

Risk-Averse Control of Undiscounted Transient Markov Models

Özlem Çavuş* Andrzej Ruszczyński†

December 1, 2012

Abstract

We use Markov risk measures to formulate a risk-averse version of the undiscounted total cost problem for a transient controlled Markov process. Using the new concept of a multikernel, we derive conditions for a system to be risk-transient, that is, to have finite risk over an infinite time horizon. We derive risk-averse dynamic programming equations satisfied by the optimal policy and we describe methods for solving these equations. We illustrate the results on an optimal stopping problem and an organ transplant problem.

Keywords: Dynamic Risk Measures; Markov Risk Measures; Multikernels; Stochastic Shortest Path; Optimal Stopping; Randomized Policy

1 Introduction

The optimal control problem for transient Markov processes is a classical model in Operations Research (see Veinott [48], Pliska [33], Bertsekas and Tsitsiklis [7], Hernandez-Lerma and Lasserre [18], and the references therein). The research is focused on the expected total undiscounted cost model, with increased state and control space generality.

Our objective is to consider a risk-averse model. So far, risk-averse problems for transient Markov models were based on the arrival probability criteria (see, e.g., Nie and Wu [25] and Ohtsubo [27]) and utility functions (see Denardo and Rothblum [11] and Patek [31]). We plan to use the recent theory of dynamic risk measures (see Scandolo [43], Riedel [35], Ruszczyński and Shapiro [40, 42], Cheridito, Delbaen and Kupper [8], Artzner et. al. [3], Klöppel and Schweizer [22], Pflug and Römisch [32], and the references therein) to develop and solve new risk-averse formulations of the stochastic optimal control problem for transient Markov models. Specific examples of such models are stochastic shortest path problems (Bertsekas and Tsitsiklis [7]) and optimal stopping problems (*cf.* Çinlar [10], Dynkin and Yushkevich [12, 13], Puterman [34]).

Some applications of stochastic shortest path problems concerned with expected performance criteria are given in the survey paper by White [50] and the references therein. However, in many practical problems, the expected values may not be appropriate to measure performance, because they implicitly assume that the decision maker is risk-neutral. Below, we provide examples of such real-life problems which were modeled before as a discrete-time Markov decision process with expected value as the objective function. Alagoz et. al. [1] suggest a discounted, infinite horizon, and absorbing Markov decision process model to find the optimal time of liver transplant for a risk-neutral patient under the assumption that the liver is transferred from a living donor. However, referring to Chew and Ho [9], they state that the risk-neutrality of the patient is not a

*Bilkent University, Department of Industrial Engineering, Ankara, Turkey, ozlem.cavus@bilkent.edu.tr

†Rutgers University, Department of Management Science and Information Systems, Piscataway, NJ 08854, USA, rusz@rutgers.edu

realistic assumption. Kurt and Kharoufe [24] propose a discounted, infinite horizon Markov decision process model to optimal replacement time of a system for a system under Markovian deterioration and Markovian environment. So and Thomas [46] employ a discrete time Markov decision process to model profitability of credit cards.

Our theory of risk-averse control problems for transient models applies to these and many other models. Our results complement and extend the results of Ruszczyński [39], where infinite-horizon *discounted* models were considered. We consider *undiscounted* models for *transient* Markov systems. The paper is organized as follows.

In section 2 we quickly review some basic concepts of controlled Markov models. In section 3 we adapt and extend our earlier theory of Markov risk measures. In section 4 we introduce and analyze the concept of a *multikernel* (a multivalued kernel), which is essential for our theory. General assumptions and technical issues associated with measurability of decision rules are discussed in section 5. Section 6 is devoted to the analysis of a finite horizon model. The main model with infinite horizon and dynamic risk measures is analyzed in section 7. We introduce in it the concept of a *risk-transient model* and develop equations for evaluating policies in such models. In section 8 we derive risk-averse versions of dynamic programming equations for risk transient models. Section 9 compares randomized and deterministic policies. Finally, section 10 illustrates our results on risk-averse versions of an optimal stopping problem of Karlin [21] and of the organ transplant problem of Alagoz *et al.* [1].

2 Controlled Markov Processes

We quickly review the main concepts of controlled Markov models and we introduce relevant notation (for details, see [14, 17, 18]). Let \mathcal{X} be a state space, and \mathcal{U} a control space. We assume that \mathcal{X} and \mathcal{U} are Borel spaces (Borel subsets of Polish spaces). A control set is a measurable multifunction $U : \mathcal{X} \rightrightarrows \mathcal{U}$; for each state $x \in \mathcal{X}$ the set $U(x) \subseteq \mathcal{U}$ is a nonempty set of possible controls at x . A controlled transition kernel Q is a measurable mapping from the graph of U to the set $\mathcal{P}(\mathcal{X})$ of probability measures on \mathcal{X} (equipped with the topology of weak convergence).

The cost of transition from x to y , when control u is applied, is represented by $c(x, u, y)$, where $c : \mathcal{X} \times \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$. Only $u \in U(x)$ and those $y \in \mathcal{X}$ to which transition is possible matter here, but it is convenient to consider the function $c(\cdot, \cdot, \cdot)$ as defined on the product space.

A *stationary controlled Markov process* is defined by a state space \mathcal{X} , a control space \mathcal{U} , a control set U , a controlled transition kernel Q , and a cost function c .

For $t = 1, 2, \dots$ we define the space of state and control histories up to time t as $\mathcal{H}_t = \text{graph}(U)^t \times \mathcal{X}$. Each history is a sequence $h_t = (x_1, u_1, \dots, x_{t-1}, u_{t-1}, x_t) \in \mathcal{H}_t$.

We denote by $\mathcal{P}(\mathcal{U})$ and $\mathcal{P}(U(x))$ the sets of probability measures on \mathcal{U} and set of probability measures on $U(x)$. A *randomized policy* is a sequence of measurable functions $\pi_t : \mathcal{H}_t \rightarrow \mathcal{P}(\mathcal{U})$, $t = 1, 2, \dots$, such that $\pi_t(h_t) \in \mathcal{P}(U(x_t))$ for all $h_t \in \mathcal{H}_t$. In words, the distribution of the control u_t is supported on a subset of the set of feasible controls $U(x_t)$. A *Markov policy* is a sequence of measurable functions $\pi_t : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{U})$, $t = 1, 2, \dots$, such that $\pi_t(x) \in \mathcal{P}(U(x))$ for all $x \in \mathcal{X}$. The function $\pi_t(\cdot)$ is called the *decision rule* at time t . A Markov policy is *stationary* if there exists a function $\pi : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{U})$ such that $\pi_t(x) = \pi(x)$, for all $t = 1, 2, \dots$ and all $x \in \mathcal{X}$. Such a policy and the corresponding decision rule are called *deterministic*, if for every $x \in \mathcal{X}$ there exists $u(x) \in U(x)$ such that the measure $\pi(x)$ is supported on $\{u(x)\}$. In this paper we focus on deterministic policies.

Consider the canonical sample space $\Omega = \mathcal{X}^\infty$ with the product σ -algebra \mathcal{F} . Let P_1 be the initial distribution of the state $x_1 \in \mathcal{X}$. Suppose we are given a deterministic policy $\Pi = \{\pi_t\}_{t=1}^\infty$. The Ionescu Tulcea theorem (see, *e. g.*, [6]) states that there exists a unique probability measure P^Π on (Ω, \mathcal{F}) such that

for every measurable set $B \subset \mathcal{X}$ and all $h_t \in \mathcal{H}_t$, $t = 1, 2, \dots$,

$$\begin{aligned} P^\Pi(x_1 \in B) &= P_1(B); \\ P^\Pi(x_{t+1} \in B | h_t) &= Q(B | x_t, \pi_t(h_t)). \end{aligned}$$

To simplify our notation, from now on we assume that the initial state x_1 is fixed. It will be obvious how to modify our results for a random initial state. For a stationary decision rule π , we write Q^π to denote the corresponding transition kernel.

Our interest is in *transient* Markov models. We assume that some *absorbing state* $x_A \in \mathcal{X}$ exists, such that $Q(\{x_A\} | x_A, u) = 1$ and $c(x_A, u, x_A) = 0$ for all $u \in U(x_A)$. Thus, after the absorbing state is reached, no further costs are incurred.¹ To analyze such Markov models, it is convenient to consider the effective state space $\tilde{\mathcal{X}} = \mathcal{X} \setminus \{x_A\}$, and the effective controlled substochastic kernel \tilde{Q} whose arguments are restricted to $\tilde{\mathcal{X}}$ and whose values are nonnegative measures on $\tilde{\mathcal{X}}$, so that $\tilde{Q}(B | x, u) = Q(B | x, u)$, for all Borel sets $B \subset \tilde{\mathcal{X}}$, all $x \in \tilde{\mathcal{X}}$, and all $u \in U(x)$.

Our point of departure is the *expected total cost problem*, which is to find a policy $\Pi = \{\pi_t\}_{t=1}^\infty$ so as to minimize the expected cost until absorption:

$$\min_{\Pi} \mathbb{E}^\Pi \left[\sum_{t=1}^{\infty} c(x_t, u_t, x_{t+1}) \right].$$

Here $\mathbb{E}^\Pi[\cdot]$ denotes the expected value with respect to the measure P^Π . Under appropriate assumptions, the problem has a solution in form of a stationary Markov policy (see, e.g., [18, sec. 9.6]). The optimal policy can be found by solving appropriate dynamic programming equations.

Our intention is to introduce risk aversion to the problem, and to replace the expected value operator by a dynamic risk measure. We do not assume that the costs are nonnegative, and thus our approach applies also, among others, to stochastic longest path problems and optimal stopping problems with positive rewards.

3 Markov Risk Measures

Suppose T is a fixed time horizon. Each policy $\Pi = \{\pi_1, \pi_2, \dots\}$ results in a cost sequence $Z_t = c(x_{t-1}, u_{t-1}, x_t)$, $t = 2, \dots, T+1$ on the probability space $(\Omega, \mathcal{F}, P^\Pi)$. We define the σ -subalgebras \mathcal{F}_t on \mathcal{X}^t , and vector spaces \mathcal{Z}_t^Π of \mathcal{F}_t -measurable random variables on Ω , $t = 1, \dots, T$.

To evaluate risk of this sequence we use a dynamic time-consistent risk measure of the following form:

$$\begin{aligned} J_T(\Pi, x_1) &= \rho_1^\Pi \left(c(x_1, \pi_1(x_1), x_2) + \rho_2^\Pi \left(c(x_2, \pi_2(x_2), x_3) + \dots \right. \right. \\ &\quad \left. \left. + \rho_{T-1}^\Pi (c(x_{T-1}, \pi_{T-1}(x_{T-1}), x_T) + \rho_T^\Pi (c(x_T, \pi_T(x_T), x_{T+1}))) \dots \right) \right). \end{aligned} \quad (1)$$

Here, $\rho_t^\Pi : \mathcal{Z}_{t+1}^\Pi \rightarrow \mathcal{Z}_t^\Pi$, $t = 1, \dots, T$, are one-step conditional risk measures. Ruszczyński [39, sec. 3] derives the nested formulation (1) from general properties of monotonicity and time-consistency of dynamic measures of risk.

It is convenient to introduce vector spaces $\mathcal{Z}_{t,\theta}^\Pi = \mathcal{Z}_t^\Pi \times \mathcal{Z}_{t+1}^\Pi \times \dots \times \mathcal{Z}_\theta^\Pi$, where $1 \leq t \leq \theta \leq T+1$ and the conditional risk measures $\rho_{t,\theta}^\Pi : \mathcal{Z}_{t,\theta}^\Pi \rightarrow \mathcal{Z}_t^\Pi$ defined as follows:

$$\rho_{t,\theta}^\Pi(Z_t, \dots, Z_\theta) = Z_t + \rho_t^\Pi \left(Z_{t+1} + \rho_{t+1}^\Pi (Z_{t+2} + \dots + \rho_{\theta-1}^\Pi(Z_\theta) \dots) \right). \quad (2)$$

¹The case of a larger class of absorbing states easily reduces to the case of one absorbing state.

As indicated in [39], the fundamental difficulty of formulation (1) is that at time t the value of $\rho_t^\Pi(\cdot)$ is \mathcal{F}_t -measurable and is allowed to depend on the entire history h_t of the process. In order to overcome this difficulty, in [39, sec. 4] a new construction of a one-step conditional measure of risk was introduced. Its arguments are functions on the state space \mathcal{X} , rather than on the probability space Ω . We adapt this construction to our case, with a slightly more general form of the cost function.

Let $\mathcal{V} = \mathcal{L}_p(\mathcal{X}, \mathcal{B}, P_0)$, where \mathcal{B} is the σ -field of Borel sets on \mathcal{X} , P_0 is some reference probability measure on \mathcal{X} , and $p \in [1, \infty)$. It is convenient to think of the dual space \mathcal{V}' as the space of signed measures m on $(\mathcal{X}, \mathcal{B})$, which are absolutely continuous with respect to P_0 , with densities (Radon–Nikodym derivatives) lying in the space $\mathcal{L}_q(\mathcal{X}, \mathcal{B}, P_0)$, where $1/p + 1/q = 1$. We make the following general assumption.

G0. For all $x \in \mathcal{X}$ and $u \in U(x)$ the probability measure $Q(x, u)$ is an element of \mathcal{V}' .

In the case of finite state and control spaces P_0 may be the uniform measure; in other cases P_0 should be chosen in such a way that condition (G0) is satisfied. The existence of the measure P_0 is essential for the pairing of \mathcal{V} and its dual space \mathcal{V}' , as discussed below.

We consider the set of probability measures in \mathcal{V}' :

$$\mathcal{M} = \{m \in \mathcal{V}' : m(\mathcal{X}) = 1, m \geq 0\}.$$

We also assume that the spaces \mathcal{V} and \mathcal{V}' are endowed with topologies that make them paired topological vector spaces with the bilinear form

$$\langle \varphi, m \rangle = \int_{\mathcal{X}} \varphi(y) m(dy), \quad \varphi \in \mathcal{V}, \quad m \in \mathcal{V}'.$$

The space \mathcal{V}' (and thus \mathcal{M}) will be endowed with the weak* topology. We may endow \mathcal{V} with the strong (norm) topology, or with the weak topology.

Definition 3.1 A measurable function $\sigma : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}$ is a risk transition mapping if for every $x \in \mathcal{X}$ and every $m \in \mathcal{M}$, the function $\varphi \mapsto \sigma(\varphi, x, m)$ is a coherent measure of risk on \mathcal{V} .

Recall that $\sigma(\cdot)$ is a coherent measure of risk on \mathcal{V} (we skip the other two arguments for brevity), if (see [2])

- A1.** $\sigma(\alpha\varphi + (1-\alpha)\psi) \leq \alpha\sigma(\varphi) + (1-\alpha)\sigma(\psi)$, $\forall \alpha \in (0, 1)$, $\varphi, \psi \in \mathcal{V}$;
- A2.** If $\varphi \leq \psi$ then $\sigma(\varphi) \leq \sigma(\psi)$, $\forall \varphi, \psi \in \mathcal{V}$;
- A3.** $\sigma(a + \varphi) = a + \sigma(\varphi)$, $\forall \varphi \in \mathcal{V}$, $a \in \mathbb{R}$;
- A4.** $\sigma(\beta\varphi) = \beta\sigma(\varphi)$, $\forall \varphi \in \mathcal{V}$, $\beta \geq 0$.

Example 3.1 Consider the first-order mean–semideviation risk measure analyzed by Ogryczak and Ruszczyński [28, 29], and Ruszczyński and Shapiro [41, Example 4.2], [42, Example 6.1]), but with the state and the underlying probability measure as its arguments. We define

$$\sigma(\varphi, x, m) = \langle \varphi, m \rangle + \kappa \langle (\varphi - \langle \varphi, m \rangle)_+, m \rangle, \quad (3)$$

where $\kappa \in [0, 1]$. We can verify directly that conditions (A1)–(A4) are satisfied. In a more general setting, $\kappa : \mathcal{X} \rightarrow [0, 1]$ may be a measurable function.

Example 3.2 Another important example is the Average Value at Risk (see, *inter alia*, Ogryczak and Ruszczyński [30, Sec. 4], Pflug and Römisich [32, Sec. 2.2.3, 3.3.4], Rockafellar and Uryasev [37], Ruszczyński and Shapiro [41, Example 4.3], [42, Example 6.2]), which has the following risk transition counterpart:

$$\sigma(\varphi, x, m) = \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha} \langle (\varphi - \eta)_+, m \rangle \right\}, \quad \alpha \in (0, 1).$$

Again, the conditions (A1)–(A4) can be verified directly. In a more general setting $\alpha : \mathcal{X} \rightarrow [\alpha_{\min}, \alpha_{\max}] \subset (0, 1)$ may be a measurable function.

We shall use the property of *law invariance* of a risk transition mapping. For a function $\varphi \in \mathcal{V}$ and a probability measure $\mu \in \mathcal{M}$ we can define the distribution function $F_\varphi^\mu : \mathbb{R} \rightarrow [0, 1]$ as follows

$$F_\varphi^\mu(\eta) = \mu\{y \in \mathcal{X} : \varphi(y) \leq \eta\}.$$

Definition 3.2 A risk transition mapping $\sigma : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}$ is law invariant, if for all $\varphi, \psi \in \mathcal{V}$ and all $\mu, \nu \in \mathcal{M}$ such that $F_\varphi^\mu \equiv F_\psi^\nu$, we have $\sigma(\varphi, x, \mu) = \sigma(\psi, x, \nu)$ for all $x \in \mathcal{X}$.

The concept of law invariance corresponds to a similar concept for coherent measures of risk, but here we additionally need to take into account the variability of the probability measure. The risk transition mappings of Examples 3.1 and 3.2 are law invariant.

The concept of law invariance is important in the context of Markov decision processes, where the model essentially defines the distribution of the state process, for every policy Π . It also greatly simplifies the analysis of specific problems, as illustrated in section 10.1.

Risk transition mappings allow for convenient formulation of risk-averse preferences for controlled Markov processes, where the cost is evaluated by formula (1). Consider a controlled Markov process $\{x_t\}$ with a deterministic Markov policy $\Pi = \{\pi_1, \pi_2, \dots\}$. For a fixed time t and a measurable function $g : \mathcal{X} \times \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$ the value of $Z_{t+1} = g(x_t, u_t, x_{t+1})$ is a random variable. We assume that g is *w-bounded*, that is,

$$|g(x, u, y)| \leq C(w(x) + w(y)), \quad \forall x \in \mathcal{X}, u \in U(x), y \in \mathcal{X},$$

for some constant $C > 0$ and for the some *weight (bounding) function* $w : \mathcal{X} \rightarrow [1, \infty)$, $w \in \mathcal{V}$ (see, [5, sec. 2.4], [18, sec. 7.2], and [49] for the role of weight functions in Markov decision processes). Then Z_{t+1} is an element of \mathcal{Z}_{t+1}^Π . Let $\rho_t^\Pi : \mathcal{Z}_{t+1}^\Pi \rightarrow \mathcal{Z}_t^\Pi$ be a family of conditional risk measures satisfying (A1)–(A4), for every deterministic policy Π . By definition, $\rho_t^\Pi(g(x_t, u_t, x_{t+1}))$ is an element of \mathcal{Z}_t^Π , that is, it is an \mathcal{F}_t -measurable function on (Ω, \mathcal{F}) . In the definition below, we restrict it to depend on the past only via the current state x_t .

Definition 3.3 A family of one-step conditional risk measures $\rho_t^\Pi : \mathcal{Z}_{t+1}^\Pi \rightarrow \mathcal{Z}_t^\Pi$ is a Markov risk measure with respect to the controlled Markov process $\{x_t\}$, if there exists a law invariant risk transition mapping $\sigma : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}$ such that for all *w-bounded measurable functions* $g : \mathcal{X} \times \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$ and for all feasible deterministic Markov policies Π we have

$$\rho_t^\Pi(g(x_t, \pi_t(x_t), x_{t+1})) = \sigma(g(x_t, \pi_t(x_t), \cdot), x_t, Q(x_t, \pi_t(x_t))), \quad \text{a.s.} \quad (4)$$

Observe that the right hand side of formula (4) is parametrized by x_t , and thus it defines a special \mathcal{F}_t -measurable function of ω , whose dependence on the past is carried only via the state x_t . The quantifier "a.s." means "almost surely with respect to the measure P^Π ."

4 Stochastic Multikernels

In order to analyze Markov measures of risk, we need to introduce the concept of a multikernel.

Definition 4.1 A multikernel is a measurable multifunction \mathfrak{M} from \mathcal{X} to the space of regular measures on $(\mathcal{X}, \mathcal{B}(X))$. It is stochastic, if its values are sets of probability measures. It is substochastic, if $0 \leq M(B|x) \leq 1$ for all $M \in \mathfrak{M}(x)$, $B \in \mathcal{B}(X)$, and $x \in \mathcal{X}$. It is convex (closed), if for all $x \in \mathcal{X}$ its value $\mathfrak{M}(x)$ is a convex (closed) set.

The concept of a multikernel is thus a multivalued generalization of the concept of a kernel. A measurable selector of a stochastic multikernel \mathfrak{M} is a stochastic kernel M such that $M(x) \in \mathfrak{M}(x)$ for all $x \in \mathcal{X}$. We symbolically write $M \prec \mathfrak{M}$ to indicate that M is a measurable selector of \mathfrak{M} .

Recall that a composition $M_1 M_2$ of (sub-) stochastic kernels M_1 and M_2 is given by the formula:

$$[M_1 M_2](B|x) = \int_{\mathcal{X}} M_2(B|y) M_1(dy|x), \quad B \in \mathcal{B}(\mathcal{X}), \quad x \in \mathcal{X}. \quad (5)$$

It is also a (sub-) stochastic kernel. Multikernels, in particular substochastic multikernels, can be composed in a similar fashion.

Definition 4.2 *If \mathfrak{M}_1 and \mathfrak{M}_2 are multikernels, then their composition $\mathfrak{M}_1 \mathfrak{M}_2$ is defined as follows:*

$$[\mathfrak{M}_1 \mathfrak{M}_2](B|x) = \left\{ [M_1 M_2](B|x) : M_i \prec \mathfrak{M}_i, \quad i = 1, 2 \right\}.$$

It follows from Definition 4.2, that a composition of (sub-) stochastic multikernels is a (sub-) stochastic multikernel. We may compose a substochastic multikernel \mathfrak{M} with itself several times, to obtain its ‘‘power’’:

$$(\mathfrak{M})^k = \underbrace{\mathfrak{M} \mathfrak{M} \dots \mathfrak{M}}_{k \text{ times}}.$$

Multikernels can be *added* by employing the Minkowski sum of their values:

$$[\mathfrak{M}_1 + \mathfrak{M}_2](x) = \mathfrak{M}_1(x) + \mathfrak{M}_2(x) = \left\{ \mu : \mu = \mu_1 + \mu_2, \quad \mu_i \in \mathfrak{M}_i(x), \quad i = 1, 2 \right\}, \quad x \in \mathcal{X}.$$

The sum of stochastic multikernels is a multikernel with non-negative values.

The concept of a multikernel and the composition operation arise in a natural way in the context of Markov risk measures. If $\sigma(\cdot, \cdot, \cdot)$ is a risk transition mapping, then the function $\sigma(\cdot, x, m)$ is lower semicontinuous for all $x \in \mathcal{X}$ and $m \in \mathcal{M}$ (see Ruszczyński and Shapiro [41, Proposition 3.1]). Then it follows from [41, Theorem 2.2] that for every $x \in \mathcal{X}$ and $m \in \mathcal{M}$ a closed convex set $\mathcal{A}(x, m) \subset \mathcal{M}$ exists, such that for all $\varphi \in \mathcal{V}$ we have

$$\sigma(\varphi, x, m) = \max_{\mu \in \mathcal{A}(x, m)} \langle \varphi, \mu \rangle. \quad (6)$$

In fact, we also have

$$\mathcal{A}(x, m) = \partial_{\varphi} \sigma(0, x, m), \quad (7)$$

that is, $\mathcal{A}(x, m)$ is the subdifferential of $\sigma(\cdot, x, m)$ at 0 (for the foundations of conjugate duality theory, see [36]). In many cases, the multifunction $\mathcal{A} : \mathcal{X} \times \mathcal{M} \rightrightarrows \mathcal{M}$ can be described analytically.

Example 4.1 For the mean-semideviation model of Example 3.1, following the derivations of [41, Example 4.2], we have

$$\mathcal{A}(x, m) = \left\{ \mu \in \mathcal{M} : \exists (h \in \mathcal{L}_{\infty}(\mathcal{X}, \mathcal{B}, P_0)) \frac{d\mu}{dm} = 1 + h - \langle h, m \rangle, \quad \|h\|_{\infty} \leq \kappa, \quad h \geq 0 \right\}. \quad (8)$$

Similar formulas can be derived for higher order measures.

Example 4.2 For the Conditional Average Value at Risk of Example 3.2, following the derivations of [41, Example 4.3], we obtain

$$\mathcal{A}(x, m) = \left\{ \mu \in \mathcal{M} : \frac{d\mu}{dm} \leq \frac{1}{\alpha} \right\}. \quad (9)$$

Consider formula (4) and suppose that $g(x_t, u_t, x_{t+1}) = v(x_{t+1})$ for some measurable w -bounded function $v : \mathcal{X} \rightarrow \mathbb{R}$. Using the representation (6) we can write it as follows:

$$\rho_t^\Pi(v(x_{t+1})) = \max_{\mu \in \mathcal{A}(x_t, Q(x_t, \pi_t(x_t)))} \int_{\mathcal{X}} v(y) M(dy), \quad \text{a.s.} \quad (10)$$

Suppose policy Π is stationary and $\pi_t = \pi$ for all t . For every $x \in \mathcal{X}$ we can define the set of probability measures:

$$\mathfrak{M}^\pi(x) = \mathcal{A}(x, Q(x, \pi(x))), \quad x \in \mathcal{X}. \quad (11)$$

The multifunction $\mathfrak{M}^\pi : \mathcal{X} \rightrightarrows \mathcal{P}(\mathcal{X})$, assigning to each $x \in \mathcal{X}$ the set $\mathfrak{M}^\pi(x)$, is a closed convex stochastic multikernel. We call it a *risk multikernel*, associated with the risk transition mapping $\sigma(\cdot, \cdot, \cdot)$, the controlled kernel Q , and the decision rule π . Its measurable selectors $M^\pi \ll \mathfrak{M}^\pi$ are transition kernels.

It follows that formula (10) for stationary policies Π can be rewritten as follows:

$$\rho_t^\Pi(v(x_{t+1})) = \max_{M \in \mathfrak{M}^\pi(x_t)} \int_{\mathcal{X}} v(y) M(dy). \quad (12)$$

In the risk-neutral case we have

$$\rho_t^\Pi(v(x_{t+1})) = \mathbb{E}^\Pi[v(x_{t+1})|x_t] = \int_{\mathcal{X}} v(y) Q(dy|x_t, \pi(x_t)).$$

The comparison of the last two displayed equations reveals that in the risk-neutral case we have

$$\mathfrak{M}^\pi(x) = \{Q(x, \pi(x))\}, \quad x \in \mathcal{X}, \quad (13)$$

that is, the risk multikernel \mathfrak{M}^π is single-valued, and its only selector is the kernel $Q(\cdot, \pi(\cdot))$. In the risk-averse case, the risk multikernel \mathfrak{M}^π is a closed convex-valued multifunction, whose measurable selectors are transition kernels. It is evident that properties of this multifunction are germane for our analysis. We return to this issue in section 7, where we calculate some examples of risk multikernels.

Remark 4.1 *If $m \in \mathcal{A}(x, m)$ for all $x \in \mathcal{X}$ and $m \in \mathcal{M}$, then it follows from equation (11) that $Q(\cdot, \pi(\cdot))$ is a measurable selector of \mathfrak{M}^π . Moreover, it follows from (6) that for any function $\varphi \in \mathcal{V}$ we have*

$$\rho_t^\Pi(\varphi(x_{t+1})) \geq \int_{\mathcal{X}} \varphi(y) Q(dy|x_t, \pi(x_t)) = \mathbb{E}^\Pi[\varphi(x_{t+1})|x_t].$$

It follows that the dynamic risk measure (1) is bounded from below by the expected value of the total cost.

The condition $m \in \mathcal{A}(x, m)$ is satisfied by the measures of risk in Examples 4.1 and 4.2.

Interestingly, uncertain transition matrices were used by Nilim and El Ghaoui in [26] to increase robustness of control rules for Markov models. There is also an intriguing connection to Markov games (see, e.g., [16, 20]). In our theory, controlled multikernels arise in a natural way in the analysis of risk-averse preferences.

5 General Assumptions. Semicontinuity and Measurability

We call the controlled kernel Q *setwise (strongly) continuous*, if for all Borel sets $B \subset \mathcal{X}$ and all convergent sequences $\{(x_k, u_k)\}, k = 1, 2, \dots,$

$$\lim_{k \rightarrow \infty} Q(B|x_k, u_k) = Q(B|x, u),$$

where $x = \lim_{k \rightarrow \infty} x_k$ and $u = \lim_{k \rightarrow \infty} u_k$. We call Q *weakly* continuous*, if for all functions $v \in \mathcal{V}$

$$\lim_{k \rightarrow \infty} \int_{\mathcal{X}} v(y) Q(dy|x_k, u_k) = \int_{\mathcal{X}} v(y) Q(dy|x, u).$$

Under condition (G0), setwise and weak* continuity concepts are equivalent, because the set of piecewise constant functions is dense in \mathcal{V} .

In the product space $\mathcal{X} \times \mathcal{M}$ we always consider the product topology of strong convergence in \mathcal{X} and weak* convergence in \mathcal{M} . In all our analysis we make the following assumptions:

- G1.** The transition kernel $Q(\cdot, \cdot)$ is setwise continuous;
- G2.** The multifunction $\mathcal{A}(\cdot, \cdot) \equiv \partial_\phi \sigma(0, \cdot, \cdot)$ is lower semicontinuous;
- G3.** The function $c(\cdot, \cdot, \cdot)$ is measurable, w -bounded, and $c(\cdot, \cdot, y)$ is lower semicontinuous for all $y \in \mathcal{X}$;
- G4.** The multifunction $U(\cdot)$ is measurable and compact-valued.

We need the following semicontinuity property of a risk transition mapping.

Proposition 5.1 *Suppose (G0)–(G3) and let $v \in \mathcal{V}$. Then the mapping $(x, u) \mapsto \sigma(c(x, u, \cdot) + v(\cdot), x, Q(x, u))$ is lower semicontinuous on $\text{graph}(U)$.*

Proof. Let $\varphi(x, u, y) = c(x, u, y) + v(y)$. Consider the dual representation (6) of the risk transition mapping

$$\sigma(\varphi(x, u, \cdot), x, Q(x, u)) = \max_{\mu \in \mathcal{A}(x, Q(x, u))} \int_{\mathcal{X}} \varphi(x, u, y) \mu(dy). \quad (14)$$

By (G0), (G1), and (G2), the multifunction $(x, u) \mapsto \mathcal{A}(x, Q(x, u))$ is lower semicontinuous. Owing to condition (G3), the function $(x, u, \mu) \mapsto \int_{\mathcal{X}} \varphi(x, u, y) \mu(dy)$ is lower semicontinuous on $\text{graph}(U) \times \mathcal{M}$. The assertion follows now from [4, Theorem 1.4.16], whose proof remains valid in our setting as well. \square

Some comments on the assumptions of Proposition 5.1 are in order. Continuity assumptions of the kernel Q are standard in the theory of risk-neutral Markov control processes (see, e.g., [17, App. C], [44]). If the risk transition mapping $\sigma(\cdot, \cdot, \cdot)$ is continuous, then its subdifferential (7) is upper semicontinuous. However, in Proposition 5.1 we assume *lower* semicontinuity of the mapping $(x, m) \mapsto \partial_\phi \sigma(0, x, m)$, which is not trivial and should be verified for each case.

Example 5.1 Let us verify the lower semicontinuity assumption for the multifunction \mathcal{A} given in (8). Consider an arbitrary $\mu \in \mathcal{A}(x, m)$ and suppose $x_k \rightarrow x$, $m_k \rightarrow m$, as $k \rightarrow \infty$. We need to find $\mu_k \in \mathcal{A}(x_k, m_k)$ such that $\mu_k \rightarrow \mu$. Let h be the function, for which, according to (8), $\frac{d\mu}{dm} = 1 + h - \int h(z) m(dz)$. We define μ_k by specifying their Radon–Nikodym derivatives: $\frac{d\mu_k}{dm_k} = 1 + h - \int h(z) m_k(dz)$. By construction, $\mu_k \in \mathcal{A}(x_k, m_k)$. Then, for any function $v \in \mathcal{V}$ we obtain

$$\begin{aligned} \int_{\mathcal{X}} v(y) \mu_k(dy) &= \int_{\mathcal{X}} v(y) \left(1 + h(y) - \int_{\mathcal{X}} h(z) m_k(dz) \right) m_k(dy) \\ &= \int_{\mathcal{X}} v(y) (1 + h(y)) m_k(dy) - \int_{\mathcal{X}} h(z) m_k(dz) \int_{\mathcal{X}} v(y) m_k(dy). \end{aligned}$$

As $m_k \rightarrow m$, we conclude that for all $v \in \mathcal{V}$

$$\lim_{k \rightarrow \infty} \int_{\mathcal{X}} v(y) \mu_k(dy) = \int_{\mathcal{X}} v(y) (1 + h(y)) m(dy) - \int_{\mathcal{X}} h(z) m(dz) \int_{\mathcal{X}} v(y) m(dy) = \int_{\mathcal{X}} v(y) \mu(dy),$$

which is the weak* convergence of μ_k to μ .

In the following result we use the concept of a *normal integrand*, that is, a function $f : \mathcal{X} \times U \rightarrow \mathbb{R} \cup +\infty$ such that its epigraphical mapping $x \mapsto \{(u, \alpha) \in U \times \mathbb{R} : f(x, u) \leq \alpha\}$ is a closed-valued and measurable multifunction (see Rockafellar and Wets [38, sec. 14.D]).

Proposition 5.2 *Suppose (G0)–(G4) and let $v \in \mathcal{V}$. Then the function*

$$\psi(x) = \inf_{u \in U(x)} \sigma(c(x, u, \cdot) + v(\cdot), x, Q(x, u)), \quad x \in \mathcal{X},$$

is measurable and w -bounded, and a measurable selector $\pi \prec U$ exists, such that

$$\psi(x) = \sigma(c(x, \pi(x), \cdot) + v(\cdot), x, Q(x, \pi(x))), \quad \forall x \in \mathcal{X}.$$

Proof. Consider the function $f : \mathcal{X} \times U \rightarrow \mathbb{R} \cup +\infty$ defined as follows:

$$f(x, u) = \begin{cases} \sigma(c(x, u, \cdot) + v(\cdot), x, Q(x, u)) & \text{if } u \in U(x), \\ +\infty & \text{otherwise.} \end{cases}$$

Owing to Proposition 5.1, $f(\cdot, \cdot)$ is lower semicontinuous, and is thus a normal integrand [38, Ex. 14.31]. It follows from [38, Thm. 14.37] that the function $\psi(x) = \inf_u f(x, u)$ is measurable and that the optimal solution mapping $\Psi(x) = \{u \in U : \psi(x) = f(x, u)\}$ is measurable. By (G4), the set $U(x)$ is compact, and thus $\Psi(x) \neq \emptyset$ for all $x \in \mathcal{X}$. Ψ is also compact-valued. By virtue of [23], a measurable selector $\pi \prec \Psi$ exists. Let us recall the dual representation (14) again:

$$\psi(x) = \max_{\mu \in \mathcal{A}(x, Q(x, \pi(x)))} \int_{\mathcal{X}} \varphi(x, \pi(x), y) \mu(dy),$$

with $\varphi(x, u, y) = c(x, u, y) + v(y)$. As the set $\mathcal{A}(x, Q(x, \pi(x)))$ contains only probability measures, and the function $\varphi(\cdot, \cdot, \cdot)$ is w -bounded, the function $\psi(\cdot)$ is w -bounded as well. \square

6 Finite Horizon Problem

We consider the Markov model at times $1, 2, \dots, T+1$ under deterministic policies $\Pi = \{\pi_1, \pi_2, \dots, \pi_T\}$. The cost at the last stage is given by a function $v_{T+1}(x_{T+1})$. Consider the problem

$$\min_{\Pi} J_T(\Pi, x_1), \tag{15}$$

with $J_T(\Pi, x_1)$ defined by formula (1), with Markov conditional risk measures $\rho_t^\Pi, t = 1, \dots, T$:

$$J_T(\Pi, x_1) = \rho_1^\Pi \left(c(x_1, u_1, x_2) + \rho_2^\Pi \left(c(x_2, u_2, x_3) + \dots + \rho_T^\Pi \left(c(x_T, u_T, x_{T+1}) + v_{T+1}(x_{T+1}) \right) \dots \right) \right). \tag{16}$$

This means that every one-step measure has the form (4), with some risk transition mapping $\sigma(\cdot, \cdot, \cdot)$.

Theorem 6.1 *Assume that the general conditions (G0)–(G4) are satisfied, and that the function $v_{T+1}(\cdot)$ is measurable and w -bounded. Then problem (15) has an optimal solution and its optimal value $v_1(x)$ is the solution of the following dynamic programming equations:*

$$v_t(x) = \min_{u \in U(x)} \sigma(c(x, u, \cdot) + v_{t+1}(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}, \quad t = T, \dots, 1. \tag{17}$$

Moreover, an optimal Markov policy $\hat{\Pi} = \{\hat{\pi}_1, \dots, \hat{\pi}_T\}$ exists and satisfies the equations:

$$\hat{\pi}_t(x) \in \operatorname{argmin}_{u \in U(x)} \sigma(c(x, u, \cdot) + v_{t+1}(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}, \quad t = T, \dots, 1. \tag{18}$$

Conversely, every solution of equations (17)–(18) defines an optimal Markov policy $\hat{\Pi}$.

Proof. Our proof is based on the ideas of the proof of Ruszczyński [39, Thm. 2], but with refinements rectifying some technical inaccuracies.²

Using the monotonicity condition (A2) for $t = 1, \dots, T$, we can rewrite problem (15) as follows:

$$\inf_{\pi_1, \dots, \pi_T} \left\{ \rho_1^\Pi \left(c(x_1, u_1, x_2) + \dots + \rho_T^\Pi \left(c(x_T, u_T, x_{T+1}) + v_{T+1}(x_{T+1}) \right) \dots \right) \right\} = \\ \inf_{\pi_1, \dots, \pi_{T-1}} \left\{ \rho_1^\Pi \left(c(x_1, u_1, x_2) + \dots + \inf_{\pi_T} \rho_T^\Pi \left(c(x_T, u_T, x_{T+1}) + v_{T+1}(x_{T+1}) \right) \dots \right) \right\}.$$

Owing to the Markov structure of the conditional risk measure ρ_T , the innermost optimization problem can be rewritten as follows:

$$\inf_{\pi_T} \sigma \left(c(x_T, \pi_T(x_T), \cdot) + v_{T+1}(\cdot), x_T, Q(x_T, \pi_T(x_T))) \right) \\ = \inf_{u \in U(x_T)} \sigma \left(c(x_T, u, \cdot) + v_{T+1}(\cdot), x_T, Q(x_T, u) \right). \quad (19)$$

The problem becomes equivalent to (17) for $t = T$, and its solution is given by (18) for $t = T$. By Proposition 5.2, a measurable selector $\hat{\pi}_T(\cdot)$ exists, such that $\hat{\pi}_T(x_T)$ is the minimizer in (19) for any x_T . Finally, the optimal value in (19), which we denote by $v_T(x_T)$, is measurable and w -bounded.

After that, the horizon $T + 1$ is decreased to T , and the final cost becomes $v_T(x_T)$. Proceeding in this way for $T, T - 1, \dots, 1$ we obtain the assertion of the theorem. \square

It follows from our proof that the functions $v_t(\cdot)$ calculated in (17) are the optimal values of tail subproblems formulated for a fixed $x_t = x$ as follows:

$$v_t(x) = \min_{\pi_t, \dots, \pi_T} \rho_t^\Pi \left(c(x_t, \pi_t(x_t), x_{t+1}) + \rho_{t+1}^\Pi \left(c(x_{t+1}, \pi_{t+1}(x_{t+1}), x_{t+2}) + \dots \right. \right. \\ \left. \left. + \rho_T^\Pi \left(c(x_T, \pi_T(x_T), x_{T+1}) + v_{T+1}(x_{T+1}) \right) \dots \right) \right).$$

We call them *value functions*, as in risk-neutral dynamic programming. It is obvious that we may have non-stationary costs, transition kernels, and risk transition mappings in this case. Also, the assumption that the process is transient is not needed.

Equations (17)–(18) provide a computational recipe for solving finite horizon problems.

7 Evaluation of Stationary Markov Policies in Infinite Horizon Problems

Consider a stationary policy $\Pi = \{\pi, \pi, \dots\}$ and define the cost until absorption as follows:

$$J_\infty(\Pi, x_1) = \lim_{T \rightarrow \infty} J_T(\Pi, x_1), \quad (20)$$

where each $J_T(\Pi, x_1)$ is defined by the formula

$$J_T(\Pi, x_1) = \rho_1^\Pi \left(c(x_1, \pi(x_1), x_2) + \rho_2^\Pi \left(c(x_2, \pi(x_2), x_3) + \dots + \rho_T^\Pi \left(c(x_T, \pi(x_T), x_{T+1}) \right) \dots \right) \right) \\ = \rho_{1, T+1}^\Pi \left(0, c(x_1, \pi(x_1), x_2), c(x_2, \pi(x_2), x_3), \dots, c(x_T, \pi(x_T), x_{T+1})) \right), \quad (21)$$

²In [39, Thm. 2] we missed the measurability condition on $U(\cdot)$ and the assumptions of *joint* continuity (lower semicontinuity) of the kernel and the cost functions. *Omnia principia parva sunt.*

with Markov conditional risk measures ρ_t^Π , $t = 1, \dots, T$, sharing the same risk transition mapping $\sigma(\cdot, \cdot, \cdot)$. We assume all conditions of Theorem 6.1.

The first question to answer is when this cost is finite. This question is nontrivial, because even for uniformly bounded costs $Z_t = c(x_{t-1}, \pi(x_{t-1}), x_t)$, $t = 2, 3, \dots$, and for a transient finite-state Markov chain, the limit in (20) may be infinite, as the following example demonstrates.

Example 7.1 Consider a transient Markov chain with two states and with the following transition probabilities: $Q_{11} = Q_{12} = \frac{1}{2}$, $Q_{22} = 1$. Only one control is possible in each state, the cost of each transition from state 1 is equal to 1, and the cost of the transition from 2 to 2 is 0. Clearly, the time until absorption is a geometric random variable with parameter $\frac{1}{2}$. Let $x_1 = 1$. If the limit (20) is finite, then (skipping the dependence on Π) we have

$$J_\infty(1) = \lim_{T \rightarrow \infty} J_T(1) = \lim_{T \rightarrow \infty} \rho_1(1 + J_{T-1}(x_2)) = \rho_1(1 + J_\infty(x_2)).$$

In the last equation we used the continuity of $\rho_1(\cdot)$. Clearly, $J_\infty(2) = 0$.

Suppose that we are using the Average Value at Risk from Example 3.2, with $0 < \alpha \leq \frac{1}{2}$, to define $\rho_1(\cdot)$. Using standard identities for the Average Value at Risk (see, e.g., [45, Thm. 6.2]), we obtain

$$\begin{aligned} J_\infty(1) &= \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha} \mathbb{E}[(1 + J_\infty(x_2) - \eta)_+] \right\} \\ &= 1 + \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha} \mathbb{E}[(J_\infty(x_2) - \eta)_+] \right\} = 1 + \frac{1}{\alpha} \int_{1-\alpha}^1 F^{-1}(\beta) d\beta, \end{aligned} \quad (22)$$

where $F(\cdot)$ is the distribution function of $J_\infty(x_2)$. As all β -quantiles of $J_\infty(x_2)$ for $\beta \geq \frac{1}{2}$ are equal to $J_\infty(1)$, the last equation yields $J_\infty(1) = 1 + J_\infty(1)$, a contradiction. It follows that a composition of average values at risk has no finite limit, if $0 < \alpha \leq \frac{1}{2}$.

On the other hand, if $\frac{1}{2} < \alpha < 1$, then

$$F^{-1}(\beta) = \begin{cases} J_\infty(2) = 0 & \text{if } 1 - \alpha \leq \beta < \frac{1}{2}, \\ J_\infty(1) & \text{if } \frac{1}{2} \leq \beta \leq 1. \end{cases}$$

Formula (22) then yields $J_\infty(1) = 1 + \frac{1}{2\alpha} J_\infty(1)$. This equation has a solution $J_\infty(1) = 2\alpha/(2\alpha - 1)$.

If we use the mean-semideviation model of Example 3.1, we obtain

$$\begin{aligned} J_\infty(1) &= \mathbb{E}[1 + J_\infty(x_2)] + \kappa \mathbb{E} \left[\left(1 + J_\infty(x_2) - \mathbb{E}[1 + J_\infty(x_2)] \right)_+ \right] \\ &= 1 + \frac{1}{2} J_\infty(1) + \kappa \frac{1}{2} \left(J_\infty(1) - \frac{1}{2} J_\infty(1) \right) = 1 + \frac{2 + \kappa}{4} J_\infty(1). \end{aligned}$$

Thus $J_\infty(1) = 4/(2 - \kappa)$, which is finite for all $\kappa \in [0, 1]$, that is, for all values of κ for which the model defines a coherent measure of risk.

It follows that deeper properties of the measures of risk and their interplay with the transition kernel need to be investigated to answer the question about finiteness of the dynamic measure of risk in this case.

Recall that with every risk transition mapping $\sigma(\cdot, \cdot, \cdot)$, every controlled kernel Q , and every decision rule π , a multikernel \mathfrak{M}^π is associated, as defined in (11). Similarly to the expected value case, it is convenient to consider the effective state space $\tilde{\mathcal{X}} = \mathcal{X} \setminus \{x_A\}$, and the *effective substochastic multikernel* $\tilde{\mathfrak{M}}^\pi$ whose arguments are restricted to $\tilde{\mathcal{X}}$ and whose values are convex sets of nonnegative measures on $\tilde{\mathcal{X}}$ defined by the identity: $\tilde{\mathfrak{M}}^\pi(B|x) \equiv \mathfrak{M}^\pi(B|x)$, for all $B \in \mathcal{B}(\tilde{\mathcal{X}})$ and $x \in \tilde{\mathcal{X}}$.

A function $v \in \mathcal{V}$ with $v(x_A) = 0$ can be identified with a function \tilde{v} on $\tilde{\mathcal{X}}$; we shall write $\|\tilde{v}\|$ for the norm $\|v\|$ in \mathcal{V} ; we shall also write $\tilde{v} \in \mathcal{V}$ to indicate that the corresponding extension v is an element of \mathcal{V} . Recall that the norm $\|\cdot\|_w$ associated with a weight function w is defined as follows:

$$\|v\|_w = \sup_{x \in \tilde{\mathcal{X}}} \frac{v(x)}{w(x)}.$$

The corresponding operator norm $\|A\|_w$ of a substochastic kernel A is defined as follows:

$$\|A\|_w = \sup_{x \in \tilde{\mathcal{X}}} \frac{1}{w(x)} \int_{\tilde{\mathcal{X}}} w(y) A(dy|x).$$

Definition 7.1 We call the Markov model with a risk transition mapping $\sigma(\cdot, \cdot, \cdot)$ and with a stationary Markov policy $\{\pi, \pi, \dots\}$ risk-transient if a weight function $w : \tilde{\mathcal{X}} \rightarrow [1, \infty)$, $w \in \mathcal{V}$, and a constant K exist such that

$$\|M\|_w \leq K \quad \text{for all } M \ll \sum_{j=1}^T (\tilde{\mathfrak{M}}^\pi)^j \quad \text{and all } T \geq 0. \quad (23)$$

If the estimate (23) is uniform for all Markov policies, the model is called uniformly risk-transient.

In the special case of a risk-neutral model, owing to the equation (13), Definition 7.1 reduces to the condition that

$$\left\| \sum_{j=1}^{\infty} (\tilde{Q}^\pi)^j \right\|_w \leq K, \quad (24)$$

which has been analyzed by Pliska [33] and [18, sec. 9.6].

Example 7.2 Consider the simple transient chain of Example 7.1 with the Average Value at Risk from Examples 3.2 and 4.2, where $0 < \alpha \leq 1$. From (9) we obtain

$$\mathcal{A}(i, m) = \left\{ (\mu_1, \mu_2) : 0 \leq \mu_j \leq \frac{m_j}{\alpha}, j = 1, 2; \mu_1 + \mu_2 = 1 \right\}.$$

As only one control is possible, formula (11) simplifies to

$$\mathfrak{M}(i) = \left\{ (\mu_1, \mu_2) : 0 \leq \mu_j \leq \frac{Q_{ij}}{\alpha}, j = 1, 2; \mu_1 + \mu_2 = 1 \right\}, \quad i = 1, 2.$$

The effective state space is just $\tilde{\mathcal{X}} = \{1\}$, and we conclude that the effective multikernel is the interval

$$\tilde{\mathfrak{M}} = \left[0, \min \left(1, \frac{1}{2\alpha} \right) \right].$$

For $0 < \alpha \leq \frac{1}{2}$ we can select $\tilde{M} = 1 \in \tilde{\mathfrak{M}}$ to show that $1 \in (\tilde{\mathfrak{M}})^j$ for all j , and thus condition (23) is not satisfied. On the other hand, if $\frac{1}{2} < \alpha \leq 1$, then for every $\tilde{M} \in \tilde{\mathfrak{M}}$ we have $0 \leq \tilde{M} < 1$, and condition (23) is satisfied.

Consider now the mean-semideviation model of Examples 3.1 and 4.1. From (8) we obtain

$$\begin{aligned} \mathcal{A}(i, m) &= \left\{ (\mu_1, \mu_2) : \mu_j = m_j (1 + h_j - (h_1 m_1 + h_2 m_2)), 0 \leq h_j \leq \kappa, j = 1, 2 \right\}, \\ \mathfrak{M}(i) &= \left\{ (\mu_1, \mu_2) : \mu_j = Q_{ij} (1 + h_j - (h_1 Q_{i1} + h_2 Q_{i2})), 0 \leq h_j \leq \kappa, j = 1, 2 \right\}, \quad i = 1, 2. \end{aligned}$$

Calculating the lowest and the largest possible values of μ_1 we conclude that

$$\tilde{\mathfrak{M}} = \left[\frac{1}{2} \left(1 - \frac{\kappa}{2} \right), \frac{1}{2} \left(1 + \frac{\kappa}{2} \right) \right].$$

For every $\kappa \in [0, 1]$, Definition 7.1 is satisfied.

We can now provide sufficient conditions for the finiteness of the limit (20).

Theorem 7.1 *Suppose a stationary policy $\Pi = \{\pi, \pi, \dots\}$ is applied to a the controlled Markov model with a risk transition mapping $\sigma(\cdot, \cdot, \cdot)$. If the model satisfies conditions (G0)–(G3), is risk-transient for the policy Π , and the cost function $c(\cdot, \cdot, \cdot)$ is w -bounded, then the limit*

$$J_\infty(\Pi, \cdot) = \lim_{T \rightarrow \infty} J_T(\Pi, \cdot), \quad (25)$$

exists in \mathcal{V} and is w -bounded. If the model is additionally uniformly risk-transient, then $\|J_\infty(\Pi, \cdot)\|_w$ is uniformly bounded for all Π and the limit function $(\pi, x) \mapsto J_\infty(\Pi, x)$ is lower semicontinuous.

Proof. By Conditions (A1)–(A4), each conditional risk measure $\rho_{1,T}(\cdot)$ is convex and positively homogeneous, and thus subadditive. For any $1 < T_1 < T_2$ we obtain the following estimate of (21):

$$\begin{aligned} J_{T_2-1}(\Pi, x_1) &= \rho_{1,T_2}^\Pi(0, Z_2, \dots, Z_{T_2}) \\ &\leq \rho_{1,T_2}^\Pi(0, Z_2, \dots, Z_{T_1}, 0, \dots, 0) + \rho_{1,T_2}^\Pi(0, \dots, 0, Z_{T_1+1}, \dots, Z_{T_2}) \\ &= \rho_{1,T_1}^\Pi(0, Z_2, \dots, Z_{T_1}) + \rho_{1,T_2}^\Pi(0, \dots, 0, Z_{T_1+1}, \dots, Z_{T_2}) \\ &= J_{T_1-1}(\Pi, x_1) + \rho_{1,T_2}^\Pi(0, \dots, 0, Z_{T_1+1}, \dots, Z_{T_2}). \end{aligned} \quad (26)$$

As the cost function is w -bounded, $Z_{j+1} \leq C(\bar{w}(x_j) + \bar{w}(x_{j+1}))$, where $\bar{w}(x) = w(x)$ if $x \in \tilde{\mathcal{X}}$, and $\bar{w}(x_A) = 0$. Owing to the monotonicity and positive homogeneity of the conditional risk mappings,

$$\begin{aligned} \rho_{1,T_2}^\Pi(0, \dots, 0, Z_{T_1+1}, \dots, Z_{T_2}) &\leq 2C\rho_{1,T_2}^\Pi(0, \dots, 0, \bar{w}(x_{T_1+1}), \dots, \bar{w}(x_{T_2})) \\ &= 2C\rho_1^\Pi \left(\rho_2^\Pi \left(\dots \rho_{T_1}^\Pi \left(\bar{w}(x_{T_1+1}) + \rho_{T_1+1}^\Pi(\bar{w}(x_{T_1+1}) + \dots + \rho_{T_2-1}^\Pi(\bar{w}(x_{T_2})) \dots) \right) \dots \right) \right). \end{aligned} \quad (27)$$

If $x_{T_2-1} \neq x_A$, applying (12) to the innermost expression, we obtain

$$\rho_{T_2-1}^\Pi(\bar{w}(x_{T_2})) = \max_{m \in \tilde{\mathfrak{M}}^\pi(x_{T_2-1})} \int_{\tilde{\mathcal{X}}} w(y) m(dy). \quad (28)$$

It is a function of x_{T_2-1} , which we denote as $v_{T_2-1}(x_{T_2-1})$. Restricting the domains of the functions to $\tilde{\mathcal{X}}$, we may write the relation:

$$v_{T_2-1} = \tilde{M}_{T_2-1} w, \quad \tilde{M}_{T_2-1} \leq \tilde{\mathfrak{M}}^\pi, \quad (29)$$

where the selector \tilde{M}_{T_2-1} has values $\tilde{M}_{T_2-1}(x_{T_2-1})$ equal to the maximizers in (28). The maximizers exist owing to the weak* compactness of the values of the multikernel $\tilde{\mathfrak{M}}^\pi$. One step earlier, we obtain

$$\begin{aligned} \rho_{T_2-2}^\Pi(\bar{w}(x_{T_2-1}) + \rho_{T_2-1}^\Pi(\bar{w}(x_{T_2}))) &= \rho_{T_2-2}^\Pi(\bar{w}(x_{T_2-1}) + v_{T_2-1}(x_{T_2-1})) \\ &= \max_{m \in \tilde{\mathfrak{M}}^\pi(x_{T_2-2})} \int_{\mathcal{X}} [w(y) + v_{T_2-1}(y)] m(dy). \end{aligned} \quad (30)$$

Again, the maximizers $\tilde{M}_{T_2-2}(x_{T_2-2})$ in (30) exist, and they can be chosen in a measurable way. Denoting the optimal value by $v_{T_2-2}(x_{T_2-2})$, we obtain a relation similar to (29):

$$v_{T_2-2} = \tilde{M}_{T_2-2}(w + v_{T_2-1}) = (\tilde{M}_{T_2-2} + \tilde{M}_{T_2-2}\tilde{M}_{T_2-1})w, \quad \tilde{M}_{T_2-2} \leq \tilde{\mathfrak{M}}^\pi, \quad \tilde{M}_{T_2-1} \leq \tilde{\mathfrak{M}}^\pi. \quad (31)$$

Proceeding in this way, we can calculate the function

$$v_{T_1}(x_{T_1}) = \rho_{T_1}^\Pi \left(\bar{w}(x_{T_1+1}) + \rho_{T_1+1}^\Pi(\bar{w}(x_{T_1+1}) + \rho_{T_1+2}^\Pi(\bar{w}(x_{T_1+1}) + \dots + \rho_{T_2-1}^\Pi(\bar{w}(x_{T_2})) \dots) \right)$$

on $\tilde{\mathcal{X}}$ as follows:

$$v_{T_1} = (\tilde{M}_{T_1} + \tilde{M}_{T_1} \tilde{M}_{T_1+1} + \cdots + \tilde{M}_{T_1} \tilde{M}_{T_1+1} \cdots \tilde{M}_{T_2-1})w,$$

with $\tilde{M}_j \ll \tilde{\mathfrak{M}}^{\pi}$, $j = T_1, \dots, T_2 - 1$. In the formula above, we restrict the domains of the functions to $\tilde{\mathcal{X}}$; at x_A their values are zero. Finally, defining

$$v_1(x_1) = \rho_1^{\Pi} \left(\rho_2^{\Pi} \left(\cdots \rho_{T_1}^{\Pi} \left(\bar{w}(x_{T_1+1}) + \rho_{T_1+1}^{\Pi} \left(\bar{w}(x_{T_1+1}) + \cdots + \rho_{T_2-1}^{\Pi} \left(\bar{w}(x_{T_2}) \right) \right) \right) \right) \right),$$

we obtain the representation

$$v_1 = \tilde{M}_1 \tilde{M}_2 \cdots \tilde{M}_{T_1-1} (\tilde{M}_{T_1} + \tilde{M}_{T_1} \tilde{M}_{T_1+1} + \cdots + \tilde{M}_{T_1} \tilde{M}_{T_1+1} \cdots \tilde{M}_{T_2-1})w, \quad (32)$$

with $\tilde{M}_j \ll \tilde{\mathfrak{M}}^{\pi}$, $j = 1, \dots, T_2 - 1$. This combined with (26)–(27) yields an estimate:

$$J_{T_2-1}(\Pi, \cdot) - J_{T_1-1}(\Pi, \cdot) \leq 2C \tilde{M}_1 \tilde{M}_2 \cdots \tilde{M}_{T_1-1} (\tilde{M}_{T_1} + \tilde{M}_{T_1} \tilde{M}_{T_1+1} + \cdots + \tilde{M}_{T_1} \tilde{M}_{T_1+1} \cdots \tilde{M}_{T_2-1})w. \quad (33)$$

Consider now the sequence of costs $Z_1, \dots, Z_{T_1}, -Z_{T_1+1}, \dots, -Z_{T_2}$, in which we flip the sign of the costs $Z_{t+1} = c(x_t, u_t, x_{t+1})$ for $t \geq T_1$. From subadditivity, similarly to (26), we obtain

$$\rho_{1, T_2}^{\Pi}(0, Z_2, \dots, Z_{T_1}, -Z_{T_1+1}, \dots, -Z_{T_2}) \leq \rho_{1, T_1}^{\Pi}(0, Z_2, \dots, Z_{T_1}) + \rho_{1, T_2}^{\Pi}(0, \dots, 0, -Z_{T_1+1}, \dots, -Z_{T_2}). \quad (34)$$

By convexity of $\rho_{1, T_2}(\cdot)$,

$$2\rho_{1, T_1}^{\Pi}(0, Z_2, \dots, Z_{T_1}) \leq \rho_{1, T_2}^{\Pi}(0, Z_2, \dots, Z_{T_1}, Z_{T_1+1}, \dots, Z_{T_2}) + \rho_{1, T_2}^{\Pi}(0, Z_2, \dots, Z_{T_1}, -Z_{T_1+1}, \dots, -Z_{T_2}).$$

Substituting the estimate (34), we deduce that

$$\rho_{1, T_2}^{\Pi}(0, Z_2, \dots, Z_{T_2}) \geq \rho_{1, T_1}^{\Pi}(0, Z_2, \dots, Z_{T_1}) - \rho_{1, T_2}^{\Pi}(0, \dots, 0, -Z_{T_1+1}, \dots, -Z_{T_2}).$$

As $|Z_{t+1}|$ are bounded by $C(\bar{w}(x_t) + \bar{w}(x_{t+1}))$, the estimate (27) applies to the last element on the right hand side. We obtain

$$\begin{aligned} J_{T_2-1}(\Pi, x_1) - J_{T_1-1}(\Pi, x_1) &= \rho_{1, T_2}^{\Pi}(0, Z_2, \dots, Z_{T_2}) - \rho_{1, T_1}^{\Pi}(0, Z_2, \dots, Z_{T_1}) \\ &\geq -2C \rho_1^{\Pi} \left(\rho_2^{\Pi} \left(\cdots \rho_{T_1}^{\Pi} \left(\bar{w}(x_{T_1+1}) + \rho_{T_1+1}^{\Pi} \left(\bar{w}(x_{T_1+1}) + \cdots + \rho_{T_2-1}^{\Pi} \left(\bar{w}(x_{T_2}) \right) \right) \right) \right) \right) = -2C v_1(x_1), \end{aligned}$$

where $v_1(\cdot)$ has representation (32). This combined with (33) yields

$$|J_{T_2-1}(\Pi, x_1) - J_{T_1-1}(\Pi, x_1)| \leq 2C |v_1(x_1)|, \quad x_1 \in \tilde{\mathcal{X}}.$$

This pointwise estimate implies the relations between the norms:

$$\|J_{T_2-1}(\Pi, \cdot) - J_{T_1-1}(\Pi, \cdot)\|_w \leq 2C \|v_1\|_w.$$

In view of representation (32), we obtain the estimate

$$\|J_{T_2-1}(\Pi, \cdot) - J_{T_1-1}(\Pi, \cdot)\|_w \leq 2C \|\tilde{M}_1 \tilde{M}_2 \cdots \tilde{M}_{T_1-1} (\tilde{M}_{T_1} + \tilde{M}_{T_1} \tilde{M}_{T_1+1} + \cdots + \tilde{M}_{T_1} \tilde{M}_{T_1+1} \cdots \tilde{M}_{T_2-1})w\|_w.$$

By Definition 7.1, $\|\tilde{M}_{T_1} + \tilde{M}_{T_1} \tilde{M}_{T_1+1} + \cdots + \tilde{M}_{T_1} \tilde{M}_{T_1+1} \cdots \tilde{M}_{T_2-1}\|_w \leq K$, and thus

$$\|J_{T_2-1}(\Pi, \cdot) - J_{T_1-1}(\Pi, \cdot)\|_w \leq 2CK \|\tilde{M}_1 \tilde{M}_2 \cdots \tilde{M}_{T_1-1}\|_w. \quad (35)$$

Observe that $\tilde{M}_1 \tilde{M}_2 \cdots \tilde{M}_{T_1-1} \ll (\tilde{\mathfrak{M}}^{\pi})^{T_1-1}$. It follows from Definition 7.1 that for any sequence of selectors $A_j \ll (\tilde{\mathfrak{M}}^{\pi})^j$ we have $\|\sum_{j=1}^{\infty} A_j\|_w \leq K$. Therefore, $\|A_j\|_w \rightarrow 0$, as $j \rightarrow \infty$. Consequently, the right

hand side of (35) converges to 0, when $T_1, T_2 \rightarrow \infty$, $T_1 < T_2$. Hence, the sequence of functions $J_T(\Pi, \cdot)$, $T = 1, 2, \dots$ is convergent to some w -bounded limit $J_\infty(\Pi, \cdot) \in \mathcal{V}$. The convergence is w -uniform, that is,

$$\lim_{T \rightarrow \infty} \sup_{x \in \tilde{\mathcal{X}}} \frac{|J_T(\Pi, x) - J_\infty(\Pi, x)|}{w(x)} = 0.$$

If the model is uniformly risk-transient, then the estimate (35) is the same for all Markov policies Π , and thus $\|J_\infty(\Pi, \cdot)\|_w$ is uniformly bounded. Moreover,

$$\lim_{T \rightarrow \infty} \sup_{\substack{x \in \tilde{\mathcal{X}} \\ \Pi \in \Pi^{\text{DM}}}} \frac{|J_T(\Pi, x) - J_\infty(\Pi, x)|}{w(x)} = 0,$$

where Π^{DM} is the set of all stationary deterministic Markov policies. As each of the functions $(\pi, x) \mapsto J_T(\Pi, x)$ is lower semicontinuous, so is the limit function $(\pi, x) \mapsto J_\infty(\Pi, x)$. \square

Remark 7.1 *It is clear from the proof of Theorem 7.1, that*

$$J_\infty(\Pi, x_1) = \lim_{T \rightarrow \infty} \rho_{1,T}^\Pi(0, Z_2, \dots, Z_T + f(x_T)), \quad (36)$$

for any w -bounded measurable function $f: \mathcal{X} \rightarrow \mathbb{R}$, because $c(x_{T-1}, u_t, x_T) + f(x_T)$ is still w -bounded.

This analysis allows us to derive policy evaluation equations for the infinite horizon problem, in the case of a fixed Markov policy.

Theorem 7.2 *Suppose a controlled Markov model with a risk transition mapping $\sigma(\cdot, \cdot, \cdot)$ is risk-transient for the stationary Markov policy $\Pi = \{\pi, \pi, \dots\}$, with some weight function $w(\cdot)$. If condition (G3) is satisfied, then a w -bounded function $v \in \mathcal{V}$ satisfies the equations*

$$v(x) = \sigma(c(x, \pi(x), \cdot) + v(\cdot), x, Q(x, \pi(x))), \quad x \in \tilde{\mathcal{X}}, \quad (37)$$

$$v(x_A) = 0, \quad (38)$$

if and only if $v(x) = J_\infty(\Pi, x)$ for all $x \in \mathcal{X}$.

Proof. Denote $Z_t = c(x_{t-1}, u_{t-1}, x_t)$. Suppose a w -bounded function $v \in \mathcal{V}$ satisfies the equations (37)–(38). By (G3), $c(x, \pi(x), \cdot) \in \mathcal{V}$, and thus the right-hand side of (37) is well-defined. By iteration of (37), we obtain for all $x_1 \in \mathcal{X}$ the following equation:

$$v(x_1) = \rho_1^\Pi \left(c(x_1, u_1, x_2) + \rho_2^\Pi \left(c(x_2, u_2, x_3) + \dots + \rho_T^\Pi \left(c(x_T, u_T, x_{T+1}) + v(x_{T+1}) \right) \dots \right) \right).$$

Denote $Z_t = c(x_{t-1}, u_{t-1}, x_t)$. Using subadditivity and monotonicity of the conditional risk measures we deduce that:

$$\begin{aligned} v(x_1) &= \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1} + v(x_{T+1})) \leq \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1}) + \rho_{1,T+1}^\Pi(0, 0, \dots, v(x_{T+1})) \\ &\leq \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1}) + \rho_{1,T+1}^\Pi(0, 0, \dots, |v(x_{T+1})|). \end{aligned} \quad (39)$$

By convexity of $\rho_{1,T+1}^\Pi(\cdot)$,

$$\begin{aligned} 2\rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1}) &\leq \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1} + v(x_{T+1})) + \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1} - v(x_{T+1})) \\ &= v(x_1) + \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1} - v(x_{T+1})). \end{aligned} \quad (40)$$

Similar to (39),

$$\begin{aligned} \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1} - v(x_{T+1})) &\leq \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1}) + \rho_{1,T+1}^\Pi(0, 0, \dots, -v(x_{T+1})) \\ &\leq \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1}) + \rho_{1,T+1}^\Pi(0, 0, \dots, |v(x_{T+1})|). \end{aligned}$$

Substituting into (40) we obtain

$$v(x_1) \geq \rho_{1,T+1}^\Pi(0, Z_2, \dots, Z_{T+1}) - \rho_{1,T+1}^\Pi(0, 0, \dots, |v(x_{T+1})|).$$

Combining this estimate with (39), we conclude that

$$|v(x_1) - J_T(\Pi, x_1)| \leq \rho_{1,T+1}^\Pi(0, 0, \dots, |v(x_{T+1})|). \quad (41)$$

Consider the function

$$d_{1,T}(x_1) = \rho_{1,T+1}^\Pi(0, 0, \dots, |v(x_{T+1})|).$$

Proceeding exactly as in the proof of Theorem 7.1, we obtain a representation similar to (32):

$$d_{1,T} = \tilde{M}_1 \cdots \tilde{M}_T |v|,$$

with $\tilde{M}_j < \tilde{\mathfrak{M}}^{\pi}$, $j = 1, \dots, T-1$. Thus, $d_{1,T} = A_T |v|$, with $A_T < (\tilde{\mathfrak{M}}^{\pi})^T$. By Definition 7.1, for any sequence of selectors $A_t < (\tilde{\mathfrak{M}}^{\pi})^t$, $t = 1, 2, \dots$, we have $\|\sum_{t=1}^{\infty} A_t\|_w \leq K$. Therefore, $\|A_t\|_w \rightarrow 0$ and $\|d_{1,t}\|_w \rightarrow 0$, as $t \rightarrow \infty$. Using this in (41) we conclude that $v(\cdot) \equiv J_\infty(\Pi, \cdot)$, as postulated.

To prove the converse implication we can use the fact that all conditional risk measures $\rho_t^\Pi(\cdot)$ share the same risk transition mapping $\sigma(\cdot, \cdot, \cdot)$ to rewrite (21) as follows:

$$J_T(\Pi, x_1) = \sigma(c(x_1, \pi(x_1), \cdot) + J_{T-1}(\Pi, \cdot), x_1, Q(x_1, \pi(x_1))).$$

The function $\sigma(\cdot, x_1, \mu)$, as a finite-valued coherent measure of risk on a Banach lattice \mathcal{V} , is continuous (see [41, Prop. 3.1]). By Theorem 7.1, the sequence $\{J_T(\Pi, \cdot)\}$ is convergent to $J_\infty(\Pi, \cdot)$ in the space \mathcal{V} , and $J_\infty(\Pi, \cdot)$ is w -bounded. Therefore,

$$\lim_{T \rightarrow \infty} J_T(\Pi, x_1) = \sigma(c(x_1, \pi(x_1), \cdot) + \lim_{T \rightarrow \infty} J_{T-1}(\Pi, \cdot), x_1, Q(x_1, \pi(x_1))).$$

This is identical with equation (37) with $v(\cdot) \equiv J_\infty(\Pi, \cdot)$. Equation (38) is obvious. \square

8 Dynamic Programming Equations for Infinite Horizon Problems

We shall now focus on the optimal value function

$$J^*(x) = \inf_{\Pi \in \Pi^{\text{DM}}} J_\infty(\Pi, x), \quad x \in \mathcal{X}, \quad (42)$$

where Π^{DM} is the set of all stationary deterministic Markov policies. To simplify notation, we define the operators $\mathfrak{D} : \mathcal{V} \rightarrow \mathcal{V}$ and $\mathfrak{D}_\pi : \mathcal{V} \rightarrow \mathcal{V}$ as follows:

$$[\mathfrak{D}v](x) = \min_{u \in U(x)} \sigma(c(x, u, \cdot) + v(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}, \quad (43)$$

$$[\mathfrak{D}_\pi v](x) = \sigma(c(x, \pi(x), \cdot) + v^k(\cdot), x, Q(x, \pi(x))), \quad x \in \mathcal{X}, \quad (44)$$

where $\pi < U$. Owing to the monotonicity of $\sigma(\cdot, x, \mu)$, both operators are nondecreasing. By construction, $\mathfrak{D}v \leq \mathfrak{D}_\pi v$ for all $v \in \mathcal{V}$ and all $\pi < U$.

Theorem 8.1 *Assume that conditions (G0)–(G4) are satisfied and that the model is uniformly risk-transient. Then a measurable w -bounded function $v : \mathcal{X} \rightarrow \mathbb{R}$ satisfies the equations*

$$v(x) = \inf_{u \in U(x)} \sigma(c(x, u, \cdot) + v(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}, \quad (45)$$

$$v(x_A) = 0, \quad (46)$$

if and only if $v(x) = J^(x)$ for all $x \in \mathcal{X}$. Moreover, a measurable minimizer $\pi^*(x)$, $x \in \mathcal{X}$, on the right hand side of (45) exists and defines an optimal deterministic Markov policy $\Pi^* = \{\pi^*, \pi^*, \dots\}$.*

Proof. Consider a sequence of Markov deterministic policies $\Pi^k = \{\pi^k, \pi^k, \dots\}$, $k = 1, 2, \dots$ constructed in the following way. We choose any $\pi^1 \in U$. Its value $v^1(\cdot) = J_\infty(\Pi^1, \cdot)$ is then given by equations (37)–(38). For $k = 1, 2, \dots$ we determine $\pi^{k+1}(\cdot)$ as the measurable solution of the problem

$$\min_{u \in U(x)} \sigma(c(x, u, \cdot) + v^k(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}, \quad (47)$$

which exists by Proposition 5.2. The corresponding value of the policy $\Pi^{k+1} = \{\pi^{k+1}, \pi^{k+1}, \dots\}$ is the function $v^{k+1}(\cdot) = J_\infty(\Pi^{k+1}, \cdot)$, and the iteration continues. By construction, the sequences $\{\pi^k\}$ and $\{v^k\}$ satisfy the relations:

$$\mathfrak{D}_{\pi^{k+1}} v^k = \mathfrak{D} v^k \leq \mathfrak{D}_{\pi^k} v^k = v^k. \quad (48)$$

Applying the operator $\mathfrak{D}_{\pi^{k+1}}$ to this relation, we deduce from its monotonicity that

$$[\mathfrak{D}_{\pi^{k+1}}]^{T-1} v^k \leq \mathfrak{D}_{\pi^{k+1}} v^k = \mathfrak{D} v^k \leq v^k, \quad T = 2, 3, \dots \quad (49)$$

Relation (49) can be equivalently written as

$$\rho_{1,T}^{\Pi^{k+1}}(0, Z_2, \dots, Z_T + v^k(x_T)) \leq [\mathfrak{D} v^k](x_1) \leq v^k(x_1),$$

where $Z_t = c(x_{t-1}, u_{t-1}, x_t)$, $t = 2, 3, \dots, T$, is the cost sequence resulting from the policy Π^{k+1} . Passing to the limit with $T \rightarrow \infty$, from Remark 7.1 we conclude that the sequence $\{v^k\}$ is nonincreasing:

$$v^{k+1}(x) = J_\infty(\Pi^{k+1}, x) \leq [\mathfrak{D} v^k](x) \leq v^k(x), \quad x \in \mathcal{X}, \quad k = 0, 1, 2, \dots \quad (50)$$

Since $v^k \geq J^*$, the sequence $\{v^k\}$ is monotonically convergent to some limit $v^\infty \geq J^*$. By Lebesgue Theorem, it is also convergent in the space \mathcal{V} . As the function $\sigma(\cdot, x, \mu)$ is a coherent measure of risk, it follows from [41, Prop. 3.1] that it is continuous on \mathcal{V} and thus

$$\sigma(c(x, u, \cdot) + v^k(\cdot), x, Q(x, u)) \downarrow \sigma(c(x, u, \cdot) + v^\infty(\cdot), x, Q(x, u)), \quad \text{as } k \rightarrow \infty, \quad \forall u \in U(x). \quad (51)$$

The left inequality in (50) also implies that

$$v^{k+1}(x) \leq \sigma(c(x, u, \cdot) + v^k(\cdot), x, Q(x, u)), \quad \forall u \in U(x). \quad (52)$$

Passing to the limit with $k \rightarrow \infty$ on both sides of (52) and using (51), we conclude that

$$v^\infty(x) \leq \sigma(c(x, u, \cdot) + v^\infty(\cdot), x, Q(x, u)), \quad \forall u \in U(x).$$

Since this is true for all $x \in \mathcal{X}$ and all $u \in U(x)$, it follows that

$$v^\infty(x) \leq [\mathfrak{D} v^\infty](x) = \min_{u \in U(x)} \sigma(c(x, u, \cdot) + v^\infty(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}. \quad (53)$$

Iterating this inequality, we conclude that for every feasible decision rule π and every $T = 1, 2, \dots$

$$v^\infty \leq [\mathfrak{D}]^T v^\infty \leq [\mathfrak{D}_\pi]^T v^\infty.$$

Owing to Remark 7.1, the right hand side of this inequality is convergent to $J_\infty(\Pi, \cdot)$, when $T \rightarrow \infty$. Since Π was arbitrary, $v^\infty \leq J^*$, and thus $v^\infty = J^*$. It remains to show that v^∞ satisfies equation (45). It follows from the monotonicity of the operator \mathfrak{D} and relation (48), that

$$\mathfrak{D}v^\infty \leq \mathfrak{D}v^k \leq v^k.$$

Passing to the limit with $k \rightarrow \infty$ we conclude that $\mathfrak{D}v^\infty \leq v^\infty$. This combined with (53) yields $v^\infty = \mathfrak{D}v^\infty$, which is equation (45). Denoting by $\pi^*(x)$ the (measurable) minimizer on the right hand side of (53), we also see that $v^\infty = \mathfrak{D}_{\pi^*}v^\infty$. By Theorem 7.2, $v^\infty(\cdot) = J_\infty(\Pi^*, \cdot)$.

To prove the converse implication, suppose $v \in \mathcal{V}$ satisfies (45)–(46) and $\|v\|_w < \infty$. By Proposition 5.2, a measurable minimizer $\hat{\pi}(x)$ on the right hand side of (45) exists. We obtain the equation

$$v(x) = \sigma(c(x, \hat{\pi}(x), \cdot) + v(\cdot), x, Q(x, \hat{\pi}(x))), \quad x \in \mathcal{X}.$$

Due to Theorem 7.2,

$$v(x) = J_\infty(\hat{\Pi}, x) \geq J^*(x), \quad x \in \mathcal{X}, \quad (54)$$

where $\hat{\Pi} = \{\hat{\pi}, \hat{\pi}, \dots\}$. On the other hand, it follows from (45) that for any stationary deterministic Markov policy $\Pi = \{\pi, \pi, \dots\}$ we have

$$v(x) \leq \sigma(c(x, \pi(x), \cdot) + v(\cdot), x, Q(x, \pi(x))), \quad x \in \mathcal{X}. \quad (55)$$

The risk transition mapping σ is nondecreasing with respect to the first argument. Therefore, iterating inequality (55) we obtain the following inequality:

$$v(x_1) \leq \rho_{1,T}^{\Pi}(0, Z_2, \dots, Z_T + v(x_T)),$$

Passing to the limit with $T \rightarrow \infty$ and applying Remark 7.1, we obtain for all stationary deterministic Markov policies $\Pi = \{\pi, \pi, \dots\}$ the inequality $v(\cdot) \leq J_\infty(\Pi, \cdot)$. The last estimate together with (54) implies that $v(\cdot) \equiv J^*(\cdot)$ and that the stationary policy $\hat{\Pi}$ is optimal. \square

Remark 8.1 *The construction of the sequence of stationary Markov policies $\{\Pi^k\}$ and their corresponding values $\{v^k\}$ employed in the first part of the proof of Theorem 8.1 can be interpreted as a risk-averse version of the policy iteration method.*

We can now address the case of non-stationary deterministic policies. For a deterministic policy $\Lambda = \{\lambda_1, \lambda_2, \dots\}$ we define

$$J_\infty(\Lambda, x_1) = \liminf_{T \rightarrow \infty} J_T(\Lambda, x_1) \quad \text{and} \quad \hat{J}(x_1) = \inf_{\Lambda} J_\infty(\Lambda, x_1).$$

In the theorem below, we make an additional technical assumption that the function $\hat{J}(\cdot)$ is measurable. It is obviously satisfied for a finite or countable state space. In general, its verification, even in the expected value case, requires additional assumptions (see, e. g., [18, sec. 9.3]). We discuss some sufficient conditions after the theorem.

Theorem 8.2 *Assume that conditions (G0)–(G4) are satisfied and that the model is uniformly risk-transient. Additionally, assume that the function $\hat{J}(\cdot)$ is measurable and a constant C exists such that $J_\infty(\Lambda, x) \geq -Cw(x)$ for all $x \in \mathcal{X}$ and for all policies Λ . Then a w -bounded function $v \in \mathcal{V}$ satisfies the equations (45)–(46) if and only if $v(x) = \hat{J}(x)$ for all $x \in \mathcal{X}$. Moreover, a measurable minimizer $\pi^*(x)$, $x \in \mathcal{X}$, on the right hand side of (45) exists and defines an optimal deterministic policy $\Pi^* = \{\pi^*, \pi^*, \dots\}$.*

Proof. As for stationary Markov policies Π we have $\|J_\infty(\Pi, \cdot)\|_w < \infty$, in view of the additional assumption we have $\|\hat{J}\|_w < \infty$. Denote $\Lambda^2 = \{\lambda_2, \lambda_3, \dots\}$. Due to the monotonicity and continuity of $\rho_1^T(\cdot)$, we obtain the chain of relations

$$\begin{aligned} \hat{J}(x_1) &= \inf_{\lambda_1, \lambda_2, \dots} \liminf_{T \rightarrow \infty} \rho_1^{\lambda_1} (c(x_1, \lambda_1(x_1), x_2) + J_{T-1}(\Lambda^2, x_2)) \\ &\geq \inf_{\lambda_1, \lambda_2, \dots} \liminf_{T \rightarrow \infty} \rho_1^{\lambda_1} (c(x_1, \lambda_1(x_1), x_2) + \inf_{\tau \geq T-1} J_\tau(\Lambda^2, x_2)) \\ &= \inf_{\lambda_1, \lambda_2, \dots} \lim_{T \rightarrow \infty} \rho_1^{\lambda_1} (c(x_1, \lambda_1(x_1), x_2) + \inf_{\tau \geq T-1} J_\tau(\Lambda^2, x_2)) \\ &= \inf_{\lambda_1, \lambda_2, \dots} \rho_1^{\lambda_1} \left(c(x_1, \lambda_1(x_1), x_2) + \liminf_{T \rightarrow \infty} J_{T-1}(\Lambda^2, x_2) \right) \\ &= \inf_{\lambda_1, \lambda_2, \dots} \rho_1^{\lambda_1} (c(x_1, \lambda_1(x_1), x_2) + J_\infty(\Lambda^2, x_2)), \end{aligned}$$

Owing to the monotonicity of $\rho_1(\cdot)$, we can move the minimization with respect to Λ^2 inside:

$$\hat{J}(x_1) \geq \inf_{\lambda_1} \rho_1^{\lambda_1} \left(c(x_1, \lambda_1(x_1), x_2) + \inf_{\Lambda^2} J_\infty(\Lambda^2, x_2) \right) = \inf_{\lambda_1} \rho_1^{\lambda_1} (c(x_1, \lambda_1(x_1), x_2) + \hat{J}(x_2)).$$

Observe that we assumed that $\hat{J}(\cdot)$ is measurable, and thus the right hand side of the last inequality is well-defined. Thus $\hat{J}(\cdot)$ satisfies the inequality:

$$\hat{J}(x) \geq \inf_{u \in U(x)} \sigma(c(x, u, \cdot) + \hat{J}(\cdot), x, Q(x, u)), \quad x \in \mathcal{X}. \quad (56)$$

We can now repeat the argument from the proof of Theorem 8.1. Denote by $\hat{\lambda}(x)$ the minimizer in (56), which exists by Proposition 5.2. Iterating inequality (56), and passing to the limit we conclude that

$$\hat{J}(x) \geq J_\infty(\hat{\Lambda}, x), \quad x \in \mathcal{X},$$

where $\hat{\Lambda} = \{\hat{\lambda}, \hat{\lambda}, \dots\}$ is a stationary Markov policy. Therefore, optimization with respect to stationary Markov policies is sufficient, and the result follows from Theorem 8.1. \square

To prove that our additional technical assumption that $\hat{J}(\cdot)$ is measurable is true for nonnegative costs $c(\cdot, \cdot, \cdot)$, we shall represent $\hat{J}(\cdot)$ as a limit of functions $v^k(\cdot)$ defined by the *value iteration method*:

$$v^{k+1} = \mathfrak{D}v^k, \quad k = 0, 1, 2, \dots, \quad (57)$$

where $v^0 = 0$.

Proposition 8.1 *Assume that conditions (G0)–(G4) are satisfied, the mapping $\sigma(\cdot, x, \cdot)$ is continuous, $c(\cdot, \cdot, \cdot)$ is nonnegative, and the model is uniformly risk-transient. Then the sequence $\{v^k\}$ defined by the value iteration method satisfies the relation*

$$\hat{J}(x) = \lim_{k \rightarrow \infty} v^k(x), \quad x \in \mathcal{X}.$$

Moreover, $\hat{J} \in \mathcal{V}$.

Proof. Owing to the non-negativity of the costs and to the monotonicity of $\sigma(\cdot, x, m)$, the sequence $\{v^k\}$ satisfies the inequalities:

$$v^k \leq v^{k+1} \leq \hat{J} \leq J^*, \quad k = 0, 1, 2, \dots, \quad (58)$$

where J^* is the best value of stationary Markov policies. Hence, this sequence has a pointwise limit v^∞ , which is also (by Lebesgue Theorem) a limit in \mathcal{V} . As $\sigma(\cdot, x, \cdot)$ is continuous, due to [4, Thm. 1.4.16], the operator \mathfrak{D} is continuous. It follows that

$$v^\infty = \lim_{k \rightarrow \infty} v^{k+1} = \lim_{k \rightarrow \infty} \mathfrak{D}v^k = \mathfrak{D}v^\infty.$$

By Theorem 8.1 and relations (58), $v^\infty = \hat{J} = J^*$. □

9 Randomized Decision Rules

Although we focus on deterministic policies in this paper, it may be of use to provide preliminary discussion of models with randomized policies. If randomized policies are allowed, we need to revise our definition of the risk transition mapping, to account for the fact that for a given state $x_t = x$, both u_t and x_{t+1} are random. In this case, it is convenient to consider functions on the product space $\mathcal{U} \times \mathcal{X}$ equipped with its product Borel σ -algebra \mathcal{B} .

Suppose the current state is x and we use a randomized decision rule π . This control, together with the transition kernel Q defines a probability measure $[\pi \circ Q]_x$ on the product space $\mathcal{U} \times \mathcal{X}$ as follows:

$$[\pi \circ Q]_x(B_u \times B_y) = \int_{B_u} Q(B_y|x, u) \pi(du|x), \quad B_u \in \mathcal{B}(U), \quad B_y \in \mathcal{B}(\mathcal{X}). \quad (59)$$

The measure is extended to other sets in \mathcal{B} in a usual way.

The cost incurred at the current stage is given by the function $c(x, \cdot, \cdot)$ on the product space $\mathcal{U} \times \mathcal{X}$.

The construction below parallels our earlier presentation in section 3, but with notational complication resulting from the need to deal with the product space $\mathcal{U} \times \mathcal{X}$. We define $\bar{\mathcal{V}} = \mathcal{L}_p(\mathcal{U} \times \mathcal{X}, \mathcal{B}, P_0)$, where $p \in [1, \infty)$, where \bar{P}_0 is some reference probability measure on $\mathcal{U} \times \mathcal{X}$. The dual space $\bar{\mathcal{V}}'$ is the space of signed measures m on $(\mathcal{U} \times \mathcal{X}, \mathcal{B})$, which are absolutely continuous with respect to \bar{P}_0 , with densities (Radon–Nikodym derivatives) lying in the space $\mathcal{L}_q(\mathcal{U} \times \mathcal{X}, \mathcal{B}, P_0)$, where $1/p + 1/q = 1$. We consider the set of probability measures in $\bar{\mathcal{V}}'$:

$$\mathcal{M} = \{m \in \bar{\mathcal{V}}' : m(\mathcal{U} \times \mathcal{X}) = 1, m \geq 0\}.$$

We also assume that the spaces $\bar{\mathcal{V}}$ and $\bar{\mathcal{V}}'$ are endowed with topologies that make them paired topological vector spaces with the bilinear form

$$\langle \varphi, m \rangle = \int_{\mathcal{U} \times \mathcal{X}} \varphi(u, y) m(du \times dy), \quad \varphi \in \bar{\mathcal{V}}, \quad m \in \bar{\mathcal{V}}'.$$

We can now define a *risk transition mapping* as a measurable function $\sigma : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}$, such that for every $x \in \mathcal{X}$ and every $m \in \mathcal{M}$, the function $\varphi \mapsto \sigma(\varphi, x, m)$ is a coherent measure of risk on \mathcal{V} .

Consider a controlled Markov process $\{x_t\}$ with some randomized Markov policy $\Pi = \{\pi_1, \pi_2, \dots\}$. For a fixed time t and a measurable w -bounded function $g : \mathcal{X} \times \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$, the value of $Z_{t+1} = g(x_t, u_t, x_{t+1})$, where u_t is distributed according to $\pi_t(x_t)$, is a random variable.

We define now a *Markov one-step conditional risk measure* $\rho_t(\cdot)$ by requiring that

$$\rho_t(g(x_t, u_t, x_{t+1})) = \sigma(g(x_t, \cdot, \cdot), x_t, [\pi_t \circ Q]_{x_t}), \quad \text{a.s.}$$

This is analogous to (4).

With these modifications, we can repeat the analysis presented in sections 3–8, with formal adjustments accounting for the use of randomized policies. For example, the dynamic programming equations (45)–(46) would take on the form:

$$v(x) = \inf_{\lambda(x) \in \mathcal{P}(U(x))} \sigma(c(x, \cdot, \cdot) + v(\cdot), x, [\lambda \circ Q]_x), \quad x \in \mathcal{X}, \quad (60)$$

$$v(x_\Delta) = 0. \quad (61)$$

Justification of these equations, though, is straightforward only in the case of *finite control sets* $U(x)$. Otherwise, serious technical difficulties arise, associated with the nonlinearity of the mapping in (60). With the present state of our knowledge, though, we may already make an observation that a randomized policy may be strictly better than a deterministic policy.

Observe that the mapping $\lambda(x) \mapsto \sigma(c(x, \cdot, \cdot) + v(\cdot), x, [\lambda \circ Q]_x)$, which plays the key role in the dynamic programming equation (60), is nonlinear, in general, as opposed to the expected value model, where

$$\sigma(c(x, \cdot, \cdot) + v(\cdot), x, [\lambda \circ Q]_x) = \int_{U(x)} \int_{\mathcal{X}} (c(x, u, y) + v(y)) Q(dy|x, u) \lambda(du|x).$$

In the expected value case, it is sufficient to consider only the extreme points of the set $\mathcal{P}(U(x))$, which are the measures assigning unit mass to one of the controls $u \in U(x)$:

$$\inf_{\lambda \in \mathcal{P}(U(x))} \int_{U(x)} \int_{\mathcal{X}} (c(x, u, y) + v(y)) Q(dy|x, u) \lambda(du|x) = \inf_{u \in U(x)} \int_{\mathcal{X}} (c(x, u, y) + v(y)) Q(dy|x, u).$$

In the risk averse case, this simplification is not justified and a randomized policy may be strictly better than any deterministic policy.

A question arises whether it is possible to identify cases in which deterministic policies are sufficient. It turns out that we can prove this for the conditional Average Value at Risk from Example 3.2, which in our setting takes on the following form:

$$\sigma(\varphi, x, \mu) = \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha} \int_{U(x) \times \mathcal{X}} (\varphi(u, y) - \eta)_+ \mu(du \times dy) \right\}, \quad \alpha \in (0, 1). \quad (62)$$

Lemma 9.1 *If the risk transition mapping has the form (62) then the dynamic programming equations (60) have a solution in deterministic decision rules.*

Proof. Interchanging the integration and the infimum in (62), we obtain a lower bound

$$\begin{aligned} \sigma(\varphi, x, [\lambda \circ Q]_x) &= \inf_{\eta \in \mathbb{R}} \int_{U(x)} \int_{\mathcal{X}} \{ \eta + (\varphi(u, y) - \eta)_+ \} Q(dy|x, u) \lambda(du|x) \\ &\geq \int_{U(x)} \inf_{\eta \in \mathbb{R}} \int_{\mathcal{X}} \{ \eta + (\varphi(u, y) - \eta)_+ \} Q(dy|x, u) \lambda(du|x). \end{aligned}$$

The above inequality becomes an equation for every Dirac measure λ . On the right-hand side of (60) we have

$$\inf_{\lambda(x) \in \mathcal{P}(U(x))} \sigma(c(x, \cdot, \cdot) + v(\cdot), x, [\lambda \circ Q]_x) \geq \inf_{\lambda \in \mathcal{P}(U(x))} \int_{U(x)} \inf_{\eta \in \mathbb{R}} \int_{\mathcal{X}} \{ \eta + (\varphi(u, y) - \eta)_+ \} Q(dy|x, u) \lambda(du|x).$$

As the right hand side achieves its minimum over $\lambda \in \mathcal{P}(U(x))$ at a Dirac measure concentrated at an extreme point of $U(x)$, and both sides coincide in this case, the minimum of the left hand side is also achieved at such measure. Consequently, for risk transition mappings of form (62) deterministic Markov policies are optimal. \square

In general, randomized policies may be better, as the example in section 10.2 illustrates.

10 Illustrative Examples

We illustrate our models and results on two examples.

10.1 Asset Selling

Let us at first consider the classical example of asset selling originating from Karlin [21]. Offers S_t arriving in time periods $t = 1, 2, \dots$ are independent integer-valued integrable random variables, distributed according to measure P . At each time we may accept the highest offer received so far, or we may wait, in which case a waiting cost c_0 is incurred. Denoting the random stopping time by τ we see that the total “cost” equals $Z = c_0\tau - \max_{0 \leq t \leq \tau} S_t$. The problem is an example of an *optimal stopping problem*, a structure of considerable theoretical and practical relevance (see, e.g., Çinlar [10], Dynkin and Yushkevich [12, 13], and Puterman [34]).

Formally, we introduce the state space $\mathcal{X} = \{x_A\} \cup \{0, 1, 2, \dots\}$, where x_A is the absorbing state reached after the transaction, and the other states represent the highest offer received so far. The control space is $\mathcal{U} = \{0, 1\}$, with 0 representing “wait” and 1 representing “sell.” The state evolves according to the equation

$$x_{t+1} = \begin{cases} \max(x_t, S_{t+1}) & \text{if } u_t = 0, \\ x_A & \text{if } u_t = 1. \end{cases}$$

Denoting by $F_S(\cdot)$ the distribution function of S , we can write the controlled transition kernel Q as follows:

$$Q(y|x, 0) = \begin{cases} P(y) & \text{if } y > x, \\ F_S(x) & \text{if } y = x, \\ 0 & \text{if } y < x, \end{cases} \quad Q(y|x, 1) = \begin{cases} 1 & \text{if } y = x_A, \\ 0 & \text{if } y \neq x_A. \end{cases}$$

The cost function is

$$c(x, u, y) = \begin{cases} c_0 & \text{if } u = 0, \\ -x & \text{if } u = 1. \end{cases}$$

The expected value version of this problem has a known solution: accept the first offer greater than or equal to the solution \hat{x} of the equation

$$c_0 = \sum_{s=0}^{\infty} (s - \hat{x})_+ P(s). \quad (63)$$

We shall solve the risk-averse version of the problem. We choose $P_0(y) = P(y)/2$ for $y \in \mathbb{N}$, and $P_0(x_A) = 1/2$. The space \mathcal{V} is the space of P_0 -integrable functions $v : \mathcal{X} \rightarrow \mathbb{R}$. The space \mathcal{M} is the space of probability measures μ on \mathcal{X} , for which

$$\sup \left\{ \frac{\mu(y)}{P_0(y)} : P_0(y) > 0, y \in \mathcal{X} \right\} < \infty.$$

Observe that all measures $Q(\cdot, x, u)$ are elements of \mathcal{M} .

Suppose $\sigma : \mathcal{V} \times \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}$ is a risk transition mapping. Owing to (A3), and to the fact that $v(x_A) = 0$, we have

$$\sigma(c(x, u, \cdot) + v(\cdot), x, Q(x, u)) = \begin{cases} c_0 + \sigma(v(\cdot), x, Q(x, 0)) & \text{if } u = 0, \\ -x & \text{if } u = 1. \end{cases}$$

Equation (45) takes on the form:

$$v(x) = \min \left\{ -x, c_0 + \sigma(v(\cdot), x, Q(x, 0)) \right\}, \quad x = 0, 1, 2, \dots \quad (64)$$

Suppose σ is law invariant (Definition 3.2). As the distribution of $v(\cdot)$ with respect to the measure $Q(x, 0)$ on \mathbb{N} is the same as the distribution of $v(\max(x, S))$ under the measure P of S , we obtain

$$\sigma(v(\cdot), x, Q(x, 0)) = \sigma(v(\max(x, \cdot)), x, P).$$

Suppose our attitude to risk does not depend on the current state, that is, $\sigma(\cdot, \cdot, \cdot)$ does not depend on its second argument. Using (6), we may rewrite the last equation as follows:

$$\sigma(v(\cdot), x, Q(x, 0)) = \max_{\mu \in \mathcal{A}} \sum_{s=0}^{\infty} v(\max(x, s)) \mu(s).$$

The convex closed set of probability measures \mathcal{A} is fixed, because $\sigma(\cdot, \cdot, \cdot)$ does not depend on its second argument, and the third argument, P , is now fixed. Equation (64) takes on the form

$$v(x) = \min \left\{ -x, c_0 + \max_{\mu \in \mathcal{A}} \sum_{s=0}^{\infty} v(\max(x, s)) \mu(s) \right\}, \quad x = 0, 1, 2, \dots \quad (65)$$

Observe that $v(x) \leq -x$ and thus $v(\max(x, s)) \leq -\max(x, s)$. The last displayed inequality implies that

$$v(x) \leq \min \left\{ -x, c_0 + \max_{\mu \in \mathcal{A}} \sum_{s=0}^{\infty} [-\max(x, s)] \mu(s) \right\} = \min \left\{ -x, c_0 - \min_{\mu \in \mathcal{A}} \sum_{s=0}^{\infty} \max(x, s) \mu(s) \right\}, \quad x = 0, 1, 2, \dots$$

If the offer at level x is accepted, then $v(x) = -x$. After simplifications, we obtain the inequality:

$$\min_{\mu \in \mathcal{A}} \sum_{s=0}^{\infty} (s - x)_+ \mu(s) \leq c_0.$$

This suggests the solution: *accept any offer that is greater or equal to the solution x^* of the equation*

$$\min_{\mu \in \mathcal{A}} \sum_{s=0}^{\infty} (s - x)_+ \mu(s) = c_0; \quad (66)$$

if $x < x^$, then wait.* The corresponding value function is $v^*(x) = -\max(x, x^*)$. Equation (65) can be verified by direct substitution.

Observe that the solution (66) of the risk-averse problem is closely related to the solution (63) of the expected value problem. The only difference is that we have to account for the least favorable distribution of the offers. If $P \in \mathcal{A}$ (which is the typical situation), then the critical level $x^* \leq \hat{x}$.

10.2 Organ Transplant

We illustrate our results on a risk-averse version of a simplified organ transplant problem discussed in Alagoz et. al. [1]. We consider the discrete-time absorbing Markov chain depicted in Figure 1. State S, which is the initial state, represents a patient in need of an organ transplant. State L represents life after a successful transplant. State D (absorbing state) represents death. Two control values are possible in state S: W (for “Wait”), in which case transition to state D or back to state S may occur, and T (for “Transplant”), which results in a transition to states L or D. The probability of death is lower for W than for T, but successful transplant may result in a longer life, as explained below. In other two states only one (formal) control value is possible: “Continue”.

The rewards collected at each time step are months of life. In state S a reward equal to 1 is collected, if the control is W; otherwise, the immediate reward is 0. In state L the reward $r(L)$ is collected, representing the sure equivalent of the random length of life after transplant. In state D the reward is 0.

Generally, in a cost minimization problem, the value of a dynamic measure of risk (1) is the “fair” sure charge one would be willing to incur, instead of a random sequence of costs. In our case, which will be a maximization problem, we shall work with the negatives of the months of life as our “costs.” The value of

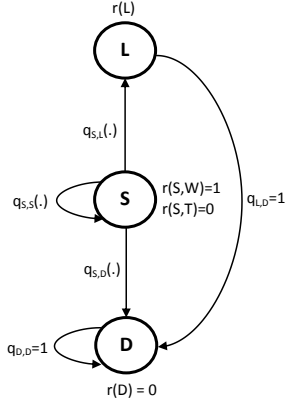


Figure 1: The organ transplant model.

the measure of risk, therefore, can be interpreted as the negative of a sure life length which we consider to be equivalent to the random life duration faced by the patient.

Let us start from describing the way the deterministic equivalent length of life $r(L)$ at state L is calculated. The state L is in fact an aggregation of n states in a survival model representing months of life after transplant, as depicted in Figure 2.

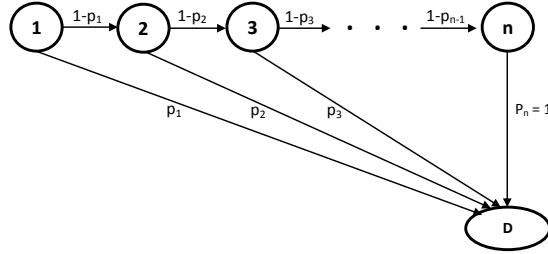


Figure 2: The survival model.

In state $i = 1, \dots, n$, the patient dies with probability p_i and survives with probability $1 - p_i$. The probability $p_n = 1$. The reward collected at each state $i = 1, \dots, n$ is equal to 1. In order to follow the notation of our paper, we define the cost $c(\cdot) = -r(\cdot)$. For illustration, we apply the mean-semideviation model of Example 3.1 with $\kappa = 1$.

The risk transition mapping has the form:

$$\sigma(\varphi, i, \nu) = \underbrace{\mathbb{E}_\nu[\varphi]}_{\text{expected value}} + \kappa \underbrace{\mathbb{E}_\nu[(\varphi - \mathbb{E}_\nu[\varphi])_+]}_{\text{semideviation}}. \quad (67)$$

Owing to the monotonicity property (A2), $\sigma(\varphi, i, \nu) \leq 0$, whenever $\varphi(\cdot) \leq 0$.

In (67), the measure ν is the transition kernel at the current state i , and the function $\varphi(\cdot)$ is the cost incurred at the current state and control plus the value function at the next state. At each state $i = 1, \dots, n - 1$ two transitions are possible: to D with probability p_i and $\varphi = -1$, and to $i + 1$ with probability $1 - p_i$ and $\varphi = -1 + v_{i+1}(i + 1)$. At state $i = n$ the transition to D occurs with probability 1, and $\varphi = -1$. Therefore, $v_n(n) = -1$.

The survival problem is a finite horizon problem, and thus we apply equation (17). As there is no control to choose, the minimization operation in is eliminated. The equation has the form:

$$v_i(i) = \sigma(\varphi, i, Q_i), \quad i = 1, \dots, n - 1,$$

with φ and Q_i as explained above. By induction, $v_i(i) \leq 0$, for $i = n-1, n-2, \dots, 1$.

Let us calculate the mean and semideviation components of (67) at states $i = 1, \dots, n-1$:

$$\begin{aligned}\mathbb{E}_{Q_i}[\varphi] &= -p_i + (1-p_i)(-1 + v_{i+1}(i+1)) = -1 + (1-p_i)v_{i+1}(i+1), \\ \mathbb{E}_{Q_i}[(\varphi - \mathbb{E}_{Q_i}[\varphi])_+] &= \mathbb{E}_{Q_i}[(\varphi + 1 - (1-p_i)v_{i+1}(i+1))_+] \\ &= p_i(-1 + 1 - (1-p_i)v_{i+1}(i+1))_+ + (1-p_i)(-1 + v_{i+1}(i+1) + 1 - (1-p_i)v_{i+1}(i+1))_+ \\ &= p_i(-(1-p_i)v_{i+1}(i+1))_+ + (1-p_i)(p_i v_{i+1}(i+1))_+ \\ &= -p_i(1-p_i)v_{i+1}(i+1).\end{aligned}$$

In the last equation we used the fact that $v_{i+1}(i+1) \leq 0$. For $i = 1, \dots, n-1$, the dynamic programming equations (17) take on the form:

$$v_i(i) = \underbrace{-1 + (1-p_i)v_{i+1}(i+1)}_{\text{expected value}} - \kappa \underbrace{p_i(1-p_i)v_{i+1}(i+1)}_{\text{semideviation}}, \quad i = n-1, n-2, \dots, 1.$$

The value $v(1)$ is the negative of the risk-adjusted length of life with new organ. For $\kappa = 0$ the above formulas give the negative of the expected length of life with new organ.

In our calculations we used the transition data provided in Table 1. They have been chosen for purely illustrative purposes and do not correspond to any real medical situation.

Table 1: Transition probabilities from state S.

Control	S	L	D
W	0.99882	0	0.00118
T	0	0.90782	0.09218

For the survival model, we used the distribution function, $F(x)$, of lifetime of the American population from Jasiulewicz [19]. It is a mixture of Weibull, lognormal, and Gompertz distributions:

$$F(x) = w_1 \left(1 - \exp\left(-\left(\frac{x}{\delta}\right)^\beta\right) \right) + w_2 \Phi\left(\frac{\log x - m}{\sigma}\right) + w_3 \left(1 - \exp\left(-\frac{b}{\alpha}(e^{\alpha x} - 1)\right) \right), \quad x \geq 0.$$

The values of the parameters and weights, provided by Jasiulewicz [19], are given in Table 2.

Table 2: Values of parameters for $F(x)$.

Distribution	Parameters	Weights
Weibull	$\delta = 0.297, \beta = 0.225$	$w_1 = 0.0170$
Lognormal	$m = 3.11, \sigma = 0.218$	$w_2 = 0.0092$
Gompertz	$b = 0.0000812, \alpha = 0.0844$	$w_3 = 0.9737$

Then, we calculated the probability of dying at age k (in months) as follows:

$$p_k = \frac{F(k/12 + 1/24) - F(k/12 - 1/24)}{1 - F(k/12 - 1/24)}, \quad k = 1, 2, \dots$$

The maximum lifetime of the patient was taken to be 1200 months, and that the patient after transplant has survival probabilities starting from $k = 300$. Therefore, $n = 900$ in the survival model used for calculating $r(L)$.

Let $\lambda = (\lambda_W, \lambda_T)$ be the randomized policy in the state S and let $\Lambda = \{\lambda \in \mathbb{R}^2 : \lambda_W + \lambda_T = 1, \lambda \geq 0\}$. The dynamic programming equation (45) at S takes on the form

$$v(S) = \min_{\lambda \in \Lambda} \left\{ \underbrace{\lambda_W [q_{S,S}(W)(v(S) - 1) + q_{S,D}(W)(v(D) - 1)] + \lambda_T [q_{S,L}(T)v(L) + q_{S,D}(T)v(D)]}_{\text{expected value } \mu} + \kappa \left(\underbrace{\lambda_W [q_{S,S}(W)(v(S) - 1 - \mu)_+ + q_{S,D}(W)(v(D) - 1 - \mu)_+]_{\text{semideviation } \dots}} + \lambda_T \underbrace{[q_{S,L}(T)(v(L) - \mu)_+ + q_{S,D}(T)(v(D) - \mu)_+]_{\text{... semideviation}}} \right) \right\}.$$

In the semideviation parts, we wrote μ for the expectation of the value function in the next state, which is given by the first underbraced expression, and which is also dependent on λ . Of course, the above expression can be simplified, by using the fact that $v(L) < v(S) < v(D) = 0$, but we prefer to leave it in the above form to illustrate the way it has been developed.

We compared two optimal control models for this problem. The first one was the expected value model ($\kappa = 0$), which corresponds to the expected reward $r(L) = 610.46$ in the survival model. Standard dynamic programming equations were solved, and the optimal decision in state S turned out to be W.

The second model was the risk-averse model using the mean–semideviation risk transition mapping with $\kappa = 1$. This changed the reward at state L to 515.35. We considered two versions of this model. In the first version, we restricted the feasible policies to be deterministic. In this case, the optimal action in state S was T. In the second version, we allowed randomized policies, as in our general model. Then the optimal policy in state S was W with probability $\lambda_W = 0.9873$ and T with probability $\lambda_T = 0.0127$.

How can we interpret these results? The optimal randomized policy results in a random waiting time before transplanting the organ. This is due to the fact that immediate transplant entails a significant probability of death, and a less risky policy is to “dilute” this probability in a long waiting time. This cannot be derived from an expected value model, no matter what the data, because deterministic policies are optimal in such a model: either transplant immediately or never.

References

- [1] Alagoz, O., L. M. Maillart, A. J. Schaefer and M. S. Roberts, The optimal timing of living-donor liver transplantation, *Management Science*, 50, 1420–1430, 2004.
- [2] Artzner, P., F. Delbaen, J. M. Eber and D. Heath, Coherent measures of risk, *Mathematical Finance*, 9, 203–228, 1999.
- [3] Artzner, P., F. Delbaen, J.-M. Eber, D. Heath, and H. Ku, Coherent multiperiod risk adjusted values and Bellmans principle, *Annals of Operations Research* 152, 5–22, 2007.
- [4] Aubin, J.-P., and H. Frankowska, *Set-Valued Analysis*, Birkhäuser, Boston, 1990.
- [5] N. Bäuerle, U. Rieder, *Markov Decision Processes with Applications to Finance*, Springer, Berlin, 2011.
- [6] Bertsekas, D., and S. E. Shreve, *Stochastic Optimal Control. The Discrete Time Case*, Academic Press, New York, 1978.
- [7] Bertsekas, D. P., and Tsitsiklis J. N., An analysis of stochastic shortest-path problems, *Mathematics Of Operations Research*, 16 (3), 580–595, 1991.

- [8] Cheridito, P., F. Delbaen, and M. Kupper, Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability* 11, 57–106, 2006.
- [9] Chew, S. H. and J. L. Ho, Hope: An empirical study of attitude toward the timing of uncertainty resolution, *Journal of Risk and Uncertainty* 8 (3), 267–288, 1994.
- [10] Çinlar, E., *Introduction to Stochastic Processes*, Prentice-Hall, Englewood Cliffs, 1975.
- [11] Denardo, E. V. and U. G. Rothblum, Optimal stopping, exponential utility, and linear programming, *Mathematical Programming* 16, 228–244, 1979.
- [12] Dynkin, E.B., and A.A. Yushkevich, *Markov Processes: Theory and Problems*, Plenum, New York, 1969.
- [13] Dynkin, E.B., and A.A. Yushkevich, *Controlled Markov Processes*, Springer-Verlag, New York, 1979.
- [14] Feinberg, E. A., and A. Shwartz, *Handbook of Markov Decision Processes: Methods and Applications*, Kluwer, Dordrecht, 2002.
- [15] Frittelli, M., and G. Scandolo, Risk measures and capital requirements for processes. *Mathematical Finance*, 16, 589–612, 2006.
- [16] González-Trejo, J. I., O. Hernández-Lerma, and L. F. Hoyos-Reyes, Minimax control of discrete-time stochastic systems, *SIAM J. Control Optim.*, 41, 1626-1659, 2003.
- [17] Hernández-Lerma, O., and J. B. Lasserre, *Discrete-Time Markov Control Processes. Basic Optimality Criteria*, Springer, New York, 1996.
- [18] Hernández-Lerma, O., and J. B. Lasserre, *Further Topics on Discrete-Time Markov Control Processes*, Springer, New York, 1999.
- [19] Jasiulewicz, H., Application of mixture models to approximation of age-at-death distribution, *Mathematics and Economics*, 19, 237–241, 1997.
- [20] A. Jaśkiewicz, A. S. Nowak, Stochastic games with unbounded payoffs: applications to robust control in economics, *Dyn. Games Appl.* 1, 253279, 2011.
- [21] Karlin, S., Stochastic models and optimal policies for selling an asset, in: *Studies in Applied Probability and Management Science*, K. J. Arrow, S. Karlin, and S. Scarf (Eds.), Stanford University Press, Palo Alto, 1962, pp. 148–158.
- [22] Klöppel, S., and M. Schweizer, Dynamic indifference valuation via convex risk measures, *Math. Finance*, 17, 599–627, 2007.
- [23] K. Kuratowski, C. Ryll-Nardzewski, A general theorem on selectors, *Bulletin de l'Académie Polonaise de Sciences, Série Mathématiques, Astronomie et Physique* 13, 397–411, 1965.
- [24] Kurt, M., and J. P. Kharoufeh, Monotone optimal replacement policies for a Markovian deteriorating system in a controllable environment, *Operations Research Letters*, 38, 273–279, 2010.
- [25] Nie, Y., and Wu, X., Shortest path problem considering on-time arrival probability, *Transportation Research B*, 43 (6), 597–613, 2009.
- [26] Nilim, A., and El Ghaoui, L., Robust control of Markov decision processes with uncertain transition matrices, *Operations Research* 53 (5), 780798, 2005.

- [27] Ohtsubo, Y., Minimizing risk models in stochastic shortest path problems, *Mathematical Methods of Operations Research*, 57 (1), 79–88, 2003.
- [28] Ogryczak, W., and A. Ruszczyński, From stochastic dominance to mean-risk models: Semideviations as risk measures, *European Journal of Operational Research*, 116, 33–50, 1999.
- [29] Ogryczak, W., and A. Ruszczyński, On consistency of stochastic dominance and mean-semideviation models, *Mathematical Programming*, 89, 217–232, 2001.
- [30] Ogryczak, W., and A. Ruszczyński, Dual stochastic dominance and related mean-risk models, *SIAM Journal on Optimization*, 13(1), 60–78, 2002.
- [31] Patek, S. D., On terminating Markov decision processes with a risk averse objective function, *Automatica*, 37(9), 1379–1386, 2001.
- [32] Pflug, G. Ch., and W. Römisch. *Modeling, Measuring and Managing Risk*. World Scientific, Singapore, 2007.
- [33] Pliska, S. R., On the transient case for Markov decision chains with general state spaces, in: *Dynamic Programming and Its Applications*, M. L. Puterman (ed.), Academic Press, New York, 1979, pp. 335–349.
- [34] Puterman, M. L., *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, New York, 1994.
- [35] Riedel, F., Dynamic coherent risk measures, *Stochastic Processes and Their Applications*, 112, 185–200, 2004.
- [36] R. T. Rockafellar. *Conjugate duality and optimization*. Society for Industrial and Applied Mathematics, Philadelphia, Pa., 1974. Lectures given at the Johns Hopkins University, Baltimore, Md., June, 1973, Conference Board of the Mathematical Sciences Regional Conference Series in Applied Mathematics, No. 16.
- [37] R. T. Rockafellar and S. P. Uryasev, Conditional value-at-risk for general loss distributions, *Journal of Banking and Finance*, 26, 1443–1471, 2002.
- [38] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*, Springer, Berlin, 1998.
- [39] Ruszczyński, A., Risk-averse dynamic programming for Markov decision processes, *Mathematical Programming*, Series B, 125, 235–261, 2010.
- [40] Ruszczyński, A. and A. Shapiro, Optimization of risk measures, In *Probabilistic and Randomized Methods for Design under Uncertainty*, G. Calafiore and F. Dabbene (Eds.), Springer, London, 2005.
- [41] Ruszczyński, A. and A. Shapiro, Optimization of convex risk functions, *Mathematics of Operations Research*, 31, 433–452, 2006.
- [42] Ruszczyński, A. and A. Shapiro, Conditional risk mappings, *Mathematics of Operations Research*, 31, 544–561, 2006.
- [43] Scandolo, G., *Risk Measures in a Dynamic Setting*, PhD Thesis, Università degli Studi di Milano, Milan, 2003.
- [44] Schäl, M., Conditions for optimality in dynamic programming and for the limit of n-stage optimal policies to be optimal, *Probability Theory and Related Fields* 32, 179–196, 1975.

- [45] Shapiro, A., D. Dentcheva and A. Ruszczyński, *Lectures on Stochastic Programming*, SIAM Publications, Philadelphia 2009.
- [46] So, M.M.C. and L.C.. Thomas, Modelling the profitability of credit cards by Markov decision processes, *European Journal of Operational Research*, 212, 123–130, 2011.
- [47] Tapiero, C. S. and I. Venezia, A Mean Variance Approach to the Optimal Machine Maintenance and Replacement Problem, *The Journal of the Operational Research Society*, 30, 457–466, 1979.
- [48] Veinott, A. F., Discrete dynamic programming with sensitive discount optimality criteria, *Annals of Mathematical Statistics*, 40, 1635–1660, 1969.
- [49] Wessels, J., Markov programming by successive approximations with respect to weighted supremum norms, *J. Math. Anal. Appl.* 58, 326-335, 1977.
- [50] White, D. J., A survey of applications of Markov decision processes, *Journal of Operational Research Society*, 44, 1073–1096, 1993.