

Convergence of Finite Element Methods for Singular Stochastic Control

Martin G. Vieten (University of Wisconsin-Milwaukee, mgvieten@uwm.edu)
Richard H. Stockbridge (University of Wisconsin-Milwaukee)

Abstract: A numerical method is proposed for a class of stochastic control problems including singular behavior. This method solves an infinite-dimensional linear program equivalent to the stochastic control problem using a finite element type approximation, which results in a solvable finite-dimensional program. The discretization scheme as well as the necessary assumptions are discussed, and a detailed convergence analysis for the discretization scheme is given. Its performance is illustrated by two examples featuring a long-term average cost criterion.

Keywords: singular stochastic control, finite element method, linear programming, relaxed controls

AMS classification: 93E20, 93E25

1 Introduction

1.1 Motivation and Literature

This paper considers singular stochastic control problems for a process X whose dynamics are initially specified by a stochastic differential equation (SDE)

$$dX_t = b(X_t, u_t)dt + \sigma(X_t, u_t)dW_t + d\xi_t, \quad X_0 = x_0, \quad (1.1)$$

where W is a Brownian motion process and ξ is another stochastic process that evolves singularly in time. The process u represents the control influencing the evolution of X . Given two cost functions \tilde{c}_0 and \tilde{c}_1 , u has to be chosen from a set of admissible controls in such a way that it minimizes either a long-term average cost criterion

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t \tilde{c}_0(X_s, u_s) ds + \int_0^t \tilde{c}_1(X_s, u_s) d\xi_s \right] \quad (1.2)$$

or a discounted infinite horizon cost criterion

$$\mathbb{E} \left[\int_0^\infty e^{-\alpha s} \tilde{c}_0(X_s, u_s) ds + \int_0^\infty e^{-\alpha s} \tilde{c}_1(X_s, u_s) d\xi_s \right], \quad (1.3)$$

for some discounting rate $\alpha > 0$. Such control problems are considered in a relaxed sense by using a martingale problem formulation involving the infinitesimal generators of X , and an equivalent infinite-dimensional linear program for the expected occupation measures of both the process X and the control u . Approximate solutions to this linear program are attained by discretizing the infinite-dimensional constraint space of functions using a finite element approach, and discrete approximations of the expected occupation measures. The ϵ -optimality of approximate solutions is shown and the method is applied to two example problems.

The classical analytic approach to stochastic control problems is given by methods based on the dynamic programming principle, as presented in Fleming and Rishel [6] or Fleming and Soner [7]. Central to these methods is the solution of the so-called Hamilton-Jacobi-Bellman (HJB) equation. Numerical methods can be derived by solving a control problem for an approximate, discrete Markov chain, as extensively discussed in Kushner and Dupuis [15], or by using discrete methods to approximate the solution to the HJB equation, frequently considering viscosity solutions. An example is given in Kumar and Muthuraman [12]. Another numerical technique using dynamic programming was analyzed in Anselmi et. al. [1].

As an alternative, linear programming approaches have been instrumental in the analytic treatment of various stochastic control problems. The first example is given in Manne [18], where an ergodic Markov chain for an inventory problem under long-term average costs is analyzed. Bhatt and Borkar [2] as well as Kurtz and Stockbridge [13] investigated the linear programming approach for solutions of controlled martingale problems using long-term average and discounted cost criteria for infinite horizon problems, as well as finite horizon and first exit problems for absolutely continuous control. Taksar [22] establishes equivalence between

a linear program and a stochastic control problem for a multi-dimensional diffusion with singular control. Jump diffusions of Levy-Ito type are considered by Serrano [21].

To provide an alternative to numerical techniques based on the dynamic programming principle, the linear programming approach has been exploited using various discretization techniques. A very general setting can be found in Mendiondo and Stockbridge [19]. Moment-based approaches have been used in a line of publications, as can be seen in Helmes et. al. [9] and Lasserre and Prieto-Rumeau [16]. Recent research by Kaczmarek et. al. [11] and Rus [20] has been investigating a novel approximation technique for the linear programming formulation by borrowing ideas from the finite element method used for solving partial differential equations. A discretization of the occupation measures (by discretizing their densities) and the linear constraints with a finite set of basis functions gives a solvable finite-dimensional linear program. Kaczmarek et. al. [11] indicated that a finite element discretization approach may outperform Markov chain approximation methods as well as a finite difference approximation to the Hamilton-Jacobi-Bellman equation stemming from the dynamic programming approach. However, no analytic treatment of the convergence properties was provided.

The present paper closes this gap by providing a modified finite element based approximation scheme for which convergence of the computed cost criterion values can be guaranteed. To this end, the approximation scheme is split up in several steps which either deal with the discretization of the measures or the constraints. The separate steps are set up in such a way that convergence of the discrete optimal solutions to the analytic optimal solution can be proven. The proofs are, on one hand, based on the concept of weak convergence of measures, and on a detailed analysis of discretized approximations of the measures on the other hand.

The paper is structured as follows. The next subsection presents the notation and formally introduces the linear programming formulation for singular stochastic control problems, along with a review of important results from the literature. The approximation scheme is discussed in Section 2. Then, we provide the convergence proof for this scheme in Section 3 and illustrate the performance of the numerical method on two examples in Section 4. A short outlook on possible research directions concludes this paper. Additional proofs needed to prove the results from Section 3 are given in Appendix A.

1.2 Notation and Formalism

The natural numbers are denoted by \mathbb{N} , and the non-negative integers are $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. The symbol used for the real numbers is \mathbb{R} , and that for the non-negative real numbers is \mathbb{R}_+ . The space of n -dimensional vectors is \mathbb{R}^n , and the space of n by m matrices is $\mathbb{R}^{n \times m}$.

The set of continuous functions on a topological space S is denoted by $C(S)$. The set of twice differentiable functions on S is denoted by $C^2(S)$, while its subset of twice differentiable functions with compact support is referred to by $C_c^2(S)$. The space of uniformly continuous, bounded functions is denoted by $C_b^u(S)$. On a function space, $\|\cdot\|_\infty$ refers to the uniform norm of functions. On \mathbb{R}^n , $\|\cdot\|_\infty$ refers to the maximum norm of vector components, while on $\mathbb{R}^{n \times m}$, it refers to the maximum absolute row sum norm. The space of Lebesgue integrable functions is $L^1(S)$. For any given function f , let $f^+ : E \ni x \mapsto f^+(x) := \max(f(x), 0)$ be the positive part of a function f .

In terms of measurable spaces, we use $\mathcal{B}(S)$ to describe the σ -algebra of Borel sets on a topological space S . Given a measurable space (Ω, \mathcal{F}) , the set of probability measures on Ω is $\mathcal{P}(\Omega)$, while the set of finite Borel measures is denoted by $\mathcal{M}(\Omega)$. The symbol $\delta_{\{s\}}$ denotes the Dirac measure on $s \in S$. When using the differential dx as an integrator, it is understood that this refers to integration by Lebesgue measure. When we explicitly refer to the Lebesgue measure, we use the symbol λ . A Brownian motion process is denoted by the symbol W .

Consider the SDE given by (1.1). We assume that $X_t \in E = [e_l, e_r]$, with $\infty < e_l < e_r < \infty$, and $u_t \in U = [u_l, u_r]$, with $\infty < u_l < u_r < \infty$, for all $t \geq 0$. E and U are called the state space and control space, respectively. The coefficient functions $b : E \times U \mapsto \mathbb{R}$ and $\sigma : E \times U \mapsto \mathbb{R}_+ - \{0\}$ are called the drift and diffusion functions. They are assumed to be continuous. The process ξ is a singular stochastic process stemming from the behavior of X at the boundaries of the state space e_l and e_r , and is given by either a reflection, a jump or a combination of both. The infinitesimal generators of a process solving (refintroduction:sde) are $\tilde{A} : C_c^2(E) \mapsto C(E \times U)$, called the continuous generator, and $B : C_c^2(E) \mapsto C(E \times U)$, called the singular generator. For $f \in C_c^2(E)$, \tilde{A} is defined by $\tilde{A}f(x, u) = b(x, u)f'(x) + \sigma^2(x, u)f''(x)$.

B is defined by either of

$$Bf(x, u) = \pm f'(x) \quad \text{or} \quad Bf(x, u) = f(x + u) - f(x). \quad (1.4)$$

The first form of B models a reflection process (+ forcing a reflection to the right and $-$ forcing a reflection to the left) and the second form models a jump process jumping from x to $x + u$. With these generators, a specification of the dynamics that requires

$$f(X_t) - f(x_0) - \int_0^t \tilde{A}f(X_s, u_s) ds - \int_0^t Bf(X_s, u_s) d\xi_s \quad (1.5)$$

to be a martingale for all $f \in C_c^2$ is equivalent to (1.1) in terms of weak solutions. Hence, the values of the cost criteria determined by (1.2) and (1.3) remain identical. The following relaxed formulation of (1.5) is better suited for the purpose of stochastic control.

Definition 1.1. Let X be a stochastic process with state space E , let Λ be a stochastic process taking values in $\mathcal{P}(U)$, and let Γ be a random variable taking values in the space of measures on $([0, \infty) \times E \times U)$, with $\Gamma([0, t] \times E \times U) \in \mathcal{M}([0, t] \times E \times U)$ for all t . The triplet (X, Λ, Γ) is a relaxed solution to the singular, controlled martingale problem for (\tilde{A}, B) if there is a filtration $\{\mathcal{F}_t\}_{t \geq 0}$ such that X , Λ and Γ are \mathcal{F}_t -progressively measurable and

$$f(X_t) - f(x_0) - \int_0^t \int_U \tilde{A}f(X_s, u) \Lambda_s(du) ds - \int_{[0, t] \times E \times U} Bf(x, u) \Gamma(ds \times dx \times du)$$

is an $\{\mathcal{F}_t\}_{t \geq 0}$ -martingale for all $f \in C_c^2(E)$.

The relaxation is given by the fact that the control is no longer represented by a process u , but is encoded in the random measures Λ and Γ . Assume that the cost functions \tilde{c}_0 and \tilde{c}_1 are continuous and non-negative. The cost criteria for a relaxed solution of the singular, controlled martingale problem are

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t \int_U \tilde{c}_0(X_s, u) \Lambda_s(du) ds + \int_{[0, t] \times E \times U} \tilde{c}_1(x, u) \Gamma(ds \times dx \times du) \right], \quad (1.6)$$

for the long-term average cost criterion, and for $\alpha > 0$,

$$\mathbb{E} \left[\int_0^\infty \int_U e^{-\alpha s} \tilde{c}_0(X_s, u) \Lambda_s(du) ds + \int_{[0, \infty) \times E \times U} e^{-\alpha s} \tilde{c}_1(x, u) \Gamma(ds \times dx \times du) \right], \quad (1.7)$$

for the infinite horizon discounted cost criterion. A stochastic control problem given by (1.5) together with (1.6) or (1.7) can be reformulated as an infinite dimensional linear program. To this end, we set

$$c_0(x, u) = \begin{cases} \tilde{c}_0(x, u) & \text{if } \alpha = 0 \\ \tilde{c}_0(x, u)/\alpha & \text{if } \alpha > 0 \end{cases} \quad \text{and} \quad c_1(x, u) = \begin{cases} \tilde{c}_1(x, u) & \text{if } \alpha = 0 \\ \tilde{c}_1(x, u)/\alpha & \text{if } \alpha > 0. \end{cases}$$

Furthermore, for $\alpha \geq 0$ define the operator $A : C_c^2(E) \mapsto C(E \times U)$ by

$$Af(x, u) = \tilde{A}f(x, u) - \alpha f(x) \quad (1.8)$$

and the functional $Rf = -\alpha f(x_0)$, x_0 being the starting point of the diffusion.

Definition 1.2. The infinite-dimensional linear program for a singular stochastic control problem is given by

$$\begin{aligned} & \text{Minimize} \quad \int_{E \times U} c_0 d\mu_0 + \int_{E \times U} c_1 d\mu_1 \\ & \text{Subject to} \quad \begin{cases} \int_{E \times U} Af d\mu_0 + \int_{E \times U} Bf d\mu_1 = Rf \quad \forall f \in C_c^2(E) \\ \mu_0 \in \mathcal{P}(E \times U) \\ \mu_1 \in \mathcal{M}(E \times U). \end{cases} \end{aligned}$$

The measures μ_0 and μ_1 are the expected occupation measures of X and Γ . We frequently consider the measures on $(E, \mathcal{B}(E))$ given by $\mu_{0,E}(\cdot) = \mu_0(\cdot \times U)$ and $\mu_{1,E}(\cdot) = \mu_1(\cdot \times U)$. We refer to these measures as the state space marginals of μ_0 and μ_1 , respectively. The properties of such linear programs and their relation to stochastic control problems for singular, controlled martingale problems are stated in Theorem 1.4. These results use the notion of a regular conditional probability defined as follows.

Definition 1.3. Let $(E \times U, \mathcal{B}(E \times U), \mu)$ be a measure space, and let $P : E \times U \ni (x, u) \mapsto x \in E$ be the projection map onto E . Let μ_E be the distribution of P , which is identical to the state space marginal of μ . A map $\eta : \mathcal{B}(U) \times E \mapsto [0, 1]$ is called a regular conditional probability if

- i) for each $x \in E$, $\eta(\cdot, x) : \mathcal{B}(U) \mapsto [0, 1]$ is a probability measure,
- ii) for each $V \in \mathcal{B}(U)$, $\eta(V, \cdot) : E \mapsto [0, 1]$ is a measurable function, and
- iii) for all $V \in \mathcal{B}(U)$ and all $F \in \mathcal{B}(E)$ we have

$$\mu(F \times V) = \int_F \eta(V, x) \mu_E(dx).$$

Theorem 1.4. The problem of minimizing either the long-term average cost criterion of (1.6) or the infinite horizon discounted cost criterion of (1.7) over the set of all relaxed solutions (X, Λ, Γ) to the singular, controlled martingale problem for (\bar{A}, B) is equivalent to the linear program stated in (1.2). Moreover, there exists an optimal solution (μ_0^*, μ_1^*) . Let η_0^* and η_1^* be the regular conditional probabilities of μ_0^* and μ_1^* with respect to their state space marginals. Then an optimal relaxed control is given in feedback form by $\Lambda_t^* = \eta_0^*(\cdot, X_t^*)$ and $\Gamma^*(ds \times dx \times du) = \eta_1^*(du, x) \tilde{\Gamma}^*(ds \times dx)$ for a random measure $\tilde{\Gamma}^*$ on $[0, \infty) \times E$, where $(X^*, \Lambda^*, \Gamma^*)$ is a relaxed solution to the singular, controlled martingale problem for (\bar{A}, B) having occupation measures (μ_0^*, μ_1^*) .

Proof. See Kurtz and Stockbridge [14], Theorem 2.1 and Theorem 3.3, respectively. \square

By this result, it suffices to find optimal solutions to the infinite linear program when solving a singular stochastic control problem, and approximate solutions to the linear program serve as approximate solutions to the control problem. Section 2.1 presents how we discretize the infinite dimensional linear program to a computationally attainable formulation, which is the basis for the numerical technique used in this paper. The analysis of this discretization scheme relies in part on the notion of weak convergence of *finite* measures which is defined next. Let S be a measurable space in the following, equipped with a topology.

Definition 1.5. Consider a sequence of finite measures $\{\mu_n\}_{n \in \mathbb{N}}$ and another finite measure μ on S . We say that μ_n converges weakly to μ , in symbol $\mu_n \Rightarrow \mu$, if for all $f \in C_b^u(S)$

$$\int_S f(x) \mu_n(dx) \rightarrow \int_S f(x) \mu(dx) \quad \text{as } n \rightarrow \infty$$

holds.

Note that we are considering finite measures, and not necessarily probability measures. In particular, we could encounter a situation where the sequence of numbers $\{\mu_n(S)\}_{n \in \mathbb{N}}$ is unbounded. This differs from ‘classical’ considerations of weak convergence, which for example can be found in Billingsley [3]. However, Bogachev [4] (see Chapter 8 in Volume 2) offers a discussion of the concept of weak convergence in this more general case. Central to our purposes is Theorem 1.9, which states sufficient conditions for the existence of weakly converging subsequences when considering sequences of finite measures, based on the following two concepts.

Definition 1.6. A sequence of finite measures $\{\mu_n\}_{n \in \mathbb{N}}$ on S is called tight if for each $\epsilon > 0$, there is a compact set K_ϵ in S such that

$$\mu_n(K_\epsilon^C) < \epsilon$$

holds for all $n \in \mathbb{N}$.

Remark 1.7. If S is compact, any sequence of finite measures on S is tight.

Definition 1.8. A sequence of finite measures $\{\mu_n\}_{n \in \mathbb{N}}$ on S is called *uniformly bounded* if for some $l \geq 0$, $\mu_n(S) \leq l$ holds for all $n \in \mathbb{N}$.

If a sequence of finite measures on S is tight and uniformly bounded, the existence of convergent subsequences is guaranteed by the following result.

Theorem 1.9. Let $\{\mu_n\}_{n \in \mathbb{N}}$ be a sequence of finite measures on S . Then, the following are equivalent.

- i) $\{\mu_n\}_{n \in \mathbb{N}}$ contains a weakly convergent subsequence,
- ii) $\{\mu_n\}_{n \in \mathbb{N}}$ is tight and uniformly bounded.

Proof. See Bogachev [4], Theorem 8.6.2. □

2 Approximation

We begin the presentation of the proposed method by describing the discretization scheme. Then, we discuss the assumption being necessary for the convergence of the method.

2.1 Discretization

The proposed numerical technique is based on a discretization of the infinite-dimensional linear program in three steps. First, we introduce a limit on the full mass of the measure μ_1 . Then, we restrict the number of constraint functions. Thirdly, we introduce discrete versions of the measures. In the process, several assumptions on the measure μ_0 are made. For the sake of exposition, we elaborate on these assumptions separately in Section 2.2.

Since the discretization brings forth several distinct sets of measures, we define the cost criterion using the following, general formulation.

$$J : \mathcal{P}(E \times U) \times \mathcal{M}(E \times U) \ni (\mu_0, \mu_1) \mapsto J(\mu_0, \mu_1) = \int_{E \times U} c_0 d\mu_0 + \int_{E \times U} c_1 d\mu_1 \in \mathbb{R}_+$$

We choose to consider $C_c^2(E)$ as a normed space in the right sense. Set $\|f\|_{\mathcal{D}} = \|f\|_{\infty} + \|f'\|_{\infty} + \|f''\|_{\infty}$ and define $\mathcal{D}_{\infty} = (C_c^2(E), \|\cdot\|_{\mathcal{D}})$ to designate that we consider $C_c^2(E)$ to be a specific normed space. Set

$$\mathcal{M}_{\infty} = \{(\mu_0, \mu_1) \in \mathcal{P}(E \times U) \times \mathcal{M}(E \times U) : \int A f d\mu_0 + \int B f d\mu_1 = R f \quad \forall f \in \mathcal{D}_{\infty}\}.$$

For analytical purposes, we introduce an upper bound on $\mu_1(E \times U)$. For $l > 0$ define

$$\mathcal{M}_{\infty}^l = \{(\mu_0, \mu_1) \in \mathcal{M}_{\infty} : \mu_1(E \times U) \leq l\}. \quad (2.1)$$

Remark 2.1. As l increases, more measures of \mathcal{M}_{∞} will lie in \mathcal{M}_{∞}^l . For l large enough, the optimal solution will lie in \mathcal{M}_{∞}^l , as we have that $\mu_1^* \in \mathcal{M}(E \times U)$ and hence $\mu_1^*(E) < \infty$.

Definition 2.2. The l -bounded infinite-dimensional linear program is given by

$$\min \{J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_{\infty}^l\}.$$

The set \mathcal{M}_{∞}^l features an infinite set of constraints given by all $f \in \mathcal{D}_{\infty}$ and measures μ_0 and μ_1 having an infinite number of degrees of freedom. First, we discretize the set of constraints using B-spline basis functions. To construct these basis functions, fix $q \in \mathbb{N}$ and consider a finite set of pointwise distinct grid points $\{e_k\}_{k=-3}^{q+3}$ in E , with $e_0 = e_l$, $e_r = e_q$ and $e_k < e_{k+1}$ for $k = -3, -2, \dots, q+2$.

Definition 2.3. The set of cubic B-spline basis functions for a grid $\{e_k\}_{k=-3}^{q+3}$ is defined on \mathbb{R} by

$$f_k(x) = (e_{k+4} - e_k) \sum_{i=k}^{k+4} \frac{[(e_i - x)^3]^+}{\Psi'_k(e_i)}, \quad k = -3, -2, \dots, q-1,$$

where

$$\Psi_k(x) = \prod_{i=k}^{k+4} (x - e_i), \quad k = -3, -2, \dots, q-1.$$

An analysis of these basis function is given in de Boor [5]. Provided that

$$\max_{k=-3, \dots, q+2} (e_{k+1} - e_k) \rightarrow 0 \quad \text{and} \quad \frac{\max_{k=-3, \dots, q+2} (e_{k+1} - e_j)}{\min_{k=-3, \dots, q+2} (e_{k+1} - e_j)} \rightarrow 1$$

as $n \rightarrow \infty$, Theorem 1 of Hall and Meyer [8] holds and the following statement can be shown.

Proposition 2.4. The normed space \mathcal{D}_∞ is separable and a countable basis $\{f_k\}_{k \in \mathbb{N}}$ is given by the cubic B-splines basis functions.

For fixed $\tilde{q} \in \mathbb{N}$, define a grid using the dyadic partition of E given by

$$e_k = e_l + \frac{e_r - e_l}{2^{\tilde{q}}} \cdot k, \quad k = -3 \dots, 2^{\tilde{q}} + 3.$$

and consider the $n := 2^{\tilde{q}} + 2$ B-spline basis functions $\{f_k\}_{k=1}^n$ on this grid. This allows us to define

$$\begin{aligned} \mathcal{M}_n = \{(\mu_0, \mu_1) \in \mathcal{P}(E \times U) \times \mathcal{M}(E \times U) : \\ \int A f_k d\mu_0 + \int B f_k d\mu_1 = R f_k, \quad k = 1, \dots, n\} \end{aligned}$$

and we can define \mathcal{M}_n^l in a similar manner using the mass restriction on μ_1 as seen in (2.1).

Definition 2.5. The l -bounded (n, ∞) -dimensional linear program is given by

$$\min \{J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_n^l\}.$$

Next, we discretize the measures. Theorem 1.4 reveals that it is sufficient to regard feedback controls which can be represented by regular conditional probabilities. In particular, this result states that we can consider measures (μ_0, μ_1) which can be decomposed according to $\mu_0(dx \times du) = \eta_0(dx, u)\mu_{0,E}(dx)$ and $\mu_1(dx \times du) = \eta_1(dx, u)\mu_{1,E}(dx)$ for two regular conditional probabilities η_0 and η_1 . We furthermore assume that, first, for any interval or singleton $V \subset U$, $x \mapsto \eta_0(V, x)$ is continuous almost everywhere with respect to Lebesgue measure, second, that $\mu_{0,E}$ has a density p with respect to Lebesgue measure and third, that p satisfies the constraint that $\lambda(\{x : p(x) = 0\}) = 0$. In other words, p must only be equal to zero on a set of Lebesgue measure 0. The particulars of these assumptions are discussed in Section 2.2, and we continue here with the description of the approximation scheme.

Define a sequence k_m as follows. As c_0 , b and σ are continuous over a compact set, for all $m \in \mathbb{N}$, there is a $\delta_m > 0$ such that for all $u, v \in U$ with $|u - v| \leq \delta_m$, it is true that

$$\max \left\{ |c_0(x, u) - c_0(x, v)|, |b(x, u) - b(x, v)|, \left| \frac{1}{2} \sigma^2(x, u) - \frac{1}{2} \sigma^2(x, v) \right| \right\} \leq \frac{1}{2^{m+1}}, \quad (2.2)$$

uniformly in x . Set k_m to be the smallest integer such that $\frac{u_r - u_l}{2^{k_m}} \leq \delta_m$. The parameter k_m controls the discretization of the control space U , and the specific choice enables an accurate approximate integration of the cost function c_0 and the functions $A f_k$ against the relaxed control η_0 in the convergence proof of Section 3. So, define

$$U^{(m)} = \{u_j = u_l + \frac{u_r - u_l}{2^{k_m}} \cdot j, j = 0, \dots, 2^{k_m}\}. \quad (2.3)$$

Similarly, we set

$$E^{(m)} = \{e_j = e_l + \frac{e_r - e_l}{2^m} \cdot j, j = 0, \dots, 2^m\}. \quad (2.4)$$

The union of these sets over all $m \in \mathbb{N}$ is dense in the control space and state space, respectively. The number m is called the discretization level. It determines the degrees of freedom of the discrete measures $\hat{\mu}_0$ (approximating μ_0) and $\hat{\mu}_1$ (approximating μ_1), which are defined as follows.

First, we approximate the density p of $\mu_{0,E}$. Choose a countable basis of $L^1(E)$, say $\{p_n\}_{n \in \mathbb{N}_0}$, given by indicator functions over subintervals of E . We truncate this basis to p_0, \dots, p_{2^m-1} (given by the indicator functions of the intervals of length $1/2^m$, compare (2.4)) to approximate the density p by

$$\hat{p}_m(x) = \sum_{j=0}^{2^m-1} \gamma_j p_j(x) \quad (2.5)$$

where $\gamma_j \in \mathbb{R}_+$, $j = 0, \dots, 2^m - 1$ are weights to be chosen under the constraint that $\int_E \hat{p}_m(x) dx = 1$. Set $E_j = [x_j, x_{j+1}]$ for $j = 0, 1, \dots, 2^m - 2$ and $E_{2^m-1} = [x_{2^m-1}, x_{2^m}]$ to define

$$\hat{\eta}_{0,m}(V, x) = \sum_{j=0}^{2^m-1} \sum_{i=0}^{2^{k_m}} \beta_{j,i} I_{E_j}(x) \delta_{\{u_i\}}(V), \quad (2.6)$$

where $\beta_{j,i} \in \mathbb{R}_+$, $j = 0, \dots, m-1, i = 0, \dots, k_m$ are weights to be chosen under the constraint that $\sum_{i=0}^{2^{k_m}-1} \beta_{j,i} = 1$ for $j = 0, \dots, m-1$. We approximate η_0 using (2.6), which means that this relaxed control is approximated by point masses in U -‘direction’ and piecewise constant in E -‘direction’. Then, we set $\hat{\mu}_{0,m}(du \times dx) = \hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx$.

To approximate the singular occupation measure μ_1 , we use that the process is only showing singular behavior at e_l and e_r . Thus, if we introduce the regular conditional probability η_1 and write $\mu_1(dx \times du) = \eta_1(du, x) \mu_{1,E}(dx)$, and for $F \in \mathcal{B}(E)$, we have for $F \in \mathcal{B}(E)$

$$\mu_{1,E}(F) = w_1 \delta_{\{e_l\}}(F) + w_2 \delta_{\{e_r\}}(F) \quad (2.7)$$

with $w_1, w_2 \in \mathbb{R}_+$. We approximate the relaxed control η_1 by

$$\hat{\eta}_{1,m}(V, e_l) = \sum_{i=0}^{2^{k_m}} \zeta_{1,i} \delta_{\{u_i\}}(V), \quad \hat{\eta}_{1,m}(V, e_r) = \sum_{i=0}^{2^{k_m}} \zeta_{2,i} \delta_{\{u_i\}}(V) \quad (2.8)$$

with $\sum_{i=0}^{2^{k_m}} \zeta_{j,i} = 1$ for $j = 1, 2$. So, we have $\hat{\mu}_{1,m}(dx \times du) = \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx)$. In summary, we consider measures of the form

$$(\hat{\mu}_{0,m}, \hat{\mu}_{1,m})(dx \times du) = (\hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx, \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx))$$

and we introduce the notation

$$\begin{aligned} \mathcal{M}_{n,m}^l &= \{(\mu_{0,m}, \mu_{1,m}) \in \mathcal{M}_n^l : (\mu_{0,m}, \mu_{1,m})(du, dx) \\ &= (\hat{\eta}_m(du, x) \hat{p}_m(x) dx, \hat{\eta}_m(du, x) \mu_{1,E}(dx))\}. \end{aligned}$$

This finalizes the discretization of the measures and leaves us with the following linear program.

Definition 2.6. *The l -bounded (n, m) -dimensional linear program is given by*

$$\inf \{J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_{n,m}^l\}.$$

This linear program is linear in the coefficients given by the products $\beta_{j,i} \cdot \gamma_j$ and $\zeta_{j,i} \cdot \alpha_j$, and the cost functional can as well be expressed as a linear combination of these coefficients.

Up to this point, we introduced four sets of measures, \mathcal{M}_∞ , \mathcal{M}_∞^l , \mathcal{M}_n^l and $\mathcal{M}_{n,m}^l$, and we later on will use ϵ -optimal solutions in $\mathcal{M}_{n,m}^l$ to approximate the optimal solution in \mathcal{M}_∞ . However the relations between those sets are $\mathcal{M}_\infty^l \subset \mathcal{M}_\infty$, $\mathcal{M}_n^l \supset \mathcal{M}_\infty^l$ and $\mathcal{M}_{n,m}^l \subset \mathcal{M}_n^l$. As this does not provide a clear nested structure, it has to be carefully analyzed how optimal solutions in these sets relate to each other. This is presented in Section 3.

2.2 Assumptions

Before we move to the presentation of the convergence argument, we elaborate on the assumptions on μ_0 which were made in Section 2.1. These assumptions restrict the set of feasible measures considered in the linear program given by (1.2) to measures which allow the approximation to converge. Albeit technical, the imposed restrictions do not curtail the set of feasible measures beyond what can be considered to be ‘implementable’ solutions, in other words, the set of measures will still be large enough to include any type of control that could be used in a real-world application.

First, we assume that the state space marginal $\mu_{0,E}$ of the expected occupation measures μ_0 has a density p with respect to the Lebesgue measure. As shown in [23], Section II.2, this is guaranteed when certain assumptions on the regular conditional probability η_0 of the continuous occupation measure μ_0 are fulfilled. To be precise, we have to assert that the functions

$$x \mapsto \int_U b(x, u) \eta_0(du, x), \quad x \mapsto \int_U \sigma^2(x, u) \eta_0(du, x) \quad (2.9)$$

are continuous everywhere except for finitely many points in E . On the one hand, this is satisfied for controls of the form given by (2.6), which includes the important class of so-called bang-bang controls. Bang-bang controls put full mass on either of the end points u_l and u_r of the control space U . Usually, when the cost function c_0 does not depend on the control value u , the optimal solution is given by a bang-bang control. If this is not the case, optimal controls are frequently given in the form of a continuous function $v : E \mapsto U$ and a control satisfying $\eta_0(\{v(x)\}, x) = 1$. It is easy to see that in both cases the two functions defined in (2.9) are continuous except for finitely many points.

Secondly, we assume that p must be equal to zero only on a set of Lebesgue measure 0. The analysis in [23], Section II.2 shows that the densities encountered when using both the long-term average cost criterion and the discounted infinite horizon criterion satisfy this assumption. Thirdly, we assume that for any set $V \subset U$ which is either an interval or a singleton, the function $x \mapsto \eta_0(V, x)$ is continuous almost everywhere with respect to Lebesgue measure. This allows us to approximate the function $x \mapsto \eta_0(V, x)$ uniformly by a function which is piecewise constant over intervals, and the approximate function values on these intervals are given by the values of $\eta_0(V, x)$ at the left endpoints of the intervals. This makes the statement of (3.8) of the convergence argument true. Controls of the form given in (2.6) satisfy this requirement. However, if η_0 fulfills $\eta_0(\{v(x)\}, x) = 1$ for some continuous function v on E , we have to assert more regularity on v , according to the following definition.

Definition 2.7. *A continuous function $v : E \mapsto U$ is said to have finitely many modes if there are finitely many points $e_l = \hat{y}_1 < \hat{y}_2 < \dots < \hat{y}_k = e_r$ such that for all $2 \leq i \leq k - 1$, there are points a_i and b_i with $\hat{y}_{i-1} < a_i < \hat{y}_i < b_i < \hat{y}_{i+1}$ and either of the following statements hold:*

- i) *v is strictly increasing on (a_i, \hat{y}_i) and strictly decreasing on (\hat{y}_i, b_i) as well as increasing on $(\hat{y}_{i-1}, \hat{y}_i)$ and decreasing on $(\hat{y}_i, \hat{y}_{i+1})$,*
- ii) *v is strictly decreasing on (a_i, \hat{y}_i) and strictly increasing on (\hat{y}_i, b_i) as well as decreasing on $(\hat{y}_{i-1}, \hat{y}_i)$ and increasing on $(\hat{y}_i, \hat{y}_{i+1})$.*

We assume that v only has finitely many modes in the following. The rationale behind this assumption is as follows. Obviously, $x \mapsto \eta_0(V, x)$ is piecewise constant, either 0 or 1. The fact that v ‘oscillates’ only finitely many times between its modes ensures that $x \mapsto \eta_0(V, x)$ does not switch from 0 to 1 or from 1 to 0 more than finitely many times, and hence it is discontinuous on a set that has measure 0 with respect to Lebesgue measure.

3 Convergence

The first part of this section gives an overview of the convergence argument, illustrating the main ideas of the analysis. The proofs of the propositions and corollary are given in the second part.

3.1 Statement of the main results

The l -bounded (n, m) -dimensional linear program introduced in Section 2.1 is a finite dimensional linear program that can be solved with standard solvers that are available in numerical libraries, and hence optimal solutions are attainable. We proceed to show that the optimal solution to the l -bounded (n, m) -dimensional linear program is an ϵ -optimal solution to the infinite dimensional program for l, n and m large enough. We use the notations

$$\begin{aligned} J^* &= \inf \{ J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_\infty^l \} \\ J_n^* &= \inf \{ J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_n^l \} \\ J_{n,m}^* &= \inf \{ J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_{n,m}^l \}. \end{aligned}$$

For l large enough, J^* is indeed the optimal solution to the unbounded problem, as stated in Remark 2.1, in other words, $J^* = \min \{ J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_\infty \}$

Since an infimum might not be computationally attainable in \mathcal{M}_∞^l and \mathcal{M}_n^l , we withdraw to the slightly relaxed optimization problem of finding an ϵ -optimal solution, in other words, we try to find a pair of measures $(\mu_0^\epsilon, \mu_1^\epsilon) \in \mathcal{M}_\infty^l$ such that

$$J(\mu_0^\epsilon, \mu_1^\epsilon) - J(\mu_0, \mu_1) \leq \epsilon \quad \forall (\mu_0, \mu_1) \in \mathcal{M}_\infty^l.$$

Note that trivially, $J(\mu_0^\epsilon, \mu_1^\epsilon) - J^* \geq 0$. The ϵ -optimality for \mathcal{M}_n^l is defined analogously. The following convergence analysis proves that we can find ϵ -optimal measures in \mathcal{M}_∞^l using the approximation proposed in Section 2. The outline of the proof is as follows. First, it is shown that it suffices to find an ϵ -optimal solution in \mathcal{M}_n^l .

Proposition 3.1. *For each $n \in \mathbb{N}$, assume that $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) \in \mathcal{M}_n^l$ and that for $n \in \mathbb{N}$, $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$ is an ϵ -optimal solution for the l -bounded, (n, ∞) -dimensional linear program. Then, for $\delta > 0$, there exists an $N(\delta)$ such that*

$$J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) - J^* \leq 2\epsilon + \delta.$$

for all $n \geq N(\delta)$.

Next, we establish that ϵ -optimal solutions in \mathcal{M}_n^l can be obtained using the discretization introduced in Section 2.1. The central result reads as follows.

Proposition 3.2. *For $(\mu_0, \mu_1) \in \mathcal{M}_n^l$ and each $\epsilon > 0$, there is an m_0 such that for all $m \geq m_0$ there exists a $(\hat{\mu}_{0,m}, \hat{\mu}_{1,m}) \in \mathcal{M}_{n,m}^l$, with*

$$|J(\mu_0, \mu_1) - J(\hat{\mu}_{0,m}, \hat{\mu}_{1,m})| < \epsilon.$$

This result shows that arbitrary (not necessarily optimal) measures in \mathcal{M}_n^l can be approximated, in terms of their cost criterion, by measures in $\mathcal{M}_{n,m}^l$. Regarding optimal measures, the following statement is an easy consequence from Proposition 3.2.

Corollary 3.3. *For each $m \in \mathbb{N}$, assume that $(\mu_{0,n,m}^*, \mu_{1,n,m}^*) \in \mathcal{M}_{n,m}^l$ and that for $m \in \mathbb{N}$, $(\mu_{0,n,m}^*, \mu_{1,n,m}^*)$ is an optimal solution to the l -bounded, (n, m) -dimensional linear program. Then, the sequence of numbers $\{J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)\}_{m \in \mathbb{N}}$ converges to J_n^* as $m \rightarrow \infty$.*

In conjunction, the preceding results allow us to prove the following theorem.

Theorem 3.4. *For any $\epsilon > 0$, there is an $l > 0$, an $N \in \mathbb{N}$ and an $M \in \mathbb{N}$ such that*

$$|J^* - J_{n,m}^*| < \epsilon$$

holds for all $n \geq N$ and $m \geq M$.

Proof. Pick any $\epsilon > 0$. Choose l large enough such that $J^* = \min \{J(\mu_0, \mu_1) | (\mu_0, \mu_1) \in \mathcal{M}_\infty\}$. Pick $\hat{\epsilon}$ and $\delta > 0$ in such a way that $2\hat{\epsilon} + \delta < \epsilon$. By Proposition 3.1, choose $N \in \mathbb{N}$ large enough such that for all $n \geq N$, an $\hat{\epsilon}$ -optimal solution $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$ to the l -bounded, (n, ∞) -dimensional program is an $2\hat{\epsilon} + \delta$ -optimal solution to the l -bounded infinite-dimensional linear program. Now, using Corollary 3.3, choose $M \in \mathbb{N}$ large enough such that for all $m \geq M$, the optimal solution $(\mu_{0,n,m}^*, \mu_{1,n,m}^*)$ to the l -bounded (n, m) -dimensional linear program is an $\hat{\epsilon}$ -optimal solution to the l -bounded (n, ∞) -dimensional program. But then,

$$|J^* - J_{n,m}^*| \equiv |J^* - J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)| < 2\hat{\epsilon} + \delta < \epsilon$$

□

3.2 Proofs of the main results

The proof of Proposition 3.1 is a rather straight forward application of the theory of weak convergence as introduced in Section 1.2. The proofs of Proposition 3.2 and Corollary 3.3 on the other hand require an in-depth analysis of the approximation properties of the proposed discretization scheme. We begin with the proof of Proposition 3.1, stated again for the sake of exposition.

Proposition 3.1. *For each $n \in \mathbb{N}$, assume that $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) \in \mathcal{M}_n^l$ and that for $n \in \mathbb{N}$, $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$ is an ϵ -optimal solution for the l -bounded, (n, ∞) -dimensional linear program. Then, for $\delta > 0$, there exists an $N(\delta)$ such that*

$$J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) - J^* \leq 2\epsilon + \delta.$$

for all $n \geq N(\delta)$.

Proof. Assume first that $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$ converges weakly to some $(\mu_0^\epsilon, \mu_1^\epsilon)$. Then, for $g \in \mathcal{D}_\infty$, and $g_k \rightarrow g$ in \mathcal{D}_∞ we have

$$\begin{aligned} \int Agd\mu_0 + \int Bgd\mu_1 &= \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \int Ag_k d\mu_{0,n}^\epsilon + \int Bg_k d\mu_{1,n}^\epsilon \\ &= \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} Rg_k = \lim_{k \rightarrow \infty} Rg_k = Rg. \end{aligned}$$

Thus, $(\mu_0^\epsilon, \mu_1^\epsilon) \in \mathcal{M}_\infty$. Now observe that since c_1 and c_2 are continuous,

$J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) \rightarrow J(\mu_0^\epsilon, \mu_1^\epsilon)$. Assume there would be an $(\hat{\mu}_0, \hat{\mu}_1) \in \mathcal{M}_\infty^l$ such that $J(\mu_0^\epsilon, \mu_1^\epsilon) > J(\hat{\mu}_0, \hat{\mu}_1) + \epsilon$, that is, assume $(\mu_0^\epsilon, \mu_1^\epsilon)$ would not be ϵ -optimal. For n large enough, we would have $J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) > J(\hat{\mu}_0, \hat{\mu}_1) + \epsilon$. But as $\mathcal{M}_n \supset \mathcal{M}_\infty$, this would imply that $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$ is not ϵ -optimal. So, $(\mu_0^\epsilon, \mu_1^\epsilon)$ is ϵ -optimal.

In general, we cannot guarantee weak convergence of $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$, but since $E \times U$ is compact and the full mass of $\mu_{0,n}^\epsilon$ and $\mu_{1,n}^\epsilon$ is uniformly bounded by 1 and l , respectively, the existence of a convergent subsequence is given by Theorem 1.9. Consider two different subsequences of ϵ -optimal measures $(\mu_{0,n_1}^\epsilon, \mu_{1,n_1}^\epsilon)$ and $(\mu_{0,n_2}^\epsilon, \mu_{1,n_2}^\epsilon)$. Set $z_1 = \lim_{n_1 \rightarrow \infty} J(\mu_{0,n_1}^\epsilon, \mu_{1,n_1}^\epsilon)$ and $z_2 = \lim_{n_2 \rightarrow \infty} J(\mu_{0,n_2}^\epsilon, \mu_{1,n_2}^\epsilon)$. Both z_1 and z_2 are ϵ -optimal cost criterion values in \mathcal{M}_∞^l . Hence we can conclude that $|z_1 - z_2| < \epsilon$. In particular, for $z \in \mathbb{R}$ such that $J(\mu_0^\epsilon, \mu_1^\epsilon) \in [z - \frac{\epsilon}{2}, z + \frac{\epsilon}{2}]$ for any weak limit $(\mu_0^\epsilon, \mu_1^\epsilon)$ of a sequence of ϵ -optimal measures. This means that for $\delta > 0$, there is an $N \equiv N(\delta)$ large enough such that for all $n \geq N$, $J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) \in (z - \frac{\epsilon}{2} - \delta, z + \frac{\epsilon}{2} + \delta)$. Now assume that $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$ does not converge weakly, and that $J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) \notin (z - \frac{\epsilon}{2} - \delta, z + \frac{\epsilon}{2} + \delta)$ infinitely often. This would allow for the construction of a subsequence $(\mu_{0,n_3}^\epsilon, \mu_{1,n_3}^\epsilon)$ with $J(\mu_{n_3}, \mu_{1,n_3}) \notin (z - \frac{\epsilon}{2} - \delta, z + \frac{\epsilon}{2} + \delta) \forall n_3 \in \mathbb{N}$. However, by the tightness argument of Theorem 1.9, that sequence contains a converging sub-subsequence $(\mu_{0,n_4}, \mu_{1,n_4})$. This means that $J(\mu_{0,n_4}, \mu_{1,n_4}) \in (z - \frac{\epsilon}{2} - \delta, z + \frac{\epsilon}{2} + \delta)$ eventually, contradicting the preceding assumption. So, there exists an $N \in \mathbb{N}$ such that $J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) \in (z - \frac{\epsilon}{2} - \delta, z + \frac{\epsilon}{2} + \delta) \forall n \geq N$. Consider a limit $(\mu_0^\epsilon, \mu_1^\epsilon)$ of a convergent subsequence of $(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon)$. By its ϵ -optimality and the properties of the infimum,

$$\begin{aligned} J^* + \epsilon &\geq J(\mu_0^\epsilon, \mu_1^\epsilon) \geq z - \frac{\epsilon}{2} \\ \Leftrightarrow J^* &\geq z - \frac{3\epsilon}{2}. \end{aligned}$$

Thereby, for $n \geq N$,

$$J(\mu_{0,n}^\epsilon, \mu_{1,n}^\epsilon) - J^* \leq z + \frac{\epsilon}{2} + \delta - \left(z - \frac{3\epsilon}{2}\right) = 2\epsilon + \delta$$

from which the claim follows. \square

We proceed to analyze the discretization scheme of Section 2.1 and its approximation properties in order to prove Proposition 3.2. From here on, we consider a fixed n , assuming that it is large enough to guarantee that the statement of Proposition 3.1 holds. In particular, we consider the space \mathcal{D}_n , which is spanned by the basis functions f_1, f_2, \dots, f_n .

Proposition 3.2 states that we can approximate the cost criterion of a pair of measures $(\mu_0, \mu_1) \in \mathcal{M}_n^l$ arbitrarily closely by a pair of measures $(\mu_{m,0}, \mu_{m,1}) \in \mathcal{M}_{n,m}^l$. In the following, we consider a fixed pair of measures (μ_0, μ_1) . If η_0 and η_1 are the regular conditional probabilities of μ_0 and μ_1 , respectively, we use the following choice of the coefficients $\beta_{j,i}$ and $\zeta_{j,i}$ in (2.6) and (2.8), respectively. We use the mesh points in E and U as defined in (2.4) and (2.3). Set

$$U_i = [u_i, u_{i+1}) \text{ for } 0 \leq i \leq 2^{k_m} - 1, \quad U_{2^{k_m}} = \{u_{2^{k_m}}\} \quad (3.1)$$

and with that,

$$\beta_{j,i} = \eta_0(U_i, x_j) = \int_{U_i} \eta_0(du, x_j), \quad j = 0, \dots, 2^m - 1, i = 0, \dots, 2^{k_m}, \quad (3.2)$$

$$\zeta_{1,i} = \eta_1(U_i, e_l) = \int_{U_i} \eta_1(du, e_l), \quad \zeta_{2,i} = \eta_1(U_i, e_r) = \int_{U_i} \eta_1(du, e_r), \quad (3.3)$$

$$i = 0, \dots, 2^{k_m}.$$

The following two facts can easily be derived using the uniform continuity of c_0, c_1 as well as f_1, \dots, f_n (recall that these are continuous functions on a compact set), and the specific forms of the generators A and B as in seen (1.8) and (1.4), respectively.

Lemma 3.6. *For $\epsilon > 0$, there is a $\delta > 0$ such that, uniformly in $x \in E$,*

$$\max \{|c_0(x, u) - c_0(x, v)|, |Af_1(x, u) - Af_1(x, v)|, \dots, |Af_n(x, u) - Af_n(x, v)|\} < \epsilon$$

holds whenever $|u - v| < \delta$.

Lemma 3.7. *For $\epsilon > 0$, there is a $\delta > 0$ such that for $s = e_l$ or $s = e_r$,*

$$\max \{|c_1(s, u) - c_1(s, v)|, |Bf_1(s, u) - Bf_1(s, v)|, \dots, |Bf_n(s, u) - Bf_n(s, v)|\} < \epsilon$$

holds whenever $|u - v| < \delta$.

The following two results ensure that we can approximate the cost criterion of a pair of measures $(\mu_0, \mu_1) \in \mathcal{M}_n^l$ arbitrarily closely, and that the our approximate measures ‘almost’ satisfy the linear constraints.

Proposition 3.8. *Consider a regular conditional probability η_0 and a probability density function p stemming from a continuous occupation measure μ_0 such that $\mu_0(dx \times du) = \eta_0(du, x)p(x)dx$. Let $g(x, u) = c_0(x, u)$ or $g(x, u) = Af_k(x, u)$ for any $k = 1, 2, \dots, n$. For $\epsilon > 0$, there exists an $m_0 \in \mathbb{N}$ such that for all $m \geq m_0$,*

$$\left| \int_E \int_U g(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U g(x, u) \hat{\eta}_{0,m}(du, x) p(x) dx \right| < \epsilon, \quad (3.4)$$

where $\hat{\eta}_{0,m}$ is of the form (2.6), using the coefficients specified in (3.2).

Proof. Observe that given (2.6),

$$\begin{aligned} |I| &\equiv \left| \int_E \int_U g(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U g(x, u) \hat{\eta}_{0,m}(du, x) p(x) dx \right| \\ &= \left| \int_E \int_U g(x, u) \eta_0(du, x) p(x) dx - \int_E \left(\sum_{j=0}^{2^m-1} \sum_{i=0}^{2^{k_m}} \beta_{j,i} I_{E_j}(x) g(x, u_i) \right) p(x) dx \right|. \end{aligned}$$

By the definition of $\beta_{j,i}$ in (3.2) and the triangle inequality it follows that

$$\begin{aligned} |I| &\leq \left| \int_E \int_U g(x, u) \eta_0(du, x) p(x) dx - \int_E \left(\sum_{i=0}^{2^{k_m}} g(x, u_i) \int_{U_i} \eta_0(du, x) \right) p(x) dx \right| \\ &\quad + \left| \int_E \left(\sum_{i=0}^{2^{k_m}} g(x, u_i) \int_{U_i} \eta_0(du, x) - \sum_{j=0}^{2^m-1} \sum_{i=0}^{2^{k_m}} \int_{U_i} \eta_0(du, x_j) I_{E_j}(x) g(x, u_i) \right) p(x) dx \right| \\ &\equiv |I_1| + |I_2|. \end{aligned}$$

Observe that

$$|I_1| = \left| \int_E \left(\sum_{i=0}^{2^{k_m}} \int_{U_i} (g(x, u) - g(x, u_i)) \eta_0(du, x) \right) p(x) dx \right|.$$

By Lemma 3.6, there is a $\delta > 0$ such that for all $|u - v| < \delta$, we have that $|g(x, u) - g(x, v)| < \frac{\epsilon}{2}$, uniformly in $x \in E$. Choose m_1 large enough such that for all $m \geq m_1$ it is true that $\frac{1}{2^{k_m}} < \delta$. Then

$$|I_1| < \left| \int_E \left(\sum_{i=0}^{2^{k_m}} \int_{U_i} \frac{\epsilon}{2} \eta_0(du, x) \right) p(x) dx \right| = \frac{\epsilon}{2}. \quad (3.5)$$

We now examine the term

$$|I_2| = \left| \int_E \sum_{i=0}^{2^{k_m}} g(x, u_i) \left(\int_{U_i} \eta_0(du, x) - \sum_{j=0}^{2^m-1} I_{E_j}(x) \int_{U_i} \eta_0(du, x_j) \right) p(x) dx \right|. \quad (3.6)$$

Regard I_2 as a sequence with two indices, say $I_2(a, b) \equiv I_2(a, k_b) \equiv I_2$, where a and b are two discretization levels, with a slight abuse of notation.

Our first claim is that $I_2(a, b)$ is a Cauchy sequence in b . To see this, we analyze two successive elements of the sequence. Consider

$$\begin{aligned} |I_2(a, b+1) - I_2(a, b)| &= |I_2(a, k_{b+1}) - I_2(a, k_b)| \\ &= \left| \int_E \left[\sum_{i=0}^{2^{k_{b+1}}} g(x, \tilde{u}_i) \left(\eta_0(\tilde{U}_i, x) - \sum_{j=0}^{2^a-1} I_{E_j}(x) \eta_0(\tilde{U}_i, x_j) \right) \right. \right. \\ &\quad \left. \left. - \sum_{i=0}^{2^{k_b}} g(x, u_i) \left(\eta_0(U_i, x) - \sum_{j=0}^{2^a-1} I_{E_j}(x) \eta_0(U_i, x_j) \right) \right] p(x) dx \right| \end{aligned}$$

where \tilde{U}_i and \tilde{u}_i are used to indicate the partition of U and the points of the discrete set in U of the discretization level $b+1$, defined analogously to (2.3) and (3.1). Due to the additivity of measures, the two sums over i , if regarded as a Riemann-type approximation to an integral, only differ by a more accurate choice of the ‘rectangle height’ $g(x, u_i)$ and $g(x, \tilde{u}_i)$ in the Riemann sum. To formalize this, for $i \in \{0, \dots, 2^{k_{b+1}}\}$

let $\pi(i) \in \{0, \dots, 2^{k_b}\}$ be the index such that $\tilde{U}_i \subset U_{\pi(i)}$. Observe that $\sum_{i=0}^{2^{k_b+1}} |\eta_0(U_i, x) - \eta_0(U_i, x_j)| \leq 2$, independently of x_j , and thus independently of our choice of $E^{(a)}$. This is due to the fact regular conditional probabilities are indeed probability measures with a full mass of 1. Then,

$$\begin{aligned}
& |I_2(a, b+1) - I_2(a, b)| \\
&= \left| \int_E \left[\sum_{i=0}^{2^{k_b+1}} (g(x, \tilde{u}_i) - g(x, u_{\pi(i)})) \cdot \left(\eta_0(\tilde{U}_i, x) - \sum_{j=0}^{2^a-1} I_{E_j}(x) \eta_0(\tilde{U}_i, x_j) \right) \right] p(x) dx \right| \\
&\leq \int_E \left[\sum_{i=0}^{2^{k_b+1}} |g(x, \tilde{u}_i) - g(x, u_{\pi(i)})| \cdot \left| \eta_0(\tilde{U}_i, x) - \sum_{j=0}^{2^a-1} I_{E_j}(x) \eta_0(\tilde{U}_i, x_j) \right| \right] p(x) dx \\
&\leq K \left(\frac{1}{2} \right)^{b+1} \int_E \left[\sum_{i=0}^{2^{k_b+1}} \left| \eta_0(\tilde{U}_i, x) - \sum_{j=0}^{2^a-1} I_{E_j}(x) \eta_0(\tilde{U}_i, x_j) \right| \right] p(x) dx \\
&\leq K \left(\frac{1}{2} \right)^{b+1} \cdot 2 = K \left(\frac{1}{2} \right)^b
\end{aligned}$$

by the fact that $|g(x, \tilde{u}_i) - g(x, u_{\pi(i)})|$ is uniformly bounded by $K \left(\frac{1}{2} \right)^{b+1}$, with $K = 1$ if $g(x, u) = c_0(x, u)$, and $K = \max\{\|f_1\|_{\mathcal{D}}, \dots, \|f_k\|_{\mathcal{D}}\}$ if $g(x, u) = Af_k(x, u)$ by our choice of $U^{(k_m)}$, compare (2.2).

Now, for some $\vartheta > 0$, choose b large enough such that $\sum_{j=b}^{\infty} \left(\frac{1}{2} \right)^j < \frac{\vartheta}{K}$. Then, for all $b_1 > b_2 \geq b$, we have

$$\begin{aligned}
|I_2(a, b_1) - I_2(a, b_2)| &= \left| \sum_{j=b_2}^{b_1-1} I_2(a, j+1) - I_2(a, j) \right| \\
&\leq \sum_{j=b_2}^{b_1-1} |I_2(a, j+1) - I_2(a, j)| \\
&\leq K \sum_{j=b_2}^{b_1-1} \left(\frac{1}{2} \right)^j < \vartheta,
\end{aligned} \tag{3.7}$$

revealing that I_2 is Cauchy in b , which is the same as saying it is Cauchy in k_b . The bound on the increment of I_2 in k_b (given by (3.7)) is independent of a , so it does not depend on the choice of $E^{(a)}$. Given this result, choose $m_2 \geq m_1$ such that for all $m \geq m_2$, we have that k_m is large enough to make $|I_2(m, b_1) - I_2(m, b_2)| < \frac{\epsilon}{4}$ hold for all $b_1, b_2 \geq k_m$. An application of the dominated convergence theorem together with the assumption that $x \mapsto \eta(U_i, x)$ is continuous almost everywhere reveals that the choice of coefficients in (3.2) ensures that for k_m fixed and for each $i \in \{0, 1, \dots, 2^{k_m}\}$, there is a $m^{(1,i)}$ large enough such that for all $m^{(i)} \geq m^{(1,i)}$, we have

$$\begin{aligned}
& \int_E \left| \eta_0(U_i, x) - \sum_{j=0}^{2^{m^{(i)}}-1} I_{E_j}(x) \eta_0(U_i, x_j) \right| p(x) dx \\
&< \frac{\epsilon}{4 \max\{\|c_0\|_{\infty}, \|Af_1\|_{\infty}, \dots, \|Af_n\|_{\infty}\} (2^{k_m+1})}
\end{aligned} \tag{3.8}$$

Set $\tilde{m} = \max\{\max_{i=0, \dots, 2^{k_m}} m^{(1,i)}, m_2\}$. Then,

$$\begin{aligned}
I_2(\tilde{m}, k_m) &\leq \int_E \|g\|_{\infty} \sum_{i=0}^{2^{k_m}} \left| \eta_0(U_i, x) - \sum_{j=0}^{2^{\tilde{m}}-1} I_{E_j}(x) \eta_0(U_i, x_j) \right| p(x) dx \\
&\leq \|g\|_{\infty} \sum_{i=0}^{2^{k_m}} \int_E \left| \eta_0(U_i, x) - \sum_{j=0}^{2^{\tilde{m}}-1} I_{E_j}(x) \eta_0(U_i, x_j) \right| p(x) dx \\
&\leq \frac{\epsilon}{4}.
\end{aligned}$$

Note that $I_2(\tilde{m}, k_m)$ is decreasing in \tilde{m} , which again is revealed using the dominated convergence theorem. Also, for $l \geq k_m$, we have

$$\begin{aligned} |I_2(\tilde{m}, l)| &\leq |I_2(\tilde{m}, l) - I_2(\tilde{m}, k_m)| + |I_2(\tilde{m}, k_m)| \\ &< \frac{\epsilon}{4} + \frac{\epsilon}{4} = \frac{\epsilon}{2} \end{aligned}$$

which means that $I_2 \equiv I_2(\tilde{m}, k_m)$ does not exceed $\frac{\epsilon}{2}$ when \tilde{m} or k_m increase. Choose $m_0 = \max\{\tilde{m}, m_2\}$. Then, for all $m \geq m_0$, we have that $I_2 < \frac{\epsilon}{2}$. \square

Proposition 3.9. *Consider a singular occupation measure μ_1 that decomposes into $\mu_1(dx \times du) = \eta_1(du, x)\mu_{1,E}(dx)$. Let $g(x, u) = c_1(x, u)$ or $g(x, u) = Bf_k(x, u)$ for any $k = 1, 2, \dots, n$. For $\epsilon > 0$, there exists an $m_0 \in \mathbb{N}$ such that for all $m \geq m_0$,*

$$\left| \int_E \int_U g(x, u) \eta_1(du, x) \mu_{1,E}(dx) - \int_E \int_U g(x, u) \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx) \right| < \epsilon, \quad (3.9)$$

where $\hat{\eta}_{1,m}$ is of the form (2.8), using the coefficients specified in (3.3).

Proof. We only have to show that

$$\left| \int_U g(s, u) \eta_1(du, s) - g(s, u) \hat{\eta}_{1,m}(du, s) \right| < \frac{\epsilon}{\mu_{1,E}(E)} \quad (3.10)$$

uniformly for $s = e_l$ or $s = e_r$, since $\mu_{1,E}$ only puts mass on these two points. By (3.3),

$$\begin{aligned} &\left| \int_U g(s, u) \eta_1(du, s) - \int_U g(s, u) \hat{\eta}_{1,m}(du, s) \right| \\ &= \left| \sum_{i=0}^{2^{k_m}} \int_{U_i} g(s, u) \eta_1(du, s) - \int_{U_i} g(s, u) \hat{\eta}_{1,m}(du, s) \right| \\ &= \left| \sum_{i=0}^{2^{k_m}} \int_{U_i} g(s, u) \eta_1(du, s) - g(s, u_i) \zeta_{j,i} \right| \\ &= \left| \sum_{i=0}^{2^{k_m}} \int_{U_i} g(s, u) \eta_1(du, s) - g(s, u_i) \int_{U_i} \eta_1(du, s) \right| \\ &\leq \sum_{i=0}^{2^{k_m}} \int_{U_i} |g(s, u) - g(s, u_i)| \eta_1(du, s) \end{aligned}$$

holds. According to Lemma 3.7, there is a $\delta > 0$ such that $|g(s, u) - g(s, v)| < \frac{\epsilon}{\mu_{1,E}(E)}$ whenever $|u - v| < \delta$. Hence it suffices to choose m_0 large enough such that for all $m \geq m_0$ it is true that $\frac{1}{2^{k_m}} < \delta$ to ensure that $|u - u_i| < \delta$. \square

Proposition 3.8 and Proposition 3.9 have established the approximation properties of the coefficient choices made by (3.2) and (3.3), without paying any respect to the constraints defining $\mathcal{M}_{n,m}^l$. We proceed to link these approximate controls to measures that actually fulfill those constraints. To do so, we need to be able to quantify how far away a given approximation is from satisfying the constraints. This motivates the following definitions.

Definition 3.10. *Let η_0 be any relaxed control. For $n, m \in \mathbb{N}$, define the constraint matrix $C^{(m)} \in \mathbb{R}^{n+1, 2^m}$ by*

$$\begin{aligned} C_{k,j}^{(m)} &= \int_E \int_U Af_k(x, u) \eta_0(du, x) p_j(x) dx, \quad k = 1, 2, \dots, n, \quad j = 0, 1, \dots, 2^m - 1 \\ C_{n+1,j}^{(m)} &= \int_E p_j(x) dx \quad j = 0, 1, \dots, 2^m - 1. \end{aligned}$$

Definition 3.11. For $m \in \mathbb{N}$ and $\hat{\eta}_{0,m}$ given by (3.2), $\hat{\eta}_{1,m}$ given by (3.3) and some \tilde{p} in the span of $\{p_0, p_1, \dots, p_{2^m-1}\}$, with $\mu_{1,E}$ being the state space marginal of the previously fixed measure μ_1 , the constraint error $d^{(m)}(\tilde{p}) \in \mathbb{R}^{n+1}$ is defined for $k = 1, \dots, n$ by

$$d_k^{(m)}(\tilde{p}) = Rf_k - \int_E \int_U Af_k(x, u) \hat{\eta}_{0,m}(du, x) \tilde{p}(x) dx \\ - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx),$$

and for $k = n+1$ by $d_{n+1}^{(m)} = 1 - \int_E \tilde{p}(x)(dx)$.

One can increase m to the point that the constraint matrix $C^{(m)}$ has full rank. This will help us find adjustments to the coefficients of \tilde{p} which will let the constraint error vanish. The following results show that we can attain an arbitrarily small constraint error using the proposed approximation. Their proofs are rather technical, and are given in Appendix A.

Lemma 3.12. Let p be a probability density function with $\lambda(\{x : p(x) = 0\}) = 0$. Then, for any $\epsilon > 0$ and $D_1 > 0$, there exists an $\hat{\epsilon}_1 < \epsilon$ and an m_0 such that for all $m \geq m_0$, there is a \tilde{p}_m in the span of $\{p_0, p_1, \dots, p_{2^m-1}\}$ with $\|p - \tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_1}{D_1}$ and $\tilde{p}_m \geq \hat{\epsilon}_1$ on E .

Lemma 3.13. Consider a pair of measures $(\mu_0, \mu_1) \in \mathcal{M}_{n,\infty}$, and let $\mu_0(dx \times du) = \eta_0(du, x) p(x) dx$ as well as $\mu_1(dx \times du) = \eta_1(du, x) \mu_{1,E}(dx)$. Let $\bar{A} = \max_{k=1,\dots,n} \|Af_k\|_\infty$. For $\delta > 0$ and $D_2 \geq 1$, there exists an $\hat{\epsilon}_2 \leq \delta$ and an $m_0 \in \mathbb{N}$ such that for all $m \geq m_0$, there is a function \tilde{p}_m in the span of $\{p_0, p_1, \dots, p_{2^m-1}\}$ with $\|d^{(m)}(\tilde{p}_m)\|_\infty < \hat{\epsilon}_2$, where $d^{(m)}(\tilde{p}_m)$ is the constraint error using the approximations $\hat{\eta}_{0,m}(du, x)$ and $\hat{\eta}_{1,m}(du, x)$ of the given controls η_0 and η_1 defined by the coefficients given in (3.2) and (3.3). In particular, $\|p - \tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_2}{3 \max\{\bar{A}, 1\}}$ as well as $\tilde{p}_m \geq D_2 \cdot \hat{\epsilon}_2$ holds.

Next, we establish that we can find a ‘correction’ term y for the coefficients of \tilde{p} , which can be used to define a measure that satisfies the constraints, with a maximum norm that does not exceed a given bound ϵ . For the proof of the following statement, we again refer to Appendix A.

Lemma 3.14. Consider a pair of measures $(\mu_0, \mu_1) \in \mathcal{M}_{n,\infty}$, and let $\mu_0(dx \times du) = \eta_0(du, x) p(x) dx$ as well as $\mu_1(dx \times du) = \eta_1(du, x) \mu_{1,E}(dx)$. For any $\vartheta > 0$, there is a $\hat{\vartheta} < \vartheta$ and an $m_0 \in \mathbb{N}$ such that for all $m \geq m_0$, there is a $\tilde{p}_m \in \text{span}\{p_0, p_1, \dots, p_{2^m-1}\}$ such that the equation $C^{(m)}y = -d^{(m)}(\tilde{p}_m)$ has a solution \tilde{y} with $\|\tilde{y}\|_\infty \leq \hat{\vartheta}$. In particular, $\tilde{p}_m \geq \hat{\vartheta}$ and $\|p - \tilde{p}_m\|_{L^1(E)} < \hat{\vartheta}$ hold.

At this point, we have gathered the results to prove Proposition 3.2, stated again for the sake of presentation.

Proposition 3.2. For $(\mu_0, \mu_1) \in \mathcal{M}_n^l$ and each $\epsilon > 0$, there is an m_0 such that for all $m \geq m_0$ there exists a $(\hat{\mu}_{0,m}, \hat{\mu}_{1,m}) \in \mathcal{M}_{n,m}^l$, with

$$|J(\mu_0, \mu_1) - J(\hat{\mu}_{0,m}, \hat{\mu}_{1,m})| < \epsilon.$$

Proof. Fix $\epsilon > 0$. For $(\mu_0, \mu_1) \in \mathcal{M}_n^l$, let $\mu_{0,E}$ be the state space marginal of μ_0 and let η_0 be the regular conditional probability such that $\mu_0(dx \times du) = \eta_0(du, x) \mu_{0,E}(dx)$. Similarly, let $\mu_1(dx \times du) = \eta_1(du, x) \mu_{1,E}(dx)$. Define $\hat{\eta}_{0,m}$ and $\hat{\eta}_{1,m}$ using the coefficients given by (3.2) and (3.3), respectively. First, by Proposition 3.9, we have that there is an m_1 such that $\forall m \geq m_1$,

$$\left| \int_E \int_U c_1(x, u) \eta_1(du, x) \mu_{1,E}(dx) - \int_E \int_U c_1(x, u) \hat{\eta}_{1,m}(du, x) \hat{\mu}_{1,E}(dx) \right| < \frac{\epsilon}{2}. \quad (3.11)$$

Now, we consider the approximation of the cost accrued by c_0 . We will show that

$$\left| \int_E \int_U c_0(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U c_0(x, u) \hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx \right| < \frac{\epsilon}{2}, \quad (3.12)$$

for the given choice of $\hat{\eta}_{0,m}$ and a choice of \hat{p}_m to be identified. This will be done by a successive application of the triangle inequality. First, observe that

$$\begin{aligned} & \left| \int_E \int_U c_0(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U c_0(x, u) \hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx \right| \\ & \leq \left| \int_E \int_U c_0(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U c_0(x, u) \hat{\eta}_{0,m}(du, x) p(x) dx \right| \\ & \quad + \left| \int_E \int_U c_0(x, u) \hat{\eta}_{0,m}(du, x) p(x) dx - \int_E \int_U c_0(x, u) \hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx \right| \\ & \equiv |I_1| + |I_2|. \end{aligned}$$

Now set

$$\vartheta = \min \left\{ \frac{\epsilon}{8\|c_0\|_\infty \max\{1, (e_r - e_l)\}}, \frac{3\epsilon \max\{\bar{A}, 1\}}{8\|c_0\|_\infty} \right\}. \quad (3.13)$$

By Lemma 3.14 we can choose an $m_2 \geq m_1$ such that for all $m \geq m_2$, there is a function $\tilde{p}_m = \sum_{i=0}^{2^m-1} \tilde{\gamma}_i p_i \in \text{span}\{p_0, p_1, \dots, p_{2^m-1}\}$ that allows for a solution \tilde{y} to $C^{(m)}y = -d^{(m)}(\tilde{p}_m)$ with $\|\tilde{y}\|_\infty \leq \hat{\vartheta} < \vartheta$ for some $\hat{\vartheta} < \vartheta$, but $\tilde{p}_m \geq \hat{\vartheta}$. This m_2 is also large enough to approximate p by \tilde{p}_m with an accuracy of $\frac{\epsilon}{8\|c_0\|_\infty}$ for all $m \geq m_2$. Define new coefficients $\gamma_i = \tilde{\gamma}_i - \tilde{y}_i$ and set $\hat{p}_m = \sum_{i=0}^{2^m-1} \gamma_i p_i$. Then, for all $k = 1, 2, \dots, n$,

$$\begin{aligned} d_k^{(m)}(\hat{p}_m) &= Rf_k - \left(C^{(m)}\gamma\right)_k - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m} \mu_{1,E}(dx) \\ &= Rf_k - \left(C^{(m)}(\tilde{\gamma} - \tilde{y})\right)_k - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m} \mu_{1,E}(dx) \\ &= Rf_k - \left(C^{(m)}\tilde{\gamma}\right)_k - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m} \mu_{1,E}(dx) + \left(C^{(m)}\tilde{y}\right)_k \\ &= d_k^{(m)}(\tilde{p}_m) - d_k^{(m)}(\tilde{p}_m) = 0 \end{aligned}$$

and

$$d_{n+1}^{(m)}(\hat{p}_m) = 1 - (C^{(m)}\tilde{\gamma})_{n+1} + (C^{(m)}\tilde{y})_{n+1} = d_{n+1}^{(m)}(\tilde{p}_m) - d_{n+1}^{(m)}(\tilde{p}_m) = 0$$

which shows that \hat{p}_m fulfills the constraints. But also, $\hat{p}_m \geq 0$. So $(\hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx, \hat{\eta}_{0,m}(du, x) \mu_{1,E}(dx)) \in \mathcal{M}_{n,m}^l$. Furthermore,

$$\|p - \hat{p}_m\|_{L^1(E)} \leq \|p - \tilde{p}_m\|_{L^1(E)} + \|\tilde{p}_m - \hat{p}_m\|_{L^1(E)} \leq \frac{\epsilon}{8\|c_0\|_\infty} + \frac{\epsilon}{8\|c_0\|_\infty} = \frac{\epsilon}{4\|c_0\|_\infty}.$$

This shows that

$$|I_2| \leq \int_E \|c_0\|_\infty \int_U \hat{\eta}_{0,m}(du, x) |p(x) - \hat{p}_m(x)| dx < \|c_0\|_\infty \|p - \hat{p}_m\|_{L^1} < \frac{\epsilon}{4},$$

since $\int_U \hat{\eta}_{0,m}(du, x) = 1$. Turning to $|I_1|$, we have seen in Proposition 3.8 that there is an $m_3 \geq m_2$ such that for all $m \geq m_3$, $|I_1| < \frac{\epsilon}{4}$. To sum up,

$$\begin{aligned} & \left| \int_E \int_U c_0(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U c_0(x, u) \hat{\eta}_{0,m}(du, x) \hat{p}_m(x) dx \right| \\ & \leq |I_1| + |I_2| \leq \frac{\epsilon}{4} + \frac{\epsilon}{4} = \frac{\epsilon}{2} \end{aligned}$$

But this gives us the assertion, and finishes the proof, setting $m_0 = m_3$. \square

So far, we have investigated the approximation of arbitrary measures. Proposition 3.2 is instrumental in proving the next important result, Corollary 3.3, which analyzes how optimal solutions in $\mathcal{M}_{n,m}^l$ relate to ϵ -optimal solution in \mathcal{M}_n^l . The following lemma will as well be needed. Its proof is similar to an argument used in (3.1), and thus is omitted.

Lemma 3.15. *Let $\{\mu_{0,n,m}, \mu_{1,n,m}\}$ be a sequence of measures such that for each m , $(\mu_{0,n,m}, \mu_{1,n,m}) \in \mathcal{M}_{n,m}^l$. Assume that $\mu_{0,n,m} \Rightarrow \hat{\mu}_{0,n}$ and $\mu_{1,n,m} \Rightarrow \hat{\mu}_{1,n}$ as $m \rightarrow \infty$. Then, $(\hat{\mu}_{0,n}, \hat{\mu}_{1,n}) \in \mathcal{M}_n^l$.*

Now we can prove Corollary 3.3 from Section 3.1.

Corollary 3.3. *For each $m \in \mathbb{N}$, assume that $(\mu_{0,n,m}^*, \mu_{1,n,m}^*) \in \mathcal{M}_{n,m}^l$ and that for $m \in \mathbb{M}$, $(\mu_{0,n,m}^*, \mu_{1,n,m}^*)$ is an optimal solution to the l -bounded, (n, m) -dimensional linear program. Then, the sequence of numbers $\{J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)\}_{m \in \mathbb{N}}$ converges to $J_n^* = \inf_{(\mu_{0,n}, \mu_{1,n}) \in \mathcal{M}_n^l} J(\mu_{0,n}, \mu_{1,n})$ as $m \rightarrow \infty$.*

Proof. First, observe that if $\mu_{0,n,m}^* \Rightarrow \hat{\mu}_{0,n}^*$ and $\mu_{1,n,m}^* \Rightarrow \hat{\mu}_{1,n}^*$ as $m \rightarrow \infty$ for some $(\hat{\mu}_{0,n}^*, \hat{\mu}_{1,n}^*) \in \mathcal{M}_n^l$, it follows that $J(\hat{\mu}_{0,n}^*, \hat{\mu}_{1,n}^*) = J_n^*$. The proof of this claim is as follows. Assume the opposite. Then, there is a pair of measures $(\hat{\mu}_{0,n}, \hat{\mu}_{1,n}) \in \mathcal{M}_n^l$ with $J(\hat{\mu}_{0,n}^*, \hat{\mu}_{1,n}^*) - J(\hat{\mu}_{0,n}, \hat{\mu}_{1,n}) > \epsilon$ for some $\epsilon > 0$. Select m_0 large enough such that for all $m \geq m_0$, $|J(\hat{\mu}_{0,n}^*, \hat{\mu}_{1,n}^*) - J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)| < \frac{\epsilon}{2}$ and hence $J(\mu_{0,n,m}^*, \mu_{1,n,m}^*) - J(\hat{\mu}_{0,n}, \hat{\mu}_{1,n}) > \frac{\epsilon}{2}$ for all $m \geq m_0$. By Proposition 3.2, select $m \geq m_0$ large enough that there is a pair of measures $(\hat{\mu}_{0,n,m}, \hat{\mu}_{1,n,m}) \in \mathcal{M}_{n,m}^l$ with $|J(\hat{\mu}_{0,n}, \hat{\mu}_{1,n}) - J(\hat{\mu}_{0,n,m}, \hat{\mu}_{1,n,m})| < \frac{\epsilon}{2}$. But then, $J(\hat{\mu}_{0,n,m}, \hat{\mu}_{1,n,m}) < J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)$, contradicting that $(\mu_{0,n,m}^*, \mu_{1,n,m}^*)$ is the optimal solution in $\mathcal{M}_{n,m}^l$.

Also, $\{J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)\}_{m \in \mathbb{N}}$ is a decreasing sequence which is bounded from below, so it converges. As $\{\mu_{0,n,m}^*\}_{m \in \mathbb{N}}$ and $\{\mu_{1,n,m}^*\}_{m \in \mathbb{N}}$ are sequences of measures over a compact space, they are tight, and the full mass of $\{\mu_{1,n,m}^*\}_{m \in \mathbb{N}}$ is uniformly bounded by l . So there is a convergent subsequence $\{(\mu_{0,n,m_1}^*, \mu_{1,n,m_1}^*)\}_{m_1 \in \mathbb{N}}$ with

$$\mu_{0,n,m_1}^* \Rightarrow \hat{\mu}_{0,n}^* \quad \text{and} \quad \mu_{1,n,m_1}^* \Rightarrow \hat{\mu}_{1,n}^* \quad (3.14)$$

for some $(\hat{\mu}_{0,n}^*, \hat{\mu}_{1,n}^*) \in \mathcal{M}_n^l$. By the first part of this proof and because c_0 and c_1 are bounded and uniformly continuous,

$$\begin{aligned} J_n^* &= J(\hat{\mu}_{0,n}^*, \hat{\mu}_{1,n}^*) = \int c_0 d\hat{\mu}_{0,n}^* + \int c_1 d\hat{\mu}_{1,n}^* \\ &= \lim_{m_1 \rightarrow \infty} \left(\int c_0 d\mu_{0,n,m_1}^* + \int c_1 d\mu_{1,n,m_1}^* \right) \\ &= \lim_{m_1 \rightarrow \infty} J(\mu_{0,n,m_1}^*, \mu_{1,n,m_1}^*), \end{aligned}$$

but $\{J(\mu_{0,n,m}^*, \mu_{1,n,m}^*)\}_{m \in \mathbb{N}}$ converges, and any subsequence has to converge to its very limit. So,

$$\lim_{m \rightarrow \infty} J(\mu_{0,n,m}^*, \mu_{1,n,m}^*) = J_n^*$$

□

4 Examples

4.1 Modified Bounded Follower

Consider a stochastic control problem with state space $E = [0, 1]$ such that the process is governed by the SDE

$$dX_t = u(X_t)dt + \sigma dW_t + d\xi_t, \quad X_0 = x_0$$

in which $u(x) \in U = [-1, 1]$, and ξ is a process that captures the singular behavior of X . The latter is given by a reflection to the right at $\{0\}$ and a jump from $\{1\}$ to $\{0\}$. We use the relaxed martingale formulation, compare Definition 1.1, and retain the coefficient functions $b(x, u) = u$ and $\sigma(x, u) \equiv \sigma$. We adopt the long-term average cost criterion, with cost functions $c_0(x, u) = x^2$, $c_1(e_r, u) \equiv c_1$ at the right endpoint for some $c_1 \in \mathbb{R}_+$ and $c_1(e_l, u) = 0$ at the left endpoint. This problem is known as the modified bounded follower in the literature. According to [10], the optimal control for this problem is a degenerate relaxed control η_0 with $\eta_0(\{u_a(x)\}, x) = 1$, where u_a is of the form

$$u_a(x) = \begin{cases} -1 & x < a \\ +1 & x \geq a. \end{cases}$$




Figure 4.1: Computed optimal relaxed control, coarse grid, modified bounded follower

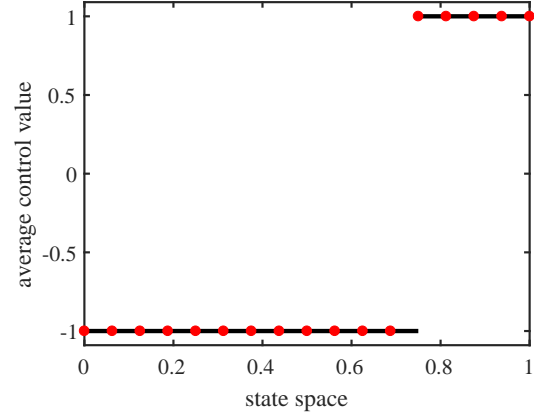


Figure 4.2: Computed optimal average control, coarse grid, modified bounded follower

The ‘switching point’ a depends on the coefficient and cost functions. Furthermore, the state space marginal $\mu_{0,E}$ under the optimal control has the density

$$p_a(x) = \frac{\int_x^1 \exp\left(\int_x^y -\frac{2}{\sigma^2} u_a(z) dz\right) dy}{\int_0^1 \int_x^1 \exp\left(\int_x^y -\frac{2}{\sigma^2} u_a(z) dz\right) dy dx}.$$

We will compare the performance of the proposed numerical method against this analytic solution. Table 4.1 shows the configuration of the problem, along with the optimal switching point a under this configuration and the weights of the occupation measure μ_1 , capturing the singular behavior on the left boundary $\{0\}$ and on the right boundary $\{1\}$, denoted w_1 and w_2 , respectively. It also shows the value of the cost criterion J^* under the optimal control. Table 4.2 and Table 4.3 show the results and performance measures for various

Table 4.1: Configuration and analytic solution, modified bounded follower

x_0	σ	c_1	a	w_1	w_2	J^*
0.1	$\sqrt{2}$	0.01	0.7512	2.4659	1.5555	0.1540

Table 4.2: Results (1), modified bounded follower

n	m	T	$J_{n,m}^*$	e_a	e_r	e_{L^1}
3	3	0.0089	0.15400	$9.798 \cdot 10^{-6}$	$6.363 \cdot 10^{-5}$	$9.082 \cdot 10^{-2}$
4	4	0.0091	0.15399	$1.199 \cdot 10^{-6}$	$7.789 \cdot 10^{-6}$	$4.563 \cdot 10^{-2}$
5	5	0.0101	0.15399	$5.447 \cdot 10^{-7}$	$3.537 \cdot 10^{-6}$	$2.287 \cdot 10^{-2}$
6	6	0.0124	0.15399	$4.809 \cdot 10^{-7}$	$3.123 \cdot 10^{-6}$	$1.145 \cdot 10^{-2}$
7	7	0.0173	0.15399	$4.713 \cdot 10^{-7}$	$3.061 \cdot 10^{-6}$	$5.73 \cdot 10^{-3}$
8	8	0.0338	0.15399	$4.694 \cdot 10^{-7}$	$3.048 \cdot 10^{-6}$	$2.866 \cdot 10^{-3}$
9	9	0.0844	0.15399	$3.306 \cdot 10^{-7}$	$2.147 \cdot 10^{-6}$	$1.451 \cdot 10^{-3}$
10	10	0.2754	0.15399	$2.655 \cdot 10^{-7}$	$1.724 \cdot 10^{-6}$	$7.258 \cdot 10^{-4}$
11	11	0.9586	0.15399	$3.550 \cdot 10^{-7}$	$2.305 \cdot 10^{-6}$	$4.509 \cdot 10^{-4}$

discretization levels n and m . To achieve higher accuracy, we added another mesh point for the choice of basis functions for p by cutting the interval in the middle of the state space in half. As the cost function does not depend on u we expect the optimal solution to be a bang-bang control. Hence it suffices to choose $k_m = 0$, which means the optimization has to choose between two possible control values $\{-1, 1\}$.

Table 4.2 shows the result for the approximate cost criterion $J_{n,m}^*$. The column e_a refers to the absolute

error between J^* and $J_{n,m}^*$, the column e_r to the relative error between J^* and $J_{n,m}^*$ and the column e_{L^1} to the $L^1(E)$ -distance between \hat{p}_m and p . T is the execution time, which is an average time taken from 1000 repetitions of the same optimization run. In Table 4.3, \hat{w}_1 and \hat{w}_2 refer to the approximate values for w_1 and w_2 . The discretization levels m for table 4.3 are the same as in table 4.2 for respective n , and as before, e_a refers to the absolute error and e_r refers to the relative error of these quantities. Note that the

Table 4.3: Results (2), modified bounded follower

n	\hat{w}_1	e_a	e_r	\hat{w}_2	e_a	e_r
3	2.4667	$7.887 \cdot 10^{-4}$	$3.199 \cdot 10^{-4}$	1.5577	$2.157 \cdot 10^{-3}$	$1.387 \cdot 10^{-3}$
4	2.4661	$1.984 \cdot 10^{-4}$	$8.047 \cdot 10^{-5}$	1.5560	$5.440 \cdot 10^{-4}$	$3.480 \cdot 10^{-4}$
5	2.4659	$4.969 \cdot 10^{-5}$	$2.015 \cdot 10^{-5}$	1.5556	$1.355 \cdot 10^{-4}$	$8.700 \cdot 10^{-5}$
6	2.4659	$1.243 \cdot 10^{-5}$	$5.040 \cdot 10^{-6}$	1.5555	$3.388 \cdot 10^{-5}$	$2.178 \cdot 10^{-5}$
7	2.4659	$3.107 \cdot 10^{-6}$	$1.260 \cdot 10^{-6}$	1.5555	$8.471 \cdot 10^{-6}$	$5.446 \cdot 10^{-6}$
8	2.4659	$7.768 \cdot 10^{-7}$	$3.150 \cdot 10^{-7}$	1.5555	$2.118 \cdot 10^{-6}$	$1.362 \cdot 10^{-6}$
9	2.4650	$8.577 \cdot 10^{-4}$	$3.478 \cdot 10^{-4}$	1.5532	$2.331 \cdot 10^{-3}$	$1.499 \cdot 10^{-3}$
10	2.4655	$4.286 \cdot 10^{-4}$	$1.738 \cdot 10^{-4}$	1.5543	$1.167 \cdot 10^{-3}$	$7.500 \cdot 10^{-4}$
11	2.4631	$2.833 \cdot 10^{-3}$	$1.149 \cdot 10^{-3}$	1.5543	$1.165 \cdot 10^{-3}$	$7.489 \cdot 10^{-4}$

method produces already fairly accurate approximations in almost negligible time for $n = 5$ or $n = 6$. The over-proportional increase in computing time for higher discretization levels ($n = 10$ and $n = 11$) is due to longer execution time of the linear program solver, and might indicate that the approximate problem is becoming ill-conditioned. For $n = 12$ and $m = 12$, no reliable solution could be produced. In this case, the linear programming solver could find no point satisfying the constraints, which can be circumvented by increasing the discretization level m without increasing the number of constraints n . However, this did not show better performance than the presented cases. The absolute error for $n = 11$ is on a comparable level to results obtained in [20]. Both the error of the cost criterion value and the L^1 -error of the state space density are steadily decreasing, which is a strong indication of a convergent method, together with the presented convergence results. The inferior approximation quality at $n = 11$ compared to $n = 10$ is believed to be due to the problem becoming ill-conditioned.

Figure 4.1 shows the computed relaxed control $\hat{\eta}_0$ for $n = 4$ and $m = 4$. Figure 4.2 shows the average

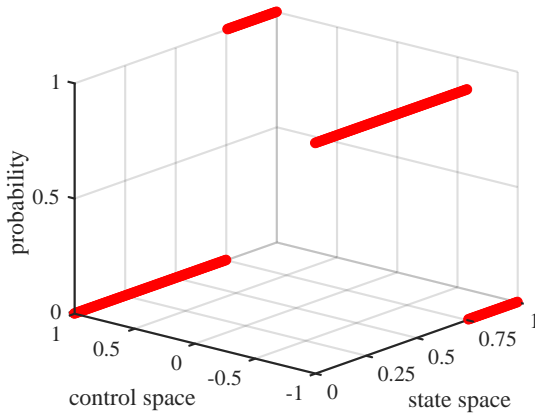


Figure 4.3: Computed optimal relaxed control, fine grid, modified bounded follower




Figure 4.4: Computed optimal average control, fine grid, modified bounded follower

control value specified by this relaxed control. These figures have to be understood as follows. Figure 4.1 displays the full relaxed control, specifying the probability to pick a certain control value u when the process is in a certain state x . This state x is found on the x -axis of the plot, labeled 'state space' and the choice of a control value u corresponds to the y -axis of the plot, labeled 'control space', while the probability $\hat{\eta}_0(\{u\}, x)$ of picking this control value u is presented on the z -axis, labeled 'probability'. For example, at $x = 0.25$,

the control value $u = -1$ is chosen with probability 1, and the control value $u = 1$ is chosen with probability 0. We can see that for any possible value of x , $\hat{\eta}_0$ assigns full mass on either one of the two possible control values $u = -1$ and $u = 1$. Hence, $\hat{\eta}_0$ can be represented by its average control function, which is given by $x \mapsto \int_U u \hat{\eta}_0(du, x)$. It is shown in Figure 4.2. In both Figure 4.1 and Figure 4.2, the red dots represent the mesh points of the mesh $E^{(m)}$ as defined in (2.4). The switching point a at $x = 0.75$, where the control switches from -1 to $+1$ is clearly visible in both figures.

The approximate state space density for $n = 4$ and $m = 4$, as displayed in blue in Figure 4.5, clearly shows the features inherited from the piecewise constant basis functions we use to approximate p . Its irregular pattern is due to the fact that we introduced an additional mesh point in the middle of the state space. Figure 4.5 also shows the exact solution displayed in red.

For a finer grid with parameters $n = 10$ and $m = 10$, Figure 4.3 shows the computed relaxed control $\hat{\eta}_0$. Figure 4.4 shows the average control function. The switching point a again is clearly visible. The red dots indicating the mesh points lie so dense that they form a solid line in both plots. Figure 4.6 shows the approximate state space density for the parameter choice of $n = 10$ and $m = 10$. The exact solution could not be visually distinguished from the approximate solution and is thus omitted from the figure. One can also see a change in concavity of the state space density at roughly $x = 0.75$, which is where the control switches its behavior from selecting $u = -1$ to $u = 1$.




Figure 4.5: State space density, coarse grid, modified bounded follower




Figure 4.6: State space density, fine grid, modified bounded follower

4.2 Simple Particle Problem with Costs of Control

To illustrate the performance of the numerical method on a different type of problem, consider a stochastic control problem with state space $E = [-1, 1]$ such that the process is governed by the SDE

$$dX_t = u(X_t)dt + \sigma dW_t + d\xi_t, \quad X_0 = x_0$$

in which $u(x) \in U = [-1, 1]$. ξ models reflections at both -1 (to the right) and 1 (to the left), keeping the process inside of E . X can be viewed as a particle that randomly diffuses inside a confined space, and bounces off at the boundaries. Again, we adopt the long-term average cost criterion and use the relaxed martingale formulation, compare Definition 1.1. We retain the coefficient functions $b(x, u) = u$ and $\sigma(x, u) \equiv \sigma$. To differentiate this example from the previous one, consider a cost structure given by $c_0 = x^2 + u^2$ and $c_1(x, u) \equiv c_1$ for some $c_1 \in \mathbb{R}_+$ at both left endpoint $x = -1$ and right endpoint $x = 1$. In particular, this means that using the control induces a cost. In contrast to the modified bounded follower of Section 4.1, we will see a different structure of the control since choosing the maximal or minimal control values might not be optimal any longer, as this introduces additional costs. For this problem, no analytical solution is known to the authors.

We examine the influence of the cost of the reflection c_1 on the optimal control. All subsequent calculations

use $\sigma = \sqrt{2}/2$, $n = 9$, $m = 9$ and $k_m = m + 3 = 12$. The latter is needed to attain a sufficient approximation of the cost function, compare (2.2).

Figure 4.7 shows the average control $x \mapsto \int_U u \hat{\eta}_0(du, x)$ for a cost of reflections given by $c_1 = 0.01$. We




Figure 4.7: Average of optimal control, $c_1 = 0.01$, simple particle problem




Figure 4.8: State space density, $c_1 = 0.01$, simple particle problem

chose to show a plot of the average control function rather than the full relaxed control since the numerical solutions were degenerate relaxed controls, putting full mass on the values attained by $x \mapsto \int_U u \hat{\eta}_0(du, x)$, with the exception of small rounding errors. Moreover, a full visualization of the relaxed control as seen in Section 4.1 is infeasible due to the high number of possible control values. Figure 4.8 shows the state space density associated with the control of Figure 4.7. The computed optimality criterion is $J_{n,m}^* = 0.30259$.




Figure 4.9: Average of optimal control, $c_1 = 1$, simple particle problem




Figure 4.10: State space density, $c_1 = 1$, simple particle problem




Figure 4.11: Average of optimal control, $c_1 = 6$, simple particle problem




Figure 4.12: State space density, $c_1 = 6$, simple particle problem

The interesting observation from this simulation is that the optimal control favors using the effect of the reflection to efficiently ‘push’ the process back into the interior. As the penalty for the reflection with $c_1 = 0.01$ is rather mild compared to the cost of using the control at full scale, increasingly less influence is enacted by the control as we move closer to both boundaries -1 and 1 . Figure 4.9 shows the optimal control when the costs of the reflection is increased to $c_1 = 1$. It reveals that with a higher penalty for the reflection, it is beneficial to use the control more extensively, although a similar pattern as in the previous case can be observed when the process approaches the boundaries of the state space. The control is used slightly less in this area to benefit from the reflection in direction of the origin. The overall heavier use of the control results in a state space density (Figure 4.10) which is more concentrated around the origin than the one from the previous case, see Figure 4.8. The value of the optimality criterion is given by $J_{n,m}^* = 0.42745$.

To illustrate an extreme case, we show a third example with $c_1 = 6$. Figure 4.11 shows the optimal control in this setting, Figure 4.12 displays the state space density. In contrast to the previous two cases, the optimal control tries to avoid a reflection under all circumstances by using its full force pushing back to the origin when the process approaches the boundaries of the state space. Still, a trade-off is made when the process is close to the origin, and the control is used with less than full force to avoid the costs induced by c_0 . The state space density concentrates even more around the origin in this setting. The value of the cost criterion is given by $J_{n,m}^* = 0.66399$.

5 Outlook

The considerations presented in the present paper can be extended in several ways. From a numerical analysis point of view, it is highly interesting how the numerical scheme behaves if higher-order basis elements are used to approximate the density p of the state space marginal of the occupation measure μ_0 . The analytic solution described in Section 4.1 has a density that is infinitely differentiable everywhere but at one point, thus justifying the use of, for example, piecewise linear basis functions. However, this would require an adaption of the presented convergence proof, in particular regarding the analysis leading up to the proof of Lemma 3.14. One aspect to be addressed is the fact that as soon as we use standard elements with a order larger than 1, for example, quadratic Lagrange elements, the non-negativity of the approximate density cannot be guaranteed by restricting the coefficients to be non-negative.

Another topic to research would be the introduction of adaptive meshing techniques for both state and control space. Analytic or heuristic error estimator could guide a successive refinement of the meshes, leading to a increase in accuracy without significantly higher computation time.

From a modeling point of view, on the one hand, an adaption of the discretization scheme for models featuring an unbounded state space would enhance the number of applications for this numerical scheme. Several control problems in finance and economics feature an unbounded state space, and are well suited for

the linear programming approach. Initial investigations show that models with an unbounded state space can be approximated using a bounded state space with reflection boundaries. A full analysis of this approach would allow us to use the methods presented in this paper in order to solve such models.

On the other hand, problems with finite time horizon or even optimal stopping problems could also be solved with similar numerical techniques. While the analytic linear programming approach to address such problems is well studied, the discretization techniques presented in this chapter would have to be enhanced to reflect the time dependency of both constraint functions and measures. A numerical analysis of such techniques was conducted in [17], but a convergence analysis remained unconsidered.

A Additional proofs

This appendix provides the proofs of Lemma 3.12, Lemma 3.13 and Lemma 3.14.

Proof of Lemma 3.12. Find $\hat{\epsilon}_1 < \epsilon$ such that $\lambda(\{x : p(x) \leq \hat{\epsilon}_1\}) < \frac{1}{2D_1}$, which is possible due to the continuity from above of measures. Define

$$\bar{p}(x) = \begin{cases} p(x) & p(x) > \hat{\epsilon}_1 \\ \hat{\epsilon}_1 & p(x) \leq \hat{\epsilon}_1 \end{cases}.$$

Then, $\|p - \bar{p}\|_{L^1(E)} \leq \hat{\epsilon}_1 \cdot \lambda(\{x : p(x) \leq \hat{\epsilon}_1\}) \leq \frac{\hat{\epsilon}_1}{2D_1}$. Now, choose m_0 large enough such that for all $m \geq m_0$, there is a $\tilde{p}_m \in \text{span}(p_0, p_1, \dots, p_{2^m-1})$ with $\|\bar{p} - \tilde{p}_m\|_{L^1(E)} \leq \frac{\hat{\epsilon}_1}{2D_1}$ and $\tilde{p}_m \geq \hat{\epsilon}_1$ (note there is no point in choosing $\tilde{p}_m < \hat{\epsilon}_1$ when approximating \bar{p}). Then,

$$\|p - \tilde{p}_m\|_{L^1(E)} \leq \|p - \bar{p}\|_{L^1(E)} + \|\bar{p} - \tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_1}{2D_1} + \frac{\hat{\epsilon}_1}{2D_1} = \frac{\hat{\epsilon}_1}{D_1}$$

holds. □

Proof of Lemma 3.13. Fix $\delta > 0$. Since $(\mu_0, \mu_1) \in \mathcal{M}_{n,\infty}$, we have that for each $k = 1, 2, \dots, n$

$$Rf_k = \int_E \int_U Af_k(x, u) \eta_0(du, x) p(x) dx + \int_E \int_U Bf_k(x, u) \mu_1(dx \times du)$$

and thereby for any \tilde{p}_m in the span of $\{p_0, p_1, \dots, p_{2^m-1}\}$

$$\begin{aligned} d_k^{(m)}(\tilde{p}_m) &= Rf_k - \int_E \int_U Af_k(x, u) \hat{\eta}_{0,m}(du, x) \tilde{p}_m(x) dx - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx) \\ &= \int_E \int_U Af_k(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U Af_k(x, u) \hat{\eta}_{0,m}(du, x) \tilde{p}_m(x) dx \\ &\quad + \int_E \int_U Bf_k(x, u) \eta_1(du, x) \mu_{1,E}(dx) - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx) \end{aligned}$$

holds. The triangle inequality reveals that

$$\begin{aligned} |d_k^{(m)}(\tilde{p}_m)| &\leq \left| \int_E \int_U Af_k(x, u) \eta_0(du, x) p(x) dx - \int_E \int_U Af_k(x, u) \hat{\eta}_{0,m}(du, x) p(x) dx \right| \\ &\quad + \left| \int_E \int_U Af_k(x, u) \hat{\eta}_{0,m}(du, x) p(x) dx - \int_E \int_U Af_k(x, u) \hat{\eta}_{0,m}(du, x) \tilde{p}_m(x) dx \right| \\ &\quad + \left| \int_E \int_U Bf_k(x, u) \eta_1(du, x) \mu_{1,E}(dx) - \int_E \int_U Bf_k(x, u) \hat{\eta}_{1,m}(du, x) \mu_{1,E}(dx) \right| \\ &\equiv |d_{k,1}^{(m)}| + |d_{k,2}^{(m)}| + |d_{k,3}^{(m)}|. \end{aligned}$$

Apply Lemma 3.12 with $\epsilon = \delta$ and $D_1 = D_2 \cdot 3 \cdot \max\{\bar{A}, 1\}$. Take $\hat{\epsilon}_1$ and m_1 from this result. Set $\hat{\epsilon}_2 = \hat{\epsilon}_1/D_2$. Then, $\hat{\epsilon}_2 \leq \delta$ and for all $m \geq m_1$, there is a $\tilde{p}_m \in \text{span}\{p_0, p_1, \dots, p_{2^m-1}\}$ such that $\|p - \tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_2}{3 \cdot \max\{\bar{A}, 1\}}$ as well as $\tilde{p}_m \geq D_2 \cdot \hat{\epsilon}_2$ holds. Also,

$$\left| d_{k,2}^{(m)} \right| \equiv \left| \int_E \int_U A f_k(x, u) \hat{\eta}_{0,m}(du, x) (p(x) - \tilde{p}_m(x)) dx \right| \leq \bar{A} \|p - \tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_2}{3}.$$

By Proposition 3.8, we can choose $m_2 \geq m_1$ such that for all $m \geq m_2$, $|d_{k,1}^m|$ is bounded by $\frac{\hat{\epsilon}_2}{3}$. By Proposition 3.9, we can choose $m_3 \geq m_2$ such that $|d_{k,3}^{(m)}|$ is bounded by $\frac{\hat{\epsilon}_2}{3}$ for all $m \geq m_3$, which shows that $|d_k^{(m)}(\tilde{p}_m)| < \hat{\epsilon}_2$ for $k = 1, 2, \dots, n$. For $k = n+1$, since p is a probability density,

$$\|\tilde{p}_m\|_{L^1(E)} \leq \|\tilde{p}_m - p\|_{L^1(E)} + \|p\|_{L^1(E)} < \frac{\hat{\epsilon}_2}{3 \max\{\bar{A}, 1\}} + 1 < \hat{\epsilon}_2 + 1.$$

Now assume that $\|\tilde{p}_m\|_{L^1(E)} < 1 - \hat{\epsilon}_2$. Then,

$$\|p\|_{L^1(E)} \leq \|p - \tilde{p}_m\|_{L^1(E)} + \|\tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_2}{3 \max\{\bar{A}, 1\}} + 1 - \hat{\epsilon}_2 < \hat{\epsilon}_2 + 1 - \hat{\epsilon}_2 = 1,$$

a contradiction, and we have that

$$1 - \hat{\epsilon}_2 \leq \|\tilde{p}_m\|_{L^1(E)} \leq 1 + \hat{\epsilon}_2.$$

Hence, $|d_{n+1}^{(m)}(\tilde{p}_m)| < \hat{\epsilon}_2$, which completes the proof, upon setting $m_0 = m_3$. \square

Now we can show that the statement of Lemma 3.14 is true.

Proof of Lemma 3.14. Fix $\vartheta > 0$. Select $m_1 \in \mathbb{N}$ large enough such that for all $m \geq m_1$, $C^{(m)}$ has full rank and thus $n+1$ independent columns. For any $m \geq m_1$, let $\bar{C}^{(m)} \in \mathbb{R}^{n+1, n+1}$ be a matrix consisting of $n+1$ independent columns of $C^{(m)}$. Set

$$\delta = \frac{\vartheta}{\max\left\{1, \left\|(\bar{C}^{(m_1)})^{-1}\right\|_{\infty}\right\}}$$

and by Lemma 3.13, with δ and $D_2 = \max\{1, \|(\bar{C}^{(m_1)})^{-1}\|_{\infty}\}$, find $m_2 \geq m_1$ such that for all $m \geq m_2$, there is a \tilde{p}_m , with $\|d^{(m)}(\tilde{p}_m)\|_{\infty} < \hat{\epsilon}_2 \leq \delta$, for some $\hat{\epsilon}_2 > 0$, satisfying

$$\tilde{p}_m \geq \max\left\{1, \left\|(\bar{C}^{(m_1)})^{-1}\right\|_{\infty}\right\} \cdot \hat{\epsilon}_2 \geq \left\|(\bar{C}^{(m_1)})^{-1}\right\|_{\infty} \cdot \hat{\epsilon}_2$$

as well as $\|p - \tilde{p}_m\|_{L^1(E)} < \frac{\hat{\epsilon}_2}{3 \cdot \max\{\bar{A}, 1\}}$. Set $\hat{\vartheta} = \max\left\{1, \left\|(\bar{C}^{(m_1)})^{-1}\right\|_{\infty}\right\} \cdot \hat{\epsilon}_2$ and note that $\frac{\hat{\epsilon}_2}{3 \cdot \max\{\bar{A}, 1\}} < \hat{\vartheta}$.

Consider the solution $\tilde{y} \in \mathbb{R}^{2^{m_1}}$ for $C^{(m_1)}y = -d^{(m_2)}(\tilde{p}_m)$ that is given by injecting $\tilde{y} = (\bar{C}^{(m_1)})^{-1}(-d^{(m_2)}(\tilde{p}_m)) \in \mathbb{R}^{n+1}$ into $\mathbb{R}^{2^{m_1}}$. Then,

$$\|\tilde{y}\|_{\infty} = \|\tilde{y}\|_{\infty} = \left\|(\bar{C}^{(m_1)})^{-1} d^{(m_2)}(\tilde{p}_m)\right\|_{\infty} < \left\|(\bar{C}^{(m_1)})^{-1}\right\|_{\infty} \|d^{(m_2)}(\tilde{p}_m)\|_{\infty} \leq \hat{\vartheta}.$$

We now show that there is a solution \tilde{y} to $C^{(m_2)}y = d^{(m_2)}(\tilde{p}_m)$ that satisfies $\|\tilde{y}\|_{\infty} \leq \hat{\vartheta}$. By the definition of the constraint matrix, for any $m \in \mathbb{N}$, we have that for $k = 1, 2, \dots, n+1$ and $i = 0, 1, \dots, 2^m - 1$,

$$C_{k,2i}^{(m+1)} + C_{k,2i+1}^{(m+1)} = C_{k,i}^{(m)}$$

holds. Indeed, since for $1 \leq k \leq n$, by the choice of basis functions $\{p_0, p_1, \dots, p_{2^m-1}\}$ as indicator functions over dyadic intervals, the entries of $C_{k,i}^{(m)}$ are given by integration of the functions $A f_k$ over intervals that

are cut in half, and if $k = n + 1$, the entries are simply given by the interval lengths $(x_{j+1} - x_j)$ since $p_j = 1$ on $[x_{j+1}, x_j)$. Hence, if y is a solution to $C^{(m)}y = -d$, the vector $\bar{y} \in \mathbb{R}^{2^{m+1}}$ with components

$$\bar{y}_{2i+1} = \bar{y}_{2i} = y_i$$

where $i = 0, 1, \dots, 2^m - 1$, satisfies $C^{(m+1)}y = -d$, and $\|y\|_\infty = \|\bar{y}\|_\infty$ holds. Inductively, this reveals that for any $m \geq m_1$, there is a solution \tilde{y} to $C^{(m)}y = -d^{(m_2)}(\tilde{p}_m)$ which satisfies $\|\tilde{y}\|_\infty = \|y\|_\infty \leq \hat{\vartheta}$. In particular, this means that there is a solution \tilde{y} to $C^{(m_2)}y = -d^{(m_2)}(\tilde{p}_m)$, with $\|\tilde{y}\|_\infty = \|y\|_\infty < \hat{\vartheta}$. For any $m \geq m_2$, this analysis can be conducted similarly, showing the result for $m_0 = m_2$. \square

References

- [1] J. ANSELM, F. DUFOUR, AND T. PRIETO-RUMEAU, *Computable approximations for continuous-time Markov decision processes on Borel spaces based on empirical measures*, J. Math. Anal. Appl., 443 (2016), pp. 1323–1361.
- [2] A. G. BHATT AND V. S. BORKAR, *Occupation measures for controlled Markov processes: characterization and optimality*, Ann. Probab., 24 (1996), pp. 1531–1562.
- [3] P. BILLINGSLEY, *Convergence of probability measures*, Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons, Inc., New York, Second ed., 1999. A Wiley-Interscience Publication.
- [4] V. I. BOGACHEV, *Measure theory. Vol. I, II*, Springer-Verlag, Berlin, 2007.
- [5] C. DE BOOR, *A practical guide to splines*, vol. 27 of Applied Mathematical Sciences, Springer-Verlag, New York, revised ed., 2001.
- [6] W. H. FLEMING AND R. W. RISHEL, *Deterministic and stochastic optimal control*, Springer-Verlag, Berlin-New York, 1975. Applications of Mathematics, No. 1.
- [7] W. H. FLEMING AND H. M. SONER, *Controlled Markov processes and viscosity solutions*, vol. 25 of Stochastic Modelling and Applied Probability, Springer, New York, second ed., 2006.
- [8] C. A. HALL AND W. W. MEYER, *Optimal error bounds for cubic spline interpolation*, J. Approximation Theory, 16 (1976), pp. 105–122.
- [9] K. HELMES, S. RÖHL, AND R. H. STOCKBRIDGE, *Computing moments of the exit time distribution for Markov processes by linear programming*, Oper. Res., 49 (2001), pp. 516–530.
- [10] K. HELMES AND R. H. STOCKBRIDGE, *Determining the optimal control of singular stochastic processes using linear programming*, in Markov processes and related topics: A Festschrift for Thomas G. Kurtz, vol. 4 of Inst. Math. Stat. (IMS) Collect., Inst. Math. Statist., Beachwood, OH, 2008, pp. 137–153.
- [11] P. KACZMAREK, S. T. KENT, G. A. RUS, R. H. STOCKBRIDGE, AND B. A. WADE, *Numerical solution of a long-term average control problem for singular stochastic processes*, Math. Methods Oper. Res., 66 (2007), pp. 451–473.
- [12] S. KUMAR AND K. MUTHURAMAN, *A numerical method for solving singular stochastic control problems*, Operations Research, 52 (2004), pp. 563–582.
- [13] T. G. KURTZ AND R. H. STOCKBRIDGE, *Existence of Markov controls and characterization of optimal Markov controls*, SIAM J. Control Optim., 36 (1998), pp. 609–653.
- [14] T. G. KURTZ AND R. H. STOCKBRIDGE, *Linear programming formulations of singular stochastic processes*, [arXiv:1707.09209], (2017).

- [15] H. J. KUSHNER AND P. DUPUIS, *Numerical methods for stochastic control problems in continuous time*, vol. 24 of Applications of Mathematics (New York), Springer-Verlag, New York, Second ed., 2001. Stochastic Modelling and Applied Probability.
- [16] J.-B. LASSERRE AND T. PRIETO-RUMEAU, *SDP vs. LP relaxations for the moment approach in some performance evaluation problems*, Stoch. Models, 20 (2004), pp. 439–456.
- [17] M. LUTZ, *A linear programming approach for optimal exercise and valuation of American lookback options*, Master's thesis, University of Wisconsin-Milwaukee, 2007.
- [18] A. S. MANNE, *Linear programming and sequential decisions*, Management Sci., 6 (1960), pp. 259–267.
- [19] M. S. MENDIONDO AND R. H. STOCKBRIDGE, *Approximation of infinite-dimensional linear programming problems which arise in stochastic control*, SIAM J. Control Optim., 36 (1998), pp. 1448–1472.
- [20] G. A. RUS, *Finite element methods for control of singular stochastic processes*, PhD thesis, 2009. The University of Wisconsin - Milwaukee.
- [21] R. SERRANO, *On the lp formulation in measure spaces of optimal control problems for jump-diffusions*, Systems & Control Letters, 85 (2015), pp. 33–36.
- [22] M. I. TAKSAR, *Infinite-dimensional linear programming approach to singular stochastic control*, SIAM Journal on Control and Optimization, 35 (1997), pp. 604–625.
- [23] M. G. VIETEN, *Numerical Solution of Stochastic Control Problems Using the Finite Element Method*, PhD thesis, 2018. The University of Wisconsin - Milwaukee.