

OPTIMAL CONTROL OF CONDITIONAL VALUE-AT-RISK IN
CONTINUOUS TIME*

CHRISTOPHER W. MILLER† AND INSOON YANG‡

Abstract. We consider continuous-time stochastic optimal control problems featuring conditional value-at-risk (CVaR) in the objective. The major difficulty in these problems arises from time inconsistency, which prevents us from directly using dynamic programming. To resolve this challenge, we convert to an equivalent bilevel optimization problem in which the inner optimization problem is standard stochastic control. Furthermore, we provide conditions under which the outer objective function is convex and differentiable. We compute the outer objective's value via a Hamilton–Jacobi–Bellman equation and its gradient via the viscosity solution of a linear parabolic equation, which allows us to perform gradient descent. The significance of this result is that we provide an efficient dynamic-programming-based algorithm for optimal control of CVaR without lifting the state space. To broaden the applicability of the proposed algorithm, we propose convergent approximation schemes in cases where our key assumptions do not hold and characterize relevant suboptimality bounds. In addition, we extend our method to a more general class of risk metrics, which includes mean variance and median deviation. We also demonstrate a concrete application to portfolio optimization under CVaR constraints. Our results contribute an efficient framework for solving time-inconsistent CVaR-based sequential optimization.

Key words. conditional value-at-risk, risk measures, time inconsistency, stochastic optimal control, Hamilton–Jacobi–Bellman equations, viscosity solutions, dynamic programming

AMS subject classifications. 49L20, 49L25, 90C39, 93E20, 91G80

DOI. 10.1137/16M1058492

1. Introduction. Conditional value-at-risk (CVaR) has received significant attention over the past two decades as a tool for managing risk. CVaR measures the expected value conditional upon being within some percentage of the worst-case loss scenarios. More specifically, the CVaR of a random variable X , whose distribution has no probability atoms, is defined as

$$\text{CVaR}_\alpha(X) := \mathbb{E}[X \mid X \geq \text{VaR}_\alpha(X)], \quad \alpha \in (0, 1),$$

where the value-at-risk (VaR) of X (with the cumulative distribution function F_X) is given by

$$\text{VaR}_\alpha(X) := \inf\{x \in \mathbb{R} \mid F_X(x) \geq \alpha\}.$$

In other words, VaR is equal to the $(1 - \alpha)$ worst-case quantile of a loss distribution, while CVaR equals the conditional expectation of the loss within that quantile. When the distribution has a probability atom, the definition of CVaR should be further refined (see [54]). Note that both functions penalize only when “bad events” occur.

While both VaR and CVaR are risk measures, only CVaR is *coherent* in the sense of Artzner et al. [2]. In addition, CVaR takes into account the possibility of tail

*Received by the editors January 26, 2016; accepted for publication (in revised form) January 20, 2017; published electronically March 21, 2017.

<http://www.siam.org/journals/sicon/55-2/M105849.html>

Funding: The first author was supported in part by NSF GRFP under grant DGE 1106400. The second author was supported in part by NSF CRII: CPS under CNS 1657100.

†Department of Mathematics, University of California, Berkeley, San Francisco, CA 94108 (miller@math.berkeley.edu).

‡Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 (insoonya@usc.edu).

events where losses exceed VaR. In fact, one common criticism of VaR stems from its incapability of distinguishing situations beyond VaR [54]. Due to the superior mathematical properties and practical implications, CVaR has gained popularity in risk management [59]. In particular, *static* or single-stage optimization with CVaR functions can be efficiently performed via convex and linear programming methods [53, 40]. With the advances in optimization algorithms for CVaR, this risk measure has shown to be useful in various finance and engineering applications.

Dynamic or sequential optimization of CVaR is often of interest when decisions can be made at multiple stages. In such an optimal control setting, we can optimize a control action at a certain time based on the information from observations up to that time. This dynamic control approach enjoys an effective usage of information gathered in the process of making decisions under uncertainty. The need for efficient optimal control tools with CVaR is also motivated by emerging dynamic risk management problems in engineering and finance (e.g., [49, 63]).

The major challenge in optimal control involving CVaR arises from its *time inconsistency* [3]. For example, an optimal strategy for tomorrow constructed today is no longer optimal when considered tomorrow because CVaR is not a time-consistent risk measure. Mathematically, this time inconsistency prevents us from directly applying dynamic programming, in contrast with problems involving Markov risk measures [55, 14, 57] or risk-sensitive criteria [30, 23]. To overcome this difficulty, several methods have been proposed. A state-space lifting approach for dynamic programming with a discrete-time and discrete-state Markov decision process (MDP) setting is first proposed in [6].¹ Another lifting method and relevant algorithms are developed in [48, 16], relying on a so-called CVaR decomposition theorem [48]. This approach uses a dual representation of CVaR and hence requires optimization over a space of probability densities when solving an associated Bellman equation. This optimization problem can be effectively solved in discrete-time and finite discrete-state MDPs. However, it becomes computationally intractable in (uncountable) continuous-state optimal control problems as the space of densities is infinite dimensional. In [27], a different approach is developed for risk-aware discrete-time finite-state MDPs, which is based on occupation measures. Due to the nonconvexity of the resulting infinite-dimensional optimization problem, this method uses a successive linear approximation procedure.

In this paper, we propose a new method to solve continuous-time and continuous-space optimal control problems involving CVaR. By using a so-called *extremal* representation of CVaR originally proposed in [53], we reformulate the optimal control problem as a bilevel optimization problem in which the outer optimization problem is convex and the inner optimization problem is standard stochastic optimal control. To avoid lifting state space, we develop a gradient-descent-based method to solve the outer optimization problem. Specifically, we prove the differentiability of the outer objective function and provide a probabilistic interpretation of its gradient under a certain semiconcave approximation.

To develop a computationally efficient and stable gradient-descent-based method, it is essential to be able to compute the objective's value and gradient. The outer objective value can be computed by solving the inner problem: we demonstrate a dynamic programming or, equivalently, a Hamilton–Jacobi–Bellman (HJB) method

¹This lifting approach was recently generalized to semi-Markov decision processes in [29] and is simplified in the application of manufacturing systems [1]. Note that our method is different in the sense that we avoid extending the state space by solving an associated outer optimization problem via gradient descent and therefore reduce computational complexity in general.

to solve the inner problem. More importantly, we show that the gradient of the outer objective function can be obtained as the viscosity solution of an associated linear parabolic equation, which we call the *gradient* partial differential equation (PDE), under certain conditions. These two PDE characterizations complete the proposed gradient-descent-based method to solve the bilevel optimization problem.

We use two important assumptions while proving the main results. One is the semiconcavity of the outer objective function, while the other is the uniform parabolicity of the HJB equation. We construct convergent approximation schemes which relax each of these assumptions when needed. For the theoretical (and also practical) implication of the approximations, we provide bounds on the gap between the optimal outer objective value and that of the perturbed problem.

In the final section of the paper, we demonstrate a practical implementation of our methodology in an optimal investment problem subject to CVaR constraints. To our knowledge, this is the first solution of a dynamic portfolio optimization problem subject to tail-risk constraints in continuous time. The closest comparisons to our results are given by approximate equilibrium solutions [20], mean-field control approaches [47], or in mean-variance frameworks [45].

The rest of this paper is organized as follows: In section 2, we introduce optimal control problems involving a class of risk metrics including CVaR. We present the main results of this paper in section 3, which are used to construct a gradient-descent-based method to solve the reformulated bilevel optimization problem. Relevant assumptions imposed in the main results are explicitly relaxed in section 4 using convergent approximation schemes. Last, we demonstrate the performance of the proposed method through an example of mean-CVaR portfolio selection in section 5.

2. Problem setup.

2.1. Controlled process. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting a standard d -dimensional Brownian motion W with an associated filtration $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ satisfying the usual conditions. Let \mathbb{A} be a compact and finite-dimensional set of controls. Define the set \mathcal{A} of admissible control strategies as the collection of all \mathcal{F}_t -progressively measurable processes which are valued in \mathbb{A} almost surely.

The control $A \in \mathcal{A}$ affects a system state of interest through the following SDE:

$$(2.1) \quad \begin{aligned} dX_t^A &= \mu(X_t^A, A_t) dt + \sigma(X_t^A, A_t) dW_t, \\ X_0^A &= \mathbf{x}_0 \in \mathbb{R}^n. \end{aligned}$$

We assume that $\mu : \mathbb{R}^n \times \mathbb{A} \rightarrow \mathbb{R}^n$ and $\sigma : \mathbb{R}^n \times \mathbb{A} \rightarrow \mathbb{R}^{n \times d}$ are continuous functions such that, for some $K > 0$,

$$\begin{aligned} \|\mu(x, a) - \mu(x', a)\| + \|\sigma(x, a) - \sigma(x', a)\| &\leq K\|x - x'\|, \\ \|\mu(x, a)\| + \|\sigma(x, a)\| &\leq K(1 + \|x\| + \|a\|) \end{aligned}$$

for all $x, x' \in \mathbb{R}^n$ and for all $a \in \mathbb{A}$. Under these conditions, for each control $A \in \mathcal{A}$, there exists a unique strong solution, X^A , of the SDE (2.1).

2.2. Optimal control with a class of risk measures. The main goal of this paper is to provide an efficient algorithm for solving the following stochastic optimal

control problem with a nonstandard objective:²

$$(2.2) \quad \inf_{A \in \mathcal{A}} \rho(g(X_T^A)),$$

where $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is a given convex cost function and $\rho : L^2(\Omega) \rightarrow \mathbb{R}$ is of a class of risk measures defined below.

DEFINITION 2.1. *A function $\rho : L^2(\Omega) \rightarrow \mathbb{R}$ is said to be an extremal risk measure if there exists a convex function $f : \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}$ such that*

$$\rho(\xi) = \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi, y)].$$

For simplicity, we assume that f is convex with at most quadratic growth and that g is convex and Lipschitz continuous. However, we note that these growth conditions can easily be relaxed on a case-by-case basis by noting that $X_T^A \in L^p(\Omega)$ for all $p < \infty$.

The primary motivation of this definition is the following extremal formula involving CVaR of a random variable ξ , whose distribution has no probability atom³ [53]:

$$(2.3) \quad \text{CVaR}_\alpha[\xi] = \inf_{y \in \mathbb{R}} \mathbb{E} \left[y + \frac{1}{1-\alpha} (\xi - y)^+ \right], \quad \alpha \in (0, 1).$$

The intuition behind this equality is that the optimal y is equal to VaR at probability α . Then, CVaR is equal to VaR plus the expected losses exceeding VaR divided by the probability of these losses occurring, $1 - \alpha$.

We show that several additional risk metrics of interest in application can be written in this form. Some of these are not (coherent) risk measures in the sense of Artzner et al. [2]. However, we justify our nomenclature of “extremal risk measure” by providing simple conditions on f under which ρ is a coherent risk measure in Appendix A. We emphasize that for ease of exposition, we will generally refer to ρ as an extremal risk measure in this paper even when these conditions are not satisfied.

Example 1. The variance of $\xi \in L^2(\Omega)$ can be expressed as

$$(2.4) \quad \text{Var}[\xi] = \inf_{y \in \mathbb{R}} \mathbb{E}[(\xi - y)^2].$$

Note that, with $f(\xi, y) := (\xi - y)^2$, variance is of the form of Definition 2.1. Although variance is not a risk measure, our method can handle control problems with variance-related criteria including mean-variance optimal control: $\mathbb{E}[\xi] + \lambda \text{Var}[\xi] = \inf_{y \in \mathbb{R}} \mathbb{E}[\xi + \lambda(\xi - y)^2]$.

Similarly, the median absolute deviation (MAD) of $\xi \in L^2(\Omega)$ can be expressed as

$$(2.5) \quad \text{MAD}[\xi] = \mathbb{E}[|\xi - \text{Med}[\xi]|] = \inf_{y \in \mathbb{R}} \mathbb{E}[|\xi - y|].$$

Again, with $f(\xi, y) := |\xi - y|$, MAD is of the form of Definition 2.1. MAD is a deviation risk measure (e.g., [44]).

²Note that this formulation includes the following problems with running and terminal costs: $\inf_{A \in \mathcal{A}} \rho(\int_0^T r(X_{1,t}^A, A_t) dt + q(X_{1,T}^A))$. For such problems, we introduce a new state, $X_{2,t} := \int_0^t r(X_{1,s}^A, A_s) ds$, and rewrite the optimization problem as $\inf_{A \in \mathcal{A}} \rho(X_{2,T}^A + q(X_{1,T}^A)) =: \inf_{A \in \mathcal{A}} \rho(g(X_T^A))$, where $X_t := (X_{1,t}, X_{2,t})$ and $g(x) = x_2 + q(x_1)$.

³In our problem setting, if $g(X_T^A)$ has an atom, (2.3) is interpreted as CVaR⁻, called the lower CVaR [53].

Example 2. Our method can handle control problems with criteria including various combinations of (2.3), (2.4), and (2.5). For example,

- mean-CVaR:

$$\mathbb{E}[\xi] + \lambda \text{CVaR}_\alpha[\xi] = \inf_{y \in \mathbb{R}} \mathbb{E} \left[\xi + \lambda \left(y + \frac{1}{1-\alpha} (\xi - y)^+ \right) \right];$$

- variance-CVaR tradeoff:

$$\text{Var}[\xi] + \lambda \text{CVaR}_\alpha[\xi] = \inf_{y \in \mathbb{R}^2} \mathbb{E} \left[(\xi - y_1)^2 + \lambda \left(y_2 + \frac{1}{1-\alpha} (\xi - y_2)^+ \right) \right];$$

- CVaR_α - CVaR_β tradeoff:

$$\text{CVaR}_\alpha[\xi] + \lambda \text{CVaR}_\beta[\xi] = \inf_{y \in \mathbb{R}^2} \mathbb{E} \left[\begin{pmatrix} 1 \\ \lambda \end{pmatrix} \cdot \begin{pmatrix} y_1 + (1-\alpha)^{-1} (\xi - y_1)^+ \\ y_2 + (1-\beta)^{-1} (\xi - y_2)^+ \end{pmatrix} \right].$$

We will provide a more concrete example in section 5.

One important feature of the optimization problem (2.2) is that, in general, it is a time-inconsistent problem. This makes it impossible to directly apply dynamic programming because there is not an analogue of the law of iterated expectations for non-Markov risk measures. To tackle this issue, one stream of research efforts focuses on developing *dynamic risk measures* with which sequential optimization or optimal control can be performed in a time-consistent manner (e.g., [51, 15, 25, 36, 3, 56, 55]). Another class of research activities acknowledges time inconsistency as an inherent property and proposes two different solution concepts. The first approach is to reinterpret the problem as a game against one's future self. This approach is used, for example, in [8] and leads to a PDE system. The second method is to rewrite the original problem in a form where we can apply dynamic programming in an indirect way. This approach has been used to reduce the problem to dynamic programming in a higher-dimensional state space or to a sequence of iterated standard control problems in [6, 41, 31, 7, 48, 63]. The rationale in this paper is similar in spirit to this so-called *indirect* dynamic programming method. However, one key advantage of the proposed approach is that the structure of the extremal risk measures allows us to perform optimization over an extra variable where the objective function can be evaluated by dynamic programming involving no additional state variables. We construct optimal time-inconsistent controls by solving an equivalent bilevel optimization problem, without lifting the state space.

2.3. Outline of main results. We summarize the main results of this paper by the following:

1. We demonstrate an equivalence between the time-inconsistent stochastic control problem involving extremal risk measures and the bilevel optimization problem in Proposition 3.1.
2. We provide conditions under which the outer optimization problem is convex in Theorem 3.2. Furthermore, under additional conditions, we prove differentiability and provide a probabilistic interpretation of the gradient in Theorem 3.4.
3. We demonstrate a dynamic programming approach to solving the inner problem and provide conditions under which we have a PDE characterization of the gradient of the outer problem in Theorem 3.6. This allows the use of gradient descent in solving the bilevel optimization problem.

4. Finally, we provide convergent approximations which relax two of the key assumptions imposed on the problem. In Theorems 4.5 and 4.8, we provide suboptimality bounds under each approximation scheme.

Using these results, we can solve time-inconsistent optimal control problems involving extremal risk measures via a gradient descent solution of a bilevel optimization problem. As an example of how the approximation schemes and gradient descent are utilized, consider the explicit problem

$$\inf_{A \in \mathcal{A}} \text{CVaR}_\alpha [X_T^A].$$

We can solve this problem via the following procedure:

- Step 1: Recall $f^{\text{CVaR}}(\xi, y) := y + (1 - \alpha)^{-1}(\xi - y)^+$. This does not satisfy our uniform semiconcavity assumption and hence we first apply inf-convolution to have

$$f_\epsilon(\xi, y) := \inf_{z \in \mathbb{R}} \left[f^{\text{CVaR}}(\xi, z) + \frac{(y - z)^2}{2\epsilon} \right],$$

which is uniformly semiconcave in y (ϵ -semiconcave approximation).

- Step 2: We note that we can reformulate the perturbed optimal control problem involving f_ϵ as the following bilevel optimization problem:

$$\inf_{A \in \mathcal{A}} \inf_{y \in \mathbb{R}} \mathbb{E} [f_\epsilon(X_T^A, y)] = \inf_{y \in \mathbb{R}} \left[V_\epsilon(y) := \inf_{A \in \mathcal{A}} \mathbb{E} [f_\epsilon(X_T^A, y)] \right],$$

where the inner problem is a standard stochastic control problem.

- Step 3: We solve the outer problem using a gradient descent algorithm in y . At each iteration with $y \in \mathbb{R}^m$, we solve the inner problem via dynamic programming by solving an HJB equation to evaluate $V(y)$. In addition, we compute its gradient value $DV(y)$ by solving a linear parabolic equation, which we call the *gradient PDE*.
- Step 4: When the gradient descent algorithm converges, we obtain a minimizer y_ϵ^* of the outer problem and its associated minimizer A_ϵ^* of the inner problem. We quantify a suboptimality bound of this solution using Theorem 4.5. The gap between this solution and an optimal solution tends to zero as $\epsilon \rightarrow 0$.

In this high-level description of the proposed method, we notice its several advantages. First of all, the proposed gradient descent approach does not require us to lift the state-space in the case of terminal cost problems. This is a considerable advantage over existing methods which regard y as another state variable because the computational complexity of dynamic programming increases exponentially with the dimension of the state space. Second, the gradient PDE provides a systematic numerical approximation of $DV(y)$ at any $y \in \mathbb{R}^m$. Therefore, it avoids using derivative-free optimization algorithms which are not convergent in general or automatic differentiation tools such as finite differencing which can be inaccurate when the objective function contains (numerical) noise [43]. Last, the analytical studies to relax the assumptions used to prove the main results broaden the applicability of the proposed method. In particular, we show that the approximate solution obtained from an ϵ -semiconcave approximation converges to the true optimal solution as ϵ tends to 0.

We summarize the complete series of approximations by stating the following general version of our results. We emphasize that significantly stronger results are provided throughout the paper when allowing various additional assumptions. We

break the statement of the following theorem into three pieces. These correspond to (1) rewriting the optimal control as a bilevel problem, (2) introducing an approximation scheme where proximal supergradients of the top-level problem can be computed in terms of a system of PDEs, and (3) a convergence result for the approximation scheme.

THEOREM 2.2. *The following statements hold:*

(1) *Let $f : \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}$ both be convex and Lipschitz continuous. Define an extremal risk measure $\rho : L^2(\Omega) \rightarrow \mathbb{R}$ as*

$$\rho(\xi) := \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi, y)].$$

Furthermore, define

$$V(y) := \inf_{A \in \mathcal{A}} \mathbb{E}[f(g(X_T^A), y)].$$

Then

$$p^* := \inf_{A \in \mathcal{A}} \rho(g(X_T^A)) = \inf_{y \in \mathbb{R}^m} V(y).$$

(2) *For any $\epsilon, \eta > 0$, define*

$$f_\epsilon(x, y) := \inf_{z \in \mathbb{R}^m} \left[f(x, z) + \frac{\|y - z\|^2}{2\epsilon} \right].$$

Let \hat{W} be an independent n -dimensional Brownian motion, and let $\hat{\mathcal{F}}$ be the filtration generated by the joint process (W, \hat{W}) . Let $\hat{\mathcal{A}}$ represent the collection of all $\hat{\mathcal{F}}$ -adapted processes valued in \mathbb{A} . For each control $A \in \hat{\mathcal{A}}$, define $\hat{X}^{A, \eta}$ as the unique strong solution of the perturbed SDE

$$d\hat{X}_t^{A, \eta} = \mu(\hat{X}_t^{A, \eta}, A_t) dt + \sigma(\hat{X}_t^{A, \eta}, A_t) dW_t + \eta d\hat{W}_t.$$

Define

$$V_{\epsilon, \eta}(y) := \inf_{A \in \hat{\mathcal{A}}} \mathbb{E}[f_\epsilon(g(\hat{X}_T^{A, \eta}), y)]$$

and denote by $\partial V_{\epsilon, \eta}(y)$ the collection of proximal supergradients of $V_{\epsilon, \eta}$ at y .

Then we have the following results:

(a) *For any $y \in \mathbb{R}^m$, if we define*

$$DV_{\epsilon, \eta}(y) := \mathbb{E}[D_y f_\epsilon(g(\hat{X}_T^{A, \eta}), y)],$$

then $DV_{\epsilon, \eta}(y) \in \partial V_{\epsilon, \eta}(y)$.

(b) *For any $y \in \mathbb{R}^m$, $V_{\epsilon, \eta}(y)$ and $DV_{\epsilon, \eta}(y)$ may be computed in terms of the solutions of two PDEs. The value function is related to an HJB equation, while the proximal supergradient is related to a formal linearization of the HJB equation.*

(3) *Let y^* be a minimizer of V and $y_{\epsilon, \eta}^*$ be a minimizer of $V_{\epsilon, \eta}$. Then*

$$V(y_{\epsilon, \eta}^*) \leq V(y^*) + C(\epsilon + \eta),$$

where $C > 0$ is a constant depending only on the Lipschitz constants of f , g , μ , and σ . That is, $V(y_{\epsilon, \eta}^) \rightarrow p^*$ linearly as $\epsilon, \eta \rightarrow 0$.*

To summarize, this theorem provides a complete convergent algorithm for solving the optimal control of a $\rho(g(X_T^A))$. First, we show how to rewrite this as a bilevel minimization problem, the lower-level of which is a standard optimal stochastic control problem. There is, in general, no special structure in the upper-level problem (e.g., convexity), so we demonstrate a very general semiconcave approximation, the objective function and proximal supergradient of which can be computed in terms of solutions of a PDE system. Last, we demonstrate how a solution of the approximate problem converges to an optimal solution of the original problem.

We will present a proof of Theorem 2.2 in section 4.4 after developing the relevant machinery and approximations in section 3 and sections 4.1–4.3.

3. Main results: Gradient descent and viscosity solutions.

3.1. Equivalent bilevel optimization. Recall that our goal is to solve the generalized stochastic control problem

$$(3.1) \quad \inf_{A \in \mathcal{A}} \rho(g(X_T^A)),$$

where ρ is a fixed extremal risk measure and $g(X_T^A)$ is the state-dependent cost when the control A is executed.

In general, (3.1) is a time-inconsistent nonlinear stochastic optimal control problem to which we cannot apply dynamic programming. However, we show how to use the structure of ρ as an extremal risk measure to convert this into an equivalent bilevel optimization problem.

PROPOSITION 3.1 (bilevel optimization). *We can write the problem of dynamic optimization over an extremal risk measure as*

$$(3.2) \quad \inf_{A \in \mathcal{A}} \rho(g(X_T^A)) = \inf_{y \in \mathbb{R}^m} V(y),$$

where V is defined via a standard stochastic optimal control problem of the form

$$(3.3) \quad V(y) := \inf_{A \in \mathcal{A}} \mathbb{E}[f(g(X_T^A), y)].$$

The proof is straightforward by the definition of extremal risk measures.

Remark 1. The value $V(y)$ depends on the initial value x_0 of X^A . For simplicity, however, we suppress the dependency and implicitly assume that we fix the initial value as x_0 for the rest of this paper.

At this point, we have converted the time-inconsistent stochastic control problem (3.1) to an equivalent bilevel optimization problem (3.2) involving a standard stochastic control problem (3.3). For convenience, we call the right-hand sides of (3.2) and (3.3) the *outer optimization* problem and the *inner optimization* problem, respectively.

3.2. A note on assumptions. We record three main assumptions. These are used to obtain various properties of the bilevel optimization problem along the way, and we discuss, where applicable, how they may be relaxed in section 4.

Assumption 1 (uniform semiconcavity). The function $y \mapsto f(x, y)$ is uniformly semiconcave for all $x \in \mathbb{R}$. That is, there exists $M > 0$ such that

$$f(x, y + \xi) \leq f(x, y) + \xi \cdot D_y f(x, y) + \frac{1}{2} M \|\xi\|^2$$

for all $(x, y) \in \mathbb{R} \times \mathbb{R}^m$ and for all $\xi \in \mathbb{R}^m$.

It is important to note that the convexity of f together with Assumption 1 guarantees that the map $y \mapsto f(x, y)$ is continuously differentiable for each $x \in \mathbb{R}^m$. This regularity will ultimately carry over to various smoothness results of the outer optimization problem in our bilevel optimization problem, i.e., the right-hand side of (3.2).

Remark 2. Relaxing Assumption 1 is particularly important from the application perspective. Note, for example, that CVaR and MAD do not satisfy this uniform semiconcavity assumption. For the clarity of presentation, however, we begin with this seemingly restrictive assumption but later relax it by using the inf-convolution operator in section 4.1.

Assumption 2 (uniform parabolicity). There exists an $\epsilon > 0$ such that

$$\sigma(x, a)\sigma(x, a)^\top - \epsilon I_n \text{ is positive semidefinite}$$

for all $(x, a) \in \mathbb{R}^n \times \mathbb{A}$.

The main use of Assumption 2 is to construct optimal controls to the stochastic control problem (3.3). In particular, this assumption guarantees that the viscosity solution to the HJB equation is smooth, [24, 39]. In section 4.2, we relax this constraint by adding additional independent Brownian motions to the dynamics of X^A .

Assumption 3 (convexity). The control set \mathbb{A} is convex and the map $(A, y) \mapsto f(g(X_T^A), y)$ is jointly convex, almost surely.

The primary use of Assumption 3 is to guarantee convexity of the outer optimization problem in our bilevel optimization formulation. In particular, this allows us to implement a gradient descent algorithm with guaranteed convergence to a global minimum. When Assumption 3 does not hold, we can still compute so-called proximal supergradients and run a descent algorithm which converges to a local minimizer.

Example 3. We emphasize the following three sufficient conditions, each of which guarantees Assumption 3 holds (e.g., [11]):

- $A \mapsto g(X_T^A)$ is affine,
- $A \mapsto g(X_T^A)$ is convex⁴ and $x \mapsto f(x, y)$ is nondecreasing convex for each $y \in \mathbb{R}^m$, or
- $A \mapsto g(X_T^A)$ is concave and $x \mapsto f(x, y)$ is nonincreasing convex for each $y \in \mathbb{R}^m$.

Recall that $f^{\text{CVaR}}(\xi, y) := y + (1 - \alpha)^{-1}(\xi - y)^+$, $x \mapsto f^{\text{CVaR}}(x, y)$ is nondecreasing convex. Therefore, if $A \mapsto g(X_T^A)$ is convex, Assumption 3 holds.

We note that Assumption 3 is quite strong. However, it proves verifiable in some practical applications in engineering and finance such as risk-aware demand response, electric vehicle charging control, inventory control, and portfolio management (e.g., [50, 62]). As a concrete example, consider the mean-CVaR portfolio optimization problem considered in section 5. For clarity of exposition, we retain this as a main assumption but emphasize the analogous results which hold even when this assumption is broken.

3.3. Analytical properties of the outer objective function V . In this section we investigate some analytical properties of the outer objective function V . We begin by showing the convexity of V . Then, we present a semiconcavity estimate of

⁴We acknowledge that checking the convexity of $A \mapsto X_T^A$ beyond the criteria proposed in [4, 17] is often a nontrivial task. It is a topic of future research.

V at points where there exists an optimal control. Furthermore, we use this estimate and convexity to show the differentiability of V at such points and to provide a probabilistic representation of the gradient.

THEOREM 3.2 (convexity). *Suppose that Assumption 3 holds. Then, the outer objective function V is convex.*

The convexity of V motivates us to use a (sub)gradient-descent-type algorithm for solving the outer optimization problem. Under the semiconcavity assumption (Assumption 1), we justify gradient descent approaches by proving the differentiability of V .

To that end, we first record a semiconcavity estimate. It is important to note that this estimate does not depend on the convexity of V or Assumption 3. This result connects the semiconcavity of f to that of V .

PROPOSITION 3.3. *Suppose that Assumption 1 holds. For any fixed $y \in \mathbb{R}^m$, we assume there exists $A \in \mathcal{A}$ such that⁵*

$$V(y) = \mathbb{E} [f(g(X_T^A), y)].$$

Then, we have

$$V(y + \xi) \leq V(y) + \xi \cdot \mathbb{E} [D_y f(g(X_T^A), y)] + \frac{1}{2} M \|\xi\|^2$$

for all $\xi \in \mathbb{R}^m$.

Proof. Fix $y \in \mathbb{R}^m$. By Assumption 1, there exists $M > 0$ such that

$$(3.4) \quad f(x, y + \xi) \leq f(x, y) + \xi \cdot D_y f(x, y) + \frac{1}{2} M \|\xi\|^2$$

for all $x \in \mathbb{R}$ and $\xi \in \mathbb{R}^m$. Let $A \in \mathcal{A}$ be a control such that

$$V(y) = \mathbb{E} [f(g(X_T^A), y)].$$

Note that such a control depends on the choice of y . If we apply inequality (3.4) pointwise and take expectations, we see

$$V(y + \xi) \leq \mathbb{E} [f(g(X_T^A), y + \xi)] \leq \mathbb{E} [f(g(X_T^A), y)] + \xi \cdot \mathbb{E} [D_y f(g(X_T^A), y)] + \frac{1}{2} M \|\xi\|^2$$

for all $\xi \in \mathbb{R}^m$. Then, the result holds. \square

In the following result, we combine the convexity and semiconcavity estimates of V to show that V is in fact differentiable. In particular, we provide a probabilistic representation of the gradient at each point.

THEOREM 3.4 (differentiability). *Suppose that Assumptions 1 and 3 hold. For any fixed $y \in \mathbb{R}^m$, we assume that there exists $A \in \mathcal{A}$ such that*

$$V(y) = \mathbb{E} [f(g(X_T^A), y)].$$

Then, V is differentiable at y and its gradient can be computed as

$$(3.5) \quad DV(y) = \mathbb{E} [D_y f(g(X_T^A), y)].$$

⁵We assume the existence of an optimal control in this proposition. However, we will show that this assumption is valid in the next subsection by constructing an optimal control from an associated HJB equation under Assumption 2.

We emphasize that, while the value of the gradient seems the natural consequence of an envelope theorem, we do not know that V is differentiable a priori. This theorem proves differentiability as a direct consequence of convexity and semiconcavity, while identifying the gradient in the process.

Proof. The function V is convex by Theorem 3.2. Because of the quadratic growth of f , the Lipschitz assumptions on g , μ , and σ , and the compactness of \mathbb{A} , it is a standard result that $V(y)$ is finite for every $y \in \mathbb{R}^m$ (see [22]). Therefore, its subdifferential is nonempty at each point [52]. Fix $y \in \mathbb{R}^m$ and suppose $A \in \mathcal{A}$ is a control such that

$$V(y) = \mathbb{E} [f(g(X_T^A), y)].$$

Let $z \in \partial V(y)$ be an arbitrary subgradient of V at y , i.e.,

$$V(y + \xi) \geq V(y) + \xi \cdot z \quad \forall \xi \in \mathbb{R}^m.$$

By Proposition 3.3, we also have the inequality

$$V(y + \xi) \leq V(y) + \xi \cdot \mathbb{E} [D_y f(g(X_T^A), y)] + \frac{1}{2} M \|\xi\|^2 \quad \forall \xi \in \mathbb{R}^m.$$

Putting these together, we obtain

$$\xi \cdot (z - \mathbb{E} [D_y f(g(X_T^A), y)]) \leq \frac{1}{2} M \|\xi\|^2 \quad \forall \xi \in \mathbb{R}^m.$$

Choosing $\xi := M^{-1}(z - \mathbb{E} [D_y f(g(X_T^A), y)])$, we have

$$M \|\xi\|^2 = \xi \cdot (z - \mathbb{E} [D_y f(g(X_T^A), y)]) \leq \frac{1}{2} M \|\xi\|^2,$$

which is a contradiction unless $z = \mathbb{E} [D_y f(g(X_T^A), y)]$.

Because $\partial V(y)$ is single valued, we conclude that V is differentiable at y and also that

$$DV(y) = \mathbb{E} [D_y f(g(X_T^A), y)]$$

as desired. \square

Due to Theorem 3.4, we can solve the outer optimization problem using a gradient descent algorithm given that the function value $V(y)$ and its gradient $DV(y)$ are provided. We notice that $V(y)$ can be computed by dynamic programming. It is worth mentioning that the overall problem is still time inconsistent but our bilevel decomposition allows us to solve the inner optimization problem using dynamic programming. We investigate the inner optimization problem and provide a constructive approach to solve the overall problem in the following subsection.

3.4. PDE characterization of V and DV . Recall that $V(y)$ is the optimal value of the inner optimization problem (3.3) given the initial value x_0 of the state. Note that this problem can be solved by dynamic programming; we first compute $V(y)$ in terms of the viscosity solution of an associated HJB equation. In the process we construct an optimal control, which guarantees V is differentiable by Theorem 3.4. Furthermore, we can compute $DV(y)$ in terms of the viscosity solution of a linear parabolic equation.

PROPOSITION 3.5. *Given $y \in \mathbb{R}^m$, let $v : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the viscosity solution of the HJB equation*

$$(3.6) \quad \begin{aligned} v_t + \inf_{a \in \mathbb{A}} \left[\frac{1}{2} \text{tr}(\sigma(x, a) \sigma(x, a)^\top D_x^2 v) + \mu(x, a) \cdot D_x v \right] &= 0 && \text{in } [0, T) \times \mathbb{R}^n, \\ v(T, x) &= f(g(x), y) && \text{on } \{t = T\} \times \mathbb{R}^n. \end{aligned}$$

Then, we have

$$V(y) = v(0, \mathbf{x}_0),$$

where \mathbf{x}_0 is the initial value of the SDE (2.1).

This is a standard result. The important point is that, under Assumption 2, (3.6) is uniformly parabolic and concave in $D_x^2 v$. Then we know from regularity results for HJB equations that the value function v is twice differentiable in space (see [21, 12, 26]). This allows us to compute $DV(y)$ by solving a linear equation.

THEOREM 3.6 (gradient PDE). *Suppose that Assumptions 2–3 hold. Given $y \in \mathbb{R}^m$, let v be the viscosity solution of (3.6).*

(1) *An optimal control, $A_t^* := a^*(t, X_t^{A^*})$, exists, where $a^* : [0, T) \times \mathbb{R}^n \rightarrow \mathbb{A}$ satisfies*

$$a^*(t, x) \in \arg \min_{a \in \mathbb{A}} \left[\frac{1}{2} \text{tr}(\sigma(x, a) \sigma(x, a)^\top D_x^2 v) + \mu(x, a) \cdot D_x v \right] \\ \forall (t, x) \in [0, T) \times \mathbb{R}^n.$$

(2) *Let $w : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ be defined as*

$$(3.7) \quad w(t, x) := \mathbb{E}_t \left[D_y f \left(g(X_T^{A^*}), y \right) \middle| X_t^{A^*} = x \right].$$

Then, $w := (w^{(1)}, \dots, w^{(m)})$ is a viscosity solution of the decoupled system of linear equations

$$(3.8) \quad \begin{aligned} w_t^{(k)} + \frac{1}{2} \text{tr} \left(\sigma(x, a^*(t, x)) \sigma(x, a^*(t, x))^\top D_x^2 w^{(k)} \right) \\ + \mu(x, a^*(t, x)) \cdot D_x w^{(k)} &= 0 && \text{in } [0, T) \times \mathbb{R}^n, \\ w^{(k)}(T, x) &= [D_y f(g(x), y)]_k && \text{on } \{t = T\} \times \mathbb{R}^n \end{aligned}$$

for $k = 1, \dots, m$. Furthermore, we have

$$DV(y) = w(0, \mathbf{x}_0),$$

where \mathbf{x}_0 is the initial value of the SDE (2.1).

Before providing a rigorous proof, we note the intuition behind the result. It is clear that w defined in (3.7) satisfies

$$(3.9) \quad w(t, x) = \mathbb{E}_t \left[w \left(t + h, X_{t+h}^{A^*} \right) \middle| X_t^{A^*} = x \right]$$

for all $(t, x) \in [0, T) \times \mathbb{R}^n$ and $h \in (0, T - t)$. This follows from the law of iterated expectations. If w were smooth, we could apply Ito's lemma for small $h > 0$ and conclude that w satisfies (3.8). However, the coefficients of (3.8) are not necessarily continuous and hence the solution w is potentially discontinuous. That is, we must consider this within the framework of discontinuous viscosity solutions. For more on this topic, see [32, 61].

Proof.

- (1) The existence of such an optimal control follows from a standard argument due to the uniform parabolicity of the HJB (3.6) (e.g., [24]).
- (2) Without loss of generality, we can assume $m = 1$ and suppress the superscript of w because the PDE system (3.8) is decoupled. Our goal is then to show that the function w defined in (3.7) is a discontinuous viscosity solution of (3.8). The key property of w is listed in (3.9). For convenience, we introduce the notation

$$\begin{aligned} (\mathcal{L}^{a^*} \phi)(t, x) := & \phi_t(t, x) + \operatorname{tr}(\sigma(x, a^*(t, x)) \sigma(x, a^*(t, x))^\top D_x^2 \phi(t, x)) \\ & + \mu(x, a^*(t, x)) \cdot D_x \phi(t, x), \end{aligned}$$

where ϕ is an arbitrary smooth function.

As a reminder, we define the lower- and upper-semicontinuous envelopes of a locally bounded function ψ as

$$\psi_*(x) := \liminf_{y \rightarrow x} \psi(y) \text{ and } \psi^*(x) := \limsup_{y \rightarrow x} \psi(y),$$

respectively.

Fix $(\bar{t}, \bar{x}) \in [0, T) \times \mathbb{R}$ and let $\phi : [0, T) \times \mathbb{R} \rightarrow \mathbb{R}$ be a smooth function satisfying

$$0 = (w_* - \phi)(\bar{t}, \bar{x}) = \min_{[0, T) \times \mathbb{R}^n} (w_* - \phi).$$

That is, ϕ touches the lower-semicontinuous envelope of w from below at (\bar{t}, \bar{x}) . Our goal is to show that

$$(\mathcal{L}^{a^*} \phi)_*(\bar{t}, \bar{x}) \leq 0.$$

Towards that end, let $\{(t_k, x_k)\}_{k=0}^\infty$ be a sequence such that, as $k \rightarrow \infty$,

$$(t_k, x_k) \rightarrow (\bar{t}, \bar{x}) \text{ and } w(t_k, x_k) \rightarrow w_*(\bar{t}, \bar{x}).$$

Since ϕ is smooth, $\delta_k := w(t_k, x_k) - \phi(\bar{t}, \bar{x}) \rightarrow 0$. We also define

$$h_k := \sqrt{\delta_k} 1_{\{\delta_k \neq 0\}} + k^{-1} 1_{\{\delta_k = 0\}}$$

and take k large enough that $h_k \in (0, T - t_k)$. Then, we have

$$\begin{aligned} w(t_k, x_k) &= \mathbb{E}_{t_k} \left[w(t_k + h_k, X_{t_k+h_k}^{A^*}) \middle| X_{t_k}^{A^*} = x_k \right] \\ &\geq \mathbb{E}_{t_k} \left[\phi(t_k + h_k, X_{t_k+h_k}^{A^*}) \middle| X_{t_k}^{A^*} = x_k \right] \\ &= \phi(t_k, x_k) + \mathbb{E}_{t_k} \left[\int_{t_k}^{t_k+h_k} (\mathcal{L}^{a^*} \phi)(s, X_s^{A^*}) ds \middle| X_{t_k}^{A^*} = x_k \right]. \end{aligned}$$

Rearranging this, we conclude

$$\frac{\delta_k}{h_k} \geq \mathbb{E}_{t_k} \left[\frac{1}{h_k} \int_{t_k}^{t_k+h_k} (\mathcal{L}^{a^*} \phi)(s, X_s^{A^*}) ds \middle| X_{t_k}^{A^*} = x_k \right].$$

Sending $k \rightarrow \infty$, the left-hand side converges to zero, while the right-hand side dominates the lower-semicontinuous envelope of $\mathcal{L}^{a^*}\phi$ almost surely. Therefore, we conclude by the dominated convergence theorem that

$$(\mathcal{L}^{a^*}\phi)_*(\bar{t}, \bar{x}) \leq 0.$$

The opposite inequality proceeds exactly the same. We also notice that w satisfies the boundary condition of (3.8) pointwise. Therefore, w is a viscosity solution of the linear PDE (3.8).

Finally, we put this together with Theorem 3.4 to conclude

$$w(0, \mathbf{x}_0) = \mathbb{E} \left[D_y f \left(g(X_T^{A^*}), y \right) \right] = DV(y). \quad \square$$

We note that some care must be taken when arguing that the viscosity solution to the PDE in (3.8) is unique. It is well known that nondivergence form linear PDEs may not have a unique viscosity solution in dimensions greater than two if the diffusion coefficients are not continuous (see [58, 37, 42]). However, there has been subsequent work on finding structural conditions on the coefficients which guarantee weak uniqueness in higher dimensions. For fully nonlinear elliptic equations, [33] provides assumptions on the nonlinearity, not including continuous coefficients, which guarantee unique upper-semicontinuous viscosity solutions in the sense of Theorem 3.6. Recent results focus on bounds on the mean-oscillation of the coefficients [38, 34, 19], as well as on the set of discontinuities being small [60, 35].

We emphasize that in many practical problems the linear parabolic equation (3.8) has a unique viscosity solution. In the next proposition, we provide two conditions which can be easily checked.

PROPOSITION 3.7. *Define $F : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{S}^n \rightarrow \mathbb{R}$ as*

$$F(x, p, R) := \inf_{a \in \mathbb{A}} \left[\frac{1}{2} \text{tr}(\sigma(x, a)\sigma(x, a)^\top R) + \mu(x, a) \cdot p \right].$$

If either (i) $n \leq 2$ or (ii) both $D_p F$ and $D_R F$ exist and are continuous, then there is a unique viscosity solution of (3.8).

Proof. If $n \leq 2$, then this follows from Theorem 2.17 of [37]. If $D_p F$ and $D_R F$ exist and are continuous, then we note that the gradient PDE (3.8) can be rewritten in the form

$$w_t^{(k)} + \text{tr} \left(D_R F(x, D_x u, D_x^2 u)^\top D_x^2 w^{(k)} \right) + D_p F(x, D_x u, D_x^2 u) \cdot D_x w^{(k)} = 0.$$

Recall u has continuous second derivatives in space, so the coefficients are all continuous functions. Then, the uniqueness of viscosity solutions follows as usual. \square

Using Proposition 3.5 and Theorem 3.6, we can calculate $V(y)$ and $DV(y)$ at each $y \in \mathbb{R}^m$ by (numerically) solving the PDEs (3.6) and (3.8). Therefore, we can solve the outer optimization problem using a gradient-descent-type algorithm and find a globally optimal solution due to the convexity of V . We will not discuss numerical optimization algorithms as they are not the focus of this paper (we refer to, for example, [43] for detailed algorithms).

One natural question regarding the computation of the gradient is whether it is possible to utilize other numerical methods, most obviously Monte Carlo methods.

Theorem 3.4 can of course be used directly for this. The practitioner would obtain a numerical solution of (3.6), then generate random optimal trajectories via Monte Carlo simulations, and estimate $DV(y)$ via the sample expected value corresponding to (3.5). This however, introduces extra numerical (sampling) error in the calculation of the gradient, which is often a source of instability.

Furthermore, this gradient descent approach gains a dimensionality reduction by m , which is the dimension of y , as opposed to a dynamic programming method over the lifted state space of (x, y) . Even when $m = 1$, this gain is considerable because the computational complexity of dynamic programming increases exponentially with the dimension of the state space.⁶

4. Relaxation of assumptions. The goal of this section is to explicitly relax Assumptions 1 and 2, and to a lesser extent Assumption 3. In the case of Assumptions 1 and 2, we provide convergent approximation schemes including suboptimality bounds.

4.1. Uniformly semiconcave approximation. In this section, we consider an approximation scheme for the case that f does not satisfy the semiconcavity assumption (Assumption 1). The idea is to modify f via inf-convolution to obtain a function f_ϵ for some small $\epsilon > 0$. We show that this new function f_ϵ satisfies Assumption 1. We then prove that the resulting perturbed value function converges to the unperturbed problem as $\epsilon \rightarrow 0$. Relaxing this semiconcavity assumption is particularly important for problems with CVaR objectives.

First, we recall the definition of inf-convolution and the key properties of the resulting function.

PROPOSITION 4.1 (semiconcave approximation). *Recall that $f : \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}$ is convex. For $\epsilon > 0$, define the inf-convolution f_ϵ as*

$$(4.1) \quad f_\epsilon(x, y) := \inf_{z \in \mathbb{R}^m} \left[f(x, z) + \frac{\|y - z\|^2}{2\epsilon} \right].$$

Then, $f_\epsilon : \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}$ has the following properties:

- (1) f_ϵ is convex,
- (2) $y \mapsto f_\epsilon(x, y)$ is ϵ -semiconcave for all $x \in \mathbb{R}^n$, and
- (3) $f_\epsilon \rightarrow f$ uniformly as $\epsilon \rightarrow 0$.

For a proof of Proposition 4.1, see [32, 52]. In particular, this implies that f_ϵ satisfies Assumption 1. We also show that approximation by inf-convolution will not break the convexity required by Assumption 3.

PROPOSITION 4.2. *If f satisfies Assumption 3, then so does f_ϵ for each $\epsilon > 0$.*

Proof. See Appendix C for the proof. □

We now define the ϵ -perturbed outer objective function as

$$V_\epsilon(y) := \inf_{A \in \mathcal{A}} \mathbb{E} [f_\epsilon(g(X_T^A), y)].$$

⁶One can further alleviate the computational burden by employing advanced numerical methods such as sparse grids and multigrid techniques (e.g., [28, 10]). It is difficult to make precise statements about the complexity of the entire algorithm without discussing certain implementation choices. However, it is essentially that of a gradient-based unconstrained concave minimization algorithm, where each step requires the solution of an $(n+1)$ -dimensional HJB and m linear PDEs of dimension $(n+1)$.

For any $\epsilon > 0$, the perturbed function f_ϵ satisfies Assumption 1. Then, we can apply the results of section 3 to minimize V_ϵ . The goal of this subsection is to show convergence as $\epsilon \rightarrow 0$.

First, we record an estimate involving inf-convolution.

PROPOSITION 4.3. *Suppose that $y \mapsto f(x, y)$ is uniformly Lipschitz continuous for all $x \in \mathbb{R}$. Fix $\epsilon > 0$ and $(x, y) \in \mathbb{R} \times \mathbb{R}^m$. Then, we have*

$$f_\epsilon(x, y) \leq f(x, y) \leq f_\epsilon(x, y) + C\epsilon$$

for a constant C depending only on f .

Proof. Taking $z = y$ in the inf-convolution operator (4.1), we immediately have that

$$f_\epsilon(x, y) \leq f(x, y).$$

Next, let $z^* \in \mathbb{R}^m$ be a minimizer in f_ϵ . Using this and the uniform Lipschitz continuity of f , we have

$$f_\epsilon(x, y) = f(x, z^*) + \frac{\|y - z^*\|^2}{2\epsilon} \geq f(x, y) - L\|y - z^*\| + \frac{\|y - z^*\|^2}{2\epsilon}$$

for some $L > 0$. On the other hand, by the Cauchy–Schwarz inequality, we have

$$L\|y - z\| \leq \frac{1}{2}\epsilon L^2 + \frac{\|y - z\|^2}{2\epsilon}.$$

Putting these inequalities together, we then obtain

$$f_\epsilon(x, y) \geq f(x, y) - \frac{1}{2}\epsilon L^2. \quad \square$$

Next, we prove a bound relating the perturbed and unperturbed outer objective functions.

PROPOSITION 4.4. *Suppose that $y \mapsto f(x, y)$ is uniformly Lipschitz continuous for all $x \in \mathbb{R}$. Fix $\epsilon > 0$ and $y \in \mathbb{R}^m$. Then, we have*

$$V_\epsilon(y) \leq V(y) \leq V_\epsilon(y) + C\epsilon$$

for a constant C depending only on f .

Proof. By Proposition 4.3, there exists a constant C such that

$$f_\epsilon(g(X_T^A), y) \leq f(g(X_T^A), y) \leq f_\epsilon(g(X_T^A), y) + C\epsilon \quad \text{almost surely}$$

for all $A \in \mathcal{A}$. Taking expectation and using the fact that $A \in \mathcal{A}$ is arbitrary, we have

$$V_\epsilon(y) \leq V_0(y) \leq V_\epsilon(y) + C\epsilon.$$

Noting that $V_0(y) = V(y)$, the result follows. \square

Finally, we use this result to prove a suboptimality bound when using this approximation.

THEOREM 4.5 (convergence and suboptimality bound I). *Suppose that $y \mapsto f(x, y)$ is uniformly Lipschitz continuous for all $x \in \mathbb{R}$. Fix $\epsilon > 0$ and let $y_\epsilon^* \in \mathbb{R}^m$ be a minimizer of V_ϵ . Let $y^* \in \mathbb{R}^m$ be a minimizer of the unperturbed value function V . Then, we have*

$$|V(y^*) - V_\epsilon(y_\epsilon^*)| \leq C\epsilon$$

for a constant C depending only on f .

Proof. Suppose that instead

$$V(y^*) - V_\epsilon(y_\epsilon^*) > C\epsilon.$$

Combining this with the minimality of y^* , we have

$$V_\epsilon(y_\epsilon^*) + C\epsilon < V(y^*) \leq V(y_\epsilon^*),$$

which contradicts Proposition 4.4. A similar argument using the minimality of y_ϵ^* is contradictory to the possibility that

$$V_\epsilon(y_\epsilon^*) - V(y^*) > C\epsilon.$$

Then, the result holds. \square

Note that the constant C is equal to $\frac{1}{2}L^2$. Therefore, once we obtain an approximation y_ϵ^* , we can explicitly compute a bound, $C\epsilon$, of the gap between the approximately optimal objective value $V_\epsilon(y_\epsilon^*)$ and the optimal value $V(y^*)$. This bound is useful to determine an appropriate resolution ϵ of the semiconcave approximation. Furthermore, the bound is linear in ϵ and hence the suboptimality gap converges to zero as $\epsilon \rightarrow 0$ (see Figure 1 in section 5.3 for a numerical experiment result).

4.2. Uniformly parabolic approximation. Relaxing the uniform parabolicity assumption (Assumption 2) is particularly important when $0 \in \mathbb{A}$. To relax Assumption 2, we consider an approximation scheme for the case that the dynamics of X do not satisfy the assumption. The key idea is to perturb the dynamics of X with extra sources of risk (or uncertainty) so that uniform parabolicity holds. We then show the resulting perturbed value function converges to that of the unperturbed problem as $\eta \rightarrow 0$.

Let \hat{W} be an independent n -dimensional Brownian motion, and let $\hat{\mathcal{F}}$ be the filtration generated by the joint process (W, \hat{W}) . Let $\hat{\mathcal{A}}$ represent the collection of all $\hat{\mathcal{F}}$ -adapted processes valued in \mathbb{A} . For each control $A \in \hat{\mathcal{A}}$ and $\eta \in \mathbb{R}$, define $\hat{X}^{A,\eta}$ as the unique strong solution of the perturbed SDE

$$d\hat{X}_t^{A,\eta} = \mu\left(\hat{X}_t^{A,\eta}, A_t\right) dt + \sigma\left(\hat{X}_t^{A,\eta}, A_t\right) dW_t + \eta d\hat{W}_t.$$

We define the following η -perturbed value function:

$$V_\eta(y) := \inf_{A \in \hat{\mathcal{A}}} \mathbb{E} \left[f \left(g(\hat{X}_T^{A,\eta}), y \right) \right].$$

For any $\eta > 0$, the dynamics of the perturbed process $\hat{X}^{A,\eta}$ satisfy Assumption 2. Then, we can apply the results of section 3 to minimize V_η . The goal of this subsection is to show convergence to the unperturbed case as $\eta \rightarrow 0$.

First, we have the following estimate involving the value function under the perturbed dynamics.

PROPOSITION 4.6. *Suppose that f is Lipschitz continuous. Fix $y \in \mathbb{R}^m$. Then, for any $\eta, \eta' \geq 0$, we have the estimate*

$$|V_\eta(y) - V_{\eta'}(y)| \leq C|\eta - \eta'|$$

for some constant C which depends only on $\mu, \sigma, \mathbb{A}, f$, and g .

Proof. Fix $y \in \mathbb{R}^m$ and $\eta, \eta' \geq 0$. For any $\delta > 0$, let $A \in \hat{\mathcal{A}}$ be a δ -suboptimal control such that

$$(4.2) \quad V_\eta(y) + \delta \geq \mathbb{E} \left[f \left(g(\hat{X}_T^{A,\eta}), y \right) \right].$$

By a standard argument using the Lipschitz constants of the perturbed dynamics, we have

$$\mathbb{E} \left[\left\| \hat{X}_T^{A,\eta} - \hat{X}_T^{A,\eta'} \right\|^2 \right] \leq C(\eta - \eta')^2,$$

for some constant C which depends only on the Lipschitz constants of μ and σ and the set \mathbb{A} (for example, see [61]). Because f and g are Lipschitz continuous, we have

$$(4.3) \quad \begin{aligned} \mathbb{E} \left[\left| f \left(g(\hat{X}_T^{A,\eta'}), y \right) - f \left(g(\hat{X}_T^{A,\eta}), y \right) \right| \right] &\leq C \mathbb{E} \left[\left\| \hat{X}_T^{A,\eta'} - \hat{X}_T^{A,\eta} \right\| \right] \\ &\leq C \left(\mathbb{E} \left[\left\| \hat{X}_T^{A,\eta'} - \hat{X}_T^{A,\eta} \right\|^2 \right] \right)^{1/2} \\ &\leq C |\eta - \eta'| \end{aligned}$$

for a new constant C . Combining (4.2) and (4.3), we conclude

$$V_\eta(y) + \delta \geq \mathbb{E} \left[f \left(g(\hat{X}_T^{A,\eta'}), y \right) \right] - C |\eta - \eta'| \geq V_{\eta'}(y) - C |\eta - \eta'|.$$

Because δ was arbitrary and, by symmetry, between η and η' , the desired result follows. \square

Remark 3. In this proof we crucially assume that f is Lipschitz continuous. Therefore, if we are to apply this approximation scheme to mean-variance optimization, we would need to attempt to modify the proof.

Next, we prove the intuitive statement that enlarging the filtration \mathbb{F} to $\hat{\mathbb{F}}$ does not affect the value function when $\eta = 0$.

PROPOSITION 4.7. *For any $y \in \mathbb{R}^m$, we have $V_0(y) = V(y)$.*

Proof. Fix $y \in \mathbb{R}^m$ and consider the following two functions:

$$\begin{aligned} u(t, x) &:= \inf_{A \in \mathcal{A}} \mathbb{E}_t \left[f \left(g(X_T^A), y \right) \middle| X_t^A = x \right], \\ \hat{u}(t, x) &:= \inf_{A \in \hat{\mathcal{A}}} \mathbb{E}_t \left[f \left(g(\hat{X}_T^A), y \right) \middle| \hat{X}_t^A = x \right]. \end{aligned}$$

By Proposition 3.5, both u and \hat{u} are the viscosity solutions of the same HJB equation. Furthermore,

$$V(y) = u(0, x) \text{ and } V_0(y) = \hat{u}(0, x).$$

Then, the result follows by the uniqueness of viscosity solutions of (3.6). \square

Finally, combining these two results, we show a suboptimality bound when using this uniformly parabolic approximation.

THEOREM 4.8 (convergence and suboptimality bound II). *Suppose that f is Lipschitz continuous. Fix $\eta > 0$ and let $y_\eta^* \in \mathbb{R}^m$ be a minimizer of V_η . Let $y^* \in \mathbb{R}^m$ be a minimizer of the unperturbed value function V . Then, we have*

$$|V(y^*) - V_\eta(y_\eta^*)| \leq C\eta$$

for some constant C which depends only on $\mu, \sigma, \mathbb{A}, f$, and g .

Proof. Suppose that

$$V(y^*) - V_\eta(y_\eta^*) > C\epsilon.$$

Note that y^* is a minimizer; we have

$$V_\eta(y_\eta^*) + C\eta < V(y^*) = V_0(y^*) \leq V_0(y_\eta^*),$$

which is contradictory to Proposition 4.6. A similar argument using the minimality of y_η^* contradicts the possibility that

$$V_\eta(y_\eta^*) - V(y^*) > C\eta.$$

Then, the result follows. \square

Note that we can obtain a solution to the original problem by choosing $\eta > 0$ sufficiently small, depending on the constant C , and applying the gradient descent algorithm to minimize V_η . In practice, the constant C can be computed explicitly in terms of the size of \mathbb{A} and the Lipschitz constants on μ , σ , f , and g by following the steps in the proofs above. Consequently, using Theorem 4.8, we can explicitly compute a bound on the suboptimality gap between the obtained solution and the global minimizer y^* . As $\eta \rightarrow 0$, the suboptimality gap tends to zero.

4.3. Nonconvex case. In the previous two subsections, we showed convergent approximation schemes when either the uniform semiconcavity or the uniform parabolicity assumption did not hold. In the case of the convexity assumption (Assumption 3), we do not provide a direct approximation scheme, but instead note analogies of the results in section 3 which lead to a proximal supergradient descent algorithm.

In particular, we show that in the absence of Assumption 3, all results in section 3 hold with the caveat that V is semiconcave instead of convex.

COROLLARY 4.9. *Suppose that, for any fixed $y \in \mathbb{R}^m$, there exists $A \in \mathcal{A}$ such that*

$$V(y) = \mathbb{E} [f(g(X_T^A), y)].$$

If we denote by $\partial V(y)$ the collection of proximal supergradients of V at y , we have

$$\mathbb{E} [D_y f(g(X_T^A), y)] \in \partial V(y).$$

This follows immediately from the inequality proved in Proposition 3.3. Furthermore, because of the PDE results of section 3.4, we know that there exists an optimal control at each $y \in \mathbb{R}^m$. We make the following conclusion.

COROLLARY 4.10. *The function V is semiconcave. For each $y \in \mathbb{R}^m$, we can compute $V(y)$ and a proximal supergradient at y by solving PDEs (3.6) and (3.8), respectively.*

Consequently, the globally convergent gradient descent algorithm may be replaced with a proximal supergradient descent algorithm which converges to a local minimum. For more on tools from nonsmooth analysis and semiconcavity, see [13, 32].

4.4. Proof of Theorem 2.2. At this point, the proof of Theorem 2.2 is just a matter of piecing together the various results from the previous two sections. We provide a brief outline of where each step is used in the following.

Proof. Recall the statement of the theorem is broken into three pieces. We follow this separation in this proof.

- (1) The first section of the theorem is Proposition 3.1.
- (2) Applying the results of sections 4.1–4.3, it follows that the approximate problem satisfies Assumptions 1 and 2. Assumption 1 follows from Proposition 4.1, while Assumption 2 follows from directly from the dynamics of the approximate problem. Then the results of the second section of the theorem follow from the discussion in section 4.3.
- (3) Last, we consider the convergence result. By combining Theorems 4.5 and 4.8, we immediately conclude that

$$|V(y^*) - V_{\epsilon,\eta}(y_{\epsilon,\eta}^*)| \leq C(\epsilon + \eta)$$

for some $C > 0$ which depends only on the Lipschitz constants of f , g , μ , and σ . Similarly, applying Propositions 4.4–4.7, we conclude that

$$V_{\epsilon,\eta}(y_{\epsilon,\eta}^*) \leq V_{\epsilon,\eta}(y^*) + C(\epsilon + \eta)$$

for some (possibly different) $C > 0$, which also depends only on the Lipschitz constants above. Last, combining these two inequalities, we conclude the stated convergence result. \square

5. Example: Mean-CVaR portfolio optimization. In this section, we illustrate a practical use of our main results and approximation methods in an application to portfolio optimization under a mean-CVaR objective. Our goal is to ultimately use this methodology to compute the efficient frontier representing the trade-off between maximizing expected log-return and minimizing the CVaR of losses. We emphasize that dynamic optimization can significantly reduce CVaR while maintaining the same expected return as compared to optimal static investment strategies.

5.1. Problem formulation. Consider a market consisting of n risky assets evolving via the SDE

$$\frac{dS_t^{(i)}}{S_t^{(i)}} = \mu_i dt + \Sigma_{i,j}^{1/2} dW_t^{(j)}$$

for each $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, d\}$. Here $\mu \in \mathbb{R}^n$ is a vector of drifts and Σ is the covariance matrix of returns. We also assume that there exists a risk-free asset with drift r .

We assume that we choose a control A , which is a progressively measurable process lying in some compact set \mathbb{A} , representing the percent of the portfolio exposed to each of the n risky assets. For example, we might choose $\mathbb{A} := \{a \in \mathbb{R}^n \mid a^\top \Sigma a \leq l\}$ for a constant l corresponding to a hard portfolio risk cap.

With this setup, our portfolio value Z evolves via the SDE

$$\frac{dZ_t^A}{Z_t^A} = [r + A_t^\top (\mu - r 1)] dt + A_t^\top \Sigma^{1/2} dW_t.$$

For simplicity, we consider the log value of the portfolio, $X_t^A := \log Z_t^A$, which can be seen to solve

$$dX_t^A = \left[r + A_t^\top (\mu - r 1) - \frac{1}{2} A_t^\top \Sigma A_t \right] dt + A_t^\top \Sigma^{1/2} dW_t.$$

Without loss of generality, we assume $Z_0^A = S_0 = 1$. Then, $X_0^A = 0$ and we can interpret X_t^A as the log-returns of the portfolio up to time t .

In this section, we consider the problem of minimizing a mean-CVaR objective:

$$(5.1) \quad \inf_{A \in \mathcal{A}} [\mathbb{E} [-X_T^A] + \lambda \text{CVaR}_\alpha [-X_T^A]]$$

for fixed $\lambda > 0$ and $\alpha \in (0, 1)$. By varying λ , we can compute a subset of the efficient frontier between expected log-return and the CVaR of losses.

5.2. Solution via gradient descent. From the results of section 3, the problem (5.1) is equivalent to the bilevel optimization

$$\inf_{y \in \mathbb{R}} V(y),$$

where

$$V(y) := \inf_{A \in \mathcal{A}} \mathbb{E} [f(g(X_T^A), y)].$$

Here, we take $g(X_T^A) = -X_T^A$ and

$$f(x, y) := x + \lambda \left(y + \frac{1}{1-\alpha} (x - y)^+ \right).$$

We can check that $A \mapsto g(X_T^A)$ is convex almost surely and f is convex and non-decreasing in x . Therefore, Assumption 3 is satisfied, i.e., $(A, y) \mapsto f(g(X_T^A), y)$ is jointly convex. However, both Assumption 1 and Assumption 2 are potentially violated. Therefore, we need to apply the approximations of section 4 before we can proceed with the proposed gradient descent method.

THEOREM 5.1. *For any $\epsilon > 0$, we have the inf-convolution of $f(g(x), y)$ is*

$$f_\epsilon(g(x), y) := \begin{cases} -x + \lambda \left(y - \frac{x+y}{1-\alpha} \right) - \frac{1}{2} \left(\frac{\alpha}{1-\alpha} \lambda \right)^2 \epsilon & \text{for } x + y < -\frac{\alpha}{1-\alpha} \lambda \epsilon, \\ \frac{1}{2\epsilon} (y+x)^2 - (1+\lambda)x & \text{for } -\frac{\alpha}{1-\alpha} \lambda \epsilon \leq x + y \leq \lambda \epsilon, \\ -x + \lambda y - \frac{1}{2} \lambda^2 \epsilon & \text{for } x + y > \lambda \epsilon. \end{cases}$$

If we also consider the perturbed dynamics

$$d\hat{X}_t^{A, \epsilon} = \left[r + A_t^\top (\mu - r 1) - \frac{1}{2} A_t^\top \Sigma A_t \right] dt + A_t^\top \Sigma^{1/2} dW_t + \epsilon d\hat{W}_t$$

and the perturbed value function

$$V_\epsilon(y) := \inf_{A \in \hat{\mathcal{A}}} \mathbb{E} [f_\epsilon(g(\hat{X}_T^{A, \epsilon}), y)],$$

then Assumptions 1–3 all hold and we can apply the gradient descent method from section 3 to minimize V_ϵ . Furthermore, there exists a constant C which depends only on μ , r , Σ , λ , α , and \mathbb{A} such that

$$|V(y^*) - V_\epsilon(y_\epsilon^*)| \leq C\epsilon$$

for any $y_\epsilon^* \in \mathbb{R}^m$ which minimizes V_ϵ and $y^* \in \mathbb{R}^m$ which minimizes V .

The proof is elementary but tedious. Verifying the proposed expression for f_ϵ is an exercise in minimizing piecewise functions. We check that f_ϵ satisfies Assumption 1, the new dynamics clearly satisfy Assumption 2, and together they satisfy Assumption 3 by checking the sufficient conditions in Example 3. Finally, the error bounds result from combining the error bounds from sections 4.1 and 4.2.




FIG. 1. The relative error $|V(y^*) - V_\epsilon(y_\epsilon^*)|$ in the objective. We compare a numerical error (solid) and an explicit calculation of the theoretical error bounds (dashed) provided in section 4.

Remark 4. In this theorem we use a single parameter, ϵ , for both approximation schemes. In practice, we can modify this result to use separate parameters for each scheme. This could be useful for separating out the numerical error due to each approximation.

5.3. Numerical results: Computation of an efficient frontier. In this section we consider a concrete example involving selection between a single risky asset, representing a US stock index, and a risk-free asset. We compute an efficient frontier representing the trade-off between expected log-return and CVaR when using optimal dynamic strategies. For comparison, we compare with an efficient frontier when restricting to static strategies, i.e., strategies where A is constant over time, representing a fixed leverage ratio.

For our example, we choose $\mu = 11\%$, $\sigma = 20\%$,⁷ and $r = 1\%$ as market parameters. We take our time horizon as $T = 1$ and constrain our leverage ratio to lie within the range $A := [-6, +6]$.⁸ Finally, we consider CVaR at the $\alpha = 95\%$ threshold.

For each fixed $\lambda > 0$, we solve the corresponding dynamic mean-CVaR optimization problem using the techniques outlined in section 5.2.⁹ An example of the convergence of the approximation scheme described in Theorem 5.1, compared with the corresponding theoretical error bounds, is shown in Figure 1. This illustrates the linear decrease in relative error from approximation as a function of the approximation parameter ϵ .

⁷This choice corresponds, roughly, to the historical arithmetic mean and standard deviation of annual returns on the S&P 500, including dividend reinvestment, over the period 1928–2014. However, we emphasize that the exact choice of parameters should not be taken too seriously in this example.

⁸We choose this range to correspond, roughly, to the maximum leverage a qualifying US investor can achieve with a portfolio margin policy, as described at <http://www.finra.org/industry/portfolio-margin-faq>. In practice, the exact constraints depend upon the type of investor and financial instruments used for investment. We emphasize that this choice is meant for illustration only.

⁹We use finite-difference solvers for PDEs (3.6) and (3.8). We apply upwinding methods to obtain a monotone scheme when solving the gradient PDE [18, 5]. It is difficult to make any quantitative statements regarding the choice of discretization parameters. However, we note that the approximation parameter, ϵ , should not be too small compared to the mesh spacing used in the finite-difference method or the difference between f and f_ϵ will become numerically irresolvable.




FIG. 2. The efficient frontier of mean-CVaR portfolio optimization, representing the possible trade-off between maximizing expected log-returns and minimizing CVaR, computed by varying $\lambda \in (0, 1]$.

We compute points on the efficient frontier between expected log-return and CVaR by varying λ over the interval $(0, 1]$. At each value of λ , we compute the expected log-return under the optimal control by solving a linear parabolic equation analogous to (3.8). We then compute the corresponding CVaR using our computation of expected log-return and the mean-CVaR objective. The resulting frontier is shown in Figure 2 (solid).

For comparison, we consider the same optimization problem when restricted to a subcollection of static controls, defined as

$$\mathcal{A}_{\text{static}} := \{A \in \mathcal{A} \mid \exists a \in \mathbb{A} \text{ such that } A(t) = a \text{ for all } t \in [0, T] \text{ a.s.}\}.$$

These strategies represent constant leverage portfolios. An important example of these is the “buy-and-hold” strategy, e.g., $A(t) \equiv 1$. Under this class of controls, X_T^A is normally distributed. Therefore, we can directly compute optimal strategies and construct the efficient frontier.

In Figure 2, we illustrate a comparison between the efficient frontier under our dynamic strategies and under static strategies. We see that by employing strategies with dynamic leverage, we can significantly reduce CVaR at the 95% quantile while maintaining the same expected log-return, as compared to a static leverage strategy. Similarly, we can increase expected log-return while maintaining the same CVaR using a dynamic strategy. For example, the static buy-and-hold strategy, $A(t) \equiv 1$, has an expected log-return of 9% and CVaR of approximately 32%. By employing strategies with dynamic leverage, we can reduce CVaR by approximately 50% while maintaining the same expected log-return or, alternatively, increase expected log-return by approximately 30% while maintaining the same CVaR.

We next turn our attention to an examination of statistical and qualitative properties of the optimal dynamic control and resulting returns. In Figure 3, we illustrate the cumulative distribution function (CDF) of X_T^A under the optimal dynamic control corresponding an expected log-return of 9%. We compare this to the CDF of X_T^A under the buy-and-hold strategy, which follows a normal distribution. While both of these distributions have the same expected value, the one corresponding to the optimal dynamic strategy has significantly fatter (right) tails on the upside and an effective (left) floor on losses on the downside. We attribute this to a (de)leveraging




FIG. 4. (a) A sample path of stock prices and the corresponding portfolio log-return process (X^{A^*}) and (b) the corresponding optimal leverage process (A^*).

effect of the dynamic strategy whereby it increases leverage significantly once it has “locked in” gains and will deleverage only as needed to discourage losses exceeding a certain threshold.

This qualitative tendency of the optimal strategy to increase in leverage once it has locked in gains is emphasized further by sample paths illustrated in Figure 4. Here, we illustrate a particular sample path of stock prices (quoted as log-return), as well as the corresponding optimal dynamic leverage process, A^* , and the resulting portfolio log-return process, X^{A^*} . Note that the stock price corresponds to the log-return under the static buy-and-hold strategy, $A(t) \equiv 1$. We observe that early on in the period, the leverage process increases or decreases in sync with overall portfolio returns. However, as it becomes later in the period and the portfolio return is positive,




FIG. 5. (a) A sample path of stock prices and the corresponding portfolio log-return process (X^{A^*}) and (b) the corresponding optimal leverage process (A^*).

the optimal leverage increases significantly before being capped at a fixed value. The optimal strategy generally does not appear to decrease leverage late in the period, even with stock price declines, unless it is risking falling below the loss threshold seen in the jump in Figure 3.

In Figure 5, we illustrate an alternative sample path which emphasizes how the increasing leverage can lead to large returns on the upside. In this path, the leverage process, A^* , initially decreases to lower risk as the portfolio takes initial losses. However, in the latter half of the period, as stock prices rise, the increasing leverage leads to a return on the portfolio which significantly exceeds that of the buy-and-hold strategy. It is this transition from low leverage when avoiding tail losses to high leverage when locking in gains which allows the strategy to maintain a low CVaR while maximizing expected log-return.

The tendency of the optimal dynamic strategy to keep leverage higher than the static strategy unless it is facing losses also helps explain the skew seen in Figure 3. Because the dynamic strategy has the option to decrease its leverage to stop losses, it can achieve a significantly lower CVaR while maintaining a preference for high leverage, which contributes to large returns in positive outcomes. However, there is no such thing as a free lunch; in neutral outcomes, the positive correlation between log-returns and leverage leads to decay in portfolio value from convexity [46]. In this sense, the optimal dynamic strategy shares many qualitative features with constant proportion portfolio insurance (CPPI) strategies [9]. This makes sense as CPPI strategies are generally employed to limit downside losses, while maintaining upside gains, using dynamic trading.

In this example, we chose to consider only a single risky asset for ease of interpretation of the strategy. However, it is clear from the generality of section 5.2 that we could perform the same computations with multiple assets without increasing the size of the state space in the stochastic optimal control subproblem. In the multiple

asset case, the extra complexity appears in the need to solve a quadratic program with more decision variables at each step in the solution of the HJB in (3.6).

Appendix A. Coherent extremal risk measures.

PROPOSITION A.1 (coherence). *Suppose the following properties of $f : \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}$ hold:*

- *Positive-homogeneity and normalization: $f(ax, ay) = af(x, y)$ for $a > 0$ and $\inf_{y \in \mathbb{R}^m} f(0, y) = 0$.*
- *Monotonicity: $x \mapsto f(x, y)$ is nondecreasing for each $y \in \mathbb{R}^m$.*
- *Subadditivity: $f(x_1 + x_2, y_1 + y_2) \leq f(x_1, y_1) + f(x_2, y_2)$ for $x_1, x_2 \in \mathbb{R}$ and $y_1, y_2 \in \mathbb{R}^m$.*
- *Translation: For every $a \in \mathbb{R}$, there exists an invertible function $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^m$ such that for every $(x, y) \in \mathbb{R} \times \mathbb{R}^m$ we have $f(x + a, y) = f(x, \phi(y)) + a$.*

Then, the function $\rho : L^2(\Omega) \rightarrow \mathbb{R}$ defined as

$$(A.1) \quad \rho(\xi) := \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi, y)]$$

is a coherent (extremal) risk measure.

Proof. Recall, the four properties of a coherent risk measure are positive homogeneity, monotonicity, subadditivity, and translational invariance [2]. We show below how each of these comes from the assumptions on f .

1. Positive-homogeneity: We compute directly

$$\rho(a\xi) = \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(a\xi, y)] = \inf_{ay \in \mathbb{R}^m} \mathbb{E}[f(a\xi, ay)] = a \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi, y)] = a\rho(\xi)$$

for $a > 0$. When $a = 0$,

$$\rho(0) = \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(0, y)] = \inf_{y \in \mathbb{R}^m} f(0, y) = 0.$$

2. Monotonicity: Let $\xi, \xi' \in L^2(\Omega)$ such that $\xi \leq \xi'$ almost surely. Then,

$$\rho(\xi) = \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi, y)] \leq \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi', y)] = \rho(\xi').$$

3. Subadditivity: Fix $\xi_1, \xi_2 \in L^2(\Omega)$. We compute

$$\begin{aligned} \rho(\xi_1 + \xi_2) &= \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi_1 + \xi_2, y)] \\ &= \inf_{y_1, y_2 \in \mathbb{R}^m} \mathbb{E}[f(\xi_1 + \xi_2, y_1 + y_2)] \\ &\leq \inf_{y_1, y_2 \in \mathbb{R}^m} \mathbb{E}[f(\xi_1, y_1) + f(\xi_2, y_2)] = \rho(\xi_1) + \rho(\xi_2). \end{aligned}$$

4. Translational-invariance: Let $\xi \in L^2(\Omega)$ and $a \in \mathbb{R}$. Then

$$\rho(\xi + a) = \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi + a, y)] = \inf_{y \in \mathbb{R}^m} \mathbb{E}[f(\xi, \phi(y)) + a] = \rho(\xi) + a. \quad \square$$

Appendix B. Proof of Theorem 3.2.

Proof. Let $y, y' \in \mathbb{R}^m$ and $\theta \in [0, 1]$. For any $\epsilon > 0$, let $A, A' \in \mathcal{A}$ be ϵ -suboptimal controls such that

$$V(y) + \epsilon \geq \mathbb{E}[f(g(X_T^A), y)] \text{ and } V(y') + \epsilon \geq \mathbb{E}[f(g(X_T^{A'}), y')].$$

By Assumption 3, we obtain

$$\begin{aligned} V(\theta y + (1 - \theta)y') &\leq \mathbb{E} \left[f \left(g \left(X_T^{\theta A + (1 - \theta)A'} \right), \theta y + (1 - \theta)y' \right) \right] \\ &\leq \mathbb{E} \left[\theta f \left(g(X_T^A), y \right) + (1 - \theta) f \left(g(X_T^{A'}), y' \right) \right] \\ &\leq \theta V(y) + (1 - \theta)V(y') + \epsilon. \end{aligned}$$

Because ϵ was arbitrary, the convexity of V follows. \square

Appendix C. Proof of Proposition 4.2.

Proof. Recall that Assumption 3 states that the map

$$(A, y) \mapsto f(g(X_T^A), y)$$

is jointly convex almost surely. We observe that the map

$$(A, y, z) \mapsto f(g(X_T^A), z) + \frac{\|y - z\|^2}{2\epsilon}$$

is almost surely jointly convex as a sum of convex functions. Therefore,

$$(A, y) \mapsto f_\epsilon(g(X_T^A), y) = \inf_{z \in \mathbb{R}^m} \left[f(g(X_T^A), z) + \frac{\|y - z\|^2}{2\epsilon} \right]$$

is almost surely jointly convex by the same type of argument as in the proof of Theorem 3.2. \square

Acknowledgment. We would like to thank Prof. Lawrence C. Evans for useful discussions on the uniqueness of viscosity solutions.

REFERENCES

- [1] A. AHMADI-JAVID AND R. MALHAMÉ, *Optimal control of a multistate failure-prone manufacturing system under a conditional value-at-risk cost criterion*, J. Optim. Theory Appl., 167 (2015), pp. 716–732.
- [2] P. ARTZNER, F. DELBAEN, J.-M. EBER, AND D. HEATH, *Coherent measures of risk*, Math. Finance, 9 (1999), pp. 203–228.
- [3] P. ARTZNER, F. DELBAEN, J.-M. EBER, D. HEATH, AND H. KU, *Coherent multiperiod risk adjusted values and Bellman's principle*, Ann. Oper. Res., 152 (2007), pp. 5–22.
- [4] V. AZHMYAKOV AND J. RAISCH, *Convex control systems and convex optimal control problems with constraints*, IEEE Trans. Automat. Control, 53 (2008), pp. 993–998.
- [5] G. BARLES AND P. E. SOUGANIDIS, *Convergence of approximation schemes for fully nonlinear second order equations*, Asymptot. Anal., 4 (1991), pp. 271–283.
- [6] N. BÄUERLE AND J. OTT, *Markov decision processes with average-value-at-risk criteria*, Math. Methods Oper. Res., 74 (2011), pp. 361–379.
- [7] E. BAYRAKTAR AND C. W. MILLER, *Distribution-Constrained Optimal Stopping*, preprint, arXiv:1604.03042, 2016.
- [8] T. BJORK AND A. MURGOJI, *A General Theory of Markovian Time Inconsistent Stochastic Control Problems*, preprint, SSRN 1694759, 2010.
- [9] F. BLACK AND A. F. PEROLD, *Theory of constant proportion portfolio insurance*, J. Econom. Dynam. Control, 16 (1992), pp. 403–426.
- [10] O. BOKANOWSKI, J. GARCKE, M. GRIEBEL, AND I. KLOMPMAKER, *An adaptive sparse grid semi-Lagrangian scheme for first order Hamilton-Jacobi Bellman equations*, J. Sci. Comput., 55 (2013), pp. 575–605.
- [11] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, Cambridge, 2004.
- [12] L. A. CAFFARELLI AND X. CABRÉ, *Fully Nonlinear Elliptic Equations*, AMS, Providence, RI, 1995.

- [13] P. CANNARSA AND C. SINESTRARI, *Semicconcave Functions, Hamilton–Jacobi Equations, and Optimal Control*, Birkhäuser, Boston, 2004.
- [14] Ö ÇAVUŞ AND A. RUSZCZYŃSKI, *Risk-averse control of undiscounted transient Markov models*, SIAM J. Control Optim., 52 (2014), pp. 3935–3966.
- [15] P. CHERIDITO, F. DELBAEN, AND M. KUPPER, *Dynamic monetary risk measures for bounded discrete-time processes*, Electron. J. Probab., 11 (2006), pp. 57–106.
- [16] Y. CHOW, A. TAMAR, S. MANNOR, AND M. PAVONE, *Risk-sensitive and robust decision-making: A CVaR optimization approach*, in NIPS, MIT Press, Cambridge, MA, 2015, pp. 1522–1530.
- [17] P. COLANERI, R. H. MIDDLETON, Z. CHEN, D. CAPORALE, AND F. BLANCHINI, *Convexity of the cost functional in an optimal control problem for a class of positive switched systems*, Automatica J. IFAC, 50 (2014), pp. 1227–1234.
- [18] R. COURANT, E. ISAACSON, AND M. REES, *On the solution of nonlinear hyperbolic differential equations by finite differences*, Comm. Pure. Appl. Math., 5 (1952), pp. 243–255.
- [19] H. DONG AND D. KIM, *Elliptic and parabolic equations with measurable coefficients in weighted Sobolev spaces*, Adv. Math., 274 (2015), pp. 681–735.
- [20] Y. DONG AND R. SIRCAR, *Time-inconsistent portfolio investment problems*, in Stochastic Analysis and Applications, Springer, Heidelberg, 2014, pp. 239–281.
- [21] L. C. EVANS, *Classical solutions of the Hamilton–Jacobi–Bellman equation for uniformly elliptic operators*, Trans. Amer. Math. Soc., 275 (1983), pp. 245–255.
- [22] L. C. EVANS, *An Introduction to Stochastic Differential Equations*, AMS, Providence, RI, 2013.
- [23] W. H. FLEMING AND W. M. MCENEANEY, *Risk-sensitive control on an infinite time horizon*, SIAM J. Control Optim., 33 (1995), pp. 1881–1915.
- [24] W. H. FLEMING AND H. M. SONER, *Controlled Markov Processes and Viscosity Solutions*, 2nd ed., Springer, New York, 2006.
- [25] M. FRITELLI AND G. SCANDOLO, *Risk measures and capital requirements for processes*, Math. Finance, 16 (2006), pp. 589–612.
- [26] D. GILBARG AND N. S. TRUDINGER, *Elliptic Partial Differential Equations of Second Order*, 2nd ed., Springer-Verlag, 2001.
- [27] W. B. HASKELL AND R. JAIN, *A convex analytic approach to risk-aware Markov decision processes*, SIAM J. Control Optim., 53 (2015), pp. 1569–1598.
- [28] R. H. W. HOPPE, *Multi-grid methods for Hamilton–Jacobi–Bellman equations*, Numer. Math., 49 (1986), pp. 239–254.
- [29] Y. HUANG AND X. GUO, *Minimum average value-at-risk for finite horizon semi-Markov decision processes in continuous time*, SIAM J. Optim., 26 (2016), pp. 1–28.
- [30] M. R. JAMES, J. S. BARAS, AND R. J. ELLIOTT, *Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems*, IEEE Trans. Automat. Control, 39 (1994), pp. 780–792.
- [31] C. KARNAM, J. MA, AND J. ZHANG, *Dynamic Approaches for Some Time Inconsistent Problems*, preprint, arXiv:1604.03913, 2016.
- [32] N. KATZOURAKIS, *An Introduction to Viscosity Solutions for Fully Nonlinear PDE with Applications to Calculus of Variations in L^∞* , Springer, Cham, Switzerland, 2015.
- [33] B. KAWOHL AND N. KUTEV, *Strong maximum principle for semicontinuous viscosity solutions of nonlinear partial differential equations*, Arch. Math. (Basel), 70 (1998), pp. 470–478.
- [34] D. KIM, *Parabolic equations with measurable coefficients. II*, J. Math. Anal. Appl., 334 (2007), pp. 534–548.
- [35] D. KIM, *Second order elliptic equations in \mathbb{R}^d with piecewise continuous coefficients*, Potential Anal., 26 (2007), pp. 189–212.
- [36] S. KLÖPPEL AND M. SCHWEIZER, *Dynamic indifference valuation via convex risk measures*, Math. Finance, 17 (2007), pp. 599–627.
- [37] N. V. KRYLOV, *On weak uniqueness for some diffusions with discontinuous coefficients*, Stochastic Process. Appl., 113 (2004), pp. 37–64.
- [38] N. V. KRYLOV, *Parabolic and elliptic equations with VMO coefficients*, Comm. Partial Differential Equations, 32 (2007), pp. 453–475.
- [39] P.-L. LIONS, *On the Hamilton–Jacobi–Bellman equations*, Acta Appl. Math., 1 (1983), pp. 17–41.
- [40] R. MANSINI, W. OGRYCZAK, AND M. G. SPERANZA, *Conditional value at risk and related linear programming models for portfolio optimization*, Ann. Oper. Res., 152 (2007), pp. 227–256.
- [41] C. W. MILLER, *Non-linear PDE approach to time-inconsistent optimal stopping*, SIAM J. Control Optim., to appear.
- [42] N. NADIRASHVILI, *Nonuniqueness in the Martingale problem and the Dirichlet problem for uniformly elliptic operators*, Ann. Sc. Norm. Sup. Pisa Cl. Sci. (5), 24 (1997), pp. 537–549.

- [43] J. NOCEDAL AND S. WRIGHT, *Numerical Optimization*, Springer, Berlin, 2006.
- [44] W. OGRYCZAK AND A. RUSZCZYŃSKI, *From stochastic dominance to mean-risk models: Semideviations as risk measures*, European J. Oper. Res., 116 (1999), pp. 33–50.
- [45] J. L. PEDERSEN AND G. PESKIR, *Optimal mean-variance portfolio selection*, Math. Financ. Econ., 11 (2017), pp. 137–160.
- [46] A. F. PEROLD AND W. F. SHARPE, *Dynamic strategies for asset allocation*, Financ. Anal. J., 44 (1988), pp. 16–27.
- [47] L. PFEIFFER, *Risk-averse Merton's portfolio problem*, in Second IFAC Workshop on Control of Systems Governed by Partial Differential Equations, Elsevier, New York, 2016, pp. 266–271.
- [48] G. C. PFLUG AND A. PICHLER, *Time-inconsistent multistage stochastic programs: Martingale bounds*, European J. Oper. Res., 249 (2016), pp. 155–163.
- [49] J. QIN, H.-I. SU, AND R. RAJAGOPAL, *Storage in risk limiting dispatch: Control and approximation*, in American Control Conference, IEEE, Piscataway, NJ, 2013, pp. 4202–4208.
- [50] R. RAJAGOPAL, E. BITAR, W. WU, AND P. VARAIYA, *Risk limiting dispatch of wind power*, in American Control Conference, IEEE, Piscataway, NJ, 2012, pp. 4417–4422.
- [51] F. RIEDEL, *Dynamic coherent risk measures*, Stochastic Process. Appl., 112 (2004), pp. 185–200.
- [52] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [53] R. T. ROCKAFELLAR AND S. URYASEV, *Optimization of conditional value-at-risk*, J. Risk, 2 (2000), pp. 21–42.
- [54] R. T. ROCKAFELLAR AND S. URYASEV, *Conditional value-at-risk for general loss distribution*, J. Bank. Finan., 26 (2002), pp. 1443–1471.
- [55] A. RUSZCZYŃSKI, *Risk-averse dynamic programming for Markov decision processes*, Math. Program., 125 (2010), pp. 235–261.
- [56] A. RUSZCZYŃSKI AND A. SHAPIRO, *Optimization of convex risk functions*, Math. Oper. Res., 31 (2006), pp. 433–452.
- [57] A. RUSZCZYŃSKI AND J. YAO, *Discrete-Time Approximation of Risk-Averse Control Problems for Diffusion Processes*, preprint, arXiv:1508.05316 [math.OC], 2015.
- [58] M. V. SAFONOV, *Nonuniqueness for second-order elliptic equations with measurable coefficients*, SIAM J. Math. Anal., 30 (1999), pp. 879–895.
- [59] S. SARYKALIN, G. SERRAINO, AND S. URYASEV, *Value-at-risk vs. conditional value-at-risk in risk management and optimization*, in Tutorials in Operations Research, INFORMS, Hanover, MD, 2008, pp. 270–294.
- [60] P. SORAVIA, *Uniqueness results for fully nonlinear degenerate elliptic equations with discontinuous coefficients*, Comm. Pure Appl. Anal., 5 (2006), pp. 213–240.
- [61] N. TOUZI, *Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE*, Springer, New York, 2013.
- [62] I. YANG, D. S. CALLAWAY, AND C. J. TOMLIN, *Indirect load control for electricity market risk management via risk-limiting dynamic contracts*, in American Control Conference, IEEE, Piscataway, NJ, 2015, pp. 3025–3031.
- [63] I. YANG, D. S. CALLAWAY, AND C. J. TOMLIN, *Variance-constrained risk sharing in stochastic systems*, IEEE Trans. Automat. Control, to appear.