Models with a finite or denumerable state space are considered in [5,7,9], or [21], whereas MDPs on a general state space are analyzed in [10–12,15]. One reason for the interest in the risk-sensitive average criterion stems from its connection with large deviations [2,16], and mathematical finance [22,3,17]. Markov decision models with a notion of risk more general than the one used in this paper are studied, for instance, in [4] or [20].

The results on the problem posed above, which are stated in Section 3 as Theorems 3.1 and 3.2, can be described as follows:

(i) For each risk-sensitivity coefficient, it is shown that an appropriate normalization of the discounted value function converges to the average cost as the discount factor increases to 1.

Moreover, it is shown that

(ii) the convergence is uniform for risk-sensitivity coefficients in a compact set.

This last fact is used to prove that:

(iii) If the classical risk-neutral normalization is applied in the present risk-sensitive context, then the normalized discounted value function converges to the mean value of the average cost over the interval joining zero and the risk-sensitivity parameter.

As already noted, in the present context of a finite state space, these conclusions extend the results by Di Masi and Stettner [11] and Cavazos-Cadena and Fernández-Gaucherand [6]. The ideas in this last paper will be useful in the present context, and *the approach* used in this note to establish the results described above relies on the following consequence of Hölder's inequality: For any fixed discount factor and initial state, the product of the risk-sensitivity parameter and the discounted value function is a convex mapping.

The organization of the paper is as follows: In Section 2 the risk-sensitive average and discounted criteria are introduced, and their basic properties are stated, whereas in Section 3 the main conclusions of the paper are formulated in Theorems 3.1 and 3.2, and their difference with respect to other results already available is briefly discussed in Remark 3.1. Next, Section 4 contains the technical tools that are used in Sections 5 and 6 to prove the main theorems. Finally, the exposition concludes in Section 7 with some brief comments.

Notation. The set of all non-negative integers is denoted by $\mathbb{N}$. For a real valued function f, the corresponding supremum norm is given by

$$\|f\| := \sup\{|f(y)| : y \text{ belongs to the domain of } f\},$$