



HAL
open science

A Gradient Descent-Ascent Method for Continuous-Time Risk-Averse Optimal Control

Riccardo Bonalli, Velho Gabriel, Jean Auriol

► **To cite this version:**

Riccardo Bonalli, Velho Gabriel, Jean Auriol. A Gradient Descent-Ascent Method for Continuous-Time Risk-Averse Optimal Control. 2024. hal-04381732

HAL Id: hal-04381732

<https://hal.science/hal-04381732>

Preprint submitted on 9 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Gradient Descent-Ascent Method for Continuous-Time Risk-Averse Optimal Control

Gabriel Velho^a, Jean Auriol^a, Riccardo Bonalli^a

^a*Université Paris-Saclay, CentraleSupélec, CNRS, Laboratoire des signaux et systèmes, Gif-sur-Yvette, France*

Abstract

In this paper, we consider continuous-time stochastic optimal control problems where the cost is evaluated through a coherent risk measure. We provide an explicit gradient descent-ascent algorithm which applies to problems subject to non-linear stochastic differential equations. More specifically, we leverage duality properties of coherent risk measures to relax the problem via a smooth min-max reformulation which induces artificial strong concavity in the max subproblem. We then formulate necessary conditions of optimality for this relaxed problem which we leverage to prove convergence of the gradient descent-ascent algorithm to candidate solutions of the original problem. Finally, we showcase the efficiency of our algorithm through numerical simulations involving trajectory tracking problems and highlight the benefit of favoring risk measures over classical expectation.

Keywords: Gradient descent-ascent, Coherent risk measures, Stochastic optimal control, Non-linear control, Risk-averse control

1. Introduction

Risk-averse stochastic optimal control is a powerful tool for designing control laws that feature robustness against uncertainties. Relevant applications of this theory range from safe financial investment to safe control of autonomous systems, as evidenced by the monographs [1, 2] and references therein. For example, when controlling non-linear stochastic systems in robotics or space applications, robustness against uncertainties is often leveraged to ensure safety during motion. The problem also arises in finance, where, to mitigate unrecoverable losses, an agent may want to invest safe assets uniquely, even in exchange for a smaller average gain. This is in contrast

with “classical” stochastic optimal control, which rather aims at minimizing the expectation of a stochastic cost. Although high reward on average might be attained, these *risk-neutral* methods result in optimal controls that may generate trajectories that may cause catastrophic losses with small, though non-zero, probability. As such, these solutions are considered unacceptable in terms of safety in many applications [3].

To safely control systems whose dynamics are subject to uncertainties, various methods have been developed. A first possible approach is robust control, where optimal control problems are directly solved in the worst-case scenarios [4]. Although robust control theoretically guarantees resilience to unfavorable situations, it can be over-conservative and does not apply to dynamics with unbounded uncertainties, e.g., stochastic differential equations [3]. A more flexible approach is risk-averse control, which aims at minimizing a suitable risk measure of the stochastic cost, offering a good trade-off between minimizing the expected cost and mitigating catastrophic scenarios. Many risk-averse control methods involve discretizing the dynamics before proceeding with optimization [2]. For discrete problems, risk-averse Model Predictive Control (MPC) is a popular choice [5, 6]. It operates similarly to standard MPC but optimizes a risk measure instead of the expectation of the cost. Continuous-time risk-averse control methods are less common. Although risk-averse analogs of the Hamilton-Jacobi-Bellman (HJB) equation [7] and the Pontryagin Maximum Principle (PMP) [8] have been established, these methods do not directly provide tractable numerical schemes.

Designing a risk-averse method that scales with many different risk measures is a challenging task. Therefore, the focus has mostly been on developing methods that minimize specific risk measures. A significant hindrance of risk measures is their lack of *time consistency*, meaning that the optimal solution at the current time may not solely depend on the current state’s values [9]. A class of time-consistent risk measures exists for which dynamic programming can be employed to solve the associated Markov Decision Process (MDP) [10, 11]. Nevertheless, such methods do not apply to the widely adopted Conditional Value at Risk (CV@R), a risk measure that is not time consistent [1]. To address this issue, a state-space lifting approach has been proposed in [12]. Although effective, these approaches are limited to discrete-time and finite discrete space MDP. A continuous-time and continuous space approach for the CV@R has been proposed in [13], where a representation of the CV@R is used to re-write the problem into a double minimization that is solved with a gradient descent approach. Another important class of risk

measures is the family of risk measures that can be expressed as a linear function of the state’s probability distribution [1, Chapter 6]. To minimize these risk measures, a mean field game approach can be applied. In [14] and [15], the Frank-Wolf algorithm is used to solve the problem, with proof of convergence properties to a necessary optimality condition. Yet, these approaches work with very specific risk measures.

In this paper, we propose a gradient-based method to solve continuous-time, non-convex risk-averse control problems where the corresponding risk measure is only coherent, a fairly general requirement encompassing all the aforementioned risk measures. We will only consider the case of a diffusion independent of the control to keep the computation of gradients tractable. This is a reasonable assumption for applications such as aerospace and robotics, where the diffusion process often models external perturbations (e.g., gusts of wind on a robot), which are then uncorrelated from the control. The motivation to pursue this research direction came from the fact that gradient descent methods have already seen a lot of success in the more classical setting of average-oriented minimization [16], non-stochastic control [17], as well as some risk-averse methods [13, 15, 18]. Minimizing directly the risk measure is challenging. We, therefore, leverage the duality properties of coherent risk measures to transform the minimization problem into a min-max problem of a smooth, non-convex/linear function. Min-max optimization has been widely studied in the convex-concave setting [19]. Still, recent applications to machine learning (particularly Generative Adversarial Networks) have pushed studies toward the non-convex / strongly concave setting [20]. The gradient descent-ascent method has proven to be a successful technique for addressing these problems in finite-dimensional spaces. In [21], the proposed algorithm reaches an ϵ distance from a critical point within a finite time frame. In [22], the algorithm’s convergence to the critical point is guaranteed, given a specific geometric property of the cost function.

Our contribution consists of applying gradient descent ascent to solve the min-max problem, which stems from dual representations of coherent risk measures, by extending the work in [22]. This is not an easy task in that our problem is not smooth nor strongly concave. To address such an issue, we propose to apply gradient descent-ascent to a modified version of the problem which is smooth, non-convex / strongly concave, and arbitrarily “close” to the original formulation. Via appropriate proofs of convergence, we then show that the gradient of the cost converges to 0. By requiring more regularity in the dynamics of the system, we can also prove that the sequence

of controls generated by our algorithm converges (in an appropriate sense) to the set of critical points of the original problem. We finally implement the algorithm on a realistic robotic application to showcase the efficiency of our easy-to-use method.

The paper is organized as follows. In Section 2, we begin by providing a brief summary of some basic concepts of stochastic calculus and risk measures. Section 3 outlines the specific settings of dynamics and control strategies under consideration. Next, we present in Section 4 our algorithm, providing details for the computation of the gradient. In Section 5, we prove various convergence properties under different assumptions on the dynamics. Finally, we analyze in Section 6 numerical convergence, and we showcase the effectiveness of our approach through a series of simulations.

2. Notations and Preliminary Results

In this section, we briefly present some basic concepts and results of stochastic calculus we need in our work. More details can be found in [23] and [24, Chapter 1.6]. From now on, we fix $n, m, d \in \mathbb{N}$, a finite time horizon $T > 0$, and $p \in [1, +\infty)$. Let E be a vector space. We denote E^* its dual. As a result of Riesz theorem, if E is also a Hilbert space, then it is isomorphic to its dual space E^* . Consequently, throughout the paper, elements in $(L^2)^*$ are implicitly assumed as elements in L^2 . We denote $\|\cdot\|$ as the Euclidian norm. Other norms will be properly introduced as they are needed. We consider random variables in a probability space $(\Omega, \mathcal{G}, \mathbb{P})$. For any sub-sigma algebra $\mathcal{S} \subset \mathcal{G}$, we denote by $L_{\mathcal{S}}^p(\Omega, \mathbb{R}^n)$ the Banach space of random variables $z : \Omega \rightarrow \mathbb{R}^n$ that are \mathcal{S} -measurable and satisfy:

$$\|z\|_{L^p} \triangleq \mathbb{E} [\|z\|^p]^{1/p} < \infty.$$

It is a standard consequence of Riesz theorem that $L_{\mathcal{S}}^p(\Omega, \mathbb{R}^n)^* \cong L_{\mathcal{S}}^q(\Omega, \mathbb{R}^n)$, with q the conjugate exponent of p . Let $(W_t)_{t \in [0, T]}$ be a d -dimensional Wiener process, generating a filtration

$$\mathcal{F} \triangleq (\mathcal{F}_t)_{t \in [0, T]} = (\sigma(W_s, 0 \leq s \leq t))_{t \in [0, T]}.$$

We denote by $L_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^n)$ the Banach space of processes $x : [0, T] \times \Omega \rightarrow \mathbb{R}^n$ which are progressive with respect to \mathcal{F} and such that:

$$\|x\|_{L_{\mathcal{F}}^p} \triangleq \mathbb{E} \left[\int_0^T \|x\|^p(t) dt \right]^{1/p} < \infty.$$

In addition, we denote by $C_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^n)$ the Banach space of \mathcal{F} -adapted processes x which have continuous sample paths and finite sup norm

$$\|x\|_{C_{\mathcal{F}}^p} \triangleq \mathbb{E} \left[\sup_{0 \leq t \leq T} \|x\|^p(t) \right]^{1/p} < \infty.$$

In particular, $C_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^n) \subset L_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^n)$. In what follows, we often use the notation $x(t)$, $t \in [0, T]$ to denote progressive processes. A \mathcal{F} -adapted process $x : [0, T] \times \Omega \rightarrow \mathbb{R}^n$ such that $x(t) \in L_{\mathcal{F}_t}^1$ for every $t \in [0, T]$ is called a martingale provided that

$$\mathbb{E}[x(t)|\mathcal{F}_s] = x(s),$$

for all $0 \leq s \leq t \leq T$. We then say that a martingale $x : [0, T] \times \Omega \rightarrow \mathbb{R}^n$ is uniformly bounded in $L_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^n)$ if there exists a constant $C > 0$ such that

$$\sup_{t \in [0, T]} \|x(t)\|_{L^p} \leq C.$$

For every $x = (x^1 | \dots | x^d) \in L_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^{n \times d})$, we write

$$y(t) = \int_0^T x(t) dW_t \triangleq \sum_{i=1}^d \int_0^T x^i(t) dW_t^i,$$

for the Itô integral of x with respect to W , and recall that y is then a martingale in $C_{\mathcal{F}}^p([0, T] \times \Omega, \mathbb{R}^n)$. We also recall the *Burkholder-Davis-Gundy* inequality for stochastic integrals:

$$\mathbb{E} \left[\sup_{0 \leq t \leq T} \|y(t)\|^p \right] \leq C_p \mathbb{E} \left[\left(\int_0^T \|x(t)\|^2 dt \right)^{p/2} \right], \quad (1)$$

where C_p is a constant which depends on p uniquely.

As outlined in the introduction, our work is centered on controlling the extreme values or tail of the distribution of a stochastic cost rather than controlling its expectation uniquely. To estimate these extreme values, we leverage risk measures. Unlike expectation, risk measures are not always linear or Fréchet-differentiable. However, they possess several interesting properties we discuss below. We refer to [1, Chapter 6] for all the results presented below on risk measures.

Definition 2.1. A finite risk measure ρ is a mapping from a space $\mathcal{Z} = L^p(\Omega, \mathbb{R})$ with $p \in [1, +\infty)$ to \mathbb{R} . A finite risk measure is said to be coherent if it verifies the following properties:

Convexity:

$$\forall (Z, Z') \in \mathcal{Z} \quad \forall \lambda \in [0, 1], \quad \rho(\lambda Z + (1 - \lambda)Z') \leq \lambda \rho(Z) + (1 - \lambda)\rho(Z').$$

*Monotonicity*¹:

$$\forall (Z, Z') \in \mathcal{Z}, \quad Z \leq Z' \Rightarrow \rho(Z) \leq \rho(Z').$$

Translation equivariance:

$$\forall Z \in \mathcal{Z}, \quad \forall a \in \mathbb{R}, \quad \rho(Z + a) = \rho(Z) + a.$$

Positive homogeneity:

$$\forall Z \in \mathcal{Z}, \quad \forall t \in \mathbb{R}^+, \quad \rho(tZ) = t\rho(Z).$$

Definition 2.1 yields some interesting properties that are used throughout our work.

Theorem 2.2 ([1]). A finite coherent risk measure $\rho : \mathcal{Z} \rightarrow \mathbb{R}$ verifies the following properties:

1. The convex subdifferential of ρ at 0, denoted by $\partial\rho(0)$, is a nonempty closed and bounded subset of \mathcal{Z}^* . It is, in particular, weakly-* compact.
2. For every $Z \in \mathcal{Z}$, the risk measure can be represented as

$$\rho(Z) = \max_{\zeta \in \partial\rho(0)} \mathbb{E}[\zeta Z].$$

3. For every $Z \in \mathcal{Z}$, $\partial\rho(Z)$ is a nonempty closed and bounded subset of \mathcal{Z}^* , and it can be expressed as

$$\partial\rho(Z) = \arg \max_{\zeta \in \partial\rho(0)} \mathbb{E}[\zeta Z].$$

¹for $(Z, Z') \in \mathcal{Z}$, it is said that $Z \leq Z'$ if $\mathbb{P}(\{\omega : Z(\omega) \leq Z'(\omega)\}) = 1$.

4. For every $Z, H \in \mathcal{Z}$, the mapping ρ is Hadamard directionally (sub-) differentiable at Z along H , and the differential can be expressed as

$$\rho'(Z, H) = \max_{\zeta \in \partial\rho(Z)} \mathbb{E}[\zeta H].$$

The most commonly adopted risk measure is the *Conditional-Value-at-Risk*, denoted by CV@R, which we will regularly consider in our examples. Intuitively, CV@R represents the average of the highest values that a random variable can attain. Consequently, it is a highly valuable metric for assessing potential losses in a worst-case scenario.

Definition 2.3. For Z in $L^2(\Omega, \mathbb{R})$ and $\alpha \in (0, 1]$, the *Conditional-Value-at-Risk* of parameter α of Z is given by:

$$CV@R_\alpha[Z] = \inf \left\{ t + \frac{\mathbb{E}[\max(0, Z - t)]}{\alpha} : t \in \mathbb{R} \right\}.$$

The parameter α sets up the number of “worst-case scenarios”. For instance, setting $\alpha = 0.05$ means that $CV@R_\alpha[Z]$ corresponds to the average of the top 5% values that Z can attain. Note that the subdifferential at 0 of the *Conditional-Value-at-Risk* can be explicitly calculated, which is extremely helpful for numerical simulations.

Proposition 2.4. The subdifferential of $CV@R_\alpha$ at 0 is given by:

$$\partial(CV@R_\alpha)(0) = \{\zeta \in L^2(\Omega, \mathbb{R}) : \mathbb{E}[\zeta] = 1, \zeta \in [0, \alpha^{-1}] \text{ a.s.}\}. \quad (2)$$

3. Problem Formulation

In what follows, we detail the settings in which we study the controlled stochastic dynamics. We then devise the optimization problem and present some necessary conditions for optimality an optimal solution must satisfy. As already mentioned, we focus on two different stochastic dynamics:

1. **Setting A.** A general non-linear stochastic system with no control term in the diffusion, which is coherent with our framework of modeling exterior perturbations affecting a mechanical system. The proposed algorithm can be run when considering such dynamics, although only partial theoretical guarantees of convergence may be derived. Nevertheless, through numerical simulations, we show that reliable solutions are often generated.

2. **Setting B.** A control affine stochastic system with no control term in the diffusion. This system is a particular case of the general dynamics which are often encountered in many applications, such as aerospace and robotics [25, 26, 27, 28]. Importantly, under such dynamics, it is possible to theoretically prove the convergence of our algorithm.

From now on, we assume state variables take values in \mathbb{R}^n , whereas control variables take values in \mathbb{R}^m . For the sake of clarity in the exposition and without loss of generality, from now on, we assume the dynamics are perturbed through a weighted one-dimensional Wiener process W , as below.

3.1. *Setting A: The general non-linear control system*

In this general setting, we consider non-linear stochastic dynamics in the form

$$\begin{cases} dX_u(t) &= b(t, X_u(t), u(t))dt + \sigma(t, X_u(t))dW_t, \\ X_u(0) &= X_0, \end{cases} \quad (3)$$

where $X_0 \in \mathbb{R}^n$ is a deterministic initial condition (note that extending our results where the initial condition is a random variable is straightforward). In particular, we consider diffusion terms σ independent from the control input. This is generally the case in many applications such as aerospace and robotics [29]. We also consider the following classical assumptions on the mappings $b : [0, T] \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $\sigma : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ to hold true:

1. The mappings

$$b(\cdot, x, u) : [0, T] \rightarrow \mathbb{R}^n \quad \sigma(\cdot, x) : [0, T] \rightarrow \mathbb{R}^n$$

are continuous for every $(x, u) \in \mathbb{R}^n \times \mathbb{R}^m$.

2. For all t in $[0, T]$, the mappings

$$b(t, \cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n \quad \sigma(t, \cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^n$$

are differentiable with bounded L -Lipschitz gradient. That is, there

exists a constant $L > 0$ such that

$$\begin{aligned} \left\| \frac{\partial b}{\partial x}(t, x, u) \right\| + \left\| \frac{\partial b}{\partial u}(t, x, u) \right\| + \left\| \frac{\partial \sigma}{\partial x}(t, x) \right\| &\leq L, \\ \left\| \frac{\partial b}{\partial x}(t, x, u) - \frac{\partial b}{\partial x}(t, y, v) \right\| &\leq L(|x - y| + |u - v|), \\ \left\| \frac{\partial b}{\partial u}(t, x, u) - \frac{\partial b}{\partial u}(t, y, v) \right\| &\leq L(|x - y| + |u - v|), \\ \left\| \frac{\partial \sigma}{\partial x}(t, x) - \frac{\partial \sigma}{\partial x}(t, y) \right\| &\leq L|x - y|. \end{aligned}$$

In particular, given a progressive control input $u \in L^2_{\mathcal{F}}([0, T] \times \Omega, \mathbb{R}^m)$, the assumptions above guarantee the existence and uniqueness of a solution $X_u \in C^2_{\mathcal{F}}([0, T] \times \Omega, \mathbb{R}^n)$ to (3), see, e.g., [24, Chapter 1.6]. The additional assumption on Lipschitz gradients, in particular, will be required to compute the Fréchet derivative of X_u with respect to u (see Section 4.2).

We consider cost functions of the form

$$J(u) \triangleq \int_0^T f(t, X_u(t), u(t)) dt + g(T, X_u(T)), \quad \text{with } X_u \text{ solution of (3),}$$

where f and g are positive continuous functions, such that $f(t, \cdot, \cdot)$ and $g(t, \cdot)$ are C^1 with Lipschitz gradient in the state and control variables for all t in $[0, T]$. Note that these functions cover a variety of costs that are commonly used in many applications. For instance, in trajectory tracking f consists of the quadratic difference between a reference trajectory and the output trajectory [25].

3.2. Setting B: control affine non-linear systems

In this setting, we consider the particular case of dynamics of the form

$$\begin{cases} dX_u(t) &= [b_0(t, X_u(t)) + \sum_{i=0}^m u_i(t)b_i(t, X_u(t))] dt + \sigma(t, X_u(t))dW_t, \\ X_u(0) &= X_0, \end{cases} \quad (4)$$

where the mappings $b_i : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\sigma : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ are C^1 with Lipschitz gradient in the state variables for all t in $[0, T]$. Again, these assumptions ensure existence and uniqueness of $X_u \in C^2_{\mathcal{F}}([0, T] \times \Omega, \mathbb{R}^n)$ solution to (4). We additionally require the functions b_i to be uniformly

bounded, which still allows for a wide variety of dynamics. Finally, for this setting, we consider slightly less generic stochastic cost functions of the form

$$J(u) \triangleq C_u \|u\|_{\mathcal{U}}^2 + \int_0^T f(t, X_u(t)) dt + g(T, X_u(T)),$$

with X_u solution of (4),

where f and g follow the same assumptions as above. Additionally, we require that f and g satisfy

$$\left\| \frac{\partial(f, g)}{\partial x}(t, x) - \frac{\partial(f, g)}{\partial x}(t, y) \right\| \leq L \|x - y\|, \quad \left\| \frac{\partial(f, g)}{\partial x}(t, 0) \right\| \leq L,$$

for some constant $L > 0$. These assumptions, for instance, enable the use of quadratic functionals. Note that such cost functions are still relevant for many applications, including trajectory tracking-type problems [25].

3.3. Optimization problem

We seek optimal controls in the space $\mathcal{U} = L^2([0, T], \mathbb{R}^m) \times \mathbb{R}^q$ for some fixed $q \in \mathbb{N}$, which parametrizes feedback controls in $L^2_{\mathcal{F}}([0, T] \times \Omega, \mathbb{R}^m)$. More specifically, to each control parameter $u = (v, \lambda) \in \mathcal{U}$ we associate a feedback control in the following manner:

$$u_{v, \lambda}(t, x) = v(t) + K_{\lambda}(x), \tag{5}$$

where $K_{\lambda} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a vector-valued function that is smoothly parametrized by the vector λ . We assume that

$$\begin{aligned} \forall x, y \in \mathbb{R}^n, \forall \lambda \in \mathbb{R}^q, \quad & \|K_{\lambda}(x) - K_{\lambda}(y)\| \leq C(1 + \|\lambda\|)\|x - y\| \\ \forall x \in \mathbb{R}^n, \forall \lambda, \mu \in \mathbb{R}^q, \quad & \|K_{\lambda}(x) - K_{\mu}(x)\| \leq C(1 + \|x\|)\|\lambda - \mu\|. \end{aligned} \tag{6}$$

We also assume the function $(x, \lambda) \mapsto K_{\lambda}(x)$ has a lipschitz gradient, meaning

$$\|\nabla_{\lambda, x} K_{\lambda}(x) - \nabla_{\lambda, x} K_{\mu}(y)\| \leq L(\|x - y\| + \|\lambda - \mu\|) \tag{7}$$

for some constant $L > 0$. In our implementations, we use a feedback law that is linear in λ and x , which therefore automatically verifies both assumptions, but more generic choices are clearly possible. We denote by $X_{v, \lambda} \in C^2_{\mathcal{F}}([0, T] \times \Omega, \mathbb{R}^n)$ the unique solution to (3) (or (4)) which is generated by the control $u_{v, \lambda}$. These control laws will be composed of a deterministic term to steer the trajectory and a stochastic term to compensate for the

diffusion. This choice of controls is motivated by the fact that feedback-type controls are easy to implement and are widely used in practice for risk-averse control [3]². One easily verifies that under (6), the controlled systems under these controls still satisfy the assumptions of Setting A or Setting B. In what follows, we endow the space \mathcal{U} with the scalar product

$$\langle (v_1, \lambda_1), (v_2, \lambda_2) \rangle_{\mathcal{U}} \triangleq \int_0^T v_1(t) \cdot v_2(t) dt + \lambda_1 \cdot \lambda_2.$$

From now on, since clear from the context we equivalently denote $u = (v, \lambda) = u_{v,\lambda}$ and $X_u = X_{v,\lambda} = X_{u_{v,\lambda}}$ for simplicity.

Let us now state and reformulate our optimal control problem as an infinite-dimensional min-max optimization problem using the coherent measure properties described in Theorem 2.2. In particular, let ρ be a finite coherent risk measure defined on $L^1_{\mathcal{F}_T}(\Omega, \mathbb{R})$. The main objective of this paper is to find a local optimum of the problem

$$\min_{u \in \mathcal{U}} \rho(J(u)). \tag{8}$$

This problem is not easily solvable as ρ is not always Fréchet differentiable. Given the properties of coherent risk measures we listed previously, we can rewrite (8) via the following min-max formulation:

$$\min_{u \in \mathcal{U}} \max_{\zeta \in \partial \rho(0)} \mathbb{E}[\zeta J(u)]. \tag{9}$$

Let us remark that $\partial \rho(0)$ is a subset of $L^\infty_{\mathcal{F}_T}(\Omega, \mathbb{R}) \subset L^2_{\mathcal{F}_T}(\Omega, \mathbb{R})$. Using this formulation, we can leverage saddle point algorithms to solve problem (8). Saddle point algorithms have been extensively investigated for min-max optimization, motivating our approach [19, 20, 21, 22].

Before introducing our method to solve (9) (and therefore (8)), we characterize its optimal solutions through appropriate necessary conditions for optimality.

3.4. Necessary conditions for optimality

In what follows, we assume our problem is in the most general Setting A. As of common use in min-max optimization settings, we propose to solve

²Note that in this work, we assume that the state variables are completely observable, leaving the investigation of a more general setting as a future research direction.

Problem (9) by finding extremal points [30]. For this, let us first introduce appropriate first-order necessary conditions for optimality which are satisfied by any optimal solution u^* to (8). These conditions involve the differential of the cost function J that we first need to establish clearly.

Lemma 3.1. *The cost function $J : \mathcal{U} \rightarrow L^2_{\mathcal{F}_T}$ is Fréchet differentiable. Moreover, let $\omega \in \Omega$ and $u \in \mathcal{U}$. Thanks to the Riesz representation theorem in the Hilbert space \mathcal{U} , we can represent the differential $D_u J(u, \omega)$ by an element of \mathcal{U} denoted $\nabla J(u, \omega)$ such that*

$$\forall h \in \mathcal{U}, \quad D_u J(u, \omega) \cdot h = \langle \nabla J(u, \omega), h \rangle_{\mathcal{U}}.$$

The proof can be found in [Appendix A.1](#), along with an explicit formula of ∇J .

Additional notations: In what follows, we call **gradient of J** the function $\nabla J(\cdot, \cdot)$, that we see as a mapping $\nabla J : \mathcal{U} \rightarrow L^2_{\mathcal{F}_T}(\Omega, \mathcal{U})$. We may need to differentiate the two components for clarity, in which case we will write $\nabla J = (\nabla_v J, \nabla_\lambda J)$, with $\nabla_v J(u) \in L^2_{\mathcal{F}_T}(\Omega, L^2([0, T], \mathbb{R}^m))$ and $\nabla_\lambda J(u) \in L^2_{\mathcal{F}_T}(\Omega, \mathbb{R}^q)$. Additionally, for an element z in $L^2_{\mathcal{F}_T}(\Omega, \mathcal{U})$ and u in \mathcal{U} , we will denote by $\langle z, u \rangle_{\mathcal{U}}$ the random variable in $L^2_{\mathcal{F}_T}(\Omega, \mathbb{R})$ such that

$$\forall \omega \in \Omega, \quad \langle z, u \rangle_{\mathcal{U}}(\omega) = \langle z(\omega), u \rangle_{\mathcal{U}}.$$

In particular, we can “invert” the expectation and the scalar product as follows:

$$\mathbb{E}[\langle z, u \rangle_{\mathcal{U}}] = \langle \mathbb{E}[z], u \rangle_{\mathcal{U}}. \quad (10)$$

With these properties in mind, we can establish a necessary condition any optimal solution to (9) must verify as follows:

Proposition 3.2. *Let ρ be a finite coherent risk measure and let u^* be an optimal solution to (8). There exists $\zeta^* \in \partial\rho(0)$ such that*

$$\langle \mathbb{E}[\zeta^* \nabla J(u^*)], h \rangle_{\mathcal{U}} = 0, \quad \text{for every } h \in \mathcal{U}. \quad (11)$$

The proof of Proposition 3.2 is quite similar to the proof given in [31] or [8]. Nevertheless, we report the full proof of Proposition 3.2 in [Appendix A.2](#) for the sake of completeness.

3.5. The modified problem

Seeking extremals of min-max problems is generally a hard task, and not that many algorithms which enjoy convergence guarantees exist in the literature. Notably, for finite-dimensional problems with strongly concave costs, some convergence results are provided in [22]. Inspired by this latter work, we propose an algorithm that works on a modified version of problem (9) whose cost is strongly concave and whose solutions are arbitrarily close to the solutions of the original problem. More specifically, the modified problem takes the following form:

$$\min_{u \in \mathcal{U}} \max_{\zeta \in \partial \rho(0)} \mathbb{E}[(\zeta - \gamma \zeta^2)J(u)], \quad (12)$$

where $\gamma > 0$ is a “concavifying” parameter, which controls how close (12) is to the original problem. Note that the term $\mathbb{E}[\zeta^2 J(u)]$ is well defined as $\zeta \in \partial \rho(0) \subset L_{\mathcal{F}_T}^\infty(\Omega, \mathbb{R})$. In Section 5, we discuss in what sense the solutions to (12) converge to solutions to (9) for $\gamma \rightarrow 0$, endowing our methodology with guarantees of success. Also, from a numerical point of view, we show later that $\gamma > 0$ may be chosen in the order of the machine precision.

4. Ascent-Descent Algorithm

To solve the min-max problem (9), we first introduce in this section a gradient descent-ascent-based approach to solve the modified problem (12). A gradient descent-ascent is an iterative optimization algorithm that can be used to find saddle points. Descent-ascent algorithms represent arguably sound paradigms to seek extremals for general min-max problems. Importantly, under appropriate assumptions, convergence to extremals may be theoretically proven, as we will show later. Until further notice, we start by considering the setting of general non-linear dynamics (Setting A).

4.1. Update rule

The gradient descent-ascent algorithm is an iterative optimization algorithm for min-max problems. At each iteration, it computes the Riesz representation of the gradient of the function being maximized and the gradient of the function being minimized. It then updates the estimates by moving in the direction of their respective gradients. With obvious notations, our

gradient descent-ascent algorithm updates at each iteration n our estimates ζ_n and (v_n, K_{λ_n}) with the following rule:

$$\begin{aligned} v_{n+1} &= v_n - \eta \mathbb{E}[\zeta_n \nabla_v J(v_n, K_{\lambda_n})], \\ \lambda_{n+1} &= \lambda_n - \eta \mathbb{E}[\zeta_n \nabla_\lambda J(v_n, K_{\lambda_n})], \\ \zeta_{n+1} &= \mathcal{P}_{\partial\rho(0)}(\zeta_n + \beta J(v_n, K_{\lambda_n})), \end{aligned} \tag{13}$$

where $\mathcal{P}_{\partial\rho(0)}$ is the projection on the closed convex set $\partial\rho(0)$, and η and β are two positive constants that represent the step size of the gradient. To increase convergence speed, one could make these steps vary from iteration to iteration based on the change of direction of the gradient [32]. We will keep them constant for the sake of simplicity. All convergence results should be valid for a range of step sizes that are both upper and lower bounded, although proof of this improvement is left for future work. We recall that $\partial\rho(0)$ can be viewed as a subset of $L_{\mathcal{F}_T}^\infty(\Omega, \mathbb{R})$, which justifies the above mathematical operations between ζ_n and the other terms. Note that it is important to project $\zeta_n + \beta J(u_n)$ onto $\partial\rho(0)$, as there is no a priori guarantee that $\zeta_n + \beta J(u_n)$ stays inside $\partial\rho(0)$.

Unfortunately, through the scheme (13), no clear convergence result is available. To develop theoretical guarantees of convergence, we propose to rather leverage the following alternative rule, which essentially represents an update rule for Problem (12):

$$\begin{aligned} v_{n+1} &= v_n - \eta \mathbb{E}[\zeta_n(1 - \gamma\zeta_n) \nabla_v J(v_n, K_{\lambda_n})], \\ \lambda_{n+1} &= \lambda_n - \eta \mathbb{E}[\zeta_n(1 - \gamma\zeta_n) \nabla_\lambda J(v_n, K_{\lambda_n})], \\ \zeta_{n+1} &= \mathcal{P}_{\partial\rho(0)}(\zeta_n + \beta(1 - 2\gamma\zeta_n)J(v_n, K_{\lambda_n})). \end{aligned} \tag{14}$$

The algorithms described by (13) and (14) offer the advantage of being easily implementable, provided explicit expressions of the gradient ∇J are available. We provide such expressions in the following section.

4.2. Computation of gradients

The practical implementation of (13) and (14) requires computing the gradient of the cost, and in particular of the trajectory, with respect to both the variables ζ and u . Below, we summarize the expressions of this gradient, referring the reader to Appendix A.1 for their proof. Let us clarify that we denote $K'_\lambda(x)$ the derivative of $K_\lambda(x)$ with respect to x .

Proposition 4.1. *Fix a control $u = (v, \lambda) \in \mathcal{U}$. Under the assumptions of Setting A, the Riesz representation of the gradient of J at $u = (v, \lambda)$ is*

$$\begin{aligned}\nabla_v J(v, \lambda)(t) &= I_u(t)\phi_u^{-1}(t)\partial_u b(t, X_u(t), u(t)) \\ \nabla_\lambda J(v, \lambda) &= \int_0^T I_u(s)\phi_u^{-1}(s)\partial_u b(s, X_u(s), u(s))\nabla_\lambda K_\lambda(X_u(s))ds.\end{aligned}$$

where

$$\begin{aligned}I_u(s) &\triangleq \partial_x g(T, X_u(T))\phi_u(T) \\ &+ \int_s^T [\partial_x f(t, X_u(t), u(t)) + \partial_u f(t, X_u(t), u(t))K'_\lambda(X_u(t))]\phi_u(t)dt.\end{aligned}$$

Here, $\phi_u(\cdot)$ is the fundamental matrix solution to the linearized problem, i.e., the unique solution to

$$\begin{cases} d\phi(t) = [\partial_x b(t, X_u(t), u(t)) + \partial_u b(t, X_u(t), u(t))K'_\lambda(X_u(t))]\phi(t)dt \\ \quad + \partial_x \sigma(t, X_u(t))\phi(t)dW_t, \\ \phi(0) = Id_n. \end{cases} \quad (15)$$

Thus, the primary challenge of implementing this method lies in simulating the trajectories of X_u and ϕ_u and computing expectations of random variables that involve terms from those trajectories. To deal with these technical details, we used a classical approach presented in [33, Chapter 10] and refer to it for more extensive elaboration. The contents of this section furnish the necessary components to implement the algorithm and evaluate its performance for the general Setting A.

In the next section, we establish some convergence results that may be of particular significance for Setting B, as they provide insight into the convergence of the algorithm towards a critical point that satisfies the necessary optimality condition (11).

5. Convergence Results

We leverage the method introduced in [22] to examine the convergence of the smooth non-convex/strong concave gradient descent-ascent algorithm in the most general Setting A. In what follows, we adapt the approach to an infinite-dimensional space. With some minor adjustments, we show that the gradient of the cost converges to 0 and that the cost function converges

to a finite value. While the study in [22] also establishes that in finite-dimensional spaces, every converging subsequence of the control estimates tends to a critical point, we were unable to prove this result in our general setting. Nevertheless, in the slightly more restrictive Setting B, any weakly converging subsequence of our control estimates converges to a point that is a “ γ -critical point.” That is, a control \bar{u}_γ that verifies

$$\exists \bar{\zeta}, \langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} = \|h\|_{\mathcal{U}} O(\gamma) \quad \forall h \in \mathcal{U}. \quad (16)$$

This important result shows that by progressively selecting smaller values of γ , we can obtain a control that verifies the necessary optimality condition (11). In addition, taking $\gamma \rightarrow 0$ yields critical points for the original problem, as we state at the end of this section.

5.1. Convergence of the gradient in Setting A

Let (u_n, ζ_n) be the estimates given by algorithm (14), in what follows, we prove the convergence of $\|u_{n+1} - u_n\|_{\mathcal{U}}$ to 0 in Setting A, as well as the convergence of the γ -approximated cost, denoted $\Phi_\gamma : \mathcal{U} \rightarrow \mathbb{R}$ and defined as

$$\Phi_\gamma(u) \triangleq \max_{\zeta \in \partial \rho(0)} \mathbb{E}[\zeta(1 - \gamma\zeta)J(u_n)].$$

The proof of the convergence closely follows the one presented in [22], and we, therefore, do not explicit all of the technical details. We indicate where changes were needed and develop the modifications in the Appendix for completeness. For the sake of clarity, we denote

$$\zeta_u \triangleq \arg \max_{\zeta \in \partial \rho(0)} \mathbb{E}[\zeta(1 - \gamma\zeta)J(u)]$$

which is well defined as $\zeta \mapsto \mathbb{E}[\zeta(1 - \gamma\zeta)J(u)]$ is a weakly continuous and strongly concave function over a weakly compact set.

Proposition 5.1. *Let $L > 0$ be such that J has L -Lipschitz gradient, and $\mu > 0$ such that $\zeta \mapsto \mathbb{E}[\zeta(1 - \gamma\zeta)\nabla J(u_n)]$ is μ -strongly concave. We define $K = \frac{L}{\mu}$. If the gradient rates η and β verify $\eta < \frac{1}{K^3(L+3)^2}$ and $\beta < \frac{1}{L}$, then the sequences (u_n, ζ_n) given by the update rule (14) verify:*

$$\|u_{n+1} - u_n\|_{\mathcal{U}} \rightarrow 0, \quad \|\zeta_{n+1} - \zeta_n\|_{L^2_{\mathcal{F}_T}} \rightarrow 0, \quad \|\zeta_n - \zeta_{u_n}\|_{L^2_{\mathcal{F}_T}} \rightarrow 0,$$

Proof. As mentioned above, we generalise the proof from [22, Proposition 2] with $f(u, \zeta) = \mathbb{E}[\zeta(1 - \gamma\zeta)J(u)]$, $g = 0$ and $h = \delta_{\partial\rho(0)}$, the convex characteristic function of $\partial\rho(0)$. The key argument is to consider the function

$$H(u, \zeta) = \Phi_\gamma(u) + \left(1 - \frac{1}{4K^2}\right) \|\zeta - \zeta_u\|_{L^2_{\mathcal{F}_T}}^2 \quad (17)$$

and show that $H(u_n, \zeta_n)$ is decreasing over the iterations of the algorithm. To prove it, we must show the Lipschitzianity of the functions $u \mapsto \zeta_u$ and $u \mapsto \nabla\Phi_\gamma(u)$. The first one is straightforward, as shown in [22, Proposition 1]. The second one however needs some adjustments to our infinite-dimensional setting. Using the version of the Danskin theorem found in [34] with the weak topology, we can first prove the following lemma.

Lemma 5.2. Φ_γ is Fréchet differentiable and its gradient is given by

$$\nabla\Phi_\gamma = \mathbb{E}[\zeta_u(1 - \gamma\zeta_u)\nabla J(u)].$$

The Lipschitzianity of $\nabla\Phi_\gamma(u)$ can be deduced as follows:

$$\begin{aligned} \|\Phi(u_2) - \Phi(u_1)\| &= \|\nabla_u \mathcal{J}_\gamma(u_2, \zeta_{u_2}) - \nabla_u \mathcal{J}_\gamma(u_1, \zeta_{u_1})\| \\ &\leq L(\|u_2 - u_1\| + \|\zeta_{u_2} - \zeta_{u_1}\|) \\ &\leq L(1 + K)\|u_2 - u_1\| \end{aligned}$$

With these two Lipschitzianity results, we can execute the same computations as in [22, Proposition 2] and prove the inequality

$$H(u_{n+1}, \zeta_{n+1}) \leq H(u_n, \zeta_n) - 2\|u_{n+1} - u_n\|_{\mathcal{U}}^2 - \frac{1}{4K^2} (\|\zeta_{n+1} - \zeta_{u_{n+1}}\|_{L^2_{\mathcal{F}_T}}^2 + \|\zeta_n - \zeta_{u_n}\|_{L^2_{\mathcal{F}_T}}^2).$$

From this inequality we can deduce the sequence $H(u_n, \zeta_n)$ is decreasing and therefore converges to H^* . By summing the inequality over the iterations, we obtain a telescoping series that yields

$$\begin{aligned} &\sum_{n=0}^N \left[2\|u_{n+1} - u_n\|_{\mathcal{U}}^2 - \frac{1}{4K^2} (\|\zeta_{n+1} - \zeta_{u_{n+1}}\|_{L^2_{\mathcal{F}_T}}^2 + \|\zeta_n - \zeta_{u_n}\|_{L^2_{\mathcal{F}_T}}^2) \right] \\ &\leq \sum_{n=0}^N H(u_n, \zeta_n) - H(u_{n+1}, \zeta_{n+1}) \leq H(u_0, \zeta_0) - H^*. \end{aligned}$$

This shows that the series on the left is convergent, meaning that its terms tend to zero, which proves our proposition. \square

This result additionally shows some other interesting properties of the algorithm, which are of straightforward proof.

Corollary 5.3.

1. Since $\|\zeta_n - \zeta_{u_n}\|_{L^2_{\mathcal{F}_T}}^2$ tends to zero and H decreases, the sequence $\Phi_\gamma(u_n)$ converges to a finite limit.
2. If the cost J is coercive, the sequence u_n is bounded in \mathcal{U} and is thus in a weakly compact set.

To go further, one would wish to find a generalization of [22, Theorem 1] by showing that all weak limit points of our sequence u_n are "gamma-critical points" and verify equation (16). Unfortunately, we could not prove this in the general case, as a weakly converging sequence of controls u_n does not necessarily yield a converging sequence of states x_{u_n} . Nonetheless, we can demonstrate the result in the slightly more restrictive setting of linear affine controls, as shown in the next subsection.

5.2. Convergence of the gradient in Setting B

As mentioned previously, restricting to control affine systems allows us to use important theoretical properties and improve convergence guarantees for our gradient ascent-descent algorithm. The most important property we take advantage of is the compactity of the input-output mapping $u \mapsto X_u$, as shown in the following lemma:

Lemma 5.4. *Let u_n be a sequence controls in \mathcal{U} , and $p \geq 2$. Under the hypothesis of Setting B, we have that*

$$u_n \xrightarrow{\mathcal{U}} \bar{u} \Rightarrow X_{u_n} \xrightarrow{C_X^p} X_{\bar{u}}.$$

Consequently, if we decompose the cost J of Setting B as $J(u) = C_u \|u\|_{\mathcal{U}}^2 + J_X(X_u)$, we have

$$u_n \xrightarrow{\mathcal{U}} \bar{u} \Rightarrow \nabla_u(J_X(X_{u_n})) \xrightarrow{C_X^2} \nabla_u(J_X(X_{\bar{u}})).$$

The proof of this lemma requires some long and technical computations that we detailed in [Appendix A.3](#). This compactness result, in turn, allows us to prove the following:

Proposition 5.5. *Let \bar{u} be a weak limit of a subsequence of u_n , then for any $\bar{\zeta} \in \partial\rho(0)$ that is a weak limit of a subsequence of ζ_n we have that*

$$\langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} = \|h\|_{\mathcal{U}} O(\gamma) \quad \forall v \in \mathcal{U}. \quad (18)$$

Proof. Let us first use our result on the convergence of the general algorithm. We can see that

$$\begin{aligned} v_{n+1} - v_n &= -\eta \mathbb{E}[(\zeta_n - \gamma \zeta_n^2) \nabla_v J(u_n)], \\ &= -\eta \mathbb{E}[\zeta_n \nabla_v J(u_n)] + \eta \gamma \mathbb{E}[\zeta_n^2 \nabla_v J(u_n)]. \end{aligned}$$

Since ζ_n is bounded in $L^\infty(\Omega, \mathbf{R})$, $\nabla_v J(u_n)$ is bounded in $L^2(\Omega, \mathbf{R})$, and $(v_{n+1}, \lambda_{n+1}) - (v_n, \lambda_n)$ tends to 0, there exists a constant $C > 0$ such that

$$\|\mathbb{E}[\zeta_n \nabla_v J(u_n)]\|_{L^2(0,T)} \leq C\gamma = O(\gamma). \quad (19)$$

Below, we will implicitly overload the constant C . Similarly, we have

$$|\mathbb{E}[\zeta_n \nabla_\lambda J(u_n)]| \leq C\gamma = O(\gamma). \quad (20)$$

We now go back to the equality we want to prove. For this, let h be a control in \mathcal{U} , and let \bar{u} and $\bar{\zeta}$ be weak limits of respectively u_n and ζ_n . We may compute

$$\begin{aligned} \langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} &= \langle \mathbb{E}[(\bar{\zeta} - \zeta_n) \nabla J(\bar{u})], h \rangle_{\mathcal{U}} \\ &\quad + \langle \mathbb{E}[\zeta_n (\nabla J(\bar{u}) - \nabla J(u_n))], h \rangle_{\mathcal{U}} \\ &\quad + \langle \mathbb{E}[\zeta_n \nabla J(u_n)], h \rangle_{\mathcal{U}}. \end{aligned}$$

Using equations (19) and (20), we have that

$$|\langle \mathbb{E}[\zeta_n \nabla J(u_n)], h \rangle_{\mathcal{U}}| = \|h\|_{\mathcal{U}} O(\gamma) \quad \forall v \in \mathcal{U}.$$

and since ζ_n weakly converges to ζ , with an abuse of notation we may write

$$\langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} = \langle \mathbb{E}[\zeta_n (\nabla J(\bar{u}) - \nabla J(u_n))], h \rangle_{\mathcal{U}} + \|h\|_{\mathcal{U}} O(\gamma).$$

By further decomposing $J(u) = C_u \|u\|_{\mathcal{U}}^2 + J_X(X_u)$, we additionally compute

$$\begin{aligned} \langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} &= 2C_u \langle \mathbb{E}[\zeta_n (\bar{u} - u_n)], h \rangle_{\mathcal{U}} \\ &\quad + \langle \mathbb{E}[\zeta_n (\nabla_u J_X(X_{\bar{u}}) - \nabla_u J_X(X_{u_n}))], h \rangle_{\mathcal{U}} + \|h\|_{\mathcal{U}} O(\gamma). \end{aligned}$$

Using Lemma 5.4 and the fact that $u = (v, \lambda)$ is deterministic, again with an abuse of notation, we can write

$$\langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} = 2C_u \mathbb{E}[\zeta_n] \langle \bar{u} - u_n, h \rangle_{\mathcal{U}} + \|h\|_{\mathcal{U}} O(\gamma),$$

and thanks to the weak convergence of u_n to u^* we finally obtain that

$$\langle \mathbb{E}[\bar{\zeta} \nabla J(\bar{u})], h \rangle_{\mathcal{U}} = \|h\|_{\mathcal{U}} O(\gamma).$$

□

This result implies that, as we make γ tend to 0, the solutions obtained will get closer to our original problem, as summarized in the following corollary.

Corollary 5.6. *Let \bar{u}^γ be a weak limit of a sequence of controls u_n^γ obtained with the update rule (14) with parameter γ . For any decreasing sequence $(\gamma_k)_{k \in \mathbb{N}}$ tending to 0, we can define a sequence $(\bar{u}^{\gamma_k})_{k \in \mathbb{N}}$ of controls whose weak limits all verify the original optimality condition (11).*

This corollary indicates that if we find with our algorithm a γ -critical point for smaller and smaller values of γ , we can find a control law that verifies the necessary optimality condition. Theoretically, lowering γ should imply lowering the gradient step η to satisfy the assumptions in Proposition 5.1, implying a slower convergence in practice. However, the numerical results we present in the next section show that the two convergences, i.e., with $\gamma = 0$ or $\gamma \sim 0$, are essentially equivalent, thus there is no downside in taking $\gamma > 0$ very small, e.g., in the order of the epsilon machine.

6. Numerical Simulations

In this section, we present some simulation results. Our algorithm was implemented and tested on two systems, one in Setting A, and the other in Setting B. We respectively denote them as System A and System B. By testing our algorithm on these two different systems, we can evaluate its effectiveness and robustness in solving optimal control problems in both settings. System A is an academic non-linear system that simulates a steering problem, where a vehicle is controlled through its angular velocity while moving at a constant speed. It is subject to perturbations in the state which, for instance, may model measurement errors. On the other hand, System B is

a control affine system that models a real-world robot called AstroBee [35], designed to operate in a zero-gravity environment, such as a space station. The diffusion term may, for instance, model sudden gusts of air randomly generated by onboard air purification systems, which destabilizes the robot. For simplicity, we designed and leveraged a very simple stochastic model to model the disturbances; note that other models exist in the literature, e.g., [36, 37].

6.1. System presentation

For both systems, the objective is to move from one point to another in a corridor while following a reference trajectory and avoiding obstacles. The reference trajectory is given by a reference control that provides a satisfying solution to the deterministic problem. However, when we add perturbations to the system, the trajectories resulting from the reference control often collide with obstacles. To overcome this challenge, our algorithm searches for a control strategy that minimizes a risk measure of the stochastic cost function. The cost function is composed of two terms: a trajectory-tracking term that tracks the reference trajectory and a term $J_{obstacles}(X)$ that heavily penalizes trajectories that collide with obstacles. This *penalization method* as presented in [38] is a classic optimal control approach to get rid of the constraints and is widely used in the literature [39, 40, 41].

$$J(u) = \|u\|_{\mathcal{U}}^2 + \int_0^T \|X_u(t) - X_{ref}(t)\|^2 dt + J_{obstacles}(X_u).$$

The term $J_{obstacles}(X)$ acts as a potential that is extremely high in the obstacles and thus forces the trajectories away. The risk measure significantly weighs the worst-case scenarios, which means that the cost of a control strategy will increase substantially if some random trajectories produced by the control collide with an obstacle. A control minimizing the risk measure should therefore avoid collisions with high probability.

6.2. Results

As shown in Figure 1, our algorithm converges arbitrarily close to solutions for both System A and System B, as long as we use a sufficiently small step size. Additionally, implementing the algorithm with an infinitesimal value of γ results in the same level of convergence as when using a non-negligible value of γ . In particular, although theoretical proofs typically

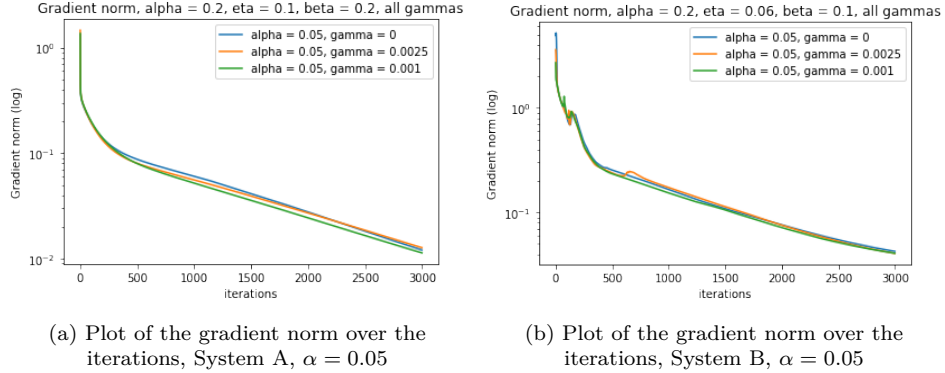


Figure 1: Convergence rates obtained for $\alpha = 0.05$ and different values of γ , for both the steering system (left) and the AstroBee system (right).

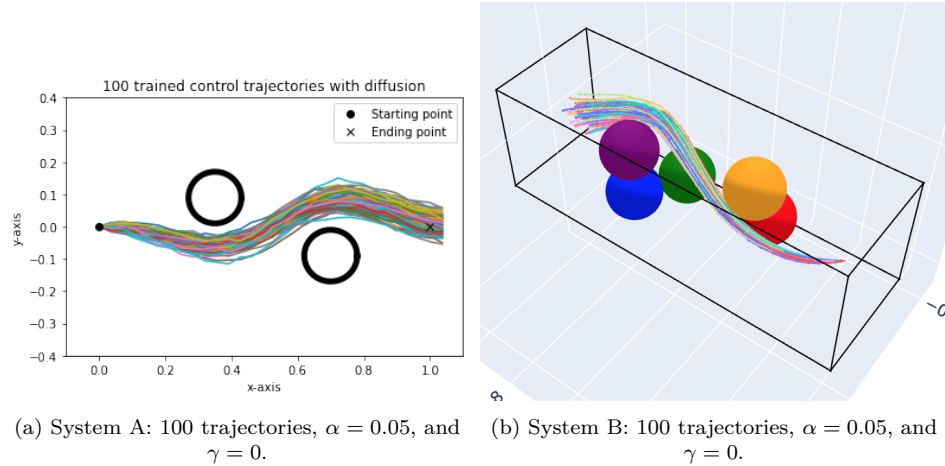


Figure 2: 100 trajectories given by a trained control for both systems.

require γ to be strictly greater than 0, our numerical results suggest that it may not always be necessary to achieve convergence.

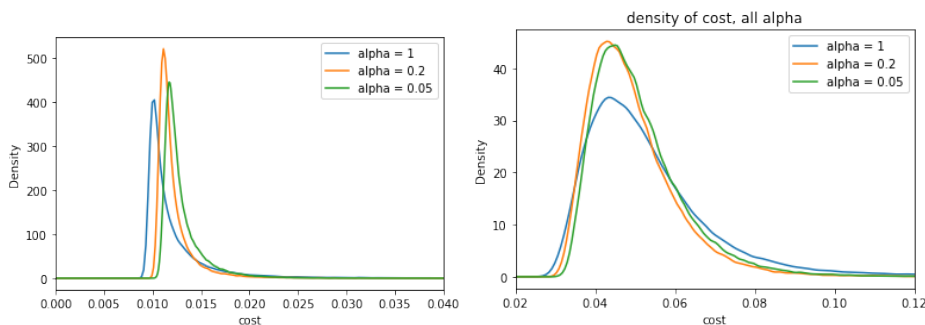
In Table 1, we evaluate the effectiveness of our algorithm under three different values of α : $\alpha = 0.05$, $\alpha = 0.2$, and $\alpha = 1$ (the latter value corresponds to the minimization of the mean). By testing our algorithm under these different values of α , we can assess the impact of different risk levels on the control strategy. Our main metric is the percentage of trajectories produced by the control strategy that resulted in collisions with obstacles. We compare these percentages between strategies obtained with different values of α to assess its impact on the level of risk in the control strategy.

Systems	$\alpha = 1$	$\alpha = 0.2$	$\alpha = 0.05$
System A	0.353%	0.112%	0.011%
System B	0.049%	0.007%	0.004%

Table 1: Percentage of collisions with obstacles for different α and $\gamma = 0$. Control strategies trained with higher values of α yield significantly more trajectories colliding with obstacles.

To evaluate the robustness of the control strategy, we can increase the size of the obstacles to see how much the probability of collision increases. We rapidly re-trained the controls in this modified problem, with obstacles of a bigger size. This allows us to assess the ability of the control strategy to handle more challenging and uncertain conditions. From Table 4, for both System A and System B, we observe the control solutions that minimized the risk measure with a low value of α are the most conservative. This results in trajectories that keep a larger distance from obstacles but also deviate more from the reference trajectory.

To further assess the impact of the parameter α on the trade-off between mean scores and worst-case scenarios, we computed the probability densities of the costs of the obtained control strategies. By comparing the probability densities of the costs under different values of α , we can identify the optimal level of risk for each system and obtain a control strategy that balances the trade-off between mean scores and worst-case scenarios. As expected, density distributions have higher means but smaller tails for lower values of α . This suggests that a more conservative risk estimate leads to a more robust control strategy with lower worst-case scenarios but also a lower mean score.



(a) Density of the cost for System A.

(b) Density of the cost for System B.

Figure 3: Densities of the cost for different α and $\gamma = 0$.

Obstacles bonus size	$\alpha = 1$	$\alpha = 0.2$	$\alpha = 0.05$
+0%	0.353%	0.112%	0.011%
+5%	0.367%	0.111%	0.011%
+10%	0.470%	0.145%	0.013%
+15%	0.600%	0.177%	0.015%
+20%	0.735%	0.225%	0.020%

Table 2: Percentage of collisions with obstacles in the steering problem.

Obstacles bonus size	$\alpha = 1$	$\alpha = 0.2$	$\alpha = 0.05$
+0%	0.049%	0.007%	0.004%
+5%	0.044%	0.006%	0.004%
+10%	0.048%	0.006%	0.006%
+15%	0.054%	0.006%	0.006%
+20%	0.066%	0.008%	0.006%

Table 3: Percentage of collisions with obstacles in the AstroBee system.

Table 4: Percentage of collisions with obstacles for different α in the two systems.

7. Conclusion and future directions

In this paper, we develop an algorithm to solve risk-averse stochastic optimal control problems subject to non-linear stochastic differential equations, where a risk measure replaces expectation in the cost. By leveraging duality results for coherent risk measures, we recast our minimization problem into a min-max problem that we solve with a gradient descent-ascent approach. We prove convergence properties for this algorithm, which we showcase through appropriate numerical implementation on non-trivial control systems. Our results show our algorithm genuinely converges under very general assumptions and settings. Importantly, our approach of minimizing the cost under a risk measure instead of the expectation yields more robust control strategies, which are capable of considerably better mitigating the uncertainties generated by the diffusion term in the stochastic differential equation.

Several exciting avenues for further exploration and testing are listed hereafter. One direction would consist in extending our algorithm to sys-

tems that are modeled through more complex stochastic processes, such as processes with control variables in the diffusion or processes with jumps. Control-dependent diffusion would allow us to model systems where the uncertainties are caused by inputs as well. While putting a control term in the diffusion is already theoretically feasible, we found it to be computationally very expensive as it requires computing conditional expectations, calling for deeper investigation. In particular, recent techniques found in the literature [42] may allow us to efficiently compute these conditional expectations. Another direction consists of testing our approach for other risk measures than just the CV@R. Additionally, we are interested in extending our approach to more sophisticated agents, e.g., soft robots. Soft robots have recently received a particular surge of interest as they appear in an increasing number of applications. An emerging model for the dynamics of soft robots hinges upon infinite-dimensional dynamics ruled by Partial Differential Equations (PDEs) [43]. It would be beneficial to extend our approach to handle stochastic PDEs: we look forward to future developments in this area.

Appendix A. Technical proofs

Appendix A.1. Proof of Lemma 3.1

In this appendix, we give the proof of Lemma 3.1. We want to show the differentiability of the cost function. We assume to be dealing with the cost and dynamics in Setting A. We prove first that the function $J : \mathcal{U} \rightarrow L^2_{\mathcal{F}_T}$ is Fréchet differentiable. We also retrieve a numerically tractable formula for the Riesz representation of the differential, that we denote $\nabla J(u) = (\nabla_v J(u), \nabla_\lambda J(u))$. For the sake of clarity, we recall that if $u = (v, \lambda)$ is a control in \mathcal{U} , we denote $u(t)$ the value $u(t) \triangleq v(t) + K_\lambda(X_u(t))$, where X_u uniquely solves (3) with control (5). We also denote $K'_\lambda(x)$ the differential of $K_\lambda(x)$ with respect to x . We first start giving two preliminary lemmas.

Lemma Appendix A.1. *Under assumptions of Setting A, let $u = (v_u, \lambda_u)$ and $h = (v_h, \lambda_h)$ be two controls in \mathcal{U} . There exists $C_{u,h}$ a positive continuous function of $\|u\|_{\mathcal{U}}$ and $\|h\|_{\mathcal{U}}$ such that*

$$\|X_u - X_h\|_{C^2_{\mathcal{F}}}^2 \leq C_{u,h} \|u - h\|_{\mathcal{U}}^2. \quad (\text{A.1})$$

Proof. Let $u = (v_u, \lambda_u)$ and $h = (v_h, \lambda_h)$ be two controls in \mathcal{U} . The difference

in the two trajectories X_u and X_h is given by

$$\begin{aligned} X_u(s) - X_h(s) &= \int_0^s [b(r, X_u(r), u(r)) - b(r, X_h(r), h(r))] dr \\ &\quad + \int_0^s [\sigma(r, X_u(r)) - \sigma(r, X_h(r))] dW_r. \end{aligned}$$

where we denote the control value $u(r) = v_u(r) + K_{\lambda_u}(X_u(r))$. If we note $C > 0$ a positive constant (that will be overloaded in the rest of the proof), we can bound the $C_{\mathcal{F}}^2$ norm of the difference above by

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] &\leq C \mathbb{E} \left[\sup_{0 \leq s \leq t} \left\| \int_0^s [\sigma(r, X_u(r)) - \sigma(r, X_h(r))] dW_r \right\|^2 \right] \\ &\quad + C \mathbb{E} \left[\sup_{0 \leq s \leq t} \left\| \int_0^s [b(r, X_u(r), u(r)) - b(r, X_h(r), h(r))] dr \right\|^2 \right]. \end{aligned}$$

Using Cauchy-Schwartz in the deterministic integral and the *Burkholder-Davis-Gundy* inequality in the stochastic integral, we get:

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] &\leq C \mathbb{E} \left[\int_0^t \|\sigma(r, X_u(r)) - \sigma(r, X_h(r))\|^2 dr \right] \\ &\quad + C \mathbb{E} \left[\int_0^t \|b(r, X_u(r), u(r)) - b(r, X_h(r), h(r))\|^2 dr \right]. \end{aligned}$$

By using the lipschitzianity of the dynamics and taking the sup in the integral, we get

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] &\leq C \mathbb{E} \left[\int_0^t \sup_{0 \leq r \leq s} \|X_u(r) - X_h(r)\|^2 ds \right] \\ &\quad + C \mathbb{E} \left[\int_0^t \|u(s) - h(s)\|^2 ds \right]. \end{aligned}$$

We can write

$$\begin{aligned} u(s) - h(s) &= v_u(s) - v_h(s) + K_{\lambda_u}(X_u(s)) - K_{\lambda_h}(X_u(s)) \\ &\quad + K_{\lambda_h}(X_u(s)) - K_{\lambda_h}(X_h(s)), \end{aligned}$$

and therefore have

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] &\leq C \mathbb{E} \left[\int_0^t \sup_{0 \leq r \leq s} \|X_u(r) - X_h(r)\|^2 ds \right] \\ &\quad + C \mathbb{E} \left[\int_0^t \|v_u(s) - v_h(s)\|^2 ds \right] + C \mathbb{E} \left[\int_0^t \|K_{\lambda_u}(X_u(s)) - K_{\lambda_h}(X_u(s))\|^2 ds \right] \\ &\quad + C \mathbb{E} \left[\int_0^t \|K_{\lambda_h}(X_u(s)) - K_{\lambda_h}(X_h(s))\|^2 ds \right]. \end{aligned}$$

Using the assumptions on $K_\lambda(X)$ in (6), we obtain that

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] &\leq C \mathbb{E} \left[\int_0^t \sup_{0 \leq r \leq s} \|X_u(r) - X_h(r)\|^2 ds \right] \\ &+ C \int_0^t \|v_u(s) - v_h(s)\|^2 ds + C \left(1 + \mathbb{E} \left[\int_0^t \|X_u(s)\|^2 ds \right] \right) \|\lambda_u - \lambda_h\|^2 \\ &+ C(1 + \|\lambda_h\|^2) \mathbb{E} \left[\int_0^t \|X_u(s) - X_h(s)\|^2 ds \right]. \end{aligned}$$

The dynamics of the SDE (3) are continuous in $u = (v, \lambda)$ and the solution X_u changes continuously with the dynamics [24], implying that $\mathbb{E} \left[\int_0^t \|X_u(s)\|^2 ds \right]$ is a continuous function of u . We may therefore introduce $C_{u,h}$, an overloaded constant that depends continuously on u and h . From this remark, by taking the sup in the integral above, we may infer that

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] &\leq C_{u,h} \mathbb{E} \left[\int_0^t \sup_{0 \leq r \leq s} \|X_u(r) - X_h(r)\|^2 ds \right] \\ &+ C_{u,h} \left[\int_0^t \|v_u(s) - v_h(s)\|^2 ds + \|\lambda_u - \lambda_h\|^2 \right]. \end{aligned}$$

Finally, the Gronwall lemma applied to $t \mapsto \sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2$ yields

$$\mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_u(s) - X_h(s)\|^2 \right] \leq C_{u,h} \|u(s) - h(s)\|_{\mathcal{U}}^2.$$

□

In what follows, we denote $y_{u,h}$ the unique element of $L^2_{\mathcal{F}}(\Omega \times [0, T], \mathbf{R}^n)$ that verifies the linearized SDE

Lemma Appendix A.2. *Let $u = (v_u, \lambda_u)$ and $h = (v_h, \lambda_h)$ be two controls in \mathcal{U} , let us consider the following linear SDE with stochastic dynamics:*

$$\begin{cases} dY(t) &= [\partial_x b(t, X_u(t), u(t)) + \partial_u b(t, X_u(t), u(t)) K'_{\lambda_u}(X_u(t))] Y(t) \\ &+ \partial_u b(t, X_u(t), u(t)) [v_h(t) + \nabla_\lambda K_{\lambda_u}(X_u(t)) \lambda_h] dt \\ &+ \partial_x \sigma(t, X_u(t)) Y(t) dW_t, \\ Y(0) &= 0. \end{cases}$$

This equation has a unique solution denoted $y_{u,h} \in L^2_{\mathcal{F}}(\Omega \times [0, T], \mathbf{R}^n)$. This solution can be written explicitly as

$$y_{u,h}(t) = \phi_u(t) \int_0^t \phi_u^{-1}(s) \partial_u b(s, X_u(s), u(s)) [v_h(s) + \nabla_\lambda K_\lambda(X_u(s)) \lambda_h] ds$$

where ϕ_u , defined in (15), is the resolvent of the linear SDE (15). Consider ϵ a positive constant and denote $\delta X_{u,\epsilon,h} \triangleq X_{u+\epsilon h} - X_u$. Then,

$$\|\delta X_{u,\epsilon,h} - \epsilon y_{u,h}\|_{C_{\mathcal{F}}^2} \leq C(\epsilon)\epsilon\|h\|_{L_{\mathcal{F}}^2} = o(\epsilon)$$

where $C(\epsilon) > 0$ is a real positive function of ϵ that tends to zero as $\epsilon \rightarrow 0$.

Proof. The process Y is the solution of a linear equation of the form

$$dY(t) = (A(t)Y(t) + r(t))dt + C(t)Y(t)dW_t,$$

where $A(t), C(t)$ and $r(t)$ are \mathcal{F}_t -adapted processes. If $A(t)$ and $C(t)$ are $L_{\mathcal{F}}^\infty(\Omega \times [0, T])$, and $r(t)$ is $L_{\mathcal{F}}^2(\Omega \times [0, T])$, then [44] ensures us that there exists a unique strong solution that can be written explicitly. The assumptions made on $K_\lambda(x)$ ensure us that

$$\|K'_\lambda(x)\| \leq C(1 + \|\lambda\|), \quad \|\nabla_\lambda K_\lambda(x)\| \leq C(1 + \|x\|),$$

which in turn implies

$$\begin{aligned} \|A\|_\infty &\leq \|\partial_x b\|_\infty + C\|\partial_u b\|_\infty(1 + \|\lambda_u\|) \\ \|C\|_\infty &\leq \|\partial_x \sigma\|_\infty \\ \|r\|_{L^2} &\leq \|\partial_u b\|_\infty(\|v_h\|_{L^2} + C(1 + \|X_u\|_{C_{\mathcal{F}}^2})\|\lambda_h\|) \end{aligned}$$

The solution therefore exists and can be written explicitly.

As for the approximation, the functions $\delta X_{u,\epsilon,h}$ and $y_{u,h}$ follow the same SDE up to a $o(\epsilon)$ term. Therefore, the proof follows the exact same steps as the proof of Lemma [Appendix A.1](#), which we then avoid detailing here. \square

We can now move to the proof of Lemma [3.1](#). To prove differentiability, we start by estimating the difference

$$\begin{aligned} \delta J_{u,\epsilon,h} &= J(u + \epsilon h) - J(u) \\ &= \int_0^T [f(t, X_{u+\epsilon h}(t), v_u(t) + \epsilon v_h(t) + K_{\lambda_u + \epsilon \lambda_h}(X_{u+\epsilon h}(t))) \\ &\quad - f(t, X_u(t), v_u(t) + K_{\lambda_u}(X_u(t)))]dt + g(T, X_{u+\epsilon h}(T)) - g(T, X_u(T)) \\ &= \int_0^T \partial_x f(t, X_u(t), u(t))\delta X_{u,\epsilon,h}(t)dt \\ &\quad + \int_0^T \partial_u f(t, X_u(t), u(t))[\epsilon v_h(t) + K_{\lambda_u + \epsilon \lambda_h}(X_{u+\epsilon h}(t)) - K_{\lambda_u}(X_u(t))]dt \\ &\quad + \partial_x g(T, X_u(T))\delta X_{u,\epsilon,h}(T) + o(\epsilon) + o(\delta X_{u,\epsilon,h}). \end{aligned}$$

where $\delta X_{u,\epsilon,h} \triangleq X_{u+\epsilon h} - X_u$. Due to Lemma [Appendix A.1](#), we have

$$\|X_{u+\epsilon h} - X_u\|_{C_{\mathcal{F}}^2}^2 \leq C\epsilon^2 \|h\|_{\mathcal{U}}^2,$$

where C is a positive constant, meaning that $o(\delta X_{u,\epsilon,h}) = o(\epsilon)$. Using this latter fact, and the assumption that $(\lambda, x) \mapsto K_\lambda(x)$ has lipschitz gradient, meaning it verifies [\(7\)](#), we may use the Taylor expansion of $K_\lambda(x)$ and write

$$K_{\lambda_u+\epsilon\lambda_h}(X_{u+\epsilon h}(t)) - K_{\lambda_u}(X_u(t)) = K'_{\lambda_u}(X_u(t))\delta X_{u,\epsilon,h}(t) + \nabla_\lambda K_{\lambda_u}(X_u(t))\epsilon\lambda_h + o(\epsilon),$$

and then finally compute

$$\begin{aligned} \delta J_{u,\epsilon,h} &= \int_0^T [\partial_x f(t, X_u(t), u(t)) + \partial_u f(t, X_u(t), u(t))K'_\lambda(X_u(t))] \delta X_{u,\epsilon,h}(t) dt \\ &\quad + \int_0^T \partial_u f(t, X_u(t), u(t)) \epsilon [v_h(t) + \nabla_\lambda K_\lambda(X_u(t))\lambda_h] dt \\ &\quad + \partial_x g(T, X_u(T)) \delta X_{u,\epsilon,h}(T) + o(\epsilon). \end{aligned}$$

We now need an estimate for $\delta X_{u,\epsilon,h}$ that is linear in ϵh . In order to obtain this latter, we first notice that $\delta X_{u,\epsilon,h}$ solves the following SDE

$$\begin{aligned} d\delta X_{u,\epsilon,h}(t) &= [\partial_x b(t, X_u(t), u(t)) + \partial_u b(t, X_u(t), u(t))K'_\lambda(X_u(t))] \delta X_{u,\epsilon,h}(t) dt \\ &\quad + \partial_u b(t, X_u(t), u(t)) \epsilon [v_h(t) + \nabla_\lambda K_\lambda(X_u(t))\lambda_h] dt + o(\epsilon) dt \\ &\quad + [\partial_x \sigma(t, X_u(t)) \delta X_{u,\epsilon,h}(t) + o(\epsilon)] dW_t. \end{aligned}$$

Lemma [Appendix A.2](#) yields that the solution of the linearized SDE is indeed a good approximation. The gradients of f and g are bounded, we can therefore replace $\delta X_{u,\epsilon,h}$ by the new estimate $\epsilon y_{u,h}$ and obtain

$$\begin{aligned} \delta J_{u,\epsilon,h} &= \int_0^T [\partial_x f(t, X_u(t), u(t)) + \partial_u f(t, X_u(t), u(t))K'_\lambda(X_u(t))] \epsilon y_{u,h}(t) dt \\ &\quad + \int_0^T \partial_u f(t, X_u(t), u(t)) \epsilon [v_h(t) + \nabla_\lambda K_\lambda(X_u(t))\lambda_h] dt \\ &\quad + \partial_x g(T, X_u(T)) \epsilon y_{u,h}(T) + o(\epsilon). \end{aligned}$$

By injecting the explicit formula of $y_{u,h}$ and inverting the integral signs, we finally obtain

$$\begin{aligned} \delta J_{u,\epsilon,h} &= \epsilon \int_0^T I_u(s) \phi_u^{-1}(s) \partial_u b(s, X_u(s), u(s)) [v_h(s) + \nabla_\lambda K_\lambda(X_u(s))\lambda_h] ds \\ &\quad + o(\epsilon), \end{aligned} \tag{A.2}$$

where I_u is the term

$$I_u(s) \triangleq \partial_x g(T, X_u(T)) \phi_u(T) + \int_s^T [\partial_x f(t, X_u(t), u(t)) + \partial_u f(t, X_u(t), u(t)) K'_\lambda(X_u(t))] \phi_u(t) dt.$$

The function $u = (v, \lambda) \mapsto J$ is therefore Fréchet differentiable and we can compute the Riesz representation of its gradient, denoted $\nabla J(u)$. We recall that $\nabla J(u) = (\nabla_v J(u), \nabla_\lambda J(u))$ is the random variable such that

$$\lim_{\epsilon \rightarrow 0} \frac{\delta J_{u, \epsilon, h}(\omega)}{\epsilon} = \langle \nabla J(u)(\omega), h \rangle_{\mathcal{U}} \quad \forall h \in \mathcal{U}.$$

From the expression of $\delta J_{u, \epsilon, h}$ in (A.2) we can directly identify the gradient as

$$\begin{aligned} \nabla_v J(v, \lambda)(t) &= I_u(t) \phi_u^{-1}(t) \partial_u b(t, X_u(t), u(t)) \\ \nabla_\lambda J(v, \lambda) &= \int_0^T I_u(s) \phi_u^{-1}(s) \partial_u b(s, X_u(s), u(s)) \nabla_\lambda K_\lambda(X_u(s)) ds. \end{aligned}$$

Appendix A.2. Proof of Proposition 3.2

Let ρ be a finite coherent risk measure, \mathcal{U} and $\partial\rho(0)$ are both convex subsets of a Hilbert space. As shown in 3.1, $u \mapsto J(u)$ is also Fréchet differentiable. Additionally, the risk function is Hadamard differentiable (H-differentiable) as stated in Theorem 2.2, and thus $u \mapsto \rho(J(u))$ is H-differentiable. From [45, Chapter 3.1], we obtain the following optimality condition for H-differentiable functions:

Proposition Appendix A.3. *Let X be a convex subset of a Banach space and $f : X \mapsto \mathbf{R}$ an H-differentiable function. If f has a minimum on X reached at x^* , it is necessary for x^* to verify:*

$$f'(x^*) \cdot (x - x^*) \geq 0, \quad \forall x \in X.$$

Let u^* be a solution to Problem (8). Therefore, the previous proposition yields:

$$-\rho'(J(u^*)) \cdot D_u J(u^*) \cdot (u - u^*) \leq 0, \quad \forall u \in \mathcal{U}.$$

In particular, from the formula of the differential of ρ given in Theorem 2.2 and from the chain rule for directional differentials, we obtain that:

$$\inf_{\zeta \in \partial\rho(J(u^*))} \mathbb{E}[-\zeta \cdot D_u J(u^*) \cdot (u - u^*)] \leq 0, \quad \forall u \in \mathcal{U}.$$

By using the Riesz representation of $D_u J(u^*)$ and using equation (10) we get

$$\inf_{\zeta \in \partial \rho(J(u^*))} \langle \mathbb{E}[-\zeta \nabla J(u^*)], u - u^* \rangle_{\mathcal{U}} \leq 0, \quad \forall u \in \mathcal{U}. \quad (\text{A.3})$$

By taking the supremum of (A.3) over $u \in \mathcal{U}$, we infer that

$$\sup_{u \in \mathcal{U}} \inf_{\zeta \in \partial \rho(J(u^*))} \langle \mathbb{E}[-\zeta \nabla J(u^*)], u - u^* \rangle_{\mathcal{U}} \leq 0. \quad (\text{A.4})$$

Now, we need to invert the sup with the inf in the above expression, and for this, we use the following version of Sion min-max theorem [46]:

Theorem Appendix A.4. *Let X and Y be two convex subsets of Hausdorff topological spaces with X being compact, and consider a continuous map f such that*

- $x \mapsto f(x, y)$ is convex in X for all $y \in Y$,
- $y \mapsto f(x, y)$ is concave in Y for all $x \in X$.

Then, it holds that

$$\sup_{y \in Y} \inf_{x \in X} f(x, y) = \inf_{x \in X} \sup_{y \in Y} f(x, y).$$

In our setting, we select $Y = \mathcal{U}$, which is convex, and $X = \partial \rho(J(u^*))$, which is convex, weakly-* compact and therefore weakly compact since it is a subset of a Hilbert space. Also, by definition, the mapping

$$(\zeta, u) \mapsto \langle \mathbb{E}[-\zeta \nabla J(u^*)], u - u^* \rangle_{\mathcal{U}} \quad (\text{A.5})$$

is weakly continuous and concave in u , as well as weakly continuous and convex in ζ . Therefore, Theorem Appendix A.4 can be applied to equation (A.4), yielding

$$\inf_{\zeta \in \partial \rho(J(u^*))} \sup_{u \in \mathcal{U}} \langle \mathbb{E}[-\zeta \nabla J(u^*)], u - u^* \rangle_{\mathcal{U}} \leq 0. \quad (\text{A.6})$$

Since the supremum of a family of lower semi-continuous functions is also lower semi-continuous, the mapping

$$\zeta \mapsto \sup_{u \in \mathcal{U}} \langle \mathbb{E}[-\zeta \nabla J(u^*)], u - u^* \rangle_{\mathcal{U}}$$

is lower semi-continuous, and therefore it attains a minimum which we denote by $\zeta^* \in \partial \rho(J(u^*))$. Thanks to this latter remark, from (A.6) we infer that

$$\langle \mathbb{E}[\zeta^* \nabla J(u^*)], u - u^* \rangle_{\mathcal{U}} \geq 0, \quad \forall u \in \mathcal{U},$$

and the conclusion follows from the fact that \mathcal{U} is a vector space.

Appendix A.3. Proof of Lemma 5.4 (compactness of the output of a control affine system)

We recall we are under the hypothesis of Setting B.

Proof of the compactness of $u \mapsto X_u$: Let u_n be a sequence of controls in \mathcal{U} weakly converging to \bar{u} . Let $p \geq 2$, we prove below the more general convergence

$$\|X_{u_n} - X_{\bar{u}}\|_{C_{\mathcal{F}}^p}^p \rightarrow 0$$

The theorem of existence of strong solutions of SDEs in [24] ensures us that, since our initial condition is deterministic, $X_u \in C_{\mathcal{F}}^p$. Below, we denote by $C > 0$ a constant that will be implicitly overloaded throughout the proof. By definition, for every $n \in \mathbb{N}$ we may compute

$$\begin{aligned} \|X_{u_n} - X_{\bar{u}}\|_{C_{\mathcal{F}}^p}^p &\leq C \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r (b_0(s, X_{u_n}(s)) - b_0(s, X_{\bar{u}}(s))) ds \right\|^p \right] \\ &+ C \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r (\sigma(s, X_{u_n}(s)) - \sigma(s, X_{\bar{u}}(s))) dW_s \right\|^p \right] \\ &+ C \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r \sum_{i=1}^m [(v_n)^{(i)}(s) - (\bar{v})^{(i)}(s)] b_i(s, X_{\bar{u}}(s)) ds \right\|^p \right] \\ &+ C \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r \sum_{i=1}^m [K_{\lambda_n}(X_{u_n}(s))^{(i)} - K_{\bar{\lambda}}(X_{\bar{u}}(s))^{(i)}] b_i(s, X_{\bar{u}}(s)) ds \right\|^p \right] \\ &+ C \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r \sum_{i=1}^m u_n^{(i)}(s) (b_i(s, X_{u_n}(s)) - b_i(s, X_{\bar{u}}(s))) ds \right\|^p \right]. \end{aligned}$$

Let us denote

$$h_n(t, \omega) \triangleq \int_0^t \sum_{i=1}^m [(v_n)^{(i)}(s) - (\bar{v})^{(i)}(s)] b_i(s, X_{\bar{u}}(s, \omega)) ds.$$

We now appropriately bound the different integral terms. Hölder's inequality yields

$$\begin{aligned} \sup_{0 \leq r \leq t} \left\| \int_0^r (b_0(s, X_{u_n}(s)) - b_0(s, X_{\bar{u}}(s))) ds \right\|^p &\leq \\ &\leq C \int_0^t \|b_0(s, X_{u_n}(s)) - b_0(s, X_{\bar{u}}(s))\|^p ds \\ &\leq C \int_0^t \|X_{u_n}(s) - X_{\bar{u}}(s)\|^p ds \leq C \int_0^t \sup_{0 \leq s' \leq s} \|X_{u_n}(s') - X_{\bar{u}}(s')\|^p ds. \end{aligned}$$

Using that the functions b_i are bounded and the assumptions on the feedback (6) paired with Hölder's inequality, we have that

$$\begin{aligned}
& \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r \sum_{i=1}^m [K_{\lambda_n}(X_{u_n}(s))^{(i)} - K_{\bar{\lambda}}(X_{\bar{u}}(s))^{(i)}] b_i(s, X_{\bar{u}}(s)) ds \right\|^p \right] \\
& \leq C \mathbb{E} \left[\int_0^t \sum_{i=1}^m \|K_{\lambda_n}(X_{u_n}(s))^{(i)} - K_{\lambda_n}(X_{\bar{u}}(s))^{(i)}\|^p \|b_i(s, X_{\bar{u}}(s))\|^p ds \right] \\
& + C \mathbb{E} \left[\int_0^t \sum_{i=1}^m \|K_{\lambda_n}(X_{\bar{u}}(s))^{(i)} - K_{\bar{\lambda}}(X_{\bar{u}}(s))^{(i)}\|^p \|b_i(s, X_{\bar{u}}(s))\|^p ds \right] \\
& \leq C \mathbb{E} \left[\int_0^t (1 + \|\lambda_n\|^p) \|X_{u_n}(s) - X_{\bar{u}}(s)\|^p ds \right] \\
& + C \mathbb{E} \left[\int_0^t (1 + \|X_{\bar{u}}(s)\|^p) \|\lambda_n - \bar{\lambda}\|^p ds \right].
\end{aligned}$$

Since λ_n converges, it is bounded. Using the fact that $X_{\bar{u}}$ is in $C_{\mathcal{F}}^p$, by taking the sup we finally have that

$$\begin{aligned}
& \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r \sum_{i=1}^m [K_{\lambda_n}(X_{u_n}(s))^{(i)} - K_{\bar{\lambda}}(X_{\bar{u}}(s))^{(i)}] b_i(s, X_{\bar{u}}(s)) ds \right\|^p \right] \\
& \leq C \mathbb{E} \left[\int_0^t \sup_{0 \leq s' \leq s} \|X_{u_n}(s') - X_{\bar{u}}(s')\|^p ds \right] + C \|\lambda_n - \bar{\lambda}\|^p.
\end{aligned}$$

For the stochastic integral, thanks to Burkholder-Davis-Gundy inequality (1) combined with Hölder we may compute

$$\begin{aligned}
& \mathbb{E} \left[\sup_{0 \leq r \leq t} \left\| \int_0^r (\sigma(s, X_{u_n}(s)) - \sigma(s, X_{\bar{u}}(s))) dW_s \right\|^p \right] \leq \\
& \leq C \mathbb{E} \left[\int_0^t \|\sigma(s, X_{u_n}(s)) - \sigma(s, X_{\bar{u}}(s))\|^2 ds \right]^{\frac{p}{2}} \\
& \leq C \mathbb{E} \left[\int_0^t \|X_{u_n}(s) - X_{\bar{u}}(s)\|^p ds \right] \leq C \int_0^t \mathbb{E} \left[\sup_{0 \leq s' \leq s} \|X_{u_n}(s') - X_{\bar{u}}(s')\|^p \right] ds.
\end{aligned}$$

Gathering all the previous bounds yields

$$\begin{aligned}
\mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_{u_n}(s) - X_{\bar{u}}(s)\|^p \right] & \leq C \int_0^t \mathbb{E} \left[\sup_{0 \leq s' \leq s} \|X_{u_n}(s') - X_{\bar{u}}(s')\|^p \right] ds \\
& + C \|\lambda_n - \bar{\lambda}\|^p + C \mathbb{E} \left[\sup_{0 \leq s \leq t} \|h_n(s, \omega)\|^p \right],
\end{aligned}$$

and an application of Grönwall inequality allows us to infer that

$$\mathbb{E} \left[\sup_{0 \leq s \leq t} \|X_{u_n}(s) - X_{\bar{u}}(s)\|^p \right] \leq C \left[\mathbb{E} \left[\sup_{0 \leq s \leq t} \|h_n(s, \omega)\|^p \right] + \|\lambda_n - \bar{\lambda}\|^p \right].$$

To conclude, we now need to prove that $\mathbb{E} \left[\sup_{0 \leq s \leq t} \|h_n(s, \omega)\|^p \right]$ tends to 0. The following lemma gives us the convergence of $\sup_{0 \leq s \leq t} \|h_n(s, \omega)\|^p$ for ω fixed in Ω (see [25] for proof).

Lemma Appendix A.5. *Let $a, b \in \mathbf{R}$ and let E be a normed vector space. For all $n \in \mathbf{N}$, let $f_n : [a, b] \rightarrow E$ be uniformly α -Hölder, that is*

$$\exists \alpha, K > 0, \quad \forall n \in \mathbf{N}, \quad \forall x, y \in [a, b] \quad \|f_n(x) - f_n(y)\| \leq K \|x - y\|^\alpha \quad (\text{A.7})$$

If the sequence f_n converges simply to an application f , then it converges uniformly.

For $\omega \in \Omega$, $h_n(s, \omega)$ converges to 0 by weak convergence of u_n . Let us recall that the functions b_i are bounded by assumption. We can now show that the sequence $h_n(\cdot, \omega)$ is uniformly Hölder thanks to Cauchy Schwartz and to the fact that v_n is uniformly bounded in $L^2(0, T)$:

$$\begin{aligned} |h_n(t_2, \omega) - h_n(t_1, \omega)| &= \left| \int_{t_1}^{t_2} \sum_{i=1}^m [(v_n)^{(i)}(s) - (\bar{v})^{(i)}(s)] b_i(s, X_{\bar{u}}(s, \omega)) ds \right| \\ &\leq C \sum_{i=1}^m \left[\left(\int_{t_1}^{t_2} (|(v_n)^{(i)}(s)|^2 + |(\bar{v})^{(i)}(s)|^2) ds \right)^{\frac{1}{2}} \left(\int_{t_1}^{t_2} \|b_i\|_\infty^2 ds \right)^{\frac{1}{2}} \right] \\ &\leq C |t_2 - t_1|^{\frac{1}{2}} \end{aligned}$$

Using Lemma [Appendix A.5](#), we have that sequence of random variables $g_n(\omega) := \sup_{0 \leq s \leq t} \|h_n(s, \omega)\|^p$ converges pointwise to 0. The random variable g_n can be uniformly bounded by the deterministic value

$$|g_n(\omega)| \leq \left(Tm \int_0^T \sum_{i=1}^m (M_u + \|(\bar{v})^{(i)}(s)\|^2) \|b_i\|_\infty^2 ds \right)^{\frac{p}{2}}$$

with M_u a uniform bound of the sequence of controls. We therefore have by dominated convergence that $\mathbb{E}[g_n]$ tends to 0, which in turn gives

$$\|X_{u_n} - X_{\bar{u}}\|_{C_{\mathcal{F}}^p}^p \rightarrow 0.$$

Proof of the compactity of $u \mapsto \nabla_u J_X(X_u)$:

For more clarity, let us consider a cost where $g = 0$. Let u be a control in \mathcal{U} , we can write the gradient as

$$\nabla_u J_X(X_u)^{(i)}(t) = \left[\int_t^T \partial_x f(s, X_u(s)) \phi_u(s) ds \right] \phi_u^{-1}(t) \partial_x b_i(t, X_u(t))$$

with ϕ_u satisfying for $t \in [0, T]$ the linear system

$$\begin{cases} d\phi_u(t) &= [\partial_x b_0(t, X_u(t)) + \sum_{i=1}^m u_i(t) \partial_x b_i(t, X_u(t))] \phi_u(t) dt \\ &+ [\sum_{i=1}^m b_i(t, X_u(t))] K'_{\lambda_u}(X_u(t)) \phi_u(t) dt \\ &+ \partial_x \sigma(t, X_u(t)) \phi_u(t) dW_t, \\ \phi(0) &= Id_N. \end{cases} \quad (\text{A.8})$$

Since functions $b_i, \partial_x b_i, \partial_x \sigma_i$ and K'_λ are all uniformly bounded, we can prove in a similar way as previously the compactity of $u \mapsto \phi_u$ in $C_{\mathcal{F}}^p$. We can also prove the compactity of $u \mapsto \phi_u^{-1}$ as it satisfies a similar SDE.

We then conclude using the Cauchy-Schwartz inequality twice on the expectation of the four terms and obtaining

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq t \leq T} \|\nabla_u J_X(X_u)^{(i)}(t) - \nabla_u J_X(X_{u_n})^{(i)}(t)\|^2 \right] &\leq \\ C \left(\int_0^T \mathbb{E} \left[\|\partial_x f(s, X_{\bar{u}}(s)) - \partial_x f(s, X_{u_n}(s))\|^8 \right] ds \right) &\|\phi_{\bar{u}}\|_{C_{\mathcal{F}}^8} \|\phi_{\bar{u}}^{-1}\|_{C_{\mathcal{F}}^8} \|\partial_x b_i\|_{\infty}^8 \\ + C \|\partial_x f(\cdot, X_{u_n}(\cdot))\|_{C_{\mathcal{F}}^8} \|\phi_{\bar{u}} - \phi_{u_n}\|_{C_{\mathcal{F}}^8} \|\phi_{\bar{u}}^{-1}\|_{C_{\mathcal{F}}^8} \|\partial_x b_i\|_{\infty}^8 & \\ + C \|\partial_x f(\cdot, X_{u_n}(\cdot))\|_{C_{\mathcal{F}}^8} \|\phi_{u_n}\|_{C_{\mathcal{F}}^8} \|\phi_{\bar{u}}^{-1} - \phi_{u_n}^{-1}\|_{C_{\mathcal{F}}^8} \|\partial_x b_i\|_{\infty}^8 & \\ + C \|\partial_x f(\cdot, X_{u_n}(\cdot))\|_{C_{\mathcal{F}}^8} \|\phi_{u_n}\|_{C_{\mathcal{F}}^8} \|\phi_{u_n}^{-1}\|_{C_{\mathcal{F}}^8} \|\partial_x b_i(\cdot, X_{\bar{u}}(\cdot)) - \partial_x b_i(\cdot, X_{u_n}(\cdot))\|_{C_{\mathcal{F}}^8} & \end{aligned}$$

The lipschitzianity of $\partial_x f$ ensures the convergence of the first term to 0 and implies that

$$\|\partial_x f(s, x)\| \leq C(1 + \|x\|)$$

which in turn yields

$$\|\partial_x f(\cdot, X_{u_n}(\cdot))\|_{C_{\mathcal{F}}^8} \leq C(1 + \|X_{u_n}\|_{C_{\mathcal{F}}^8}),$$

which is bounded thanks to the convergence of the sequence X_{u_n} . Finally, we conclude with the lipschitzianity of $\partial_x b_i$ which implies that the last term tends to 0. We therefore have

$$\|\nabla_u(J_X(X_{u_n})) - \nabla_u(J_X(X_{\bar{u}}))\|_{C_{\mathcal{F}}^2}^2 \rightarrow 0.$$

References

- [1] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*. Society for Industrial and Applied Mathematics, Jan. 2009.
- [2] M. P. Chapman, R. Bonalli, K. M. Smith, I. Yang, M. Pavone, and C. J. Tomlin, “Risk-Sensitive Safety Analysis Using Conditional Value-at-Risk,” *IEEE Transactions on Automatic Control*, vol. 67, pp. 6521–6536, Dec. 2022. Conference Name: IEEE Transactions on Automatic Control.
- [3] Y. Wang and M. P. Chapman, “Risk-averse autonomous systems: A brief history and recent developments from the perspective of optimal control,” *Artificial Intelligence*, vol. 311, p. 103743, Oct. 2022.
- [4] M. V. Kothare, V. Balakrishnan, and M. Morari, “Robust constrained model predictive control using linear matrix inequalities,” *Automatica*, vol. 32, pp. 1361–1379, Oct. 1996.
- [5] Y.-L. Chow and M. Pavone, “A framework for time-consistent, risk-averse model predictive control: Theory and algorithms,” in *2014 American Control Conference*, (Portland, OR, USA), pp. 4204–4211, IEEE, June 2014.
- [6] P. Sopasakis, M. Schuurmans, and P. Patrinos, “Risk-averse risk-constrained optimal control,” in *2019 18th European Control Conference (ECC)*, (Naples, Italy), pp. 375–380, IEEE, June 2019.
- [7] A. Ruszczyński and J. Yao, “A Risk-Averse Analog of the Hamilton–Jacobi–Bellman Equation,” in *2015 Proceedings of the Conference on Control and its Applications* (C. Bonnet, B. Pasik-Duncan, H. Ozbay, and Q. Zhang, eds.), pp. 462–468, Philadelphia, PA: Society for Industrial and Applied Mathematics, July 2015.
- [8] R. Bonalli and B. Bonnet, “First-Order Pontryagin Maximum Principle for Risk-Averse Stochastic Optimal Control Problems,” May 2023. arXiv:2204.03036 [math], accepted in SICON, to appear shortly.
- [9] A. Shapiro, “On a time consistency concept in risk averse multistage stochastic programming,” *Operations Research Letters*, vol. 37, pp. 143–147, May 2009.

- [10] A. Ruszczyński, “Risk-averse dynamic programming for Markov decision processes,” *Mathematical Programming*, vol. 125, pp. 235–261, Oct. 2010.
- [11] O. Çavuş and A. Ruszczyński, “Risk-Averse Control of Undiscounted Transient Markov Models,” *SIAM Journal on Control and Optimization*, vol. 52, pp. 3935–3966, Jan. 2014. Publisher: Society for Industrial and Applied Mathematics.
- [12] N. Bäuerle and J. Ott, “Markov Decision Processes with Average-Value-at-Risk criteria,” *Mathematical Methods of Operations Research*, vol. 74, pp. 361–379, Dec. 2011.
- [13] C. W. Miller and I. Yang, “Optimal Control of Conditional Value-at-Risk in Continuous Time,” *SIAM Journal on Control and Optimization*, vol. 55, pp. 856–884, Jan. 2017.
- [14] L. Pfeiffer, “Risk-averse Merton’s Portfolio Problem,” *IFAC-PapersOnLine*, vol. 49, no. 8, pp. 266–271, 2016.
- [15] P. Laurent, “Optimality conditions in variational form for non-linear constrained stochastic control problems,” *Mathematical Control and Related Fields*, vol. 10, pp. 493–526, Aug. 2020. Publisher: Mathematical Control and Related Fields.
- [16] R. Archibald, “A Stochastic Gradient Descent Approach for Stochastic Optimal Control,” *East Asian Journal on Applied Mathematics*, vol. 10, pp. 635–658, June 2020.
- [17] K. Konolige, “A gradient method for realtime robot control,” in *Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000) (Cat. No.00CH37113)*, vol. 1, (Takamatsu, Japan), pp. 639–646, IEEE, 2000.
- [18] S. Garreis, T. M. Surowiec, and M. Ulbrich, “An Interior-Point Approach for Solving Risk-Averse PDE-Constrained Optimization Problems with Coherent Risk Measures,” *SIAM Journal on Optimization*, vol. 31, pp. 1–29, Jan. 2021. Publisher: Society for Industrial and Applied Mathematics.

- [19] A. Nedić and A. Ozdaglar, “Subgradient Methods for Saddle-Point Problems,” *Journal of Optimization Theory and Applications*, vol. 142, pp. 205–228, July 2009.
- [20] A. Mokhtari, A. Ozdaglar, and S. Pattathil, “A Unified Analysis of Extra-gradient and Optimistic Gradient Methods for Saddle Point Problems: Proximal Point Approach,” in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pp. 1497–1507, PMLR, June 2020. ISSN: 2640-3498.
- [21] T. Lin, C. Jin, and M. I. Jordan, “On gradient descent ascent for nonconvex-concave minimax problems,” in *Proceedings of the 37th International Conference on Machine Learning*, vol. 119 of *ICML’20*, pp. 6083–6093, JMLR.org, July 2020.
- [22] Z. Chen, Y. Zhou, T. Xu, and Y. Liang, “Proximal Gradient Descent-Ascent: Variable Convergence under KL Geometry,” Jan. 2021.
- [23] J.-F. Le Gall, *Brownian Motion, Martingales, and Stochastic Calculus*, vol. 274 of *Graduate Texts in Mathematics*. Cham: Springer International Publishing, 2016.
- [24] J. Yong and X. Y. Zhou, *Stochastic controls: Hamiltonian systems and HJB equations*. No. 43 in *Applications of mathematics*, New York: Springer, 1999.
- [25] E. Trélat, “Contrôle optimal : théorie & applications,” *HAL*, vol. 2005, no. 0, 2005.
- [26] R. Bonalli, B. Hérisse, and E. Trélat, “Solving optimal control problems for delayed control-affine systems with quadratic cost by numerical continuation,” in *2017 American Control Conference (ACC)*, pp. 649–654, May 2017. ISSN: 2378-5861.
- [27] R. Bonalli, B. Herisse, and E. Trelat, “Optimal Control of Endoatmospheric Launch Vehicle Systems: Geometric and Computational Issues,” *IEEE Transactions on Automatic Control*, vol. 65, pp. 2418–2433, June 2020.

- [28] R. Bonalli, *Optimal Control of Aerospace Systems with Control-State Constraints and Delays*. phdthesis, Sorbonne Université, UPMC University of Paris 6, Laboratoire Jacques-Louis Lions ; ONERA – The French Aerospace Lab, Département TIS, Unité NGPA ; Inria Paris, Equipe CAGE, July 2018.
- [29] R. Bonalli, T. Lew, and M. Pavone, “Sequential convex programming for non-linear stochastic optimal control,” *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 28, p. 64, 2022.
- [30] M. Benzi, G. H. Golub, and J. Liesen, “Numerical solution of saddle point problems,” *Acta Numerica*, vol. 14, pp. 1–137, May 2005.
- [31] D. P. Kouri and T. M. Surowiec, “Existence and Optimality Conditions for Risk-Averse PDE-Constrained Optimization,” *SIAM/ASA Journal on Uncertainty Quantification*, vol. 6, pp. 787–815, Jan. 2018. Publisher: Society for Industrial and Applied Mathematics.
- [32] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” Jan. 2017. arXiv:1412.6980 [cs].
- [33] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*. Berlin, Heidelberg: Springer, 1992.
- [34] P. Bernhard and A. Rapaport, “On a theorem of Danskin with an application to a theorem of Von Neumann-Sion,” *Nonlinear Analysis: Theory, Methods & Applications*, vol. 24, pp. 1163–1181, Apr. 1995.
- [35] T. Smith, J. Barlow, M. Bualat, T. Fong, C. Provencher, H. Sanchez, and E. Smith, “Astrobee: A New Platform for Free-Flying Robotics on the International Space Station,” (Beijing), June 2016. NTRS Author Affiliations: NASA Ames Research Center, SGT, Inc., Universities Space Research Association NTRS Report/Patent Number: ARC-E-DAA-TN31584 NTRS Document ID: 20160007769 NTRS Research Center: Ames Research Center (ARC).
- [36] L. Siurna, G. Bragg, and G. Reusing, “Transient solutions to a stochastic model of ventilation,” *Building and Environment*, vol. 24, pp. 265–277, Jan. 1989.

- [37] A. Thavlov and H. Madsen, “A non-linear stochastic model for an office building with air infiltration,” *International Journal of Sustainable Energy Planning and Management*, pp. 55–66 Pages, Nov. 2015. Artwork Size: 55-66 Pages Publisher: International Journal of Sustainable Energy Planning and Management.
- [38] J.-F. Bonnans, J. C. Gilbert, C. Lemarechal, and C. A. Sagastizábal, *Numerical Optimization: Theoretical and Practical Aspects*. Springer Science & Business Media, Oct. 2006. Google-Books-ID: 88X5vN3ybvGC.
- [39] E. Gilbert and D. Johnson, “Distance functions and their application to robot path planning in the presence of obstacles,” *IEEE Journal on Robotics and Automation*, vol. 1, pp. 21–30, Mar. 1985. Conference Name: IEEE Journal on Robotics and Automation.
- [40] Y. Rasekhipour, A. Khajepour, S.-K. Chen, and B. Litkouhi, “A Potential Field-Based Model Predictive Path-Planning Controller for Autonomous Road Vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, pp. 1255–1267, May 2017. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [41] R. Bonalli, A. Cauligi, A. Bylard, and M. Pavone, “GuSTO: Guaranteed Sequential Trajectory optimization via Sequential Convex Programming,” in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 6741–6747, May 2019. ISSN: 2577-087X.
- [42] J. Chessari, R. Kawai, Y. Shinozaki, and T. Yamada, “Numerical methods for backward stochastic differential equations: A survey,” *Probability Surveys*, vol. 20, pp. 486–567, Jan. 2023. Publisher: Institute of Mathematical Statistics and Bernoulli Society.
- [43] C. Della Santina, C. Duriez, and D. Rus, “Model-Based Control of Soft Robots: A Survey of the State of the Art and Open Challenges,” *IEEE Control Systems Magazine*, vol. 43, pp. 30–65, June 2023. Conference Name: IEEE Control Systems Magazine.
- [44] J. Yong, “Linear forward-backward stochastic differential equations with random coefficients,” *Probability Theory and Related Fields*, vol. 135, pp. 53–83, May 2006.

- [45] J. F. Bonnans and A. Shapiro, *Perturbation Analysis of Optimization Problems*. New York, NY: Springer New York, 2000.
- [46] S. Simons, “Minimax Theorems and Their Proofs,” in *Minimax and Applications* (P. Pardalos, R. Horst, D.-Z. Du, and P. M. Pardalos, eds.), vol. 4, pp. 1–23, Boston, MA: Springer US, 1995. Series Title: Nonconvex Optimization and Its Applications.