Calhoun: The NPS Institutional Archive

Faculty and Researcher Publications	Faculty and Researcher Publications Collection

2016

Optimal Control of Uncertain Systems Using Sample Average Approximations

Phelps, Chris

Society for Industrial and Applied Mathematics

OPTIMAL CONTROL OF UNCERTAIN SYSTEMS USING SAMPLE AVERAGE APPROXIMATIONS*

CHRIS PHELPS†, JOHANNES O. ROYSET‡, AND QI GONG†

Abstract. In this paper, we introduce the *uncertain optimal control problem* of determining a control that minimizes the expectation of an objective functional for a system with parameter uncertainty in both dynamics and objective. We present a computational framework for the numerical solution of this problem, wherein an independently drawn random sample is taken from the space of uncertain parameters, and the expectation in the objective functional is approximated by a sample average. The result is a sequence of approximating standard optimal control problems that can be solved using existing techniques. To analyze the performance of this computational framework, we develop necessary conditions for both the original and approximate problems and show that the approximation based on sample averages is consistent in the sense of Polak [*Optimization: Algorithms and Consistent Approximations*, Springer, New York, 1997]. This property guarantees that accumulation points of a sequence of global minimizers (stationary points) of the approximate problem are global minimizers (stationary points) of the original problem. We show that the uncertain optimal control problem can further be approximated in a consistent manner by a sequence of nonlinear programs under mild regularity assumptions. In numerical examples, we demonstrate that the framework enables the solution of optimal search and optimal ensemble control problems.

Key words. optimal control, numerical methods, parameter uncertainty

AMS subject classifications. Primary, 49M25; Secondary, 49K45

DOI. 10.1137/140983161

1. Introduction. In this paper we consider an extension of nonlinear optimal control for problems with Mayer-type objective functional to a setting with parameter uncertainty. We introduce the *uncertain optimal control problem* (UOCP), where the objective functional and system dynamics depend on stochastic parameters and the goal is to find a control that minimizes the expected value of the objective. The UOCP addresses a number of emerging applications in optimal control that require design of an open-loop control for an uncertain system such as those arising in optimal search or ensemble control. In the optimal search problem, the goal is to design a search plan to maximize the probability of detecting a moving target with unknown location or goals [30]. In the ensemble control problem, the goal is to determine a single open-loop control for a large number of structurally identical systems with parameter variation [27], which can be viewed as a single system with stochastic parameters [40]. In addition, existing control problems such as trajectory optimization may benefit from a problem formulation that incorporates inherent uncertainty in dynamical models, environment [17, 46], and behavior of other agents [11, 37, 44]. In this paper, we develop a computational framework for the UOCP as well as necessary conditions for validation and verification of solutions.

Specifically, the UOCP is the following problem: Find an initial state and control

pair $\eta = (\xi^\eta, u^\eta)$ that minimizes the objective functional

$$(1.1) \qquad J(\eta) = \mathbb{E}^P[F(x^\eta(1, \omega), \omega)],$$

where $\mathbb{E}^P$ is the expectation on the complete probability space (Ω, Σ, P) and $\omega \in \Omega$. Furthermore, $x^\eta(t, \omega)$ is the solution to the uncertain dynamical system

$$(1.2) \qquad \dot{x}^\eta(t, \omega) = f(x^\eta(t, \omega), u^\eta(t), \omega), \qquad x^\eta(0, \omega) = \xi^\eta + \iota(\omega)$$

almost surely. Here, $\xi^\eta \in \mathbb{R}^n$, $u^\eta : [0,1] \mapsto \mathbb{R}^m$, $x^\eta : [0,1] \times \Omega \mapsto \mathbb{R}^n$, $\iota : \Omega \mapsto \mathbb{R}^n$, $f : \mathbb{R}^n \times \mathbb{R}^m \times \Omega \mapsto \mathbb{R}^n$, and $F : \mathbb{R}^n \times \Omega \mapsto \mathbb{R}$. We note that for a fixed $\omega \in \Omega$, (1.2) is a standard deterministic dynamical system, and therefore the existence and uniqueness of the solution are guaranteed under suitable regularity conditions. Such conditions then also ensure that the objective functional (1.1) is well defined.

The UOCP is an extension of the constrained nonlinear optimal control problem with a Mayer objective functional to scenarios with parameter uncertainty. That is, the standard constrained nonlinear optimal control problem with no parameter uncertainty can be viewed as a special case of the UOCP, in which the space Ω of stochastic parameters contains only a single element. In the last few decades, a variety of *direct methods* have been developed for solving constrained nonlinear optimal control problems. These methods are based on techniques for approximating a continuous-time optimal control problem by a discretized finite-dimensional nonlinear program, and have achieved substantial practical success in many areas [6, 7, 9, 29]. A number of numerical schemes have been employed for this purpose, including Euler [35, Chapter 4], Runge–Kutta [25, 42], and pseudospectral [19, 26, 38]. However, the apparent simplicity of such discretization schemes belies deep theoretical issues involved in the approximation of an optimization problem. When employing a direct approximation method it is necessary to perform careful analysis in order to demonstrate that the numerical scheme provides a valid approximation to the original problem. Indeed, even for standard optimal control problems, there are counterexamples showing that an inappropriately designed discretization may not be convergent [10]. The aim of this work is to extend the direct methods to the new problem setting with parameter uncertainty. Using the theory of consistent approximations developed by Polak [35, Chapter 3], we demonstrate that a numerical scheme based on sample average approximations produces a meaningful approximation to the UOCP.

A number of related consistency results have recently been developed for two problems related to the UOCP, including the *conditionally deterministic optimal search problem* (CDOSP). The CDOSP is an optimal control problem in which the goal is to determine the optimal search trajectory for an agent attempting to detect a target whose motion depends on some unknown parameter, such as initial position or starting time; see [9, 14, 15]. This problem is a special case of the UOCP where the unknown parameter appears only in the objective functional and not in the system dynamics. In [9], a numerical algorithm based on the approximation of the expectation by a left-point numerical integration rule is provided for a special case of searching for a target moving at constant velocity in a channel. In [14, 15] the authors use a composite-Simpson integration scheme to discretize a two-dimensional parameter space and develop a computational method for solving the CDOSP. They also analyze the performance of the computational method using Polak's consistent approximation theory [35, section 3.3]. In [32, 33] the authors show that this approach can be applied to other parameter spaces and numerical integration schemes.

Furthermore, they provide a Pontryagin-like minimum principle for solutions that are accumulation points of a sequence of global minimizers of the approximate problem.

A similar approach has been taken in the solution of the *ensemble control problem*, where the goal is to determine a single open-loop controller that simultaneously steers a large number of similar systems from an initial state to a desired final state. The ensemble is represented in this framework as a continuum of structurally identical dynamical systems indexed by parameters. Interest in this problem arises from the difficulty of designing compensating pulses for magnetic resonance experiments. Necessary and sufficient conditions for ensemble controllability are provided for a number of linear and bilinear systems in [27, 29]; however, such conditions for the general nonlinear ensemble are absent. Even when the ensemble is controllable, determination of the desired control is difficult, and it is not surprising that such a control is available in closed form for only a small number of systems. Therefore a new class of computational methods must be developed to provide numerical solutions to ensemble control problems.

One approach to overcoming the difficulty of ensemble control is to leverage existing results on computational optimal control. By developing an objective functional for the ensemble control problem, it is possible to formulate an optimal control problem such that the calculated control achieves the desired state transfer of the ensemble. A formulation for such an *optimal ensemble control problem* (OECP) is suggested in [28], wherein the goal is to minimize the expectation of the square error of the final state. By interpreting the ensemble as a single dynamical system with stochastic parameters, this problem can be fit into the UOCP framework [40]. In [39, 40, 41] the authors extend the pseudospectral optimal control method [19, 20] to the OECP using a quadrature scheme. In this approach, the parameter space is discretized according to an Legendre–Gauss–Lobatto (LGL)-quadrature scheme, and the time domain is discretized using an LGL-pseudospectral method, resulting in a high-dimensional nonlinear program for which a variety of robust and efficient numerical algorithms are available. The approach is applied to nuclear magnetic resonance [39, 40] and neuroscience applications [41] with one- and two-dimensional parameter spaces.

Although the approach of using quadrature schemes to approximate the objective functional has been successfully applied to a number of systems for both the OECP and the CDOSP, it is inherently limited to problems with few parameters, as quadrature-based numerical integration methods are known to be computationally expensive when applied to high-dimensional spaces. In the literature, other techniques have been used to approximate an uncertain dynamical system by a large but finite number of decoupled deterministic systems. The stochastic collocation [5] method is similar to the aforementioned quadrature scheme in that a finite number of nodes are selected from the parameter space, and the dynamical system is approximated by a tensor product of polynomials. The polynomial chaos method approximates the random state vector by using a Galerkin projection onto a set of orthogonal polynomials [45]. This approach has also been used in an optimal control setting with parameter uncertainty for special cases of the UOCP [12, 13, 23, 24]. However, these techniques suffer from the same curse of dimensionality as the quadrature approach. Indeed, as the number of parameters increases, the size of the approximate optimal control problem increases exponentially, and numerical solution of the approximate problem quickly becomes intractable. Still, an algorithm to solve the UOCP for high-dimensional parameter spaces is desirable, as the inclusion of a large number of parameters allows the formulation to be used for applications with more robust physical models.

In this paper, we develop a computational framework for the solution of the UOCP based on sample average approximations. In our approach, a random, independently distributed sample is taken from the space of stochastic parameters. Then, the expectation in the objective functional of the control problem is approximated by the sample average. The resulting family of standard optimal control problems can be solved using existing optimal control algorithms, such as Euler (see [35, Chapter 4]), Runge–Kutta [21, 42], and pseudospectral [26, 38] methods. We refer the reader to [2, 4] for early work on sample average approximations to stochastic optimization, which also provides our foundation. For a treatment of cases in finite dimensions, see [43]. The use of Monte Carlo sampling methods for integration avoids computational issues inherent in quadrature-based approaches, as the number of nodes, and therefore the dimension of the approximate control problem, does not depend on the number of stochastic parameters. Consequently, as the number of stochastic parameters grows, Monte Carlo methods become more favorable. Since the rate of convergence of these methods is only 1 over the square root of the sample size, the proposed framework might still involve substantial computational cost. However, we believe further algorithmic studies, beyond the scope of the present paper, might mitigate this somewhat, for example, through the development of variance reduction techniques. As with all direct methods for optimal control, it is essential to demonstrate that the proposed discretization scheme provides a valid approximation to the UOCP. It is therefore necessary to extend consistency results for standard optimal control problems [20, 35, 42] to the UOCP in order to assure that the computational framework based on sample averages provides meaningful results.

We provide three classes of results for sample average approximations of the UOCP. First, we employ an extension of the strong law of large numbers to random lower semicontinuous functions to address the convergence properties of the objective functional. We demonstrate that accumulation points of a sequence of global minimizers of the approximate problem are global minimizers of the original problem. This property shows that the sample average approximation is meaningful in the sense that if the sequence of optimal solutions to the approximate problem is seen to converge as the sample size approaches infinity, then the limit must be an optimal solution to the original problem. Second, we provide necessary optimality conditions for the UOCP and its approximation using optimality functions based on the $\mathcal{L}_2$-Frechet derivative of the objective functional, which may be used for validation and verification of numerical solutions. By again using the extension of the strong law of large numbers, we show that accumulation points of a sequence of stationary points of the approximate problem are stationary points of the original problem. Therefore, the sample average approximation scheme for the UOCP is consistent in the sense of Polak [35, section 3.3]. The above results deal with the approximation of the UOCP by sample averages. Since such approximations are standard optimal control problems, further approximations in the time domain in the usual manner are required. Third, we provide a consistency result for the approximate problem obtained by sample average approximations *as well as* time discretization using Euler's method. All of these results apply to both a constrained and an unconstrained formulation of the UOCP.

The paper is organized as follows: Section 2 formulates the UOCP. Section 3 introduces the approximation scheme based on sample averages and evaluates the convergence properties of the approximate problem as the sample size tends to infinity. Section 4 introduces optimality functions for the original and approximate problems, which provide necessary conditions for optimality. Section 5 demonstrates

a convergence property for the optimality functions, showing that the approximation based on sample averages is consistent in the sense of Polak. Section 6 shows that an approximate problem that is further subject to time discretization using Euler's method leads to a nonlinear program that is in some sense consistent with the UOCP. Section 7 applies the computational framework to two examples involving damping a harmonic oscillator with unknown natural frequency and detecting an intruder in a channel.

2. Formulation of the uncertain control problem. A number of recently considered applications in computational optimal control theory, such as optimal search [32, 33, 34] and ensemble control [39, 40, 41], require the calculation of an open-loop control for a system with stochastic parameters which will minimize the expectation of a predetermined cost functional. The UOCP, defined in section 1, is to find an initial state and control pair $\eta = (\xi^\eta, u^\eta)$ from a given admissible set to minimize the objective functional (1.1) subject to the uncertain dynamical system (1.2). In this work we develop a computational framework to solve two versions of the UOCP for which the admissible sets differ based on the nature of the control constraint.

Before we define these versions, we introduce the spaces on which we conduct our analysis. To develop optimality conditions, we make use of an inner product on the space of decision variables. Therefore we work in the $\mathcal{L}_2$ topology. Let $\mathcal{L}_2^m[0,1]$ be the space of all functions $v : [0,1] \mapsto \mathbb{R}^m$ such that $\int_0^1 \|v(t)\|^2 \, dt < \infty$. We carry out our analysis in a subspace of the Hilbert space

$$H_2 = \mathbb{R}^n \times \mathcal{L}_2^m[0,1],$$

where the inner product and norm on H_2 are defined for any $\eta = (\xi^\eta, u^\eta), \eta' = (\xi^{\eta'}, u^{\eta'}) \in H_2$ by

$$\langle \eta, \eta' \rangle_{H_2} = \langle \xi^\eta, \xi^{\eta'} \rangle + \langle u^\eta, u^{\eta'} \rangle_2.$$

Therefore the norm in H_2 is given by

$$\|\eta\|_{H_2}^2 = \|\xi^\eta\|^2 + \|u^\eta\|_2^2.$$

In this paper we address the two cases of the UOCP, where the control $u(t)$ is constrained to be in either a compact convex set or an open convex set in $\mathbb{R}^m$ for almost every $t \in [0,1]$. We therefore define the admissible sets for each of these problems as follows: Given compact, convex sets $\Xi_C \subset \mathbb{R}^n$ and $U_C \subset \mathbb{R}^m$, we define the set of admissible controls

$$\mathbf{U}_C = \{u \in L_2^m[0,1] \,|\, u(t) \in U_C \text{ for almost every } t \in [0,1]\}.$$

The set of all admissible state-control pairs for this problem is then given by $\mathbf{H}_C = \Xi_C \times \mathbf{U}_C$. Similarly, given bounded, open, convex sets $\Xi_O \subset \mathbb{R}^n$ and $U_O \subset \mathbb{R}$, we define the set of admissible controls

$$\mathbf{U}_O = \{u \in L_2^m[0,1] \,|\, u(t) \in U_O \text{ for almost every } t \in [0,1]\}.$$

The set of all admissible state-control pairs for this problem is then given by $\mathbf{H}_O = \Xi_O \times \mathbf{U}_O$.

The sets $\mathbf{H}_C$ and $\mathbf{H}_O$ are subsets of the pre-Hilbert space $H_{\infty,2} = \{(\xi, u) \in H_2 \,|\, \|u\|_\infty < \infty\}$. For mathematical convenience, we assume $\Xi_O \subset \Xi_C$ and $U_O \subset$

U_C so that $\mathbf{H}_O \subset \mathbf{H}_C$. We observe that in this work we define the admissible set differently from how it is done in Polak [35, Chapter 4], which requires the pointwise control constraint be satisfied for all $t \in [0,1]$. Let $U \subset \mathbb{R}^m$. We note that for each $u \in \mathcal{L}_2^m[0,1]$ with $u(t) \in U$ for almost every $t \in [0,1]$, there is a member $\tilde{u}$ of its equivalence class such that $\tilde{u}(t) \in U$ for every $t \in [0,1]$. Therefore for any given constraint set U, we can apply the standard results from the theory of differential equations to controls from our admissible set.

In developing optimality conditions, we evaluate derivatives with respect to the decision variable η. In order to guarantee that these derivatives exist, we work on a space $\mathbf{H}$ which is slightly larger than $\mathbf{H}_C$. To define the space $\mathbf{H}$, let $\rho_1, \rho_2 \in \mathbb{R}$ be constants large enough so that $\|\xi^\eta\| < \rho_1$, $\|u^\eta\|_\infty < \rho_2$ for all $\eta \in \mathbf{H}_C$. The existence of these constants is guaranteed by the compactness of Ξ_C and U_C. Now let $\mathbf{H} = \{(\xi, u) \in \mathbb{R}^n \times \mathcal{L}_2^m[0,1]|\, \|\xi\| < \rho_1, \|u\|_\infty < \rho_2\}$. The space $\mathbf{H}$ is open in the $\mathcal{L}_\infty$ topology, and the inclusion $\mathbf{H}_O \subset \mathbf{H}_C \subset \mathbf{H}$ holds. The reader should note that all convergence results on the sets $\mathbf{H}_O, \mathbf{H}_C$, and $\mathbf{H}$ are with respect to the $\mathcal{L}_2$ topology.

With the appropriate function spaces defined, we now state the UOCPs that are the focus of this work.

Problem B_C: Find an initial state and control pair $\eta = (\xi^\eta, u^\eta) \in \mathbf{H}_C$ to minimize the objective functional (1.1) subject to the uncertain dynamical system (1.2).

Problem B_O: Find an initial state and control pair $\eta = (\xi^\eta, u^\eta) \in \mathbf{H}_O$ to minimize the objective functional (1.1) subject to the uncertain dynamical system (1.2).

The set U_C determines a pointwise control constraint for Problem B_C. Because U_C is closed, this formulation can be used to approach uncertain optimal control problems with pointwise inequality control constraints, as long as the set of points which satisfy these constraints is convex. Problem B_O can be used to approach unconstrained optimal control problems by making U_O large enough that all reasonable controls lie in the admissible set. To conduct an analysis of Problems B_C and B_O we need the following regularity assumptions.

ASSUMPTION 1. *There exists a compact set $X_0 \subset \mathbb{R}^m$ such that for each $\eta \in \mathbf{H}$, $x^\eta(t, \omega) \in X_0$ for all $t \in [0,1]$, $\omega \in \Omega$, where x^η is the solution to (1.2) for $\eta = (\xi^\eta, u^\eta)$.*

This assumption essentially requires that there does not exist $\omega \in \Omega$ such that the dynamical system given by $f(\cdot, \cdot, \omega)$ has a finite escape time. This assumption will be valid for a number of dynamical systems frequently encountered in control problems, for example, input-to-state stable systems and systems for which f is globally Lipschitz or satisfies a linear growth condition in the state variable.

ASSUMPTION 2. *For the set X_0 defined in Assumption 1 and the set $V = \{v \in \mathbb{R}^m|\, \|v\| < \rho_2\}$, for each $\omega \in \Omega$ the function $f(\cdot, \cdot, \omega)$ is continuously differentiable on $X_0 \times V$, and for each $x \in X_0$, $v \in V$, $f(x, v, \cdot)$, it is measurable and bounded on Ω. Furthermore, there exists a measurable function $L_f : \Omega \mapsto [1, \infty)$ such that for all $x', x'' \in X_0$ and $v', v'' \in V$, the following inequalities hold for every $\omega \in \Omega$:*

$$\|f(x', v', \omega) - f(x'', v'', \omega)\| \le L_f(\omega) \left[\|x' - x''\| + \|v' - v''\|\right],$$
$$\|f_x(x', v', \omega) - f_x(x'', v'', \omega)\| \le L_f(\omega) \left[\|x' - x''\| + \|v' - v''\|\right],$$
$$\|f_u(x', v', \omega) - f_u(x'', v'', \omega)\| \le L_f(\omega) \left[\|x' - x''\| + \|v' - v''\|\right].$$

ASSUMPTION 3. *For the set X_0 defined in Assumption 1, $F(\cdot, \omega)$ is continuously differentiable on X_0 for each $\omega \in \Omega$, and $F(x, \cdot), F_x(x, \cdot)$ are measurable for each $x \in X_0$. Furthermore, there exists a measurable function $L_F : \Omega \mapsto [1, \infty)$ such that*

for any $x', x'' \in X_0$, the following inequalities hold for every $\omega \in \Omega$:

$$|F(x', \omega) - F(x'', \omega)| \leq L_F(\omega) \|x' - x''\|,$$
$$\|F_x(x', \omega) - F_x(x'', \omega)\| \leq L_F(\omega) \|x' - x''\|.$$

Assumptions 2–3 require the differentiability of the functions in the problem formulation with respect to the states and controls, as well as measurability and integrability of the Lipschitz constant with respect to the stochastic parameter ω. These assumptions will be valid for a variety of problem frameworks in physical and other applications. For instance, in previously considered optimal search [32, 33] and ensemble control settings [39, 40, 41], the parameter space is a compact subspace of $\mathbb{R}^n$, and the functions in the problem formulation are sufficiently smooth; therefore Assumptions 2–3 are valid in these cases. These assumptions are used later to establish convergence properties and optimality conditions for the UOCPs.

In order to facilitate the analysis of the computational framework for the UOCP which is the focus of this work, we first state the following results on uncertain dynamical systems.

PROPOSITION 2.1. *Suppose that Assumptions* 1–2 *are satisfied. Then, for any $\eta \in \mathbf{H}$, the uncertain dynamical systems* (1.2) *have a unique solution $x^\eta(\cdot, \omega)$ for each $\omega \in \Omega$.*

Proof. The proof follows directly from Proposition 5.6.5 of [35]. □

PROPOSITION 2.2 (Carathéodory functions are jointly measurable [1, Lemma 4.51]). *Let (S, Σ) be a measurable space, X a separable metric space, and Y a metrizable space. Let $f : X \times S \mapsto Y$ be a function such that*
  (i) *for each $x \in X$, $f(x, \cdot) : S \mapsto Y$ is measurable;*
  (ii) *for each $s \in S$, $f(\cdot, s) : X \mapsto Y$ is continuous.*
Then f is called a Carathéodory function, and $f : X \times S \mapsto Y$ is jointly measurable.

LEMMA 2.3 (see [31, Lemma 4.1.3]). *Suppose that Assumption* 1 *is satisfied, and let V be the set defined in Assumption* 2. *Let $\kappa : \mathbb{R}^l \times V \times \Omega \mapsto \mathbb{R}^l$ be such that $\kappa(\cdot, \cdot, \omega)$ is continuously differentiable for each $\omega \in \Omega$, and $\kappa(x, u, \cdot)$ is measurable for each $x \in \mathbb{R}^l, v \in V$. Suppose also that there exists a measurable function $K : \Omega \mapsto [1, \infty)$ such that for every $x, x' \in \mathbb{R}^n$, $v, v' \in V$, and $\omega \in \Omega$,*

$$\|\kappa(x, v, \omega) - \kappa(x', v', \omega)\| \leq K(\omega) [\|x - x'\| + \|v - v'\|].$$

For each $\eta = (\xi^\eta, u^\eta) \in \mathbf{H}$, $\omega \in \Omega$, let $\chi^\eta : [0, 1] \times \Omega \to \mathbb{R}^l$ be the solution to

$$\dot{\chi}^\eta(t, \omega) = \kappa(\chi^\eta(t, \omega), u^\eta(t), \omega), \qquad \chi(0) = \xi^\eta.$$

Then χ^η is measurable, and for each $\omega \in \Omega$ we have

$$\left\| x^{\eta'}(t, \omega) - x^{\eta''}(t, \omega) \right\| \leq \sqrt{2} K(\omega) e^{K(\omega)} \|\eta' - \eta''\|_{H_2}.$$

3. Approximation of the UOCP using a sample average scheme. In this section we introduce the approximate optimal control problem based on a sample average scheme. Sample average approximations have been successfully applied to a wide variety of problems from the field of stochastic optimization with finite-dimensional decision spaces [43]. In the sample average approach, a random sample of parameter values is drawn from the parameter space, and the expectation in the objective functional is approximated by the sample mean. When the sample average approximation is applied to a stochastic programming problem with a finite-dimensional decision

space, this process results in a sequence of approximating nonlinear programming (NLP) problems. In this work we use the sample average method to approximate the UOCP, which has an infinite-dimensional decision space. The resulting approximate problem is a standard optimal control problem that can be solved using existing techniques from the field of control theory [8]. In addition, we use an extension of the strong law of large numbers (see [2, 4]) to analyze the convergence properties of such an approximation.

To apply this approximation scheme, for a given sample size M, we take an independent P-distributed sample $\{\omega_1, \omega_2, \ldots, \omega_M\}$ from the parameter space Ω and approximate the objective functional (1.1) by the sample average

$$(3.1) \qquad J^M(\eta) \doteq \frac{1}{M} \sum_{i=1}^{M} F\big(x^\eta(1, \omega_i), \omega_i\big).$$

The approximate uncertain optimal control problems can then be stated as follows.

Problem B_C^M: Find $\eta \in \mathbf{H}_C$ to minimize the objective functional (3.1), where x^η is the solution to the uncertain dynamical system (1.2).

Problem B_O^M: Find $\eta \in \mathbf{H}_O$ to minimize the objective functional (3.1), where x^η is the solution to the uncertain dynamical system (1.2).

Problems B_C^M and B_O^M are standard optimal control problems which can be solved using existing techniques. We discuss the convergence properties of Problems B_C^M and B_O^M in the context of epiconvergence of the objective functionals. The concept of epiconvergence provides a natural framework to analyze the approximation of an optimization problem, as it allows us to discuss the convergence of the inf and arg min operators.

3.1. Preliminary results on epiconvergence and random lower semi-continuous functions. To leverage existing results on the convergence of sample average approximations, we recall the concepts of epiconvergence and random lower semicontinuous functions.

DEFINITION 3.1 (see [3]). *Let (X, d) be a separable complete metric space. Consider the sequence of lower semicontinuous functions $f_M : X \mapsto \mathbb{R}$. We say that f_M epiconverges to f, denoted $f_M \to^{\mathrm{epi}} f$, if and only if*

(i) *$\liminf f_M(x_M) \geq f(x)$ whenever $x_M \to x$,*

(ii) *$\lim f_M(x_M) = f(x)$ for at least one sequence $x_M \to x$.*

DEFINITION 3.2 (see [4]). *Let (X, d) be a separable complete metric space with $\mathcal{B}$ the Borel sigma-field. Let P be a probability measure on the measurable space (Ω, Σ) such that Σ is P-complete. A function $f : X \times \Omega \mapsto \mathbb{R}$ is a random lower semicontinuous function if and only if*

(i) *for all $\omega \in \Omega$, the function $x \mapsto f(x, \omega)$ is lower semicontinuous,*

(ii) *$(x, \omega) \mapsto f(x, \omega)$ is $\mathcal{B} \otimes \Sigma$ measurable.*

In probability theory, the strong law of large numbers guarantees the almost sure convergence of the sample average as the number of samples drawn approaches infinity. The following proposition extends this result to random lower semicontinuous functions.

PROPOSITION 3.3 (see [2, Theorem 2.3]). *Let (Ω, Σ, P) be a probability space such that Σ is P-complete. Let (X, d) be a separable complete metric space. Suppose that the function $f : X \times \Omega \mapsto \mathbb{R}$ is a random lower semicontinuous function and there exists an integrable function $a_0 : \Omega \mapsto \mathbb{R}$ such that $f(x, \omega) \geq a_0(\omega)$ almost surely. Let*

$\{\omega_1, \ldots, \omega_M\}$ *be an independent P-distributed random draw, and define*

$$\hat{f}(\cdot, \omega_1, \ldots, \omega_M) = \frac{1}{M} \sum_{i=1}^{M} f(\cdot, \omega_i).$$

Then, as $M \to \infty, \hat{f}(\cdot, \omega_1, \ldots, \omega_M)$ *epiconverges almost surely to* $\mathbb{E}^P f(\cdot, \omega)$.

In general, when approximating an optimization problem, it is difficult to guarantee that a global minimizer of the original problem will be an accumulation point of a sequence of global minimizers of the approximate problem. However, the following result provides a useful convergence property for the approximate arg min and inf operators when the sequence of approximating objective functionals is epiconvergent.

PROPOSITION 3.4 (see [3, Theorem 2.5]). *Let* (X, d) *be a separable complete metric space. Consider a sequence of lower semicontinuous functions* $f_M : X \mapsto \mathbb{R}$ *such that* f_M *epiconverges to* f. *If* $\{x^M\}_{M \in \mathbb{N}} \subset X$ *is a sequence of global minimizers of* f_M, *and* $\hat{x}$ *is any accumulation point of this sequence (along a subsequence indexed by a set* $K \subset \mathbb{N}$), *then* $\hat{x}$ *is a global minimizer of* f *and* $\lim_{M \in K} \inf_{x \in X} f_M(x^M) = \inf_{x \in X} f(x)$.

3.2. Epiconvergence of J^M to J. In the previous section we introduced epiconvergence as a means to determine the convergence properties of an approximate optimization problem. In this section we analyze the convergence of our sample average scheme by demonstrating the epiconvergence of the approximate objective functional J^M to the original objective functional J. To do so, we show that J can be written as the expectation of a random lower semicontinuous function. To this end we introduce $T : \mathbf{H} \times \Omega \mapsto \mathbb{R}$ given by

$$T(\eta, \omega) \doteq F(x^{\eta}(1, \omega), \omega).$$

The following lemma establishes that T is a random lower semicontinuous function.

LEMMA 3.5. *Suppose Assumptions* 1–3 *hold, and let* L_F, L_f *be defined as in Assumptions* 2–3. *For each* $\omega \in \Omega$, *the function* $T(\cdot, \omega)$ *is Lipschitz continuous with Lipschitz constant* $L_T(\omega) = \sqrt{2} L_F(\omega) L_f(\omega) e^{L_f(\omega)}$. *Furthermore,* T *is* $\mathcal{B} \otimes \Sigma$ *measurable, where* $\mathcal{B}$ *is the Borel sigma-field generated by the open sets of* $\mathbf{H}$.

Proof. From Assumption 2 and Lemma 5.6.7 of [35], it is known that for each $\eta, \eta' \in \mathbf{H}$ and $\omega \in \Omega$,

$$\left\| x^{\eta}(1, \omega) - x^{\eta'}(1, \omega) \right\| \leq \sqrt{2} L_f(\omega) e^{L_f(\omega)} \left\| \eta - \eta' \right\|_{H_2}.$$

It follows from Assumption 3 that

$$|T(\eta, \omega) - T(\eta', \omega)| \leq \sqrt{2} L_F(\omega) L_f(\omega) e^{L_f(\omega)} \left\| \eta - \eta' \right\|_{H_2}.$$

Here $F : \mathbb{R}^n \times \Omega \mapsto \mathbb{R}$ is measurable by Assumption 3 and Proposition 2.2. For each $\eta \in \mathbf{H}$, $x^{\eta}(1, \cdot)$ is measurable by Lemma 2.3, so that $T(\eta, \cdot) = F(x^{\eta}(1, \cdot), \cdot)$ is measurable. T is therefore $\mathcal{B} \times \Sigma$ measurable by Proposition 2.2. $\square$

We can now write the objective functional J and approximate objective functional J^M in terms of the random lower semicontinuous function T:

$$J(\eta) = \mathbb{E}^P \big[T(\eta, \omega) \big], \qquad\qquad J^M(\eta) = \frac{1}{M} \sum_{i=1}^{M} T(\eta, \omega_i).$$

Before we establish the epiconvergence $J^M \to^{\text{epi}} J$ using Proposition 3.3, we recall the following fact.

LEMMA 3.6 (see [31, Lemma 4.2.2]). *The space $\mathbf{H}_C$ is a complete, separable metric space.*

We can now demonstrate the epiconvergence of the approximate objective functional using the following assumption.

ASSUMPTION 4. *Let $L_T : \Omega \mapsto [1, \infty)$ be defined as in Lemma 3.5. Then $L_T \in \mathcal{L}^1(\Omega)$.*

Note that this assumption is valid when Ω is a compact subset of $\mathbb{R}^d$, and the functions f and F are continuously differentiable with respect to ω.

THEOREM 3.7. *Suppose that Assumptions 1–4 hold. Then J^M epiconverges almost surely to J on $\mathbf{H}_C$ and $\mathbf{H}_O$ as $M \to \infty$.*

Proof. By Lemma 3.6, $\mathbf{H}_C$ is a separable complete metric space. By Lemma 3.5, T is $\mathcal{B} \otimes \Sigma$ measurable, and there exist scalars a and b such that $T(\eta, \omega) \geq a + b L_T(\omega)$ for all $\eta \in \mathbf{H}_C$. By Assumption 4 the function $a + b L_T(\omega)$ is integrable. The convergence $J^M|_{\mathbf{H}_C} \to^{\text{epi}} J|_{\mathbf{H}_C}$ almost surely then follows from Proposition 3.3. This convergence, together with the fact that $J^M(\eta) \to J(\eta)$ almost surely for all $\eta \in \mathbf{H}_O$, establishes the convergence $J^M|_{\mathbf{H}_O} \to^{\text{epi}} J|_{\mathbf{H}_O}$ almost surely. $\square$

Theorem 3.7 and Proposition 3.4 show that our sample average scheme has the property that accumulation points of a sequence of global minimizers of the approximate problem are global minimizers of the original problem.

4. Optimality conditions. Absent convexity, it is not generally possible to determine whether a numerically computed solution to an optimal control problem is a global minimizer. Necessary conditions, such as Pontryagin's minimum principle [22, 36], provide a method to assess the optimality of a numerically computed solution. Polak [35, Chapter 4] provides necessary conditions for the standard nonlinear optimal control problem in terms of optimality functions, which determine the stationary points of the objective functional. In this section we apply this approach to derive optimality functions for the nonstandard uncertain Problems B_C, B_O, B_C^M, and B_O^M which are based on the $\mathcal{L}_2$-Frechet derivative of the objective functional.

DEFINITION 4.1. *An upper semicontinuous function $\theta : X \mapsto \mathbb{R}$ is an optimality function for a problem B*
 (i) *if $\theta(x) \leq 0$ for all $x \in X$,*
 (ii) *if x is a local minimizer of B, then $\theta(x) = 0$.*

To establish the Frechet derivatives of the objective functionals, we first state the Frechet derivative of T.

PROPOSITION 4.2. *Suppose that Assumptions 1–3 are satisfied.*
 (i) *For any $\omega \in \Omega$, $\eta \in \mathbf{H}$, and $\delta\eta \in H_{\infty,2}$, $T(\cdot, \omega)$ has a Frechet derivative at η given by*

$$DT(\eta; \delta\eta; \omega) = \langle \nabla_\eta T(\eta, \omega), \delta\eta \rangle_{H_2}.$$

The gradient $\nabla_\eta T(\eta, \omega) = (\nabla_\xi T(\eta, \omega), \nabla_u T(\eta, \omega))^T \in H_{\infty,2}$ is given by

$$(4.1) \qquad \nabla_\xi T(\eta, \omega) = p^\eta(0, \omega),$$

$$(4.2) \qquad \nabla_u T(\eta, \omega)(s) = f_u^T(x^\eta(s, \omega), u^\eta(s), \omega) p^\eta(s, \omega),$$

and $p^\eta(s, \omega)$ is the solution to the adjoint equation

$$(4.3)$$
$$\dot{p}^\eta(s, \omega) = -f_x^T(x^\eta(s, \omega), u(s), \omega) p^\eta(s, \omega) \quad \text{for } s \in [0, 1], \quad p^\eta(1, \omega) = F_x(x^\eta(1, \omega), \omega).$$

(ii) *The gradient $\nabla_\eta T(\cdot, \omega)$ is Lipschitz continuous on $\mathbf{H}_C$.*

(iii) *For any $\eta \in \mathbf{H}$ and $\delta\eta \in H_{\infty,2}$, $T(\cdot, \omega)$ has a Frechet differential $DT(\eta; \delta\eta; \omega)$ at η.*

    Proof. The proposition follows directly from Corollary 5.6.9 of [35].   □

    The existence of the Frechet derivative in Proposition 4.2 allows us to introduce the Frechet derivatives of J and J^M by employing Fubini's theorem.

    LEMMA 4.3. *Suppose that Assumptions 1–4 are satisfied. Then for any $\eta \in \mathbf{H}$, $\delta\eta \in H_{\infty,2}$, the following hold:*

(i) *J has a Frechet differential $DJ(\eta; \delta\eta)$ at η given by $DJ(\eta; \delta\eta) = \langle \nabla J(\eta), \delta\eta \rangle_{H_2}$ with the gradient given by*

$$(4.4) \qquad \nabla J(\eta) = \mathbb{E}^P \left[\nabla_\eta T(\eta, \omega) \right].$$

(ii) *The gradient ∇J is Lipschitz continuous on $\mathbf{H}_C$.*

(iii) *J^M has a Gateaux differential $DJ^M(\eta; \delta\eta)$ at η given by $DJ^M(\eta; \delta\eta) = \langle \nabla J^M(\eta), \delta\eta \rangle_{H_2}$, with the gradient given by*

$$(4.5) \qquad \nabla J^M(\eta) = \frac{1}{M} \sum_{i=1}^M \nabla_\eta T(\eta, \omega_i).$$

(iv) *The gradient ∇J^M is Lipschitz continuous on $\mathbf{H}_C$.*

    Proof. We prove (i) and (ii); (iii) and (iv) follow by an identical argument with Ω replaced by $\{\omega_1, \ldots, \omega_M\}$ and P replaced by the counting measure normalized to 1.

    Proof of (i): Let $\delta\eta \in H_{\infty,2}$, $\eta \in \mathbf{H}$, $\omega \in \Omega$. Because $\mathbf{H}$ is open in the L_∞ topology there exists a $\lambda^* > 0$ such that $\eta + \lambda\delta\eta \in \mathbf{H}$ for all $\lambda \in [0, \lambda^*]$. From Lemma 3.5, $T(\cdot, \omega)$ is Lipschitz continuous in η with Lipschitz constant $L_T(\omega)$ for each $\omega \in \Omega$, and by Assumption 4 we have $L_T(\omega) \in \mathcal{L}^1(\Omega)$. From this fact we have

$$|T(\eta + \lambda\delta\eta, \omega) - T(\eta, \omega)| \le \left(L_T(\omega) \left\| \delta\eta \right\|_{H_2} \right) \lambda.$$

Therefore for each $\omega \in \Omega$, $\eta \in \mathbf{H}$, $\lambda \in [0, \lambda^*]$,

$$\left| \frac{T(\eta + \lambda\delta\eta, \omega) - T(\eta, \omega)}{\lambda} \right| \le L_T(\omega) \left\| \delta\eta \right\|_{H_2}.$$

Then the Gateaux derivative of J is given by

$$
\begin{aligned}
DJ(\eta; \delta\eta) &= \lim_{\lambda \downarrow 0} \frac{\mathbb{E}^P \left[T(\eta + \lambda\delta\eta, \omega) \right] - \mathbb{E}^P \left[T(\eta, \omega) \right]}{\lambda} \\
&= \lim_{\lambda \downarrow 0} \mathbb{E}^P \left[\frac{T(\eta + \lambda\delta\eta, \omega) - T(\eta, \omega)}{\lambda} \right] \\
&= \mathbb{E}^P \left[\lim_{\lambda \downarrow 0} \frac{T(\eta + \lambda\delta\eta, \omega) - T(\eta, \omega)}{\lambda} \right] \\
&= \mathbb{E}^P \left[DT(\eta, \delta\eta; \omega) \right],
\end{aligned}
$$

where we have used the dominated convergence theorem. Let $\delta\eta = (\xi^{\delta\eta}, u^{\delta\eta})$. Note that

$$\mathbb{E}^{\mathbb{P}} \left[\int_0^1 \left\langle \left[\nabla_u T(\eta, \omega)(t) \right], u^{\delta\eta}(t) \right\rangle dt \right] \le \mathbb{E}^{\mathbb{P}} \left[\left\| \nabla T(\eta, \omega) \right\|_{H_2} \left\| \delta\eta \right\|_{H_2} \right]$$

is bounded, so that we can write

$$
\begin{aligned}
DJ(\eta; \delta\eta) &= \mathbb{E}^P\left[\langle \nabla_\xi T(\eta, \omega), \xi^{\delta\eta}\rangle\right] + \mathbb{E}^P\left[\int_0^1 \langle \nabla_u T(\eta, \omega)(t), u^{\delta\eta}(t)\rangle\, dt\right] \\
&= \mathbb{E}^P\left[\langle \nabla_\xi T(\eta, \omega), \xi^{\delta\eta}\rangle\right] + \int_0^1 \mathbb{E}^P\left[\langle \nabla_u T(\eta, \omega)(t), u^{\delta\eta}(t)\rangle\right] dt \\
&= \langle \mathbb{E}^P\left[\nabla_\xi T(\eta, \omega)\right], \xi^{\delta\eta}\rangle + \int_0^1 \langle \mathbb{E}^P\left[\nabla_u T(\eta, \omega)(t)\right], u^{\delta\eta}(t)\rangle\, dt \\
&= \langle \mathbb{E}^P\left[\nabla_\eta T(\eta, \omega)\right], \delta\eta\rangle_{H_2},
\end{aligned}
$$

where we have used Fubini's theorem. To demonstrate that the Gateaux derivative DJ is the Frechet derivative of J, consider the quantity

$$
\begin{aligned}
&\lim_{\|\delta\eta\|_{H_2}\to 0} \frac{\|J(\eta + \delta\eta) - J(\eta) - DJ(\eta; \delta\eta)\|_{H_2}}{\|\delta\eta\|_{H_2}} \\
&= \lim_{\|\delta\eta\|_{H_2}\to 0} \frac{\left\|\mathbb{E}^P\left[T(\eta + \delta\eta, \omega) - T(\eta, \omega) - DT(\eta; \delta\eta; \omega)\right]\right\|_{H_2}}{\|\delta\eta\|_{H_2}} \\
&\leq \lim_{\|\delta\eta\|_{H_2}\to 0} \mathbb{E}^P\left[\frac{\|T(\eta + \delta\eta, \omega) - T(\eta, \omega) - DT(\eta; \delta\eta; \omega)\|_{H_2}}{\|\delta\eta\|_{H_2}}\right] \\
&= \mathbb{E}^P\left[\lim_{\|\delta\eta\|_{H_2}\to 0} \frac{\|T(\eta + \delta\eta, \omega) - T(\eta, \omega) - DT(\eta; \delta\eta; \omega)\|_{H_2}}{\|\delta\eta\|_{H_2}}\right] \\
&= 0,
\end{aligned}
$$

where we have used the dominated convergence theorem.

The proof of (ii) follows directly from the Lipschitz continuity of $\nabla_\eta T(\eta, \omega)$. $\square$

We now introduce nonpositive optimality functions for Problems B_C, B_C^M, B_O, and B_O^M, based on the Frechet derivatives defined in Lemma 4.3:

$$
\text{(4.6)} \qquad \theta_C(\eta) \doteq \min_{\eta' \in \mathbf{H}_C} DJ(\eta; \eta' - \eta) + \frac{1}{2}\|\eta' - \eta\|_{H_2}^2,
$$

$$
\text{(4.7)} \qquad \theta_C^M(\eta) \doteq \min_{\eta' \in \mathbf{H}_C} DJ^M(\eta; \eta' - \eta) + \frac{1}{2}\|\eta' - \eta\|_{H_2}^2,
$$

$$
\text{(4.8)} \qquad \theta_O(\eta) \doteq -\frac{1}{2}\|\nabla J(\eta)\|_{H_2}^2,
$$

$$
\text{(4.9)} \qquad \theta_O^M(\eta) \doteq -\frac{1}{2}\left\|\nabla J^M(\eta)\right\|_{H_2}^2.
$$

PROPOSITION 4.4. *Suppose that Assumptions* 1–4 *hold. Then the functions* θ_C, θ_C^M, θ_O, *and* θ_O^M *are continuous optimality functions for Problems* B_C, B_C^M, B_O, *and* B_O^M, *respectively.*

Proof. The proof for functions θ_C, θ_C^M follows directly from Lemma 4.3 and the arguments used in the proof of Theorem 4.2.3a in [35], with J or J^M replacing f^0, H_C replacing $\mathbf{H}$, and Lemma 4.3 replacing Corollary 5.6.9. The proof for functions θ_O, θ_O^M follows directly from Lemma 4.3 and the arguments used in the proof of Theorem 4.2.3c in [35], with J or J^M replacing f^0, H replacing H^O, and Lemma 4.3 replacing Corollary 5.6.9. $\square$

In general, the necessary condition based on the $\mathcal{L}_2$ variation of the objective functional will not be equivalent to Pontryagin's minimum principle except in the case

where the Hamiltonian is convex in u. However, it can be shown that for Problem B_O, under certain regularity conditions, the necessary condition $\theta_O(\eta) = \|\nabla J(\eta)\|_{H_2}^2 = 0$ is equivalent to the stationarity of the Hamiltonian given by

$$\mathcal{H}(x, \lambda, u, t) = \mathbb{E}^P \big[f(x(t, \omega), u(t), \omega)^T p(t, \omega) \big],$$

where p is the adjoint to the state variable x. To see this, suppose the initial condition is fixed, and note that the stationarity of the Hamiltonian implies that

$$\frac{\partial}{\partial u} \mathcal{H}(x, \lambda, u, t) = \mathbb{E}^P \big[f_u(x(t, \omega), u(t), \omega)^T p(t, \omega) \big] = 0$$

for almost all t when f is sufficiently smooth. Therefore $\|\nabla J(\eta)\|_{H_2}^2 = \int_0^1 \|\mathbb{E}^P [f_u(x(t), u(t), \omega)p(t, \omega)]\|^2 dt = 0$.

For Problem B_C, the approach of [32, 33] can be extended to produce a Pontryagin-like necessary condition for global minimizers that are accumulation points of global minimizers of the approximate Problem B_C^M. A direct extension of Pontryagin's minimum principle to the UOCP is desirable, as it may lead to insights into new optimization algorithms, but this approach is not pursued here. For work relating to this topic see [16, pp. 80–82].

5. Consistency of the approximation. In section 3 we analyzed the convergence of the approximation scheme for Problems B_C and B_O using the concept of epiconvergence. Epiconvergence of the objective functionals guarantees that accumulation points of a sequence of global minimizers of the approximate Problems B_C^M and B_O^M will be global minimizers of the original Problems B_C and B_O. However, epiconvergence is not sufficient to guarantee that accumulation points of a sequence of stationary points of the approximate problem are stationary points of the original problem. In this section, we demonstrate such a property, thus showing that the approximation scheme based on sample averages is consistent in the sense of Polak [35, section 3.3].

DEFINITION 5.1 (see [35]). *Let X be a complete separable metric space, and let $G^M : X \mapsto \mathbb{R}$, $G : X \mapsto \mathbb{R}$ be lower semicontinuous functions. Let Problem C^M be the problem of finding $x \in X$ to minimize G^M, and let Problem C be the problem of finding $x \in X$ to minimize G. Let the corresponding optimality functions for Problems C^M and C be given by $\Gamma^M : X \mapsto \mathbb{R}$ and $\Gamma : X \mapsto \mathbb{R}$. We say that the sequence $(C^M, \Gamma^M)_{M \in \mathbb{N}}$ is a consistent approximation to the pair (C, Γ) if the following hold:*
   (i) *G_M epiconverges to G.*
  (ii) *If $\{x_M\}_{M=1}^{\infty}$ is a sequence converging to x, then $\limsup_{M \to \infty} \Gamma_M(x_M) \leq \Gamma(x)$.*

We have already shown the almost sure epiconvergence of the approximate objective functional J^M to the objective functional J in Theorem 3.7. Recall that $L_f(\omega)$ is the Lipschitz constant for the dynamics f, and $L_F(\omega)$ is the Lipschitz constant for the objective functional F. To establish the convergence properties of the optimality function θ_C^M, we introduce the following assumption.

ASSUMPTION 5. *There exist constants $L_f', L_F' \in [1, \infty)$ such that $L_f(\omega) \leq L_f'$ and $L_F(\omega) \leq L_F'$ almost surely.*

Note that this assumption will be valid in the case that Ω is a compact subset of $\mathbb{R}^{n_\omega}$ and f, F are continuously differentiable. Therefore the assumption is satisfied for previously considered applications such as optimal search [32, 33] and ensemble control [39, 40, 41].

The following lemma addresses the measurability and continuity of the gradient of the objective functional.

LEMMA 5.2. *Suppose that Assumptions* 1–5 *hold. Let* $\eta \in \mathbf{H}$. *Then the following are true:*

(i) $\nabla_\eta T(\eta, \cdot)(\cdot) : \Omega \times [0,1] \to \mathbb{R}$ *is measurable.*

(ii) *There exists a compact set* $U_0 \subset \mathbb{R}^m$ *such that* $\nabla_u T(\eta, \omega)(t) \in U_0$ *for all* $\eta \in \mathbf{H}$, $\omega \in \Omega$, $t \in [0,1]$.

(iii) *There exists* $L'_{\nabla T} \in [1, \infty)$ *such that* $L_{\nabla T}(\omega) \leq L'_{\nabla T}$ *almost surely, where* $L_{\nabla T}(\omega)$ *is the Lipschitz constant of* $\nabla T(\cdot, \omega)$.

(iv) *For each* M, $L_{\nabla J^M} \leq L'_{\nabla T}$ *almost surely, where* $L_{\nabla J^M}$ *is the Lipschitz constant of* $\nabla_\eta J^M$.

Proof. Part (i) follows directly from (4.1)–(4.2) and the application of Lemma 2.3 to the adjoint system (4.3). Part (ii) follows from Lipschitz continuity of f_u (Assumption 2) and p (Lemma 2.3) and the boundedness of the set $\mathbf{H}$. Part (iii) follows from Assumption 5 and the proof of Lemma 5.6.9b of [35]. Part (iv) follows from (iii) and the fact that $\nabla J^M(\cdot, \omega) = \frac{1}{M} \nabla T(\cdot, \omega_i^M)$, where $\{\omega_i^M\}_{i=1}^M$ is an independent P-distributed random draw from Ω. □

To simplify notation, for a given $\eta^* \in \mathbf{H}_C$, we introduce the following functions:

(i) $\kappa_{\eta^*}^M : \mathbf{H}_C \mapsto \mathbb{R}; \eta \mapsto \left\langle \nabla J^M(\eta^*), \eta \right\rangle_{H_2}$,

(ii) $\kappa_{\eta^*} : \mathbf{H}_C \mapsto \mathbb{R}; \eta \mapsto \left\langle \nabla J(\eta^*), \eta \right\rangle_{H_2}$,

(iii) $\mu_{\eta^*}^M : \mathbf{H}_C \mapsto \mathbb{R}; \eta \mapsto \left\langle \nabla J^M(\eta), \eta^* \right\rangle_{H_2}$,

(iv) $\mu_{\eta^*} : \mathbf{H}_C \mapsto \mathbb{R}; \eta \mapsto \left\langle \nabla J(\eta), \eta^* \right\rangle_{H_2}$.

LEMMA 5.3. *Suppose that Assumptions* 1–4 *are satisfied. Then the following hold:*

(i) $\kappa_{\eta^*}^M \to \kappa_{\eta^*}$ *uniformly almost surely for each* η^* *in* $\mathbf{H}_C$.

(ii) $\mu_{\eta^*}^M \to^{\mathrm{epi}} \mu_{\eta^*}$ *almost surely for each* η^* *in* H_2.

Proof. Proof of (i): For a given $t \in [0,1]$, because the $\nabla_u T(\eta, \omega_i)(t)$ for $i = 1, \ldots, M$ are identically distributed, the strong law of large numbers, (4.4), and (4.5) imply that $\nabla J^M(\eta^*)(t) \to \nabla J(\eta^*)(t)$ almost surely. Therefore $\nabla J^M(\eta^*) \to \nabla J(\eta^*)$ pointwise almost surely as $M \to \infty$. Recall that $\|\eta\|_{H_2} \leq \rho_1 + \rho_2$ for all $\eta \in \mathbf{H}_C$. Therefore for each $\epsilon > 0$, there exists $K \in \mathbb{N}$ such that for each $M > K$, we have $\left\| \nabla J^M(\eta^*) - \nabla J(\eta^*) \right\|_{H_2} < \frac{\epsilon}{\rho_1 + \rho_2}$ by the dominated convergence theorem. Then

$$
\begin{aligned}
\left| \kappa_{\eta^*}^M(\eta) - \kappa_{\eta^*}(\eta) \right| &= \left| \langle \nabla J^M(\eta^*) - \nabla J(\eta^*), \eta \rangle_{H_2} \right| \leq \left\| \nabla J^M(\eta^*) - \nabla J(\eta^*) \right\|_{H_2} \|\eta\|_{H_2} \\
&< \frac{\epsilon}{\rho_1 + \rho_2} (\rho_1 + \rho_2) = \epsilon.
\end{aligned}
$$

Proof of (ii): First note that by Lemma 5.2, $\langle \nabla_\eta T(\eta, \omega), \eta^* \rangle_{H_2}$ is continuous in η and measurable in ω and therefore is a random lower semicontinuous function by Proposition 2.2 and Definition 3.2. Because $\mu_{\eta^*}(\cdot) = \mathbb{E}^P \langle \nabla_\eta T(\cdot, \omega), \eta^* \rangle_{H_2}$ by the proof of Lemma 4.3, μ_{η^*} is the expectation of a random lower semicontinuous function and is bounded by Lemma 5.2(iii). The result then follows from (4.4), (4.5), and Proposition 3.3. □

5.1. Consistency of the approximation of Problem B_C. Lemma 5.3 allows us to establish the almost sure consistent approximation of Problem B_C.

THEOREM 5.4. *Suppose that Assumptions* 1–5 *hold. Then the sequence* $(B_C^M, \theta_C^M)_{M \in \mathbb{N}}$ *is almost surely a consistent approximation to the pair* (B_C, θ_C) *on the decision space* $\mathbf{H}_C$.

Proof. The almost sure epiconvergence of J^M to J on $\mathbf{H}_C$ is established in Theorem 3.7.

It remains to show that $\limsup_{M \to \infty} \theta_C^M(\eta^M) \le \theta_C(\eta)$ whenever $\eta^M \to \eta$. Suppose that $\eta^M \in \mathbf{H}_C$ and $\eta^M \to \eta$. First we write

$$
\begin{aligned}
\theta_C^M(\eta^M) &= \min_{\eta' \in \mathbf{H}_C} \left\{ \left\langle \nabla J^M(\eta^M), \eta' - \eta^M \right\rangle_{H_2} + \frac{1}{2} \left\| \eta' - \eta^M \right\|_{H_2}^2 \right\} \\
&= \min_{\eta' \in \mathbf{H}_C} \left\{ \left\langle \nabla J^M(\eta^M), \eta' \right\rangle_{H_2} + \frac{1}{2} \left\| \eta' - \eta^M \right\|_{H_2}^2 \right\} - \left\langle \nabla J^M(\eta^M), \eta^M \right\rangle_{H_2} \\
&= \min_{\eta' \in \mathbf{H}_C} \left\{ \left\langle \nabla J^M(\eta^M) - \nabla J^M(\eta), \eta' \right\rangle_{H_2} + \left\langle \nabla J^M(\eta), \eta' \right\rangle_{H_2} + \frac{1}{2} \left\| \eta' - \eta^M \right\|_{H_2}^2 \right\} \\
&\quad - \left\langle \nabla J^M(\eta^M), \eta^M \right\rangle_{H_2} \\
&= \min_{\eta' \in \mathbf{H}_C} \left\{ \left\langle \nabla J^M(\eta^M) - \nabla J^M(\eta), \eta' \right\rangle_{H_2} + \kappa_\eta^M(\eta') + \frac{1}{2} \left\| \eta' - \eta^M \right\|_{H_2}^2 \right\} \\
&\quad - \left\langle \nabla J^M(\eta^M), \eta^M - \eta \right\rangle_{H_2} - \mu_\eta^M(\eta^M).
\end{aligned}
$$
(5.1)

Similarly,

$$
\theta_C(\eta) = \min_{\eta' \in \mathbf{H}_C} \left[\kappa_\eta(\eta') + \frac{1}{2} \left\| \eta' - \eta \right\|_{H_2}^2 \right] - \mu_\eta(\eta).
$$
(5.2)

We examine the behavior of $\limsup_{M \to \infty} \theta_C^M(\eta^M)$ by looking at each expression in (5.1).

Note that $\mathbf{H}_C$ is bounded; therefore we have by Lemma 5.2(iv)

$$
\begin{aligned}
\left\langle \nabla J^M(\eta^M) - \nabla J^M(\eta), \eta' \right\rangle_{H_2} &\le \left\| \nabla J^M(\eta^M) - \nabla J^M(\eta) \right\|_{H_2} \left\| \eta' \right\|_{H_2} \\
&\le L_{\nabla J^M} \left\| \eta^M - \eta \right\|_{H_2} \left\| \eta' \right\|_{H_2} \to 0
\end{aligned}
$$

uniformly in η' on $\mathbf{H}_C$. Similarly, because

$$
\left\| \eta' - \eta^M \right\|_{H_2}^2 - \left\| \eta' - \eta \right\|_{H_2}^2 = \left\| \eta^M \right\|_{H_2}^2 - \left\| \eta \right\|_{H_2}^2 + 2 \langle \eta - \eta^M, \eta' \rangle_{H_2} \to 0
$$

uniformly in η' on $\mathbf{H}_C$, we have $\left\| \eta' - \eta^M \right\|_{H_2}^2 \to \left\| \eta' - \eta \right\|_{H_2}^2$ uniformly in η'. This, combined with the uniform convergence $\kappa_\eta^M \to \kappa_\eta$, shows that

$$
\begin{aligned}
&\min_{\eta' \in \mathbf{H}_C} \left\langle \nabla J^M(\eta^M) - \nabla J^M(\eta), \eta' \right\rangle_{H_2} + \kappa_\eta^M(\eta') + \frac{1}{2} \left\| \eta' - \eta^M \right\|_{H_2}^2 \\
&\quad \to \min_{\eta' \in \mathbf{H}_C} \kappa_\eta(\eta') + \frac{1}{2} \left\| \eta' - \eta \right\|_{H_2}^2.
\end{aligned}
$$
(5.3)

Because $\langle \nabla J^M(\eta^M), \eta^M - \eta \rangle_{H_2} \to 0$ almost surely and μ_η^M epiconverges to μ_η, we have from (5.1)–(5.3)

$$
\lim_{M \to \infty} \sup \theta_C^M(\eta^M) \le \theta_C(\eta) \text{ almost surely.} \qquad \square
$$

5.2. Consistency of the approximation of Problem B_O. We now demonstrate the almost sure consistent approximation of Problem B_O.

THEOREM 5.5. *Suppose that Assumptions 1–4 hold. Then the sequence $(B_O^M, \theta_O^M)_{M \in \mathbb{N}}$ is almost surely a consistent approximation to the pair (B_O, θ_O) on the decision space $\mathbf{H}_O$.*

Proof. The almost sure epiconvergence of J^M to J on $\mathbf{H}_O$ was established in Theorem 3.7. It remains to establish the convergence properties of the optimality

functions. Suppose that $\eta^M \in \mathbf{H}_O$ and $\eta^M \to \eta \in \mathbf{H}_O$. Recall that H_2 is a complete Hilbert space. By Lemma 5.3(ii) and the Riesz representation theorem, for each $f \in H_2^*$ we have $\liminf_{M \to \infty} f(\nabla J^M(\eta^M)) \geq f(\nabla J(\eta))$ almost surely. By the Hahn–Banach theorem there exists $f^* \in H_2^*$ such that $\|f^*\|_{H_2^*} = 1$ and $f^*(\nabla J(\eta)) = \|\nabla J(\eta)\|_{H_2}$. Furthermore, for each M, we have

$$f^*(\nabla J^M(\eta^M)) \leq \|f^*\|_{H_2^*} \left\| \nabla J^M(\eta^M) \right\|_{H_2} = \left\| \nabla J^M(\eta^M) \right\|_{H_2}.$$

Therefore

$$\|\nabla J(\eta)\|_{H_2} = f^*(\nabla J(\eta)) \leq \lim \inf_{M \to \infty} f^*(\nabla J^M(\eta^M)) \leq \lim \inf_{M \to \infty} \left\| \nabla J^M(\eta^M) \right\|_{H_2}.$$

Consequently $\limsup_{M \to \infty} \theta_O^M(\eta^M) \leq \theta_O(\eta)$ almost surely.   □

6. The time-discretized problem. In sections 3–5 we analyzed a computational framework for the UOCP based on sample average approximations. This process creates a sequence of approximating standard optimal control problems which can be solved using existing techniques. In this section, we address the convergence properties of the method which uses the Euler discretization to solve the approximate Problem B_C^M. This approach can be generalized to other direct discretization algorithms such as Runge–Kutta [25, 42], as well as to the unconstrained Problem B_O^M.

First we introduce the framework with which we will perform our discrete approximation, taken from [35, Chapter 5]. This will involve an approximation of the admissible set as well as an approximation of the objective functional. For $k \in \{0, 1, \ldots, N-1\}$, let

$$(6.1) \qquad \pi_{N,k}(t) = \begin{cases} \sqrt{N} & \text{for all } t \in [k/N, (k+1)/N) \text{ if } k \leq N-1, \\ 0 & \text{otherwise.} \end{cases}$$

For any integer $N \geq 1$, we define the subspace $\mathcal{L}_N \subset \mathcal{L}_{\infty,2}^m[0,1]$ by

$$\mathcal{L}_N = \left\{ u \in \mathcal{L}_{\infty,2}^m[0,1] \middle| u(t) = \sum_{k=0}^{N-1} u_k \pi_{N,k}(t) \right\},$$

and

$$H_N = \mathbb{R}^n \times \mathcal{L}_N \subset H_{\infty,2}.$$

We then define the admissible set for the approximate problem as

$$\mathbf{H}_{C,N} = \mathbf{H}_C \cap H_N.$$

$\mathbf{H}_{C,N}$ is the set of all admissible initial state and control pairs for Problem B_C, with the additional requirement that the control be constant on each interval $[\frac{i}{N}, \frac{i+1}{N})$ for $i \in \{0, \ldots, N-1\}$.

For each $\omega \in \Omega$ and $\eta \in H_N$, we approximate the dynamics (1.2) using the Euler integration formula:

$$x_N^\eta \left(\frac{k+1}{N}, \omega \right) - x_N^\eta \left(\frac{k}{N}, \omega \right) = \frac{1}{N} f \left(x_N^\eta \left(\frac{k}{N}, \omega \right), u^\eta \left(\frac{k}{N} \right) \right), \quad k \in \{0, \ldots, N-1\},$$

$$x_N^\eta(0, \omega) = \xi^\eta + \iota(\omega).$$

For a detailed derivation of this approximation scheme and its relation to the NLP problem, see Polak [35, Chapter 5].

Recall that the objective functional to the UOCP is given by $J(\eta) = \mathbb{E}^P[T(\eta, \omega)]$, where $T(\eta, \omega) = F(x^\eta(1, \omega), \omega)$. Let $T^N : \mathbf{H} \times \Omega \mapsto \mathbb{R}$ be the time-discretized approximation to T, i.e., $T^N(\eta, \omega) = F(x_N^\eta(1, \omega), \omega)$. Given a random P-distributed draw $\{\omega_1, \ldots, \omega_M\}$ from Ω, we can define the sample average and time-discretized approximation to the objective functional J by

$$(6.2) \qquad J^{MN} = \frac{1}{M} \sum_{i=1}^{M} T^N(\eta, \omega_i).$$

Combining this objective functional with the discretized dynamics

$$x_N^\eta\left(\frac{k+1}{N}, \omega_i\right) - x_N^\eta\left(\frac{k}{N}, \omega_i\right) = \frac{1}{N} f\left(x_N^\eta\left(\frac{k}{N}, \omega_i\right), u^\eta\left(\frac{k}{N}\right)\right),$$
$$k \in \{0, \ldots, N-1\}, \quad i \in \{1, \ldots, M\},$$
$$(6.3) \qquad x_N^\eta(0, \omega_i) = \xi^\eta + \iota(\omega_i), \quad i \in \{1, \ldots, M\},$$

we can define the fully discretized problem.

Problem B_C^{MN}: Find an initial state and control pair $\eta = (u^\eta, \xi^\eta) \in \mathbf{H}_{C,N}$ to minimize the objective functional (6.2) subject to the constraints (6.3).

In order to approximate Problem B_C by Problem B_C^{MN}, our desire is to assign to each sample size $M \in \mathbb{N}$ a number $\mathcal{N}(M)$ of time discretization nodes in such a way that $J^{M\mathcal{N}(M)} \to^{\text{epi}} J$. To this end we introduce the following assumption.

ASSUMPTION 6. *For the function* $\mathcal{N} : \mathbb{N} \mapsto \mathbb{N}$, *we have* $\mathcal{N}(M) \to \infty$ *as* $M \to \infty$.

In section 3 we showed that $J^M \to^{\text{epi}} J$ as $M \to \infty$. It is well known that $J^{MN} \to^{\text{epi}} J^M$ as $N \to \infty$ (see [35, Chapter 4]). However, these conditions are not sufficient to guarantee that $J^{M\mathcal{N}(M)} \to^{\text{epi}} J$ as $M \to \infty$ for arbitrary assignments $\mathcal{N} : \mathbb{N} \to \mathbb{N}$. We demonstrate such a property by analyzing the error introduced by the time discretization approximation. Our approach is based on the fact that the effect of such a time discretization on a standard optimal control problem (which we can consider as a special case in which the value of the parameter ω is fixed) is known and is determined by L_f, L_F. That is, we can use existing results to uniformly bound (in both η and ω) the error introduced to the UOCP by approximating $T(\eta, \omega)$ by $T^N(\eta, \omega)$.

PROPOSITION 6.1. *Suppose that Assumptions* 1–6 *are satisfied. Then there exists a* K_T *such that*

$$(6.4) \qquad |T^N(\eta, \omega) - T(\eta, \omega)| \leq K_T/N$$

for every $\eta = (\xi, u) \in \mathbf{H}_C$ *almost surely.*

Proof. The proof follows from Assumption 5, the boundedness of the set $\mathbf{H}_C$, and the proofs of Theorems 5.6.23 and 5.6.24 in [35]. $\quad \square$

The fact that this convergence is uniform in both η and ω allows us to address the convergence $J^{MN} \to^{\text{epi}} J$.

THEOREM 6.2. *Suppose Assumptions* 1–6 *are satisfied. Then* $J^{M\mathcal{N}(M)} \to^{\text{epi}} J$ *almost surely.*

Proof. In order to establish epiconvergence we must show that
  (i) $\liminf J^{M\mathcal{N}(M)}(\eta_M) \geq J(\eta)$ whenever $\eta_M \to \eta$,
  (ii) $\lim J^{M\mathcal{N}(M)}(\eta_M) = J(\eta)$ for at least one sequence $\eta_M \to \eta$.

To do so, note that Assumption 5 implies the existence of a constant $L_T \in [1, \infty)$ such that $|T(\eta, \omega) - T(\eta', \omega)| \leq L_T \|\eta - \eta'\|$ for all $\eta, \eta' \in \mathbf{H}_C$ almost surely. Then consider the difference

$$
\begin{aligned}
\left| J(\eta) - J^{M\mathcal{N}(M)}(\eta_M) \right| &\leq \left| J(\eta) - J^M(\eta_M) \right| + \left| J^M(\eta_M) - J^{M\mathcal{N}(M)}(\eta_M) \right| \\
&= \left| J(\eta) - J^M(\eta_M) \right| + \left| \frac{1}{M} \sum_{i=1}^{M} T(\eta, \omega_i) - T^{\mathcal{N}(M)}(\eta_M, \omega_i) \right| \\
&\leq \left| J(\eta) - J^M(\eta_M) \right| + \left| \frac{1}{M} \sum_{i=1}^{M} T(\eta, \omega_i) - T(\eta_M, \omega_i) \right| \\
&\quad + \left| \frac{1}{M} \sum_{i=1}^{M} T(\eta_M, \omega_i) - T^{\mathcal{N}(M)}(\eta_M, \omega_i) \right| \\
&\leq \left| J(\eta) - J^M(\eta_M) \right| + L_T \|\eta - \eta_M\| + \frac{K_T}{\mathcal{N}}(M).
\end{aligned}
$$

The result then follows from Assumption 6 and the almost sure epiconvergence of J^M to J. $\square$

   This result establishes that the UOCP can be approximated by a sequence of high-dimensional NLP problems by using a sample average scheme to approximate the expectation over the parameter space and an Euler scheme to discretize the time domain. The resulting numerical solutions will be meaningful in the sense that an accumulation point of a sequence of global minimizers of the approximate problem will be a global minimizer of the original problem. In order to establish a similar result for stationary points, we must develop optimality conditions for the approximate problem and analyze the approximation of the adjoint variables. Such a result is beyond the scope of this work.

   7. Numerical examples. In sections 2–6 we proposed a computational framework for the UOCP and demonstrated that it can be approximated by a sequence of high-dimensional NLP problems under mild regularity assumptions. In this section we provide two examples which demonstrate this process. The first example deals with a simple harmonic oscillator and is included for the purpose of illustrating the main ideas in an accessible manner. Of course, the oscillator can be stabilized using a standard closed-loop control. However, the example highlights the ability to handle ensemble control problems where it might not be practical to construct a closed-loop controller due to sensor or actuator limitations. The second example involves 10 stochastic parameters distributed according to a complex joint probability distribution and illustrates the framework in a context where numerical integration by quadrature rules is computational costly and/or difficult to implement.

   The resulting problems are approximated numerically by taking a random sample of size M from the parameter space using a Monte Carlo method and approximating the objective functional using the sample average. The resulting standard optimal control problem is discretized using a LGL-pseudospectral method with N nodes in the time domain. This yields an MNn_u-dimensional NLP problem (here n_u is the dimension of the control input) which then is solved using the sequential quadratic programming package SNOPT [18]. Although a consistency result for the time- and space-discretized problem using an LGL-pseudospectral scheme in the time domain is beyond the scope of this work, such a technique is used in this section as it provides

faster convergence for the adjoint variables and can therefore better demonstrate the convergence of the optimality functions.

7.1. Harmonic oscillator. Consider the problem of designing a controller to stabilize a harmonic oscillator with natural frequency ω uniformly distributed on $[\delta_0, \delta_1]$. The oscillator in question is modelled by the uncertain dynamical system

$$(7.1) \qquad \left[\begin{array}{c} \dot{x}_1 \\ \dot{x}_2 \end{array}\right] = \left[\begin{array}{cc} 0 & -\omega \\ \omega & 0 \end{array}\right] \left[\begin{array}{c} x_1 \\ x_2 \end{array}\right] + \left[\begin{array}{c} u_1 \\ u_2 \end{array}\right], \qquad \left[\begin{array}{c} x_1(0) \\ x_2(0) \end{array}\right] = \left[\begin{array}{c} 1 \\ 0 \end{array}\right]$$

for all $\omega \in [\delta_0, \delta_1]$, $t \in [0, t_f]$. The difficulty in this problem lies in the fact that a control which stabilizes the oscillator for a specific value of the parameter, such as the worst case scenario, may cause dispersion in the other states. An example of this dispersion is illustrated in Figure 1, which shows a sample of end states $x(t_f, \omega), \omega \in [0, 2\pi]$ for a control designed using Pontryagin's minimum principle to calculate the optimal solution for the problem with a single parameter value.

We address this problem using an optimal control framework. We introduce a UOCP where the goal is to find a control which minimizes the cost functional

$$J(u) = \mathbb{E}^P \left[\beta \big[(x_1(t_f, \omega))^2 + (x_2(t_f, \omega))^2 \big] + \gamma \int_0^{t_f} \big[(u_1(t))^2 + (u_2(t))^2 \big] dt \right]$$

$$(7.2) \qquad = \beta \, \mathbb{E}^P \left[(x_1(t_f, \omega))^2 + (x_2(t_f, \omega))^2 \right] + \gamma \int_0^{t_f} \big[(u_1(t))^2 + (u_2(t))^2 \big] dt.$$

Here β and γ are scale factors which weight the priority of minimizing the error of the final state against that of minimizing the expended control energy. This objective functional can be used to design a control which achieves an end state in a desired neighborhood of zero.

We use the computational framework to numerically calculate an optimal control for this UOCP for two scenarios, with and without the presence of control constraints.

Problem S_O: Find a control $u : [0,1] \mapsto (-1000, 1000) \times (-1000, 1000)$ to minimize the objective functional (7.2) subject to the uncertain dynamical system (7.1), where $t_f = 1$, $\delta_0 = 0$, $\delta_1 = 2\pi$, $\beta = 1000$, $\gamma = 1$.

This problem approximates the unconstrained problem by allowing the admissible controls to take values in a large open subset of $\mathbb{R}^2$. Note that because the initial state is fixed, for notational convenience we will take the decision space to be $\mathbf{U}_O$, the set of all $u \in \mathcal{L}_2^2[0,1]$ such that $u(t) \in (-1000, 1000) \times (-1000, 1000)$ for almost every $t \in [0,1]$. We use the computational framework proposed in this paper to calculate a control which stabilizes the system in the face of this state dispersion. A sample computed trajectory for $M = 52$ with 54 time discretization nodes is shown in Figure 2, and the optimal control and sample of optimal final states are shown in Figure 3. When this result is compared to the method of using Pontryagin's minimum principle to stabilize a single parameter value (see Figure 1), it is clear that the UOCP method reduces the dispersion of the end state while keeping the control energy within reasonable bounds. Tuning the parameters β and γ in the objective functional can further reduce dispersion of the end state at the cost of increased control energy.

The antisymmetry of the state dynamics and the quadratic form of the cost functional give this problem an easily verifiable necessary condition. First we cast the problem in the form of section 2. We introduce the auxiliary state x_3 and define the state dynamics by The antisymmetry of the state dynamics and the quadratic form of the cost functional give this problem an easily verifiable necessary condition. First

(a)

(b)
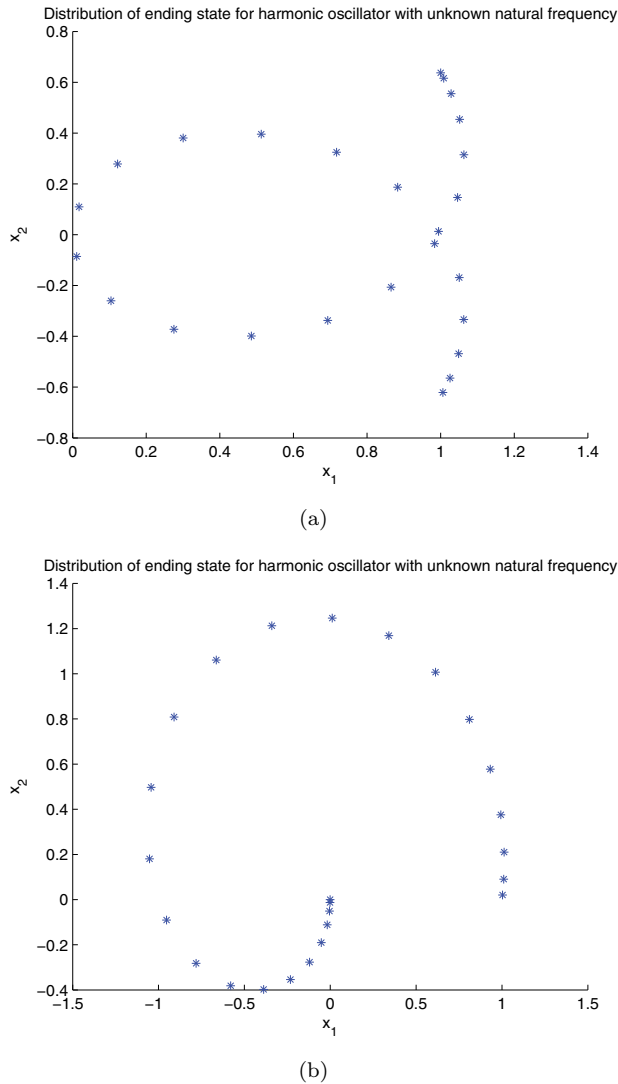
FIG. 1. *Dispersion of end states of the uncertain harmonic oscillator with* $\Omega = [0, 2\pi]$ *for a control designed using Pontyagin's minimum principle to stabilize* (a) *the median case* $(\omega = \pi)$ *and* (b) *the worst case* $(\omega = 2\pi)$.

we cast the problem in the form of section 2. We introduce the auxiliary state x_3 and define the state dynamics by

$$(7.3) \quad \begin{bmatrix} \dot{x}_1(t, \omega) \\ \dot{x}_2(t, \omega) \\ \dot{x}_3(t, \omega) \end{bmatrix} = \begin{bmatrix} 0 & -\omega & 0 \\ \omega & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1(t, \omega) \\ x_2(t, \omega) \\ x_3(t, \omega) \end{bmatrix} + \begin{bmatrix} u_1(t) \\ u_2(t) \\ \gamma(u_1(t))^2 + (u_2(t))^2 \end{bmatrix},$$
$$\begin{bmatrix} x_1(0, \omega) \\ x_2(0, \omega) \\ x_3(0, \omega) \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$
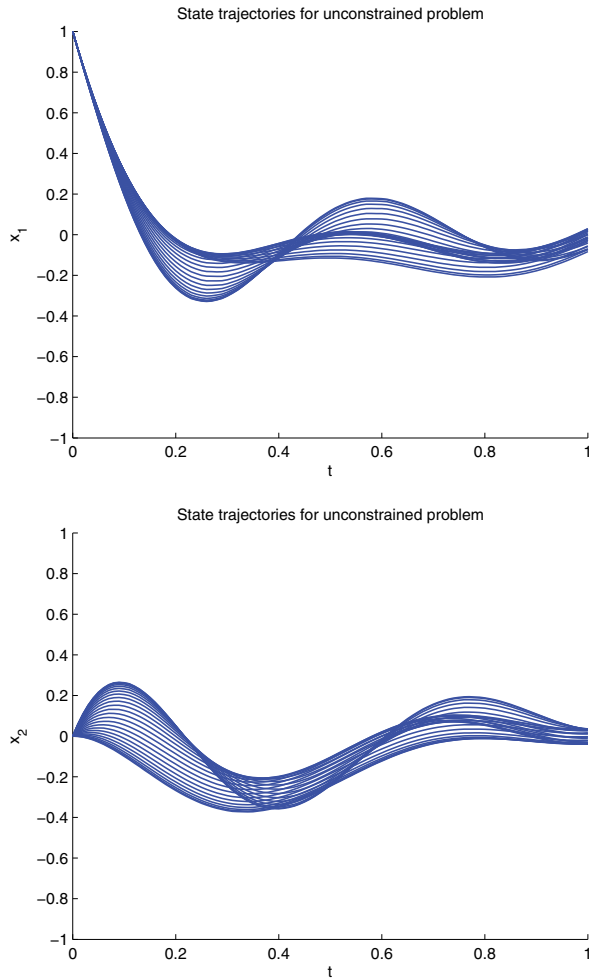
FIG. 2. *A sample of optimal state trajectories for the problem of stabilizing a harmonic oscillator with unknown natural frequency, with no control constraint. The optimal control is computed using the numerical framework proposed in this paper together with the NLP package SNOPT.*

The cost functional is then given by

$$(7.4) \qquad J(u) = \mathbb{E}^P[F(x(1,\omega),\omega)], \qquad\qquad F(x,\omega) = \beta x_1^2 + \beta x_2^2 + \gamma x_3.$$

Finally, the adjoint equation defined in (4.3) is given by

$$(7.5)$$
$$\begin{bmatrix} \dot{p}_1(t,\omega) \\ \dot{p}_2(t,\omega) \\ \dot{p}_3(t,\omega) \end{bmatrix} = \begin{bmatrix} 0 & -\omega & 0 \\ \omega & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} p_1(t,\omega) \\ p_2(t,\omega) \\ p_3(t,\omega) \end{bmatrix}, \qquad \begin{bmatrix} p_1(1,\omega) \\ p_2(1,\omega) \\ p_3(1,\omega) \end{bmatrix} = \begin{bmatrix} 2\beta x_1(1,\omega) \\ 2\beta x_2(1,\omega) \\ \gamma \end{bmatrix}.$$

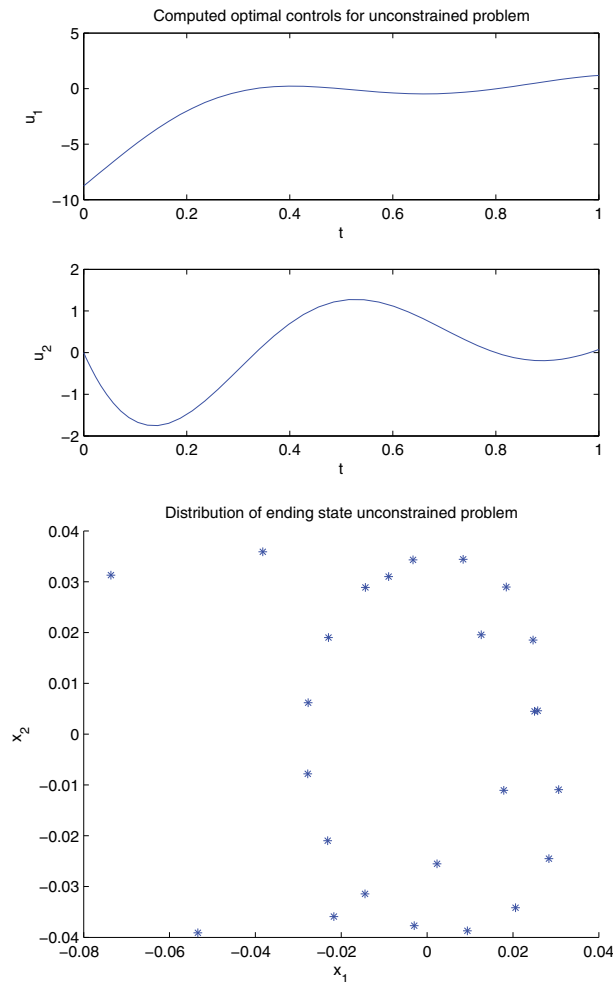The necessary condition defined in section 4 then requires that for an optimal solution

FIG. 3. *The optimal control and a sample of final values for the optimal state for the unconstrained problem.*

we have $-\frac{1}{2}\left\|\nabla J(u)\right\|_{H_2} = 0$, where ∇J is the Frechet derivative given by

$$
\begin{aligned}
\nabla J(u) &= \mathbb{E}^P[f_u^T(x(t,\omega), u(t), \omega)p(t,\omega)] \\
&= \mathbb{E}^P \left[\begin{array}{c} p_1(t,\omega) + 2\gamma u_1(t) \\ p_2(t,\omega) + 2\gamma u_2(t) \end{array} \right] \\
&= \left[\begin{array}{c} 2\gamma u_1(t) \\ 2\gamma u_2(t) \end{array} \right] + \mathbb{E}^P \left[\begin{array}{c} p_1(t,\omega) \\ p_2(t,\omega) \end{array} \right].
\end{aligned}
$$

(7.6)

The optimality function $-\frac{1}{2}\left\|\nabla J(u)\right\|_{H_2}$ provides a necessary condition which can be used for validation and verification of solutions. That is, if a sequence of numerically computed solutions to the approximate problem converge, but their associated optimality values diverge, the limit point cannot be an optimal solution to the original problem. For the harmonic oscillator UOCP considered in this section, we fix the number of time discretization nodes at 54 and increase the sample size M to analyze
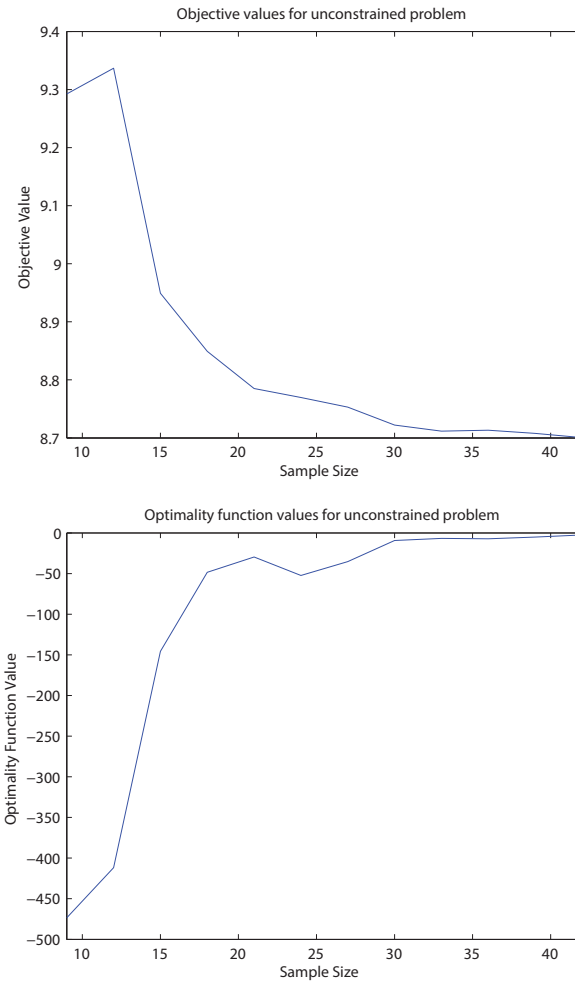
Fig. 4. *The objective functional and optimality function for the unconstrained problem as a function of the sample size M.*

the behavior of these values as the sample size increases. For a given sample size M, we determine the optimal control u_M^* for the approximate Problem B^M, then calculate the objective and optimality values $J(u_M^*)$ and $\theta_O(u_M^*)$. The optimality and objective values, as a function of the sample size M, are shown in Figure 4. We next consider a pointwise control-constrained UOCP for the harmonic oscillator.

Problem S_C: Find a control $u : [0,1] \mapsto [-3,3] \times [-3,3]$ to minimize the objective functional (7.2) subject to the uncertain dynamical system (7.1), where $t_f = 1$, $\delta_0 = 0$, $\delta_1 = 2\pi$, $\beta = 1000$, $\gamma = 1$.

Because the initial state is fixed, for notational convenience we will take the decision space to be $\mathbf{U}_C$, the set of all $u \in \mathcal{L}_2^2[0,1]$ such that $u(t) \in [-3,3] \times [-3,3]$ for almost every $t \in [0,1]$. As with the unconstrained problem, the optimal control is computed numerically using the framework proposed in this paper, with an LGL-pseudospectral discretization with 54 nodes in the time domain. A sample of computed state trajectories for $M = 42$ is shown in Figure 5. The corresponding
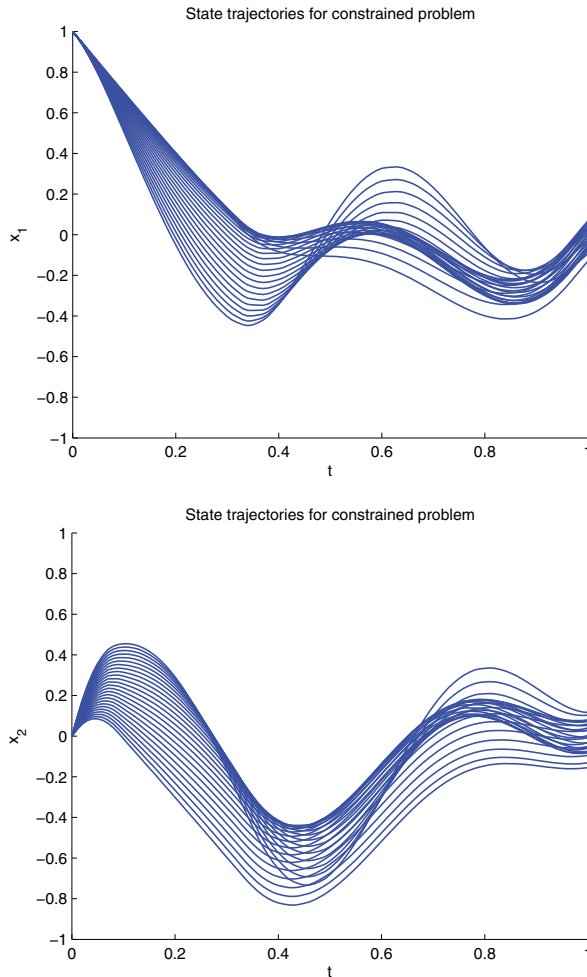
FIG. 5. *A sample of optimal state trajectories for the problem of stabilizing a harmonic oscillator with unknown natural frequency, in the presence of a pointwise control constraint.*

optimal controls, as well as a sample of ending states, are shown in Figure 6.

In section 4 it is shown that an optimal solution must satisfy the necessary condition $\theta_C(u) = 0$, where θ_C is given by (4.6). By substituting (7.6) we have

$$\theta_C(u) = \min_{u \in \mathbf{U}_C} \left(\langle \nabla J, u' - u \rangle_2 + \frac{1}{2} \|u' - u\|_2^2 \right)$$

$$= -\langle \nabla J, u \rangle_2 + \frac{1}{2} \|u\|_2^2 + \min_{u' \in \mathbf{U}_C} \left(\langle \nabla J - u, u' \rangle + \frac{1}{2} \|u'\|_2^2 \right).$$

The value of the objective functional $J(u_M^*)$ and optimality function $\theta_C(u_M^*)$ for a number of sample sizes M is shown in Figure 7. The state variables $x(t, \omega)$ and adjoint variables $p(t, \omega)$ are calculated by computing an optimal control for 54 LGL-pseudospectral nodes in the time domain and solving the resulting state-adjoint system. The value of θ_C is then approximated using the MATLAB quadratic programming package `quadprog`.
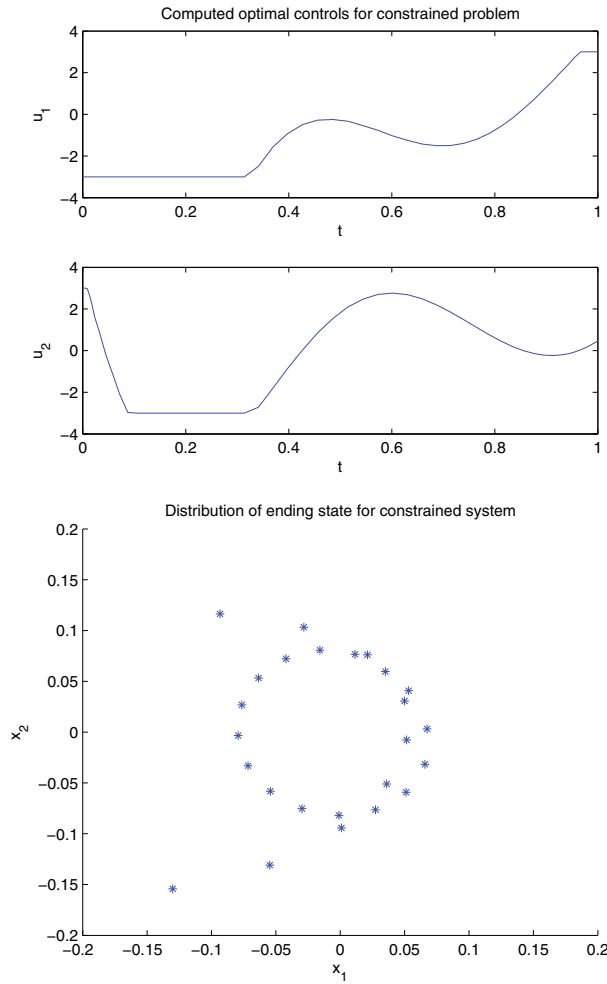
FIG. 6. *The optimal control and a sample of final values for the optimal state for the constrained problem.*

7.2. Intruder detection in a channel. We next consider an intruder detection problem inspired by [9]. A single searcher is attempting to detect a nonevading target moving down a channel. We assume the searcher has imperfect sensors and a turn-rate constraint. The objective is to find a trajectory for the searcher which maximizes the probability of detecting the target in the time horizon [0,75]. The searcher is assumed to be a Dubins vehicle with known constant velocity v. The dynamics of the searcher are given by

$$(7.7) \qquad \dot{x}_1(t) = v \cos x_3(t), \qquad\qquad \dot{x}_2(t) = v \sin x_3(t),$$
$$\dot{x}_3(t) = u(t), \qquad\qquad |u(t)| \leq K \quad \text{for all } t \in [0, 75],$$

where (x_1, x_2) represents the position of the searcher and x_3 is the heading angle. The control, u, is the turning rate of the vehicle. In the simulation, we set $v = 1$ and $K = 0.25$. We let the channel be given by the rectangle $R = [-20, 20] \times [-10, 10]$.
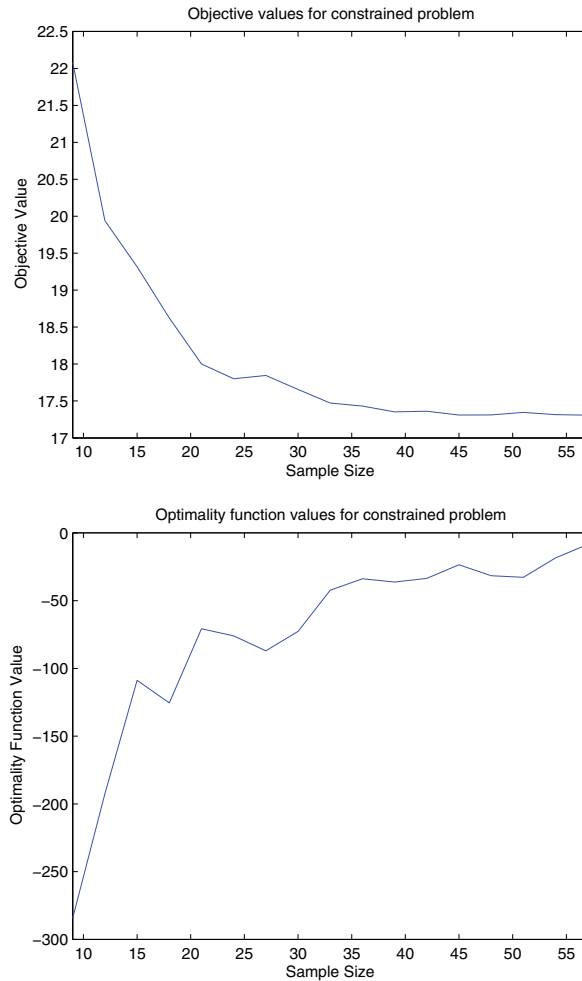
FIG. 7. *The value of the objective functional and optimality function for the constrained problem as a function of the sample size M.*

For each $\omega = (\omega_1, \omega_2, \ldots, \omega_{10}) \in \mathbb{R}^{10}$, we define the target trajectory $y(t, \omega) \in \mathbb{R}^2$ by

$$y_1(t) = \omega_1 + \omega_2 t + \frac{1}{2}\omega_3 t^2 + \frac{1}{6}\omega_4 t^3 + \frac{1}{24}\omega_5 t^4,$$

$$y_2(t) = \omega_6 + \omega_7 t + \frac{1}{2}\omega_8 t^2 + \frac{1}{6}\omega_9 t^3 + \frac{1}{24}\omega_{10} t^4.$$

Let $A \subset \mathbb{R}^{10}$ be the rectangle defined by $\omega_1 \in [0, 20]$, $\omega_5 \in [-10, 10]$, $\omega_2, \omega_7 \in [-\frac{1}{4}, \frac{1}{4}]$, $\omega_3, \omega_7 \in [-\frac{1}{40}, \frac{1}{40}]$, $\omega_4, \omega_9 \in [-\frac{1}{800}, \frac{1}{800}]$, $\omega_5, \omega_{10} \in [-\frac{1}{20000}, \frac{1}{20000}]$, and let

$$B \subset \mathbb{R}^{10} = \{\omega \in \mathbb{R}^{10} | y(t, \omega) \in R \text{ for all } t \in [0, 75],$$
$$\dot{y}_1(t, \omega) < 0 \text{ for all } t \in [0, 75]\}.$$

Note that B is the set of all parameter values for which the corresponding target trajectory is in the channel and is moving from right to left for all times $t \in [0, 75]$. We set $\Omega = A \cap B$.

The effectiveness of the search is given by

$$\tilde{r}(x(t), y(t,\omega)) = \beta\,\Phi\Big(\frac{F^k - D\,\|x(t) - y(t,\omega)\|^2 - b}{\sigma}\Big),$$

where Φ is the standard normal cumulative distribution function, β is the scan opportunity rate, F^k is the figure of merit, and σ reflects the variability in the signal excess. In the simulation we use the values $\beta = 1$, $F^k = 20$, $b = 20$, $D = 1$, and $\sigma = 10$. For more information about the formulation of this model, see [9, 14, 15, 30, 31]. The problem then becomes to minimize the functional

$$(7.8) \qquad \mathbb{E}^P\left[\exp\left(-\int_0^{75} \tilde{r}(x(t), y(t,\omega))dt\right)\right]$$

subject to the dynamics (7.7), where P is the uniform distribution on Ω. It is easily seen that this problem can be transformed into the form (1.1).

Due to the irregular shape of the parameter space Ω, this problem would be particularly challenging if we were to apply quadrature-based methods. However, the proposed framework is easily implemented with sampling carried out using the acceptance-rejection method. Using 54 nodes in the time domain and $M = 5000$, we obtain the searcher trajectory in Figure 8. We note that the figure shows only 10 of the 5000 target trajectories.
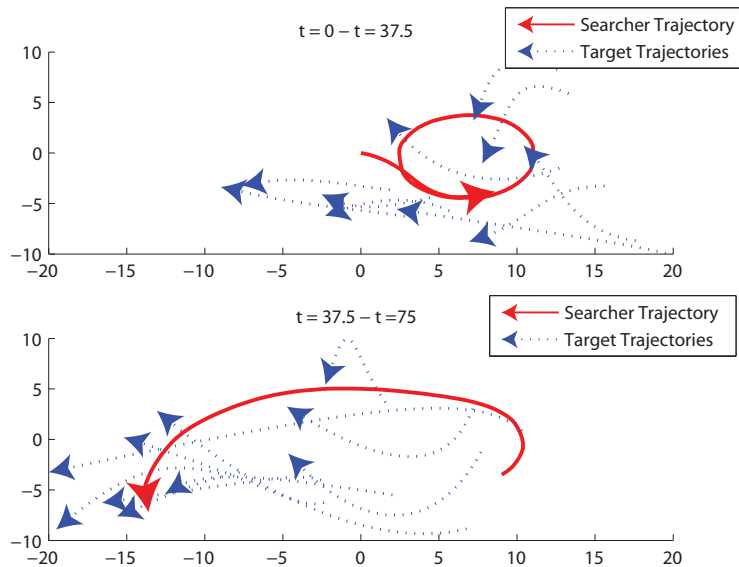

FIG. 8. *Computed trajectory for a searcher attempting to detect an intruder in the channel for $M = 5000$. For reference, 10 possible target trajectories are shown. The target moves right to left down the channel, and the searcher starts at $(0,0)$ at time $t = 0$. The arrows in the figure indicate the orientation of the trajectories.*

In this section we demonstrate that the numerical method proposed in this paper can be used to control a system with stochastic parameters either with or without pointwise control constraints. In addition, we assess the validity of the numerically

computed solutions using the necessary conditions developed in this work. The problem setting opens a wide variety of possible application areas for this method, including optimal search and ensemble control. This technique is based on a sample average scheme which scales favorably as the number of stochastic parameters increases. Therefore it can be applied to systems with a large number of stochastic parameters previously beyond the reach of techniques such as those in [32, 33, 39, 40, 41]. Even though the computational cost remains reasonable in these examples (60 minutes for the first example and 35 minutes for the second, using Intel Core i7–4700HQ 2.40 GHz with 16 GB RAM), the problem of optimal control under uncertainty remains challenging, and we anticipate much future algorithmic work to improve this further.

REFERENCES

[1] C. D. ALIPRANTIS AND K. C. BORDER, *Infinite Dimensional Analysis: A Hitchhiker's Guide*, 3rd ed., Springer, Berlin, 2006.

[2] Z. ARTSTEIN AND R. J.-B. WETS, *Consistency of minimizers and the SLLN for stochastic programs*, J. Convex Anal., 2 (1995), pp. 1–17.

[3] H. ATTOUCH AND R. J.-B. WETS, *A convergence theory for saddle functions*, Trans. Amer. Math. Soc., 280 (1981), pp. 1–41.

[4] H. ATTOUCH AND R. J.-B. WETS, *Epigraphical processes: Laws of large numbers for random LSC functions*, Sém. Anal. Convexe, 20 (1990), 13.

[5] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM J. Numer. Anal., 45 (2007), pp. 1005–1034.

[6] N. S. BEDROSSIAN, S. BHATT, W. KANG, AND I. M. ROSS, *Zero-propellant maneuver guidance: Rotating the international space station with computational dynamic optimization*, IEEE Control Syst. Mag., 29 (2009), pp. 53–73.

[7] N. BEDROSSIAN, M. KARPENKO, AND S. BHATT, *Overclock my satellite*, IEEE Spectr., 49 (2012), pp. 54–62.

[8] A. E. BRYSON AND Y. C. HO, *Applied Optimal Control*, Hemisphere, New York, 1975.

[9] H. CHUNG, E. POLAK, J. O. ROYSET, AND S. SASTRY, *On the optimal detection of an underwater intruder in a channel using unmanned underwater vehicles*, Naval Res. Logist., 58 (2011), pp. 804–820.

[10] J. CULLUM, *Finite-dimensional approximations of state-constrained continuous optimal control problems*, SIAM J. Control, 10 (1972), pp. 649–670.

[11] X. C. DING, A. RAHMANI, AND M. EGERSTEDT, *Optimal multi-UAV convoy protection*, in Proceedings of the Second International Conference on Robot Communication and Coordination (ROBOCOMM '09), Odense, Denmark, 2009.

[12] J. FISHER AND R. BHATTACHARYA, *Linear quadratic regulation of systems with stochastic parameter uncertainties*, Automatica J. IFAC, 45 (2009), pp. 2831–2841.

[13] J. FISHER AND R. BHATTACHARYA, *Optimal trajectory generation with probabilistic system uncertainty using polynomial chaos*, J. Dyn. Syst. Meas. Control, 133 (2010), 014501.

[14] J. FORAKER, J. O. ROYSET, AND I. KAMINER, *Search-trajectory optimization: Part I, Formulation and theory*, J. Optim. Theory Appl., to appear, doi:10.1007/s10957-015-0768-y.

[15] J. FORAKER, J. O. ROYSET, AND I. KAMINER, *Search-trajectory optimization: Part II, Algorithms and computations*, J. Optim. Theory Appl., to appear, doi:10.1007/s10957-015-0770-4.

[16] R. GABASOV AND F. M. KIRILLOVA, *Principi maksimuma v teorii optimal'novo upravleniya*, Izdat. Nauka i Tekhnika, Minsk, 1974.

[17] A. S. GADRE, S. DU, AND D. J. STILWELL, *A topological map based approach to long range operation of an unmanned surface vehicle*, in Proceedings of the American Control Conference (ACC), 2012, pp. 5401–5407.

[18] P. E. GILL, W. MURRAY, AND M. A. SAUNDERS, *SNOPT: An SQP algorithm for large-scale constrained optimization*, SIAM Rev., 47 (2005), pp. 99–131.

[19] Q. GONG, W. KANG, AND I. M. ROSS, *A pseudospectral method for the optimal control of constrained feedback linearizable systems*, IEEE Trans. Automat. Control, 51 (2006), pp. 1115–1129.

[20] Q. GONG, I. M. ROSS, W. KANG, AND F. FAHROO, *Connections between the covector mapping theorem and convergence of pseudospectral methods for optimal control*, Comput. Optim. Appl., 41 (2008), pp. 307–335.

[21] W. W. HAGER, *Runge-Kutta methods in optimal control and the transformed adjoint system*, Numer. Math., 87 (2000), pp. 247–282.

[22] R. F. HARTL, S. P. SETHI, AND R. G. VICKSON, *A survey of the maximum principles for optimal control problems with state constraints*, SIAM Rev., 37 (1995), pp. 181–218.

[23] F. S. HOVER, *Gradient dynamic optimization with Legendre chaos*, Automatica J. IFAC, 44 (2008), pp. 135–140.

[24] F. S. HOVER AND M. S. TRIANTAFYLLOU, *Application of polynomial chaos in stability and control*, Automatica J. IFAC, 42 (2006), pp. 789–795.

[25] S. KAMESWARAN AND L. T. BIEGLER, *Convergence rates for direct transcription of optimal control problems using collocation at Radau points*, Comput. Optim. Appl., 41 (2008), pp. 81–126.

[26] W. KANG, *Rate of convergence for the Legendre pseudospectral optimal control of feedback linearizable systems*, J. Control Theory Appl., 8 (2010), pp. 391–405.

[27] J.-S. LI, *Ensemble control of finite-dimensional time-varying linear systems*, IEEE Trans. Automat. Control, 56 (2011), pp. 345–357.

[28] J.-S. LI AND N. KHANEJA, *Ensemble control of linear systems*, in Proceedings of the 2007 46th IEEE Conference on Decision and Control, New Orleans, LA, 2007, pp. 3768–3773.

[29] J.-S. LI, J. RUTHS, AND D. STEFANATOS, *A pseudospectral method for optimal control of open quantum systems*, J. Chem. Phys., 131 (2009), 164110.

[30] M. MANGEL, *Search theory: A differential equations approach*, in Search Theory: Some Recent Developments, D. V. Chudnovsky and G. V. Chudnovsky, eds., Dekker, New York, 1989, pp. 55–101.

[31] C. PHELPS, *Computational Optimal Control of Nonlinear Systems with Parameter Uncertainty*, Ph.D. thesis, University of California, Santa Cruz, Santa Cruz, CA, 2015.

[32] C. PHELPS, Q. GONG, J. O. ROYSET, AND I. KAMINER, *Consistent approximation of an optimal search problem*, in Proceedings of the 2012 IEEE 51st Annual Conference on Decision and Control (CDC 2012), Maui, HI, 2012, pp. 630–637.

[33] C. PHELPS, Q. GONG, J. O. ROYSET, C. WALTON, AND I. KAMINER, *Consistent approximation of a nonlinear optimal control problem with uncertain parameters*, Automatica J. IFAC, 50 (2014), pp. 2987–2997.

[34] C. PHELPS, J. O. ROYSET, AND Q. GONG, *Sample average approximations in optimal control of uncertain systems*, in Proceedings of the 2013 IEEE 52nd Annual Conference on Decision and Control (CDC 2013), Florence, Italy, 2013, pp. 1958–1965.

[35] E. POLAK, *Optimization: Algorithms and Consistent Approximations*, Springer, New York, 1997.

[36] L. S. PONTRYAGIN, V. G. BOLTYANSKII, R. V. GAMKRELIDZE, AND E. F. MISHCHENKO, *The Mathematical Theory of Optimal Processes*, Interscience Publishers, John Wiley & Sons, New York, London, 1962.

[37] S. A. P. QUINTERO, F. PAPI, D. J. KLEIN, AND L. CHISCI, *Optimal UAV coordination for target tracking using dynamic programming*, in Proceedings of the 2010 49th IEEE Conference on Decision and Control, Atlanta, GA, 2010, pp. 4541–4546.

[38] I. M. ROSS AND M. KARPENKO, *A review of pseudospectral optimal control: From theory to flight*, Annu. Rev. Control, 36 (2012), pp. 182–197.

[39] J. RUTHS AND J.-S. LI, *Optimal ensemble control of open quantum systems with a pseudospectral method*, in Proceedings of the 2010 49th IEEE Conference on Decision and Control, Atlanta, GA, 2010, pp. 3008–3013.

[40] J. RUTHS AND J.-S. LI, *Optimal control of inhomogenous ensembles*, IEEE Trans. Automat. Control, 57 (2012), pp. 2021–2032.

[41] J. RUTHS, A. ZLOTNIK, AND J.-S. LI, *Convergence of a pseudospectral method for optimal control of complex dynamical systems*, in Proceedings of the 2011 50th IEEE Conference on Decision and Control and European Control Conference, Orlando, FL, 2011, pp. 5553–5558.

[42] A. SCHWARTZ AND E. POLAK, *Consistent approximations for optimal control problems based on Runge–Kutta integration*, SIAM J. Control Optim., 34 (1996), pp. 1235–1269.

[43] A. SHAPIRO, D. DENTCHEVA, AND A. RUSZCZYŃSKI, *Lectures on Stochastic Programming: Modeling and Theory*, 2nd ed., SIAM, Philadelphia, 2014.

[44] C. L. WALTON, Q. GONG, I. KAMINER, AND J. O. ROYSET, *Optimal motion planning for searching for uncertain targets*, in Proceedings of the 19th IFAC World Congress, Cape Town, South Africa, 2014, pp. 8977–8982.

[45] D. XIU AND G. E. KARNIADAKIS, *The Wiener–Askey polynomial chaos for stochastic differential equations*, SIAM J. Sci. Comput., 24 (2002), pp. 619–644.

[46] M. ZABARANKIN, S. URYASEV, AND P. PARDALOS, *Optimal risk path algorithms*, Appl. Optim., 66 (2002), pp. 273–296.