August 2023

The Vanishing Discount Method for Stochastic Control: A Linear Programming Approach

Brian Hospital
University of Wisconsin-Milwaukee

THE VANISHING DISCOUNT METHOD FOR STOCHASTIC CONTROL: A LINEAR PROGRAMMING APPROACH

by

Brian J. Hospital

A Dissertation Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

in

Mathematics

at

The University of Wisconsin-Milwaukee

August 2023

ABSTRACT

THE VANISHING DISCOUNT METHOD FOR STOCHASTIC CONTROL: A LINEAR PROGRAMMING APPROACH

by

Brian J. Hospital

The University of Wisconsin-Milwaukee, 2023
Under the Supervision of Professor Richard H. Stockbridge

Under consideration are convergence results between optimality criteria for two infinite-horizon stochastic control problems: the long-term average problem and the α-discounted problem, where $\alpha \in (0, 1]$ is a given discount rate. The objects under control are those stochastic processes that arise as (relaxed) solutions to a controlled martingale problem; and such controlled processes, subject to a given budget constraint, comprise the feasible sets for the two stochastic control problems.

In this dissertation, we define and characterize the expected occupation measures associated with each of these stochastic control problems, and then reformulate each problem as an equivalent linear program over a space of such measures. We then establish sufficient conditions under which the long-term average linear program can be "asymptotically approximated" by the α-parameterized family of (suitably normalized) α-discounted linear programs as $\alpha \downarrow 0$. This approach is what can be referred to as the vanishing discount method. To state these conditions precisely, our analysis turns to set-valued mappings called correspondences. In particular, once we establish the appropriate framework, we see that our main results can be stated in a manner similar to that of Berge's Theorem.

First of all, I wish to thank my advisor, Professor Richard H. Stockbridge, without whom this dissertation would—almost surely—not exist.

I would like to further thank Professors Chao Zhu, David Spade, Kevin McLeod, and Peter Hinow for agreeing to serve on my dissertation committee and for the various insightful and encouraging comments they have provided.

I also wish to thank my parents for their persistent support of virtually every pursuit of mine, including—but certainly not limited to—my interest in mathematics.

Finally, I would like to thank my wife, Sarah C. Schaefer, for her guidance, companionship, and understanding during my time as a graduate student.

Brian J. Hospital

TABLE OF CONTENTS

LIST OF FIGURES

INTRODUCTION

I.1 Motivation and Overview

The objective of this dissertation is to analyze the structural relationship between two infinite-horizon stochastic control problems: the *long-term average problem* and the α-*discounted problem*, where $\alpha \in (0, 1]$ represents a given *discount rate* (or *discount factor*). Solutions to these problems are obtained by identifying optimal control *policies* that minimize, respectively, a *long-term average expected cost* and an α-*discounted expected cost.* We further augment the canonical formulation of these control problems by considering an additional *budget constraint* (or *resource constraint*). Note that one can equivalently seek to *maximize* a long-term average (or α-discounted) expected *reward* in an equivalent manner; our use of the minimization framework throughout this dissertation can be considered essentially arbitrary.

The long-term average expected cost can, informally, be defined as a limit of *finite-horizon* average expected costs as time extends to the infinite horizon. For a given time $t > 0$, the corresponding finite-horizon average expected cost is simply the expected cumulative cost the controller has obtained (under the adoption of some policy) up until time t divided by t (and, hence, this is more precisely described as a cost *rate* that we wish to minimize). By then passing to a limit (or, if necessary, a limit superior) as $t \to \infty$, we arrive at the long-term average expected cost.

The α-discounted expected cost is perhaps best understood in economic terms—that is, as the *present (discounted) value* of the total costs accumulated by the controller over all time. Note that the role of the discount rate $\alpha > 0$ is integral since, otherwise, unbounded total costs would be realized under too large a class of (potentially suboptimal) policies; moreover, in many practical settings (e.g., financial applications), a discount rate is a realistic means of accounting for inflationary effects and devaluation. For example, a controller's accumulation of wealth over time is a natural interpretation of the "rewards" accrued by the process in the *maximization* formulation of this scheme. Furthermore, under this interpretation, the inclusion of a budget/resource constraint carries a degree practical significance; e.g., a firm may wish to find production levels of a commodity that result in the largest profit without exceeding the supply of that commodity. However, this approach is not without its limitations. Indeed, applications have been identified in which such discounting results in suboptimal policies because of myopia. For example, the optimal solution to the discounted problem in certain harvesting applications can result in premature extinction of a species, as demonstrated in, e.g., Alvarez and Shepp (1998) and Song et al. (2011).

As we wish to analyze and compare the structure of these two problems, the question as to when (or if) they, in fact, yield the same optimal values (or optimal policies) naturally arises. The relationship between the optimal values is well known and is a consequence of the so-called *Abelian theorems* in which the optimal value of the long-term average problem is obtained as the limit of the (suitably normalized) α-discounted optimal values as $\alpha \downarrow 0$; and similar results regarding the optimal (and feasible) policies of the problems can be established under additional conditions, which we discuss in detail below. These facts comprise what can be referred to as the *vanishing discount method*. This approach has been extensively studied and documented since the initial developments of Markov decision processes in the middle of the twentieth century; see, for example, Taylor (1965) and Veinott (1966). In essence, this approach allows for the optimal values for one problem to be approximated by the optimal values for the other (in, say, instances where one problem is more tractable than the other). A rigorous analysis of this method will be among the primary concerns of this dissertation.

While the vanishing discount method provides a way of relating certain structural properties of the long-term average problem and the α-discounted problem (including their optimality criteria), in practice, at least one of these problems must, naturally, be solved. Perhaps the most well-developed and popular ways of doing so has been via dynamic programming techniques. In particular, under suitable conditions, the optimal value of the long-term average problem can be realized as a solution to the *Hamilton-Jacobi-Bellman (HJB) equation*, as detailed in Fleming and Rishel (1975), and more recently in Pham (2009). As a second-order elliptic partial differential equation, the Hamilton-Jacobi-Bellman equation can often be quite difficult to solve explicitly—often admitting only viscosity solutions rather than classical solutions; see Kawaguchi and Morimoto (2007). An alternative approach involves expressing solutions to each of the problems in terms of so-called *expected occupation measures* and then reformulating the stochastic problems as equivalent *linear programs* over (possibly infinite-dimensional) spaces of these measures. This approach has its history in the early development of decision processes as presented in, e.g., Manne (1960); but the more recent developments described in Kurtz and Stockbridge (1998) will serve as the basis for our approach to solving the stochastic control problems considered throughout this dissertation.

In order to characterize the appropriate expected occupation measures—and, hence, construct the desired linear programs—we first identify those controlled processes that arise as (relaxed) solutions to a *controlled martingale problem*. For this analysis, we introduce the *generator* L of a controlled process, a linear operator on a space of functions that characterizes the dynamics of a stochastic process through its instantaneous transitional behavior. The use of this operator will both allow for more efficient exposition and for the contribution of the techniques of operator theory to our analysis. In particular, we will see that each collection of expected occupation measures can be fully characterized as those measures that satisfy a particular *adjoint equation* for the operator L.

By formulating the long-term average and α-discounted problems in terms of their respective linear programs, our task is then to use these linear programs to determine exactly what conditions must be imposed on our model in order for the vanishing discount method to be applicable. Since each $\alpha \in (0, 1]$ corresponds to an α-discounted linear program, application of the vanishing discount method in this context involves analysis of the asymptotic behavior of an α-parameterized family of linear programs as $\alpha \downarrow 0$. To obtain the desired convergence, however, we will see that one must first *normalize* each α-discounted expected occupation measure by its associated discount rate α to then obtain an *α-normalized (discounted) expected occupation measure*. Consequently, much of our analysis will focus on the (α-parameterized) family of *α-normalized linear programs* as $\alpha \downarrow 0$. To be precise, it will be necessary for us to identify the appropriate *feasible set* (a set of α-normalized expected occupation measures), the set of *optimal solutions* (a subset of the feasible set), and the *optimal value* for each α-normalized linear program. We then seek to establish conditions under which each of these objects "converges" to, respectively, the feasible set, the set of optimal solutions, and the optimal value for the long-term average linear program.

Since we will need to examine the asymptotic behavior of families of *sets* (i.e., the feasible sets and the sets of optimal solutions for the α-normalized linear programs), we will require a suitable notion of convergence that agrees with the mechanics of the vanishing discount method. One such notion of convergence can be expressed in terms of continuity of set-valued functions called *correspondences*, the theory of which is extensively detailed in Aliprantis and Border (2006). Indeed: By defining a correspondence on the interval $[0, 1]$ that maps each $\alpha \in (0, 1]$ to the associated set of optimal solutions for the α-normalized linear program, and maps $\alpha = 0$ to the set of optimal solutions for the long-term average linear program, verifying that the vanishing discount method can be applied then becomes a matter of verifying that this correspondence is continuous at the point $\alpha = 0$. Here, we note that a correspondence is said to be *continuous* at a point (in a topological space) if it is both *lower hemicontinuous* and *upper hemicontinuous* at that point. As this is a matter that is of great importance to our results, we have devoted a section of this dissertation to discuss the concept of hemicontinuity in detail.

The application of correspondences to our problem not only provides us with a convenient language and an elegant framework with which to carry out our analysis, it also allows us to exploit a number of powerful results associated with the theory of correspondences. In particular, we will see that *Berge's Theorem* serves as a rather useful model in this setting. Among its conclusions is, for example, the statement that the "value function" (which assigns, to each $\alpha \in [0,1]$, the optimal value for the corresponding linear program) is continuous. As the vanishing discount method relies upon stability of the optimal values as the discount rate α "vanishes," the objective of establishing conditions that ensure this value function is continuous at $\alpha = 0$ will naturally play a significant role in our analysis.

After establishing the main theoretical results of this dissertation, we present a number of examples intended to demonstrate both the utility of the vanishing discount method and the subtleties with which one must contend before applying it. Among our examples is an illuminating "geometric" example appearing in Hernández-Lerma and Prieto-Rumeau (2010) that we have recast using our linear programming framework, and then expanded upon to emphasize some of the nuances of our results. In fact, this paper and its accompanying example served as one of the primary inspirations for the research conducted for this dissertation, as it is one of the few to consider the vanishing discount method for *constrained* processes. We then turn our attention to a more classical problem from control theory in which one considers a controlled diffusion process. For this example, we analyze solutions to a typical stochastic differential equation (SDE), characterize the solutions to the SDE in terms of the appropriate generator, and then identify the corresponding long-term average and α-normalized linear programs. In each example presented, we analyze the hypotheses and conclusions of our main results to demonstrate how they may perhaps be applied in practice.

The template used to produce this manuscript was modified from Vieten (2018) and, accordingly, we use virtually the same numbering scheme for definitions, equations, lemmas, propositions, theorems, corollaries, remarks, and examples. As such, we will simply reference—essentially verbatim, with only minor syntactical and stylistic changes—the following description of this scheme found in that document: We use a continuous numbering scheme for equations, lemmas, propositions, theorems, corollaries, remarks and examples in the following. These objects will be referenced by a roman numeral indicating the chapter, followed by an arabic number indicating the section and a second arabic number representing the consecutive number of the object in that section, separated by a period, respectively. To give an example, "(III.2.1)" is the first object, an equation, appearing in the second section of Chapter III. The next object in this section is a definition—and is therefore referenced by "Definition III.2.2." However, for the sake of readability, the roman numerals are omitted when the object lies in the same chapter in which it is referenced. So, the aforementioned "Definition III.2.2" will appear as "Definition 2.2" in Chapter III, and as "Definition III.2.2" in any other chapter.

The Appendix is handled much like a "Chapter," however the *number* that usually indicates a particular section is instead a capital *letter*. For example, "Lemma B.7" refers to the seventh object, a lemma, appearing in section B of the Appendix. Note that, since this manner of enumeration is unique to the Appendix, there is no need to prepend each object with the corresponding "Chapter" number when objects in the Appendix are referenced in the main body of the document.

Lemmas, propositions, theorems, and corollaries are presented in italicized text, whereas remarks and examples are not. Definitions are not italicized, except for the term being defined. Consequently, we use the symbol $\diamond$ to terminate remarks, examples, and definitions to make it apparent where the main body of text is meant to be resumed. In a similar manner, we use the usual symbol $\square$ to indicate the end of a proof.

I.2 Dynamics of the Controlled Process

To describe the control models of interest, we consider a linear operator $L : \mathscr{D}(L) \to \mathscr{R}(L)$, called a *generator*, and those pairs $(X, \Lambda) = \{(X_t, \Lambda_t) : t \in \mathbb{R}^+\}$ such that the process $M := \{M_t : t \in \mathbb{R}^+\}$ defined by

$$M_t := f(X_t) - f(X_0) - \int_0^t \int_G Lf(X_s, u)\, \Lambda_s(du)\, ds \tag{2.1}$$

is a martingale for every $f \in \mathscr{D}(L)$. In (2.1), the process $X = \{X_t : t \in \mathbb{R}^+\}$ is an E-valued *state process* and $\Lambda = \{\Lambda_t : t \in \mathbb{R}\}$ is a $\mathcal{P}(G)$-valued *control policy* (or, simply, *control*), where $\mathcal{P}(S)$ denotes the set of Borel probability measures on the measurable space S. Both E and G are assumed to be locally compact, complete, separable, metric spaces (noting that, as a consequence, $E \times G$ is a locally compact, complete, separable, metric space). Being locally compact, we let $E^\Delta := E \cup \{\Delta\}$ denote the one-point compactification of E, where we assume that $\Delta \notin E$.

To characterize the domain $\mathscr{D}(L)$ of the operator L, we let $C(S)$ denote the set of continuous real-valued functions on the set S, and let $\widehat{C}(S)$ denote the set of $f \in C(S)$ that vanish at infinity. The set $\mathscr{D}(L)$ is taken to be a subset of $\widehat{C}(E)$ for which the following conditions are satisfied:

(D1) $\mathscr{D}(L)$ is dense in $\widehat{C}(E)$.

(D2) $L_u f := Lf(\cdot, u) \in \widehat{C}(E)$ for each $f \in \mathscr{D}(L)$ and each $u \in G$.

(D3) $\limsup_{x \to \Delta} \left(\sup\{|Lf(x,u)| : u \in K\}\right) = 0$ for each $f \in \mathscr{D}(L)$ and each compact $K \subset G$.

(D4) $L_u f$ satisfies the positive maximum principle for each $u \in G$; i.e.,

$$L_u f(x_0) \leq 0 \quad \text{whenever} \quad \sup\{f(x) : x \in E\} = f(x_0) \geq 0.$$

(D5) $\mathscr{D}(L)$ is an algebra.

(D6) There exists a $\psi \in C(E \times G)$ with $\psi \geq 1$ such that, for each $f \in \mathscr{D}(L)$, there is a constant $a_f \in \mathbb{R}^+$ satisfying

$$|Lf(x,u)| \leq a_f \psi(x,u), \qquad \forall (x,u) \in E \times G.$$

In accordance with the conditions imposed on $\mathscr{D}(L)$, we can view the range $\mathscr{R}(L)$ as a subset of $C(E \times G)$. As discussed in Kurtz and Stockbridge (1998), the above conditions guarantee the existence (at least in E^Δ) of *relaxed solutions* to the *controlled martingale problem for* L. Such solutions are precisely those processes (X, Λ) satisfying (2.1) above, the description of which we formalize by means of the following definition.

Definition 2.2. We say that an $E \times \mathcal{P}(G)$-valued process $(X, \Lambda) = \{(X_t, \Lambda_t) : t \in \mathbb{R}^+\}$ is a *relaxed solution* to the *controlled martingale problem for* L if the following hold.

(M1) There exists a filtration $\{\mathcal{F}_t\} \coloneqq \{\mathcal{F}_t : t \in \mathbb{R}^+\}$ such that (X, Λ) is $\{\mathcal{F}_t\}$-progressive.

(M2) For each $f \in \mathscr{D}(L)$, the process M in (2.1) is an $\{\mathcal{F}_t\}$-martingale.

The notation $\mathscr{M} \coloneqq \mathscr{M}(L)$ will be used to refer to the set of such solutions. $\diamond$

Remark 2.3. Recall that a process (X, Λ) is $\{\mathcal{F}_t\}$*-progressive* (or *progressively measurable with respect to* $\{\mathcal{F}_t\}$) if, for each $t \in \mathbb{R}^+$ and each $\Gamma \in \mathcal{B}(E \times \mathcal{P}(G))$, the set

$$\{(s, \omega) \in [0, t] \times \Omega : (X_s(\omega), \Lambda_s(\omega)) \in \Gamma\}$$

belongs to the σ-algebra $\mathcal{B}([0, t]) \times \mathcal{F}_t$. It is clear that an $\{\mathcal{F}_t\}$-progressive process is also $\{\mathcal{F}_t\}$-adapted (i.e., "progressive" is a strictly stronger property than "adapted"). Moreover, as shown in Chung and Doob (1965) or Meyer (1966), every $\{\mathcal{F}_t\}$-adapted process has a $\{\mathcal{F}_t\}$-progressive modification; and, as shown in Karatzas and Shreve (1991), every $\{\mathcal{F}_t\}$-adapted process with càdlàg paths is $\{\mathcal{F}_t\}$-progressive. Our requirement that each $(X, \Lambda) \in \mathscr{M}$ be $\{\mathcal{F}_t\}$-progressive (rather than simply $\{\mathcal{F}_t\}$-adapted) is included in order to guarantee that, for each $\{\mathcal{F}_t\}$-stopping time τ, the corresponding stopped process $(X, \Lambda)^\tau$ defined by $(X_t, \Lambda_t)^\tau \coloneqq (X_{t\wedge\tau}, \Lambda_{t\wedge\tau})$ is $\{\mathcal{F}_t\}$-measurable, as shown in Ethier and Kurtz (1986). $\diamond$

Now, for each $(X, \Lambda) \in \mathscr{M}$, we furthermore wish to consider only those controls that are "available" when the state process X occupies a particular state (i.e., for each $s \in \mathbb{R}^+$, we wish to consider only those controls that have a positive probability of being chosen when $X_s = x$). We can formalize this notion in the following manner.

Let $\mathcal{U}$ be a closed subset of $E \times G$ and, for each $x \in E$, define the set $G(x)$ by

$$G(x) \coloneqq \{u \in G : (x, u) \in \mathcal{U}\}.$$

We then say that the control Λ is *admissible* if

$$\int_0^t \int_G I_{\mathcal{U}}(X_s, u)\, \Lambda_s(du)\, ds = t, \qquad \forall t \in \mathbb{R}^+,$$

where I_S denotes the indicator function of the set S. In essence, this admissibility condition ensures that, for each $s \in \mathbb{R}^+$, the support of the measure $\Lambda_s \in \mathcal{P}(G)$ is a subset of $G(x)$ when $X_s = x$. Hence, for virtually all intents and purposes, the relation

$$G = \bigcup_{x \in E} G(x)$$

may be considered a suitable characterization of the control space G. However, for clarity of exposition, we will usually use $E \times G$ and $\mathcal{U}$ interchangeably throughout this dissertation.

I.3 Optimality Criteria

The main purpose of this dissertation is to analyze the structural relationship between the *long-term average stochastic problem* and, for each $\alpha \in (0, 1]$, the *α-discounted stochastic problem*. To describe each of these control problems, we begin with a given *constrained control model* $\{E \times G, L, c, c_1\}$. In this model, $E \times G$ and L are as above; and $c, c_1 \in M(E \times G)$ are, respectively, a given *cost rate* function and a given *budget rate* function, where $M(S)$ denotes the set of $\mathcal{B}(S)$-measurable real-valued functions on the set S. Note that, under this formulation, an *unconstrained* control model can be denoted by $\{E \times G, L, c\}$.
Let us now briefly recall the following fairly standard definitions, the first of which comes (with some augmentation) from Section 2.10 in Aliprantis and Border (2006).

Definition 3.1. Let $f : S \to [-\infty, +\infty]$ be a function on a topological space S.

- f is *lower semicontinuous* if, for each $\lambda \in \mathbb{R}$, the set $\{a \in S : f(a) \leq \lambda\}$ is closed (or, equivalently, the set $\{a \in S : f(a) > \lambda\}$ is open). We can also say that f is *lower semicontinuous* at $a_0 \in S$ if and only if

$$f(a_0) \leq \liminf_{a \to a_0} f(a).$$

- *upper semicontinuous* if, for each $\lambda \in \mathbb{R}$, the set $\{a \in S : f(a) \geq \lambda\}$ is closed (or, equivalently, the set $\{a \in S : f(a) < \lambda\}$ is open). We can also say that f is *upper semicontinuous* at $a_0 \in S$ if and only if

$$\limsup_{a \to a_0} f(a) \leq f(a_0).$$

Note that such a function f is continuous if and only if it is both lower semicontinuous and upper semicontinuous. ◇

The next definition comes from Ricceri (2016).

Definition 3.2. A function $f : S \to \mathbb{R}$ on a topological space S is *inf-compact* if, for each $\lambda \in \mathbb{R}$, the set $\{a \in S : f(a) \leq \lambda\}$ is compact. Observe that an inf-compact function on a Hausdorff space (such as $E \times G$) is necessarily lower semicontinuous. ◇

We then assume that the functions c and c_1 satisfy the following conditions.

(C1) c and c_1 are bounded below; i.e., there exist constants $\kappa, \kappa_1 \in \mathbb{R}^+$ such that

$$-\kappa < c(x,u) \quad \text{and} \quad -\kappa_1 < c_1(x,u), \qquad \forall (x,u) \in E \times G.$$

(C2) c and c_1 are inf-compact (and, hence, lower semicontinuous).

(C3) At least one of the following hold:

- There exist constants $a, b \in \mathbb{R}^+$ and $\beta \in (0,1)$ such that

$$\psi(x,u) \leq ac^{\beta}(x,u) + b, \qquad \forall (x,u) \in E \times G.$$

- There exist constants $a_1, b_1 \in \mathbb{R}^+$ and $\beta_1 \in (0,1)$ such that

$$\psi(x,u) \leq a_1 c_1^{\beta_1}(x,u) + b_1, \qquad \forall (x,u) \in E \times G.$$

Note that ψ is the same function introduced in condition (D6) above.

To be precise, we interpret the function c^{β} (and, *mutatis mutandis*, $c_1^{\beta_1}$) as follows:

$$c^{\beta}(x,u) := \begin{cases} -|c(x,u)|^{\beta}, & \text{if } c(x,u) < 0; \\ (c(x,u))^{\beta}, & \text{if } c(x,u) \geq 0. \end{cases}$$

Remark 3.3. Recall that we will define each optimization problem as a *minimization* problem. However, our model can be easily modified to accommodate *maximization* problems by replacing the cost rate function c with, say, a *reward rate* function r that is required to satisfy appropriately modified versions of the conditions (C1)-(C3) above (e.g., r is *upper* semicontinuous, r is bounded *above*, etc.). Of course, in this case, one can simply opt to convert such a maximization problem into a minimization problem by instead considering the function $-r$. ◇

When given such a constrained control model, we are then primarily interested in two (infinite-horizon) optimality criteria: the *(minimum) long-term average expected cost* and, given $\alpha > 0$, the *(minimum) α-discounted expected cost.* We define the former criterion as the optimal (i.e., minimum) value of the *long-term average stochastic problem*, and the latter criterion as the optimal (i.e., minimum) value of the *α-discounted stochastic problem.* We now describe each of these problems in detail.

I.3.1 The Long-Term Average Stochastic Problem

Let $\{E \times G, L, c, c_1, \theta\}$ be a constrained control model as described above, where the additional element $\theta \in \mathbb{R}$ is a given *constraint constant.* For the *long-term average stochastic problem*, our objective is then to minimize the *long-term average expected cost*

$$J(X, \Lambda) \coloneqq \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \int_G c(X_t, u)\, \Lambda_s(du)\, ds\right]$$

over all $(X, \Lambda) \in \mathscr{M}$ satisfying the *budget constraint*

$$Q_1(X, \Lambda) \coloneqq \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \int_G c_1(X_t, u)\, \Lambda_s(du)\, ds\right] \leq \theta.$$

Thus, the *minimum long-term average expected cost* is given by

$$J^* \coloneqq \inf\{J(X, \Lambda) \in \mathbb{R} : (X, \Lambda) \in \mathscr{M},\, Q_1(X, \Lambda) \leq \theta\}.$$

Remark 3.4. In principle, the constraint constant θ can of course be considered as a vector

$$\theta = (\theta_1, \theta_2, \ldots, \theta_n) \in \mathbb{R}^n,$$

where n given budget rate functions $c_1, c_2, \ldots, c_n$ are required to satisfy $Q_i(X, \Lambda) \leq \theta_i$ for each $i = 1, 2, \ldots, n$. As this is a trivial generalization that adds little more than notational complications, we treat only the $n = 1$ case. $\diamond$

For this problem, we will assume that there exists at least one $(X, \Lambda) \in \mathscr{M}$ satisfying $J(X, \Lambda) < +\infty$ and $Q_1(X, \Lambda) \leq \theta$. Indeed, this assumption is justified, since otherwise the problem itself is not well-posed.

I.3.2 The α-Discounted Stochastic Problem

Given a *discount rate* $\alpha \in (0,1]$, let $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ be a constrained control model as described above, where the additional elements $\theta_\alpha \in \mathbb{R}$ and $\nu_0 \in \mathcal{P}(G)$ are, respectively, a constraint constant (that may depend on α) and an initial distribution on the state space E. The objective of the *α-discounted stochastic problem* is then to minimize the *α-discounted expected cost*

$$J_\alpha(X, \Lambda; \nu_0) := \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t} \int_G c(X_t, u)\, \Lambda_t(du)\, dt\right]$$

over all $(X, \Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$, subject to the budget constraint

$$Q_1^\alpha(X, \Lambda; \nu_0) := \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t} \int_G c_1(X_t, u)\, \Lambda_t(du)\, dt\right] \leq \theta_\alpha.$$

Note in the above that $\mathbb{E}_{\nu_0}[\cdot] := \mathbb{E}[\cdot | X_0 \sim \nu_0]$.

Remark 3.5. It may perhaps be undesirable to consider the constraint constant θ_α to be *given* in the model $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ since one of our primary goals will be to "asymptotically approximate" the long-term average stochastic problem with the (α-indexed) family of α-discounted stochastic problems; and so, we would prefer the liberty to *choose* a suitable family $\{\theta_\alpha : \alpha \in (0,1]\}$ of associated constraint constants to accomplish this goal. In particular, we will see that such a family $\{\theta_\alpha : \alpha \in (0,1]\}$ will be required to satisfy the condition

$$\lim_{\alpha \downarrow 0} \alpha \theta_\alpha = \theta \tag{T1}$$

in order to apply the vanishing discount method. ◇

The *minimum total α-discounted expected cost* (i.e., optimum value) for this problem is then given by

$$J_\alpha^* := J_\alpha^*(\nu_0) := \inf\{J_\alpha(X, \Lambda; \nu_0) : (X, \Lambda) \in \mathscr{M},\ Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta_\alpha\}.$$

Note that this problem does, on its own, depend on the choice of an initial distribution ν_0 for the process X, but we will see that this dependence "vanishes" in the limit as $\alpha \downarrow 0$. As with the long-term average stochastic problem, we assume that there exists at least one $(X, \Lambda) \in \mathscr{M}$ with $J_\alpha(X, \Lambda; \nu_0) < +\infty$ and $Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta_\alpha$.

Remark 3.6. Our choice of parameter space for the discount rates is not *per se* limited to the half-open interval $(0, 1]$. Indeed, it would make no difference were we to choose $(0, 1)$ or $(0, 2]$, or in fact any interval of the form $(0, \lambda)$ or $(0, \lambda]$, for $\lambda \in (0, +\infty)$. However, our analysis only concerns itself with values of α near 0; so our choice of $(0, 1]$ is simply a matter of convenience. ◇

It is primarily for historical and contextual reasons that we discuss the α-discounted stochastic problem in this dissertation, as this is a problem of considerable interest in its own right. Notice, however, that there is a disparity in mass that the long-term average stochastic problem and the α-discounted stochastic problem assign to "time" in the infinite horizon; i.e., the former assigns unit mass to time in this sense, whereas the latter assigns a mass of α^{-1}. Consequently, when analyzing the linear programs with an application of the *vanishing discount method* in mind, we will instead consider the asymptotic behavior of the (α-parameterized) family of *α-normalized* (discounted) linear programs as $\alpha \downarrow 0$. The details of this approach will be made precise below.

LINEAR PROGRAMMING FORMULATIONS

In this chapter, our objective is to express the long-term average and α-discounted stochastic problems introduced above as linear programs over (infinite-dimensional) spaces of so-called *occupation measures*. As such, we first define the occupation measures corresponding to each stochastic problem, and then we derive adjoint equations (in terms of the generator L) that fully characterize these collections of measures. These adjoint equations will then provide the conditions the occupations measures must satisfy in order to be considered feasible (and, hence, optimal) for their respective linear programs. We will then be well-positioned to formulate the desired linear programs, and furthermore show that these linear programs are, in fact, equivalent to their associated stochastic programs.

II.1 Occupation Measures

II.1.1 The Long-Term Average Expected Occupation Measure

Before stating our definition of a long-term average expected occupation measure, we must first define a related measure: a *finite-horizon average expected occupation measure*. Note that $\mathcal{B}(S)$ denotes the Borel σ-algebra on the set S, and recall that $\mathcal{P}(S)$ denotes the set of probability measures on the measurable space $(S, \mathcal{B}(S))$.

Definition 1.1. Given a process $(X, \Lambda) \in \mathscr{M}$ and a time $t > 0$, we define the associated *finite-horizon average expected occupation measure* $\rho_t \in \mathcal{P}(E \times G)$ by

$$\rho_t(\Gamma) := \frac{1}{t}\mathbb{E}\left[\int_0^t \int_G I_\Gamma(X_s, u)\, \Lambda_s(du)\, ds\right], \qquad \forall \Gamma \in \mathcal{B}(E \times G).$$

Note that it will often prove useful to express the above as

$$\rho_t(\Gamma) = \int_{E\times G} I_\Gamma(x,u)\,\rho_t(dx,du) = \frac{1}{t}\mathbb{E}\left[\int_0^t \int_G I_\Gamma(X_s,u)\,\Lambda_s(du)\,ds\right].$$

◇

Our definition of a long-term average expected occupation measure further relies upon the notion of *weak convergence* of measures, which merits some discussion. In much of what follows, we take as our source Bogachev (2007), though the definition of weak convergence that appears there is considerably more general than is necessary for the scope of this dissertation. In particular, the investigative reader will notice that much of Bogachev (2007) concerns itself with the distinction between the *Baire* σ-algebra and the Borel σ-algebra on a given topological space (the latter of which being our primary σ-algebra of interest). Fortunately, however, in a *perfectly normal* space (e.g., any metric space) such as $E \times G$, these two σ-algebras happen to coincide.

Note in the following that $\mathcal{M}(S)$ denotes the set of positive (i.e., nonnegative) finite measures on the measurable space $(S, \mathcal{B}(S))$, and $\bar{C}(S)$ denotes the set of bounded, continuous, real-valued functions on S.

Definition 1.2. A sequence $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{M}(E \times G)$ is said to be *weakly convergent* to a measure $\mu \in \mathcal{M}(E \times G)$ if

$$\int_{E\times G} \xi(x,u)\,\mu(dx,du) = \lim_{n\to\infty} \int_{E\times G} \xi(x,u)\,\mu_n(dx,du), \qquad \forall \xi \in \bar{C}(E\times G).$$

We use the notation $\mu_n \Rightarrow \mu$ to denote such convergence. If we wish to specify the index over which the limit is taken, we will write $\mu_n \Rightarrow_n \mu$. ◇

It is typical in the literature to consider only weak convergence of *probability* measures, as in our following definition of long-term average expected occupation measures.

Definition 1.3. If $\{\rho_t : t > 0\} \subset \mathcal{P}(E \times G)$ is a family of finite-horizon average expected occupation measures, then $\rho \in \mathcal{P}(E \times G)$ is called a *long-term average expected occupation measure* if $\rho_{t_n} \Rightarrow \rho$ for some subsequence $\{\rho_{t_n} : n \in \mathbb{N}\}$ of $\{\rho_t : t > 0\}$ with $t_n \to \infty$ as $n \to \infty$. In other words, the measure ρ is a limit point of the set $\{\rho_t : t > 0\}$ in the topology of weak convergence. $\diamond$

We now wish to establish a suitable characterization of the collection of long-term average expected occupation measures for the control models we are considering. To do so, we first show that any long-term average expected occupation measure ρ necessarily annihilates the function Lf for each $f \in \mathscr{D}(L)$. Having shown this, we will then see that, for any ρ satisfying this condition, there exists a *stationary* $(X, \Lambda) \in \mathscr{M}$ that has ρ as its one-dimensional distribution (and, hence, its weak limit). We will make this statement more precise later on; but, for now, let us consider the following proposition in which we characterize the collection of long-term average expected occupation measures as those probability measures that satisfy the aforementioned annihilation condition.

Proposition 1.4. *Suppose that $\{E \times G, L, c, c_1, \theta\}$ is a constrained control model in which L satisfies (D1)-(D6), and c and ψ satisfy conditions (C1)-(C3). If $(X, \Lambda) \in \mathscr{M}$ with $J(X, \Lambda) < +\infty$, then there exists a long-term average expected occupation measure ρ associated with (X, Λ) satisfying*

$$\int_{E \times G} Lf(x, u)\, \rho(dx, du) = 0, \qquad \forall f \in \mathscr{D}(L). \tag{MA1}$$

Moreover, if ρ is any long-term average expected occupation measure, then it necessarily satisfies (MA1).

We must temporarily delay our proof of the above proposition with a further digression into a number of important definitions and results.

Definition 1.5. A sequence $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{M}(E \times G)$ of measures is called *tight* if, for each $\epsilon > 0$, there is a compact subset K_ϵ of $E \times G$ such that, for every $n \in \mathbb{N}$, we have $\mu_n(K_\epsilon^{\mathsf{c}}) < \epsilon$. $\diamond$

Definition 1.6. A sequence $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$ of probability measures is called *relatively compact* if every subsequence $\{\mu_{n_k} : k \in \mathbb{N}\}$ of $\{\mu_n : n \in \mathbb{N}\}$ contains a further subsequence $\{\mu_{n_{k(i)}} : i \in \mathbb{N}\}$ such that $\mu_{k(i)} \Rightarrow_i \mu$ for some $\mu \in \mathcal{P}(E \times G)$. ◇

The main result relating tightness and relative compactness is *Prohorov's Theorem*. For our purposes, the following theorem (which one may recognize as a special case of Theorem 8.6.2 in Bogachev (2007)) is what we will often refer to as "Prohorov's Theorem."

Theorem 1.7. *For a sequence $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{M}(E \times G)$, the following conditions are equivalent:*

(a) $\{\mu_n : n \in \mathbb{N}\}$ contains a weakly convergent subsequence.

(b) $\{\mu_n : n \in \mathbb{N}\}$ is tight and $\{\mu_n(E \times G) : n \in N\} \subset \mathbb{R}^+$ is bounded.

Remark 1.8. Here we note that sequences in $\mathcal{M}(E \times G)$ satisfying condition (a) in the above theorem will be referred to as "relatively compact," even though this terminology usually only applies to sequences of probability measures throughout the literature. Also, when a sequence satisfies condition (a), we will often take the "weakly convergent subsequence" to just be the sequence itself if it makes no meaningful difference in our argument. ◇

For completeness, we also include the analogous statements for probability measures below. The following two results can be found in Section 5 of Billingsley (1999).

Theorem 1.9. *If the sequence $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$ is tight, then it is relatively compact.*

Corollary 1.10. *If the sequence $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$ is tight, and if each subsequence that converges weakly at all in fact converges to a measure $\mu \in \mathcal{P}(E \times G)$, then the entire sequence must converge weakly to μ.*

One final definition (and an accompanying example) will then provide us with the necessary machinery to tackle the proof of Proposition 1.4.

Definition 1.11. A set of functions $M \subset \bar{C}(E \times G)$ is called *separating* if whenever $\mu, \nu \in \mathcal{P}(E \times G)$ and

$$\int_{E\times G} \xi(x,u)\, \mu(dx,du) = \int_{E\times G} \xi(x,u)\, \nu(dx,du), \qquad \forall \xi \in M,$$

we have $\mu = \nu$. ◇

Example 1.12. $C_c(E \times G)$, the space of functions $\xi \in \bar{C}(E \times G)$ with compact support, is separating. ◇

The interested reader is directed to Section 4 of Chapter 3 in Ethier and Kurtz (1986) for a proof of the preceding fact. Now, as promised, we proceed with the proof of Proposition 1.4.

Proof of Proposition 1.4. Let $(X, \Lambda) \in \mathscr{M}$ with $J(X,\Lambda) = m < +\infty$; i.e.,

$$m := \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \int_G c(X_s,u)\,\Lambda_s(du)\,ds\right] < +\infty.$$

Let $\{t_n : n \in \mathbb{N}\} \subset [1,\infty)$ be a sequence of times with $t_n \to \infty$ as $n \to \infty$ and

$$\left|\frac{1}{t_n}\mathbb{E}\left[\int_0^{t_n} \int_G c(X_s,u)\,\Lambda_s(du)\,ds\right] - m\right| \le \frac{1}{n}, \qquad \forall n \in \mathbb{N}.$$

Then define the sequence $\{\rho_{t_n} : n \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$ of finite-horizon average expected occupation measures by

$$\rho_{t_n}(\Gamma) := \frac{1}{t_n}\mathbb{E}\left[\int_0^{t_n} \int_G I_\Gamma(X_s,u)\,\Lambda_s(du)\,ds\right], \qquad \forall \Gamma \in \mathcal{B}(E\times G).$$

We wish to show that the sequence $\{\rho_{t_n} : n \in \mathbb{N}\}$ is tight (and, hence, relatively compact). To this end, let $\epsilon \in (0,1)$ be given and let $M := M(m,\epsilon,\kappa) > 0$ be chosen so that $M > (1+m+\kappa)/\epsilon$, recalling that $-\kappa$ is a lower bound on c. Then define the sets

$$H := \{(x,u) \in E\times G : -\kappa \le c(x,u) < 0\}, \quad H_M := \{(x,u) \in E\times G : 0 \le c(x,u) \le M\},$$

$$\text{and} \quad K_M := \{(x,u) \in E \times G : c(x,u) \le M\},$$

noting that $K_M = H \cup H_M$ and that K_M is compact since c is assumed to be inf-compact. The sets H, H_M and K_M^{c} evidently form a partition of $E \times G$; so we have, for each $n \in \mathbb{N}$,

$$\begin{aligned}
M\rho_{t_n}(K_M^{\mathsf{c}}) &\le \int_{K_M^{\mathsf{c}}} c(x,u)\,\rho_{t_n}(dx,du) \\
&\le \int_{K_M^{\mathsf{c}}} c(x,u)\,\rho_{t_n}(dx,du) + \int_H (c(x,u)+\kappa)\,\rho_{t_n}(dx,du) + \int_{H_M} c(x,u)\,\rho_{t_n}(dx,du) \\
&\le \int_{E\times G} c(x,u)\,\rho_{t_n}(dx,du) + \kappa \\
&\le 1 + m + \kappa,
\end{aligned}$$

which implies that

$$\rho_{t_n}(K_M^{\mathsf{c}}) \le \frac{1+m+\kappa}{M} < \epsilon.$$

Since K_M is compact and $\rho_{t_n}(K_M^{\mathsf{c}}) < \epsilon$ for each $n \in \mathbb{N}$, it follows that $\{\rho_{t_n} : n \in \mathbb{N}\}$ is, indeed, tight (and, hence, relatively compact). Thus, there is a $\rho \in \mathcal{P}(E \times G)$ such that $\rho_{t_n} \Rightarrow \rho$ (see Remark 1.8). Now, let $f \in \mathscr{D}(L)$ be arbitrary. Since $(X, \Lambda) \in \mathscr{M}$ it then follows that

$$\frac{1}{t_n}\mathbb{E}[f(X_{t_n}) - f(X_0)] = \frac{1}{t_n}\mathbb{E}\left[\int_0^{t_n}\int_G Lf(X_s,u)\,\Lambda_s(du)\,ds\right] = \int_{E\times G} Lf(x,u)\,\rho_{t_n}(dx,du);$$

and since $\mathscr{D}(L) \subset \widehat{C}(E) \subset \bar{C}(E)$, we have

$$\lim_{n\to\infty}\frac{1}{t_n}\mathbb{E}[f(X_{t_n}) - f(X_0)] = 0.$$

However, we have only $Lf \in C(E \times G)$, so it is not immediately apparent that

$$\lim_{n\to\infty}\int_{E\times G} Lf(x,u)\,\rho_{t_n}(dx,du) = \int_{E\times G} Lf(x,u)\,\rho(dx,du).$$

So, recalling condition (D6), define the sequence $\{\zeta_n : n \in \mathbb{N}\} \subset \mathcal{M}(E \times G)$ by

$$\zeta_n(\Gamma) := \int_{\Gamma} \psi(x,u)\, \rho_{t_n}(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E \times G),$$

observing that ψ can then be viewed as the Radon-Nikodym derivative $\psi = d\zeta_n/d\rho_{t_n}$, for each $n \in \mathbb{N}$. We claim that the sequence $\{\zeta_n(E \times G) : n \in \mathbb{N}\} \subset \mathbb{R}^+$ is bounded.
To show this, we begin by partitioning $E \times G$ into the sets

$$F_1 := \{(x,u) \in E \times G : -\kappa \leq c(x,u) < -1\}, \qquad F_2 := \{(x,u) \in E \times G : -1 \leq c(x,u) < 1\},$$

$$\text{and} \qquad F_3 := \{(x,u) \in E \times G : c(x,u) \geq 1\}.$$

We then see that $c^\beta(x,u) < -1 < 0 \leq c(x,u) + \kappa$ for each $(x,u) \in F_1$, $c^\beta(x,u) \leq c(x,u) + 1$ for each $(x,u) \in F_2$, and $c^\beta(x,u) \leq c(x,u)$ for each $(x,u) \in F_3$. Thus, for each $n \in \mathbb{N}$, we have

$$\begin{aligned}
\zeta_n(E \times G) &= \int_{E \times G} \psi(x,u)\, \rho_{t_n}(dx,du) \\
&\leq \int_{E \times G} \big(ac^\beta(x,u) + b\big)\, \rho_{t_n}(dx,du) \\
&= \int_{F_1} ac^\beta(x,u)\, \rho_{t_n}(dx,du) + \int_{F_2} ac^\beta(x,u)\, \rho_{t_n}(dx,du) \\
&\qquad\qquad + \int_{F_3} ac^\beta(x,u)\, \rho_{t_n}(dx,du) + b \\
&\leq \int_{F_1} a(c(x,u) + \kappa)\, \rho_{t_n}(dx,du) + \int_{F_2} a(c(x,u) + 1)\, \rho_{t_n}(dx,du) \\
&\qquad\qquad + \int_{F_3} ac(x,u)\, \rho_{t_n}(dx,du) + b \\
&= a \int_{E \times G} c(x,u)\, \rho_{t_n}(dx,du) + a\kappa\rho_{t_n}(F_1) + a\rho_{t_n}(F_2) + b \\
&\leq a(m + (1/n)) + a\kappa + a + b \\
&\leq a(m + \kappa + 2) + b < +\infty.
\end{aligned}$$

Hence, the sequence $\{\zeta_n(E \times G) : n \in \mathbb{N}\}$ is bounded. We now claim that the sequence $\{\zeta_n : n \in \mathbb{N}\}$ is tight. The rescaling

$$\zeta_n' \coloneqq \frac{\zeta_n}{a(m+\kappa+2)+b}, \qquad \forall n \in \mathbb{N},$$

then defines a sequence $\{\zeta_n' : n \in \mathbb{N}\}$ of subprobability measures; and so, as in the argument above, we again have

$$\begin{aligned} M\zeta_n'(K_M^{\mathsf{c}}) &\leq \int_{K_M^{\mathsf{c}}} c(x,u)\, \zeta_n'(dx,du) \\ &\leq \int_{E\times G} c(x,u)\, \zeta_n'(dx,du) + \kappa \\ &\leq \int_{E\times G} c(x,u)\, \rho_{t_n}(dx,du) + \kappa \\ &\leq 1 + m + \kappa, \end{aligned}$$

which implies that

$$\frac{\zeta_n(K_M^{\mathsf{c}})}{a(m+\kappa+2)+b} = \zeta_n'(K_M^{\mathsf{c}}) \leq \frac{1+m+\kappa}{M} < \epsilon.$$

Since $a(m+\kappa+2)+b$ is independent of ϵ, it follows that the sequence $\{\zeta_n : n \in \mathbb{N}\}$ is tight. So, having shown that $\{\zeta_n(E \times G) : n \in \mathbb{N}\}$ is bounded and $\{\zeta_n : n \in \mathbb{N}\}$ is tight, we may conclude that $\{\zeta_n : n \in \mathbb{N}\}$ is relatively compact. Again, this yields a measure $\zeta \in \mathcal{M}(E\times G)$ satisfying $\zeta_n \Rightarrow \zeta$. Moreover, by normalizing ζ and each ζ_n via the measures $\hat{\zeta} \coloneqq \zeta/\zeta(E\times G)$ and $\hat{\zeta}_n \coloneqq \zeta_n/\zeta_n(E \times G)$, we have $\hat{\zeta} \in \mathcal{P}(E \times G)$ and $\{\hat{\zeta}_n : n \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$. To then show that $\hat{\zeta}_n \Rightarrow \hat{\zeta}$, we need only to verify that $\zeta_n(E \times G) \to \zeta(E \times G)$ as $n \to \infty$. This, however, is an immediate consequence of $\zeta_n \Rightarrow \zeta$ and integration of the constant function $1 \in \bar{C}(E \times G)$; i.e.,

$$\lim_{n\to\infty} \zeta_n(E \times G) = \lim_{n\to\infty} \int_{E\times G} 1\, d\zeta_n = \int_{E\times G} 1\, d\zeta = \zeta(E \times G).$$

We now wish to verify that

$$\zeta(\Gamma) = \int_\Gamma \psi(x,u)\,\rho(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E\times G).$$

So, let $\xi \in C_c(E\times G)$ be arbitrary. Then $\psi\xi \in C_c(E\times G)$, as well. Thus,

$$\begin{aligned}\int_{E\times G} \xi(x,u)\,\hat{\zeta}(dx,du) &= \lim_{n\to\infty}\int_{E\times G}\xi(x,u)\,\hat{\zeta}_n(dx,du)\\ &= \lim_{n\to\infty}\frac{1}{\zeta_n(E\times G)}\int_{E\times G}\xi(x,u)\,\zeta_n(dx,du)\\ &= \lim_{n\to\infty}\frac{1}{\zeta_n(E\times G)}\int_{E\times G}(\psi\xi)(x,u)\,\rho_n(dx,du)\\ &= \frac{1}{\zeta(E\times G)}\int_{E\times G}(\psi\xi)(x,u)\,\rho(dx,du).\end{aligned}$$

If we now define the measure

$$\hat{\rho}(\Gamma) := \frac{1}{\zeta(E\times G)}\int_\Gamma \psi(x,u)\,\rho(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E\times G),$$

then we see that

$$\int_{E\times G}\xi(x,u)\,\hat{\rho}(dx,du) = \frac{1}{\zeta(E\times G)}\int_{E\times G}(\psi\xi)(x,u)\,\rho(dx,du) = \int_{E\times G}\xi(x,u)\,\hat{\zeta}(dx,du).$$

Thus, since $\xi \in C_c(E\times G)$ is arbitrary and $C_c(E\times G)$ is separating, it follows that $\hat{\zeta} = \hat{\rho}$; and so,

$$\zeta(\Gamma) = \hat{\zeta}(\Gamma)\zeta(E\times G) = \hat{\rho}(\Gamma)\zeta(E\times G) = \int_\Gamma \psi(x,u)\,\rho(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E\times G),$$

as desired. So, having established that

$$\zeta(\Gamma) = \int_\Gamma \psi(x,u)\,\rho(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E\times G),$$

we may now view ψ as the Radon-Nikodym derivative $\psi = d\zeta/d\rho$. Then, since $|L/\psi| \le a_f$ by condition (D6), we have

$$\begin{aligned} 0 &= \lim_{n\to\infty} \int_{E\times G} Lf(x,u)\,\rho_{t_n}(dx,du) \\ &= \lim_{n\to\infty} \int_{E\times G} (Lf(x,u)/\psi(x,u))\,\zeta_n(dx,du) \\ &= \int_{E\times G} (Lf(x,u)/\psi(x,u))\,\zeta(dx,du) = \int_{E\times G} Lf(x,u)\,\rho(dx,du), \end{aligned}$$

which establishes condition (MA1) and, hence, proves the proposition. □

To simplify notation moving forward, we introduce the bilinear mapping $\langle\cdot,\cdot\rangle : M(E\times G)\times \mathcal{M}_s(E\times G) \to [-\infty,\infty]$ defined by

$$\langle \xi,\mu\rangle \coloneqq \int_{E\times G} \xi(x,u)\,\mu(dx,du),$$

where $\mathcal{M}_s(E\times G)$ denotes the set of finite signed measures on $E\times G$ (the choice of which is simply to ensure that the domain of $\langle\cdot,\cdot\rangle$ is a product of vector spaces—and hence $\langle\cdot,\cdot\rangle$ is, in fact, "bilinear"). Thus, condition (MA1) can be written more compactly as

$$\langle Lf,\rho\rangle = 0, \qquad \forall f\in\mathscr{D}(L). \tag{MA1}$$

Note that those probability measures on $E\times G$ that satisfy (MA1) will, on occasion, be referred to as *stationary*. In order to be more precise about this, we need the following definition.

Definition 1.13. Let $\mu\in\mathcal{M}(E\times G)$, and let μ^E denote the state marginal measure of μ. Then a mapping $\eta : E\times\mathcal{B}(G)\to[0,1]$ is called the *regular conditional probability distribution on G, given $x\in E$, relative to μ^E* (or, more succinctly, the *regular conditional distribution of μ*) if it satisfies the following:

(a) For each $x\in E$, the mapping $\eta(x,\cdot) : \mathcal{B}(G)\to[0,1]$ belongs to $\mathcal{P}(G)$.

(b) For each $\Gamma_2 \in \mathcal{B}(G)$, the mapping $\eta(\cdot, \Gamma_2) : E \to [0,1]$ belongs to $M(E)$.

(c) For each $\Gamma_1 \times \Gamma_2 \in \mathcal{B}(E \times G)$, the mapping η satisfies

$$\mu(\Gamma_1 \times \Gamma_2) = \int_{\Gamma_1} \eta(x, \Gamma_2)\, \mu^E(dx).$$

$\diamond$

Now, suppose $\rho \in \mathcal{P}(E \times G)$ satisfies (MA1), and let η be the regular conditional distribution of ρ. If X is a stationary *process* with $X_0 \sim \rho^E$, then the process $(X, \eta(X, \cdot)) := \{(X_t, \eta(X_t, \cdot)) : t \in \mathbb{R}^+\}$ is stationary and its one-dimensional distributions satisfy

$$\mathbb{E}[I_{\Gamma_1}(X_t)\eta(X_t, \Gamma_2)] = \rho(\Gamma_1 \times \Gamma_2), \qquad \forall t \in \mathbb{R}^+.$$

Note that we will have further occasion to consider relaxed controls of the form $\Lambda = \eta(X, \cdot) := \{\eta(X_t, \cdot) : t \in \mathbb{R}^+\}$, where $\Lambda_t(\cdot) = \eta(X_t, \cdot)$ for each $t \in \mathbb{R}^+$. In this case, we say that such controls are given in *feedback form*.

The following result is a consequence of Theorem 2.2 in Kurtz and Stockbridge (1998), and its statement completes our characterization of the collection of long-term average expected occupation measures for our model.

Theorem 1.14. *Suppose that $\{E \times G, L, c, c_1, \theta\}$ is a constrained control model in which L satisfies (D1)-(D6) and c and ψ satisfy (C1)-(C3). If $\rho \in \mathcal{P}(E \times G)$ satisfies (MA1) and η is the regular conditional distribution of ρ given $x \in E$, then there exists a stationary process X such that $(X, \eta(X, \cdot))$ is a stationary relaxed solution of the controlled martingale problem for L.*

II.1.2 The α-Discounted Expected Occupation Measure

Given $(X, \Lambda) \in \mathscr{M}$ and a discount rate $\alpha \in (0, 1]$, we define the associated *α-discounted expected occupation measure* $\mu_\alpha \in \mathcal{M}(E \times G)$ by

$$\mu_\alpha(\Gamma) := \mathbb{E}\left[\int_0^\infty \int_G e^{-\alpha t} I_\Gamma(X_t, u)\, \Lambda_t(du)\, dt\right], \qquad \forall \Gamma \in \mathcal{B}(E \times G),$$

where it will often prove convenient to express $\mu_\alpha(\Gamma)$ in the form

$$\mu_\alpha(\Gamma) = \int_{E \times G} I_\Gamma(x, u)\, \mu_\alpha(dx, du).$$

We now wish to establish an analogous condition to (MA1) in order to characterize, for a fixed $\alpha \in (0, 1]$, the collection of α-discounted expected occupation measures. To this end, we first define the operator $L_\alpha : \mathscr{D}(L) \to \mathscr{R}(L)$ as follows: For each $f \in \mathscr{D}(L)$,

$$L_\alpha f(x, u) := Lf(x, u) - \alpha f(x), \qquad \forall (x, u) \in E \times G.$$

Now, a lemma:

Lemma 1.15. *Suppose that $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ is a constrained control model in which L satisfies conditions (D1)-(D6) and c_1 and ψ satisfy conditions (C1)-(C3). If $(X, \Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$, and μ_α is the α-discounted expected occupation measure associated with (X, Λ), then*

$$\int_{E \times G} L_\alpha f(x, u)\, \mu_\alpha(dx, du) = \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t}\left(\int_G L_\alpha f(X_t, u)\, \Lambda_t(du)\right) dt\right], \qquad \forall f \in \mathscr{D}(L).$$

Proof. Here, we avail ourselves of the so-called "Standard Machine," as described in Williams (1991). We begin by assuming that $\xi = I_\Gamma$ for some $\Gamma \in \mathcal{B}(E \times G)$. By the definition of μ_α, we have

$$\int_{E\times G} \xi(x,u)\,\mu_\alpha(dx,du) = \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t}\left(\int_G \xi(X_t,u)\,\Lambda_t(du)\right)dt\right]. \tag{1.16}$$

Now, suppose that ξ is a nonnegative simple function on $E\times G$; i.e.,

$$\xi = \sum_{k=1}^{n} y_k I_{\Gamma_k},$$

for some $n\in\mathbb{N}$, $\{y_1,y_2,\ldots,y_n\}\subset[0,\infty]$, and $\{\Gamma_1,\Gamma_2,\ldots,\Gamma_n\}\subset\mathcal{B}(E\times G)$. Then, via the linearity of integration, we have

$$\begin{aligned}
\int_{E\times G} \xi(x,u)\,\mu_\alpha(dx,du) &= \int_{E\times G}\sum_{k=1} y_k I_{\Gamma_k}(x,u)\,\mu_\alpha(dx,du)\\
&= \sum_{k=1}^{n} y_k \int_{E\times G} I_{\Gamma_k}(x,u)\,\mu_\alpha(dx,du)\\
&= \sum_{k=1}^{n} y_k \mathbb{E}\left[\int_0^\infty e^{-\alpha t}\left(\int_G I_{\Gamma_k}(X_t,u)\,\Lambda_t(du)\right)dt\right]\\
&= \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t}\left(\int_G \sum_{k=1}^{n} y_k I_{\Gamma_k}(X_t,u)\,\Lambda_t(du)\right)dt\right]\\
&= \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t}\left(\int_G \xi(X_t,u)\,\Lambda_t(du)\right)dt\right].
\end{aligned}$$

Next, we take ξ to be a nonnegative $\mathcal{B}(E\times G)$-measurable function on $E\times G$ and, for each $r\in\mathbb{N}$, we define the rth "staircase function" $g^{(r)}:[0,\infty]\to[0,\infty]$ by

$$g^{(r)}(y) := \begin{cases} 0, & \text{if } y=0;\\ (i-1)2^{-r}, & \text{if } (i-1)2^{-r} < y \le i2^{-r} \le r,\ i\in\mathbb{N};\\ r, & \text{if } y>r. \end{cases}$$

Then $\xi^{(r)} := g^{(r)}\circ\xi$ defines a sequence $\{\xi^{(r)} : r\in\mathbb{N}\}$ of nondecreasing nonnegative simple functions on $E\times G$ that converges pointwise to ξ; and so, by the Monotone Convergence Theorem and the previous steps, we see that (1.12) holds in this case, as well.

Finally, suppose only that ξ is μ_α-integrable. Then, by simply writing $\xi = \xi^+ - \xi^-$, we see that ξ is a (finite) linear combination of nonnegative $\mathcal{B}(E \times G)$-measurable functions, each of which is μ_α-integrable. It therefore follows that (1.12) holds when ξ is only known to be μ_α-integrable.

Now, we observe that $\psi \in C(E \times G)$ and $\psi \geq 1$, so ψ is a nonnegative $\mathcal{B}(E \times G)$-measurable function. Thus, (1.12) holds when $\xi = \psi$; i.e.,

$$\int_{E\times G} \psi(x,u)\,\mu_\alpha(dx,du) = \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t}\left(\int_G \psi(X_t,u)\,\Lambda_t(du)\right)dt\right].$$

Moreover, since $\langle c_1, \mu_\alpha\rangle \leq \theta_\alpha$ for some $\theta_\alpha \in \mathbb{R}$, condition (C3) further ensures that ψ is, in fact, μ_α-integrable. Now, let $f \in \mathscr{D}(L)$ and let $\|f\|_\infty$ denote the supremum norm of f. By condition (D6), we then have

$$|L_\alpha f(x,u)| = |Lf(x,u) - \alpha f(x)| \leq |Lf(x,u)| + \alpha|f(x)| \leq a_f\psi(x,u) + \alpha\|f\|_\infty,$$

for every $(x,u) \in E \times G$, where $\|f\|_\infty < +\infty$ since $\mathscr{D}(L) \subset \widehat{C}(E)$. So, the μ_α-integrability of ψ ensures that $L_\alpha f$ is also μ_α-integrable for each $f \in \mathscr{D}(L)$. Hence,

$$\int_{E\times G} L_\alpha f(x,u)\,\mu_\alpha(dx,du) = \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t}\left(\int_G L_\alpha f(X_t,u)\,\Lambda_t(du)\right)dt\right], \qquad \forall f \in \mathscr{D}(L),$$

as desired. □

We are now in a position to state the adjoint condition that will both characterize the collection of α-discounted expected occupation measures and serve as the defining constraint of the forthcoming α-discounted linear program.

Proposition 1.17. *Suppose that $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ is a constrained control model in which L satisfies conditions (D1)-(D6) and c_1 and ψ satisfy conditions (C1)-(C3). If $(X,\Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$, and μ_α is the α-discounted expected occupation measure associ-*

ated with (X, Λ)*, then*

$$\langle L_\alpha f, \mu_\alpha \rangle = -\int_E f \, d\nu_0, \qquad \forall f \in \mathscr{D}(L). \tag{MD1}$$

Proof. Let $f \in \mathscr{D}(L)$ be given. Then a slight modification of Lemma 3.2 in Chapter 4 of Ethier and Kurtz (1986) shows that the process $\Psi := \{\Psi_t : t \in \mathbb{R}^+\}$ defined by

$$\Psi_t := e^{-\alpha t} f(X_t) - f(X_0) - \int_0^t e^{-\alpha s} \left(\int_G L_\alpha f(X_s, u) \, \Lambda_s(du) \right) ds$$

is a mean-zero martingale. So, since $f \in \mathscr{D}(L) \subset \widehat{C}(E)$, it follows that

$$0 = e^{-\alpha t} \mathbb{E}_{\nu_0}[f(X_t)] - \mathbb{E}_{\nu_0}[f(X_0)] - \mathbb{E}_{\nu_0} \left[\int_0^t e^{-\alpha s} \left(\int_G L_\alpha f(X_s, u) \, \Lambda_s(du) \right) ds \right], \qquad \forall t \in \mathbb{R}^+.$$

Since $|L_\alpha f|$ is dominated by $a_f \psi + \alpha \|f\|_\infty$ and ψ is μ_α-integrable (as in the proof of Lemma 1.11), it follows that

$$\lim_{t \to \infty} \mathbb{E}_{\nu_0} \left[\int_0^t e^{-\alpha s} \left(\int_G L_\alpha f(X_s, u) \, \Lambda_s(du) \right) ds \right] = \mathbb{E}_{\nu_0} \left[\int_0^\infty e^{-\alpha s} \left(\int_G L_\alpha f(X_s, u) \, \Lambda_s(du) \right) ds \right];$$

and so,

$$-\mathbb{E}_{\nu_0}[f(X_0)] = \mathbb{E}_{\nu_0} \left[\int_0^\infty e^{-\alpha s} \left(\int_G L_\alpha f(X_s, u) \, \Lambda_s(du) \right) ds \right].$$

Finally, noting that $\mathbb{E}_{\nu_0}[f(X_0)] = \int_E f \, d\nu_0$, we may conclude from Lemma 1.11 that

$$\langle L_\alpha f, \mu_\alpha \rangle = -\int_E f \, d\nu_0, \qquad \forall f \in \mathscr{D}(L),$$

which completes the proof. □

In analogy with our characterization of long-term average expected occupation measures, we have a similar result stating that each α-discounted expected occupation measure that satisfies (MD1) is the associated α-discounted expected occupation measure for some $(X, \Lambda) \in \mathscr{M}$. This result is a consequence of Corollary 5.3 in Kurtz and Stockbridge (1998).

Theorem 1.18. *Suppose that $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ is a constrained control model in which L satisfies conditions (D1)-(D6) and c_1 and ψ satisfy conditions (C1)-(C3). If $(X, \Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$, μ_α is the α-discounted expected occupation measure associated with (X, Λ), and η_α is the regular conditional distribution of μ_α, then there exists a process X such that $(X, \eta_\alpha(X, \cdot)) \in \mathscr{M}$.*

II.2 Equivalent Linear Programming Formulations

In the previous section, we defined and characterized the collection of long-term average expected occupation measures and, for each fixed $\alpha \in (0, 1]$, the collection of α-discounted expected occupation measures. We showed that each long-term average expected occupation measure ρ is characterized by the adjoint equation

$$\langle Lf, \rho \rangle = 0, \qquad \forall f \in \mathscr{D}(L); \tag{MA1}$$

and that, for a fixed $\alpha \in (0, 1]$, each α-discounted expected occupation measure μ_α is characterized by

$$\langle L_\alpha f, \mu_\alpha \rangle = -\alpha \int_E f \, d\nu_0, \qquad \forall f \in \mathscr{D}(L). \tag{MD1}$$

We now wish to use these characterizations to express the long-term average stochastic problem and the α-discounted stochastic problem as (potentially infinite-dimensional) linear programs over spaces of their associated occupation measures. To do so, we first define each linear program, and then show that solving each linear program is, in fact, *equivalent* to solving the corresponding stochastic problem.

II.2.1 The Long-Term Average Linear Program

Given a constrained control model $\{E \times G, L, c, c_1, \theta\}$ in which L satisfies (D1)-(D6) and c and c_1 satisfy (C1)-(C3), we define the *long-term average linear program* to be as follows:

$$\text{Minimize} \quad \langle c, \rho \rangle \qquad \text{subject to} \qquad \begin{cases} \langle Lf, \rho \rangle = 0, \quad \forall f \in \mathscr{D}(L); & \text{(MA1)} \\ \langle 1, \rho \rangle = 1; & \text{(MA2)} \\ \rho \in \mathcal{P}(E \times G); & \text{(MA3)} \\ \langle c_1, \rho \rangle \leq \theta. & \text{(CA1)} \end{cases}$$

For this linear program, let $\mathfrak{M}$ denote the set of *feasible* long-term average expected occupation measures and let $\mathfrak{M}^*$ denote the set of *optimal* long-term average expected occupations measures.

Remark 2.1. Here we acknowledge that (MA2) is technically redundant since any ρ that satisfies (MA3) necessarily satisfies (MA2); but the inclusion of (MA2) simply allows for a more consistent appearance among the linear programs under consideration. ◇

Our task now is to show that the minimum cost (i.e., optimum value) for this linear program is equal to that of the long-term average stochastic problem. That is,

$$J^* = \inf\{\langle c, \rho \rangle : \rho \in \mathfrak{M}\} = \langle c, \rho^* \rangle, \tag{2.2}$$

where ρ^* denotes an element of $\mathfrak{M}^*$. Indeed, we say that the long-term average linear program and the long-term average stochastic problem are *equivalent* if:

(a) (2.2) holds; and

(b) $\rho^* \in \mathfrak{M}^*$ if and only if the stationary process $(X^*, \Lambda^*) \in \mathscr{M}$ associated with ρ^* satisfies $J(X^*, \Lambda^*) = J^*$.

Note that, if there exists $\rho^* \in \mathfrak{M}^*$, then Theorem 1.14 guarantees the existence of a $(X, \Lambda^*) = (X, \eta^*(X, \cdot)) \in \mathscr{M}$, where η^* is the regular conditional distribution of ρ^*; but it then remains to show that $J(X, \eta^*(X, \cdot)) = J^*$. Similarly, if there exists a $(X^*, \Lambda^*) \in \mathscr{M}$ with $J(X^*, \Lambda^*) = J^*$, then there is a $\rho \in \mathcal{P}(E \times G)$ associated with (X^*, Λ^*) that satisfies the adjoint condition (MA1); but, in this case, it then remains to show that $\rho \in \mathfrak{M}^*$. Hence, it would be desirable to show that there necessarily exists such an optimal solution to either the long-term average stochastic problem or the long-term average linear program. Fortunately, under the conditions we have imposed on our model—in particular, that there is at least one $(X, \Lambda) \in \mathscr{M}$ with $J(X, \Lambda) < +\infty$ and $Q_1(X, \Lambda) \leq \theta$—this happens to be the case, as we will show shortly.

To verify this equivalence between the long-term average linear program and the long-term average stochastic problem, we rely extensively upon Theorem 6.1 in Kurtz and Stockbridge (1998), which provides such an equivalence when the budget constraint (CA1) is excluded. We state this (appropriately modified) result as the following lemma.

Lemma 2.3. *Suppose $\{E \times G, L, c\}$ is a control model in which L satisfies (D1)-(D6) and c satisfies (C1)-(C3). Then the stochastic problem of minimizing*

$$J(X, \Lambda) = \limsup_{t \to \infty} \frac{1}{t} \mathbb{E}\left[\int_0^t \int_G c(X_t, u)\, \Lambda_s(du)\, ds\right]$$

over all $(X, \Lambda) \in \mathscr{M}$ is equivalent to the linear program of minimizing

$$\langle c, \rho \rangle = \int_{E \times G} c(x, u)\, \rho(dx, du)$$

over all $\rho \in \mathcal{P}(E \times G)$ satisfying

$$\langle Lf, \rho \rangle = 0, \qquad \forall f \in \mathscr{D}(L). \tag{MA1}$$

Moreover, there exists a $\rho^* \in \mathcal{P}(E \times G)$ satisfying (MA1) and

$$\langle c, \rho^* \rangle \leq \langle c, \rho \rangle < +\infty$$

for every $\rho \in \mathcal{P}(E \times G)$ satisfying (MA1).

We can now compare the above lemma with our desired equivalence theorem below in which the budget constraint (CA1) is included. Also note this theorem shows that the long-term average linear program, in fact, has an optimal solution.

Theorem 2.4. *Suppose $\{E \times G, L, c, c_1, \theta\}$ is a constrained control model in which L satisfies (D1)-(D6) and c and c_1 satisfy (C1)-(C3). Then the long-term average stochastic problem of minimizing*

$$J(X, \Lambda) = \limsup_{t \to \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t \int_G c(X_t, u) \, \Lambda_s(du) \, ds \right]$$

over all $(X, \Lambda) \in \mathscr{M}$ with $Q_1(X, \Lambda) \leq \theta$ is equivalent to the long-term average linear program; i.e.,

$$J^* = \inf\{\langle c, \rho \rangle : \rho \in \mathfrak{M}\}.$$

Moreover, there exists an optimal $\rho^* \in \mathfrak{M}^*$ satisfying $\langle c, \rho^* \rangle = J^*$, and a corresponding stationary $(X^*, \Lambda^*) = (X^*, \eta^*(X^*, \cdot)) \in \mathscr{M}$ satisfying $J(X^*, \Lambda^*) = J^*$ and $Q_1(X^*, \Lambda^*) \leq \theta$, where η^* is the regular conditional distribution of ρ^*.

Proof. Let $(X, \Lambda) \in \mathscr{M}$ and, for each $t \in \mathbb{R}^+$, define the probability measure

$$\rho_t(\Gamma) = \frac{1}{t} \mathbb{E} \left[\int_0^t \int_G I_\Gamma(X_s, u) \, \Lambda_s(du) \, ds \right], \qquad \forall \Gamma \in \mathcal{B}(E \times G).$$

If $J(X, \Lambda) < +\infty$ and $Q_1(X, \Lambda) \leq \theta$, then—as in the proof of Proposition 1.4—the conditions on c ensure that the collection $\{\rho_t : t \in \mathbb{R}^+\}$ is relatively compact. Hence, if ρ is any limit point of $\{\rho_t : t \in \mathbb{R}^+\}$, the lower semicontinuity of c and c_1 (along with Fatou's Lemma) ensure that

$$\int_{E \times G} c(x, u) \, \rho(dx, du) \leq J(X, \Lambda) \qquad \text{and} \qquad \int_{E \times G} c_1(x, u) \, \rho(dx, du) \leq \theta.$$

Since the proof of Proposition 1.4 also shows that ρ satisfies (MA1), it then follows that

$$\inf\{\langle c,\rho\rangle : \rho\in\mathfrak{M}\} \le \inf\{J(X,\Lambda) : (X,\Lambda)\in\mathscr{M},\, Q_1(X,\Lambda)\le\theta\} = J^*.$$

Now, let $\rho\in\mathfrak{M}$. If $\langle c,\rho\rangle < +\infty$, then the conditions on c and c_1 ensure that there exists an optimal $\rho^*\in\mathfrak{M}^*$; and, by Theorem 1.14, there exists a stationary $(X^*,\Lambda^*)\in\mathscr{M}$ with marginals given by ρ^*. The conditions on c_1 further ensure that $Q_1(X^*,\Lambda^*)\le\theta$; and so,

$$J^* = J(X^*,\Lambda^*) \le \langle c,\rho^*\rangle = \inf\{\langle c,\rho\rangle : \rho\in\mathfrak{M}\},$$

which completes the proof. □

II.2.2 The α-Discounted Linear Program

Given a discount rate $\alpha\in(0,1]$ and a constrained control model $\{E\times G, L, c, c_1, \theta_\alpha, \nu_0\}$ in which L satisfies (D1)-(D6) and c_1 and c_1 satisfy (C1)-(C3), we define the *α-discounted linear program* to be as follows:

$$\text{Minimize} \quad \langle c,\mu_\alpha\rangle \qquad \text{subject to} \quad \begin{cases} \langle L_\alpha f,\mu_\alpha\rangle = -\int_E f\,d\nu_0, \quad \forall f\in\mathscr{D}(L); & \text{(MD1)}\\ \langle 1,\mu_\alpha\rangle = \alpha^{-1}; & \text{(MD2)}\\ \mu_\alpha\in\mathcal{M}(E\times G); & \text{(MD3)}\\ \langle c_1,\mu_\alpha\rangle \le \theta_\alpha. & \text{(CD1)}\end{cases}$$

For this linear program, let $\mathfrak{M}_\alpha$ denote the set of feasible α-discounted expected occupation measures and let $\mathfrak{M}^*_\alpha$ denote the set of optimal α-discounted occupation measures.

As with the long-term average linear program, we wish to show that the minimum cost for this linear program is equal to the minimum cost for the α-discounted stochastic problem; i.e.,

$$J^*_\alpha := \inf\{\langle c,\mu_\alpha\rangle : \mu_\alpha\in\mathfrak{M}^*_\alpha\} = \langle c,\mu^*_\alpha\rangle,$$

where μ_α^* denotes an element of $\mathfrak{M}_\alpha^*$. We proceed in much the same manner by first presenting a slightly modified version of Theorem 6.3 in Kurtz and Stockbridge (1998) as the following lemma. Note that the aforementioned modification is our use of unnormalized measures, as opposed to probability measures.

Lemma 2.5. *Let $\alpha \in (0,1]$ be given, and suppose $\{E \times G, L, c, \nu_0\}$ is a control model in which L satisfies (D1)-(D6) and c satisfies (C1)-(C3). Then the stochastic problem of minimizing*

$$J_\alpha(X, \Lambda; \nu_0) := \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t} \int_G c(X_t, u)\, \Lambda_t(du)\, dt\right]$$

over all $(X, \Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$ is equivalent to the linear program of minimizing

$$\langle c, \mu_\alpha \rangle = \int_{E \times G} c(x, u)\, \mu_\alpha(dx, du)$$

over all $\mu_\alpha \in \mathcal{M}(E \times G)$ with $\langle 1, \mu_\alpha \rangle = \alpha^{-1}$ satisfying $\langle c, \mu_\alpha \rangle < +\infty$ and

$$\langle L_\alpha f, \mu_\alpha \rangle = -\int_E f\, d\nu_0, \qquad \forall f \in \mathscr{D}(L). \tag{MD1}$$

What we then seek to verify is the following theorem for our constrained control model.

Theorem 2.6. *Let $\alpha \in (0,1]$ be given, and suppose $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ is a constrained control model in which L satisfies (D1)-(D6) and c and c_1 satisfy (C1)-(C3). Then the α-discounted stochastic problem of minimizing*

$$J_\alpha(X, \Lambda; \nu_0) := \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t} \int_G c(X_t, u)\, \Lambda_t(du)\, dt\right]$$

over all $(X, \Lambda) \in \mathscr{M}$ with $Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta_\alpha$ is equivalent to the α-discounted linear program; i.e.,

$$J_\alpha^* = \inf\{\langle c, \mu_\alpha \rangle : \mu_\alpha \in \mathfrak{M}_\alpha\}.$$

Moreover, there exists an optimal $\mu_\alpha^* \in \mathfrak{M}_\alpha^*$ satisfying $\langle c, \mu_\alpha^* \rangle = J_\alpha^* < +\infty$, and a corresponding $(X^*, \Lambda^*) = (X^*, \eta_\alpha^*(X^*, \cdot)) \in \mathscr{M}$ satisfying $J_\alpha(X^*, \Lambda^*; \nu_0) = J_\alpha^*$, $Q_1^\alpha(X^*, \Lambda^*; \nu_0) \leq \theta_\alpha$, and $X_0^* \sim \nu_0$, where η_α^* is the regular conditional distribution of μ_α^*.

Proof. Let $(X, \Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$, $J(X, \Lambda; \nu_0) < +\infty$, and $Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta_\alpha$, and define the measure $\hat{\mu}_\alpha \in \mathcal{P}(E \times G)$ by

$$\int_{E \times G} I_\Gamma(x, u)\, \hat{\mu}_\alpha(dx, du) = \alpha \mathbb{E}_{\nu_0}\left[\int_0^\infty e^{-\alpha t} \int_G I_\Gamma(X_t, u)\, \Lambda_t(du)\, dt\right], \qquad \forall \Gamma \in \mathcal{B}(E \times G).$$

Then by condition (C3) and the proof of Lemma 1.15, we have

$$\int_0^\infty e^{-\alpha t}\, \mathbb{E}\left[\int_G \psi(X_t, u)\, \Lambda_t(du)\right] dt < +\infty;$$

and, by the argument given in the proof of Proposition 1.17, we have

$$\langle L_\alpha f, \hat{\mu}_\alpha \rangle = -\alpha \int_E f\, d\nu_0, \qquad \forall f \in \mathscr{D}(L).$$

The definition of $\hat{\mu}_\alpha$ then implies that

$$J_\alpha(X, \Lambda; \nu_0) = \frac{1}{\alpha} \int_{E \times G} c(x, u)\, \hat{\mu}_\alpha(dx, du).$$

Thus, the rescaled measure $\mu_\alpha \in \mathcal{M}(E \times G)$ defined by $\mu_\alpha \coloneqq \alpha^{-1} \hat{\mu}_\alpha$ is an α-discounted expected occupation measure that satisfies $\langle 1, \mu_\alpha \rangle = \alpha^{-1}$ and

$$J_\alpha(X, \Lambda; \nu_0) = \int_{E \times G} c(x, u)\, \mu_\alpha(dx, du) = \langle c, \mu_\alpha \rangle.$$

Moreover, since $Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta$, the lower semicontinuity of c_1 ensures that $\langle c_1, \mu_\alpha \rangle \leq \theta$, as well. Therefore $\mu_\alpha \in \mathfrak{M}_\alpha$.

Conversely, if μ_α satisfies

$$\langle L_\alpha f, \mu_\alpha \rangle = -\int_E f \, d\nu_0, \qquad \forall f \in \mathscr{D}(L),$$

$\langle c, \mu_\alpha \rangle < +\infty$, and $\langle c_1, \mu_\alpha \rangle \leq \theta$, then condition (C3) ensures that the hypotheses of Corollary 5.3 in Kurtz and Stockbridge (1998) (or Theorem 1.18 above) are satisfied; and, hence, there exists a $(X, \Lambda) = (X, \eta_\alpha(X, \cdot)) \in \mathscr{M}$ with

$$J_\alpha(X, \Lambda; \nu_0) = \int_{E \times G} c(x, u) \, \mu_\alpha(dx, du) = \langle c, \mu_\alpha \rangle,$$

where η_α is the regular conditional distribution of μ_α. Moreover, the lower semicontinuity of c_1 ensures that $Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta$. □

We again remark that each α-discounted expected occupation measure μ_α is a member of $\mathcal{M}(E \times G)$, but μ_α is not necessarily a probability measure (unless, of course, $\alpha = 1$). Thus, to each $\mu_\alpha \in \mathfrak{M}_\alpha$ we must associate an *α-normalized expected occupation measure* $\hat{\mu}_\alpha := \alpha \mu_\alpha \in \mathcal{P}(E \times G)$ when our objective is to apply the vanishing discount method. This will be the subject of our next chapter.

THE VANISHING DISCOUNT METHOD

In this chapter, we wish to formalize our statement of the vanishing discount method. With this desire in mind, we begin by introducing the α-normalized linear program. This linear program will essentially "replace" the α-discounted linear program in what remains of this discussion. Even though it is only a simple "rescaling" of the α-discounted linear program, the α-normalized linear program nevertheless represents—to the best of our knowledge—an entirely novel contribution to the literature.

Having introduced the α-normalized linear program, we then present both a heuristic and a formal statement of the vanishing discount method using the linear programming framework we developed in the previous chapter.

III.1 The α-Normalized Linear Program

Given a discount rate $\alpha \in (0,1]$ and a constrained control model $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ in which L satisfies (D1)-(D6) and c and c_1 satisfy (C1)-(C3), we define the *α-normalized linear program* as follows:

$$\text{Minimize} \quad \langle c, \hat{\mu}_\alpha \rangle \qquad \text{subject to} \qquad \begin{cases} \langle L_\alpha f, \hat{\mu}_\alpha \rangle = -\alpha \int_E f \, d\nu_0, \quad \forall f \in \mathscr{D}(L); & \text{(MN1)} \\ \langle 1, \hat{\mu}_\alpha \rangle = 1; & \text{(MN2)} \\ \hat{\mu}_\alpha \in \mathcal{P}(E \times G); & \text{(MN3)} \\ \langle c_1, \hat{\mu}_\alpha \rangle \leq \alpha \theta_\alpha. & \text{(CN1)} \end{cases}$$

For this linear program, let $\hat{\mathfrak{M}}_\alpha$ denote the set of feasible α-normalized expected occupation measures and let $\hat{\mathfrak{M}}_\alpha^*$ denote the set of optimal α-normalized expected occupation measures.

Again we note that $\hat{\mu}_\alpha$ is an α-normalized expected occupation measure if and only if $\hat{\mu}_\alpha = \alpha\mu_\alpha$ for some α-discounted expected occupation measure μ_α. However, for sake of consistency and completeness, we include the following equivalence theorem for the α-normalized linear program, as well.

Theorem 1.1. *Let $\alpha \in (0, 1]$ be given, and suppose $\{E \times G, L, c, c_1, \theta_\alpha, \nu_0\}$ is a constrained control model in which L satisfies (D1)-(D6) and c and c_1 satisfy (C1)-(C3). Then the α-discounted linear program is equivalent to the α-normalized linear program in the sense that the optimal value of the α-normalized linear program is given by*

$$\hat{J}^*_\alpha := \alpha J^*_\alpha = \inf\{\langle c, \hat{\mu}_\alpha\rangle : \hat{\mu}_\alpha \in \hat{\mathfrak{M}}_\alpha\} = \inf\{\langle c, \alpha\mu_\alpha\rangle : \mu_\alpha \in \mathfrak{M}_\alpha\} = \langle c, \hat{\mu}^*_\alpha\rangle = \langle c, \alpha\mu^*_\alpha\rangle,$$

and that $\hat{\mu}^*_\alpha \in \hat{\mathfrak{M}}^*_\alpha$ if and only if $\mu^*_\alpha \in \mathfrak{M}^*_\alpha$, where $\hat{\mu}^*_\alpha = \alpha\mu^*_\alpha$. Moreover, there exists an optimal $\hat{\mu}^*_\alpha \in \hat{\mathfrak{M}}^*_\alpha$ satisfying $\langle c, \hat{\mu}^*_\alpha\rangle = \hat{J}^*_\alpha < +\infty$, and a corresponding $(X^*, \Lambda^*) = (X^*, \hat{\eta}^*_\alpha(X^*, \cdot)) \in \mathscr{M}$ satisfying $J_\alpha(X^*, \Lambda^*; \nu_0) = J^*_\alpha$, $Q^\alpha_1(X^*, \Lambda^*; \nu_0) \leq \theta_\alpha$, and $X^*_0 \sim \nu_0$, where $\hat{\eta}^*_\alpha$ is the regular conditional distribution of $\hat{\mu}^*_\alpha$.

Proof. Let $\mu^*_\alpha \in \mathfrak{M}^*_\alpha$ (the existence of which is guaranteed by Theorem II.2.6) and let $\hat{\mu}^*_\alpha := \alpha\mu^*_\alpha$. Since μ^*_α satisfies (MD1)-(MD3) and (CD1), a simple rescaling by the factor α makes it is clear that $\hat{\mu}^*_\alpha$ satisfies (MN1)-(MN3) and (CN1). By the very same argument, it is also clear that $\hat{\mu}^*_\alpha \in \hat{\mathfrak{M}}^*_\alpha$ and

$$\hat{J}^*_\alpha := \alpha J^*_\alpha = \alpha \inf\{\langle c, \mu_\alpha\rangle : \mu_\alpha \in \mathfrak{M}_\alpha\} = \inf\{\langle c, \alpha\mu_\alpha\rangle : \mu_\alpha \in \mathfrak{M}_\alpha\} = \inf\{\langle c, \hat{\mu}_\alpha\rangle : \hat{\mu}_\alpha \in \hat{\mathfrak{M}}_\alpha\};$$

and so

$$\hat{J}^*_\alpha = \langle c, \hat{\mu}^*_\alpha\rangle = \int_{E\times G} c(x, u)\, \hat{\mu}^*_\alpha(dx, du) = \alpha \int_{E\times G} c(x, u)\, \mu^*_\alpha(dx, du) = \langle c, \alpha\mu^*_\alpha\rangle,$$

as desired. □

III.2 Heuristic View

We can begin our discussion by analyzing the asymptotic behavior of the (α-parameterized) family of α-normalized linear programs as $\alpha \downarrow 0$. In other words, we will examine the asymptotic behavior of each of the *constraints* in the α-normalized linear program as $\alpha \downarrow 0$ and compare these results to the constraints of the long-term average linear program. Roughly speaking, we seek to verify that weak limits (as $\alpha \downarrow 0$) of α-normalized occupation measures are feasible for the long-term average linear program.

It will therefore be useful to let

$$\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$$

denote *a constrained control model for the vanishing discount method* (or, more succinctly, *a vanishing discount model*), recalling that θ is a constraint constant for the long-term average linear program and $\{\theta_\alpha\} := \{\theta_\alpha : \alpha \in (0,1]\}$ is a collection of constraint constants for the family of α-discounted linear programs (or, equivalently, $\{\alpha\theta_\alpha : \alpha \in (0,1]\}$ is a collection of constraint constants for the family of α-normalized linear programs). As a means of simplifying exposition, whenever we refer to the model $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$, it will be assumed that it satisfies the conditions we have specified up until this point. That is: E and G (and, hence, $E \times G$) are locally compact, complete, separable, metric spaces; L satisfies conditions (D1)-(D6); c and c_1 satisfy (C1)-(C3); and θ and $\{\theta_\alpha\}$ satisfy condition (T1). Note that we have intentionally omitted the initial distribution ν_0 from this model to emphasize that any dependence on this parameter, in fact, "vanishes" as $\alpha \downarrow 0$.

III.2.1 Feasibility of Weak Limits

Let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence in $(0,1]$ with $\alpha_n \to 0$ as $n \to \infty$, and let $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ be a sequence of α_n-normalized occupation measures with $\hat{\mu}_{\alpha_n} \in \hat{\mathfrak{M}}_{\alpha_n}$ for every $n \in \mathbb{N}$. Recall that, by definition,

$$\begin{aligned}\langle L_{\alpha_n} f, \hat{\mu}_{\alpha_n} \rangle &= \int_{E\times G} L_{\alpha_n} f(x,u)\, \hat{\mu}_{\alpha_n}(dx,du) \\ &= \int_{E\times G} Lf(x,u)\, \hat{\mu}_{\alpha_n}(dx,du) - \alpha_n \int_E f(x)\, \hat{\mu}^E_{\alpha_n}(dx),\end{aligned}$$

where $\hat{\mu}^E_{\alpha_n}$ denotes the state marginal measure of $\hat{\mu}_{\alpha_n}$. Thus, if $(X,\Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$ and $\hat{\mu}_\alpha$ is the α-normalized occupation measure associated with (X,Λ), then the adjoint condition (MN1) can be written

$$\int_{E\times G} Lf(x,u)\, \hat{\mu}_{\alpha_n}(dx,du) - \alpha_n \int_E f(x)\, \hat{\mu}^E_{\alpha_n}(dx) = -\alpha_n \int_E f(x)\, \nu_0(dx), \qquad \forall f \in \mathscr{D}(L).$$

So, since $\mathscr{D}(L) \subset \widehat{C}(E)$, $\hat{\mu}^E_{\alpha_n} \in \mathcal{P}(E)$, and $\alpha_n \to 0$, it is clear that

$$\lim_{n\to\infty} \left| \alpha_n \int_E f(x)\, \hat{\mu}^E_{\alpha_n}(dx) \right| \le \lim_{n\to\infty} \alpha_n \|f\|_\infty = 0, \qquad \forall f \in \mathscr{D}(L).$$

Similarly, since $\nu_0 \in \mathcal{P}(E)$, we also see that

$$\lim_{n\to\infty} \alpha_n \int_E f(x)\, \nu_0(dx) = 0, \qquad \forall f \in \mathscr{D}(L).$$

Thus, we have

$$\lim_{n\to\infty} \langle L_{\alpha_n} f, \hat{\mu}_{\alpha_n} \rangle = \lim_{n\to\infty} \langle Lf, \hat{\mu}_{\alpha_n} \rangle = 0, \qquad \forall f \in \mathscr{D}(L).$$

Now, suppose that $\mu_0 \in \mathcal{P}(E \times G)$ with $\hat{\mu}_\alpha \Rightarrow \mu_0$. Intuition then *suggests* that

$$\lim_{n\to\infty} \langle L_{\alpha_n} f, \hat{\mu}_{\alpha_n} \rangle = \langle Lf, \mu_0 \rangle, \qquad \forall f \in \mathscr{D}(L); \tag{2.1}$$

i.e., μ_0 satisfies the adjoint condition (MA1) of the long-term average linear program. Furthermore, as μ_0 is a weak limit of probability measures, it trivially satisfies the mass conditions (MA2) and (MA3) of the long-term average linear program. As for the budget constraint (CA1) of the long-term average linear program, we have the following proposition.

Proposition 2.2. *Let $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ be a constrained control model for the vanishing discount method, and let $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ be a sequence of α_n-normalized occupation measures with $\hat{\mu}_{\alpha_n} \in \hat{\mathfrak{M}}_{\alpha_n}$ for every $n \in \mathbb{N}$. If $\mu_0 \in \mathcal{P}(E \times G)$ with $\hat{\mu}_{\alpha_n} \Rightarrow \mu_0$, then*

$$\langle c_1, \mu_0 \rangle \leq \theta;$$

i.e., μ_0 satisfies the budget constraint (CA1) of the long-term average linear program.

Proof. Since c_1 is lower semicontinuous and bounded below, Corollary A.5 tells us that

$$\langle c_1, \mu_0 \rangle = \int_{E \times G} c_1(x,u)\, \mu_0(dx, du) \leq \liminf_{n \to \infty} \int_{E \times G} c_1(x,u)\, \hat{\mu}_{\alpha_n}(dx, du) = \liminf_{n \to \infty} \langle c_1, \hat{\mu}_{\alpha_n} \rangle.$$

So, since $\langle c_1, \hat{\mu}_{\alpha_n} \rangle \leq \alpha_n \theta_{\alpha_n}$ for every $n \in \mathbb{N}$, condition (T1) then tells us

$$\langle c_1, \mu_0 \rangle \leq \liminf_{n \to \infty} \langle c_1, \hat{\mu}_{\alpha_n} \rangle \leq \lim_{n \to \infty} \alpha_n \theta_n = \theta.$$

Hence, μ_0 satisfies (CA1). □

So, in order for μ_0 to be feasible for the long-term average linear program (i.e., $\mu_0 \in \mathfrak{M}$), we need only to verify that it satisfies the adjoint condition (MA1). In due time, we will see that this is, in fact, true; but we delay the formal statement and proof of this fact in order to discuss the stability of optimal solutions (and optimal values) for the family of α-normalized linear programs under passage to a limit as $\alpha \downarrow 0$.

III.2.2 Optimality

As with our discussion of weak limits of *feasible* α-normalized occupation measures, our hope is that

$$\lim_{\alpha\downarrow 0}\langle L_\alpha f, \hat{\mu}^*_\alpha\rangle = \langle Lf, \mu^*_0\rangle = 0, \qquad \forall f \in \mathscr{D}(L), \tag{2.3}$$

for any μ^*_0 that is obtained as a weak limit of a sequence $\{\hat{\mu}^*_\alpha : \alpha \in (0,1]\}$ of *optimal* α-normalized occupation measures as $\alpha \downarrow 0$. Furthermore, we wish that $\langle c, \mu^*_0\rangle \le \theta$ and

$$\lim_{\alpha\downarrow 0}\langle c, \hat{\mu}^*_\alpha\rangle = \lim_{n\to\infty} \hat{J}^*_\alpha = \langle c, \mu^*_0\rangle = J^*, \tag{2.4}$$

as this would imply that the optimal value of the long-term average linear program can be obtained by taking such a limit. Indeed, satisfaction of the conditions (2.3) and (2.4) is what one may describe as the primary objective of the vanishing discount method since, practically speaking, one is often only interested in "solving" a control problem. However, by further analyzing the behavior of the feasible measures for these linear programs, one obtains a deeper insight into the theoretical structures of these problems and how they are related.

Of course, the validity of statements (2.1), (2.3), and (2.4) remains to be verified since, for instance, the function Lf is not necessarily bounded. Furthermore, we would like to be able to verify if such weak limits μ_0 and μ^*_0 necessarily *exist*. Hence, the task before is to rigorously analyze these statements and, when necessary, establish further conditions under which each of these statements is guaranteed to hold.

III.3 Formal Statement

To now formalize the heuristic presentation above, we define the following "limiting" objects. Let $\mathfrak{M}_0$ be the set of $\mu_0 \in \mathcal{P}(E \times G)$ for which there exists a sequence $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ of *feasible* α_n-normalized expected occupation measures such that $\alpha_n \to 0$ and $\hat{\mu}_{\alpha_n} \Rightarrow \mu_0$, and let $\mathfrak{M}_0^*$ be the set of $\mu_0^* \in \mathfrak{M}_0$ for which there exists a sequence $\{\hat{\mu}_{\alpha_n}^* : n \in \mathbb{N}\}$ of *optimal* α_n-normalized expected occupation measures such that $\alpha_n \to 0$ and $\hat{\mu}_{\alpha_n}^* \Rightarrow \mu_0^*$. Finally, define

$$J_0^* := \lim_{\alpha \downarrow 0} \hat{J}_\alpha^*$$

when this limit exists.

Our goal is then to identify sufficient conditions under which the following "VDM Relations" hold for the vanishing discount model $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$:

(V1) $\emptyset \neq \mathfrak{M}_0 \subset \mathfrak{M}$.

(V2) $\emptyset \neq \mathfrak{M}_0^* \subset \mathfrak{M}^*$.

(V3) $J_0^* = J^*$.

We can describe these relations in words as follows:

(V1) There exists at least one weak limit μ_0 of some sequence $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ of *feasible* α_n-normalized occupation measures (where $\alpha_n \to 0$). Moreover, any such weak limit of any such sequence is feasible for the long-term average linear program.

(V2) There exists a weak limit μ_0^* of some sequence $\{\hat{\mu}_{\alpha_n}^* : n \in \mathbb{N}\}$ of *optimal* α_n-normalized occupation measures (where $\alpha_n \to 0$). Moreover, any such weak limit of any such sequence is optimal for the long-term average linear program.

(V3) The limit (as $\alpha \downarrow 0$) of the optimal values of the (α-indexed) family of α-normalized linear programs is equal to the optimal value of the long-term average linear program.

Of course, any control model satisfying relations (V1)-(V3) is tautologically one for which the vanishing discount method is applicable. What we must now set out to establish is verifiable *hypotheses* for our vanishing discount model whose satisfaction will ensure that (V1)-(V3) will hold. The theory of *correspondences* will provide us with a practical (and elegant) way of presenting and analyzing these hypotheses; and, hence, this theory—and its application to our formulation of the vanishing discount method—will be the subject of the next chapter.

CORRESPONDENCES AND VANISHING DISCOUNT

In this chapter, our focus will be the theory of *correspondences* and its role in our formulation of the vanishing discount method. In particular, we will see that the conditions we desire can be stated in a manner similar to that of *Berge's Theorem* (see Section B of this document's Appendix for a statement of this result), though with slightly weaker hypotheses and conclusions. Indeed, this theorem will play a crucial role in establishing the appropriate sufficient conditions for the model $\{E\times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ to ensure that the VDM Relations (V1)-(V3) hold.

Integral to this discussion will be the notion of *hemicontinuity*, the analysis of which begins this chapter. Intuitively speaking, a correspondence φ is *upper hemicontinuous* at a point α_0 in its domain if $\varphi(\alpha)$ does not "explode" as α moves away from α_0, though it may "implode;" and φ is *lower hemicontinuous* at α_0 if $\varphi(\alpha)$ does not implode as α moves away from α_0, but it may explode. Some useful illustrations of these phenomena can be found in Section 14.1 of Sydsaeter et al. (2005). The unfamiliar reader is again directed to Section B of this document's Appendix for many of the basic definitions and results regarding correspondences, hemicontinuity, and Berge's Theorem. Virtually all of these facts come from Chapter 17 in Aliprantis and Border (2006), if a much more detailed and general overview is desired. The properties of correspondences that we present in this chapter will simply be those facts that are most relevant to the framework we have hitherto developed in this dissertation.

IV.1 Hemicontinuity

For our purposes, we need only consider correspondences whose domain and codomain are, respectively, the sets $[0,1] \subset \mathbb{R}$ and $\mathcal{P}(E \times G)$, rather than general topological spaces. The interval $[0,1]$ is, of course, a metric space under the standard Euclidean metric it inherits as a subspace of $\mathbb{R}$; and since $E \times G$ is a complete, separable, metric space, it is well known that $\mathcal{P}(E \times G)$ is also a complete, separable, metric space under the Prohorov metric, which we now define for reference purposes.

Definition 1.1. Let $\mathscr{C}$ be the collection of closed subsets of $E \times G$, and let d be a metric on $E \times G$. Then the *Prohorov metric* π on $\mathcal{P}(E \times G)$ is defined by

$$\pi(\mu, \nu) := \inf \{\epsilon > 0 : \mu(\Phi) \leq \nu(\Phi^{\epsilon}) + \epsilon, \forall \Phi \in \mathscr{C}\},$$

where

$$\Phi^{\epsilon} := \{a \in E \times G : \inf\{d(a,b) < \epsilon : b \in \Phi\}\}.$$

◇

A detailed proof that π, in fact, defines a metric on $\mathcal{P}(E \times G)$ may be found in Section 1 of Chapter 3 in Ethier and Kurtz (1986). The Portmanteau Theorem (i.e., Theorem A.4) furthermore shows that weak convergence in $\mathcal{P}(E \times G)$ is equivalent to convergence in the Prohorov metric.

Now, since both $[0,1]$ and $\mathcal{P}(E \times G)$ are metric spaces, we may avail ourselves of the following sequential characterization of upper hemicontinuity (the likes of which, in fact, requires only that the domain and codomain of φ be first countable and metrizable, respectively). The next two lemmas appear in Section 17.3 of Aliprantis and Border (2006), but their proofs are left "as exercises" to the reader; and so, for sake of completeness, we include the required proofs here. Note, however, that these proofs are essentially modifications of proofs for the more general results involving *nets* found in the same section. Again, the metric space structure of both $[0, 1]$ and $\mathcal{P}(E \times G)$ allows us to work specifically with sequences instead of nets.

Recall that $\alpha \in [0, 1]$ is a *limit point* of the sequence $\{\alpha_n : n \in \mathbb{N}\} \subset [0, 1]$ if, for each $\epsilon > 0$ and each $n \in \mathbb{N}$, there exists some integer $n_k \geq n$ such that $\alpha_{n_k} \in (\alpha - \epsilon, \alpha + \epsilon)$. Further recall that α is a limit point of $\{\alpha_n : n \in \mathbb{N}\}$ if and only if α is the limit of some subsequence $\{\alpha_{n_k} : k \in \mathbb{N}\}$ of $\{\alpha_n : n \in \mathbb{N}\}$. The unconvinced reader is directed to Theorem 2.16 in Aliprantis and Border (2006) for a proof of the preceding statement.

Lemma 1.2. *If $\varphi : [0, 1] \twoheadrightarrow \mathcal{P}(E \times G)$ is a correspondence and $\alpha \in [0, 1]$, then the following statements are equivalent.*

(a) φ is upper hemicontinuous at α and $\varphi(\alpha)$ is compact.

(b) If a sequence $\{(\alpha_n, \mu_n) : n \in \mathbb{N}\}$ in $\mathrm{Gr}(\varphi)$ satisfies $\alpha_n \to \alpha$, then the sequence $\{\mu_n : n \in \mathbb{N}\}$ has a limit point μ in $\varphi(\alpha)$; i.e., there exists a subsequence $\{(\alpha_{n_k}, \mu_{n_k}) : k \in \mathbb{N}\}$ of $\{(\alpha_n, \mu_n) : n \in \mathbb{N}\}$ and a $\mu \in \varphi(\alpha)$ such that $\mu_{n_k} \Rightarrow_k \mu$.

Proof. (a) $\Longrightarrow$ (b): Assume that φ is upper hemicontinuous at $\alpha \in [0, 1]$ with $\varphi(\alpha)$ compact, and let $\{(\alpha_n, \mu_n) : n \in \mathbb{N}\}$ be a sequence in $\mathrm{Gr}(\varphi)$ that satisfies $\alpha_n \to \alpha$. Note that the upper hemicontinuity of φ at α guarantees that $\varphi(\alpha) \neq \emptyset$. Now, in pursuit of a contradiction, suppose that $\{\mu_n : n \in N\}$ has no limit point in $\varphi(\alpha)$. This implies that, for every $\mu \in \varphi(\alpha)$, there is an open neighborhood V_μ of μ and an integer N_μ such that, for all $n \geq N_\mu$, we have $\mu_n \notin V_\mu$. Since $\varphi(\alpha)$ is compact, it lies in some finite union $V := V_{\mu_1} \cup \cdots \cup V_{\mu_k}$. So choose an $N_0 \in \mathbb{N}$ such that $N_0 \geq N_{\mu_i}$ for each $i = 1, \ldots, k$. Then, for all $n \geq N_0$, we must have $\mu_n \notin V$. However, since φ is upper hemicontinuous at α, for large enough n we must have $\mu_n \in \varphi(\alpha_n) \subset V$. This is a contradiction.

(b) $\Longrightarrow$ (a): By way of contraposition, suppose that φ is not upper hemicontinuous at α. Then there exists an open neighborhood U of $\varphi(\alpha)$ such that, for large enough $n \in \mathbb{N}$, each open neighborhood $V_n := (\alpha - n^{-1}, \alpha + n^{-1}) \cap [0, 1]$ of α contains an α_n for which there exists a $\mu_n \in \varphi(\alpha_n)$ with $\mu_n \notin U$. Clearly, the sequence $\{\mu_n : n \in \mathbb{N}\}$ does not have a limit point in $\varphi(\alpha)$. This shows that φ must be upper hemicontinuous at α.

To show that $\varphi(\alpha)$ is compact, suppose that (b) holds and let $\{\mu_n : n \in \mathbb{N}\}$ be a sequence in $\varphi(\alpha)$. By choosing $\alpha_n = \alpha$ for each $n \in \mathbb{N}$, the sequence $\{\alpha_n : n \in \mathbb{N}\}$ trivially satisfies $\alpha_n \to \alpha$; and by (b), the sequence $\{\mu_n : n \in \mathbb{N}\}$ has a limit point in $\varphi(\alpha)$. Since the sequence $\{\mu_n : n \in \mathbb{N}\}$ is arbitrary, $\varphi(\alpha)$ is sequentially compact; and since $\varphi(\alpha)$ is a subset of the metric space $\mathcal{P}(E \times G)$, it follows that $\varphi(\alpha)$ is compact. □

As with upper hemicontinuity, we have the following sequential characterization of lower hemicontinuity. Observe that, unlike the conclusion of Lemma 1.2, no additional structure is obtained on the set $\varphi(\alpha)$ when φ is lower hemicontinuous at α.

Lemma 1.3. *If $\varphi : [0,1] \twoheadrightarrow \mathcal{P}(E \times G)$ is a correspondence and $\alpha \in [0,1]$, then the following statements are equivalent.*

(a) φ is lower hemicontinuous at α.

(b) If $\alpha_n \to \alpha$ then, for each $\mu \in \varphi(\alpha)$, there exists a subsequence $\{\alpha_{n_k} : k \in \mathbb{N}\}$ of the sequence $\{\alpha_n : n \in \mathbb{N}\}$ and a sequence $\{\mu_k : k \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$ with $\mu_k \in \varphi(\alpha_{n_k})$ for each $k \in \mathbb{N}$ such that $\mu_k \Rightarrow \mu$.

Proof. (a) $\implies$ (b): Assume that φ is lower hemicontinuous at α_0. Let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence that satisfies $\alpha_n \to \alpha_0$ and fix $\mu_0 \in \varphi(\alpha_0)$. Now, fix $k \in \mathbb{N}$ and define the open neighborhood

$$U_k \coloneqq \{\mu \in \mathcal{P}(E \times G) : \pi(\mu, \mu_0) < k^{-1}\}$$

of ν_0. Then $\mu_0 \in \varphi(\alpha_0) \cap U_k$ and, since φ is lower hemicontinuous at α_0, the set

$$\varphi^\ell(U_k) = \{\alpha \in [0,1] : \varphi(\alpha) \cap U_k \neq \emptyset\}$$

is an open neighborhood of α_0. So, define the set $V_k \coloneqq (\alpha_0 - k^{-1}, \alpha_0 + k^{-1}) \cap [0,1]$. Then there is a $n_k \in \mathbb{N}$ such that $\alpha_n \in V_k \cap \varphi^\ell(U_k)$ for every $n \geq n_k$; in particular, $\alpha_{n_k} \in V_k \cap \varphi^\ell(U_k)$.

We may therefore form a subsequence $\{\alpha_{n_k} : k \in \mathbb{N}\}$ of $\{\alpha_n : n \in \mathbb{N}\}$ by choosing such an α_{n_k} for each $k \in \mathbb{N}$. Moreover, the set $\varphi(\alpha_{n_k}) \cap U_k$ is nonempty for each $k \in \mathbb{N}$; and so, we may select a $\mu_k \in \varphi(\alpha_{n_k}) \cap U_k$ to form the sequence $\{\mu_k : k \in \mathbb{N}\} \subset \mathcal{P}(E \times G)$ which evidently satisfies $\mu_k \Rightarrow \mu_0$.

(b) $\implies$ (a): Assume that (b) holds and, in pursuit of a contradiction, assume that φ is not lower hemicontinuous at α. Then there exists an open set $U \subset \mathcal{P}(E \times G)$ with $\varphi(\alpha) \cap U \neq \emptyset$ such that, for any $n \in \mathbb{N}$, there is some $\alpha_n \in (\alpha - n^{-1}, \alpha + n^{-1})$ with $\varphi(\alpha_n) \cap U = \emptyset$. It is clear that the sequence $\{\alpha_n : n \in \mathbb{N}\}$ satisfies $\alpha_n \to \alpha$. So, let $\mu \in \varphi(\alpha) \cap U$. Then, by (b), we can assume (by passing to a subsequence if necessary), that there is a sequence $\{\mu_n : n \in \mathbb{N}\}$ with $\mu_n \Rightarrow \mu$ and $\mu_n \in \varphi(\alpha_n)$ for each n. Since $\varphi(\alpha_n) \cap U = \emptyset$ for each n, it follows that $\mu_n \in U^{\mathsf{c}}$ for each n. However, U^{c} is a closed set; and so, we must have $\mu \in U^{\mathsf{c}}$. This, of course, contradicts the fact that $\mu \in U$. $\square$

In what follows, we will take the above characterizations of upper and lower hemicontinuity to serve more-or-less as *definitions* of these properties in the sequel. Also, before continuing with the next section, the reader is reminded that a correspondence is said to be *continuous at* α if it is both upper hemicontinuous and lower hemicontinuous at α; and a correspondence is *continuous* if it is continuous at each α in its domain. Be aware that, in particular, one should not refer to a correspondence as "hemicontinuous."

IV.2 Vanishing Discount Results

This section contains the main result of this dissertation, the aim of which is to provide the weakest possible (sufficient) conditions under which the vanishing discount applies for the model $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$. As stated above, this result will be presented in a manner similar to that of Berge's Theorem; and, hence, the hypotheses of our result will be stated in terms of the following objects.

Define the "feasibility correspondence" $\varphi : [0,1] \twoheadrightarrow \mathcal{P}(E \times G)$ by

$$\varphi(\alpha) \coloneqq \begin{cases} \mathfrak{M} & \text{if } \alpha = 0, \\ \hat{\mathfrak{M}}_\alpha & \text{if } \alpha \in (0,1]; \end{cases} \tag{2.1}$$

and define the "objective function" $F : \mathrm{Gr}(\varphi) \to (-\infty, +\infty]$ by

$$F(\alpha, \mu_\alpha) \coloneqq \begin{cases} \langle c, \rho \rangle & \text{if } \alpha = 0, \\ \langle c, \hat{\mu}_\alpha \rangle & \text{if } \alpha \in (0,1], \end{cases} \tag{2.2}$$

recalling that ρ denotes an element of $\mathfrak{M}$ and $\hat{\mu}_\alpha$ denotes an element of $\hat{\mathfrak{M}}_\alpha$. Note that we necessarily have $F(\alpha, \mu_\alpha) > -\infty$ for every $(\alpha, \mu_\alpha) \in \mathrm{Gr}(\varphi)$ because c is bounded below.
Now define the "argmin correspondence" $\varphi^* : [0,1] \to \mathcal{P}(E \times G)$ by

$$\varphi^*(\alpha) \coloneqq \begin{cases} \mathfrak{M}^* & \text{if } \alpha = 0, \\ \hat{\mathfrak{M}}^*_\alpha & \text{if } \alpha \in (0,1]; \end{cases} \tag{2.3}$$

noting that φ^* is a *subcorrespondence* of φ (i.e., $\varphi^*(\alpha) \subset \varphi(\alpha)$ for each $\alpha \in [0,1]$); and define the "value function" $F^* : [0,1] \to (-\infty, +\infty]$ by

$$F^*(\alpha) \coloneqq \inf\{F(\alpha, \mu_\alpha) : \mu_\alpha \in \varphi(\alpha)\}, \tag{2.4}$$

noting that $\inf(\emptyset) = +\infty$.
We now wish to establish conditions on the feasibility correspondence φ in (2.1) and the objective function F in (2.2) that will guarantee satisfaction of the VDM Relations

(V1) $\emptyset \neq \mathfrak{M}_0 \subset \mathfrak{M}$,

(V2) $\emptyset \neq \mathfrak{M}^*_0 \subset \mathfrak{M}^*$, and

(V3) $J^*_0 = J^*$,

introduced in the previous chapter. Our main results accomplish just that, and we state them as the next two theorems. The first theorem provides a more general result, whereas the second theorem is specific to our vanishing discount model $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$. We omit proofs of these theorems for the moment.

Theorem 2.5. *Let the feasibility correspondence φ and the objective function F be as in (2.1) and (2.2), respectively. If φ is a continuous at $\alpha = 0$, φ has nonempty compact values, and F is continuous, then (V1)-(V3) hold.*

Theorem 2.6. *Let the feasibility correspondence φ and the objective function F be as in (2.1) and (2.2), respectively, and suppose that $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ is a constrained control model for the vanishing discount method. If φ is lower hemicontinuous at $\alpha = 0$ and F is upper semicontinuous on $\{0\} \times \varphi(0)$, then (V1)-(V3) hold.*

Thus, when considering any constrained control model satisfying the conditions we have assumed in this manuscript, one can obtain the same conclusions of Theorem 2.5, but with weaker hypotheses.

We dedicate the next three subsections to a detailed analysis of each of the hypotheses (and conclusions) of Theorem 2.5 and Theorem 2.6 to see precisely how we have arrived at these statements. In particular, we will see how certain properties possessed by the argmin correspondence φ^* in (2.3) and value function F^* in (2.4) directly affect the optimality relations (V2) and (V3). Each of these subsections also contains an analysis of the role our model assumptions play in weakening the hypotheses of our main results.

Remark 2.7. As a preliminary observation, recall our assumption that there is at least one $(X, \Lambda) \in \mathscr{M}$ with $J(X, \Lambda) < +\infty$ and $Q_1(X, \Lambda) \leq \theta$; and so, by the equivalence theorems presented in Chapter II, we are guaranteed at least one $\rho \in \mathfrak{M}$ with $\langle c, \rho \rangle < +\infty$. Similarly, our assumption that, for each $\alpha > 0$ there is at least one $(X, \Lambda) \in \mathscr{M}$ with $J_\alpha(X, \Lambda; \nu) < +\infty$ and $Q_1^\alpha(X, \Lambda; \nu_0) \leq \theta_\alpha$ guarantees that, for each $\alpha \in (0, 1]$, there is at least one $\hat{\mu}_\alpha \in \hat{\mathfrak{M}}_\alpha$

with $\langle c, \hat{\mu}_\alpha \rangle < +\infty$. In other words, *we may assume that the feasibility correspondence φ has nonempty values.* $\diamond$

IV.2.1 Analysis of Conditions for (V1)

Our first observation in this analysis is that (V1) is satisfied if (and only if) the feasibility correspondence φ is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact. We separate this claim into its necessary and sufficient conditions, and show that this property is obtained "for free" by the conditions we have imposed on our vanishing discount model.

Lemma 2.8. *If the feasibility correspondence φ in (2.1) is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact, then $\emptyset \neq \mathfrak{M}_0 \subset \mathfrak{M}$.*

Proof. To first show that $\mathfrak{M}_0 \neq \emptyset$, let $\{(\alpha_n, \hat{\mu}_{\alpha_n}) : n \in \mathbb{N}\}$ be a sequence in $\mathrm{Gr}(\varphi)$ where $\alpha_n = n^{-1}$ for each $n \in \mathbb{N}$. Then $\alpha_n \to 0$; and so, by Lemma 1.2, the sequence $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ has a weak limit $\rho \in \varphi(0)$. It then follows (from the definition of $\mathfrak{M}_0$) that $\rho \in \mathfrak{M}_0$. Thus $\mathfrak{M}_0 \neq \emptyset$.

Now we show that $\mathfrak{M}_0 \subset \mathfrak{M}$. Let $\mu_0 \in \mathfrak{M}_0$. There is then a sequence $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ with $\hat{\mu}_{\alpha_n} \in \hat{\mathfrak{M}}_{\alpha_n} = \varphi(\alpha_n)$ for each $n \in \mathbb{N}$, satisfying $\alpha_n \to 0$ and $\hat{\mu}_{\alpha_n} \Rightarrow \mu_0$ (where we may assume without loss of generality that $\alpha_n > 0$ for each $n \in \mathbb{N}$). Since φ is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact, Lemma 1.2 yields a $\rho \in \varphi(0) = \mathfrak{M}$ satisfying $\hat{\mu}_{\alpha_n} \Rightarrow \rho$. Since $\hat{\mu}_{\alpha_n} \Rightarrow \mu_0$ and $\hat{\mu}_{\alpha_n} \Rightarrow \rho$, we must therefore have $\mu_0 = \rho$ (recalling that $\mathcal{P}(E \times G)$ is a metric space). Hence, $\mathfrak{M}_0 \subset \mathfrak{M}$, as desired. $\square$

We now observe that the conditions we have imposed on our vanishing discount model $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$, in fact, guarantee that φ is upper hemicontinuous and $\varphi(0)$ is compact. Thus, the hypotheses of Lemma 2.8 are "automatically" satisfied. We formalize this claim by way of the following proposition.

Proposition 2.9. *Suppose $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ is a constrained control model for the vanishing discount method. Then the feasibility correspondence φ in (2.1) is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact.*

Proof. Let $\{(\alpha_n, \hat{\mu}_{\alpha_n}) : n \in \mathbb{N}\}$ be a sequence in $\mathrm{Gr}(\varphi)$ where $\{\alpha_n : n \in \mathbb{N}\} \subset (0, 1]$ and $\alpha_n \to 0$, and define

$$\Theta \coloneqq \sup\{\alpha_n \theta_{\alpha_n} : n \in \mathbb{N}\}.$$

We then have $-\infty < -\kappa_1 < \langle c_1, \hat{\mu}_{\alpha_n} \rangle \leq \Theta$, for each $n \in \mathbb{N}$. So, given $M > 0$, define the set

$$K_M \coloneqq \{(x, u) \in E \times G : c_1(x, u) \leq M\},$$

noting that K_M is compact since c_1 is inf-compact. Now,

$$M \hat{\mu}_{\alpha_n}(K_M^{\mathsf{c}}) \leq \int_{K_M^{\mathsf{c}}} c_1(x, u)\, \hat{\mu}_{\alpha_n}(dx, du) \leq \Theta, \qquad \forall n \in \mathbb{N}.$$

Given $\epsilon > 0$, we can then choose M sufficiently large so that $M > \Theta/\epsilon$; hence,

$$\hat{\mu}_{\alpha_n}(K_M^{\mathsf{c}}) \leq \frac{\Theta}{M} < \epsilon, \qquad \forall n \in \mathbb{N}.$$

The sequence $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ is therefore tight (and, hence, relatively compact). Thus, there is a $\mu_0 \in \mathcal{P}(E \times G)$ satisfying $\hat{\mu}_{n_k} \Rightarrow \mu_0$ for some subsequence $\{\hat{\mu}_{n_k} : k \in \mathbb{N}\}$ of $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$; i.e., $\mu_0 \in \mathfrak{M}_0 \neq \emptyset$.

We now wish to show that $\mu_0 \in \varphi(0)$. In Section 2 of Chapter III we established that μ_0 trivially satisfies (MA2) and (MA3); and Proposition III.2.2 shows that μ_0 satisfies the budget constraint (CA1). To show that μ_0 satisfies the adjoint condition

$$\langle Lf, \mu_0 \rangle = 0, \qquad \forall f \in \mathscr{D}(L), \tag{MA1}$$

we employ an argument similar to the one that appears in the proof of Proposition II.1.4. That is, we define the sequence $\{\zeta_k : k \in \mathbb{N}\} \subset \mathcal{M}(E \times G)$ by

$$\zeta_k(\Gamma) = \int_\Gamma \psi(x,u)\,\hat{\mu}_{n_k}(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E \times G).$$

Then, given $\epsilon \in (0,1)$, we use condition (C3) and a suitably large choice of $R > 0$ to show that

$$\zeta_k(K_R^{\mathsf{c}}) \leq \int_{K_R^{\mathsf{c}}} a_1 c_1(x,u)\,\hat{\mu}_{n_k}(dx,du) + b_1 \hat{\mu}_{n_k}(K_R^{\mathsf{c}}) < \epsilon, \qquad \forall k \in \mathbb{N},$$

and, hence, conclude that $\{\zeta_k : k \in \mathbb{N}\}$ is tight. By then appropriately partitioning $E \times G$, it can be shown that

$$\zeta_k(E \times G) \leq a_1(\Theta + \kappa_1 + 1) + b_1 < +\infty, \qquad \forall k \in \mathbb{N};$$

and, hence, $\{\zeta_k(E \times G) : k \in \mathbb{N}\}$ is bounded. The relative compactness of $\{\zeta_k : k \in \mathbb{N}\}$ then yields a $\zeta \in \mathcal{M}(E \times G)$ satisfying $\zeta_{k_i} \Rightarrow \zeta$ for some subsequence $\{\zeta_{k_i} : i \in \mathbb{N}\}$ of $\{\zeta_k : k \in \mathbb{N}\}$ and

$$\zeta(\Gamma) = \int_\Gamma \psi(x,u)\,\mu_0(dx,du), \qquad \forall \Gamma \in \mathcal{B}(E \times G).$$

We have shown (again, in Section 2 of Chapter III) that

$$\lim_{n\to\infty} \langle Lf, \hat{\mu}_{\alpha_n} \rangle = 0, \qquad \forall f \in \mathscr{D}(L).$$

So, we then use condition (D6) to conclude that

$$0 = \lim_{n\to\infty} \langle Lf, \hat{\mu}_{\alpha_n} \rangle = \langle Lf, \mu_0 \rangle, \qquad \forall f \in \mathscr{D}(L),$$

which shows that μ_0 satisfies constraint (MA1) for the long-term average linear program. Thus, $\mu_0 \in \mathfrak{M} = \varphi(0)$; and so, φ is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact by Lemma 1.2. □

Though it is not imperative to our main results, we nevertheless make the observation that the *converse* of Lemma 2.8 holds for our model, as well. Note that Proposition 2.9 shows that $\emptyset \neq \mathfrak{M}_0$ when its hypotheses are satisfied, so we omit this condition from the hypothesis in the corollary below.

Corollary 2.10. *Suppose $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ is a constrained control model for the vanishing discount method. Then the feasibility correspondence φ in (2.1) is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact if $\mathfrak{M}_0 \subset \mathfrak{M}$.*

Proof. Let $\{(\alpha_n, \hat{\mu}_{\alpha_n}) : n \in \mathbb{N}\}$ be a sequence in $\mathrm{Gr}(\varphi)$ with $\alpha_n \to 0$ (and assume that $\alpha_n > 0$ for each $n \in \mathbb{N}$). Proposition 2.9 then guarantees the existence of a $\mu_0 \in \mathfrak{M}_0$ with $\hat{\mu}_{\alpha_n} \Rightarrow \mu_0$. Since $\mathfrak{M}_0 \subset \mathfrak{M}$, it follows that $\mu_0 \in \mathfrak{M}$; and so, by Lemma 1.2, φ is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact. □

IV.2.2 Analysis of Conditions for (V2)

Now, in consideration of the relation (V2), we have a result that is an analog of Lemma 2.8, but with the sets of *optimizers*. This result indicates that any weak limit μ_0^* of a family of optimal α-normalized occupation measures, as $\alpha \downarrow 0$, is also an optimizer (i.e., minimizer) for the long-term average linear program. As the proof is nearly identical to that of Lemma 2.8, we state this result as a corollary.

Corollary 2.11. *If the argmin correspondence φ^* in (2.3) is upper hemicontinuous at $\alpha = 0$ and $\varphi^*(0)$ is compact, then $\emptyset \neq \mathfrak{M}_0^* \subset \mathfrak{M}^*$.*

Remark 2.12. As with Lemma 2.8, the converse of Corollary 2.11 is easily seen to hold for our vanishing discount model. ◇

Since Corollary 2.11 is a statement about the argmin correspondence φ^*, our next objective is then to establish conditions for the feasibility correspondence φ and the objective function F under which the hypotheses of this corollary will hold. To this end, we first observe that these hypotheses will be satisfied if φ and F satisfy the hypotheses of Berge's Theorem. That is, if φ is continuous (i.e., upper hemicontinuous and lower hemicontinuous) on the entire interval $[0,1]$, φ has nonempty compact values, and the objective function F is real-valued and continuous (i.e., upper semicontinuous and lower semicontinuous) on $\mathrm{Gr}(\varphi)$, then Berge's Theorem ensures that φ^* will be upper hemicontinuous on $[0,1]$ (and, hence, φ^* will be upper hemicontinuous at $\alpha = 0$) and φ^* will have compact values (and, hence, $\varphi^*(0)$ will be compact). However, we would naturally prefer to have weaker requirements for φ and F than those stipulated in Berge's Theorem; and, indeed, we should expect that these conditions can be weakened since the conclusions of Berge's Theorem are more powerful than is necessary for our purposes. What is somewhat surprising is that these hypotheses cannot be weakened as much as one might expect, as the following result demonstrates.

Lemma 2.13. *Let the feasibility correspondence φ be as in (2.1), and let the objective function F be as in (2.2). If φ is continuous at $\alpha = 0$, φ has nonempty compact values, and F is continuous, then φ^* is upper hemicontinuous at $\alpha = 0$ and φ^* has nonempty compact values (and, hence, $\varphi^*(0)$ is compact).*

Proof. Since φ has compact values and F is lower semicontinuous, Theorem B.13 tells us that, for each $\alpha \in [0,1]$, there exists a minimizer in the compact set $\varphi^*(\alpha)$. So, let $\{(\alpha_n, \hat{\mu}^*_{\alpha_n}) : n \in \mathbb{N}\}$ be a sequence in $\mathrm{Gr}(\varphi^*)$ with $\alpha_n \to 0$ (and, as usual, assume that $\alpha_n > 0$, for each $n \in \mathbb{N}$). Since φ is upper hemicontinuous at $\alpha = 0$, there is a $\rho \in \varphi(0)$ with $\hat{\mu}^*_{\alpha_n} \Rightarrow \rho$. To then show that ρ is a minimizer (i.e., $\rho \in \varphi^*(0)$), we need to show that $F(0,\rho) \leq F^*(0)$ (since the inequality $F^*(0) \leq F(0,\rho)$ is immediate). Since F is lower semicontinuous, we have

$$F(0,\rho) \leq \liminf_{n\to\infty} F(\alpha_n, \hat{\mu}^*_{\alpha_n}) = \liminf_{n\to\infty} F^*(\alpha_n).$$

Then, since φ is lower hemicontinuous at $\alpha = 0$ and F is upper semicontinuous (on $\{0\} \times \varphi(0)$), Corollary B.15 tells us that F^* is upper semicontinuous at $\alpha = 0$; i.e.,

$$\liminf_{n\to\infty} F^*(\alpha_n) \leq \limsup_{n\to\infty} F^*(\alpha_n) \leq F^*(0),$$

from which the desired result follows. □

Let us now discuss in detail each of the hypotheses in the lemma above.

We have shown that the conditions on the vanishing discount model $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ ensure that the feasibility correspondence φ is upper hemicontinuous at $\alpha = 0$ and that $\varphi(0)$ is compact; and, as per Remark 2.7, we know that φ has nonempty values. What then remains to be discussed (as pertains to φ) is the lower hemicontinuity of this correspondence at $\alpha = 0$ and the compactness of $\varphi(\alpha)$ for each $\alpha \in (0, 1]$.

As the proof of Lemma 2.13 demonstrates, the lower hemicontinuity of φ (at $\alpha = 0$) is only required here to obtain the *upper* semicontinuity of the value function F^* at $\alpha = 0$. Nevertheless, we will see that this lower hemicontinuity condition plays an indispensable role in our desired results; and moreover, unlike upper hemicontinuity, this is *not* a property that φ automatically obtains from the conditions on our vanishing discount model.

As for the compactness of $\varphi(\alpha)$ for each $\alpha \in (0, 1]$: Here, it will be sufficient to show (by Prohorov's Theorem) that, given a fixed $\alpha \in (0, 1]$, the collection $\varphi(\alpha) = \hat{\mathfrak{M}}_\alpha$ of probability measures is tight. Fortunately, this is the case. Indeed, since each $\hat{\mu}_\alpha \in \hat{\mathfrak{M}}_\alpha$ necessarily satisfies

$$-\infty < -\kappa_1 \leq \langle c_1, \hat{\mu}_\alpha \rangle \leq \alpha\theta_\alpha < +\infty,$$

and since c_1 is inf-compact, an argument similar enough to the one given in the proof of Proposition 2.9 then shows that $\varphi(\alpha) = \hat{\mathfrak{M}}_\alpha$ is, indeed, compact. We summarize this fact as the following proposition.

Proposition 2.14. *Suppose $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ is a constrained control model for the vanishing discount method. Then, for each fixed $\alpha \in (0, 1]$, the collection $\hat{\mathfrak{M}}_\alpha$ is a compact subset of $\mathcal{P}(E \times G)$; i.e., $\varphi(\alpha)$ is compact for each $\alpha \in (0, 1]$, where φ is the feasibility correspondence in (2.1).*

Lastly, we discuss the continuity of the objective function F. By the manner in which we have defined φ in (2.1) and F in (2.2), we see that there are, in a sense, two cases to consider:

(a) F is continuous at each $(0, \rho) \in \{0\} \times \varphi(0)$; and

(b) F is continuous at each $(\alpha, \hat{\mu}_\alpha) \notin \{0\} \times \varphi(0)$.

In order to discuss the continuity of such a function, it will be useful for us to specify precisely what notion of convergence we have in mind for sequences in $\mathrm{Gr}(\varphi) \subset [0, 1] \times \mathcal{P}(E \times G)$. For this purpose, choosing the so-called *taxicab metric* will be sufficient for our needs. That is, we write $(\alpha_n, \mu_{\alpha_n}) \to (\alpha, \mu_\alpha)$ if and only if, given $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that

$$|\alpha_n - \alpha| + \pi(\mu_{\alpha_n}, \mu_\alpha) < \epsilon, \qquad \forall n \geq N,$$

where $|\cdot|$ denotes the usual Euclidean metric on $[0, 1] \subset \mathbb{R}$ and π denotes the Prohorov metric on $\mathcal{P}(E \times G)$. Thus, under the taxicab metric, it is clear that $(\alpha_n, \mu_{\alpha_n}) \to (\alpha, \mu_\alpha)$ if and only if $\alpha_n \to \alpha$ and $\mu_{\alpha_n} \Rightarrow \mu_\alpha$ (again recalling part (b) of the Portmanteau Theorem). We then see that F is continuous at (α, μ_α) if, for each sequence $\{(\alpha_n, \mu_{\alpha_n}) : n \in \mathbb{N}\} \subset \mathrm{Gr}(\varphi)$ with $(\alpha_n, \mu_{\alpha_n}) \to (\alpha, \mu_\alpha)$, we have

$$F(\alpha, \mu_\alpha) = \langle c, \mu_\alpha \rangle = \lim_{n\to\infty} \langle c, \mu_{\alpha_n} \rangle = \lim_{n\to\infty} F(\alpha_n, \mu_{\alpha_n}).$$

Now, since c is lower semicontinuous and bounded below, the weak convergence $\mu_{\alpha_n} \Rightarrow \mu$ only guarantees that

$$\langle c, \rho \rangle \leq \liminf_{n\to\infty} \langle c, \hat{\mu}_{\alpha_n} \rangle.$$

In other words, the conditions on our model are sufficient to ensure that F is only *lower semicontinuous* on $\mathrm{Gr}(\varphi)$. Nevertheless, this observation is important enough to state as a proposition.

Proposition 2.15. *Suppose $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ is a constrained control model for the vanishing discount method. If φ and F are as in (2.1) and (2.2), then F is lower semicontinuous.*

Now, a careful reading of the proof of Lemma 2.13 illustrates that the hypotheses of this lemma could, in fact, be weakened to require only that F be upper semicontinuous on $\{0\} \times \varphi(0)$ and lower semicontinuous everywhere else on $\mathrm{Gr}(\varphi)$. Such a modification, however, is unlikely to yield any sort of practical utility. In any case, as with the lower hemicontinuity of the feasibility correspondence φ at $\alpha = 0$, the upper semicontinuity of the objective function F on $\{0\} \times \varphi(0)$ will have to be verified on an *ad hoc* basis when application of the vanishing discount method is desired.

IV.2.3 Analysis of Conditions for (V3)

We now turn our attention to the last, and perhaps most important, of the VDM Relations. Recall that

$$J_0^* := \lim_{\alpha \downarrow 0} \hat{J}_\alpha^* := \lim_{\alpha \downarrow 0} \Big(\inf\{\langle c, \hat{\mu}_\alpha \rangle : \hat{\mu}_\alpha \in \hat{\mathfrak{M}}_\alpha\} \Big) \qquad \text{and} \qquad J^* = \inf\{\langle c, \rho \rangle : \rho \in \mathfrak{M}\};$$

and, hence, the relation $J_0^* = J^*$ is a statement about the optimal (i.e., minimal) *values* of the linear programs in our construction.

It should be apparent from our definitions that the relation $J_0^* = J^*$ will hold if and only if the value function F^* in (2.4) is continuous (and finite) at $\alpha = 0$. However, it will prove useful to be perhaps a bit more explicit about this statement, as in the next two lemmas.

Lemma 2.16. *If the value function F^* in (2.4) is upper semicontinuous at $\alpha = 0$, then $J_0^* \le J^*$.*

Proof. Let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence in $(0,1]$ with $\alpha_n \to 0$. If F^* is upper semicontinuous at $\alpha = 0$, then

$$J_0^* \le \limsup_{n\to\infty} \hat{J}_{\alpha_n}^* = \limsup_{n\to\infty} F^*(\alpha_n) \le F^*(0) = J^*,$$

as desired. □

For the next lemma, we simplify things a bit by adding the assumption that $F^*(0) < +\infty$. This is justified, however, since the conditions on our model (see Remark 2.7) guarantee that there is at least one $\rho \in \varphi(0)$ with $F(0,\rho) < +\infty$.

Lemma 2.17. *If the value function F^* in (2.4) is lower semicontinuous at $\alpha = 0$ and $F^*(0) < +\infty$, then $J_0^* \ge J^*$.*

Proof. Let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence in $(0,1]$ with $\alpha_n \to 0$. If F^* is lower semicontinuous at $\alpha = 0$ and $F^*(0) < +\infty$, then

$$J_0^* \ge \liminf_{n\to\infty} \hat{J}_{\alpha_n}^* = \liminf_{n\to\infty} F^*(\alpha_n) \ge F^*(0) = J^*,$$

as desired. □

Remark 2.18. If, for example, one is only interested in obtaining a lower bound on the minimum value of the long-term average linear program via the vanishing discount method, Lemma 2.16 says that only upper semicontinuity of F^* at $\alpha = 0$ is required. On the other hand, if an upper bound is desired, one can look to Lemma 2.17. ◇

As in the previous analyses, we now wish to identify sufficient conditions for the feasibility correspondence φ and the objective function F that will allow for the hypotheses of Lemma 2.16 and Lemma 2.17 to be fulfilled. Again, we remark that the continuity of F^* on the entire interval $[0,1]$ is among the conclusions of Berge's Theorem; but, as above, we should not require the full strength Berge's Theorem and, hence, we should not require the full strength of its hypotheses. Accordingly, the following two lemmas, in analogy with the preceding two lemmas, state sufficient conditions for the value function F^* to be upper semicontinuous at $\alpha = 0$ and lower semicontinuous at $\alpha = 0$, respectively.

Lemma 2.19. *If the feasibility correspondence φ is lower hemicontinuous at $\alpha = 0$, $\varphi(0)$ is compact, and the objective function F is continuous, then the value function F^* is upper semicontinuous at $\alpha = 0$.*

Proof. Since F is lower semicontinuous, it attains its minimum on the compact set $\{0\} \times \varphi(0)$; i.e., there is a minimizer $\rho^* \in \varphi^*(0) \subset \varphi(0)$. So, let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence in $(0, 1]$ with $\alpha_n \to 0$. Since φ is lower hemicontinuous at $\alpha = 0$, there is some subsequence of $\{\alpha_n : n \in \mathbb{N}\}$ (which, for notational reasons, we will take to be the sequence itself) and a sequence $\{\hat{\mu}_{\alpha_n} : n \in \mathbb{N}\}$ with $\hat{\mu}_{\alpha_n} \in \varphi(\alpha_n)$ for each $n \in \mathbb{N}$ and $\hat{\mu}_{\alpha_n} \Rightarrow \rho^*$. By minimality, we must have $F^*(\alpha_n) \leq F(\alpha_n, \hat{\mu}_{\alpha_n})$ for every $n \in \mathbb{N}$; and since F is upper semicontinuous, we have

$$\limsup_{n \to \infty} F^*(\alpha_n) \leq \limsup_{n \to \infty} F(\hat{\mu}_{\alpha_n}, \alpha_n) \leq F(0, \rho^*) = F^*(0),$$

which shows that F^* is upper semicontinuous at $\alpha = 0$. □

Lemma 2.20. *If the feasibility correspondence φ is upper hemicontinuous at $\alpha = 0$, φ has nonempty compact values, and the objective function F is lower semicontinuous, then F^* is lower semicontinuous at $\alpha = 0$.*

Proof. Let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence in $(0, 1]$ satisfying $\alpha_n \to 0$. Since $\varphi(\alpha_n)$ is compact for each $n \in \mathbb{N}$, and F is lower semicontinuous, F attains a minimum on each compact section $\{\alpha_n\} \times \varphi(\alpha_n)$. We can therefore construct a sequence $\{(\alpha_n, \hat{\mu}^*_{\alpha_n}) : n \in \mathbb{N}\}$ in $\mathrm{Gr}(\varphi)$ where $\hat{\mu}^*_{\alpha_n} \in \hat{\mathfrak{M}}^*_{\alpha_n}$ for each $n \in \mathbb{N}$; and since φ is upper hemicontinuous at $\alpha = 0$, there is a $\rho \in \varphi(0)$ with $\hat{\mu}^*_{\alpha_n} \Rightarrow \rho$. We then have $F(\alpha_n, \hat{\mu}^*_{\alpha_n}) = F^*(\alpha_n)$ for each n, and $F^*(0) \leq F(0, \rho)$. So, by the lower semicontinuity of F, we obtain

$$F^*(0) \leq F(0,\rho) \leq \liminf_{n\to\infty} F(\alpha_n, \hat{\mu}^*_{\alpha_n}) = \liminf_{n\to\infty} F^*(\alpha_n),$$

which shows that F^* is lower semicontinuous at $\alpha = 0$. □

Remark 2.21. Observe that all of the hypotheses of Lemma 2.20 are conditions that are automatically satisfied by our model. ⋄

We can now state, as a consequence of the preceding two lemmas, sufficient conditions for (V3) to hold:

Theorem 2.22. *If the feasibility correspondence φ is continuous at $\alpha = 0$, φ has nonempty compact values, and F is continuous, then F^* is continuous at $\alpha = 0$. Hence, $J_0^* = J^*$.*

Remark 2.23. Notice that the hypotheses of this theorem are exactly the same as those in Lemma 2.13. Again, we can weaken these hypotheses slightly by only requiring that F be lower semicontinuous on $\mathrm{Gr}(\varphi)$ and upper semicontinuous on $\{0\} \times \varphi(0)$. ⋄

IV.2.4 Verification of Main Results

We are now able to present the proofs of Theorem 2.5 and Theorem 2.6 as immediate consequences of the analyses conducted in the preceding subsections.

Proof of Theorem 2.5. Since φ is continuous at $\alpha = 0$ and φ has nonempty compact values, φ is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact; and so, (V1) holds by Lemma 2.8. Since φ is continuous at $\alpha = 0$, φ has nonempty compact values, and F is continuous, the argmin correspondence φ^* is upper hemicontinuous at $\alpha = 0$ and $\varphi^*(0)$ is compact (by Lemma 2.13); and so, (V2) holds by Corollary 2.11. Since φ is continuous at $\alpha = 0$, φ has nonempty compact values, and F is continuous, the value function F^* is continuous at $\alpha = 0$ (by Theorem 2.22); and so, (V3) holds by Theorem 2.22. □

To then prove Theorem 2.6, we need only appeal to the above proof and the additional results that pertain specifically to our vanishing discount model.

Proof of Theorem 2.6. Let $\{E \times G, L, c, c_1, \theta, \{\theta_\alpha\}\}$ be a constrained control model for the vanishing discount method, and recall the following facts established above:

- Remark 2.7 showed that φ has nonempty values.
- Proposition 2.9 showed that φ is upper hemicontinuous at $\alpha = 0$ and $\varphi(0)$ is compact.
- Proposition 2.14 showed that $\varphi(\alpha)$ is compact for each $\alpha \in (0, 1]$.
- Proposition 2.15 showed that F is lower semicontinuous.

Thus, these facts together with the argument given in the proof of Theorem 2.5 yield the desired results. □

The next chapter explores a number of examples that provide insight into not only *when* this formulation of the vanishing discount method may be applied, but also *how* this method works.

EXAMPLES

In this chapter, we present examples that illustrate the results discussed in this dissertation. We first consider an adaptation of an example appearing in Hernández-Lerma and Prieto-Rumeau (2010) in which our constrained control model is equipped with a very basic discrete state and control spaces. The rudimentary structure of this model allows us to express the feasible sets for our linear programs as subsets of the unit simplex in $\mathbb{R}^2$. We use this geometric illustration to visualize how the feasible sets (and sets of optimizers) for the α-normalized linear programs behave as $\alpha \downarrow 0$. Using this basic model, we present two examples. The first example serves as a useful demonstration of our results, for we are able to verify directly that the hypotheses and the conclusions of Theorem IV.2.5 are satisfied. For the second example, we show that, by simply changing the constraint constants slightly, the hypotheses of Theorem IV.2.5 may fail to hold. In particular, we will see that the lower hemicontinuity of the feasibility correspondence φ may fail to be lower hemicontinuous at $\alpha = 0$ if the constraint constant is not appropriately chosen.

We then conclude with an application of our results to a controlled diffusion problem. Here, in the spirit of Chapter 11 in Øksendal (2003), we consider an *Itô process* in which the control is given in the form of the drift coefficient. As a means of comparison, we omit the budget constraint to show that our results may remain applicable in the "unconstrained" setting, as well. In this particular example, we are able to characterize the *densities* of the feasible measures for the long-term average and α-normalized linear programs in terms of the respective feedback control functions. We find that each α-normalized problem proves rather difficult to solve explicitly, so we further our analysis by solving the long-term average problem using dynamic programming methods and applying the vanishing discount method to obtain approximate solutions to the α-normalized (and, hence, α-discounted) problems.

V.1 Example: A Discrete State Space

Let $X := \{X_t : t \in \mathbb{R}^+\}$ be a controlled continuous-time Markov chain with state space $E = \{0, 1\}$ and control space $G = \{0, 1, 2, 3\}$, where $G(0) = \{0\}$ and $G(1) = \{1, 2, 3\}$. The instantaneous transition rates for X are given by

$$q_{00}(0) = -1, \qquad q_{11}(1) = -2, \qquad \text{and} \qquad q_{11}(2) = q_{11}(3) = -4,$$

where $q_{ij}(u)$ denotes the transition rate from state i to state j under control u. For this example, we consider both the cost rate function c with values

$$c(0, 0) = -1 \qquad \text{and} \qquad c(1, 1) = c(1, 2) = c(1, 3) = 0;$$

and the budget rate function c_1 with the values

$$c_1(0, 0) = 0, \qquad c_1(1, 1) = 3, \qquad c_1(1, 2) = 5, \qquad \text{and} \qquad c_1(1, 3) = 10.$$

Note that c and c_1 trivially satisfy conditions (C1)-(C3) since each has finite range and the product space $E \times G$ is finite (and hence we can simply equip $E \times G$ with the discrete topology). Now, for such a controlled process X, we assume that the controls are chosen according to a family of policies that we identify with the unit simplex $Z \subset \mathbb{R}^2$. That is,

$$Z := \{\mathbf{z} = (z_1, z_2) \in \mathbb{R}^2 : z_1 \geq 0,\ z_2 \geq 0,\ z_1 + z_2 \leq 1\},$$

where

$$\eta(1, \{1\}) = z_1, \qquad \eta(1, \{2\}) = z_2, \qquad \text{and} \qquad \eta(1, \{3\}) = z_3 := 1 - z_1 - z_2,$$

noting that $\eta(0,\{0\}) = 1$ since $G(0) = \{0\}$. Recall that $\eta(X_t, \cdot) = \Lambda_t(\cdot)$, $t \in \mathbb{R}^+$, represents a control given in feedback form.

To analyze the dynamics of each feasible $(X, \Lambda) = (X, \eta(X, \cdot)) \in \mathscr{M}$, we use the given transition rates $q_{ij}(u)$ to evaluate $Lf(i,u)$ for each $(i,u) \in \mathcal{U}$. A routine computation yields the following: For each $f \in \mathscr{D}(L)$,

$$
\begin{aligned}
Lf(0,0) &= q_{01}(0)f(1) + q_{00}(0)f(0) = f(1) - f(0),\\
Lf(1,1) &= q_{10}(1)f(0) + q_{11}(1)f(1) = 2[f(0) - f(1)],\\
Lf(1,2) &= q_{10}(2)f(0) + q_{11}(2)f(1) = 4[f(0) - f(1)], \text{ and}\\
Lf(1,3) &= q_{10}(3)f(0) + q_{11}(3)f(1) = 4[f(0) - f(1)].
\end{aligned}
$$

Note that L trivially satisfies conditions (D1)-(D6). Indeed, since E and G are finite (and, hence, compact) sets, conditions (D1)-(D5) are trivially satisfied; and since Lf is bounded for every $f \in \mathscr{D}(L)$, condition (D6) is trivially satisfied (e.g., ψ can be chosen to be a constant function).

V.1.1 Example 1A: The Long-Term Average Linear Program

The long-term average linear program in this setting is as follows:

$$
\text{Minimize} \quad \langle c, \rho \rangle \qquad \text{subject to} \quad
\begin{cases}
\langle Lf, \rho \rangle = 0, \quad \forall f \in \mathscr{D}(L); & \text{(MA1)}\\
\langle 1, \rho \rangle = 1; & \text{(MA2)}\\
\rho \in \mathcal{P}(E \times G); & \text{(MA3)}\\
\langle c_1, \rho \rangle \le \theta. & \text{(CA1)}
\end{cases}
$$

The objective function is then given by

$$
\langle c, \rho \rangle = \sum_{(i,u)\in\mathcal{U}} c(i,u)\,\rho(i,u) = -\rho(0,0),
$$

and the budget constraint (CA1) is

$$\langle c_1, \rho \rangle = \sum_{(i,u)\in\mathcal{U}} c_1(i,u)\,\rho(i,u) = 3\rho(1,1) + 5\rho(1,2) + 10\rho(1,3) \le \theta,$$

noting the abuse of notation $\rho(i,u) = \rho(\{(i,u)\})$.

For this example, we will consider the constraint constant $\theta = \frac{5}{4}$.

It can be shown that each $\rho \in \mathfrak{M}$ can be expressed in terms of a $\mathbf{z} = (z_1, z_2) \in Z$ via the parameterization

$$\rho(i,u) = \frac{4 - 2z_1}{5 - 2z_1} I_{\{(0,0)\}}(i,u) + \sum_{u=1}^{3} \frac{z_u}{5 - 2z_1} I_{\{(1,u)\}}(i,u), \qquad \forall (i,u) \in \mathcal{U},$$

recalling that $z_3 := 1 - z_1 - z_3$. The budget constraint (CA1) then implies that the feasible set $\mathfrak{M}$ for this linear program can be viewed as the *feasible region*

$$\left\{\mathbf{z} \in Z : \frac{10 - 7z_1 - 5z_2}{5 - 2z_1} \le \frac{5}{4}\right\} = \{\mathbf{z} \in Z : 18z_1 + 20z_2 \ge 15\},$$

which is the quadrilateral region in Z with vertices

$$(0, 3/4), \qquad (0,1), \qquad (1,0), \qquad (5/6, 0),$$

as depicted below:

Figure V.1: Feasible Region for the Long-Term Average LP

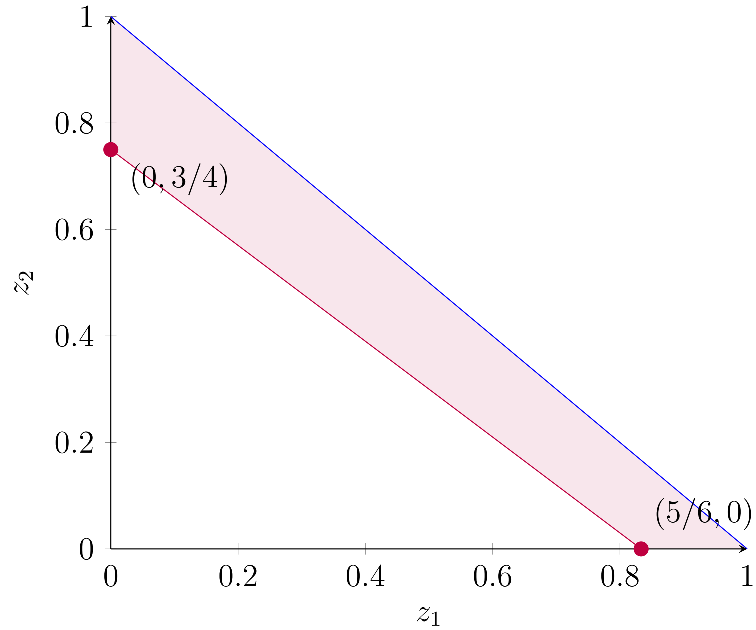

V.1.2 Example 1A: The α-Normalized Linear Program

Note that, for purposes of illustration, we will take the parameter space for our discount rate to be $(0,2]$ rather than $(0,1]$. So, let $\alpha \in (0,2]$ be given and assume that $X_0 = 0$. Recall that the α-normalized linear program in this setting is as follows:

$$\text{Minimize} \quad \langle c, \hat{\mu}_\alpha \rangle \quad \text{subject to} \quad \begin{cases} \langle L_\alpha f, \hat{\mu}_\alpha \rangle = -\alpha f(0), \quad \forall f \in \mathscr{D}(L); & \text{(MN1)} \\ \langle 1, \hat{\mu}_\alpha \rangle = 1; & \text{(MN2)} \\ \hat{\mu}_\alpha \in \mathcal{P}(E \times G); & \text{(MN3)} \\ \langle c_1, \hat{\mu}_\alpha \rangle \leq \alpha \theta_\alpha. & \text{(CN1)} \end{cases}$$

As above, the objective function is given by

$$\langle c, \hat{\mu}_\alpha \rangle = \sum_{(i,u) \in \mathcal{U}} c(i,u)\, \hat{\mu}_\alpha(i,u) = -\hat{\mu}_\alpha(0,0)$$

and the budget constraint (CA1) is

$$\langle c_1, \hat{\mu}_\alpha \rangle = \sum_{(i,u)\in\mathcal{U}} c_1(i,u)\,\hat{\mu}_\alpha(i,u) = 3\hat{\mu}_\alpha(1,1) + 5\hat{\mu}_\alpha(1,2) + 10\hat{\mu}_\alpha(1,3) \le \alpha\theta_\alpha.$$

For this example, we will consider the constraint constant $\theta_\alpha = \dfrac{5}{\alpha(\alpha+4)}$, noting that

$$\lim_{\alpha\downarrow 0} \alpha\theta_\alpha = \lim_{\alpha\downarrow 0} \frac{5}{\alpha(\alpha+4)} = \frac{5}{4} = \theta.$$

Hence, $\{\theta_\alpha : \alpha \in (0,2]\}$ satisfies condition (T1). An appropriate parameterization of each $\hat{\mu}_\alpha \in \hat{\mathfrak{M}}_\alpha$ in terms of a $\mathbf{z} \in Z$ then shows that the feasible set $\hat{\mathfrak{M}}_\alpha$ can be viewed as the feasible region

$$\left\{\mathbf{z} \in Z : \frac{10 - 7z_1 - 5z_2}{5 + \alpha - 2z_1} \le \frac{5}{\alpha+4}\right\} = \{\mathbf{z} \in Z : (18+7\alpha)z_1 + (20+5\alpha)z_2 \ge 15 + 5\alpha\},$$

which is the quadrilateral region in Z with vertices

$$\left(0, \frac{\alpha+3}{\alpha+4}\right), \qquad (0,1), \qquad (1,0), \qquad \left(\frac{5\alpha+15}{7\alpha+18}, 0\right)$$

as depicted in the plot below with $\alpha = 1.8$.

Figure V.2: Feasible Region for the α-Normalized LP ($\alpha = 1.8$)

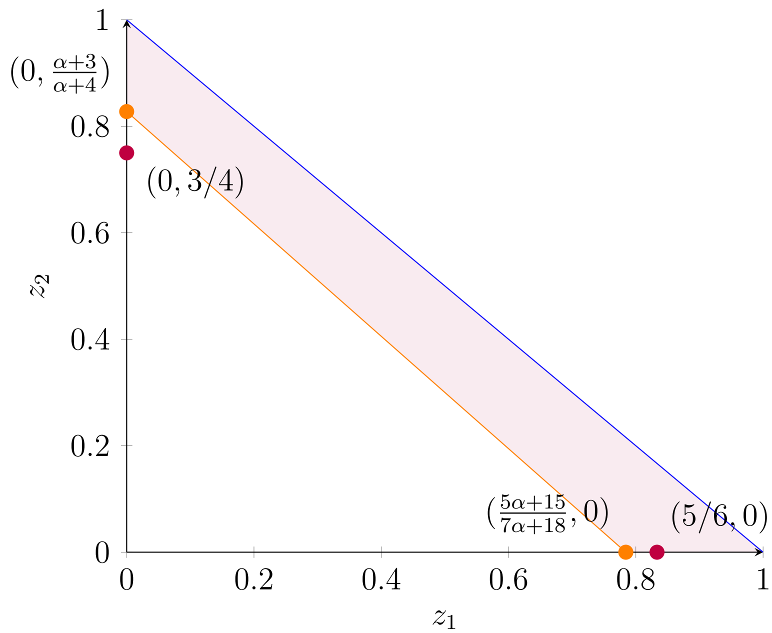

V.1.3 Example 1A: Correspondences and Optimality

The feasibility correspondence $\varphi : [0,2] \twoheadrightarrow Z$ for this example is

$$\varphi(\alpha) = \begin{cases} \{\mathbf{z} \in Z : 18z_1 + 20z_2 \geq 15\} & \text{if } \alpha = 0, \\ \{\mathbf{z} \in Z : (18+7\alpha)z_1 + (20+5\alpha)z_2 \geq 15 + 5\alpha\} & \text{if } \alpha \in (0,2]; \end{cases}$$

and the objective function $F : \mathrm{Gr}(\varphi) \to \mathbb{R}$ is

$$F(\alpha, \mathbf{z}) = \begin{cases} -(4 - z_1)/(5 - z_1) & \text{if } \alpha = 0, \\ -(4 + \alpha - 2z_1)/(5 + \alpha - 2z_1) & \text{if } \alpha \in (0,2]; \end{cases}$$

Now, the parameterization of each measure of interest in terms of a point $(z_1, z_2) \in Z \subset \mathbb{R}^2$ provides us with a convenient "visual" way of verifying that

(H1) φ is continuous at $\alpha = 0$,

(H2) φ has nonempty compact values, and

(H3) F is continuous.

Remark 1.1. It can be shown that the parameterization $\rho \mapsto \mathbf{z}$ defines an isometry between the feasible set $\mathfrak{M}$ and its corresponding feasible region; and that, for each $\alpha \in (0,1]$, the parameterization $\hat{\mu}_\alpha \mapsto \mathbf{z}$ defines an isometry between $\hat{\mathfrak{M}}_\alpha$ and its corresponding feasible region. Thus, weak convergence of measures is equivalent to convergence of points (under the usual Euclidean metric) in Z. ◇

Indeed: Let $\{(\alpha_n, \mathbf{z}_n) : n \in \mathbb{N}\}$ be an arbitrary sequence in $\mathrm{Gr}(\varphi) \subset [0,2] \times Z$ with $\alpha_n \to 0$. Then each $(\alpha_n, \mathbf{z}_n) = \left(\alpha_n, (z_1^{(n)}, z_2^{(n)})\right)$ in this sequence satisfies

$$(18 + 7\alpha_n)z_1^{(n)} + (20 + 5\alpha_n)z_2^{(n)} \geq 15 + 5\alpha_n.$$

So, we observe that

$$18z_1^{(0)} + 20z_2^{(0)} = \lim_{n\to\infty} \left((18 + 7\alpha_n)z_1^{(n)} + (20 + 5\alpha_n)z_2^{(n)}\right) \geq \lim_{n\to\infty} (15 + 5\alpha_n) = 15,$$

for some $\mathbf{z}_0 = (z_1^{(0)}, z_2^{(0)}) \in Z$, noting that Z is a compact subset of $\mathbb{R}^2$. Clearly, we have $\mathbf{z}_0 \in \varphi(0)$; and so, φ is upper hemicontinuous at $\alpha = 0$. Now, let $\{\alpha_n : n \in \mathbb{N}\}$ be a sequence in $[0,2]$ with $\alpha_n \to 0$, and let $\mathbf{z}_0 = (z_1^{(0)}, z_2^{(0)}) \in \varphi(0)$. Since $\alpha_n \to 0$, we can choose a decreasing subsequence $\{\alpha_{n_k} : k \in \mathbb{N}\}$ of $\{\alpha_n : n \in \mathbb{N}\}$ with $\alpha_{n_k} \to 0$. We now need to find a sequence $\{\mathbf{z}_k : k \in \mathbb{N}\}$ with $\mathbf{z}_k \in \varphi(\alpha_{n_k})$ for every $k \in \mathbb{N}$ and $\mathbf{z}_k \to \mathbf{z}_0$. So, observe that, if $z_1^{(0)} \geq 0.5$, we can choose the constant sequence $\mathbf{z}_k = \mathbf{z}_0$ since $\mathbf{z}_0 \in \varphi(\alpha_{n_k})$ for every $k \in \mathbb{N}$; and if $z_1^{(0)} < 0.5$, then we can define our sequence by

$$\mathbf{z}_k = \left(z_1^{(0)}, \frac{15 + 5\alpha_{n_k} - (18 + 7\alpha_{n_k})z_1^{(0)}}{20 + 5\alpha_{n_k}}\right), \qquad \forall k \in \mathbb{N}.$$

Thus, φ is lower semicontinuous at $\alpha = 0$.

Since, for each $\alpha \in [0, 2]$, the set $\varphi(\alpha)$ is easily seen to be a nonempty, closed, and bounded subset of $\mathbb{R}^2$ (in fact, a closed polygon), it is clear that φ has nonempty compact values.

Finally, the continuity of F is also clear since F is a rational function (of only the two real variables α and z_1) whose denominator is nonzero for every $(\alpha, \mathbf{z}) \in \mathrm{Gr}(\varphi)$.

Thus, having satisfied hypotheses (H1)-(H3), we should obtain the desired conclusions

(V1) $\emptyset \neq \mathfrak{M}_0 \subset \mathfrak{M}$,

(V2) $\emptyset \neq \mathfrak{M}_0^* \subset \mathfrak{M}^*$, and

(V3) $J_0^* = J^*$.

Indeed, since

$$18z_1 + 20z_2 = \lim_{\alpha \downarrow 0} \left((18 + 7\alpha)z_1 + (20 + 5\alpha)z_2\right) \geq \lim_{\alpha \downarrow 0} (15 + 5\alpha) = 15,$$

it follows that $\emptyset \neq \mathfrak{M}_0 = \mathfrak{M}$.

Now, when $\alpha = 0$, the objective function

$$F(0, \mathbf{z}) = -\frac{4 - z_1}{5 - z_1}$$

obtains it minimum at any $\mathbf{z} = (z_1, z_2) \in \varphi(0)$ with $z_1 = 0$. So,

$$\mathfrak{M}^* = \varphi^*(0) = \{\mathbf{z} \in Z : z_1 = 0,\ z_2 \geq 0.75\}$$

and

$$J^* = F^*(0) = -\frac{4 - 0}{5 - 0} = -\frac{4}{5}.$$

On the other hand, when $\alpha \in (0, 2]$, we have

$$F(\alpha, \mathbf{z}) = -\frac{4 + \alpha - z_1}{5 + \alpha - z_1},$$

which obtains it minimum at any $\mathbf{z} = (z_1, z_2) \in \varphi(\alpha)$ with $z_1 = 0$. So,

$$\hat{\mathfrak{M}}^*_\alpha = \varphi^*(\alpha) = \left\{\mathbf{z} \in Z : z_1 = 0,\ z_2 \geq \frac{3 + \alpha}{4 + \alpha}\right\}$$

and

$$\hat{J}^*_\alpha = F^*(\alpha) = -\frac{4 + \alpha - 0}{5 + \alpha - 0} = -\frac{4 + \alpha}{5 + \alpha}.$$

It is then not too difficult to see that

$$\emptyset \neq \mathfrak{M}^*_0 = \mathfrak{M}^*$$

and

$$J^*_0 = \lim_{\alpha \downarrow 0} \hat{J}^*_\alpha = \lim_{\alpha \downarrow 0} F^*(\alpha) = -\lim_{\alpha \downarrow 0} \frac{4 + \alpha}{5 + \alpha} = -\frac{4}{5} = F^*(0) = J^*,$$

as desired.

To now demonstrate the sensitivity of the hypotheses we have presented, we will consider a similar example in which all that is changed is the constraint constants.

V.1.4 Example 1B: The Long-Term Average Linear Program

We consider the same long-term average linear program as above but with the constraint constant $\theta = 1$.

Recalling the parameterization obtained in the previous example, the feasible set $\mathfrak{M}$ for this problem can now be viewed as the region

$$\left\{\mathbf{z} \in Z : \frac{10 - 7z_1 - 5z_2}{5 - 2z_1} \leq 1\right\} = \{\mathbf{z} \in Z : z_1 + z_2 \geq 1\}$$

which is the line segment in $\mathbb{R}^2$ with vertices $(0,1)$ and $(1,0)$:

Figure V.3: Feasible Region for the Long-Term Average LP

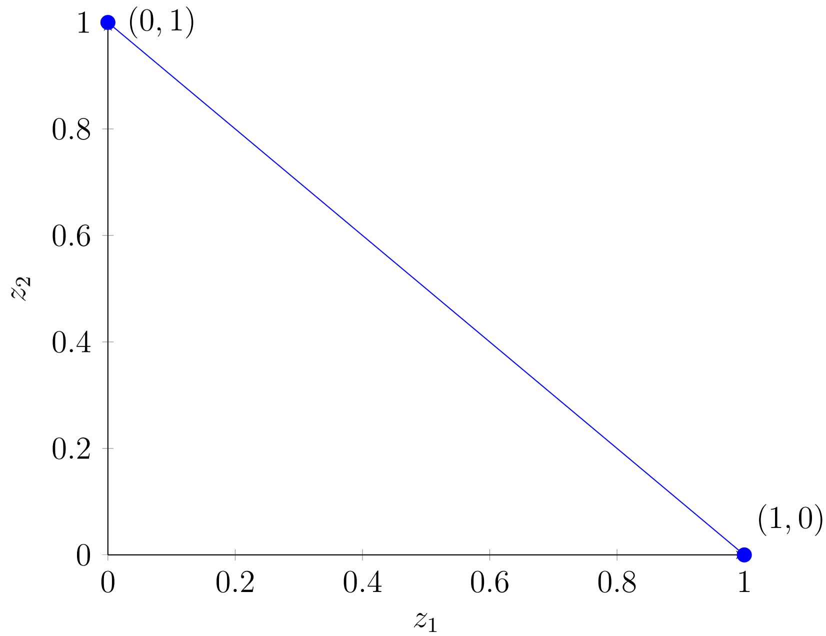

V.1.5 Example 1B: The α-Normalized Linear Program

Let $\alpha \in (0,2]$ be given and again assume that $X_0 = 0$.

We now take our constraint constant to be $\theta_\alpha = \dfrac{4}{\alpha(\alpha+4)}$, observing that

$$\lim_{\alpha \downarrow 0} \alpha\theta_\alpha = \lim_{\alpha \downarrow 0} \frac{4}{\alpha+4} = 1.$$

Hence, the family $\{\theta_\alpha : \alpha \in (0,2]\}$ again satisfies condition (T1). The feasible set $\hat{\mathfrak{M}}_\alpha$ for this linear program can be identified with the region

$$\left\{\mathbf{z} \in Z : \frac{10 - 7z_1 - 5z_2}{5 + \alpha - 2z_1} \leq \frac{4}{\alpha+4}\right\} = \{\mathbf{z} \in Z : 6\alpha + 20 \leq (7\alpha+20)z_1 + (5\alpha+20)z_2\},$$

which is the triangular region in Z with the vertices

$$(0.5, 0.5), \qquad (1, 0), \qquad \left(\frac{6\alpha + 20}{7\alpha + 20}, 0\right),$$

depicted as the shaded region below:

Figure V.4: Feasible Region for the α-Normalized LP ($\alpha = 1.8$)

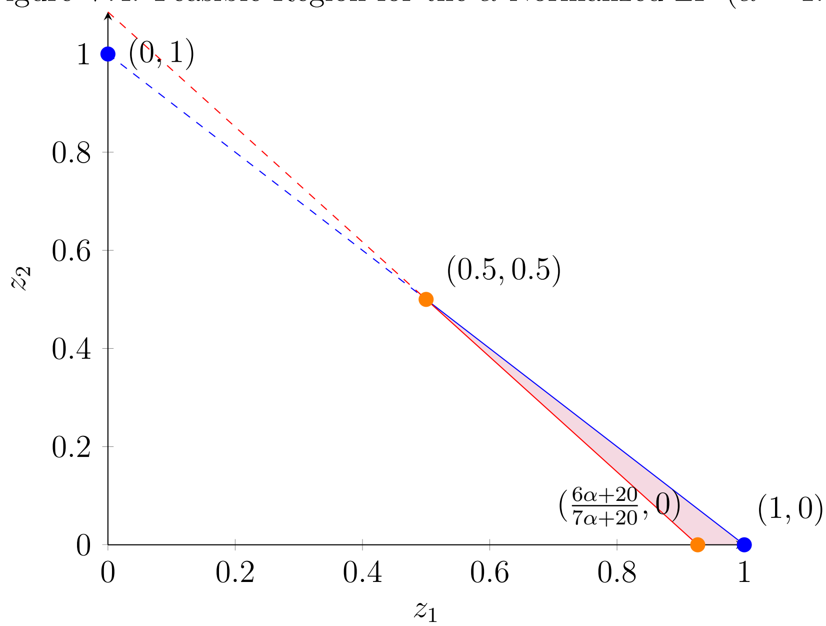

V.1.6 Example 1B: Correspondences and Optimality

The feasibility correspondence $\varphi : [0, 2] \twoheadrightarrow Z$ for this example is

$$\varphi(\alpha) = \begin{cases} \{\mathbf{z} \in Z : z_1 + z_2 \geq 1\} & \text{if } \alpha = 0, \\ \{\mathbf{z} \in Z : 6\alpha + 20 \leq (7\alpha + 20)z_1 + (5\alpha + 20)z_2\} & \text{if } \alpha \in (0, 2]; \end{cases}$$

and the objective function $F : \mathrm{Gr}(\varphi) \to \mathbb{R}$ is (as in Example 1A)

$$F(\alpha, \mathbf{z}) = \begin{cases} -(4 - z_1)/(5 - z_1) & \text{if } \alpha = 0, \\ -(4 + \alpha - 2z_1)/(5 + \alpha - 2z_1) & \text{if } \alpha \in (0, 2]; \end{cases}$$

Again, we consider the following hypotheses:

(H1) φ is continuous at $\alpha = 0$,

(H2) φ has nonempty compact values, and

(H3) F is continuous.

In this example, we see that

$$\varphi(\alpha_1) \supset \varphi(\alpha_2) \supset \{\mathbf{z} \in Z : z_1 + z_2 = 1,\ 0.5 \leq z_1 \leq 1\}$$

whenever $2 \geq \alpha_1 \geq \alpha_2 > 0$, which confirms the upper hemicontinuity of φ at $\alpha = 0$. However, our illustration makes clear that φ is not lower hemicontinuous at $\alpha = 0$. Indeed: If we consider the point $(0, 1) \in \varphi(0)$ and the set $B := \{\mathbf{z} \in Z : |\mathbf{z} - (0, 1)| < 0.5\}$, it is clear that $B \cap \varphi(\alpha) = \emptyset$ for every $\alpha \in (0, 2]$. Hence, φ cannot be lower hemicontinuous at $\alpha = 0$.

As in the previous example, (H2) is clearly satisfied. Regarding the continuity of F: The definition of F remains the same as in Example 1A, but its domain is quite different. Indeed: When $\alpha = 0$, the domain of F can be viewed as the diagonal line segment depicted in Figure V.3; and, for each $\alpha \in (0, 2]$, the domain of F can be viewed as the triangular region depicted in Figure V.4, where this region shrinks down to the diagonal line segment with vertices $(0.5, 0.5)$ and $(1, 0)$ as $\alpha \downarrow 0$. Nevertheless, since F is a rational function with no possibility of division by 0, it is indeed continuous on $\mathrm{Gr}(\varphi)$.

Let us now analyze the VDM Relations (V1)-(V3).

We first observe that $\mathfrak{M}_0$ can be identified with the region

$$\{\mathbf{z} \in Z : z_1 + z_2 = 1,\ 0.5 \leq z_1 \leq 1\};$$

and so, $\emptyset \neq \mathfrak{M}_0 \subset \mathfrak{M}$ (recalling that the upper hemicontinuity of φ at $\alpha = 0$ guarantees this by Lemma IV.2.8). Now, observe that we once again have

$$F^*(0) = \min\left\{-\frac{4-z_1}{5-z_1} : 0 \leq z_1 \leq 1\right\} = -\frac{4-0}{5-0} = -\frac{4}{5};$$

but, for each $\alpha \in (0, 1]$, we have

$$F^*(\alpha) = \min\left\{-\frac{4+\alpha-2z_1}{5+\alpha-2z_1} : 0.5 \leq z_1 \leq 1\right\} = -\frac{4+\alpha-2(0.5)}{5+\alpha-2(0.5)} = -\frac{3+\alpha}{4+\alpha}.$$

Thus

$$\mathfrak{M}_0^* = \{(0.5, 0.5)\} \qquad \text{and} \qquad \mathfrak{M}^* = \{(0, 1)\};$$

and so, (V2) cannot hold since $\mathfrak{M}_0^* \cap \mathfrak{M}^* = \emptyset$. We furthermore see that

$$J_0^* = \lim_{\alpha \downarrow 0} \hat{J}_\alpha^* = \lim_{\alpha \downarrow 0} F^*(\alpha) = -\lim_{\alpha \downarrow 0} \frac{3+\alpha}{4+\alpha} = -\frac{3}{4} > -\frac{4}{5} = F^*(0) = J^*.$$

So, (V3) does not hold; but, in agreement with Lemma IV.2.17 and Lemma IV.2.20, we do obtain the inequality $J_0^* \geq J^*$.

V.2 Example: A Controlled Diffusion Problem

We consider drifted Brownian motion $X = \{X_t : t \in \mathbb{R}^+\}$ on $E = [0, 1]$ with reflections at the boundary points 0 and 1; i.e., on $E^\circ = (0, 1)$, the process X satisfies the stochastic differential equation

$$dX_t = -u_t\, dt + \sigma\, dW_t \qquad x_0 \in (0, 1),$$

where $u_t \in G = [0, 2]$ and $\sigma > 0$ are, respectively, the *drift* and *diffusion* coefficients of X, $X_0 = x_0$ is the initial state of X, and $W = \{W_t : t \in \mathbb{R}^+\}$ is standard one-dimensional Brownian motion. Note that S° denotes the interior of the set S.

The generator $L : \mathscr{D}(L) \to \mathscr{R}(L)$ is then given by

$$Lf(x,u) = \frac{\sigma^2}{2} f''(x) - uf'(x), \qquad \forall (x,u) \in E \times G,$$

where $\mathscr{D}(L) = \{f \in C^2(E^\circ) \cap C(E) : \partial_+ f(0) = 0 = \partial_- f(1)\}$. To be precise, each $f \in \mathscr{D}(L)$ is right-differentiable at 0 and left-differentiable at 1. In what follows, however, we will simply use the notation $f'(x)$ to denote the derivative of such an f at each $x \in E$.

Observe that E and G are both compact subsets of $\mathbb{R}$, and that the product space $E \times G$ is a compact rectangle in $\mathbb{R}^2$. This being the case, it poses little difficulty to verify that L satisfies conditions (D1)-(D6).

We take

$$c(x,u) = x^2 + 2x + u^2, \qquad \forall (x,u) \in E \times G,$$

as our cost rate function for this problem, which is easily seen to satisfy conditions (C1)-(C3).

In keeping with a more classical formulation of this problem, we omit the budget rate function c_1 in this example, and demonstrate that our results can nevertheless be easily adapted when a budget constraint is perhaps excluded.

Let us now state the appropriate linear programs for this problem.

V.2.1 The Long-Term Average Linear Program

The long-term average linear program is now as follows:

$$\text{Minimize} \quad \langle c, \rho \rangle \qquad \text{subject to} \qquad \begin{cases} \langle Lf, \rho \rangle = 0, \quad \forall f \in \mathscr{D}(L); & \text{(MA1)} \\ \langle 1, \rho \rangle = 1; & \text{(MA2)} \\ \rho \in \mathcal{P}(E \times G). & \text{(MA3)} \end{cases}$$

Since we are omitting the budget constraint for this problem, we will use $\mathbb{M}$ to denote the set of feasible measures to distinguish it from our usual feasible set $\mathfrak{M}$.

The objective function is given by

$$\langle c, \rho \rangle = \int_{E\times G} \left(x^2 + 2x + u^2\right) \rho(dx, du),$$

and the adjoint condition (MA1) can be written

$$\int_{E\times G} \left[\frac{\sigma^2}{2} f''(x) - uf'(x)\right] \rho(dx, du) = 0, \qquad \forall f \in \mathscr{D}(L).$$

It will prove useful to derive a more explicit characterization of each feasible $\rho \in \mathbb{M}$ (recalling that, by assumption, $\mathbb{M} \neq \emptyset$). To this end, let $\rho \in \mathbb{M}$ be fixed and arbitrary, and let η be the regular conditional distribution of ρ given $x \in E = [0, 1]$. Note that, for each $t \in \mathbb{R}^+$, the random variable X_t is Gaussian; and so, we may assume that the state marginal measure ρ^E has a density function m. The computation in Section C of the Appendix then shows that

$$m(x) = N_{\bar{u}} \exp\left\{-\frac{2}{\sigma^2} \int_0^x \bar{u}(z)\, dz\right\}, \quad \forall x \in [0, 1]. \tag{2.1}$$

where $\bar{u}$ is defined by

$$\bar{u}(x) \coloneqq \int_G u\, \eta(x, du), \qquad \forall x \in [0, 1],$$

and $N_{\bar{u}}$ is the normalizing constant

$$N_{\bar{u}} \coloneqq \left(\int_0^1 \exp\left\{-\frac{2}{\sigma^2} \int_0^y \bar{u}(z)\, dz\right\} dy\right)^{-1}.$$

Remark 2.2. To place some restriction on the admissible class of controls, we will assume that $\bar{u}$ is a differentiable function of x. $\diamond$

What (2.1) shows is that, for each $\rho \in \varphi(0)$, the regular conditional distribution η of ρ uniquely determines a density function m for ρ^E in terms of the control $\bar{u}$ and the (given) diffusion coefficient σ; and so, when selecting an *optimal* measure $\rho \in \varphi(0)$, the optimization will take place over different choices of the function $\bar{u} : E \to G$.

V.2.2 The α-Normalized Linear Program

Let $\alpha \in (0,1]$ be given and assume that $X_0 = x_0 \in (0,1)$. The α-normalized linear program for this example is then as follows:

$$\text{Minimize} \quad \langle c, \hat{\mu}_\alpha \rangle \qquad \text{subject to} \qquad \begin{cases} \langle L_\alpha f, \hat{\mu}_\alpha \rangle = -\alpha f(x_0), \quad \forall f \in \mathscr{D}(L); & \text{(MN1)} \\ \langle 1, \hat{\mu}_\alpha \rangle = 1; & \text{(MN2)} \\ \hat{\mu}_\alpha \in \mathcal{P}(E \times G). & \text{(MN3)} \end{cases}$$

As above, we let $\hat{\mathbb{M}}_\alpha$ denote the feasible set for this problem to distinguish it from $\hat{\mathfrak{M}}_\alpha$. The objective function is

$$\langle c, \hat{\mu}_\alpha \rangle = \int_{E\times G} \left(x^2 + 2x + u^2\right) \hat{\mu}_\alpha(dx, du),$$

and the adjoint condition (MN1) can be written

$$\int_{E\times G} \left[\frac{\sigma^2}{2} f''(x) - uf'(x) - \alpha f(x)\right] \hat{\mu}_\alpha(dx, du) = -\alpha f(x_0), \qquad \forall f \in \mathscr{D}(L).$$

As with the long-term average linear program, we seek a more explicit characterization of each $\hat{\mu}_\alpha \in \hat{\mathbb{M}}_\alpha$. With this purpose in mind, we let $\hat{\mu}_\alpha \in \hat{\mathbb{M}}_\alpha$, let η_α denote the regular conditional distribution for $\hat{\mu}_\alpha$, and let m_α denote the density function for $\hat{\mu}_\alpha^E$. Another computation in Section C of the Appendix then shows that m_α must satisfy the second-order homogeneous linear ordinary differential equation

$$m_\alpha''(x) + \frac{2}{\sigma^2}\bar{u}_\alpha(x) m_\alpha'(x) + \frac{2}{\sigma^2}\left(\bar{u}_\alpha'(x) - \alpha\right) m_\alpha(x) = 0, \qquad \forall x \in [0,1], \tag{2.3}$$

where, again, $\bar{u}_\alpha$ is defined by

$$\bar{u}_\alpha(x) := \int_G u\, \eta_\alpha(x, du), \qquad \forall x \in [0,1].$$

As with the long-term average linear program, we will assume that $\bar{u}_\alpha$ is differentiable (see Remark 2.2). Note that, if we further assume $\bar{u}_\alpha \in C^1(E^\circ)$, then the basic theory of ordinary differential equations guarantees that (2.3) has a solution (see, e.g., Theorem 5.1 in Chapter 1 of Coddington and Levinson (1955)); however, this differential equation cannot be solved explicitly for the density function m_α without additional structure.

V.2.3 Correspondences and Optimality of Strict Controls

For the feasibility correspondence $\varphi : [0,1] \twoheadrightarrow \mathcal{P}(E \times G)$ in this example, we have

$$\varphi(\alpha) = \begin{cases} \mathbb{M} & \text{if } \alpha = 0, \\ \hat{\mathbb{M}}_\alpha & \text{if } \alpha \in (0,1]. \end{cases}$$

The objective function $F : \mathrm{Gr}(\varphi) \to \mathbb{R}$ is

$$F(\alpha, \mu_\alpha) = \begin{cases} \langle c, \rho \rangle & \text{if } \alpha = 0, \\ \langle c, \hat{\mu}_\alpha \rangle & \text{if } \alpha \in (0,1]; \end{cases}$$

where we may write

$$\langle c, \rho \rangle = \int_{E\times G} c(x,u)\, \rho(dx, du) = \int_0^1 \left(x^2 + 2x + \int_G u^2\, \eta(x, du) \right) m(x)\, dx$$

and

$$\langle c, \hat{\mu}_\alpha \rangle = \int_{E\times G} c(x,u)\, \hat{\mu}_\alpha(dx, du) = \int_0^1 \left(x^2 + 2x + \int_G u^2\, \eta_\alpha(x, du) \right) m_\alpha(x)\, dx.$$

Let us now make the following observation about our choice of controls. Fix an $x \in E$ and suppose that $U_x \sim \eta$. Then

$$\int_G u^2\, \eta(x, du) = \mathbb{E}[U_x^2] = \mathrm{Var}(U_x) + \left(\mathbb{E}[U_x]\right)^2 = \mathrm{Var}(U_x) + \bar{u}^2(x).$$

So, any such “random” variable U_x satisfying $\mathrm{Var}(U_x) = 0$ (i.e., a *strict* control) will have corresponding cost $\langle c, \mu_\alpha \rangle$ less than that of any choice of U_x with $\mathrm{Var}(U_x) > 0$ (i.e., a relaxed control) and the same mean. It follows that the optimization can be taken over only strict controls, in which case we can simply write, e.g., $u(x) = \bar{u}(x)$ for each $x \in E$. Thus, our costs can then be expressed as

$$\langle c, \rho \rangle = \int_0^1 \left(x^2 + 2x + u^2(x) \right) m(x)\, dx$$

and

$$\langle c, \hat{\mu}_\alpha \rangle = \int_0^1 \left(x^2 + 2x + u_\alpha^2(x) \right) m_\alpha(x)\, dx.$$

V.2.4 Verification of Main Result

Let us now consider the following hypotheses:

(H1) φ is continuous at $\alpha = 0$,

(H2) φ has nonempty compact values, and

(H3) F is continuous.

To first show that φ is upper hemicontinuous at $\alpha = 0$, we let $\{(\alpha_n, \hat{\mu}_{\alpha_n}) : n \in \mathbb{N}\}$ be a sequence in $\mathrm{Gr}(\varphi)$ with $\alpha_n \to 0$ (where we assume, as usual, that $\alpha_n > 0$ for each $n \in \mathbb{N}$). Since $\hat{\mu}_\alpha \in \varphi(\alpha_n)$ for every $n \in \mathbb{N}$, it follows that

$$\langle L_{\alpha_n} f, \hat{\mu}_\alpha \rangle = -\alpha_n f(x_0), \qquad \forall f \in \mathscr{D}(L).$$

Since, for every $f \in \mathscr{D}(L)$, the function Lf is a continuous function on the compact set $E \times G$, we obtain (via the argument given in Section 2 of Chapter III)

$$\langle Lf, \mu_0 \rangle = 0, \qquad \forall f \in \mathscr{D}(L),$$

for any μ_0 satisfying $\hat{\mu}_{\alpha_n} \Rightarrow \mu_0$. Hence, $\mu_0 \in \varphi(0)$; and so, φ is upper hemicontinuous at $\alpha = 0$. As for the lower hemicontinuity of φ at $\alpha = 0$, a more involved approach is required. Let $\{\alpha_n : n \in \mathbb{N}\} \subset (0, 1]$ be a sequence satisfying $\alpha_n \to 0$, and fix $\rho_0 \in \varphi(0)$ arbitrarily. Recall that the density m_0 of the state marginal ρ_0^E is given by

$$m_0(x) = N_{\bar{u}_0} \exp\left\{-\frac{2}{\sigma^2}\int_0^x \bar{u}_0(z)\,dz\right\}, \quad \forall x \in [0, 1]. \tag{2.4}$$

Now, let $\hat{\mu}_\alpha \in \varphi(\alpha)$ and recall that the density function m_α for the state marginal measure $\hat{\mu}_\alpha^E$ must satisfy the differential equation (2.3); i.e.,

$$m_\alpha''(x) + \frac{2}{\sigma^2}\bar{u}_\alpha(x)m_\alpha'(x) + \frac{2}{\sigma^2}\left(\bar{u}_\alpha'(x) - \alpha\right)m_\alpha(x) = 0, \qquad \forall x \in [0, 1].$$

However, we observe that, by substituting m_0 into the left-hand side of (2.3), we obtain

$$m_0''(x) + \frac{2}{\sigma^2}\bar{u}_0(x)m_0'(x) + \frac{2}{\sigma^2}\left(\bar{u}_0'(x) - \alpha\right)m_0(x) = -\alpha\frac{2}{\sigma^2}m_0(x), \qquad \forall x \in [0, 1].$$

So, when $\alpha = 0$, the differential equation (2.3) has the unique solution m_0. Thus, by Theorem 4.1 in Chapter 2 of Coddington and Levinson (1955) (see Section C of the Appendix), there exists a $\delta > 0$ such that, for any fixed $\alpha \in (0, 1]$ with $\alpha < \delta$, every solution m_α of (2.3) exists and, moreover, $m_\alpha \to m_0$ uniformly over $[0, 1]$. Hence, given a sequence $\{\alpha_n : n \in \mathbb{N}\} \subset (0, 1]$ with $\alpha_n \to 0$, we can construct a strictly decreasing subsequence $\{\alpha_{n_k} : k \in \mathbb{N}\} \subset (0, \delta)$ such that, for each $k \in \mathbb{N}$, we can choose a measure $\mu_k \in \mathcal{P}(E \times G)$ whose regular conditional distribution is η_0 and whose state marginal μ_k^E has density function $m_{\alpha_{n(k)}}$. It then follows that $\mu_k \in \varphi(\alpha_{n_k})$ for each $k \in \mathbb{N}$ and $\mu_k \Rightarrow \rho_0$. Therefore, φ is lower hemicontinuous at $\alpha = 0$.

Now recall that, for our constrained optimization problems, we take as one of our basic assumptions that $\mathfrak{M} \neq \emptyset$ and $\hat{\mathfrak{M}}_\alpha \neq \emptyset$ for each $\alpha \in (0, 1]$. However, we can observe that such an assumption is unnecessary for this example. Indeed: Fix a $(X, \Lambda) \in \mathscr{M}$, and let

$\{\rho_{t_n} : n \in \mathbb{N}\}$ be a sequence of finite-horizon average occupation measures associated with (X, Λ), as in the proof of Proposition II.1.4. The compactness of $E \times G$ then guarantees that this sequence is tight, which yields a measure $\rho \in \mathbb{M}$. Similarly, for a fixed $\alpha \in (0, 1]$, the proof of Proposition II.1.17 and the compactness of $E \times G$ yield a measure $\hat{\mu}_\alpha \in \hat{\mathbb{M}}_\alpha$. Hence, φ has nonempty values.

To show that $\varphi(0) = \mathbb{M}$ is compact, we need only to (once again) recognize that each $\rho \in \varphi(0)$ is a measure on the compact space $E \times G$; and so, any sequence $\{\rho_n : n \in \mathbb{N}\} \subset \varphi(0)$ necessarily has a subsequence converging (weakly) to some measure in $\varphi(0)$. The same argument applies to $\hat{\mathbb{M}}_\alpha$ for each $\alpha \in (0, 1]$. Hence, φ is compact-valued.

Finally, the continuity of F follows from the definition of weak convergence and the compactness of $E \times G$. Indeed: If $\{(\alpha_n, \mu_{\alpha_n}) : n \in \mathbb{N}\} \subset \mathrm{Gr}(\varphi)$ is a sequence with $(\alpha_n, \mu_{\alpha_n}) \to (\alpha, \mu_\alpha)$ then we have

$$\lim_{n\to\infty} F(\alpha_n, \mu_{\alpha_n}) = \lim_{n\to\infty} \int_{[0,1]^2} c(x, u)\, \mu_{\alpha_n}(dx, du) = \int_{[0,1]^2} c(x, u)\, \mu_\alpha(dx, du) = F(\alpha, \mu_\alpha)$$

since $\mu_{\alpha_n} \Rightarrow \mu$ and c is a continuous function on the compact set $E \times G$ (hence, c is bounded).

So, having satisfied hypotheses (H1)-(H3), we obtain the desired (appropriately modified) conclusions

(V1) $\emptyset \neq \mathbb{M}_0 \subset \mathbb{M}$,

(V2) $\emptyset \neq \mathbb{M}_0^* \subset \mathbb{M}^*$, and

(V3) $\tilde{J}_0^* = \tilde{J}^*$,

where

$$\tilde{J}_0^* = \lim_{\alpha \downarrow 0} \Big(\inf\{\langle c, \hat{\mu}_\alpha \rangle : \hat{\mu}_\alpha \in \hat{\mathbb{M}}_\alpha\} \Big) \qquad \text{and} \qquad \tilde{J}^* = \inf\{\langle c, \rho \rangle : \rho \in \mathbb{M}\}.$$

Moreover, from our deductions above, we see that the optimal value $\tilde{J}^*$ for the long-term average linear program admits the more explicit and tractable form

$$\tilde{J}^* = \inf \left\{ N_u \int_0^1 \left(x^2 + 2x + u^2(x)\right) \exp\left\{ -\frac{2}{\sigma^2} \int_0^x u(z)\, dz \right\} dx : u \in M(E) \right\}.$$

In fact, we can say a bit a more by further analyzing the long-term average problem using dynamic programming methods, as we now demonstrate.

V.2.5 Solution of the LTA Problem Via Dynamic Programming

We seek a function $h \in \mathscr{D}(L)$ and a constant $\lambda \in \mathbb{R}$ satisfying the HJB equation

$$\inf \{c(x,u) + Lh(x,u) : u \in G\} = \lambda, \qquad \forall x \in E.$$

So, we guess that h is of the form

$$h(x) = Ax^2 + Bx + Cg(x),$$

for some $A, B, C \in \mathbb{R}$, where

$$g(x) := \int_0^x \exp\left\{ \frac{2}{\sigma^2} \int_0^y u(z)\, dz \right\} dy,$$

and $u \in M(E)$ is yet to be determined. We note that, for a fixed feedback control function $u \in M(E)$, we can write

$$Lf(x) = \frac{\sigma^2}{2} f''(x) - u(x)f'(x), \qquad \forall (x,f) \in E \times \mathscr{D}(L).$$

Hence, we have

$$Lg(x) = \frac{\sigma^2}{2} g''(x) - u(x)g'(x) = u(x)g'(x) - u(x)g'(x) = 0, \qquad \forall x \in E.$$

Now, to determine the coefficients A, B, C, we first consider the condition $h'(0) = 0 = h'(1)$. We have

$$0 = h'(0) = 2A(0) + B + Cg'(0) = B + Cg'(0) = B + C,$$

which implies that $B = -C$. We also have

$$0 = h'(1) = 2A - C + Cg'(1) = 2A + C[g'(1) - 1],$$

which then implies that $A = -Ck/2$, where $k \coloneqq g'(1) - 1$. So,

$$h(x) = -\frac{C}{2}kx^2 - Cx + Cg(x).$$

Now, using the linearity of L, we have

$$Lh(x) = -\frac{C}{2}kL[x^2] - CL[x] + CLg(x) = -\frac{C}{2}kL[x^2] - CL[x],$$

since $Lg(x) = 0$. We then compute

$$L[x^2] = \frac{\sigma^2}{2}(2) - 2xu(x) = \sigma^2 - 2xu(x) \qquad \text{and} \qquad L[x] = \frac{\sigma^2}{2}(0) - u(x) = -u(x).$$

Now, taking u to be arbitrary, we have

$$Lh(x, u) = -\frac{C}{2}k\left(\sigma^2 - 2ux\right) - C(-u) = C(x+1)u - \frac{C}{2}k\sigma^2,$$

from which it follows that

$$c(x, u) + Lh(x, u) = x^2 + 2x + u^2 + C(x+1)u - \frac{C}{2}k\sigma^2. \tag{2.5}$$

This expression is minimized when $u = u^*(x) = -C(x+1)/2$. Substituting this back into (2.5), our task is then to determine a $\lambda \in \mathbb{R}$ that satisfies

$$\left(1 - \frac{C^2}{4}\right) x^2 + 2\left(1 - \frac{C^2}{4}\right) x - \left(\frac{C^2}{4} + \frac{C}{2} k\sigma^2\right) = \lambda, \qquad \forall x \in E.$$

This is only possible if $C = \pm 2$. The choice of $C = 2$ leads to a contradiction; so, by choosing $C = -2$, we obtain the control function

$$u^*(x) = -\frac{C}{2}(x+1) = x + 1.$$

We can verify that this control function is optimal in the following manner. Since

$$\lambda \leq c(x,u) + Lh(x,u), \qquad \forall (x,u) \in E \times G,$$

λ is a lower bound on $c(x,u) + Lh(x,u)$; and so, $\lambda \leq \tilde{J}^*$. We can write

$$-Lh(x,u) \leq c(x,u) - \lambda, \qquad \forall (x,u) \in E \times G.$$

So, for any feedback control function u and resulting process $(X, u(X)) \in \mathscr{M}$, we have, for each $t > 0$,

$$\begin{aligned} 0 &= \frac{1}{t}\mathbb{E}[h(X_t) - h(X_0)] - \frac{1}{t}\mathbb{E}\left[\int_0^t Lh(X_t, u(X))\, ds\right] \\ &\leq \frac{1}{t}\mathbb{E}[h(X_t) - h(X_0)] + \frac{1}{t}\mathbb{E}\left[\int_0^t c(X_t, u(X))\, ds\right] - \lambda. \end{aligned}$$

Since h is bounded (as it is a continuous function on the compact set E), we can take $t \to \infty$ along an appropriate subsequence and obtain

$$\lambda = \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t c(X_t, u(X))\, ds\right] = \tilde{J}(X, u(X)).$$

Then, since $J(X, u^*(X)) = \lambda$, the infimum in $\tilde{J}^*$ is obtained by the process $(X, u^*(X)) \in \mathscr{M}$; and so, $\tilde{J}^* = \lambda = \tilde{J}(X, u^*(X))$, as desired.

Now, with the choice of $u = u^*$, we compute

$$k = g'(1) - 1 = \exp\left\{\frac{2}{\sigma^2}\int_0^1 (z+1)\,dz\right\} - 1 = \exp\left\{\frac{3}{\sigma^2}\right\} - 1,$$

which yields the optimal value

$$\tilde{J}^* = \lambda = -\left(\frac{C^2}{4} + \frac{C}{2}k\sigma^2\right) = -1 + k\sigma^2 = \left(\exp\left\{\frac{3}{\sigma^2}\right\} - 1\right)\sigma^2 - 1.$$

V.2.6 Analysis of Approximate Solutions Via the VDM

By then applying the vanishing discount method in this setting, one may obtain an approximate solution to the α-*discounted* (or α-normalized) linear program in the sense that, given $\epsilon > 0$, an $\alpha_\epsilon > 0$ can be found that satisfies

$$\left|\tilde{J}^* - \alpha\tilde{J}^*_\alpha\right| = \left|\left(\exp\left\{\frac{3}{\sigma^2}\right\} - 1\right)\sigma^2 - 1 - \alpha\tilde{J}^*_\alpha\right| < \epsilon, \qquad \forall\alpha \in (0, \alpha_\epsilon),$$

where $\tilde{J}^*_\alpha$ is optimal value for the α-discounted linear program. In particular, we emphasize that, with respect to the family of α-discounted linear programs, we have

$$\left|\frac{1}{\alpha}\left[\left(\exp\left\{\frac{3}{\sigma^2}\right\} - 1\right)\sigma^2 - 1\right] - \tilde{J}^*_\alpha\right| < \frac{\epsilon}{\alpha}, \qquad \forall\alpha \in (0, \alpha_\epsilon).$$

Furthermore, by considering a family $\{\hat{\mu}_\alpha : \alpha \in (0, 1]\}$ of feasible α-normalized expected occupation measures for which each $\hat{\mu}_\alpha$ has as its regular conditional distribution the measure η^* satisfying

$$u^*(x) = x + 1 = \int_G u\,\eta^*(x, du), \qquad \forall x \in E,$$

we see that

$$\lim_{\alpha \downarrow 0} \langle c, \hat{\mu}_\alpha \rangle \geq \tilde{J}_0^* = \tilde{J}^*.$$

Thus, if μ_0 is a measure satisfying $\hat{\mu}_\alpha \Rightarrow \mu_0$ as $\alpha \downarrow 0$, then the measure $\alpha^{-1}\mu_0$ will be an approximate solution to the α-discounted linear program for α small enough, and the quantity $\langle c, \mu_0 \rangle$ provides an upper bound on $\tilde{J}_0^*$. As a practical consideration, it is also worth noting that the control $u^*(x) = x + 1$ represents a relatively simple policy to implement.

CONCLUSION

This dissertation concludes with a brief summary of the results we have presented, followed by a description of some directions towards which future research in this areas could possibly be conducted.

VI.1 Summary of Results

The results we have established in this dissertation are what we believe to be a novel application of the theory of correspondences to the vanishing discount method within the linear programming framework for stochastic optimal control as developed in Kurtz and Stockbridge (1998). Indeed, we saw that the fundamental results regarding the equivalence between the stochastic control problems and the linear programs herein remained valid when an additional budget/resource constraint is included in the model. The inclusion of such a budget constraint moreover provided us with additional structure on our feasible sets of measures that allowed for some of our desired results (e.g., Proposition IV.2.9) to be more easily verified; and by characterizing these feasible sets in terms of a feasibility correspondence, we were able to more elegantly state and prove our desired results. Of particular note is the role of *lower hemicontinuity* (of the feasibility correspondence) and *upper semicontinuity* (of the objective function) in our analysis, as these are the two conditions appearing in the hypotheses of our main results, but are not automatically obtained from our model assumptions. The examples provided in Chapter V were further evidence of a certain finicky quality these two conditions seem to possess.

Our formulation of the vanishing discount method provides a viable alternative to, for instance, the more typical dynamic programming approach for solving the long-term average and α-discounted problems. Though, as we saw in Example V.2, these two approaches may even work hand-in-hand to address more stubborn problems—or if one is simply interested in obtaining additional insight into the underlying structural relationship between these optimality criteria.

VI.2 Possible Directions of Future Research

We first note that it should be possible to weaken the assumption of *local compactness* on our state and control spaces, as in Kurtz and Stockbridge (2001). (Bhatt and Karandikar (1993) and Bhatt and Borkar (1996) also consider state spaces that are only assumed to be complete, separable, metric spaces, but the former does not include a control and the latter assumes that the control space is compact.) Ideally, of course, this would be accomplished without a need for imposing further—or more cumbersome—conditions on the generator L (or any of the other components of our model).

Also dealt with in Kurtz and Stockbridge (2001) and Kurtz and Stockbridge (2017) is the more general setting in which a *singular control* is included in the model. We believe that many of our results should be amenable to models including a singular control (e.g., jump processes). Our choice to exclude singular controls from our analysis was motivated primarily by a desire to keep exposition (and notation) as lean and clear as possible.

As with the development of any theory or methodology, we are also naturally interested in identifying further practical applications of this formulation of the vanishing discount method. In particular, given the presence of a budget constraint, this approach should be well-suited to many concrete economic applications in which a certain entity (e.g., a firm) wishes to minimize a cost (or maximize a profit) without consuming too much of a particular resource. Again, we wished not to stray too far from the general scope of this dissertation, which was intended to be primarily theoretical in nature.

We have further interest in additional applications to optimal control problems involving stochastic differential equations. The controlled diffusion problem presented in this dissertation was formulated to serve simply as a clear, tractable demonstration of our results; but there is no reason to think these results will not be applicable in more general scenarios. For instance, in Example V.2, we saw that many of our hypotheses were not too difficult to verify (except, perhaps, the lower hemicontinuity of the feasibility correspondence) because of the presence of *compact* state and control spaces. The situation often becomes markedly more complicated when the state and control spaces are only *locally compact* (e.g., $\mathbb{R}$) since, for one, the cost rate function c is no longer bounded. This being the case, it may be argued that a restriction to models with bounded cost rate functions is perhaps in order. This is likely too restrictive, however, since many of the classical stochastic control problems of this nature feature unbounded cost rate functions; e.g., the Linear-Quadratic-Gaussian control problem and others with quadratic cost criteria. Nevertheless, we believe our basic model formulation to be well-suited to such problems.

Bibliography

Aliprantis, C. D. and Border, K. C. (2006). *Infinite dimensional analysis: a hitchhiker's guide*. Springer.

Alvarez, L. H. R. and Shepp, L. A. (1998). Optimal harvesting of stochastically fluctuating populations. *J. Math. Biol.*, 37(2):155–177.

Berge, C. (1997). *Topological spaces*. Dover Publications, Inc., Mineola, NY. Including a treatment of multi-valued functions, vector spaces and convexity, Translated from the French original by E. M. Patterson, Reprint of the 1963 translation.

Bhatt, A. G. and Borkar, V. S. (1996). Occupation measures for controlled Markov processes: characterization and optimality. *Ann. Probab.*, 24(3):1531–1562.

Bhatt, A. G. and Karandikar, R. L. (1993). Invariant measures and evolution equations for Markov processes characterized via martingale problems. *Ann. Probab.*, 21(4):2246–2268.

Billingsley, P. (1999). *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, second edition. A Wiley-Interscience Publication.

Bogachev, V. I. (2007). *Measure theory. Vol. I, II*. Springer-Verlag, Berlin.

Chung, K. L. and Doob, J. L. (1965). Fields, optionality and measurability. *Amer. J. Math.*, 87:397–424.

Coddington, E. A. and Levinson, N. (1955). *Theory of ordinary differential equations*. McGraw-Hill Book Co., Inc., New York-Toronto-London.

Cohn, D. L. (2013). *Measure theory*. Birkhäuser Advanced Texts: Basler Lehrbücher. [Birkhäuser Advanced Texts: Basel Textbooks]. Birkhäuser/Springer, New York, second edition.

Edwards, C. H., Penney, D. E., and Calvis, D. T. (2005). *Linear algebra and differential equations*. Pearson, fourth edition.

Ethier, S. N. and Kurtz, T. G. (1986). *Markov processes*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York. Characterization and convergence.

Fleming, W. H. and Rishel, R. W. (1975). *Deterministic and stochastic optimal control*. Applications of Mathematics, No. 1. Springer-Verlag, Berlin-New York.

Hernández-Lerma, O. and Prieto-Rumeau, T. (2010). The vanishing discount approach to constrained continuous-time controlled markov chains. *Systems & Control Letters*, 59(8):504–509.

Karatzas, I. and Shreve, S. E. (1991). *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition.

Kawaguchi, K. and Morimoto, H. (2007). Long-run average welfare in a pollution accumulation model. *J. Econom. Dynam. Control*, 31(2):703–720.

Kurtz, T. G. and Stockbridge, R. H. (1998). Existence of Markov controls and characterization of optimal Markov controls. *SIAM J. Control Optim.*, 36(2):609–653.

Kurtz, T. G. and Stockbridge, R. H. (2001). Stationary solutions and forward equations for controlled and singular martingale problems. *Electron. J. Probab.*, 6:no. 17, 52.

Kurtz, T. G. and Stockbridge, R. H. (2017). Linear programming formulations of singular stochastic processes. *arXiv:1707.09209*.

Manne, A. S. (1960). Linear programming and sequential decisions. *Management Sci.*, 6:259–267.

Meyer, P.-A. (1966). *Probability and potentials*. Blaisdell Publishing Co. [Ginn and Co.], Waltham, Mass.-Toronto, Ont.-London.

Montes-de Oca, R. and Lemus-Rodríguez, E. (2012). An unbounded Berge's minimum theorem with applications to discounted Markov decision processes. *Kybernetika (Prague)*, 48(2):268–286.

Øksendal, B. (2003). *Stochastic differential equations*. Universitext. Springer-Verlag, Berlin, sixth edition. An introduction with applications.

Pham, H. (2009). *Continuous-time stochastic control and optimization with financial applications*, volume 61 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin.

Ricceri, B. (2016). A minimax theorem in infinite-dimensional topological vector spaces. *Linear Nonlinear Anal.*, 2(1):47–52.

Song, Q., Stockbridge, R. H., and Zhu, C. (2011). On optimal harvesting problems in random environments. *SIAM J. Control Optim.*, 49(2):859–889.

Sydsaeter, K., Hammond, P., Seierstad, A., and Strøm, A. (2005). *Further mathematics for economic analysis*. Prentice Hall, second edition.

Taylor, III, H. M. (1965). Markovian sequential replacement processes. *Ann. Math. Statist.*, 36:1677–1694.

Veinott, Jr., A. F. (1966). On finding optimal policies in discrete dynamic programming with no discounting. *Ann. Math. Statist.*, 37:1284–1294.

Vieten, M. G. (2018). *Numerical Solution of Stochastic Control Problems Using the Finite Element Method*. ProQuest LLC, Ann Arbor, MI. Thesis (Ph.D.)–The University of Wisconsin - Milwaukee.

Williams, D. (1991). *Probability with martingales*. Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge.

Appendix

A Some Essential Results from Measure Theory

Perhaps the most important and well-known theorem from measure theory is the *Monotone Convergence Theorem*, which we include here for reference. The particular statement of this theorem that we give below comes from Section 2.4 of Cohn (2013).

Theorem A.1. (Monotone Convergence Theorem) *Let $(X, \mathscr{A}, \mu)$ be a measure space, and let f and $f_1, f_2, \ldots$ be $[0, +\infty]$-valued $\mathscr{A}$-measurable functions on X. Suppose that*

$$f_1(x) \leq f_2(x) \leq \cdots$$

and

$$f(x) = \lim_{n\to\infty} f_n(x)$$

hold at μ-almost every $x \in X$. Then

$$\int_X f \, d\mu = \lim_{n\to\infty} \int_X f_n \, d\mu.$$

Another rather useful result is *Fatou's Lemma*. Again, we take the following statement of this result from Section 2.4 of Cohn (2013).

Theorem A.2. (Fatou's Lemma) *Let $(X, \mathscr{A}, \mu)$ be a measure space, and let $\{f_n : n \in \mathbb{N}\}$ be a sequence of $[0, +\infty]$-valued $\mathscr{A}$-measurable functions on X. Then*

$$\int_X \liminf_{n\to\infty} f_n \, d\mu \leq \liminf_{n\to\infty} \int_X f_n \, d\mu.$$

When our focus is *probability* measures, we will occasionally refer to the so-called *Portmanteau Theorem*. The statement of this result found below comes from Section 3 of Chapter 3 in Ethier and Kurtz (1986). This particular statement relies upon the following definition.

Definition A.3. Let $\partial\Gamma \coloneqq \overline{\Gamma} \cap \overline{\Gamma^{\mathsf{c}}}$ denote the boundary of the subset $\Gamma \subset E \times G$, where $\overline{\Gamma}$ and Γ^{c} denote the closure and complement of Γ, respectively. Then, given $\mu \in \mathcal{P}(E \times G)$, we call Γ a *μ-continuity set* if $\Gamma \in \mathcal{B}(E \times G)$ and $\mu(\partial\Gamma) = 0$. $\diamond$

Theorem A.4. (Portmanteau Theorem) *Let $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{P}(X)$ and $\mu \in \mathcal{P}(X)$, let π denote the Prohorov metric on $\mathcal{P}(X)$, and suppose that X is a complete, separable, metric space. Then the following statements are equivalent.*

(a) $\lim_{n\to\infty} \pi(\mu_n, \mu) = 0$.

(b) $\mu_n \Rightarrow \mu$.

(c) $\lim_{n\to\infty} \int_X \xi \, d\mu_n = \int_X \xi \, d\mu$ *for all uniformly continuous* $\xi \in \bar{C}(X)$.

(d) $\limsup_{n\to\infty} \mu_n(\Phi) \leq \mu(\Phi)$ *for all closed sets* $\Phi \subset X$.

(e) $\liminf_{n\to\infty} \mu_n(\Upsilon) \geq \mu(\Upsilon)$ *for all open sets* $\Upsilon \subset X$.

(f) $\lim_{n\to\infty} \mu_n(\Psi) = \mu(\Psi)$ *for all* μ*-continuity sets* $\Psi \subset X$.

We often make use of the following corollary to Fatou's Lemma and the Portmanteau Theorem. Note that this corollary is sometimes referred to as "Fatou's Lemma" itself.

Corollary A.5. *Let* $\mu \in \mathcal{P}(X)$ *and let* $\{\mu_n : n \in \mathbb{N}\} \subset \mathcal{P}(X)$ *be a sequence of probability measures with* $\mu_n \Rightarrow \mu$. *If* $f : X \to \mathbb{R}$ *is lower semicontinuous and bounded below then*

$$\int_X f \, d\mu \leq \liminf_{n\to\infty} \int_X f \, d\mu_n.$$

Proof. Since f is lower semicontinuous and bounded below (say, $-\infty < -\kappa < f(x)$ for every $x \in X$), the function $x \mapsto f(x) + \kappa$ is lower semicontinuous and nonnegative. Thus, for each $\lambda \in \mathbb{R}^+$, the set $H_\lambda := \{x \in X : f(x) + \kappa > \lambda\}$ is open; and so, by the Portmanteau Theorem, we have

$$\mu(H_\lambda) \leq \liminf_{n\to\infty} \mu_n(H_\lambda), \qquad \forall \lambda \in \mathbb{R}^+. \tag{A.6}$$

Define the function $h : \mathbb{R}^+ \to \mathbb{R}^+$ by $h(\lambda) = \mu(H_\lambda)$, and define the sequence $\{h_n : n \in \mathbb{N}\}$ of functions $h_n : \mathbb{R}^+ \to \mathbb{R}^+$ by $h_n(\lambda) = \mu_n(H_\lambda)$. Then $\{h_n : n \in \mathbb{N}\}$ is a sequence of $\mathbb{R}^+$-valued $\mathcal{B}(\mathbb{R}^+)$-measurable functions (noting that each h_n is nondecreasing) on $\mathbb{R}^+$. Integrating both sides of (A.6) and applying Fatou's Lemma then yields

$$\int_0^\infty h(\lambda) \, d\lambda \leq \int_0^\infty \liminf_{n\to\infty} h_n(\lambda) \, d\lambda \leq \liminf_{n\to\infty} \int_0^\infty h_n(\lambda) \, d\lambda,$$

which implies that

$$\int_X (f(x) + \kappa) \, \mu(dx) \leq \liminf_{n\to\infty} \int_X (f(x) + \kappa) \, \mu_n(dx).$$

The desired result then follows from the fact that μ and each μ_n are probability measures. □

B Correspondences, Continuity, and Berge's Theorem

As much of the material in this section is found in Aliprantis and Border (2006), we will—for the most part—adopt the notation used therein throughout this section of the Appendix.

B.1 Basic Definitions

Definition B.1. A *correspondence* φ from a set X to a set Y is an assignment to each $x \in X$ a subset $\varphi(x)$ of Y. We write $\varphi : X \twoheadrightarrow Y$ to distinguish a correspondence from a function from X to Y. ◇

If $\varphi : X \twoheadrightarrow Y$ is a correspondence, then we refer to X as the *domain* of φ and Y as the *codomain* of φ. The *image* of a set $A \subset X$ under φ is the set

$$\varphi(A) = \bigcup_{x \in A} \varphi(x),$$

and the *range* of φ is $\varphi(X)$. The set

$$\mathrm{Gr}(\varphi) := \{(x, y) \in X \times Y : y \in \varphi(x)\}$$

is called the *graph* of φ.

Remark B.2. Since, for each $x \in X$, $\varphi(x)$ is in fact a *subset* of Y, the careful reader may—rightfully—think it more appropriate to view the codomain of φ as the *power set* of Y. However, we believe this slight abuse of terminology and notation allows for clearer and more concise exposition. $\diamond$

We also include the following (well-known) definition(s) for convenience and reference. The reader may consult Section 2.2 of Aliprantis and Border (2006) for further details.

Definition B.3. A *neighborhood of a set* A is any set B for which there is an open set V satisfying $A \subset V \subset B$. Any open set V that satisfies $A \subset V$ is called an *open neighborhood of* A. A *neighborhood of a point* x is any set B for which there is an open set V satisfying $x \in V \subset B$. The collection of all neighborhoods of a point $x \in X$, called the *neighborhood base*, or *neighborhood system*, at x, is denoted by $\mathcal{N}_x$. $\diamond$

B.2 Continuity of Correspondences

A correspondence $\varphi : X \twoheadrightarrow Y$ is said to be *continuous* at a point $x \in X$ if it is both upper hemicontinuous at x and lower hemicontinuous at x. As with functions, we say that φ is *continuous* if it is continuous at each point $x \in X$. We provide the definitions of upper hemicontinuity and lower hemicontinuity below, as well as some discussion and equivalent characterizations. For our purposes, the most useful characterizations of hemicontinuity are the sequential ones. Accordingly, these sequential characterizations are discussed in detail within the main body of this dissertation.

Remark B.4. Some sources use the term *semicontinuity* in place of hemicontinuity—including the seminal treatment of correspondences found in Berge (1997). Indeed, there is an obvious analog between semicontinuity of *functions* and hemicontinuity of correspondences. However, we have chosen "hemicontinuity" in order to avoid the reader's possible conflation of the two concepts. $\diamond$

Definition B.5. The *upper inverse* φ^{u} (or *strong inverse*) of a subset A of Y is the set

$$\varphi^{\mathrm{u}}(A) := \{x \in X : \varphi(x) \subset A\}.$$

$\diamond$

Definition B.6. A correspondence $\varphi : X \to Y$ is *upper hemicontinuous* at the point $x \in X$ if, for every neighborhood U of $\varphi(x)$, there is a neighborhood V of x such that $z \in V$ implies $\varphi(z) \subset U$. (Equivalently, the upper inverse image $\varphi^u(U)$ is a neighborhood of x in X.) As with functions, we say that φ is upper hemicontinuous on X (abbreviated *uhc*) if it is upper hemicontinuous at every point of X. $\diamond$

In practice, we seldom have need to use the above definition directly. Instead, equivalent characterizations, such as the one below, often prove to be of greater use.

Lemma B.7. *If $\varphi : X \twoheadrightarrow Y$ is a correspondence between topological spaces, then the following statements are equivalent.*

(a) φ is upper hemicontinuous.

(b) $\varphi^u(V)$ is open for each open subset V of Y.

(c) $\varphi^\ell(F)$ is closed for each closed subset F of Y.

Proof. We first show equivalence between (a) and (b). Begin by assuming that φ is upper hemicontinuous. Let V be an open subset of Y, and let $x \in \varphi^{\mathrm{u}}(V)$. Then $\varphi(x) \subset V$, and so V is an open neighborhood of $\varphi(x)$. Since φ is upper hemicontinuous, there is an open neighborhood W of x such that $z \in W$ implies $\varphi(z) \subset V$. Hence $x \in W \subset \varphi^{\mathrm{u}}(V)$, so $\varphi^{\mathrm{u}}(V)$ is open. Now, assume (b) holds. Let $x \in X$ and let U be an open neighborhood of $\varphi(x)$ in Y. Since (b) holds, $\varphi^{\mathrm{u}}(U)$ is an open neighborhood of x. By definition, $z \in \varphi^{\mathrm{u}}(U)$ implies that $\varphi(z) \subset U$. Therefore φ is upper hemicontinuous at x; and, since $x \in X$ was arbitrary, it follows that φ is upper hemicontinuous. To show that (b) and (c) are equivalent, we first assume (c) holds. Let V be an open subset of Y, and observe that $\varphi^\ell(V^{\mathsf{c}}) = [\varphi^{\mathrm{u}}(V)]^{\mathsf{c}}$. Since (c) holds and V^{c} is closed, it follows that $\varphi^{\mathrm{u}}(V)$ is an open subset of Y. Hence (b) holds. A similar enough argument shows that (b) implies (c). $\square$

Definition B.8. The *lower inverse* φ^ℓ (or *weak inverse*) of a subset A of Y is the set

$$\varphi^\ell(A) \coloneqq \{x \in X : \varphi(x) \cap A \neq \emptyset\}.$$

$\diamond$

Definition B.9. A correspondence $\varphi : X \to Y$ is *lower hemicontinuous* at the point $x \in X$ if, for every open set U that satisfies $\varphi(x) \cap U \neq \emptyset$, there is a neighborhood V of x such that $z \in V$ implies $\varphi(z) \cap U \neq \emptyset$. (Equivalently, the lower inverse image $\varphi^\ell(U)$ is a neighborhood of x.) As above, φ is lower hemicontinuous on X (abbreviated *lhc*) if it is lower hemicontinuous at every point of X. $\diamond$

We also have the following analog of Lemma B.7 for lower hemicontinuity.

Lemma B.10. *If $\varphi : X \twoheadrightarrow Y$ is a correspondence between topological spaces, then the following statements are equivalent.*

(a) φ is lower hemicontinuous.

(b) $\varphi^\ell(V)$ is open for each open subset V of Y.

(c) $\varphi^u(F)$ is closed for each closed subset F of Y.

Proof. This proof is a trivial modification of the proof of Lemma B.7.

B.3 Berge's Theorem and Related Results

First note that a correspondence $\varphi : X \twoheadrightarrow Y$ is said to have *nonempty values* if $\varphi(x) \neq \emptyset$ for each $x \in X$. Similarly, φ is said to have *closed values* (respectively, *compact values*), or be *closed-valued* (respectively, *compact-valued*), if $\varphi(x)$ is a closed (respectively, compact) set for each $x \in X$. Here, it is important to recognize that, as defined in Aliprantis and Border (2006), a closed-valued correspondence is not necessarily a *closed* correspondence (though the converse of this statement is true). To be precise, the correspondence $\varphi : X \twoheadrightarrow Y$ is said to be *closed* if $\mathrm{Gr}(\varphi)$ is a closed subset of $X \times Y$. For example, the correspondence $\varphi : [0,1] \twoheadrightarrow [0,1]$ defined by

$$\varphi(x) = \begin{cases} \{0\} & \text{if } x > 0, \\ \{1\} & \text{if } x = 0; \end{cases}$$

is closed-valued but not closed.

We now state the *Berge Maximum Theorem* as it appears in Aliprantis and Border (2006) (with some minor changes in notation).

Theorem B.11. *Let $\varphi : X \twoheadrightarrow Y$ be a continuous correspondence between topological spaces with nonempty compact values, and suppose $F : \mathrm{Gr}(\varphi) \to \mathbb{R}$ is continuous. Define the "value function" $F^* : X \to \mathbb{R}$ by*

$$F^*(x) \coloneqq \max\{F(x,y) : y \in \varphi(x)\},$$

and the "argmax" correspondence $\varphi^* : X \twoheadrightarrow Y$ of maximizers by

$$\varphi^*(x) = \{y \in \varphi(x) : F(x,y) = F^*(x)\}.$$

Then the following conditions hold:

(a) F^* is continuous.

(b) φ^* has nonempty compact values.

(c) If either F has a continuous extension to all of $X \times Y$ or Y is Hausdorff, then φ^* is upper hemicontinuous.

Remark B.12. The Berge Maximum Theorem can be easily modified—with no change to its conclusions—to yield a similar "Berge Minimum Theorem." Indeed, one need only to change "max" to "min" and "maximizers" to "minimizers" where appropriate in the above statement. What we call *Berge's Theorem* throughout this dissertation is meant to refer to whichever of these theorems is appropriate in the given context. $\diamond$

The proof of Berge's Theorem relies upon the following two results, which we also call upon a number of times throughout this dissertation. The first of these results is a trivial (but necessary for our purposes) extension of Theorem 2.43 in Aliprantis and Border (2006), which considers only *real*-valued functions; the proof, however, is no different.

Theorem B.13. *A $(-\infty, +\infty]$-valued lower semicontinuous function on a compact space attains a minimum value, and the nonempty set of minimizers is compact. Similarly, a $(-\infty, +\infty]$-valued upper semicontinuous function on a compact space attains a maximum value, and the nonempty set of maximizers is compact.*

The result below is Lemma 2.2 in Montes-de Oca and Lemus-Rodríguez (2012), but it can be viewed as a corollary to Lemma 17.29 in Aliprantis and Border (2006).

Lemma B.14. *Suppose $\varphi : X \twoheadrightarrow Y$ is a lower hemicontinuous correspondence and $F : \mathrm{Gr}(\varphi) \to (-\infty, +\infty]$ is an upper semicontinuous function. Then the function $F^* : X \to Y$ defined by*

$$F^*(x) = \min\{F(x, y) : y \in \varphi(x)\}$$

is upper semicontinuous.

The following is then a trivial corollary to the preceding lemma, but it allows us to simplify some of our arguments.

Corollary B.15. *Suppose the correspondence $\varphi : X \twoheadrightarrow Y$ is lower hemicontinuous at $x = x_0$ and the function $F : \mathrm{Gr}(\varphi) \to (-\infty, +\infty]$ is upper semicontinuous on $\{x_0\} \times \varphi(x_0)$. Then the function $F^* : X \to Y$ defined by*

$$F^*(x) = \min\{F(x, y) : y \in \varphi(x)\}$$

is upper semicontinuous at $x = x_0$.

C Miscellaneous Results for Example V.2

In this section, we include the computations used to derive the density functions m_0 and m_α for, respectively, the state marginal measures ρ^E and $\hat{\mu}_\alpha^E$ in Example V.2. This derivation involves some basic techniques from the theory of ordinary differential equations that one encounters in an undergraduate course; see, e.g., Edwards et al. (2005). Using these techniques, we see that an explicit expression for m_0 can be obtained, but that m_α is given only implicitly as a solution to a second-order linear homogeneous ordinary differential equations.

We also provide some of the omitted details for the argument that the feasibility correspondence φ in this diffusion problem is lower hemicontinuous at $\alpha = 0$. Among these details is an important theorem from Coddington and Levinson (1955) that provides the desired existence and convergence results for the (α-parameterized) family of differential equations that implicitly characterize the α-normalized state marginal densities.

C.1 Density for the Long-Term Average Linear Program

Let $\rho_0 \in \varphi(0)$. Then we have

$$\int_{[0,1]^2} Lf(x, u)\, \rho_0(dx, du) = 0, \qquad \forall f \in \mathscr{D}(L).$$

Now let m_0 be the density function for ρ_0^E, let η_0 be the regular conditional distribution of ρ_0, let $\varsigma = \frac{1}{2}\sigma^2$, and let

$$\bar{u}_0(x) \coloneqq \int_G u\,\eta_0(x, du), \qquad \forall x \in [0, 1].$$

Then, for each $f \in \mathscr{D}(L)$, we have

$$\begin{aligned}
0 &= \int_E \left(\int_G \{\varsigma f''(x) - u f'(x)\}\, \eta_0(x, du) \right) \rho_0^E(dx) \\
&= \int_E \left(\left\{ \int_G \varsigma\, \eta_0(x, du) \right\} f''(x) - \left\{ \int_G u\, \eta_0(x, du) \right\} f'(x) \right) \rho_0^E(dx) \\
&= \int_0^1 (\varsigma f''(x) - \bar{u}_0(x) f'(x))\, m_0(x)\, dx.
\end{aligned}$$

Put $g \coloneqq \varsigma m_0$. Then integrating by parts yields

$$\begin{aligned}
\int_0^1 (\varsigma f''(x) m_0(x))\, dx &= [g(x) f'(x)]_0^1 - \int_0^1 f'(x) g'(x)\, dx \\
&= g(1) f'(1) - g(0) f'(0) - \int_0^1 f'(x) g'(x)\, dx \\
&= 0 - 0 - \left([g'(x) f(x)]_0^1 - \int_0^1 f(x) g''(x)\, dx \right) \\
&= -g'(1) f(1) + g'(0) f(0) + \int_0^1 f(x) g''(x)\, dx \\
&= -\varsigma m_0'(1) f(1) + \varsigma m_0'(0) f(0) + \int_0^1 \varsigma f(x) m_0''(x)\, dx.
\end{aligned}$$

It can furthermore be shown (via integrating by parts) that

$$\int_0^1 \bar{u}_0(x) f'(x) m_0(x)\, dx = \bar{u}_0(1) m_0(1) f(1) - \bar{u}_0(0) m_0(0) f(0) - \int_0^1 f(x) (\bar{u}_0 m_0)'(x)\, dx.$$

Thus, (MA1) can be written

$$\begin{aligned}
0 &= -\varsigma m_0'(1) f(1) + \varsigma m_0'(0) f(0) + \int_0^1 \varsigma f(x) m_0''(x)\, dx \\
&\qquad - \bar{u}_0(1) m_0(1) f(1) + \bar{u}_0(0) m_0(0) f(0) + \int_0^1 f(x) (\bar{u}_0 m_0)'(x)\, dx \\
&= [\varsigma m_0'(0) + \bar{u}_0(0) m_0(0)]\, f(0) - [\varsigma m_0'(1) + \bar{u}_0(1) m_0(1)]\, f(1) \\
&\qquad + \int_0^1 [\varsigma m_0''(x) + (\bar{u}_0 m_0)'(x)]\, f(x)\, dx.
\end{aligned}$$

Since (MA1) must hold for every f in the set

$$\{f \in \mathscr{D}(L) : f(0) = 0 = f(1);\ f(x) > 0,\ \forall x \in (0,1)\},$$

it follows that

$$\varsigma m_0''(x) + (\bar{u}_0 m_0)'(x) = 0, \qquad \forall x \in [0,1].$$

So, let $x \in [0,1]$ be arbitrary. Then

$$0 = \int_0^x [\varsigma m_0''(\lambda) + (\bar{u}_0 m_0)'(\lambda)]\, d\lambda = [\varsigma m_0'(x) + \bar{u}_0(x) m_0(x)] - [\varsigma m_0'(0) + \bar{u}_0(0) m_0(0)]\,.$$

Thus,

$$\varsigma m_0'(x) + \bar{u}_0(x) m_0(x) = \varsigma m_0'(0) + \bar{u}_0(0) m_0(0), \qquad \forall x \in [0,1];$$

i.e., the function $x \mapsto \varsigma m_0'(x) + \bar{u}_0(x) m_0(x)$ is constant on $[0,1]$. By then choosing an $f \in \mathscr{D}(L)$ satisfying $f(0) \neq 0$ and $f(x) = 0$ for $x \in (0,1]$, it follows that $\varsigma m_0'(0) + \bar{u}_0(0) m_0(0) = 0$; and so,

$$m_0'(x) + \varsigma^{-1} \bar{u}_0(x) m_0(x) = 0, \qquad \forall x \in [0,1].$$

The substitutions $y \coloneqq m_0$ and $\gamma \coloneqq \varsigma^{-1}\bar{u}_0$ then allow us to write this differential equation in the standard form

$$y' + \gamma(x) y = 0, \qquad \forall x \in [0,1]. \tag{C.1}$$

By then introducing the integrating factor

$$w(x) \coloneqq \exp\left\{\int_0^x \gamma(z)\, dz\right\}, \qquad \forall x \in [0,1]$$

the differential equation (C.1) can be written

$$(wy)'(x) = 0, \qquad \forall x \in [0,1],$$

which implies that

$$w(x) y(x) = C, \qquad \forall x \in [0,1],$$

for some constant $C \in \mathbb{R}$. Thus,

$$y(x) = C \exp\left\{-\int_0^x \gamma(z)\, dz\right\} = C \exp\left\{-\frac{1}{\varsigma}\int_0^x \bar{u}_0(z)\, dz\right\}, \qquad \forall x \in [0,1];$$

and so,

$$m_0(x) = C \exp\left\{-\frac{2}{\sigma^2}\int_0^x \bar{u}_0(z)\, dz\right\}, \qquad \forall x \in [0,1].$$

Now, since $\int_0^1 m_0(x)\, dx = 1$, we have

$$C \int_0^1 \exp\left\{-\frac{2}{\sigma^2}\int_0^x \bar{u}_0(z)\, dz\right\} dx = 1;$$

and so,

$$C = \left(\int_0^1 \exp\left\{ -\frac{2}{\sigma^2} \int_0^x \bar{u}_0(z)\, dz \right\} dx \right)^{-1}.$$

Therefore, the density function m_0 for ρ_0^E is given by

$$m_0(x) = \left(\int_0^1 \exp\left\{ -\frac{2}{\sigma^2} \int_0^y \bar{u}_0(z)\, dz \right\} dy \right)^{-1} \exp\left\{ -\frac{2}{\sigma^2} \int_0^x \bar{u}_0(z)\, dz \right\}, \qquad \forall x \in [0,1].$$

C.2 Density for the α-Normalized Linear Program

Let $\hat{\mu}_\alpha \in \hat{\mathbb{M}}_\alpha$. Then we have

$$\int_{[0,1]^2} L_\alpha f(x,u)\, \hat{\mu}_\alpha = -\alpha f(x_0), \qquad \forall f \in \mathscr{D}(L).$$

Now, let m_α be the density function for $\hat{\mu}_\alpha^E$, let η_α be the regular conditional distribution of $\hat{\mu}_\alpha$, let $\varsigma := \frac{1}{2}\sigma^2$, and let

$$\bar{u}_\alpha(x) := \int_G u\, \eta_\alpha(x, du), \qquad \forall x \in [0,1].$$

Then, for each $f \in \mathscr{D}(L)$, we have

$$\begin{aligned} -\alpha f(x_0) &= \int_E \left(\int_G \{\varsigma f''(x) - u f'(x) - \alpha f(x)\}\, \eta_\alpha(x, du) \right) \hat{\mu}_\alpha^E(dx) \\ &= \int_E \left(\varsigma f''(x) - \left\{ \int_G u\, \eta_\alpha(x, du) \right\} f'(x) - \alpha f(x) \right) \hat{\mu}_\alpha^E(dx) \\ &= \int_0^1 (\varsigma f''(x) - \bar{u}_\alpha(x) f'(x) - \alpha f(x))\, m_\alpha(x)\, dx. \end{aligned}$$

So, again, we analyze the integral

$$\int_0^1 (\varsigma f''(x) - \bar{u}_\alpha(x) f'(x) - \alpha f(x))\, m_\alpha(x)\, dx$$

term by term.

Put $g \coloneqq \varsigma m_\alpha$. Then integrating by parts yields

$$\begin{aligned}
\int_0^1 (\varsigma f''(x) m_\alpha(x))\, dx &= [g(x)f'(x)]_0^1 - \int_0^1 f'(x)g'(x)\, dx \\
&= g(1)f'(1) - g(0)f'(0) - \int_0^1 f'(x)g'(x)\, dx \\
&= 0 - 0 - \left([g'(x)f(x)]_0^1 - \int_0^1 f(x)g''(x)\, dx \right) \\
&= -g'(1)f(1) + g'(0)f(0) + \int_0^1 f(x)g''(x)\, dx \\
&= -\varsigma m_\alpha'(1)f(1) + \varsigma m_\alpha'(0)f(0) + \varsigma \int_0^1 f(x)m_\alpha''(x)\, dx.
\end{aligned}$$

Integrating by parts also yields

$$\int_0^1 \bar{u}_\alpha(x)f'(x)m_\alpha(x)\, dx = \bar{u}_\alpha(1)m_\alpha(1)f(1) - \bar{u}_\alpha(0)m_\alpha(0)f(0) - \int_0^1 f(x)(\bar{u}_\alpha m_\alpha)'(x)\, dx.$$

So, the adjoint condition (MN1) can be written

$$\begin{aligned}
-\alpha f(x_0) = &-\varsigma m_\alpha'(1)f(1) + \varsigma m_\alpha'(0)f(0) + \varsigma \int_0^1 f(x)m_\alpha''(x)\, dx \\
&- \bar{u}_\alpha(1)m_\alpha(1)f(1) + \bar{u}_\alpha(0)m_\alpha(0)f(0) + \int_0^1 f(x)(\bar{u}_\alpha m_\alpha)'(x)\, dx \\
&- \alpha \int_0^1 f(x)m_\alpha(x)\, dx \\
= &\left[\varsigma m_\alpha'(0) + \bar{u}_\alpha(0)m_\alpha(0)\right] f(0) - \left[\varsigma m_\alpha'(1) + \bar{u}_\alpha(1)m_\alpha(1)\right] f(1) \\
&\int_0^1 \left(\varsigma m_\alpha''(x) + (\bar{u}_\alpha m_\alpha)'(x) - \alpha m_\alpha(x)\right) f(x)\, dx.
\end{aligned}$$

Since this must hold for every $f \in \mathscr{D}(L)$, it follows that

$$m_\alpha''(x) + \varsigma^{-1}(\bar{u}_\alpha m_\alpha)'(x) - \varsigma^{-1}\alpha m_\alpha(x) = 0, \qquad \forall x \in [0,1]. \tag{C.2}$$

Now, assuming that $\bar{u}_\alpha$ is differentiable, we have

$$(\bar{u}_\alpha m_\alpha)' = \bar{u}_\alpha m_\alpha' + \bar{u}_\alpha' m_\alpha.$$

So, (C.2) becomes the second-order linear homogeneous ordinary differential equation

$$m_\alpha''(x) + \frac{2}{\sigma^2}\bar{u}_\alpha(x)m_\alpha'(x) + \frac{2}{\sigma^2}\left(\bar{u}_\alpha'(x) - \alpha\right) m_\alpha(x) = 0, \qquad \forall x \in [0,1]. \tag{C.3}$$

C.3 Lower Hemicontinuity of Feasibility Correspondence

The following theorem appears as Theorem 4.1 in Chapter 2 of Coddington and Levinson (1955), with some minor changes in notation.

Theorem C.4. *Let D be a domain of (x, y) space, I_α the domain $|\alpha - \alpha_0| < c, c > 0$, and D_α the set of all (x, y, α) satisfying $(x, y) \in D, \alpha \in I_\alpha$. Suppose f is a continuous function on D_α bounded by a constant M there. For $\alpha = \alpha_0$ let*

$$y' = f(x, y, \alpha), \qquad y(\tau) = \xi \tag{C.5}$$

have a unique solution m_0 on the interval $[a, b]$ where $\tau \in [a, b]$. Then there exists a $\delta > 0$ such that, for any fixed α satisfying $|\alpha - \alpha_0| < \delta$, every solution m_α of (C.5) exists over $[a, b]$ and

$$m_\alpha \to m_0 \quad \text{as} \quad \alpha \to \alpha_0$$

uniformly over $[a, b]$.

Remark C.6. Though (C.5) need not have a unique solution for $\alpha \neq \alpha_0$, its solutions are nevertheless continuous in α at α_0. ◇

So, consider the differential equation

$$m''_\alpha(x) + \frac{2}{\sigma^2}\bar{u}_\alpha(x)m'_\alpha(x) + \frac{2}{\sigma^2}\left(\bar{u}'_\alpha(x) - \alpha\right)m_\alpha(x) = 0, \qquad \forall x \in [0, 1]. \tag{C.7}$$

which we can write in the form

$$m''_\alpha(x) = \gamma(x)m'_\alpha(x) + \lambda_\alpha(x)m_\alpha(x), \qquad \forall x \in [0, 1],$$

where

$$\gamma(x) = -\frac{2}{\sigma^2}\bar{u}_\alpha(x) \qquad \text{and} \qquad \lambda_\alpha(x) = -\frac{2}{\sigma^2}\left(\bar{u}'_\alpha(x) - \alpha\right), \qquad \forall x \in [0, 1].$$

By then letting

$$y_1 \coloneqq m_\alpha, \qquad \text{and} \qquad y_2 \coloneqq m'_\alpha,$$

we obtain the following system of two first-order linear equations:

$$\begin{aligned} y'_1 &= y_2, \\ y'_2 &= \gamma(x)y_2 + \lambda_\alpha(x)y_1. \end{aligned}$$

Now, by writing $y = (y_1, y_2)$, we can express the above system in form

$$y' = (y_2, \gamma y_2 + \lambda_\alpha y_1) = f(x, y, \alpha).$$

So, taking $\alpha_0 = 0$ and $c = 1$ in the statement of Theorem C.4, and assuming that $\bar{u}_\alpha \in C^1(E^\circ)$, it is then clear that f is continuous and bounded on the set D_α (which we may assume is some compact subset of $E \times \mathbb{R}^+ \times [0, 1]$). Our derivation of the long-term average density m_0 further shows that the differential equation $y' = f(x, y, 0)$ has the unique solution m_0 on the interval $E = [0, 1]$, where the initial condition

$$y(\tau) = y(x_0) = (m_0(x_0), m_0'(x_0)) = (\xi_1, \xi_2) = \xi$$

can be considered arbitrary. Thus, the hypotheses of Theorem C.4 are satisfied.

D List of Abbreviations and Symbols

The following table provides a guide to the notation used throughout the main body of this manuscript. Occasionally, this notation may differ from that which appears in the Appendix. We make an effort to indicate when this is the case, and we furthermore attempt to keep notation consistent within each particular section of the Appendix.

$\mathbb{R}^+$	the set of nonnegative real numbers; i.e., $\mathbb{R}^+ := [0, +\infty)$
$\mathbb{N}$	the set of positive integers; i.e., $\mathbb{N} := \{1, 2, 3, \dots\}$
$\mathcal{B}(S)$	the Borel σ-algebra on the set S
$\mathcal{M}(S)$	the space of positive finite measures on the measurable space $(S, \mathcal{B}(S))$
$\mathcal{M}_s(S)$	the space of signed finite measures on the measurable space $(S, \mathcal{B}(S))$
$\mathcal{P}(S)$	the space of probability measures on the measurable space $(S, \mathcal{B}(S))$
Λ	a $\mathcal{P}(G)$-valued process; i.e., $\Lambda := \{\Lambda_t : t \in \mathbb{R}^+\}$
$\mathscr{D}(L)$	the domain of the linear operator L
$\mathscr{R}(L)$	the range of the operator L
$M(S)$	the space of Borel-measurable functions on the measurable space $(S, \mathcal{B}(S))$
$C(S)$	the space of continuous, real-valued functions on the set S
$\bar{C}(S)$	the space of continuous, bounded, real-valued functions on the set S
$\widehat{C}(S)$	the space of continuous, real-valued functions on the set S that vanish at infinity
$C_c(S)$	the space of continuous, real-valued functions on the set S with compact support
$c(x, u)$	the cost rate function c evaluated at $(x, u) \in E \times G$
$c_1(x, u)$	the budget rate function c_1 evaluated at $(x, u) \in E \times G$
$\mathbb{E}[X]$	the expectation of the random variable X
I_S	the indicator function for the set S
$\mathscr{M}$	the set of relaxed solutions (X, Λ) to the controlled martingale problem for L
$J(X, \Lambda)$	the long-term average expected cost of a solution $(X, \Lambda) \in \mathscr{M}$
$J_\alpha(X, \Lambda; \nu_0)$	the α-discounted expected cost of a solution $(X, \Lambda) \in \mathscr{M}$ with $X_0 \sim \nu_0$
$\mathfrak{M}$	the set of feasible measures for the long-term average linear program
$\mathfrak{M}_\alpha$	the set of feasible measures for the α-discounted linear program (for $\alpha > 0$)
$\hat{\mathfrak{M}}_\alpha$	the set of feasible measures for the α-normalized linear program (for $\alpha > 0$)
$\mathfrak{M}^*$	the set of optimal measures for the long-term average linear program
$\mathfrak{M}^*_\alpha$	the set of optimal measures for the α-discounted linear program (for $\alpha > 0$)
$\hat{\mathfrak{M}}^*_\alpha$	the set of optimal measures for the α-normalized linear program (for $\alpha > 0$)
ρ	a generic long-term average expected occupation measure
μ_α	a generic α-discounted expected occupation measure (for $\alpha > 0$)
$\hat{\mu}_\alpha$	a generic α-normalized expected occupation measure (for $\alpha > 0$)
J^*	the optimal value of the long-term average linear program
J^*_α	the optimal value of the α-discounted linear program (for $\alpha > 0$)
$\hat{J}^*_\alpha$	the optimal value of the α-normalized linear program (for $\alpha > 0$)
φ	a generic correspondence
π	the Prohorov metric on the space $\mathcal{P}(E \times G)$