Average Cost Markov Decision Processes with Weakly Continuous Transition Probabilities

Eugene A. Feinberg
Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, New York 11794,
eugene.feinberg@sunysb.edu

Pavlo O. Kasyanov, Nina V. Zadoianchuk
Institute for Applied System Analysis, National Technical University of Ukraine "Kyiv Polytechnic Institute," Kyiv, Ukraine
{kasyanov@i.ua, ninellll@i.ua}

This paper presents sufficient conditions for the existence of stationary optimal policies for average cost Markov decision processes with Borel state and action sets and weakly continuous transition probabilities. The one-step cost functions may be unbounded, and the action sets may be noncompact. The main contributions of this paper are: (i) general sufficient conditions for the existence of stationary discount optimal and average cost optimal policies and descriptions of properties of value functions and sets of optimal actions, (ii) a sufficient condition for the average cost optimality of a stationary policy in the form of optimality inequalities, and (iii) approximations of average cost optimal actions by discount optimal actions.

1. Introduction. This paper provides sufficient conditions for the existence of stationary optimal policies for average cost Markov decision processes (MDPs) with Borel state and action sets and weakly continuous transition probabilities. The cost functions may be unbounded, and the action sets may be noncompact. The main contributions of this paper are: (i) general sufficient conditions for the existence of stationary discount optimal and average cost optimal policies and descriptions of properties of value functions and sets of optimal actions (Theorems 1, 3, and 4), (ii) a new sufficient condition for average cost optimality based on optimality inequalities (Theorem 2), and (iii) approximations of average cost optimal actions by discount optimal actions (Theorem 5).

For infinite-horizon MDPs, there are two major criteria: average costs per unit time and expected total discounted costs. The former is typically more difficult to analyze. The so-called vanishing discount factor approach is often used to approximate average costs per unit time by normalized expected total discounted costs. The literature on average cost MDPs is vast. Most of the earlier results are surveyed in Arapostathis et al. [1]. Here, we mention just a few references.

For finite-state and action sets, Derman [10] proved the existence of stationary average cost optimal policies. This result follows from Blackwell [6] and it also was independently proved by Viskov and Shiryaev [31]. When either the state set or action set is infinite, even ϵ-optimal policies may not exist for some $\epsilon > 0$; Ross [25], Dynkin and Yushkevich [11, Chapter 7], Feinberg [12, §5]. For a finite-state set and compact action sets, optimal policies may not exist; Bather [2], Chitashvili [9], and Dynkin and Yushkevich [11, Chapter 7].

For MDPs with finite-state and action sets, there exist stationary policies satisfying optimality equations (see Dynkin and Yushkevich [11, Chapter 7], where these equations are called canonical), and, furthermore, any stationary policy satisfying optimality equations is optimal. The latter is also true for MDPs with Borel state and action sets, if the value and weight (also called bias) functions are bounded; Dynkin and Yushkevich [11, Chapter 7]. When the optimal value of average costs per unit time does not depend on the initial state (the optimal value function is constant), the pair of optimality equations become a single equation. For bounded one-step costs, Taylor [30] and Ross [23] for a countable state space and Ross [24] and Gubenko and Shtatland [17] for a Borel state space provided sufficient conditions for the validity of optimality equations with a bounded bias function; see also Dynkin and Yushkevich [11, Chapter 7].

In many applications of infinite-state MDPs, one-step costs are unbounded from above. For example, holding costs may be unbounded in queueing and inventory systems. Sennott [27, 28] (and references therein) developed a theory for countable state problems with unbounded one-step costs. For unbounded costs, optimality inequalities are used instead of optimality equations to construct a stationary average cost optimal policy. Cavazos-Cadena [7] provided an example, when optimality inequalities hold while optimality equations do not.

Schäl [26] developed a theory for Borel state spaces and compact action sets. Two types of continuity assumptions for transition probabilities are considered in Schäl [26]: the setwise and weak continuity. For a countable

state space, these assumptions coincide; see Chen and Feinberg [8, Appendix]. Setwise convergence of probability measures is stronger than weak convergence; Hernández-Lerma and Lasserre [19, p. 186]. Formally speaking, the setwise continuity assumption for MDPs is not stronger than the weak continuity assumption, because the former claims that the transition probabilities are continuous in actions, while they are jointly continuous in states and actions in the latter. However, the joint continuity of transition probabilities in states and actions often holds in applications. For example, for inventory control problems with uncountable state spaces, setwise continuity of transition probabilities takes place if demand is a continuous random variable, while weak continuity holds for arbitrarily distributed demand; see Feinberg and Lewis [14, §4]. The importance of weak convergence for practical applications is mentioned in Hernández-Lerma and Lasserre [20, p. 141].

In many applications, action sets are not compact. Hernández-Lerma [18] extended Schäl's [26] results under the setwise continuity assumptions to possibly noncompact action sets. Schäl's [26] assumptions on compactness of action sets and lower semicontinuity of cost functions in the action argument are replaced in Hernández-Lerma [18] by a more general assumption; namely, that the cost functions are inf-compact in the action argument. For weakly continuous transition probabilities and possibly noncompact action sets, Feinberg and Lewis [14] proved the existence of stationary optimal policies for MDPs with cost functions being inf-compact in both state and action arguments when, in addition to Schäl's [26] boundedness assumption on the relative discounted value at each state, the so-called local boundedness condition was assumed.

The original goal of this study was to show that the results from Feinberg and Lewis [14] hold without the local boundedness condition. However, the results of this paper are more general. This paper provides a weaker boundedness condition on the relative discounted value (Assumption ($\underline{\text{B}}$)) in §5) than Assumption (B)) introduced in Schäl [26]. It also provides a more general and natural assumption (Assumption (W*), §3) than inf-compactness of the one-step cost function in both arguments. The main result of this paper, Theorem 3, establishes the validity of optimality inequalities and the existence of stationary optimal policies under Assumptions (W*) and ($\underline{\text{B}}$).

While inf-compactness of the cost function in the action parameter is a natural assumption, inf-compactness in the state argument is a more restrictive condition. For example, when the state space is unbounded (e.g., the set of nonnegative numbers) and action sets are compact, the assumption, that the cost function is inf-compact in both arguments, does not cover the case of bounded costs functions studied by Ross [24], Gubenko and Shtatland [17], and Dynkin and Yushkevich [11, Chapter 7]. Assumption (W*) covers this case as well as unbounded costs and noncompact action sets.

As follows from the example presented in Luque-Vásquez and Hernández-Lerma [22], MDPs with lower semicontinuous cost functions may possess pathological properties, even if the one-step cost function is inf-compact in the action variable. Assumption (W*(ii)) removes this difficulty. As stated in Lemma 1, this assumption is weaker than Schäl's [26] compactness and continuity assumptions for weakly continuous transition probabilities and than inf-compactness of one-step cost functions in both arguments (state and action) assumed in Feinberg and Lewis [14].

2. Model description. For a metric space S, let $\mathscr{B}(S)$ be a Borel σ-field on S; that is, the σ-field generated by all open sets of metric space S. For a set $E \subset S$, we denote by $\mathscr{B}(E)$ the σ-field whose elements are intersections of E with elements of $\mathscr{B}(S)$. Observe that E is a metric space with the same metric as on S, and $\mathscr{B}(E)$ is its Borel σ-field. For a metric space S, we denote by $\mathbb{P}(S)$ the set of probability measures on $(S, \mathscr{B}(S))$. A sequence of probability measures $\{\mu_n\}$ from $\mathbb{P}(S)$ converges weakly to $\mu \in \mathbb{P}(S)$ if for any bounded continuous function f on S

$$\int_S f(s)\,\mu_n(ds) \to \int_S f(s)\,\mu(ds) \quad \text{as } n \to \infty.$$

Consider a discrete-time MDP with a *state space* $\mathbb{X}$, an *action space* $\mathbb{A}$, one-step costs c, and transition pobabilities q. Assume that $\mathbb{X}$ and $\mathbb{A}$ are *Borel subsets* of Polish (complete separable metric) spaces with the corresponding metrics ρ and γ. For all $x \in \mathbb{X}$, a nonempty Borel subset $A(x)$ of $\mathbb{A}$ represents the *set of actions* available at x. Define the graph of A by

$$\text{Gr}(A) = \{(x, a): x \in \mathbb{X}, a \in A(x)\}.$$

Assume also that
  (i) $\text{Gr}(A)$ is a measurable subset of $\mathbb{X} \times \mathbb{A}$; that is, $\text{Gr}(A) \in \mathscr{B}(\mathbb{X} \times \mathbb{A})$, where $\mathscr{B}(\mathbb{X} \times \mathbb{A}) = \mathscr{B}(\mathbb{X}) \otimes \mathscr{B}(\mathbb{A})$;
  (ii) there exists a measurable mapping $\phi: \mathbb{X} \to \mathbb{A}$ such that $\phi(x) \in A(x)$ for all $x \in \mathbb{X}$.

The *one-step cost*, $c(x, a) \le +\infty$, for choosing an action $a \in A(x)$ in a state $x \in \mathbb{X}$, is a *bounded below measurable* function on $\mathrm{Gr}(A)$. Let $q(B \mid x, a)$ be the *transition kernel* representing the probability that the next state is in $B \in \mathcal{B}(\mathbb{X})$, given that the action a is chosen in the state x. This means that

   (i) $q(\cdot \mid x, a)$ is a probability measure on $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$ for all $(x, a) \in \mathbb{X} \times \mathbb{A}$;

   (ii) $q(B \mid \cdot, \cdot)$ is a Borel function on $(\mathrm{Gr}(A), \mathcal{B}(\mathrm{Gr}(A)))$ for all $B \in \mathcal{B}(\mathbb{X})$.

   The decision process proceeds as follows:

   (i) at each time epoch $n = 0, 1, \ldots$, the current state $x \in \mathbb{X}$ is observed;

   (ii) a decision maker chooses an action $a \in A(x)$;

   (iii) the cost $c(x, a)$ is incurred;

   (iv) the system moves to the next state according to the probability law $q(\cdot \mid x, a)$.

As explained in the text following the proof of Lemma 2, if for each $x \in \mathbb{X}$ there exists $a \in A(x)$ with $c(x, a) < \infty$, the measurability of $\mathrm{Gr}(A)$ and inf-compactness of the cost function c in the action variable a assumed later imply that assumption (ii) holds.

Let $\mathbb{H}_n = (\mathbb{X} \times \mathbb{A})^n \times \mathbb{X}$ be the *set of histories* up to epochs $n = 0, 1, \ldots$ and $\mathcal{B}(\mathbb{H}_n) = (\mathcal{B}(\mathbb{X}) \otimes \mathcal{B}(\mathbb{A}))^n \otimes \mathcal{B}(\mathbb{X})$. A *randomized decision rule* at epoch $n = 0, 1, \ldots$ is a regular transition probability $\pi_n \colon H_n \to \mathbb{A}$ concentrated on $A(\xi_n)$; that is, (i) $\pi_n(\cdot \mid h_n)$ is a probability on $(\mathbb{A}, \mathcal{B}(\mathbb{A}))$, given the history $h_n = (\xi_0, u_0, \xi_1, u_1, \ldots, u_{n-1}, \xi_n) \in \mathbb{H}_n$ satisfying $\pi_n(A(\xi_n) \mid h_n) = 1$, and (ii) for all $B \in \mathcal{B}(\mathbb{A})$, the function $\pi_n(B \mid \cdot)$ is Borel on $(\mathbb{H}_n, \mathcal{B}(\mathbb{H}_n))$. A *policy* is a sequence $\pi = \{\pi_n\}_{n=0,1,\ldots}$ of decision rules. Moreover, π is called *nonrandomized* if each probability measure $\pi_n(\cdot \mid h_n)$ is concentrated at one point. A nonrandomized policy is called *Markov* if all of the decisions depend on the current state and time only. A Markov policy is called *stationary* if all the decisions depend on the current state only. Thus, a Markov policy ϕ is defined by a sequence $\phi_0, \phi_1, \ldots$ of Borel mappings $\phi_n \colon \mathbb{X} \to \mathbb{A}$ such that $\phi_n(x) \in A(x)$ for all $x \in \mathbb{X}$. A stationary policy ϕ is defined by a Borel mapping $\phi \colon \mathbb{X} \to \mathbb{A}$ such that $\phi(x) \in A(x)$ for all $x \in \mathbb{X}$. Let

$$\mathbb{F} = \{\phi \colon \mathbb{X} \to \mathbb{A} \colon \phi \text{ is Borel and } \phi(x) \in A(x) \text{ for all } x \in \mathbb{X}\}$$

be the *set of stationary policies*.

The Ionescu-Tulcea theorem (Bertsekas and Shreve [4, pp. 140–141] or Hernández-Lerma and Lasserre [19, p. 178]) implies that an initial state x and a policy π define a unique probability P_x^π on the set of all trajectories $\mathbb{H}_\infty = (\mathbb{X} \times \mathbb{A})^\infty$ endowed with the product of σ-field defined by Borel σ-field of $\mathbb{X}$ and $\mathbb{A}$. Let $\mathbb{E}_x^\pi$ be an expectation regarding P_x^π.

For a finite-horizon $N = 0, 1, \ldots$, let us define the *expected total discounted costs*

$$v_{N,\alpha}^\pi := \mathbb{E}_x^\pi \sum_{n=0}^{N-1} \alpha^n c(\xi_n, u_n), \quad x \in \mathbb{X}, \tag{1}$$

where $\alpha \ge 0$ is the discount factor and $v_{0,\alpha}^\pi(x) = 0$. When $N = \infty$ and $\alpha \in [0, 1)$, (1) defines an *infinite-horizon expected total discounted cost* denoted by $v_\alpha^\pi(x)$.

The *average cost per unit time* is defined as

$$w^\pi(x) := \limsup_{N \to +\infty} \frac{1}{N} v_{N,1}^\pi(x), \quad x \in \mathbb{X}. \tag{2}$$

For any function $g^\pi(x)$, including $g^\pi(x) = v_{N,\alpha}^\pi(x)$, $g^\pi(x) = v_\alpha^\pi(x)$, and $g^\pi(x) = w^\pi(x)$, define the *optimal cost*

$$g(x) := \inf_{\pi \in \Pi} g^\pi(x), \quad x \in \mathbb{X},$$

where Π is *the set of all policies*.

A policy π is called *optimal* for the respective criterion, if $g^\pi(x) = g(x)$ for all $x \in \mathbb{X}$. For $g^\pi = v_{n,\alpha}^\pi$, the optimal policy is called *n-horizon discount optimal*; for $g^\pi = v_\alpha^\pi$, it is called *discount optimal*; for $g^\pi = w^\pi$, it is called *average cost optimal*.

It is well known (see, e.g., Bertsekas and Shreve [4, Proposition 8.2]) that the functions $v_{n,\alpha}(x)$ recursively satisfy the following *optimality equations* with $v_{0,\alpha}(x) = 0$ for all $x \in \mathbb{X}$,

$$v_{n+1,\alpha}(x) = \inf_{a \in A(x)} \left\{ c(x, a) + \alpha \int_{\mathbb{X}} v_{n,\alpha}(y) q(dy \mid x, a) \right\}, \quad x \in \mathbb{X}, \ n = 0, 1, \ldots. \tag{3}$$

In addition, a Markov policy ϕ, defined at the first N steps by the mappings $\phi_0, \ldots, \phi_{N-1}$, that satisfy for all $n = 1, \ldots, N$ the equations

$$v_{n,\alpha}(x) = c(x, \phi_{N-n}(x)) + \alpha \int_{\mathbb{X}} v_{n-1,\alpha}(y) q(dy \mid x, \phi_{N-n}(x)), \quad x \in \mathbb{X}, \tag{4}$$

is optimal for the horizon N; see, e.g., Bertsekas and Shreve [4, Lemma 8.7].

It is also well known (Bertsekas and Shreve [4, Propositions 9.8 and 9.12]) that v_α, where $\alpha \in [0, 1)$ satisfies the following discounted cost optimality equation:

$$v_\alpha(x) = \inf_{a \in A(x)} \left\{ c(x, a) + \alpha \int_{\mathbb{X}} v_\alpha(y) q(dy \mid x, a) \right\}, \quad x \in \mathbb{X}, \tag{5}$$

and a stationary policy ϕ_α is discount optimal if and only if

$$v_\alpha(x) = c(x, \phi_\alpha(x)) + \alpha \int_{\mathbb{X}} v_\alpha(y) q(dy \mid x, \phi_\alpha(x)), \quad x \in \mathbb{X}. \tag{6}$$

3. General assumptions and auxiliary results. Following Schäl [26], consider the following assumption.

ASSUMPTION (G). $\quad w^* := \inf_{x \in \mathbb{X}} w(x) < +\infty.$

This assumption is equivalent to the existence of $x \in \mathbb{X}$ and $\pi \in \Pi$ with $w^\pi(x) < \infty$. If Assumption (G) does not hold, then the problem is trivial, because $w(x) = \infty$ for all $x \in \mathbb{X}$ and any policy π is average cost optimal. Define the following quantities for $\alpha \in [0, 1)$:

$$m_\alpha = \inf_{x \in \mathbb{X}} v_\alpha(x), \qquad u_\alpha(x) = v_\alpha(x) - m_\alpha,$$

$$\underline{w} = \liminf_{\alpha \uparrow 1} (1 - \alpha) m_\alpha, \qquad \bar{w} = \limsup_{\alpha \uparrow 1} (1 - \alpha) m_\alpha.$$

Observe that $u_\alpha(x) \geq 0$ for all $x \in \mathbb{X}$. According to Schäl [26, Lemma 1.2], Assumption (G) implies

$$-\infty < \underline{w} \leq \bar{w} \leq w^* < +\infty. \tag{7}$$

According to Schäl [26, Proposition 1.3], under Assumption (G), if there exists a measurable function $u \colon \mathbb{X} \to [0, +\infty)$ and a stationary policy ϕ such that

$$\underline{w} + u(x) \geq c(x, \phi(x)) + \int_{\mathbb{X}} u(y) q(dy \mid x, \phi(x)), \quad x \in \mathbb{X}, \tag{8}$$

then ϕ is *average cost optimal* and $w(x) = w^* = \underline{w} = \bar{w}$ for all $x \in \mathbb{X}$. Here, we need a different form of such a statement.

THEOREM 1. *Let Assumption* (G) *hold. If there exists a measurable function $u \colon \mathbb{X} \to [0, +\infty)$ and a stationary policy ϕ such that*

$$\bar{w} + u(x) \geq c(x, \phi(x)) + \int_{\mathbb{X}} u(y) q(dy \mid x, \phi(x)), \quad x \in \mathbb{X}, \tag{9}$$

then ϕ is average cost optimal and

$$w(x) = w^\phi(x) = \limsup_{\alpha \uparrow 1} (1 - \alpha) v_\alpha(x) = \bar{w} = w^*, \quad x \in \mathbb{X}. \tag{10}$$

PROOF. Similarly to Hernández-Lerma [18, p. 239] or Schäl [26, Proposition 1.3], since u is nonnegative, by iterating (9), we obtain

$$n\bar{w} + u(x) \geq v_{n,1}^\phi(x), \quad n \geq 1, \ x \in \mathbb{X}.$$

Therefore, after dividing the last inequality by n and setting $n \to \infty$, we have

$$\bar{w} \geq w^\phi(x) \geq w(x) \geq w^*, \quad x \in \mathbb{X}, \tag{11}$$

where the second and third inequalities follow from the definitions of w and w^*, respectively. Since $\bar{w} \geq w^*$, inequalities (7) imply that for all $\pi \in \Pi$,

$$w^* = \bar{w} \leq \limsup_{\alpha \uparrow 1} (1 - \alpha) v_\alpha(x) \leq \limsup_{\alpha \uparrow 1} (1 - \alpha) v_\alpha^\pi(x) \leq w^\pi(x), \qquad \pi \in \Pi, \ x \in \mathbb{X}.$$

Finally, we obtain that

$$w^* = \bar{w} \le \limsup_{\alpha \uparrow 1}(1-\alpha)v_\alpha(x) \le \inf_{\pi \in \Pi} w^\pi(x) = w(x) \le w^\phi(x) \le \bar{w}, \quad x \in \mathbb{X}, \tag{12}$$

where the last inequality follows from (11). Thus all the inequalities in (12) are equalities. $\square$

Let us set $\mathbb{R} = [-\infty, +\infty)$, $\mathbb{R}_+ = [0, \infty)$, and $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. For an $\bar{\mathbb{R}}$-valued function f, defined on a Borel subset U of a Polish space $\mathbb{Y}$, consider the level sets

$$\mathscr{D}_f(\lambda) = \{y \in U \colon f(y) \le \lambda\}, \tag{13}$$

$-\infty < \lambda < +\infty$. We recall that the function f is *lower semicontinuous on U* if all the level sets $\mathscr{D}_f(\lambda)$ are closed and the function is *inf-compact on U* if all these sets are compact. The level sets $\mathscr{D}_f(\lambda)$ satisfy the following properties that are used in this paper:

(a) if $\lambda_1 > \lambda$, then $\mathscr{D}_f(\lambda) \subseteq \mathscr{D}_f(\lambda_1)$;

(b) if g, f are functions on U satisfying $g(y) \ge f(y)$ for all $y \in U$, then $\mathscr{D}_g(\lambda) \subseteq \mathscr{D}_f(\lambda)$.

A set is called σ-*compact* if it is a union of a countable number of compact sets. Denote by $K(\mathbb{A})$ the *family of all nonempty compact subsets* of $\mathbb{A}$ and by $K_\sigma(\mathbb{A})$ the *family of all σ-compact subsets* of $\mathbb{A}$; $K(\mathbb{A}) \subset K_\sigma(\mathbb{A})$. Also, denote by $S(\mathbb{A})$ the set of nonempty subsets of $\mathbb{A}$.

A set-valued mapping $F \colon \mathbb{X} \to S(\mathbb{A})$ is *upper semicontinuous* at $x \in \mathbb{X}$ if, for any neighborhood G of the set $F(x)$, there is a neighborhood of x; say, $U(x)$ such that $F(y) \subseteq G$ for all $y \in U(x)$ (see, e.g., Berge [3, p. 109] or Zgurovsky et al. [32, Chapter 1, p. 7]). A set-valued mapping is called *upper semicontinuous* if it is upper semicontinuous at all $x \in \mathbb{X}$.

For weakly continuous transition probabilities, the following basic assumptions were considered in Schäl [26].

Assumption (W).
 (i) *c is lower semicontinuous and bounded below on* $\mathrm{Gr}(A)$;
 (ii) *$A(x) \in K(\mathbb{A})$ for $x \in \mathbb{X}$ and $A \colon \mathbb{X} \to K(\mathbb{A})$ is upper semicontinuous*; *and*
 (iii) *the transition probability $q(\cdot \mid x, a)$ is weakly continuous in $(x, a) \in \mathrm{Gr}(A)$.*

Weak continuity of q in (x, a) means that

$$\int_{\mathbb{X}} f(z)q(dz \mid x_k, a_k) \to \int_{\mathbb{X}} f(z)q(dz \mid x, a) \quad \text{as } k \to +\infty$$

for any sequence $\{(x_k, a_k), k \ge 0\}$ converging to (x, a), where $(x_k, a_k), (x, a) \in \mathrm{Gr}(A)$, and for any bounded continuous function $f \colon \mathbb{X} \to \mathbb{R}$. We notice that there is an additional assumption in Schäl [26]; namely, that $\mathbb{X}$ is a locally compact space with countable base. However, as follows from this paper, the assumption is not necessary here as well as in Feinberg and Lewis [14], because there exists at least one stationary policy. We also remark that the assumptions in (W) were presented in a different order here than in Schäl [26], and that it is assumed in Schäl [26] that c is nonnegative. Because for discounted and average cost criteria the cost function can be shifted by adding any constant, the boundedness and nonnegativity of c are equivalent assumptions. We consider Assumption (Wu) from Feinberg and Lewis [14] without assuming that $\mathbb{X}$ is locally compact.

Assumption (Wu).
 (i) *c is inf-compact on* $\mathrm{Gr}(A)$; *and*
 (ii) *Assumption (W(iii)) holds.*

In this paper, we consider the following more general assumption. The topological meaning of Assumption (W*(ii)) is explained in Feinberg et al. [15, Lemma 2.5].

Assumption (W*).
 (i) *Assumption (W(i)) holds*;
 (ii) *if a sequence $\{x_n\}_{n=1,2,\ldots}$ with values in $\mathbb{X}$ converges and its limit x belongs to $\mathbb{X}$, then any sequence $\{a_n\}_{n=1,2,\ldots}$ with $a_n \in A(x_n)$, $n = 1, 2, \ldots$, satisfying the condition that the sequence $\{c(x_n, a_n)\}_{n=1,2,\ldots}$ is bounded above, has a limit point $a \in A(x)$; and*
 (iii) *Assumption (W(iii)) holds.*

Lemma 1.   *The following statements hold*:
 (i) *Assumption (W) implies Assumption (W*); and*
 (ii) *Assumption (Wu) implies Assumption (W*).*

PROOF. (i) Let $x_n \to x$ as $n \to \infty$, where $x \in \mathbb{X}$ and $x_n \in \mathbb{X}$, $n = 1, \dots$. We show that under Assumption (W(ii)), any sequence $\{a_n\}_{n=1,2,\dots}$ with $a_n \in A(x_n)$ has a limit point $a \in A(x)$. Indeed, since $\mathcal{K} := (\bigcup_{n \geq 1} \{x_n\}) \cup \{x\}$ is a compact set and set-valued mapping $A \colon \mathbb{X} \to K(\mathbb{A})$ is upper semicontinuous, then Berge [3, Theorem 3, p. 110] implies that the image $A(\mathcal{K})$ is also compact. As $\{a_n\}_{n \geq 1} \subset A(\mathcal{K})$, then the sequence $\{a_n\}_{n \geq 1}$ has a limit point $a \in \mathbb{A}$. Consider a sequence $n_k \to \infty$ such that $a_{n_k} \to a$. Since $A(z) \in K(\mathbb{A})$ for all $z \in X$, the upper semicontinuous set-valued mapping A is closed, and since A is closed, $a \in A(x)$; Berge [3, Theorems 5–6, pp. 111–112].

(ii) Since c is inf-compact, it is lower semicontinuous and bounded below. We just need to show that Assumption (W*(ii)) holds. Let us consider $x_n \to x$ as $n \to +\infty$ and $a_n \in A(x_n)$, $n = 1, 2, \dots$ such that $x_n, x \in \mathbb{X}$ and for some $\lambda < \infty$, the inequality $c(x_n, a_n) \leq \lambda$ holds for all $n = 1, 2, \dots$. Then, by inf-compactness of c on $\mathrm{Gr}(A)$, the level set $\mathscr{D}_c(\lambda)$ is compact. Thus the sequence $\{x_n, a_n\}_{n \geq 1}$ has a limit point $(x, a) \in \mathscr{D}_c(\lambda) \subseteq \mathrm{Gr}(A)$. Since $(x, a) \in \mathrm{Gr}(A)$, we have $a \in A(x)$. $\square$

For any $\alpha \geq 0$ and lower semicontinuous nonnegative function $u \colon \mathbb{X} \to \bar{\mathbb{R}}$, we consider an *operation* η_u^α,

$$\eta_u^\alpha(x, a) = c(x, a) + \alpha \int_{\mathbb{X}} u(y) q(dy \mid x, a), \quad (x, a) \in \mathrm{Gr}(A). \tag{14}$$

Let $L(\mathbb{X})$ be the class of all lower semicontinuous and bounded below functions $\varphi \colon \mathbb{X} \to \bar{\mathbb{R}}$ with $\mathrm{dom}\, \varphi := \{x \in \mathbb{X} \colon \varphi(x) < +\infty\} \neq \varnothing$. Observe that $\eta_u^\alpha = \eta_{\alpha u}^1$.

LEMMA 2. *For any $x \in \mathbb{X}$, the following statements hold*:

(a) *under Assumption (W*(ii))*, *the function $c(x, \cdot)$ is inf-compact on $A(x)$*; and

(b) *under Assumptions (W*(ii, iii))*, *for any $u \in L(\mathbb{X})$ and $\alpha \geq 0$, the function $\eta_u^\alpha(x, \cdot)$ is inf-compact on $A(x)$*.

PROOF. (a) For an arbitrary $\lambda \in \mathbb{R}$ and fixed $x \in \mathbb{X}$, consider the set $\mathscr{D}_{c(x, \cdot)}(\lambda) = \{a \in A(x) \colon c(x, a) \leq \lambda\}$. Assumption (W*(ii)) means that this set is compact. Thus (a) is proved.

(b) Fix $x \in \mathbb{X}$ again. Since $u \in L(\mathbb{X})$ and q is weakly continuous in a, the second summand in (14) is a lower semicontinuous function on $A(x)$ (Hernández-Lerma and Lasserre [19, p. 185]) and it is bounded below by the same constant as u. According to statement (a), $c(x, \cdot)$ is inf-compact on $A(x)$. The sum of an inf-compact function and a bounded below lower semicontinuous function is an inf-compact function. $\square$

A measurable mapping $\phi \colon \mathbb{X} \to \mathbb{A}$ such that $\phi(x) \in A(x)$ for all $x \in \mathbb{X}$ is called a selector (or a measurable selector). In our case, selectors and decision rules are the same objects. Since we identify a stationary policy with a decision rule, selectors and stationary policies are the same objects. The existence of selector for the mapping A is the necessary and sufficient condition for the existence of a policy. Let $E \subseteq \mathbb{X} \times \mathbb{A}$ and $\mathrm{proj}_{\mathbb{X}} E = \{x \in \mathbb{X} \colon (x, a) \in E$ for some $a \in E\}$ be a projection of E on $\mathbb{X}$. A Borel map $f \colon \mathrm{proj}_{\mathbb{X}} E \to \mathbb{A}$ is called a Borel uniformization of E if $(x, f(x)) \in E$ for all $x \in \mathrm{proj}_{\mathbb{X}} E$. Let $E_x = \{a \colon (x, a) \in E\}$ be a cut of E at $x \in \mathbb{X}$.

ARSENIN-KUNUGUI THEOREM (Kechris [21, p. 297]). *If E is a Borel subset of $\mathbb{X} \times \mathbb{A}$ and $E_x \in K_\sigma(\mathbb{A})$ for all $x \in \mathbb{X}$, then there exists a Borel uniformization of E and $\mathrm{proj}_{\mathbb{X}} E$ is a Borel set.*

We remark that it is assumed in Kechris [21, p. 297] that $\mathbb{X}$ is a standard Borel space (that is, isomorphic to a Borel subset of a Polish space) and $\mathbb{A}$ is a Polish space. Here, $\mathbb{X}$ and $\mathbb{A}$ are Borel subsets of Polish spaces. These two formulations are obviously equivalent.

We recall that $\mathrm{Gr}(A)$ is assumed to be Borel and $A(x) \neq \varnothing$, $x \in \mathbb{X}$. With $E = \mathrm{Gr}(A)$, the Arsenin-Kunugui theorem implies the existence of a stationary policy under the assumption $A(x) \in K(\mathbb{A})$, $x \in \mathbb{X}$. Thus, Assumption (W) implies the existence of a policy for the MDP.

Let Assumption (W*) hold. Set $F(x) = \{a \in A(x) \colon c(x, a) < \infty\}$, $x \in \mathbb{X}$. In view of Lemma 2, $F(x) = \bigcup_{n \in \{1, 2, \dots\}} \mathscr{D}_{c(x, \cdot)}(n) \in K_\sigma(\mathbb{A})$. In addition, $\mathrm{Gr}(F) = \{(x, a) \in \mathrm{Gr}(A) \colon c(x, a) < \infty\}$ is a Borel subset of $\mathbb{X} \times \mathbb{A}$. Thus, if the function c takes only finite values, a stationary policy exists in view of the Arsenin-Kunugui Theorem.

Of course, if it is possible that $c(x, a) = \infty$, a uniformization may not exist. For example, this takes place when $c(x, a) = \infty$ for all $(x, a) \in \mathrm{Gr}(A)$ and $\mathrm{Gr}(A)$ does not have a measurable selector. However, $c(x, a) = \infty$ means from a modeling prospective that this state-action pair should be excluded, because selecting a in x leads to the worst-possible result. If there are state-action pairs (x, a) with $c(x, a) = \infty$ and $\mathrm{Gr}(A)$ does not have a uniformization, the MDP can be transformed into an MDP modeling the same problem and with a nonempty set of policies. Let us exclude the situation when $c(x, a) = \infty$ for all $(x, a) \in \mathrm{Gr}(A)$, because it is trivial: all the actions are bad. Define $X = \mathrm{proj}_{\mathbb{X}} \mathrm{Gr}(F)$ and $Y = \mathbb{X} \setminus X$. Under Assumption (W*), the Arsenin-Kunugui Theorem implies that X is Borel and there exist a Borel mapping f from X to $\mathbb{A}$ such that $f(x) \in F(x)$ for

all $x \in X$. If $Y = \varnothing$ (that is, there exists an action $a \in A(x)$ with $c(x, a) < \infty$ for each $x \in \mathbb{X}$), then $\phi = f$ is a stationary policy.

Let us consider the situation when $Y \neq \varnothing$. In such an MDP, as soon as the state is in Y, the losses are infinite and there is no reason to model the process after this. Let us transform the model by choosing any $x^* \in Y$ and any $a^* \in \mathbb{A}$ and setting the new state set $\mathbb{X}^* = X \cup \{x^*\}$, keeping the original action set $\mathbb{A}$, setting new action sets $A^*(x) = F(x)$ for $x \in X$ and $A^*(x^*) = \{a^*\}$, defining the new cost function

$$c^*(x, a) = \begin{cases} c(x, a) & \text{if } x \in Y \text{ and } a \in F(x), \\ \infty & \text{if } x = x^* \text{ and } a = a^*, \end{cases}$$

and considering new transition probabilities defined for $x \in X^*$ and $a \in A^*(x)$ by

$$q^*(B \mid x, a) = \begin{cases} q(B \mid x, a) & \text{if } B \subseteq X, \ B \in \mathscr{B}(\mathbb{X}) \text{ and } x \in X, \\ q(Y \mid x, a) & \text{if } B = \{x^*\} \text{ and } x \in X, \\ 1 & \text{if } B = \{x^*\} \text{ and } x = x^*. \end{cases}$$

The new MDP is nontrivial in the sense that the set of policies is not empty. Finding an optimal policy for this MDP is equivalent to finding a policy for the original MDP until its first exit time from X, and in both cases, the process incurs infinite losses if it leaves X. So, the original and the new MDP model are the same problem.

The following lemma is useful for establishing continuity properties of the value functions $v_{n,\alpha}(x)$ and $v_\alpha(x)$ in $x \in \mathbb{X}$; for later relevant results, see Feinberg et al. [15].

LEMMA 3. *If Assumption* (W^*) *holds and* $u \in L(\mathbb{X})$, *then the function*

$$u^*(x) := \inf_{a \in A(x)} \left[c(x, a) + \int_{\mathbb{X}} u(y) q(dy \mid x, a) \right], \quad x \in \mathbb{X}, \tag{15}$$

belongs to $L(\mathbb{X})$, *and there exists* $f \in \mathbb{F}$ *such that*

$$u^*(x) = c(x, f(x)) + \int_{\mathbb{X}} u(y) q(dy \mid x, f(x)), \quad x \in \mathbb{X}. \tag{16}$$

Moreover, infimum in (15) *can be replaced by minimum, and the nonempty sets*

$$A_*(x) = \left\{ a \in A(x) \colon u^*(x) = c(x, a) + \int_{\mathbb{X}} u(y) q(dy \mid x, a) \right\}, \qquad x \in \mathbb{X}, \tag{17}$$

satisfy the following properties:
  (a) *the graph* $\mathrm{Gr}(A_*) = \{(x, a) \colon x \in \mathbb{X}, a \in A_*(x)\}$ *is a Borel subset of* $\mathbb{X} \times \mathbb{A}$;
  (b) *if* $u^*(x) = +\infty$, *then* $A_*(x) = A(x)$, *and if* $u^*(x) < +\infty$, *then* $A_*(x)$ *is compact.*

PROOF. Under Assumption (W^*), for any lower semicontinuous on $\mathbb{X}$, bounded below function $u \colon \mathbb{X} \to \bar{\mathbb{R}}$ and $\alpha \in (0, 1]$, the function $\eta_u^\alpha(x, \cdot)$ is inf-compact on $A(x)$, $x \in \mathbb{X}$. This follows from Lemma 2. Thus, infimum in (15) can be replaced by minimum and $A^*(x)$ is nonempty for any $x \in \mathbb{X}$.

Now, we show that u^* is lower semicontinuous on $\mathbb{X}$. Let us fix an arbitrary $x \in \mathbb{X}$ and any sequence $x_n \to x$ as $n \to +\infty$. We need to prove the inequality

$$u^*(x) \le \liminf_{n \to +\infty} u^*(x_n). \tag{18}$$

If $\liminf_{n \to +\infty} u^*(x_n) = +\infty$, then (18) obviously holds. Thus we consider the case, when $\liminf_{n \to +\infty} u^*(x_n) < +\infty$. There exists a subsequence $\{x_{n_k}\}_{k \geq 1} \subseteq \{x_n\}_{n \geq 1}$ such that

$$\liminf_{n \to +\infty} u^*(x_n) = \lim_{k \to +\infty} u^*(x_{n_k}).$$

Setting $\lambda = \lim_{k \to +\infty} u^*(x_{n_k}) + 1$, we get the inequality $u^*(x_{n_k}) \le \lambda$ for all $k \geq K$, where K is some natural number. Since the function $\eta_u^1(x, \cdot)$ is inf-compact on $A(x)$, Equation (15) can be rewritten as

$$u^*(x) := \min_{a \in A(x)} \eta_u^1(x, a), \quad x \in \mathbb{X}.$$

Thus, for any $k \geq K$, there exists $a_k \in A(x_{n_k})$ such that $u^*(x_{n_k}) = \eta_u^1(x_{n_k}, a_k)$. Therefore

$$c(x_{n_k}, a_k) \leq \eta_u^1(x_{n_k}, a_k) \leq \lambda, \quad k \geq K.$$

In view of Assumption (W*(ii)), there exists a convergent subsequence $\{a_{k_m}\}_{m \geq 1}$ of the sequence $\{a_k\}_{k \geq 1}$ such that $a_{k_m} \to a \in A(x)$ as $m \to +\infty$. Due to lower semicontinuity of η_u^1 on $\mathrm{Gr}(A)$,

$$\liminf_{n \to +\infty} u^*(x_n) = \lim_{k \to +\infty} u^*(x_{n_k}) = \lim_{m \to +\infty} u^*(x_{n_{k_m}}) = \lim_{m \to +\infty} \eta_u^1(x_{n_{k_m}}, a_{k_m}) \geq \eta_u^1(x, a) \geq u^*(x).$$

Inequality (18) holds. Thus u^* is lower semicontinuous on $\mathbb{X}$.

Now, we consider the nonempty sets $A_*(x)$, $x \in \mathbb{X}$, defined in (17). The graph $\mathrm{Gr}(A_*)$ is a Borel subset of $\mathbb{X} \times \mathbb{A}$, because $\mathrm{Gr}(A_*) = \{(x, a): u^*(x) = \eta_u^1(x, a)\}$, and the functions η_u^1 and u^* are lower semicontinuous on $\mathrm{Gr}(A)$ and $\mathbb{X}$, respectively, and therefore they are Borel.

We remark that if $u^* = +\infty$, then $A_*(x) = A(x)$. If $u^*(x) < \infty$, then Lemma 2 implies that the set $A_*(x)$ is compact. Indeed, fix any $x \in \mathbb{X}_f := \{x \in \mathbb{X}: u^*(x) < \infty\}$ and set $\lambda = u^*(x)$. Then, the set $A_*(x) = \{a \in A(x): \eta_u^1(x, a) \leq \lambda\} = \mathscr{D}_{\eta_u^1(x, \cdot)}(\lambda)$ is compact, because $\eta_u^1(x, \cdot)$ is inf-compact on $A(x)$.

Let us prove the existence of $f \in \mathbb{F}$ satisfying (16). Since the function u^* is lower semicontinuous, it is Borel and the sets $X_\infty := \{x \in \mathbb{X}: u^*(x) = +\infty\}$ and $\mathbb{X}_f$ are Borel. Therefore the graph of the mapping $A_*: \mathbb{X}_f \to 2^{\mathbb{A}}$ is the Borel set $\mathrm{Gr}(A) \setminus (\mathbb{X}_\infty \times \mathbb{A})$. Since the nonempty sets $A_*(x)$ are compact for all $x \in \mathbb{X}_f$, the Arsenin-Kunugui Theorem implies the existence of a Borel selector $f_1: \mathbb{X}_f \to \mathbb{A}$ such that $f_1(x) \in A_*(x)$ for all $x \in \mathbb{X}$. Consider any Borel mapping f_2 from $\mathbb{X}$ to $\mathbb{A}$ satisfying $f_2(x) \in A(x)$ for all $x \in \mathbb{X}$ and set

$$f(x) = \begin{cases} f_1(x) & \text{if } x \in \mathbb{X}_f, \\ f_2(x) & \text{if } x \in \mathbb{X}_\infty. \end{cases}$$

Then, $f \in \mathbb{F}$ and $f(x) \in A_*(x)$ for all $x \in \mathbb{X}$. $\square$

The following Lemma 4 is formulated in Schäl [26, Lemma 2.3(ii)] without proof. Reference Serfozo [29] mentioned in Schäl [26, Lemma 2.3(ii)] contains relevant facts, but it does not contain this statement. The proofs of Lemma 4 and its generalization to the functions h_n that can take negative values are provided in Feinberg et al. [16]. To make the current paper self-contained, we provide the proof of Lemma 4 in the appendix. Recall that for a metric space S, the family of all probability measures on $(S, \mathscr{B}(S))$ is denoted by $\mathbb{P}(S)$.

LEMMA 4. *Let S be an arbitrary metric space, $\{\mu_n\}_{n \geq 1} \subset \mathbb{P}(S)$ converges weakly to $\mu \in \mathbb{P}(S)$, and $\{h_n\}_{n \geq 1}$ be a sequence of measurable nonnegative $\bar{\mathbb{R}}$-valued functions on S. Then*

$$\int_S \underline{h}(s) \, \mu(ds) \leq \liminf_{n \to +\infty} \int_S h_n(s) \, \mu_n(ds),$$

where $\underline{h}(s) = \liminf_{n \to +\infty, s' \to s} h_n(s')$, $s \in S$.

PROOF. See the appendix. $\square$

We remark that $\liminf_{n \to +\infty, s' \to s} h_n(s')$ is the least upper bound of the set of all $\lambda \in \mathbb{R}$ such that there exist $N = 1, 2, \ldots$ and a neighborhood $U(s)$ of s such that $\lambda \leq \inf\{h_n(s'): n \geq N, s' \in U(s)\}$.

4. Expected total discounted costs. In this section, we establish under Assumption (W*) the standard properties of discounted MDPs: the existence of stationary optimal policies, description of the sets of stationary optimal policies, and convergence of value iterations. Theorem 2 strengthens Feinberg and Lewis [14, Proposition 3.1], where these facts are proved under Assumption (**Wu**). In terms of applications to inventory and queuing control, Assumption (W*) does not require that holding costs increase to infinity as the inventory level (or work load, or the number of customers in queue) increases to infinity.

THEOREM 2. *Let Assumption (W*) hold. Then*

(i) *the functions $v_{n, \alpha}$, $n = 0, 1, 2, \ldots$, and v_α are lower semicontinuous on $\mathbb{X}$, and $v_{n, \alpha}(x) \to v_\alpha(x)$ as $n \to +\infty$ for all $x \in \mathbb{X}$;*

(ii)

$$v_{n+1, \alpha}(x) = \min_{a \in A(x)} \left\{ c(x, a) + \alpha \int_{\mathbb{X}} v_{n, \alpha}(y) q(dy \mid x, a) \right\}, \quad x \in \mathbb{X}, \quad n = 0, 1, \ldots, \tag{19}$$

where $v_{0, \alpha}(x) = 0$ for all $x \in \mathbb{X}$, and the nonempty sets $A_{n, \alpha}(x) := \{a \in A(x): v_{n+1, \alpha}(x) = \eta_{v_{n, \alpha}}^\alpha(x, a)\}$, $x \in \mathbb{X}$, $n = 0, 1, \ldots$, satisfy the following properties: (a) the graph $\mathrm{Gr}(A_{n, \alpha}) = \{(x, a): x \in \mathbb{X}, a \in A_{n, \alpha}(x)\}$, $n = 0, 1, \ldots$, is a Borel subset of $\mathbb{X} \times \mathbb{A}$, and (b) if $v_{n+1, \alpha}(x) = +\infty$, then $A_{n, \alpha}(x) = A(x)$, and if $v_{n+1, \alpha}(x) < +\infty$, then $A_{n, \alpha}(x)$ is compact;

(iii) *for any $N = 1, 2, \ldots$, there exists a Markov optimal N-horizon policy $(\phi_0, \ldots, \phi_{N-1})$, and if, for an N-horizon Markov policy $(\phi_0, \ldots, \phi_{N-1})$, the inclusions $\phi_{N-1-n}(x) \in A_{n,\alpha}(x)$, $x \in \mathbb{X}$, $n = 0, \ldots, N-1$ hold, then this policy is N-horizon optimal*;

(iv) *for $\alpha \in [0, 1)$*,

$$v_\alpha(x) = \min_{a \in A(x)} \left\{ c(x, a) + \alpha \int_{\mathbb{X}} v_\alpha(y) q(dy \mid x, a) \right\}, \quad x \in \mathbb{X}, \tag{20}$$

and the nonempty sets $A_\alpha(x) := \{a \in A(x): v_\alpha(x) = \eta_{v_\alpha}^\alpha(x, a)\}$, $x \in \mathbb{X}$ satisfy the following properties: (a) *the graph $\mathrm{Gr}(A_\alpha) = \{(x, a): x \in \mathbb{X}, a \in A_\alpha(x)\}$ is a Borel subset of $\mathbb{X} \times \mathbb{A}$, and* (b) *if $v_\alpha(x) = +\infty$, then $A_\alpha(x) = A(x)$, and, if $v_\alpha(x) < +\infty$, then $A_\alpha(x)$ is compact.*

(v) *for an infinite horizon, there exists a stationary discount optimal policy ϕ_α, and a stationary policy is optimal if and only if $\phi_\alpha(x) \in A_\alpha(x)$ for all $x \in \mathbb{X}$.*

(vi) (Feinberg and Lewis [14, Proposition 3.1(iv)]) *under Assumption* (**Wu**), *the functions $v_{n,\alpha}$, $n = 1, 2, \ldots$, and v_α are inf-compact on $\mathbb{X}$.*

PROOF. (i)–(v). First, we prove these statements for a nonnegative cost function c. In this case, $v_{n,\alpha}(x) \geq 0$, $n = 0, 1, \ldots$, and $v_\alpha(x) \geq 0$ for all $x \in \mathbb{X}$.

By (3) and Lemma 3, $v_{1,\alpha} \in L(\mathbb{X})$ since $v_{0,\alpha} = 0 \in L(\mathbb{X})$. By the same arguments, if $v_{n,\alpha} \in L(\mathbb{X})$, then $v_{n+1,\alpha} \in L(\mathbb{X})$. Thus $v_{n,\alpha} \in L(\mathbb{X})$ for all $n = 0, 1, \ldots$. By Lemma 2, for any $n = 1, 2, \ldots$, $x \in \mathbb{X}$, and $\lambda \in \mathbb{R}$, the set $\mathbb{D}_{\eta_{v_{n,\alpha}}^\alpha(x, \cdot)}(\lambda)$ is a compact subset of $\mathbb{A}$. By Bertsekas and Shreve [4, Proposition 9.17], $v_{n,\alpha} \uparrow v_\alpha$ as $n \to +\infty$. Since the limit of a monotone increasing sequence of lower semicontinuous functions is again a lower semicontinuous function, $v_\alpha \in L(\mathbb{X})$. Lemma 3, applied to Equations (3) and (5), implies statements (ii) and (iv), respectively. Statement (iii) follows from (4) and statement (v) follows from (6).

Now, let $c(x, a) \geq K$ for all $(x, a) \in \mathrm{Gr}(A)$ and for some $K > -\infty$. For $K \geq 0$, statements (i)–(v) are proved. For $K < 0$, consider the value functions $\tilde{c} = c - K \geq 0$. If the cost function c is substituted with $\tilde{c}$, we substitute the notation v with $\tilde{v}$. Then $v_{n,\alpha}^\pi = \tilde{v}_{n,\alpha}^\pi + ((1 - \alpha^n)/(1 - \alpha))K$, $n = 0, 1, \ldots$ for all policies π. Thus $v_{n,\alpha} = \tilde{v}_{n,\alpha} + ((1 - \alpha^n)/(1 - \alpha))K$, $n = 0, 1, \ldots$, and $v_\alpha = \tilde{v}_\alpha + K/(1 - \alpha)$. Since statements (i)–(v) hold for the shifted costs $\tilde{c}$ and the value functions $\tilde{v}_{n,\alpha}$ and $\tilde{v}_\alpha$, they also hold for the initial cost function c and the value functions $v_{n,\alpha}$ and v_α. □

We remark that the conclusions of Theorem 2 and its proof remain correct when $\alpha = 1$ and the function c is nonnegative.

5. Average costs per unit time.

In this section, we show that Assumption (W*) and "boundedness" Assumption ($\underline{\mathrm{B}}$) on the function u_α, which is weaker than the boundedness Assumption (B) introduced by Schäl [26], lead to the validity of stationary average cost optimal inequalities and the existence of stationary policies. Stronger results hold under Assumption (B).

ASSUMPTION ($\underline{\mathrm{B}}$). (i) *Assumption* (G) *holds, and* (ii) $\liminf_{\alpha \uparrow 1} u_\alpha(x) < \infty$ *for all $x \in \mathbb{X}$.*

Assumption (ULB(ii)) is weaker than the assumption $\sup_{\alpha \in [0, 1)} u_\alpha(x) < \infty$ for all $x \in \mathbb{X}$ considered in Schäl [26]. This assumption and Assumption (G) were combined in Feinberg and Lewis [14] into the following assumption.

ASSUMPTION (B). (i) *Assumption* (G) *holds, and* (ii) $\sup_{\alpha \in [0, 1)} u_\alpha(x) < \infty$ *for all $x \in \mathbb{X}$.*

It seems natural to consider the assumption $\limsup_{\alpha \uparrow 1} u_\alpha(x) < \infty$ for all $x \in \mathbb{X}$, which is stronger than Assumption ($\underline{\mathrm{B}}$(ii)) and weaker than Assumption (B(ii)). However, as the following lemma shows, under Assumption (G), this assumption is equivalent to Assumption (B(ii)).

LEMMA 5. *Let the cost function c be bounded below and Assumption* (G) *hold. Then, for each $x \in \mathbb{X}$, the following two inequalities are equivalent*:
(i) $\sup_{\alpha \in [0, 1)} u_\alpha(x) < \infty$,
(ii) $\limsup_{\alpha \uparrow 1} u_\alpha(x) < \infty$.

PROOF. Obviously, (i)→(ii). Let us prove (ii)→(i). Let (ii) hold. Assume that (i) does not hold. Since $\sup_{\alpha \in [0, 1)} u_\alpha(x) = \max\{\sup_{\alpha \in [0, \alpha^*)} u_\alpha(x), \sup_{\alpha \in [\alpha^*, 1)} u_\alpha(x)\}$ for any $\alpha^* \in [0, 1)$, there exists $\alpha^* \in [0, 1)$ such that $\sup_{\alpha \in [0, \alpha^*)} u_\alpha(x) = \infty$.

Since the function u_α remains unchanged if a finite constant is added to the cost function c, we assume without loss of generality that $c(x, a) \geq 0$ for all $(x, a) \in \mathrm{Gr}(A)$. Since $c \geq 0$, the functions $v_\alpha(x)$ and m_α are nonnegative nondecreasing functions in $\alpha \in [0, 1)$. Since $v_\alpha(x) = u_\alpha(x) + m_\alpha \geq u_\alpha(x)$, we have $\sup_{\alpha \in [0, \alpha^*)} v_\alpha(x) = \infty$, and

therefore $v_\alpha(x) = \infty$ for all $\alpha \in [\alpha^*, 1)$, because of the monotonicity of v_α in α. Thus $\limsup_{\alpha\uparrow 1}(1-\alpha)v_\alpha(x) = \infty$. However, $\limsup_{\alpha\uparrow 1}(1-\alpha)v_\alpha(x) = \limsup_{\alpha\uparrow 1}(1-\alpha)(u_\alpha(x) + m_\alpha) \leq \limsup_{\alpha\uparrow 1}(1-\alpha)u_\alpha(x) + \bar{w} < \infty$, where the last inequality follows from (ii) and (7). The obtained contradiction completes the proof. $\square$

Until the end of this section, we assume that Assumption (B) holds. Let us set

$$u(x) := \liminf_{\alpha\uparrow 1, \, y\to x} u_\alpha(y), \quad x \in \mathbb{X}, \tag{21}$$

where $\liminf_{\alpha\uparrow 1, \, y\to x} u_\alpha(y)$ is the least upper bound of the set of all $\lambda \in \mathbb{R}_+$ such that there exist $\beta \in [0, 1)$ and a neighborhood $U(x)$ of x such that $\lambda \leq \inf\{u_\alpha(y): \alpha \in [\beta, 1), y \in U(x) \cap \mathbb{X}\}$.

Also, define the following nonnegative functions on $\mathbb{X}$:

$$U_\beta(x) = \inf_{\alpha\in[\beta, 1)} u_\alpha(x), \qquad \underline{u}_\beta(x) = \liminf_{y\to x} U_\beta(y), \quad \beta \in [0, 1), \ x \in \mathbb{X}. \tag{22}$$

Observe that all three defined functions take finite values at $x \in \mathbb{X}$. Indeed,

$$\underline{u}_\beta(x) \leq U_\beta(x) \leq \sup_{\beta\in[0, 1)} \inf_{\alpha\in[\beta, 1)} u_\alpha(x) = \liminf_{\alpha\uparrow 1} u_\alpha(x) < \infty, \quad \beta \in [0, 1), \ x \in \mathbb{X}, \tag{23}$$

where the first two inequalities follow from the definitions of $\underline{u}_\beta$ and U_β, respectively, and the last inequality follows from Assumption (B). For $x \in \mathbb{X}$,

$$\begin{aligned}
u(x) &= \sup_{\beta\in[0, 1), \, R>0}\left[\inf_{\alpha\in[\beta, 1), \, y\in B_R(x)} u_\alpha(y)\right] = \sup_{\beta\in[0, 1)} \sup_{R>0} \inf_{y\in B_R(x)} \inf_{\alpha\in[\beta, 1)} u_\alpha(y) \\
&= \sup_{\beta\in[0, 1)} \sup_{R>0} \inf_{y\in B_R(x)} U_\beta(y) = \sup_{\beta\in[0, 1)} \liminf_{y\to x} U_\beta(y) = \sup_{\beta\in[0, 1)} \underline{u}_\beta(x) < \infty,
\end{aligned} \tag{24}$$

where $B_R(x) = \{y \in \mathbb{X}: \rho(y, x) < R\}$, the first equality is (21), the second equality follows from the properties of infima, the third and fifth equalities follow from (22), the fourth equality follows from the definition of $\limsup$, and the last inequality follows from (23). In view of (22), the functions $U_\beta(x)$ and $\underline{u}_\beta(x)$ are nondecreasing in β. Therefore, in view of (24),

$$u(x) = \lim_{\beta\uparrow 1} \underline{u}_\beta(x), \quad x \in \mathbb{X}. \tag{25}$$

We also set for u from (25)

$$A^*(x) := \left\{a \in A(x): \bar{w} + u(x) \geq c(x, a) + \int_{\mathbb{X}} u(y)q(dy \mid x, a)\right\}, \quad x \in \mathbb{X}, \tag{26}$$

and let $A_*(x)$, $x \in \mathbb{X}$ be the sets defined in (17) for this function u; $A_*(x) \subseteq A^*(x)$.

THEOREM 3. *Suppose Assumptions* (W*) *and* (B) *hold. There exist a stationary policy ϕ satisfying* (9) *with u defined in* (21). *Thus, equalities* (10) *hold for this policy ϕ. Furthermore, the following statements hold*:
   (a) *the function u: $\mathbb{X} \to \mathbb{R}_+$, defined in* (21), *is lower semicontinuous*;
   (b) *the nonempty sets $A^*(x)$, $x \in \mathbb{X}$, satisfy the following properties*:
      (b$_1$) *the graph $\mathrm{Gr}(A^*) = \{(x, a): x \in \mathbb{X}, a \in A^*(x)\}$ is a Borel subset of $\mathbb{X} \times \mathbb{A}$; and*
      (b$_2$) *for each $x \in \mathbb{X}$, the set $A^*(x)$ is compact*;
   (c) *a stationary policy ϕ is optimal for average costs and satisfies* (9) *with u defined in* (21) *if $\phi(x) \in A^*(x)$ for all $x \in \mathbb{X}$*;
   (d) *there exists a stationary policy ϕ with $\phi(x) \in A_*(x) \subseteq A^*(x)$ for all $x \in \mathbb{X}$; and*
   (e) *if, in addition, Assumption* (Wu) *holds, then the function u defined in* (21), *is inf-compact*.

Before the proof of Theorem 3, we establish some auxiliary facts.

LEMMA 6. *Under Assumption* (B), *the functions $u, \underline{u}_\alpha$: $\mathbb{X} \to \mathbb{R}_+$, $\alpha \in [0, 1)$ are lower semicontinuous on $\mathbb{X}$. If additionally Assumption* (W*) *holds, the functions u_α: $\mathbb{X} \to \mathbb{R}_+$, $\alpha \in [0, 1)$ are lower semicontinuous on $\mathbb{X}$. Under Assumptions* (Wu) *and* (B), *the functions $u, u_\alpha, \underline{u}_\alpha$: $\mathbb{X} \to \mathbb{R}_+$, $\alpha \in [0, 1)$ are inf-compact on $\mathbb{X}$.*

PROOF. Since $\underline{u}_\alpha(x) \geq 0$, $\alpha \in [0, 1)$, and $x \in \mathbb{X}$, the functions $\underline{u}_\alpha$, $\alpha \in [0, 1)$ are lower semicontinuous; Feinberg and Lewis [14, Lemma 3.1]. Since supremum over any set of lower semicontinuous functions is a lower semicontinuous function, the function u is lower semicontinuous.

According to (7), $\bar{w} := \limsup_{\alpha \uparrow 1}(1-\alpha)m_\alpha = \inf_{\alpha \in (0,1)} \sup_{\alpha \in [\alpha,1)}(1-\alpha)m_\alpha < \infty$. Thus, there exists $\alpha_0 \in [0,1)$ such that

$$\lambda' := \sup_{\alpha \in [\alpha_0, 1)}(1-\alpha)m_\alpha < \infty. \tag{27}$$

Let us assume that the function c is bounded below. As explained in the proof of Lemma 5, without loss of generality, we can assume that $c \geq 0$. Then m_α is a nonnegative nondecreasing function. Thus $(1-\alpha)m_\alpha \leq (1-\alpha)m_{\alpha_0} \leq \lambda'/(1-\alpha_0)$, $\alpha \in [0, \alpha_0)$, and (27) implies that

$$\lambda^* = \sup_{\alpha \in [0,1)}(1-\alpha)m_\alpha < \infty. \tag{28}$$

According to Theorem 2(i, iv, v) under Assumption (W*), the function $u_\alpha(x) = v_\alpha(x) - m_\alpha$ is lower semicontinuous, and a stationary policy ϕ_α is α-discount optimal if and only if for all $x \in \mathbb{X}$,

$$v_\alpha(x) = \min_{a \in A(x)}\left\{c(x,a) + \alpha \int_{\mathbb{X}} v_\alpha(y)q(dy \mid x,a)\right\} = c(x, \phi_\alpha(x)) + \alpha \int_{\mathbb{X}} v_\alpha(y)q(dy \mid x, \phi_\alpha(x)). \tag{29}$$

The first equality in (29) is equivalent to

$$(1-\alpha)m_\alpha + u_\alpha(x) = \min_{a \in A(x)}\left[c(x,a) + \alpha \int_{\mathbb{X}} u_\alpha(y)q(dy \mid x,a)\right], \quad x \in \mathbb{X}. \tag{30}$$

Let Assumption (Wu) hold. The function $u_\alpha(x) = v_\alpha(x) - m_\alpha$ is inf-compact by Theorem 2(vi). Consider an arbitrary $\lambda \in \mathbb{R}_+$. Since $u(x) \geq \underline{u}_{\alpha_1}(x) \geq \underline{u}_{\alpha_2}(x)$, $x \in \mathbb{X}$ for all $\alpha_1, \alpha_2 \in [0,1)$, $\alpha_1 \geq \alpha_2$, then $\mathscr{D}_u(\lambda) \subseteq \mathscr{D}_{\underline{u}_\alpha}(\lambda) \subseteq \mathscr{D}_{\underline{u}_0}(\lambda)$, $\alpha \in [0,1)$. Since the functions u and $\underline{u}_\alpha$ are lower semicontinuous, the sets $\mathscr{D}_u(\lambda)$ and $\mathscr{D}_{\underline{u}_\alpha}(\lambda)$ are closed, $\alpha \in [0,1)$. Therefore, if the set $\mathscr{D}_{\underline{u}_0}(\lambda)$ is compact, then those sets are also compact and the functions u and $\underline{u}_\alpha$, $\alpha \in [0,1)$ are inf-compact.

Observe that (28) and (30) imply that $u_\alpha(x) \geq v_0(x) - \lambda^*$, $x \in X$ for all $\alpha \in [0,1)$. This implies $U_0(x) \geq v_0(x) - \lambda^*$, $x \in X$. Since $\underline{u}_0$ is the largest lower semicontinuous function that is less than or equal to U_0 at all $x \in \mathbb{X}$, we have $\underline{u}_0(x) \geq v_0(x) - \lambda^*$, $x \in X$. Since the function $\underline{u}_0$ is lower semicontinuous, the set $\mathscr{D}_{\underline{u}_0}(\lambda)$ is closed. In addition, $\mathscr{D}_{\underline{u}_0}(\lambda) \subseteq \mathscr{D}_{v_0}(\lambda + \lambda^*)$, where the set $\mathscr{D}_{v_0}(\lambda + \lambda^*)$ is compact (cf. Theorem 2(vi)). Thus the set $\mathscr{D}_{\underline{u}_0}(\lambda)$ is compact, and the functions u and $\underline{u}_\alpha$, $\alpha \in [0,1)$ are inf-compact. $\quad\square$

COROLLARY 1. *Under Assumption* (B̲), *for every sequence* $\alpha_n \uparrow 1$ *as* $n \to +\infty$ *and for every* $x \in \mathbb{X}$,

$$\underline{u}_{\alpha_n}(x) \uparrow u(x) = \lim_{n \to \infty, y \to x} \underline{u}_{\alpha_n}(y).$$

PROOF. Let $\alpha_n \uparrow 1$ as $n \to +\infty$, and $x \in \mathbb{X}$. Similar to (24),

$$\liminf_{n \to +\infty, y \to x} \underline{u}_{\alpha_n}(y) = \sup_{n=1,2,\ldots} \sup_{R>0} \inf_{y \in B_R(x)} \inf_{m \geq n} \underline{u}_{\alpha_m}(y) = \sup_{n=1,2,\ldots} \sup_{R>0} \inf_{y \in B_R(x)} \underline{u}_{\alpha_n}(y)$$

$$= \sup_{n=1,2,\ldots} \liminf_{y \to x} \underline{u}_{\alpha_n}(y) = \lim_{n \to \infty} \underline{u}_{\alpha_n}(x) = u(x),$$

where the second equality holds because the function $\underline{u}_\alpha(y)$ is nondecreasing in α, the fourth equality holds because it is lower semicontinuous, and the last equality follows from (25). $\quad\square$

LEMMA 7. *Under Assumptions* (W*) *and* (B̲), *the following inequalities hold*:

$$\bar{w} + u(x) \geq \min_{a \in A(x)}\left[c(x,a) + \int_{\mathbb{X}} u(y)q(dy \mid x,a)\right], \quad x \in \mathbb{X}. \tag{31}$$

PROOF. Let us fix an arbitrary $\varepsilon^* > 0$. Since $\bar{w} = \limsup_{\alpha \uparrow 1}(1-\alpha)m_\alpha$, there exists $\alpha_0 \in [0,1)$ such that

$$\bar{w} + \varepsilon^* > (1-\alpha)m_\alpha, \quad \alpha \in [\alpha_0, 1). \tag{32}$$

Our next goal is to prove the inequality

$$\bar{w} + \varepsilon^* + u(x) \geq \min_{a \in A(x)}\left[c(x,a) + \alpha \int_{\mathbb{X}} \underline{u}_\alpha(y)q(dy \mid x,a)\right], \quad x \in \mathbb{X}, \ \alpha \in [\alpha_0, 1). \tag{33}$$

Indeed, by (30) and (32) for every $\alpha, \beta \in [\alpha_0, 1)$ such that $\alpha \le \beta$, and for every $x \in \mathbb{X}$,

$$\bar{w} + \varepsilon^* + u_\beta(x) > (1-\beta)m_\beta + u_\beta(x) = \min_{a \in A(x)} \left[c(x,a) + \beta \int_{\mathbb{X}} u_\beta(y) q(dy \mid x, a) \right]$$

$$\ge \min_{a \in A(x)} \left[c(x,a) + \alpha \int_{\mathbb{X}} U_\alpha(y) q(dy \mid x, a) \right].$$

As the right-hand side does not depend on $\beta \in [\alpha, 1)$, we have for all $x \in \mathbb{X}$ and for all $\alpha \in [\alpha_0, 1)$,

$$\bar{w} + \varepsilon^* + U_\alpha(x) = \inf_{\beta \in [\alpha, 1)} [\bar{w} + \varepsilon^* + u_\beta(x)] \ge \min_{a \in A(x)} \left[c(x,a) + \alpha \int_{\mathbb{X}} U_\alpha(y) q(dy \mid x, a) \right]$$

$$\ge \min_{a \in A(x)} \left[c(x,a) + \alpha \int_{\mathbb{X}} \underline{u}_\alpha(y) q(dy \mid x, a) \right] = \min_{a \in A(x)} \eta_{\underline{u}_\alpha}^\alpha(x, a).$$

By Lemma 3, the function $x \to \min_{a \in A(x)} \eta_{\underline{u}_\alpha}^\alpha(x, a)$ is lower semicontinuous on $\mathbb{X}$. Thus

$$\liminf_{y \to x} \min_{a \in A(y)} \eta_{\underline{u}_\alpha}^\alpha(y, a) \ge \min_{a \in A(x)} \eta_{\underline{u}_\alpha}^\alpha(x, a), \quad x \in \mathbb{X}, \ \alpha \in [0, 1),$$

and as by definition (22), $\underline{u}_\alpha(x) = \liminf_{y \to x} U_\alpha(y)$, we finally obtain

$$\bar{w} + \varepsilon^* + \underline{u}_\alpha(x) \ge \min_{a \in A(x)} \eta_{\underline{u}_\alpha}^\alpha(x, a), \quad x \in \mathbb{X}, \ \alpha \in [\alpha_0, 1). \tag{34}$$

Since by Corollary 1 $u(x) = \sup_{\alpha \in [\alpha_0, 1)} \underline{u}_\alpha(x)$ for all $x \in \mathbb{X}$, (34) yields (33).

To complete the proof of the lemma, we fix an arbitrary $x \in \mathbb{X}$. By Lemma 3, for any $\alpha \in [0, 1)$, there exists $a_\alpha \in A(x)$ such that $\min_{a \in A(x)} \eta_{\underline{u}_\alpha}^\alpha(x, a) = \eta_{\underline{u}_\alpha}^\alpha(x, a_\alpha)$. Since $\underline{u}_\alpha \ge 0$ for $\alpha \in [\alpha_0, 1)$, the inequality (33) can be continued as

$$\bar{w} + \varepsilon^* + u(x) \ge \eta_{\underline{u}_\alpha}^\alpha(x, a_\alpha) \ge c(x, a_\alpha). \tag{35}$$

Thus, for all $\alpha \in [\alpha_0, 1)$,

$$a_\alpha \in \mathscr{D}_{\eta_{\underline{u}_\alpha}^\alpha(x, \cdot)}(\bar{w} + \varepsilon^* + u(x)) \subseteq \mathscr{D}_{c(x, \cdot)}(\bar{w} + \varepsilon^* + u(x)) \subseteq A(x).$$

By Lemma 2, the set $\mathscr{D}_{c(x, \cdot)}(\bar{w} + \varepsilon^* + u(x))$ is compact. Thus, for every sequence $\beta_n \uparrow 1$ of numbers from $[\alpha_0, 1)$, there is a subsequence $\{\alpha_n\}_{n \ge 1}$ such that the sequence $\{a_{\alpha_n}\}_{n \ge 1}$ converges and $a_* := \lim_{n \to \infty} a_{\alpha_n} \in A(x)$.

Consider a sequence $\alpha_n \uparrow 1$ such that $a_{\alpha_n} \to a_*$ for some $a_* \in A(x)$. Due to Lemma 4 and Corollary 1,

$$\liminf_{n \to +\infty} \alpha_n \int_{\mathbb{X}} \underline{u}_{\alpha_n}(y) q(dy \mid x, a_n) \ge \int_{\mathbb{X}} u(y) q(dy \mid x, a_*). \tag{36}$$

Since the function c is lower semicontinuous, (35) and (36) imply

$$\bar{w} + \varepsilon^* + u(x) \ge \limsup_{n \to \infty} \eta_{\underline{u}_{\alpha_n}}^{\alpha_n}(x, a_{\alpha_n}) \ge c(x, a_*) + \int_{\mathbb{X}} u(y) q(dy \mid x, a_*) \ge \min_{a \in A(x)} \eta_u^1(x, a).$$

Since $\bar{w} + \varepsilon^* + u(x) \ge \min_{a \in A(x)} \eta_u^1(x, a)$ for any $\varepsilon^* > 0$, this is also true when $\varepsilon^* = 0$. $\square$

PROOF OF THEOREM 3. Lemma 6 contains statements (a) and (e). Since $\mathrm{Gr}(A^*) = \{(x, a) \in \mathrm{Gr}(A): g(x, a) \ge 0\}$, where $g(x, a) = \bar{w} + u(x) - c(x, a) - \int_{\mathbb{X}} u(y) q(dy \mid x, a)$ is a Borel function, the set $\mathrm{Gr}(A^*)$ is Borel. The sets $A^*(x), x \in \mathbb{X}$ are compact in view of Lemma 2(b). Thus the statement (**b**) is proved. The Arsenin-Kunugui Theorem implies the existence of a stationary policy ϕ such that $\phi(x) \in A^*(x)$ for all $x \in \mathbb{X}$. Statement (d) follows from Lemma 3 and the Arsenin-Kunugui Theorem. The rest follows from Theorem 1. $\square$

THEOREM 4. *Suppose Assumptions* (W*) *and* (B) *hold. Then, all the conclusions of Theorem 3 hold, and in addition, for a stationary policy ϕ satisfying* (9) *with u defined in* (21),

$$w^\phi(x) = \underline{w} = \lim_{\alpha \uparrow 1}(1-\alpha)v_\alpha(x) = \lim_{N \to \infty} \frac{1}{N} v_{N,1}^\phi(x), \quad x \in \mathbb{X}. \tag{37}$$

PROOF. Consider a sequence $\{\alpha(n)\}_{n\geq 1}$ such that $\alpha(n) \uparrow 1$ as $n \to +\infty$, and

$$\lim_{n\to +\infty} (1-\alpha(n))m_{\alpha(n)} = \underline{w}.$$

Define the following nonnegative functions on $\mathbb{X}$:

$$\tilde{U}_n(x) = \inf_{m\geq n} u_{\alpha(m)}(x), \qquad \tilde{\underline{u}}_n(x) = \liminf_{y\to x} \tilde{U}_n(y), \quad n \geq 1, \ x \in \mathbb{X},$$

and

$$\tilde{u}(x) = \sup_{n\geq 1} \tilde{\underline{u}}_n(x), \quad x \in \mathbb{X}. \tag{38}$$

Observe that

$$\tilde{\underline{u}}_n(x) \leq \tilde{U}_n(x) \leq \liminf_{m\to +\infty} u_{\alpha(m)}(x) < \infty, \quad x \in \mathbb{X}, \ n = 1, 2, \ldots, \tag{39}$$

where the first two inequalities follow from the definitions of $\tilde{\underline{u}}_n$ and $\tilde{U}_n$, respectively, and the last inequality follows from Assumption (B). As follows from (38) and (39), $\tilde{u}(x) \leq \liminf_{m\to +\infty} u_{\alpha(m)}(x) < +\infty$. According to Feinberg and Lewis [14, Lemma 3.1], the functions $\tilde{\underline{u}}_n$, $n \geq 1$ are lower semicontinuous on $\mathbb{X}$. Therefore their supremum $\tilde{u}$ is also lower semicontinuous. In addition,

$$\tilde{u}(x) = \sup_{n\geq 1} \sup_{R>0} \inf_{y\in B_R(x)} \inf_{m\geq n} u_{\alpha(m)}(y) = \liminf_{n\to +\infty,\ y\to x} u_{\alpha(n)}(y), \quad x \in \mathbb{X},$$

where the first equality follows from the definitions of $\tilde{U}_n$, $\tilde{\underline{u}}_n$, and $\tilde{u}$, and the second equality is the definition of the lim inf. Since $\tilde{U}_n(x) \uparrow$, we have $\tilde{\underline{u}}_n(x) \uparrow \tilde{u}(x)$ as $n \to \infty$ for all $x \in \mathbb{X}$.

We show next that for each $x \in \mathbb{X}$,

$$\underline{w} + \tilde{u}(x) \geq \inf_{a\in A(x)} \left[c(x, a) + \int_{\mathbb{X}} \tilde{u}(y)q(dy \mid x, a) \right]. \tag{40}$$

Indeed, let us fix any $\varepsilon^* > 0$. By the definition of $\underline{w}$, there exists a subsequence $\{\alpha(n_k)\}_{k\geq 1} \subseteq \{\alpha(n)\}_{n\geq 1}$ such that for $k = 1, 2, \ldots$

$$\underline{w} + \varepsilon^* \geq (1-\alpha(n_k))m_{\alpha(n_k)}.$$

Let $x \in \mathbb{X}$ be an arbitrary state. By Theorem 2 for each $k \geq 1$, there exists $a_{n_k} \in A_{\alpha(n_k)}(x)$ such that

$$(1-\alpha(n_k))m_{\alpha(n_k)} + u_{\alpha(n_k)}(x) = c(x, a_{n_k}) + \alpha(n_k) \int_{\mathbb{X}} u_{\alpha(n_k)}(y)q(dy \mid x, a_{n_k}).$$

Thus, similarly to the proof of Lemma 7, we get (40).

From Lemma 3 and the Arsenin-Kunugui Theorem, there exists a stationary policy $\tilde{\phi} \in \mathbb{F}$ such that for any $x \in \mathbb{X}$,

$$\underline{w} + \tilde{u}(x) \geq c(x, \tilde{\phi}(x)) + \int_{\mathbb{X}} \tilde{u}(y)q(dy \mid x, \tilde{\phi}(x)). \tag{41}$$

Thus, by Schäl [26, Proposition 1.3] described in (8), for all $x \in \mathbb{X}$,

$$\bar{w} = \underline{w} = w(x) = w^{\tilde{\phi}}(x) = \lim_{\alpha\uparrow 1}(1-\alpha)v_\alpha(x) = w^*. \tag{42}$$

Let us choose any stationary policy ϕ such that inequalities (8) and (9) hold with the function u defined in (21). Since $\bar{w} = \underline{w}$, according to Theorem 3, such a stationary policy exists. Theorem 1 implies that the stationary policy ϕ satisfies (10), and Schäl [26, Proposition 1.3] (see (8)) implies that (42) holds with $\tilde{\phi} = \phi$.

In addition, (42) with $\tilde{\phi} = \phi$ implies that for all $x \in \mathbb{X}$,

$$w^\phi(x) = \lim_{\alpha\uparrow 1}(1-\alpha)m_\alpha = \lim_{\alpha\uparrow 1}(1-\alpha)(v_\alpha(x) - u_\alpha(x)) = \lim_{\alpha\uparrow 1}(1-\alpha)v_\alpha(x),$$

where the last equality follows from Assumption (B). Thus, for all $x \in \mathbb{X}$,

$$w^\phi(x) = \limsup_{n\to\infty} \frac{1}{n}v_{n,1}^\phi(x) \geq \limsup_{\alpha\uparrow 1}(1-\alpha)v_\alpha^\phi(x) \geq \liminf_{\alpha\uparrow 1}(1-\alpha)v_\alpha^\phi(x)$$

$$\geq \lim_{\alpha\uparrow 1}(1-\alpha)v_\alpha(x) = w^\phi(x),$$

where the first inequality follows from the Tauberian theorem (see Sennott [27, §A.4] or [28, Proposition 5.7]), and the last inequality follows from $v_\alpha^\phi(x) \geq v_\alpha(x)$ and the existence of the limit. So, we have the existence of $\lim_{\alpha\uparrow 1}(1-\alpha)v_\alpha^\phi(x)$. Thus the Karamata Tauberian Theorem (Sennott [27, §A.4] or Sennott [28, Proposition 5.7]) implies $w^\phi(x) = \lim_{n\to\infty}(1/n)v_{n,1}^\phi(x)$. $\square$

COROLLARY 2. *Under Assumptions* (W^*) *and* (B), *the conclusions of Theorems* 3 *and* 4 *remain correct if the function u is substituted with the function $\tilde{u}$ defined in* (38).

PROOF. As shown in the proof of Theorem 4, there exists a stationary policy $\tilde{\phi}$ satisfying (41). The function $\tilde{u}$ is nonnegative, lower semicontinuous, and takes finite values. Thus, both Schäl [26, Proposition 1.3] (see (8)) and Theorem 1 can be applied to this function. The proof of statements (a)–(d) of Theorem 3 uses just these properties of u. Statement (e) follows from Lemma 6, whose proof remains unchanged if u is replaced by $\tilde{u}$. $\square$

6. Approximation of average cost optimal strategies by α-discount optimal strategies. For a family of sets $\{\mathrm{Gr}(A_\alpha)\}_{\alpha \in (0,1)}$, $x \in \mathbb{X}$ considered in Theorem 2, we pay our attention to its upper topological limit

$$\overline{\mathrm{Lim}}_{\alpha \uparrow 1} \mathrm{Gr}(A_\alpha) = \left\{ (x,a) \in \mathbb{X} \times \mathbb{A} : \begin{array}{c} \exists \alpha_n \uparrow 1, \ n \to +\infty, \ \exists (x_n, a_n) \in \mathrm{Gr}(A_{\alpha_n}), \ n \geq 1 \\ \text{such that } (x,a) = \lim_{n \to +\infty} (x_n, a_n) \end{array} \right\},$$

defined, for example, in Zgurovsky et al. [32, Chapter 1, p. 3]. Let us set

$$A^{\mathrm{app}}(x) := \left\{ a \in A^*(x) : (x,a) \in \overline{\mathrm{Lim}}_{\alpha \uparrow 1} \mathrm{Gr}(A_\alpha) \right\}, \quad x \in \mathbb{X}.$$

THEOREM 5. *Under Assumptions* (W^*) *and* $(\underline{B})$, *the graph* $\mathrm{Gr}(A^{\mathrm{app}})$ *is a Borel subset of* $\mathrm{Gr}(A^*)$, *and for each $x \in \mathbb{X}$, the set $A^{\mathrm{app}}(x)$ is nonempty and compact. Furthermore, there exists a stationary policy ϕ^{app} such that $\phi^{\mathrm{app}}(x) \in A^{\mathrm{app}}(x)$ for all $x \in X$, and any such policy is average cost optimal.*

PROOF. Let us fix an arbitrary $x \in \mathbb{X}$. From (21) (the definition of u), there exists $\{y_n, \alpha_n\}_{n \geq 1} \subseteq \mathbb{X} \times (0,1)$ such that $y_n \to x$, $\alpha_n \uparrow 1$, $u_{\alpha_n}(y_n) \to u(x)$, $n \to +\infty$.

Let us choose an arbitrary $\varepsilon^* > 0$ and $b_n \in A_{\alpha_n}(y_n)$, $n \geq 1$. Since $\bar{w} = \lim\sup_{\alpha \uparrow 1}(1-\alpha)m_\alpha$, there exists $N \geq 1$ such that $u(x) + \varepsilon^*/2 \geq u_{\alpha_n}(y_n)$ and $\bar{w} + \varepsilon^*/2 \geq (1-\alpha_n)m_{\alpha_n}$ for all $n \geq N$.

By definition of the sets $A_\alpha(\cdot)$ for each $n \geq N$,

$$(1-\alpha_n)m_{\alpha_n} + u_{\alpha_n}(y_n) = c(y_n, b_n) + \alpha_n \int_{\mathbb{X}} u_{\alpha_n}(y)q(dy \mid y_n, b_n) = \eta_{u_{\alpha_n}}^{\alpha_n}(y_n, b_n).$$

Thus, for all $n \geq N$,

$$\bar{w} + \varepsilon^* + u(x) \geq \eta_{u_{\alpha_n}}^{\alpha_n}(y_n, b_n) \geq \eta_{U_{\alpha_n}}^{\alpha_n}(y_n, b_n) \geq \eta_{\underline{u}_{\alpha_n}}^{\alpha_n}(y_n, b_n) \geq c(y_n, b_n).$$

Therefore, because of Assumption $(W^*(\mathrm{ii}))$, the sequence $\{b_n\}_{n \geq 1}$ has a subsequence $\{b_{n_k}\}_{k \geq 1}$ such that $b_{n_k} \to a$, as $k \to +\infty$ for some $a \in A(x)$. Thus $(x,a) \in \overline{\mathrm{Lim}}_{\alpha \uparrow 1}\mathrm{Gr}(A_\alpha)$.

Let us prove that $(x,a) \in \mathrm{Gr}(A^*)$. Indeed, as $\alpha_{n_k}\underline{u}_{\alpha_{n_k}}(\cdot) \uparrow u(\cdot)$, $k \to +\infty$, then due to Lemma 4 and Corollary 1,

$$\liminf_{k \to +\infty} \alpha_{n_k} \int_{\mathbb{X}} \underline{u}_{\alpha_{n_k}}(x)q(dy \mid y_{n_k}, b_{n_k}) \geq \int_{\mathbb{X}} u(x)q(dy \mid x, a).$$

Thus, by Lemma 3, $\bar{w} + \varepsilon^* + u(x) \geq \eta_u^1(x,a)$, and this is true for any $\varepsilon^* > 0$. This implies $\bar{w} + u(x) \geq \eta_u^1(x,a)$. This inequality means that $(x,a) \in \mathrm{Gr}(A^*)$ and $A^{\mathrm{app}}(x) \neq \varnothing$, since $(x,a) \in \overline{\mathrm{Lim}}_{\alpha \uparrow 1}\mathrm{Gr}(A_\alpha)$. The set $A^{\mathrm{app}}(x)$ is compact because $\overline{\mathrm{Lim}}_{\alpha \uparrow 1}\mathrm{Gr}(A_\alpha)$ is closed (see Zgurovsky et al. [32, Chapter 1, p. 3] and Theorem 3(b)). The second statement of the theorem follows from the Arsenin-Kunugui Theorem. $\square$

COROLLARY 3. *Under Assumptions* (W^*) *and* $(\underline{B})$, *for any stationary average cost optimal policy ϕ^{app} such that $\phi^{\mathrm{app}}(x) \in A^{\mathrm{app}}(x)$ for all $x \in \mathbb{X}$, for every $x \in \mathbb{X}$, there exist $\alpha_n(x) \uparrow 1$ and $y_n(x) \to x$ as $n \to +\infty$ such that $a_n(x) \in A_{\alpha_n(x)}(y_n(x))$, $n \geq 1$, and $\phi^{\mathrm{app}}(x) = \lim_{n \to +\infty} a_n(x)$.*

PROOF. Following Theorem 5, consider a stationary average cost optimal policy ϕ^{app} such that $\phi^{\mathrm{app}}(x) \in A^{\mathrm{app}}(x)$ for all $x \in X$. Furthermore, since $A^{\mathrm{app}}(x) \subseteq A^*(x)$ for all $x \in \mathbb{X}$, any such a policy is optimal. Let us fix an arbitrary $x \in \mathbb{X}$. By definition of $A^{\mathrm{app}}(x)$, we have that $(x, \phi^{\mathrm{app}}(x)) \in \overline{\mathrm{Lim}}_{\alpha \uparrow 1}\mathrm{Gr}(A_\alpha)$. Then, there exist $\alpha_n(x) \uparrow 1$, $n \to +\infty$, and $(y_n(x), a_n(x)) \in \mathrm{Gr}(A_{\alpha_n})$, $n \geq 1$ such that $(x, \phi^{\mathrm{app}}(x)) = \lim_{n \to +\infty}(y_n(x), a_n(x))$, i.e., $\phi^{\mathrm{app}}(x) = \lim_{n \to +\infty} a_n(x)$, where $a_n(x) \in A_{\alpha_n(x)}(y_n(x))$, $n \geq 1$, $\alpha_n(x) \uparrow 1$, and $y_n(x) \to x$ as $n \to +\infty$. $\square$

We remark that if we replace in (26) the function u with the function $\tilde{u}$ defined in (38), Theorem 5 and Corollary 3 remain correct.

Let us set

$$X_\alpha := \{x \in \mathbb{X}: v_\alpha(x) = m_\alpha\}, \quad \alpha \in [0, 1)$$

under Assumptions (G), $m_\alpha < \infty$. If Assumptions (G) and (Wu) hold, then Theorem 2 implies that X_α is a compact set for each $\alpha \in [0, 1)$. This fact is useful to establish the validity of Assumptions (G); see Feinberg and Lewis [14, Lemma 5.1] and references therein.

THEOREM 6. *Let Assumptions* (G) *and* (Wu) *hold. Then, there exists a compact set $\mathcal{K} \subseteq \mathbb{X}$ such that $X_\alpha \subseteq \mathcal{K}$ for each $\alpha \in [0, 1)$.*

PROOF. From Assumption (G) and Theorem 2, we have that for each $\alpha \in [0, 1)$,

$$\varnothing \neq X_\alpha = \{x \in \mathbb{X}: u_\alpha(x) = 0\} = \mathscr{D}_{u_\alpha}(0) \subseteq \mathscr{D}_{U_\alpha}(0) \subseteq \mathscr{D}_{\underline{u}_\alpha}(0) \subseteq \mathscr{D}_{\underline{u}_0}(0).$$

In virtue of Lemma 6, we have that $\underline{u}_0 \colon \mathbb{X} \to [0, +\infty)$ is inf-compact function on $\mathbb{X}$. Setting $\mathcal{K} = \mathscr{D}_{\underline{u}_0}(0)$, we obtain the statement of the theorem. $\square$

7. Illustrative example. The following example is from Hernández-Lerma [18]. Let

$$x_{n+1} = \gamma x_n + \beta a_n + \xi_n, \quad n = 0, 1, \dots,$$

and

$$c(x, a) = qx^2 + ra^2,$$

where (a) q and r are positive constants, γ and β are two constants satisfying $\gamma\beta > 0$, and (b) ξ_n are independent and identically distributed (iid) random variables with zero mean, finite variance, and continuous density.

This problem is solved in Hernández-Lerma [18], where a stationary average cost optimal policy is computed. This problem corresponds to an MDP with $\mathbb{X} = \mathbb{A} = \mathbb{R}$ and with setwise continuous transition probabilities. However, if ξ_n do not have a density, the transition probability may not be setwise continuous, but they are weakly continuous; see Feinberg and Lewis [13, p. 48] for details. If ξ_n are arbitrary iid random variables with zero mean and finite variance, this problem satisfies Assumption (Wu), and similarly to the case when there are densities, it satisfies Assumption (**B**). Thus, Theorem 4 can be applied. The optimal policy provided in Hernández-Lerma [18] is also optimal when ξ_n may not have a density.

Appendix. Proof of Lemma 4. First, we prove the lemma for uniformly bounded above functions h_n. Let $h_n(s) \leq K < \infty$ for all $n = 1, 2, \dots$ and all $s \in S$. For $n = 1, 2, \dots$ and $s \in S$, define

$$H_n(s) = \inf_{m \geq n} h_m(s) \qquad \text{and} \qquad \underline{h}_n(s) = \liminf_{s' \to s} H_n(s').$$

The functions $\underline{h}_n \colon S \to [0, +\infty)$, $n = 1, 2, \dots$, are lower semicontinuous; see, for example, Feinberg and Lewis [14, Lemma 3.1]). In addition, for $s \in S$,

$$\underline{h}_n(s) \uparrow \underline{h}(s) \quad \text{as } n \to \infty. \tag{43}$$

Weak convergence of $\{\mu_n\}_{n \geq 1}$ to μ is equivalent to

$$\liminf_{n \to +\infty} \mu_n(A) \geq \mu(A) \quad \text{for all } A \in \mathscr{O}, \tag{44}$$

where $\mathscr{O}$ is the family of all open subsets of the space S; Billingsley [5, Theorem 2.1].

Fix an arbitrary $t > 0$. By (43), if $\underline{h}(s) > t$, then $\underline{h}_n(s) > t$, $n = 1, 2, \dots$, and

$$\{s \in S: \underline{h}(s) > t\} = \bigcup_{n \geq 1} S_n, \tag{45}$$

where

$$S_n = \{s \in S: \underline{h}_n(s) > t\}, \quad n = 1, 2, \dots,$$

are open sets, since the functions $\underline{h}_n\colon S \to \mathbb{R}_+$ are lower semicontinuous. In addition,

$$S_n \subseteq S_{n+1}, \quad n = 1, 2, \ldots. \tag{46}$$

Thus

$$\mu(\{s \in S \colon \underline{h}(s) > t\}) = \lim_{n \to +\infty} \mu(S_n) \leq \lim_{n \to +\infty} \liminf_{m \to +\infty} \mu_m(S_n)$$
$$\leq \limsup_{n \to +\infty} \liminf_{m \to +\infty} \mu_m(S_m) = \liminf_{n \to +\infty} \mu_n(S_n) = \liminf_{n \to +\infty} \mu_n(\{s \in S \colon \underline{h}_n(s) > t\}),$$

where the first equality follows from (46) and (45), the first inequality follows from (44), and the second inequality follows from (46).

Thus Serfozo [29, Lemma 2.1] yields

$$\int_S \underline{h}(s)\,\mu(ds) \leq \liminf_{n \to +\infty} \int_S \underline{h}_n(s)\,\mu_n(ds) \leq \liminf_{n \to +\infty} \int_S h_n(s)\,\mu_n(ds),$$

where the second inequality is fulfilled due to

$$\underline{h}_n(s) \leq H_n(s) \leq h_n(s), \quad s \in S, \;\; n = 1, 2, \ldots.$$

Second, consider a sequence $\{h_n\}_{n \geq 1}$ of measurable nonnegative $\bar{\mathbb{R}}$-valued functions on S. For $\lambda > 0$, set $h_n^\lambda(s) := \min\{h_n(s), \lambda\}$, $s \in S$, $n = 1, 2, \ldots$. Since the functions h_n^λ are uniformly bounded above,

$$\int_S \underline{h}^\lambda(s)\,\mu(ds) \leq \liminf_{n \to +\infty} \int_S h_n^\lambda(s)\,\mu_n(ds) \leq \liminf_{n \to +\infty} \int_S h_n(s)\,\mu_n(ds),$$

where $\underline{h}^\lambda(s) = \liminf_{n \to +\infty,\, s' \to s} h_n^\lambda(s')$, $\lambda > 0$, $s \in S$.

Then, using Fatou's lemma,

$$\int_S \underline{h}(s)\,\mu(ds) \leq \liminf_{\lambda \to +\infty} \int_S \underline{h}^\lambda(s)\,\mu(ds).$$

References

[1] Arapostathis A, Borkar VS, Fernandez-Gaucherand E, Ghosh MK, Marcus SI (1993) Discrete-time controlled Markov processes with average cost criterion: A survey. *SIAM J. Control Optim.* 31(2):282–344.
[2] Bather J (1973) Optimal decision procedures for finite Markov chains. Part I: Examples. *Adv. Appl. Probab.* 5(2):328–339.
[3] Berge E (1963) *Topological Spaces* (Macmillan, New York).
[4] Bertsekas DP, Shreve SE (1996) *Stochastic Optimal Control: The Discrete-Time Case* (Athena Scientific, Belmont, MA).
[5] Billingsley P (1968) *Convergence of Probability Measures* (John Wiley & Sons, New York).
[6] Blackwell D (1962) Discrete dynamic programming. *Ann. Math. Statist.* 33(2):719–726.
[7] Cavazos-Cadena R (1991) A counterexample on the optimality equation in Markov decision chains with the average cost criterion. *Systems and Control Lett.* 16(5):387–392.
[8] Chen RC, Feinberg EA (2010) Compactness of the space of non-randomized policies in countable-state sequential decision processes. *Math. Methods Oper. Res.* 71(2):307–323.
[9] Chitashvili RY (1975) A controlled finite Markov chain with an arbitrary set of decisions. *Theor. Probab. Appl.* 20(4):839–847.
[10] Derman C (1962) On sequential decisions and Markov chains. *Management Sci.* 9(1):16–24.
[11] Dynkin EB, Yushkevich AA (1979) *Controlled Markov Processes* (Springer-Verlag, New York).
[12] Feinberg EA (1980) An ϵ-optimal control of a finite Markov chain. *Theoret. Probab. Appl.* 25(1):70–81.
[13] Feinberg EA, Lewis ME (2004) Optimality of four-threshold policies in inventory systems with customer returns and borrowing/storage options. *Probab. Engrg. Inform. Sci.* 19(1):45–71.
[14] Feinberg EA, Lewis ME (2007) Optimality inequalities for average cost Markov decision processes and the stochastic cash balance problem. *Math. Oper. Res.* 32(4):769–783.
[15] Feinberg EA, Kasyanov PO, Zadoianchuk NV (2012) Berge's theorem for noncompact image sets. *J. Math. Anal. Appl.* Forthcoming.
[16] Feinberg EA, Kasyanov PO, Zadoianchuk NV (2012) Fatou's lemma for weakly converging probabilities. http://arXiv:1206.4073v1.
[17] Gubenko LG, Shtatland ES (1975) On controlled, discrete-time Markov decision processes. *Theory Probab. Math. Statist.* 7:47–61.
[18] Hernández-Lerma O (1991) Averege optimality in dynamic programming on Borel spaces—Unbounded costs and controls. *Systems and Control Lett.* 17(3):237–242.
[19] Hernández-Lerma O, Lasserre JB (1996) *Discrete-Time Markov Control Processes: Basic Optimality Criteria* (Springer-Verlag, New York).
[20] Hernández-Lerma O, Lasserre JB (2000) Fatou's lemma and Lebesgue's convergence theorem for measures. *J. Appl. Math. Stochastic Anal.* 13(2):137–146.
[21] Kechris AS (1995) *Classical Descriptive Set Theory* (Springer-Verlag, New York).
[22] Luque-Vásquez F, Hernández-Lerma O (1995) A counterexample on the semicontinuity of minima. *Proc. Amer. Math. Soc.* 123(10):3175–3176.

[23] Ross SM (1968) Non-discounted denumerable Markovian decision model. *Ann. Math. Statist.* 39(2):412–424.

[24] Ross SM (1968) Arbitrary state Markovian decision processes. *Ann. Math. Statist.* 39(6):2118–2122.

[25] Ross SM (1971) On the nonexistence of ϵ-optimal randomized stationary policies in average cost Markov decision models. *Ann. Math. Statist.* 42(5):1767–1768.

[26] Schäl M (1993) Average optimality in dynamic programming with general state space. *Math. Oper. Res*. 18(1):163–172.

[27] Sennott LI (1999) *Stochastic Dynamic Programming and the Control of Queueing Systems* (John Wiley & Sons, New York).

[28] Sennott LI (2002) Average reward optimization theory for denumerable state spaces. Feinberg EA, Shwartz A, eds. *Handbook of Markov Decision Processes. Methods and Applications* (Kluwer, Boston) 153–172.

[29] Serfozo R (1982) Convergence of Lebesgue integrals with varying measures. *Sankhya: The Indian Journal of Statistics (Series A)* 44(3):380–402.

[30] Taylor III, HM (1965) Markovian sequential replacement processes. *Ann. Math. Statist.* 36(6):1677–1694.

[31] Viskov OV, Shiryaev AN (1964) On controls which reduce to optimal stationary regimes. *Trudy Mat. Inst. Steklov*. 71:35–45. [In Russian; English translation: Report Number FTD-HT-67-69, National Technical Information Service, U.S. Department of Commerce.]

[32] Zgurovsky MZ, Mel'nik VS, Kasyanov PO (2011) *Evolution Inclusions and Variation Inequalities for Earth Data Processing I.* (Springer, Berlin).