

Distribution-Constrained Optimal Stopping

Erhan Bayraktar*

Christopher W. Miller†

Abstract

We solve the problem of optimal stopping of a Brownian motion subject to the constraint that the stopping time's distribution is a given measure consisting of finitely-many atoms. In particular, we show that this problem can be converted to a finite sequence of state-constrained optimal control problems with additional states corresponding to the conditional probability of stopping at each possible terminal time. The proof of this correspondence relies on a new variation of the dynamic programming principle for state-constrained problems which avoids measurable selections. We emphasize that distribution constraints lead to novel and interesting mathematical problems on their own, but also demonstrate an application in mathematical finance to model-free superhedging with an outlook on volatility.

Key words. Optimal stopping, distribution constraints, optimal control, state constraints, robust hedging with a volatility outlook.

AMS subject classifications. 60G40, 93E20, 91G80.

1 Introduction

In this paper we consider the problem of choosing an optimal stopping time for a Brownian motion when constrained in the choice of distribution for the stopping time. While standard optimal stopping theory has focused primarily on unconstrained finite- and infinite-horizon stopping times (e.g. [22, 25]) and very recently on constraints on the first moment of the stopping time (e.g. [18, 21, 1]), there is a very limited literature on the problem of optimal stopping under distribution constraints.

It turns out that distribution-constrained optimal stopping is a difficult problem, with stopping strategies depending path-wise on the Brownian motion in general. This is to be expected because a constraint on the stopping time's distribution forces the stopper to consider what he would have done along all other paths of the Brownian motion when deciding whether to stop. The main task at hand is to identify relevant state variables and then transform the problem so that it can be analyzed by standard methods.

In this article we illustrate a solution in the special case that the target distribution consists of finitely-many atoms. Our approach consists of iterated stochastic control problems wherein we introduce controlled processes representing the conditional distribution of the stopping time. We then characterize the value function of the distribution-constrained optimal stopping problem in terms of the value functions of a finite number of state-constrained optimal control problems. This dynamic approach to the problem in terms of a controlled process with unbounded diffusion is similar in flavour to recent results in non-linear optimal stopping [18] and control of measure-valued martingales in [10].

The key mathematical contributions of this paper lie in our proof of a dynamic programming principle relating each of the sequential optimal control problems. We provide an argument

*Department of Mathematics, University of Michigan (erhan@umich.edu). Supported in part by the NSF under grant number DMS-1613170 and in part by the Susan M. Smith Professorship.

†Department of Mathematics, University of California, Berkeley (miller@math.berkeley.edu). Supported in part by NSF GRFP under grant number DGE 1106400.

which avoids the use of measurable selections, similar to the proofs of weak dynamic programming principles in [7, 6, 3]. However, we deal with state-constraints in a novel way which relies on some a priori regularity of the value functions.

While the problem of distribution-constrained optimal stopping is mathematically-interesting in its own right, we emphasize that there is room for applications in mathematical finance and optimal control theory. For instance, we demonstrate an application to model-free superhedging of financial derivatives when one has an outlook on the quadratic variation of an asset price. Here, the distribution on the quadratic variation corresponds to that of a stopping time by the martingale time-change methods utilized recently in [5, 14]. Furthermore, the problem of optimal stopping under moment constraints on the stopping time reduces to the distribution-constrained optimal stopping problem in cases where there exists a unique atomic representing measure in the truncated moment problem (e.g. [11, 17]). There also appears to be a connection between distribution-constrained optimal stopping and inverse first passage-time problems (e.g. [27, 9]). We should also mention that after the publication of our preprint, [4] gave geometric descriptions of optimal stopping times using optimal transport theory.

This paper proceeds as follows. In Section 2, we provide our solution to distribution-constrained optimal stopping of Brownian motion. In particular, we characterize the solution via a finite sequence of iterated state-constrained stochastic control problems. The main result is provided by an induction argument in Theorem 1, but the heart of the argument lies mainly in the proofs of Lemma 3 and Lemma 4. We also provide time-dependent versions of these results, which can be characterized as the viscosity solutions of associated HJB equations. The key arguments here lie in a Dynamic Programming Principle in Theorem 2. In Section 3, we demonstrate an application to model-free superhedging with an outlook on volatility. We convert this problem into a distribution-constrained optimal stopping problem where the volatility outlook corresponds to a distribution constraint for the stopping time. We demonstrate numerical results which provide some intuition for the behavior of the optimal stopping strategies. Finally, we provide complete proofs of our main results in Appendices A–D.

2 Main Results

2.1 Problem formulation

We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ supporting a standard Brownian motion W . We take $\mathbb{F} := \{\mathcal{F}_t\}_{t \geq 0}$ to be the natural filtration of W augmented to satisfy the usual properties. We consider a given payoff function $f : \mathbb{R} \rightarrow \mathbb{R}$ which is assumed to be Lipschitz continuous. We also use the notation

$$X_u^{t,x} := x + W_u - W_t$$

for any $(x, t) \in \mathbb{R} \times [0, \infty)$ and $u \in [t, \infty)$.

In this paper, we are also given a target distribution μ , which is supported on $(0, \infty)$ and assumed to consist of finitely-many atoms. Without loss of generality, we assume the following representation

$$\mu = \sum_{k=1}^r p_k \delta_{t_k}, \quad (2.1)$$

where $r \in \mathbb{N}$, $0 = t_0 < t_1 < \dots < t_r$, $p_1 + \dots + p_r = 1$, and $p_1, \dots, p_r > 0$. We also introduce the convenient notation $\Delta t_k := t_k - t_{k-1}$ for each $k \in \{1, \dots, r\}$.

The distribution-constrained optimal stopping problem we consider is

$$v^* := \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[f(X_\tau^{0,x_0}) \right] \quad (2.2)$$

s.t. $\tau \sim \sum_{k=1}^r p_k \delta_{t_k}$,

where we take \mathcal{T} to be the collection of all finite-valued \mathbb{F} -stopping times which are independent of \mathcal{F}_0 . We let $x_0 \in \mathbb{R}$ be some fixed starting value. That is, we choose a stopping time τ whose distribution is equal to μ in order to maximize the expected payoff of a stopped Brownian motion starting at x_0 .

2.2 Construction of distribution-constrained stopping times

There are multiple ways to naturally represent a stopping time satisfying a distribution constraint. In this section, we outline two particular such representations and illustrate how they immediately lead to constructions of such stopping times.

We first provide a characterization of distribution-constrained stopping times in terms of a partitioning of the path space into regions with specified measures. Later, we make a connection with controlled processes.

Lemma 1. *A stopping time τ has the distribution μ if and only if it is of the following form*

$$\tau = \sum_{k=1}^r t_k 1_{A_k},$$

almost-surely, where $\{A_1, \dots, A_r\}$ partition Ω and, for each $k \in \{1, \dots, r\}$, A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$.

Proof. It is clear from the construction that such a τ is a \mathbb{F} -stopping time and $\tau \sim \mu$. The converse follows by taking a stopping time τ such that $\tau \sim \mu$ and defining the sets $A_k := \{\tau = t_k\}$ for each $k \in \{1, \dots, r\}$. □

With this in mind, we can immediately explicitly construct a stopping time with given distribution.

Corollary 1. *There exists a stopping time τ such that $\tau \sim \mu$.*

Proof. Define a partition $\{A_1, \dots, A_r\}$ of Ω as

$$\begin{aligned} A_1 &:= \{W_{t_1} - W_0 \leq \sqrt{t_1} \Phi^{-1}(p_1)\} \\ A_2 &:= \left\{ W_{t_2} - W_{t_1} \leq \sqrt{t_2 - t_1} \Phi^{-1}\left(\frac{p_2}{p_2 + \dots + p_r}\right) \right\} \setminus A_1 \\ &\vdots \\ A_k &:= \left\{ W_{t_k} - W_{t_{k-1}} \leq \sqrt{t_k - t_{k-1}} \Phi^{-1}\left(\frac{p_k}{p_k + \dots + p_r}\right) \right\} \setminus (A_1 \cup \dots \cup A_{k-1}) \\ &\vdots \\ A_r &:= \Omega \setminus (A_1 \cup \dots \cup A_{r-1}), \end{aligned}$$

where Φ is the cumulative distribution function of the standard normal distribution. It is clear that A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$ for each $k \in \{1, \dots, r\}$. Then, by Lemma 1, $\tau := \sum_{k=1}^r t_k 1_{A_k}$ defines a stopping time with $\tau \sim \mu$. □

The proof above constructs a stopping time which roughly stops when there are events in the left-tail of a distribution. However, one could easily modify the construction to stop in right-tail events, events near the median, or on the image of any Borel set of appropriate measure under Φ .

While this construction may suggest converting the distribution-constrained optimal stopping problem into optimization over Borel sets of specified measure, we emphasize next that

there is no reason to expect the stopping times to be measurable with respect to $\sigma(W_{t_1}, \dots, W_{t_r})$. In particular, in the next example, we show a construction of a distribution-constrained stopping time which is entirely path-dependent.

Corollary 2. *There exists a stopping time τ , independent of $(W_{t_1}, \dots, W_{t_r})$, satisfying $\tau \sim \mu$.*

Proof. Define a sequence of random variables (M_1, \dots, M_r) as

$$M_k := (t_k - t_{k-1})^{-1/2} \max_{t_{k-1} \leq s \leq t_k} \left| W_s - W_{t_{k-1}} - (s - t_{k-1}) \frac{W_{t_k} - W_{t_{k-1}}}{t_k - t_{k-1}} \right|$$

for each $k \in \{1, \dots, r\}$. Then each M_k is the absolute maximum of a Brownian bridge over $[t_{k-1}, t_k]$, scaled by the length of the time interval. In particular, each M_k is \mathcal{F}_{t_k} -measurable, independent of $(W_{t_1}, \dots, W_{t_r})$, and equal in distribution to the absolute maximum of a standard Brownian bridge on $[0, 1]$, the cumulative distribution function of which we denote by Φ_{BB} .

Define a partition $\{A_1, \dots, A_r\}$ of Ω as

$$\begin{aligned} A_1 &:= \{M_1 \leq \Phi_{BB}^{-1}(p_1)\} \\ A_2 &:= \left\{ M_2 \leq \Phi_{BB}^{-1}\left(\frac{p_2}{p_2 + \dots + p_r}\right) \right\} \setminus A_1 \\ &\vdots \\ A_k &:= \left\{ M_k \leq \Phi_{BB}^{-1}\left(\frac{p_k}{p_k + \dots + p_r}\right) \right\} \setminus (A_1 \cup \dots \cup A_{k-1}) \\ &\vdots \\ A_r &:= \Omega \setminus (A_1 \cup \dots \cup A_{r-1}). \end{aligned}$$

It is clear that A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$ for each $k \in \{1, \dots, r\}$. Then, by Lemma 1, $\tau := \sum_{k=1}^r t_k 1_{A_k}$ defines a stopping time with $\tau \sim \mu$ which is independent of $(W_{t_1}, \dots, W_{t_r})$. \square

Clearly, the stopping time constructed above is an admissible stopping time in the distribution-constrained optimal stopping problem, but there is no hope to express it in terms of the value of the Brownian motion at each potential time to stop. While stopping times involving the Brownian bridge may seem unnatural at first, their use is a key idea in the proofs of Lemma 3 and Lemma 4.

Another useful result obtained from Lemma 1 is the following approximation result.

Proposition 1. *Fix $p, p' \in [0, 1]^r$ satisfying $p_1 + \dots + p_r = 1$ and $p'_1 + \dots + p'_r = 1$. For any $\tau \in \mathcal{T}$ such that*

$$\tau \sim \sum_{k=1}^r p_k \delta_{t_k},$$

there exists $\tau' \in \mathcal{T}$ such that

$$\tau' \sim \sum_{k=1}^r p'_k \delta_{t_k},$$

which satisfies

$$\mathbb{P}[\tau \neq \tau'] \leq 4^r \|p - p'\|_{\ell^1}.$$

Proof. Step 1: By Lemma 1, there exists a partition $\{A_1, \dots, A_r\}$ of Ω such that each A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$ and

$$\tau = \sum_{k=1}^r t_k 1_{A_k},$$

almost-surely. The goal is to define τ' in terms of a related partition.

We first make a key observation: for any $t > 0$, $\theta \in (0, 1)$, and \mathcal{F}_t -measurable set A such that $\mathbb{P}[A] > 0$, there exists a real number $w \in \mathbb{R}$ such that

$$\mathbb{P}[W_t \geq w \mid A] = \theta.$$

This follows immediately from the observation that the distribution of W_t has no atoms, and thus when conditioning on an event of non-zero probability, the conditional distribution cannot have atoms.

Step 2: Define an \mathcal{F}_{t_1} -measurable set A'_1 as

$$A'_1 := \begin{cases} A_1 & p_1 = p'_1 \\ A_1 \cup \{W_{t_1} \geq w_1^+\} & p_1 < p'_1 \\ A_1 \cap \{W_{t_1} \geq w_1^-\} & p_1 > p'_1 \end{cases}$$

where $w_1^+, w_1^- \in \mathbb{R}$ are chosen such that

$$\mathbb{P}[W_{t_1} \geq w_1^+ \mid \Omega \setminus A_1] = \frac{p'_1 - p_1}{1 - p_1}, \quad \mathbb{P}[W_{t_1} \geq w_1^- \mid A_1] = \frac{p'_1}{p_1}.$$

Notice, if $p_1 < p'_1$, then $p_1 < 1$ and $\mathbb{P}[\Omega \setminus A_1] > 0$, so w_1^+ is well-defined. Similarly, if $p_1 > p'_1$, then $p_1 > 0$ and $\mathbb{P}[A_1] > 0$, so w_1^- is well-defined.

It is clear that $\mathbb{P}[A'_1] = p'_1$ by construction. The key property, however, is that either $A_1 \subseteq A'_1$ or $A'_1 \subseteq A_1$. From this, we can immediately compute the measure of the symmetric difference of A_1 and A'_1 ,

$$\mathbb{P}[A_1 \triangle A'_1] = |p_1 - p'_1|.$$

Step 3: Now, suppose that we have constructed $\{A'_1, \dots, A'_{k-1}\}$ already. Define

$$q_k := \mathbb{P}[A_k \setminus (A'_1 \cup \dots \cup A'_{k-1})].$$

Define a \mathcal{F}_{t_k} -measurable set A'_k as

$$A'_k := \begin{cases} A_k \setminus (A'_1 \cup \dots \cup A'_{k-1}) & q_k = p'_k \\ (A_k \setminus (A'_1 \cup \dots \cup A'_{k-1})) \cup \{W_{t_k} \geq w_k^+\} & q_k < p'_k \\ (A_k \setminus (A'_1 \cup \dots \cup A'_{k-1})) \cap \{W_{t_k} \geq w_k^-\} & q_k > p'_k \end{cases}$$

where $w_k^+, w_k^- \in \mathbb{R}$ are chosen such that

$$\begin{aligned} \mathbb{P}[W_{t_k} \geq w_k^+ \mid (A'_1 \cup \dots \cup A'_{k-1}) \cup (\Omega \setminus A_k)] &= \frac{p'_k - q_k}{1 - q_k}, \\ \mathbb{P}[W_{t_k} \geq w_k^- \mid A_k \setminus (A'_1 \cup \dots \cup A'_{k-1})] &= \frac{p'_k}{q_k}. \end{aligned}$$

As before, the inequalities between q_k and p'_k imply that w_k^+ and w_k^- are well-defined when they are needed. Furthermore, it is clear that $\mathbb{P}[A'_k] = p'_k$.

In this case, the key property becomes that either

$$A_k \setminus (A'_1 \cup \dots \cup A'_{k-1}) \subseteq A'_k \tag{2.3}$$

or

$$A'_k \subseteq A_k \setminus (A'_1 \cup \dots \cup A'_{k-1}) \subseteq A_k. \tag{2.4}$$

First, we consider the case that (2.3) holds. Because each set in $\{A_1, \dots, A_k\}$ is disjoint, for any $\ell \in \{1, \dots, k-1\}$, we can bound the overlap between A_k and A'_ℓ by the symmetric difference of A_ℓ and A'_ℓ ,

$$\mathbb{P}[A_k \cap A'_\ell] \leq \mathbb{P}[A'_\ell \setminus A_\ell] \leq \mathbb{P}[A_\ell \triangle A'_\ell].$$

Using (2.3), we can compute

$$\mathbb{P}[A_k \Delta A'_k] = \mathbb{P}[A'_k] - \mathbb{P}[A_k] + 2\mathbb{P}[A_k \setminus A'_k] \leq |p_k - p'_k| + 2\mathbb{P}[A_k \cap (A'_1 \cup \dots \cup A'_{k-1})].$$

Using the previous inequalities we deduce

$$\mathbb{P}[A_k \Delta A'_k] \leq |p_k - p'_k| + \sum_{\ell=1}^{k-1} 2\mathbb{P}[A_k \cap A'_\ell] \leq |p_k - p'_k| + \sum_{\ell=1}^{k-1} 2\mathbb{P}[A_\ell \Delta A'_\ell].$$

In the case that (2.4) holds, this same inequality immediately follows.

Step 4: By induction on $k \in \{1, \dots, r\}$, we construct a disjoint partition $\{A'_1, \dots, A'_r\}$ of Ω such that A'_k is \mathcal{F}_{t_k} -measurable and $\mathbb{P}[A'_k] = p'_k$. Furthermore, we obtain a rough bound

$$\mathbb{P}[A_k \Delta A'_k] \leq \sum_{\ell=1}^k 4^{k-\ell} |p_\ell - p'_\ell|.$$

By Lemma 1, there exists a stopping time τ' of the form

$$\tau' = \sum_{k=1}^r t_k 1_{A'_k}.$$

Then we can immediately compute¹

$$\mathbb{P}[\tau \neq \tau'] \leq \sum_{k=1}^r \mathbb{P}[A_k \Delta A'_k] \leq \sum_{k=1}^r 4^r |p_k - p'_k| \leq 4^r \|p - p'\|_{\ell^1}.$$

□

The technical importance of this result is, of course, that it allows us to obtain continuity in the problem with respect to changes in the distribution constraint on the stopping times.

While we have demonstrated that a lot can be said about distribution-constrained stopping times using the representation in Lemma 1, it turns out that we can obtain a more manageable representation if we introduce extra controlled processes which represent the conditional probability of the stopping time taking on each possible value. This vector-valued stochastic process is a martingale in a probability simplex. In the next result, we make clear the connection between this process and a distribution-constrained stopping time.

In the remainder of the paper, we define \mathcal{A} to denote the collection of all progressively-measurable, square-integrable, \mathbb{R}^r -valued processes which are independent of \mathcal{F}_0 . We also denote

$$Y_u^{t,y,\alpha} := y + \int_t^u \alpha_s dW_s$$

for all $(t, y) \in [0, \infty) \times \mathbb{R}^r$, $u \in [t, \infty)$, and $\alpha \in \mathcal{A}$. When needed, we will denote the k th coordinate of this vector-valued process by $Y^{(k),t,y,\alpha}$. We will occasionally abuse notation and leave out superscripts when they are clearly implied by the context.

We also denote by Δ the following closed and convex set

$$\Delta := \{y = (y_1, \dots, y_r) \in [0, 1]^r \mid y_1 + \dots + y_r = 1\} \subset \mathbb{R}^r.$$

We then can state a lemma regarding a characterization of distribution-constrained stopping times in terms of a state-constrained controlled martingale.

¹We emphasize that the preceding inequalities are hardly sharp. An immediate question is whether a sharper approximation result, with a constant that scales favorably as $r \rightarrow \infty$, exists.

Lemma 2. A stopping time $\tau \in \mathcal{T}$ has the distribution μ if and only if it is of the form

$$\tau = \min_{k \in \{1, \dots, r\}} \left\{ t_k \mid Y_{t_k}^{(k), 0, p, \alpha} = 1 \right\},$$

almost-surely, for some $\alpha \in \mathcal{A}$ such that

$$Y_t^{0, p, \alpha} \in \Delta,$$

almost-surely, for all $t \geq 0$, and

$$Y_{t_k}^{(k), 0, p, \alpha} \in \{0, 1\},$$

almost-surely, for each $k \in \{1, \dots, r\}$.

Proof. Step 1: Let $\alpha \in \mathcal{A}$ be a control for which $Y_t^{0, p, \alpha} \in \Delta$, almost-surely, for all $t \geq 0$ and $Y_{t_k}^{(k), 0, p, \alpha} \in \{0, 1\}$, almost-surely, for each $k \in \{1, \dots, r\}$. Define τ as

$$\tau := \min_{k \in \{1, \dots, r\}} \left\{ t_k \mid Y_{t_k}^{(k), 0, p, \alpha} = 1 \right\}.$$

It is clear from the properties above that $Y_{t_r}^{(k), 0, p, \alpha} \in \{0, 1\}$ for every $k \in \{1, \dots, r\}$ and $Y_{t_r}^{0, p, \alpha} \in \Delta$, which implies that $\tau \leq t_r$, almost-surely. Then $\tau \in \mathcal{T}$, but we must check that it has μ as its distribution.

Fix $k \in \{1, \dots, r\}$ and note that

$$\mathbb{P}[\tau = t_k] = \mathbb{P} \left[\underbrace{\{Y_{t_1}^{(1), 0, p, \alpha} = 0\} \cap \dots \cap \{Y_{t_{k-1}}^{(k-1), 0, p, \alpha} = 0\}}_A \cap \underbrace{\{Y_{t_k}^{(k), 0, p, \alpha} = 1\}}_B \right].$$

Note that $B \subset A$ up to a set of measure zero because in the set $B \setminus A$, we have $Y_{t_k}^{(k), 0, p, \alpha} = 1$ as well as $Y_{t_\ell}^{(\ell), 0, p, \alpha} = 1$ for some $\ell < k$. Because $Y^{0, p, \alpha}$ is a martingale constrained to Δ , this implies $Y_{t_k}^{(\ell), 0, p, \alpha} = 1$, almost-surely, which contradicts $Y_{t_k}^{0, p, \alpha} \in \Delta$. Then, we can conclude

$$\mathbb{P}[\tau = t_k] = \mathbb{P} \left[Y_{t_k}^{(k), 0, p, \alpha} = 1 \right] = p_k$$

because $Y_0^{(k), 0, p, \alpha} = p_k$ and $Y_t^{(k), 0, p, \alpha}$ is a martingale taking values zero and one at t_k .

Step 2: Let $\tau \in \mathcal{T}$ be a stopping time such that $\tau \sim \mu$. Then define the $[0, 1]^r$ -valued process \bar{Y} as

$$\bar{Y}_t^{(k)} := \mathbb{E} \left[1_{\{\tau = t_k\}} \mid \mathcal{F}_t \right].$$

Note that $\bar{Y}_0 = p$. By the Martingale Representation Theorem, there exists a control $\alpha \in \mathcal{A}$ for which $Y_t^{0, p, \alpha} = \bar{Y}_t$, almost-surely, for all $t \geq 0$. We can then check that,

$$Y_t^{(1), 0, p, \alpha} + \dots + Y_t^{(r), 0, p, \alpha} = \mathbb{E} \left[1_{\{\tau = t_1\}} + \dots + 1_{\{\tau = t_r\}} \mid \mathcal{F}_t \right] = 1,$$

so $Y_t^{0, p, \alpha} \in \Delta$ for all $t \geq 0$, almost-surely. Finally, for any $k \in \{1, \dots, r\}$, we have $Y_{t_k}^{(k), 0, p, \alpha} = 1_{\{\tau = t_k\}} \in \{0, 1\}$ because $\{\tau = t_k\}$ is \mathcal{F}_{t_k} -measurable.

Define a stopping time σ as

$$\sigma := \min_{k \in \{1, \dots, r\}} \left\{ t_k \mid Y_{t_k}^{(k), 0, p, \alpha} = 1 \right\}$$

and suppose that there exists a set A of non-zero probability on which $\tau \neq \sigma$. Then, for some $k, \ell \in \{1, \dots, r\}$ such that $k \neq \ell$, the set $B := A \cap \{\tau = t_k\} \cap \{\sigma = t_\ell\}$ has non-zero probability.

Suppose that $\ell < k$. Then $Y_{t_\ell}^{(\ell),0,p,\alpha} = 1$ on B and because $Y^{0,p,\alpha}$ is a martingale constrained to Δ , it follows that $Y_{t_k}^{(\ell),0,p,\alpha} = 1$ on B , and consequently, $Y_{t_k}^{(k),0,p,\alpha} = 1_{\{\tau=t_k\}} = 0$, which contradicts $\tau = t_k$ on B . On the other hand, suppose that $\ell > k$. Then $Y_{t_k}^{(k),0,p,\alpha} \neq 1$ on B , but because $Y_{t_k}^{(k),0,p,\alpha} = 1_{\{\tau=t_k\}}$ this also contradicts $\tau = t_k$ on B . We conclude $\tau = \sigma$, almost-surely. \square

2.3 Solution via iterated stochastic control

It is convenient to define a sequence of sets which will be important in the remainder of the paper. For each $k \in \{1, \dots, r\}$, define

$$\Delta_k := \{(y_1, \dots, y_r) \in \Delta \mid y_\ell = 0 \text{ for each } \ell \in \{1, \dots, k-1\}\} \subseteq \Delta.$$

Note that each set is closed and convex and $\Delta_{k+1} \subset \Delta_k$ for each $k \in \{1, \dots, r-1\}$.

We also introduce subcollections of stopping times and controls with additional independence properties, which will be used later in proofs of the Dynamic Programming Principle. In particular, for any $t \in [0, \infty)$ we define a subcollection of stopping times

$$\mathcal{T}_t := \{\tau \in \mathcal{T} \mid \tau \text{ is independent of } \mathcal{F}_t\} \subset \mathcal{T}$$

and a subcollection of controls

$$\mathcal{A}_t := \{\alpha \in \mathcal{A} \mid \alpha \text{ is independent of } \mathcal{F}_t\} \subset \mathcal{A}.$$

We then define a sequence of iterated distribution-constrained optimal stopping problems.

Definition 1. For each $k \in \{1, \dots, r\}$, define a function $v_k : \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ as

$$v_k(x, y) := \sup_{\tau \in \mathcal{T}_{t_{k-1}}} \mathbb{E} \left[f(X_\tau^{t_{k-1}, x}) \right] \quad (2.5)$$

s.t. $\tau \sim \sum_{\ell=1}^r y_\ell \delta_{t_\ell}.$

Note that $v^* = v_1(x_0, p)$. We emphasize that while each v_k is written as a function depending on an entire tuple $y = (y_1, \dots, y_r) \in \Delta_k$, we have $y_1 = \dots = y_{k-1} = 0$ by the definition of Δ_k .

Our goal is to convert these iterated distribution-constrained optimal stopping problems into iterated state-constrained stochastic control problems.

First, we record a growth and continuity estimate for each v_k .

Proposition 2. There exists $C > 0$, which depends only on f and μ , for which

$$\begin{aligned} |v_k(x, y)| &\leq C(1 + |x|) \\ |v_k(x, y) - v_k(x', y')| &\leq C(|x - x'| + \|y - y'\|^{1/2}) \end{aligned}$$

for each $k \in \{1, \dots, r\}$ and all $(x, y), (x', y') \in \mathbb{R} \times \Delta_k$.

We do not specify the choice of norm in y because it only affects the choice of constant.

Proof. Recall that f is assumed to be Lipschitz-continuous. Let $\tau \in \mathcal{T}_{t_{k-1}}$ be an arbitrary stopping time such that $\tau \sim \sum_{k=1}^r y_k \delta_{t_k}$ (such a stopping time exists by Corollary 1). Then, we have

$$\begin{aligned} \left| \mathbb{E} \left[f(X_\tau^{t_{k-1}, x}) \right] \right| &\leq \mathbb{E} \left[\left| f(X_\tau^{t_{k-1}, x}) \right| \right] \\ &\leq |f(0)| + L(|x| + \mathbb{E} [|W_\tau - W_{t_{k-1}}|]) \\ &\leq |f(0)| + L(|x| + 2\mathbb{E} [|W_{t_r}|]) \\ &\leq |f(0)| + L \left(|x| + 2\sqrt{\frac{2}{\pi} t_r} \right). \end{aligned}$$

Next, we compute

$$\begin{aligned} v_k(x', y) &\geq \mathbb{E} \left[f(X_{\tau}^{t_{k-1}, x'}) \right] \\ &\geq \mathbb{E} \left[f(X_{\tau}^{t_{k-1}, x}) \right] - L |x - x'|. \end{aligned}$$

Lastly, by applying Proposition 1, we can construct a stopping time $\tau' \in \mathcal{T}_{t_{k-1}}$ such that $\tau' \sim \sum_{k=1}^r y_k \delta_{t_k}$ and

$$\mathbb{P} [\tau \neq \tau'] \leq 4^r \|y - y'\|.$$

Then we can compute

$$\begin{aligned} v_k(x, y') &\geq \mathbb{E} \left[f(X_{\tau'}^{t_{k-1}, x}) \right] \\ &\geq \mathbb{E} \left[f(X_{\tau}^{t_{k-1}, x}) \right] - C \mathbb{E} \left[1_{\tau \neq \tau'} \left| X_{\tau'}^{t_{k-1}, x} - X_{\tau}^{t_{k-1}, x} \right| \right] \\ &\geq \mathbb{E} \left[f(X_{\tau}^{t_{k-1}, x}) \right] - 2C \mathbb{E} [1_{\tau \neq \tau'} |W_{t_r}|] \\ &\geq \mathbb{E} \left[f(X_{\tau}^{t_{k-1}, x}) \right] - 2C 2^r t_r^{1/2} \|y - y'\|_{\ell^1}^{1/2}. \end{aligned}$$

Then the stated results hold because $(x, y) \in \mathbb{R} \times \Delta_k$ and τ were both arbitrary. \square

In the remainder of the paper, it will prove useful to consider a type of perspective map on the sets Δ_k . For each $k \in \{1, \dots, r\}$, define $P_k : \Delta_k \rightarrow \Delta_k$ as

$$P_k(y_1, \dots, y_r) := \begin{cases} (y_1, \dots, y_r) & \text{if } y_k = 1 \\ (y_{k+1} + \dots + y_r)^{-1} (0, \dots, 0, y_{k+1}, \dots, y_r) & \text{if } y_k < 1. \end{cases} \quad (2.6)$$

We note three key properties of this map.

1. For any $y \in \Delta_k \setminus \{e_k\}$, we have $P_k(y) \in \Delta_{k+1}$,
2. For any $y \in \Delta_k$, the k th coordinate of $P_k(y)$ is either zero or one, and
3. The map P_k is continuous on $\Delta_k \setminus \{e_k\}$.

We now provide a dynamic programming lemma whose proof has the same flavour of the weak dynamic programming results in [7, 6, 3]. Compared to these previous results, we have a priori continuity of the value functions on the right-hand-side, so we do not need to consider upper- and lower-semicontinuous envelopes. We extend the ideas of a countable covering of the state-space by balls, each associated with a nearly optimal stopping time. To deal with the state-constraints, we employ an argument that utilizes the compactness and convexity of Δ_k along with the continuity of v_{k+1} . The proof of this lemma is the heart of the paper, but is quite involved, so it is relegated to the appendix.

Lemma 3 (Dynamic Programming). *For every $k \in \{1, \dots, r-1\}$ and every $(x, y) \in \mathbb{R} \times \Delta_k$, we have*

$$\begin{aligned} v_k(x, y) &= \sup_{\alpha \in \mathcal{A}_{t_{k-1}}} \mathbb{E} \left[Y_{t_k}^{(k), t_{k-1}, y, \alpha} f(X_{t_k}^{t_{k-1}, x}) + (1 - Y_{t_k}^{(k), t_{k-1}, y, \alpha}) v_{k+1} \left(X_{t_k}^{t_{k-1}, x}, Y_{t_k}^{t_{k-1}, y, \alpha} \right) \right] \\ &\quad \text{s.t. } Y_u^{t_{k-1}, y, \alpha} \in \Delta_k \text{ for all } u \geq t_{k-1} \\ &\quad Y_{t_k}^{(k), t_{k-1}, y, \alpha} \in \{0, 1\}, \text{ almost-surely.} \end{aligned} \quad (2.7)$$

Proof. See Appendix A. \square

Next, we demonstrate that we may relax the terminal constraint. The proof of this idea relies on a careful construction of a perturbed martingale which satisfies the terminal constraints of the previous problem, but does not significantly change the expected payoff. The proof of this result shares many key ideas with the previous lemma. For the sake of exposition, we provide this proof in the appendix as well.

Lemma 4 (Constraint Relaxation). *For every $k \in \{1, \dots, r-1\}$ and every $(x, y) \in \mathbb{R} \times \Delta_k$, we have*

$$v_k(x, y) = \sup_{\alpha \in \mathcal{A}_{t_{k-1}}} \mathbb{E} \left[Y_{t_k}^{(k), t_{k-1}, y, \alpha} f(X_{t_k}^{t_{k-1}, x}) + (1 - Y_{t_k}^{(k), t_{k-1}, y, \alpha}) v_{k+1} \left(X_{t_k}^{t_{k-1}, x}, P_k(Y_{t_k}^{t_{k-1}, y, \alpha}) \right) \right] \\ \text{s.t. } Y_u^{t_{k-1}, y, \alpha} \in \Delta_k \text{ for all } u \geq t_{k-1}, \text{ almost-surely,} \quad (2.8)$$

where $P_k : \Delta_k \rightarrow \Delta_k$ is the perspective map defined in (2.6).

Note, even though $P_k(e_k) \notin \Delta_{k+1}$, the right-hand-side of (2.8) is well-defined because v_{k+1} is known to be bounded and continuous. Then, there is a unique continuous extension of the map $(x, y) \mapsto (1 - y_k)v_{k+1}(x, y)$ from $\Delta_k \setminus \{e_k\}$ to Δ_k . That is, taking the right-hand-side to be zero when $y = e_k$.

Proof. See Appendix B. □

With these lemmas in hand, we can now state the main result of this paper.

Theorem 1. *The function $v_r : \mathbb{R} \times \Delta_r \rightarrow \mathbb{R}$ satisfies*

$$v_r(x, y) = \mathbb{E} \left[f(X_{t_r}^{t_{r-1}, x}) \right]$$

for every $(x, y) \in \mathbb{R} \times \Delta_r$.

For each $k \in \{1, \dots, r-1\}$, the function $v_k : \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ is the value function of the following state-constrained stochastic control problem

$$v_k(x, y) = \sup_{\alpha \in \mathcal{A}_{t_{k-1}}} \mathbb{E} \left[Y_{t_k}^{(k), t_{k-1}, y, \alpha} f(X_{t_k}^{t_{k-1}, x}) + (1 - Y_{t_k}^{(k), t_{k-1}, y, \alpha}) v_{k+1} \left(X_{t_k}^{t_{k-1}, x}, P_k(Y_{t_k}^{t_{k-1}, y, \alpha}) \right) \right] \\ \text{s.t. } Y_u^{t_{k-1}, y, \alpha} \in \Delta_k \text{ for all } u \geq t_{k-1}, \text{ almost-surely,}$$

where $P_k : \Delta_k \rightarrow \Delta_k$ is defined as in (2.6).

Of course, we then have

$$v^* = v_1(x_0, p_1, \dots, p_r).$$

Proof. It is clear that v_r has the representation above because there is only one admissible stopping time. The rest follows immediately from applying Lemma 3 and Lemma 4. □

2.4 Time-dependent value functions

While, for the purposes of this paper, we may consider the results of Theorem 1 as a solution to the distribution-constrained optimal stopping problem, we can consider an additional time-dependent version of the state-constrained problem which is amenable to numerical resolution. In particular, the time-dependent value functions will correspond to viscosity solutions of Hamilton-Jacobi-Bellman (HJB) equations.

Definition 2. Define a function $w_r : [t_{r-1}, t_r] \times \mathbb{R} \times \Delta_r \rightarrow \mathbb{R}$ as

$$w_r(t, x, y) := \mathbb{E} \left[f(X_{t_r}^{t,x}) \right].$$

For each $k \in \{1, \dots, r-1\}$, define a function $w_k : [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ as

$$w_k(t, x, y) := \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1} \left(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y,\alpha}) \right) \right]$$

s.t. $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \geq t$, almost-surely,

where $P_k : \Delta_k \rightarrow \Delta_k$ is defined as in (2.6).

Remark 1. Using the properties of the time-independent auxiliary value functions v_k deduced in Section 2.3, we can consider the terminal payoff as a given Hölder continuous function. Then this is a true state-constrained optimal stochastic control problem. In particular, we note that by the argument of Remark 5.2 in [7] we can equivalently define w_k as a supremum over controls in \mathcal{A} which are not necessarily independent of \mathcal{F}_t .

We note an immediate relationship with the value functions of Section 2.3.

Corollary 3. For each $k \in \{1, \dots, r\}$ we have

$$v_k(x, y) = w_k(t_{k-1}, x, y)$$

for all $(x, y) \in \mathbb{R} \times \Delta_k$.

Proof. This result is obvious from the definition of w_k and Theorem 1. □

Before stating a Dynamic Programming Principle for the time-dependent value functions, we first investigate their regularity. In particular, we aim to demonstrate that each w_k is concave in y and jointly continuous with estimates on its modulus of continuity.

Proposition 3. For each $k \in \{1, \dots, r\}$, the function $w_k(t, x, \cdot)$ is concave for each (t, x) .

Proof. We proceed by backwards induction. Notice that the set Δ_r is a singleton, so the functions w_r and v_r are both trivially concave in y .

Suppose that v_{k+1} is concave in y for some $k \in \{1, \dots, r-1\}$. The key observation is that the map

$$\begin{aligned} \Delta_k \setminus \{e_k\} \ni y &\mapsto (1 - y_k) v_{k+1}(x, P_k(y)) \\ &= (y_{k+1} + \dots + y_r) v_{k+1} \left(x, \frac{(0, \dots, 0, y_{k+1}, \dots, y_r)}{y_{k+1} + \dots + y_r} \right) \end{aligned}$$

is concave for every $x \in \mathbb{R}$ because it is the perspective transformation of the concave map $\Delta_{k+1} \ni y \mapsto v_{k+1}(x, y)$ (See Section 3.2.6 in [8]).

With this in mind fix any $(t, x) \in [t_{k-1}, t_k] \times \mathbb{R}$, $y_1, y_2 \in \Delta_k$, and $\lambda \in [0, 1]$. Let $\alpha_1, \alpha_2 \in \mathcal{A}$ be arbitrary controls for which

$$Y_u^{t,y_1,\alpha_1}, Y_u^{t,y_2,\alpha_2} \in \Delta_k,$$

almost-surely, for all $u \geq t$. Define $\bar{y} := \lambda y_1 + (1 - \lambda) y_2$ and $\bar{\alpha}_u := \lambda \alpha_{1,u} + (1 - \lambda) \alpha_{2,u}$. Then $\bar{\alpha} \in \mathcal{A}$ and

$$Y_u^{t,\bar{y},\bar{\alpha}} \in \Delta_k,$$

almost-surely, for all $u \geq t$ by the convexity of the set Δ_k .

Then using the concavity of the perspective map, we can compute

$$\begin{aligned}
w_k(t, x, \bar{y}) &\geq \mathbb{E} \left[Y_{t_k}^{(k), t, \bar{y}, \bar{\alpha}} f(X_{t_k}^{t, x}) + (1 - Y_{t_k}^{(k), t, \bar{y}, \bar{\alpha}}) v_{k+1} \left(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, \bar{y}, \bar{\alpha}}) \right) \right] \\
&\geq \mathbb{E} \left[Y_{t_k}^{(k), t, \bar{y}, \bar{\alpha}} f(X_{t_k}^{t, x}) + \lambda (1 - Y_{t_k}^{(k), t, y_1, \alpha_1}) v_{k+1} \left(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, y_1, \alpha_1}) \right) \right. \\
&\quad \left. + (1 - \lambda) (1 - Y_{t_k}^{(k), t, y_2, \alpha_2}) v_{k+1} \left(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, y_2, \alpha_2}) \right) \right] \\
&= \lambda \mathbb{E} \left[Y_{t_k}^{(k), t, y_1, \alpha_1} f(X_{t_k}^{t, x}) + (1 - Y_{t_k}^{(k), t, y_1, \alpha_1}) v_{k+1} \left(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, y_1, \alpha_1}) \right) \right] \\
&\quad + (1 - \lambda) \mathbb{E} \left[Y_{t_k}^{(k), t, y_2, \alpha_2} f(X_{t_k}^{t, x}) + (1 - Y_{t_k}^{(k), t, y_2, \alpha_2}) v_{k+1} \left(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, y_2, \alpha_2}) \right) \right].
\end{aligned}$$

But because α_1, α_2 were arbitrary, we conclude

$$w_k(t, x, \bar{y}) \geq \lambda w_k(t, x, y_1) + (1 - \lambda) w_k(t, x, y_2).$$

Then w_k is concave in y , and hence so is v_k by Corollary 3. Then the result holds by induction. \square

We can go a step further and obtain a detailed estimate of the joint continuity of w_k .

Proposition 4. *There exists $C > 0$, which depend only on f and μ , such that for each $k \in \{1, \dots, r-1\}$, we have*

$$|w_k(t, x, y) - w_k(t', x', y')| \leq C \left(|t - t'|^{1/4} + |x - x'| + (1 + |x| + |x'|) \|y - y'\|_{\ell^2}^{1/4} \right)$$

for all $(t, x, y), (t', x', y') \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$.

The proof of this statement is relatively straightforward but long-winded, so we relegate it to the appendix. We do not claim that these Hölder exponents are sharp.

Proof. See Appendix C. \square

The upside of this representation as an optimal stochastic control problem is that we can characterize each time-dependent value function w_k as a viscosity solution of a corresponding HJB equation. At this point, we can prove a Dynamic Programming Principle for the time-dependent value functions. While these are state-constrained stochastic control problems, we can directly use the a priori continuity of w_k in y and convexity of Δ_k as in the proof of Lemma 3.

Theorem 2. *Fix $k \in \{1, \dots, r-1\}$, $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$, and any $h > 0$ such that $t+h < t_k$. Let $\{\tau^\alpha\}_{\alpha \in \mathcal{A}_t}$ be a family of stopping times independent of \mathcal{F}_t and valued in $[t, t+h]$. Then*

$$\begin{aligned}
w_k(t, x, y) &= \sup_{\alpha \in \mathcal{A}_t} \mathbb{E} \left[w_k(\tau^\alpha, X_{\tau^\alpha}^{t, x}, Y_{\tau^\alpha}^{t, y, \alpha}) \right] \\
&\quad \text{s.t. } Y_u^{t, y, \alpha} \in \Delta_k \text{ for all } u \geq t, \text{ almost-surely.}
\end{aligned}$$

Proof. See Appendix D. \square

From this result, we can immediately verify that each time-dependent value function is a viscosity solution of an HJB. Once we have the Dynamic Programming Principle in hand, this result becomes reasonably standard, so we direct the interested reader to [16, 6, 24].

We first define elliptic operators $F_k, G_k : \mathbb{S}^{1+r} \times \Delta_k \rightarrow \mathbb{R}$ as

$$\begin{aligned} F_k(R, y) &:= \sup \left\{ \left(\begin{array}{c} 1 \\ a \end{array} \right)^\top R \left(\begin{array}{c} 1 \\ a \end{array} \right) \mid a \in \mathbb{A}_k(y) \right\} \\ G_k(R, y) &:= \sup \left\{ \left(\begin{array}{c} 0 \\ a \end{array} \right)^\top R \left(\begin{array}{c} 0 \\ a \end{array} \right) \mid a \in \mathbb{A}_k(y), \|a\|_{\ell^2} = 1 \right\}, \end{aligned}$$

where

$$\mathbb{A}_k(y) := \{a \in \mathbb{R}^r \mid \exists \epsilon > 0 \text{ s.t. } y + a(-\epsilon, \epsilon) \subset \Delta_k\}.$$

The intuition behind these definitions is that $\mathbb{A}_k(y)$ encodes admissible directions in which a state-constrained martingale starting from $y \in \Delta_k$ may evolve. In particular, a martingale constrained to lie in Δ_k cannot have non-zero quadratic variation in the outer normal direction on the boundary. The elliptic operator F_k shows up naturally from applying Dynamic Programming, while G_k encodes the concavity in y .

The main properties of F_k and G_k are that

$$\begin{aligned} F_k(R, y) < +\infty &\implies G_k(R, y) \leq 0 \\ G_k(R, y) < 0 &\implies F_k(R, y) < +\infty. \end{aligned}$$

Then following the arguments of [2], we can deduce that the value function is a viscosity solution of an equation involving an envelope with G_k .

Proposition 5. *The function w_r is the unique solution of the heat equation (in reversed time),*

$$\begin{cases} u_t + \frac{1}{2}u_{xx} = 0 & \text{in } [t_{r-1}, t_r) \times \mathbb{R} \times \Delta_r \\ u = f & \text{on } \{t = t_r\} \times \mathbb{R} \times \Delta_r. \end{cases}$$

For each $k \in \{1, \dots, r-1\}$, w_k is a continuous viscosity solution of the HJB equation,

$$\begin{cases} \min \{u_t + F_k(D_{xy}^2 u, y), G_k(D_{xy}^2 u, y)\} = 0 & \text{in } [t_{k-1}, t_k) \times \mathbb{R} \times \Delta_k \\ u = y_k f(x) + (1 - y_k) v_{k+1}(x, P_k(y)) & \text{on } \{t = t_k\} \times \mathbb{R} \times \Delta_k. \end{cases} \quad (2.9)$$

The proof of this statement follows from a standard argument and an additional analysis of admissible controls on the boundaries. For more details on the introduction of the operator G_k to obtain a variational inequality, we refer the interested reader to [2] or Section 4 in [23].

The more important question, of course, is whether or not one can obtain a uniqueness result for viscosity solutions of (2.9). It is standard to show that (2.9) admits a comparison principle when we have Dirichlet conditions on the boundary of Δ_k . In the following, we demonstrate that we can prove uniqueness even with second-order boundary conditions using the special structure of the domain.

Theorem 3. *There is a unique continuous viscosity solution of (2.9) which satisfies*

$$|u(t, x, y)| \leq C(1 + |x|)$$

for some $C > 0$.

Of course the time-dependent value functions satisfy this linear growth constraint as a corollary of Proposition 4.

The key idea in this proof is that when we restrict a viscosity solution of (2.9) to the relative interior of any face F of Δ_k in the y -coordinate, the restricted function is a viscosity solution of the same equation on a smaller state-space. In particular, when restricted to a vertex the equation reduces to the heat equation, for which we immediately have uniqueness. We then apply the comparison principle to the equation restricted to an edge using the fact that we have uniqueness on the vertices to deduce uniqueness on edges. We proceed as such on higher dimensional faces until we prove uniqueness on all of Δ_k .

Sketch of Proof. Fix $k \in \{1, \dots, r-1\}$ and let $u, v : [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ be two continuous viscosity solutions of (2.9). Suppose that $u > v$ at some point. By the terminal condition, this must occur at some $t_0 < t_k$.

Suppose that there exists a vertex y_0 of the simplex Δ_k such that $u \neq v$ at some point when restricting to y_0 in the y -coordinate. By Proposition 6.9 in [26], we note that both u and v are viscosity solutions of the heat equation when restricting the vertex y_0 in the y -coordinate. But this contradicts uniqueness for the heat equation.

Let F be a minimal dimension face of the simplex Δ_k such that $u \neq v$ at some point when restricting to F in the y -coordinate. Again, by Proposition 6.9 in [26], we conclude that both u and v are viscosity solutions of the same equation on the relative interior of F . Of course, the boundary of F is the union of lower dimension faces of the simplex Δ_k , so, by the assumed minimal dimension property of F , we conclude that $u = v$ when restricting to the boundary of F in the y -coordinate. Then by applying the comparison principle to (2.9) with Dirichlet boundaries, we deduce that $u = v$ on F .

The theorem then follows by considering the k -dimensional face, Δ_k itself. □

3 Application to Superhedging with a Volatility Outlook

In this section, we consider an application of distribution-constrained optimal stopping in mathematical finance. In particular, we consider the problem of model-free superhedging of a contingent claim with payoff $f(X_T)$ using only dynamic trading in an underlying asset X .

We assume that the price process X_t is a martingale under some unknown martingale measure \mathbb{Q} , but do not specify the exact volatility dynamics. However, in this problem, we assume that we have an outlook on the volatility in the form of the distribution of the quadratic variation, $\langle X \rangle_T$.²

3.1 Model-free superhedging

We follow the model-free setting of [14, 5]. Let $\Omega := \{\omega \in C([0, T], \mathbb{R}) \mid \omega_0 = 0\}$ be the canonical space equipped with uniform norm $\|\omega\|_\infty := \sup_{0 \leq t \leq T} |\omega_t|$, B the canonical process, \mathbb{Q}_0 the Wiener measure, $\mathbb{F} := \{\mathcal{F}_t\}_{0 \leq t \leq T}$ the filtration generated by B , and $\mathbb{F}^+ := \{\mathcal{F}_t^+\}_{0 \leq t \leq T}$ the right-limit of \mathbb{F} .

Fix some initial value $x_0 \in \mathbb{R}$. Then, we denote

$$X_t := x_0 + B_t.$$

For any real-valued, \mathbb{F} -progressively measurable process α satisfying $\int_0^T \alpha_s^2 ds < \infty$, \mathbb{Q}_0 -a.s., we define the probability measure on (Ω, \mathcal{F}) ,

$$\mathbb{Q}^\alpha := \mathbb{Q}_0 \circ (X^\alpha)^{-1},$$

where

$$X_t^\alpha := x_0 + \int_0^t \alpha_r dB_r.$$

Then X^α is a \mathbb{Q}^α -local martingale. We denote by \mathcal{Q} the collection of all such probability measures \mathbb{Q} on (Ω, \mathcal{F}) under which X is a \mathbb{Q} -uniformly integrable martingale. The quadratic

²We note that, while it may seem unlikely that we have an atomic measure representing our volatility outlook, this is a reasonable starting place for two reasons. It is possible to approximate more general measures by atomic measures because it is possible to prove continuity of the value function in the Wasserstein topology (See Lemma 3.1 in [10]). Second, pricing by allowing only a finite number of scenarios, as opposed to specifying a full continuous-valued model, is standard in industry (e.g., the specification of rates, default, and prepayment scenarios in standard models for securitized products).

variation process $\langle X \rangle = \langle B \rangle$ is universally defined under any $\mathbb{Q} \in \mathcal{Q}$, and takes values in the set of all non-decreasing continuous functions from \mathbb{R}_+ to \mathbb{R}_+ .

Let μ be a given probability distribution of the form (2.1). Then we consider the problem

$$\begin{aligned} \bar{U} := & \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}^{\mathbb{Q}} [f(X_T)] \\ & \text{s.t. } \langle X \rangle_T \sim \mu, \end{aligned}$$

where \mathcal{Q} is a collection of admissible martingale measures. This corresponds to a model-free superhedging price in a sense made clear by the duality results in, for example, [5].

3.2 Equivalence to distribution-constrained optimal stopping

We show that this problem is equivalent to distribution-constrained optimal stopping of Brownian motion.

Proposition 6. *We have*

$$\begin{aligned} \bar{U} := & \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}^{\mathbb{Q}} [f(X_T)] = \sup_{\tau \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}_0} [f(X_\tau)]. \\ & \text{s.t. } \langle X \rangle_T \sim \mu \quad \text{s.t. } \tau \sim \mu. \end{aligned}$$

Proof. This argument can be found in Theorem 2.4 of [5]. For completeness, we reproduce it below.

Let $\mathbb{Q} \in \mathcal{Q}$ such that the \mathbb{Q} -distribution of $\langle X \rangle_T$ is μ . It follows by the Dambis-Dubins-Schwarz Theorem that $X_T = x + \tilde{W}_{\langle X \rangle_T}$ where \tilde{W} is a standard Brownian motion and $\tau := \langle X \rangle_T$ is a stopping time with respect to the time-changed filtration with distribution μ (See Theorem 4.6 in [15]). Then $\bar{U} \leq \sup_{\tau \sim \mu} \mathbb{E}^{\mathbb{Q}_0} [f(X_\tau)]$.

Let τ be a stopping time such that $\tau \sim \mu$. Define a process X^τ as

$$X_t^\tau := x + B_{\tau \wedge \frac{t}{T-t}}.$$

Note that X^τ is a continuous martingale on $[0, T]$ with $\langle X^\tau \rangle_T = \tau$, so X^τ induces a probability measure $\mathbb{Q} \in \mathcal{Q}$ such that $\langle X^\tau \rangle_T = \tau \sim \mu$. Then the opposite inequality holds. \square

Then one can obtain a model-free super-hedging price with a volatility outlook by solving the iterated stochastic control problem in Section 2.3.

3.3 Numerical example

In this section we obtain approximate numerical solutions of the distribution-constrained optimal stopping problem using finite-difference schemes.

In particular, we consider two potential outlooks on volatility. In the first, the binary outlook, we assume equal probability between a high- and low-volatility scenario

$$\mu_2 := \frac{1}{2} \delta_{10} + \frac{1}{2} \delta_{20}.$$

In the second, we augment the binary outlook with a third extreme volatility scenario which occurs with small probability

$$\mu_3 := \frac{9}{20} \delta_{10} + \frac{9}{20} \delta_{20} + \frac{1}{10} \delta_{100}.$$

Our goal is to compute the model-free superhedging price of a European call option under each volatility outlook. Because we do not restrict to models where the price process is non-negative, we can take the payoff to be $f(x) := x^+$ without loss of generality.

Then, as before, we define value functions for each outlook as

$$v_2(x) := \sup_{\tau \in \mathcal{T}(\mu_2)} \mathbb{E}^x [f(W_\tau)] \text{ and } v_3(x) := \sup_{\tau \in \mathcal{T}(\mu_3)} \mathbb{E}^x [f(W_\tau)].$$

We solve the problem using the iterated stochastic control approach from Section 2.3. In particular, we obtain a viscosity solution of the corresponding Hamilton-Jacobi-Bellman equation in Section 2.4 using a finite-difference scheme. It is important to emphasize that, because of potential degeneracy due to the extra state-variables in w_2 and w_3 , it is critical to use a monotone numerical scheme.

In these results, we apply a version of the wide-stencil scheme introduced in [19]. In particular, we approximate the non-linear terms in each equation by monotone finite-difference approximations of the following form

$$\sup_{\alpha \in \mathbb{R}} \left[\begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top \begin{pmatrix} u_{xx} & u_{xy} \\ u_{xy} & u_{yy} \end{pmatrix} \begin{pmatrix} 1 \\ \alpha \end{pmatrix} \right] \approx \max_{k \in \mathcal{K}(t,x,y)} \frac{u(x+h, t, y+k) - 2u(x, t, y) + u(x-h, t, y-k)}{h^2},$$

where the set $\mathcal{K}(t, x, y)$ is a collection such that $y \pm k$ lies on nearby grid-points. For a rigorous analysis of wide-stencil schemes for degenerate elliptic equations, we refer the reader to [20, 13, 19].

For comparison, we consider two main special cases, which we refer to as the ‘‘mean volatility’’ value and the ‘‘support-constrained’’ value. We define the mean volatility value as the model-free superhedging price obtained by assuming the quadratic variation will be equal to the mean of the distribution in the corresponding distribution-constrained problem. We define their corresponding value functions as \underline{v}_2 and \underline{v}_3 , respectively. On the other hand, we define the support-constrained value as the model-free superhedging price obtained when only restricting the quadratic variation to have the same support as that of the distribution in the corresponding distribution-constrained problem. We define their corresponding value functions as \bar{v}_2 and \bar{v}_3 , respectively.

We expect the following ordering:

$$f(x) \leq \underline{v}_2(x) \leq v_2(x) \leq \bar{v}_2(x)$$

and

$$f(x) \leq \underline{v}_3(x) \leq v_3(x) \leq \bar{v}_3(x).$$

Furthermore, we note that we can compute \underline{v}_2 , \bar{v}_2 , \underline{v}_3 , and \bar{v}_3 explicitly in terms of heat kernels (See Section 2.3 in [12]).

We illustrate the value function for the two- and three-atom problem in Figure 1 and Figure 2, respectively. As expected, we see a superhedging value which is increasing in the underlying asset price (or, equivalently, decreasing in the strike price) and respects the bounds implied by the support-constrained and average-volatility models. As expected, the bound provided by the support-constrained superhedging problem is particularly poor in the three-model volatility outlook, where we stipulate that the high volatility (high value) case is rare.

It is interesting to note that careful comparison of the two figures illustrates an increase in superhedging value between the two volatility outlooks which is roughly proportional to the increase in square-root of expected quadratic variation. For example, there is approximately a 25% increase in value at $x = 0$, which is essentially exactly in-line with the 25.2% increase in square-root of expected quadratic variation between the two outlooks. This matches our intuition that call option superhedging prices should be proportional to expected volatility to first order.

In Figure 3, we provide a probability density estimate of W_{10} conditional on $\tau = 10$ and $\tau = 20$ for an approximate optimal stopping time for the two-atom volatility outlook model starting from $W_0 = 0$. We obtain these estimates by performing Monte Carlo simulations




Figure 1: Comparison of the model-free superhedging values corresponding to distribution constraints on quadratic variation (v_2), support constraints on quadratic variation (\bar{v}_2), and averaged quadratic variation (\underline{v}_2). Each of these is in the two-atom (binary) volatility outlook. The distribution-constrained value corresponds with the value function of an optimal stopping problem under a two-atom distribution constraint.

with controls estimated from a numerical solution of the associated HJB equations. We use grid spacings $dx = 0.1$, $dy = 0.005$, and $dt = 0.01$. We perform 10^7 simulations and verify that relevant statistics from the Monte Carlo simulation match those from the finite-difference solutions (e.g., expected payoff, distribution and moments of the stopping time and stopped process) to within a reasonable margin of error.

The density estimates provide insight into the form of the optimal strategy. Recall that the payoff is locally-affine at all points except $x = 0$, where it is strictly convex instead. Then we expect an optimal stopping strategy to be one which maximizes local time accumulated at the origin. As expected, we find that the density of W_{10} conditional upon $\tau = 10$ is largely concentrated on points away from $x = 0$, at which the payoff process is unlikely to spend significant time as a submartingale if we were to choose not to stop.

It is interesting to note the lack of sharp cut-off between the two density estimates. One might expect the optimal strategy to be of a form where there exists a “stopping region” and a “continuation region.” On the contrary, the smooth overlap of the two density estimates is persistent even as we vary the resolution of the finite-difference solver, which suggests that the true optimal stopping strategy is not of the form $\{\tau = 10\} \subset \sigma(W_{10})$. That is, the numerics suggest that optimal stopping strategies may be path-dependent even in simple examples.

A Proof of Lemma 3

This first argument is in the spirit of proofs of the weak dynamic programming principle which avoid measurable selection, as in [7, 6, 3]. In these arguments, the authors typically use a covering argument to find a countable selection of ϵ -optimal controls on small balls of the state-space. The main difficulty here is that, while a control may be admissible for the state-constrained problem at one point in state-space, there is no reason to expect it to satisfy the state constraints starting from nearby states.

The new idea in our approach is to cover Δ_{k+1} with a finite mesh. We show that we can




Figure 2: Comparison of the model-free superhedging values corresponding to distribution constraints on quadratic variation (v_3), support constraints on quadratic variation (\bar{v}_3), and averaged quadratic variation (\underline{v}_3). Each of these is in the three-atom (trinary) volatility outlook. The distribution-constrained value corresponds with the value function of an optimal stopping problem under a three-atom distribution constraint.

replace the process Y by a modified process Y^ϵ , which lies on the mesh points almost-surely at the terminal time. We construct the new process in a measurable way using the Martingale Representation Theorem on a carefully constructed random variable. Then we show that, using the continuity of v_{k+1} , that the objective function along Y is close to that along Y^ϵ for a fine enough grid.

Once we know we can consider a perturbed process Y^ϵ which lies on a finite number of points in Δ_{k+1} at the terminal time almost-surely, we can construct ϵ -optimal stopping times using a standard covering argument in \mathbb{R} .

Proof. Fix $(x, y) \in \mathbb{R} \times \Delta_k$. For convenience of notation, define

$$A := \sup_{\alpha \in \mathcal{A}_{t_{k-1}}} \mathbb{E} \left[Y_{t_k}^{(k), t_{k-1}, y, \alpha} f(X_{t_k}^{t_{k-1}, x}) + (1 - Y_{t_k}^{(k), t_{k-1}, y, \alpha}) v_{k+1} \left(X_{t_k}^{t_{k-1}, x}, Y_{t_k}^{t_{k-1}, y, \alpha} \right) \right]$$

s.t. $Y_u^{t_{k-1}, y, \alpha} \in \Delta_k$ for all $u \geq t_{k-1}$
 $Y_{t_k}^{(k), t_{k-1}, y, \alpha} \in \{0, 1\}$, almost-surely.

In the remainder of this proof, we do not write t_{k-1} in the superscripts of X and Y because it is always fixed.

Step 1: Fix an arbitrary $\epsilon > 0$. Choose $R > 0$ large enough that

$$\mathbb{P} \left[\sup_{t_{k-1} \leq u \leq t_k} |W_u - W_{t_{k-1}}| \geq R \right] \leq \epsilon^2.$$

Because v_{k+1} is continuous on the compact set $[x - R, x + R] \times \Delta_{k+1}$, we can find $\delta > 0$ small enough that

$$|v_{k+1}(x', y') - v_{k+1}(x', y'')| \leq \epsilon$$




Figure 3: Probability density estimates of W_{10} conditional on $\tau = 10$ and $\tau = 20$ for an optimal stopping time for the two-atom volatility outlook model starting from $W_0 = 0$. Density estimates were made by Monte Carlo simulations on high-resolution solutions to the associated HJB equations. Sample size, $N = 10^7$.

for all $x' \in [x - R, x + R]$ and $y', y'' \in \Delta_{k+1}$ such that

$$\|y' - y''\|_{\ell^\infty} \leq \delta.$$

Similarly, because f is Lipschitz and v_{k+1} is Lipschitz in x uniformly in y , we can find $\delta > 0$, possibly smaller than before, such that we also have

$$|f(x') - f(x'')| + |v_{k+1}(x', y') - v_{k+1}(x'', y')| \leq \epsilon$$

for all $x', x'' \in \mathbb{R}$ and $y' \in \Delta_{k+1}$ such that

$$|x' - x''| \leq \delta.$$

Step 2: We now construct a finite mesh on Δ_{k+1} . Let $\mathcal{P} := \{y_j\}_{j=1}^N$ be a finite subset of Δ_{k+1} with the property that

- The convex hull of \mathcal{P} is Δ_{k+1} , and
- Any point $y \in \Delta_{k+1}$ can be written as a convex combination of finitely-many points in \mathcal{P} , each contained in a δ -neighborhood of y .

This is possible by compactness and convexity of Δ_{k+1} . In particular, we can define a continuous function $T : \Delta_{k+1} \rightarrow [0, 1]^N$ with the properties that

- $T_j(y) = 0$ for all $y \in \Delta_{k+1}$ such that $|y - y_j| > \delta$
- $\sum_{j=1}^N T_j(y) = 1$ for all $y \in \Delta_{k+1}$, and
- $\sum_{j=1}^N y_j T_j(y) = y$ for all $y \in \Delta_{k+1}$.

This corresponds to a continuous map from a point $y \in \Delta_{k+1}$ to a probability weighting of points in \mathcal{P} such that y is a convex combination of nearby points in \mathcal{P} . Such a map can be obtained by an ℓ^2 -minimization problem, for instance.

Step 3: Let $\{A_i\}_{i \geq 1}$ be a countable and disjoint covering of \mathbb{R} with an associated set of points $\{x_i\}$ such that the δ -ball centered at x_i contains the set A_i .

For each $i \geq 1$ and $j \in \{1, \dots, N\}$, let $\tau_{i,j} \in \mathcal{T}_{t_k}$ be a stopping time satisfying

$$\tau_{i,j} \sim \sum_{\ell=1}^r y_j^{(\ell)} \delta_{t_\ell}$$

such that

$$\mathbb{E} \left[f(X_{\tau_{i,j}}^{t_k, x_i}) \right] \geq v_{k+1}(x_i, y_j) - \epsilon.$$

Note that the above uses $y_j^{(\ell)}$ to denote the ℓ th entry of the vector y_j .

By choice of $\delta > 0$ in the first step and the definition of the sets A_i , we have

$$\begin{aligned} v_{k+1}(x_i, y_j) &\geq v_{k+1}(x, y_j) - \epsilon \\ \mathbb{E} \left[f(X_{\tau_{i,j}}^{t_k, x_i}) \right] &\geq \mathbb{E} \left[f(X_{\tau_{i,j}}^{t_k, x}) \right] - \epsilon \end{aligned}$$

for all $x \in A_i$.

Putting these inequalities together, we conclude that

$$\mathbb{E} \left[f(X_{\tau_{i,j}}^{t_k, x}) \right] \geq v_{k+1}(x, y_j) - 3\epsilon$$

for all $i \geq 1$, $j \in \{1, \dots, N\}$, and $x \in A_i$.

Step 4: Let $\alpha \in \mathcal{A}_{t_{k-1}}$ be an arbitrary control for which $Y_u^{y, \alpha} \in \Delta_k$ for $u \geq t_{k-1}$ and $Y_{t_k}^{(k), y, \alpha} \in \{0, 1\}$ almost-surely. For any $0 < h < t_k - t_{k-1}$, define two random variables, M_1 and M_2 , as

$$\begin{aligned} M_1 &:= h^{-1/2} (W_{t_k} - W_{t_k-h}) \\ M_2 &:= h^{-1/2} \max_{t_k-h \leq s \leq t_k} |W_s - W_{t_k-h} - \delta^{-1} (s - t_k + h) (W_{t_k} - W_{t_k-h})|. \end{aligned} \tag{A.1}$$

Then M_1 and M_2 are \mathcal{F}_{t_k} -measurable and independent of each other. M_1 is equal in distribution to a standard normal distribution, the cumulative distribution function of which we denote by Φ . Similarly, M_2 is equal in distribution to the absolute maximum of a standard Brownian bridge on $[0, 1]$, the cumulative distribution function of which we denote by Φ_{BB} . Furthermore, if we define

$$\mathcal{G} := \sigma(\mathcal{F}_{t_k-h} \cup \sigma(W_{t_k})),$$

then M_1 is \mathcal{G} -measurable, while M_2 is independent of \mathcal{G} .

Define a random vector \bar{Y}_{t_k} as

$$\bar{Y}_{t_k}^{(k)} := 1_{\{M_2 \leq \Phi_{BB}^{-1}(Y_{t_k-h}^{(k), y, \alpha})\}}$$

and

$$\begin{aligned} \bar{Y}_{t_k}^{(k+1):r} &:= 1_{\{M_2 > \Phi_{BB}^{-1}(Y_{t_k-h}^{(k), y, \alpha})\}} \times \\ &\quad \sum_{j=1}^N y_j 1_{\{\Phi^{-1}(\sum_{i=1}^{j-1} T_i(P_k(Y_{t_k-h}^{y, \alpha}))) < M_1 \leq \Phi^{-1}(\sum_{i=1}^j T_i(P_k(Y_{t_k-h}^{y, \alpha})))\}}, \end{aligned}$$

where we follow the conventions that $\Phi^{-1}(0) = -\infty$, $\Phi^{-1}(1) = +\infty$, and that sums over an empty set are zero. We denote the $(k+1)$ th through r th entry in the random vector by $\bar{Y}_{t_k}^{(k+1):r}$. Then $\bar{Y}_{t_k} \in \Delta_k$ is \mathcal{F}_{t_k} -measurable and is constructed to have the key property that

$$\mathbb{E} [\bar{Y}_{t_k} \mid \mathcal{F}_{t_k-h}] = Y_{t_k-h}^{y, \alpha},$$

almost-surely.

By the Martingale Representation Theorem, there exists $\alpha_\epsilon \in \mathcal{A}$ for which $Y_{t_k}^{y, \alpha_\epsilon} = \bar{Y}_{t_k}$ almost-surely. It is clear by the construction that \bar{Y}_{t_k} is independent of $\mathcal{F}_{t_{k-1}}$, so we can take $\alpha_\epsilon \in \mathcal{A}_{t_{k-1}}$. Then, by construction, $Y_u^{y, \alpha_\epsilon} \in \Delta_k$ for all $u \geq t_{k-1}$, $Y_{t_k}^{(k), y, \alpha_\epsilon} \in \{0, 1\}$, and $Y_{t_k}^{y, \alpha_\epsilon} \in \mathcal{P}$ when $Y_{t_k}^{(k), y, \alpha_\epsilon} = 0$, almost-surely.

We now perform a key computation. First note that

$$\begin{aligned} & \mathbb{E} \left[Y_{t_k}^{(k), y, \alpha_\epsilon} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha_\epsilon}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y, \alpha_\epsilon}) \right] \\ &= \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 1\}} f(X_{t_k}^x) \right] + \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, \bar{Y}_{t_k}) \right]. \end{aligned}$$

For the first term on the right-hand-side, we simply compute

$$\begin{aligned} & \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 1\}} f(X_{t_k}^x) \right] = \mathbb{E} \left[1_{\{M_2 \leq \Phi_{BB}^{-1}(Y_{t_k-h}^{(k), y, \alpha})\}} f(X_{t_k}^x) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[1_{\{M_2 \leq \Phi_{BB}^{-1}(Y_{t_k-h}^{(k), y, \alpha})\}} f(X_{t_k}^x) \mid \mathcal{G} \right] \right] \\ &= \mathbb{E} \left[Y_{t_k-h}^{(k), y, \alpha} f(X_{t_k}^x) \right]. \end{aligned}$$

We deal with the second term in a similar way, but the computation is more involved. Note that by construction we have

$$\|\bar{Y}_{t_k} - P_k(Y_{t_k-h}^{t, y, \alpha})\|_{\ell^\infty} \leq \delta$$

almost-surely in the set $\{\bar{Y}_\theta^{(k)} = 0\}$. Recall we also took δ small enough such that

$$|v_{k+1}(x', y') - v_{k+1}(x', y'')| \leq \epsilon$$

for all $x' \in [x-R, x+R]$ and $y', y'' \in \Delta_{k+1}$ such that $\|y' - y''\|_{\ell^\infty} \leq \delta$. But then we can compute

$$\begin{aligned} & \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, \bar{Y}_{t_k}) \right] \\ &= \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} 1_{\{|W_{t_k}| \leq R\}} v_{k+1}(X_{t_k}^x, \bar{Y}_{t_k}) \right] + \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} 1_{\{|W_{t_k}| \geq R\}} v_{k+1}(X_{t_k}^x, \bar{Y}_{t_k}) \right] \\ &\geq \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} 1_{\{|W_{t_k}| \leq R\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right] \\ &\quad + \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} 1_{\{|W_{t_k}| \geq R\}} v_{k+1}(X_{t_k}^x, \bar{Y}_{t_k}) \right] - \epsilon \\ &\geq \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right] \\ &\quad - \mathbb{E} \left[1_{\{|W_{t_k}| \geq R\}} \left(|v_{k+1}(X_{t_k}^x, \bar{Y}_{t_k})| + |v_{k+1}(X_{t_k}^x, Y_{t_k-h}^{y, \alpha})| \right) \right] - \epsilon \\ &\geq \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right] - 2\sqrt{\mathbb{P}[|W_{t_k}| \geq R]} C(1 + |x|) - \epsilon \\ &\geq \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right] - 2(1 + C)(1 + |x|)\epsilon, \end{aligned}$$

where $C > 0$ comes from the growth bound from Proposition 2. With this in hand, we now complete the analysis of the second term

$$\begin{aligned} & \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right] = \mathbb{E} \left[1_{\{M_2 > \Phi_{BB}^{-1}(Y_{t_k-h}^{(k), y, \alpha})\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[1_{\{M_2 > \Phi_{BB}^{-1}(Y_{t_k-h}^{(k), y, \alpha})\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \mid \mathcal{G} \right] \right] \\ &= \mathbb{E} \left[(1 - Y_{t_k-h}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y, \alpha})) \right]. \end{aligned}$$

Using the continuity of f , v_{k+1} , and P_k , along with the Dominated Convergence Theorem, we note

$$\begin{aligned} \lim_{h \rightarrow 0} \mathbb{E} \left[Y_{t_k-h}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k-h}^{(k),y,\alpha}) v_{k+1} \left(X_{t_k}^x, P_k(Y_{t_k-h}^{y,\alpha}) \right) \right] \\ = \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha}) v_{k+1} \left(X_{t_k}^x, P_k(Y_{t_k}^{y,\alpha}) \right) \right] \\ = \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha}) v_{k+1} \left(X_{t_k}^x, Y_{t_k}^{y,\alpha} \right) \right]. \end{aligned}$$

Then putting these results together, we see that for $h > 0$ small enough

$$\begin{aligned} \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha\epsilon} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha\epsilon}) v_{k+1} \left(X_{t_k}^x, Y_{t_k}^{y,\alpha\epsilon} \right) \right] \\ \geq \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha}) v_{k+1} \left(X_{t_k}^x, Y_{t_k}^{y,\alpha} \right) \right] - \epsilon - 2(1+C)(1+|x|)\epsilon. \end{aligned}$$

Step 5: Lastly, we intend to construct an ϵ -optimal stopping time using the covering from the second step. Define a stopping time τ_ϵ as

$$\tau_\epsilon := t_k + 1_{\{Y_{t_k}^{(k),t,y,\alpha\epsilon} = 0\}} \sum_{i=1}^{\infty} \sum_{j=1}^N \tau_{i,j} 1_{\{X_{t_k}^x \in A_i\}} 1_{\{Y_{t_k}^{y,\alpha\epsilon} = y_j\}}.$$

By construction, we have $\tau_\epsilon \sim \sum_{\ell=1}^r y_\ell \delta_{t_\ell}$. We proceed to make a careful computation. First, note that

$$\mathbb{E} [f(X_{\tau_\epsilon}^x)] = \mathbb{E} [1_{\{\tau_\epsilon = t_k\}} f(X_{t_k}^x)] + \mathbb{E} [1_{\{\tau_\epsilon > t_k\}} f(X_{\tau_\epsilon}^x)].$$

We focus on the second term. In particular, we have

$$\begin{aligned} \mathbb{E} [1_{\{\tau_\epsilon > t_k\}} f(X_{\tau_\epsilon}^x)] &= \sum_{i=1}^{\infty} \sum_{j=1}^N \mathbb{E} \left[1_{\{Y_{t_k}^{(k),y,\alpha\epsilon} = 0\}} 1_{\{X_{t_k}^x \in A_i\}} 1_{\{Y_{t_k}^{y,\alpha\epsilon} = y_j\}} f(X_{\tau_{i,j}}^x) \right] \\ &= \sum_{i=1}^{\infty} \sum_{j=1}^N \mathbb{E} \left[1_{\{Y_{t_k}^{(k),y,\alpha\epsilon} = 0\}} 1_{\{X_{t_k}^x \in A_i\}} 1_{\{Y_{t_k}^{y,\alpha\epsilon} = y_j\}} \mathbb{E} [f(X_{\tau_{i,j}}^x) \mid \mathcal{F}_{t_k}] \right] \\ &\geq \sum_{i=1}^{\infty} \sum_{j=1}^N \mathbb{E} \left[1_{\{Y_{t_k}^{(k),y,\alpha\epsilon} = 0\}} 1_{\{X_{t_k}^x \in A_i\}} 1_{\{Y_{t_k}^{y,\alpha\epsilon} = y_j\}} v_{k+1}(X_{t_k}^x, Y_{t_k}^{y,\alpha\epsilon}) \right] - 3\epsilon \\ &= \mathbb{E} \left[(1 - Y_{t_k}^{(k),y,\alpha\epsilon}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y,\alpha\epsilon}) \right] - 3\epsilon, \end{aligned}$$

where the inequality follows from the construction in the third step and the independence of the stopping times τ_{ij} with respect to \mathcal{F}_{t_k} .

Then we conclude

$$\mathbb{E} [f(X_{\tau_\epsilon}^x)] \geq \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha\epsilon} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha\epsilon}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y,\alpha\epsilon}) \right] - 3\epsilon.$$

Combining this with the main inequality from the previous step, we obtain

$$\begin{aligned} v_k(x, y) &\geq \mathbb{E} [f(X_{\tau_\epsilon}^x)] \\ &\geq \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha\epsilon} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha\epsilon}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y,\alpha\epsilon}) \right] - 3\epsilon \\ &\geq \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y,\alpha}) \right] - 4\epsilon - 2(1+C)(1+|x|)\epsilon. \end{aligned}$$

Because ϵ and α were arbitrary, then we conclude $A \leq v_k(x, y)$.

Step 6: Let $\tau \in \mathcal{T}_{t_k}$ be an arbitrary stopping time such that $\tau \sim \sum_{\ell=1}^r y_\ell \delta_{t_\ell}$. Define a martingale as

$$Y_t^{(i)} := \mathbb{E} [1_{\{\tau = t_i\}} \mid \mathcal{F}_t]$$

for all $t \geq 0$ and each $i \in \{1, \dots, r\}$. We can easily check that $Y_0^{(i)} = y_i$ for each $i \in \{1, \dots, r\}$ and

$$Y_t^{(1)} + \dots + Y_t^{(r)} = \mathbb{E} [1_{\{\tau=t_1\}} + \dots + 1_{\{\tau=t_r\}} \mid \mathcal{F}_t] = 1.$$

Then if we consider Y as an \mathbb{R}^r -valued martingale with $Y_t^{(i)} \equiv 0$ for all $i \in \{1, \dots, k-1\}$, then we see $Y_t \in \Delta_k$ for each $t \geq 0$. Finally, we have

$$Y_{t_k}^{(k)} = \mathbb{E} [1_{\{\tau=t_k\}} \mid \mathcal{F}_{t_k}] = 1_{\{\tau=t_k\}} \in \{0, 1\}.$$

Then by the Martingale Representation Theorem, there exists $\alpha \in \mathcal{A}_{t_k}$ for which $Y_t^{y, \alpha} = Y_t$ for all $t \geq 0$, almost-surely. We can compute

$$\begin{aligned} \mathbb{E} [f(X_\tau^x)] &= \mathbb{E} [1_{\{\tau=t_k\}} f(X_{t_k}^x) + 1_{\{\tau>t_k\}} f(X_\tau^x)] \\ &= \mathbb{E} [Y_{t_k}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha}) \mathbb{E} [f(X_\tau^x) \mid \mathcal{F}_{t_k}]]. \end{aligned}$$

On the set $\{\tau > t_k\}$, we have

$$\mathbb{P} [\tau = t_i \mid \mathcal{F}_{t_k}] = \mathbb{E} [1_{\{\tau=t_i\}} \mid \mathcal{F}_{t_k}] = Y_{t_k}^{(i)}$$

for each $i \in \{k+1, \dots, r\}$. For almost every $\omega \in \{\tau > t_k\}$, we have

$$\mathbb{E} [f(X_\tau^x) \mid \mathcal{F}_{t_k}] \leq v_{k+1}(X_{t_k}^x, Y_{t_k}^{y, \alpha})$$

by the Strong Markov Property and stationarity properties of Brownian motion. Then we conclude

$$\mathbb{E} [f(X_\tau^x)] \leq \mathbb{E} [Y_{t_k}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y, \alpha})] \leq A.$$

Because τ was an arbitrary stopping time, this implies

$$v_k(x, y) \leq A$$

□

B Proof of Lemma 4

The main idea of this argument is that we can take a controlled process Y , which does not satisfy $Y_{t_k}^{(k)} \in \{0, 1\}$, and modify it on an interval $[t_k - h, t_k]$ to a perturbed process Y^ϵ with the properties that $Y_{t_k-h} = Y_{t_k-h}^\epsilon$ and $Y_{t_k}^{\epsilon, (k)} \in \{0, 1\}$. In particular, we may do this in a way that does not appreciably change the expected payoff.

One key idea, which we draw the reader's attention towards, is the use of the Brownian bridge over $[t_k - h, t_k]$ in the construction. This construction is in the spirit of Corollary 2. While one might initially attempt a construction similar to Corollary 1, using a Brownian bridge instead of Brownian increments allows us to condition on W_{t_k} at a key point in the argument.

Proof. Fix $(x, y) \in \mathbb{R} \times \Delta_k$. For convenience of notation, define

$$\begin{aligned} A &:= \sup_{\alpha \in \mathcal{A}_{t_{k-1}}} \mathbb{E} \left[Y_{t_k}^{(k), t_{k-1}, y, \alpha} f(X_{t_k}^{t_{k-1}, x}) + (1 - Y_{t_k}^{(k), t_{k-1}, y, \alpha}) v_{k+1}(X_{t_k}^{t_{k-1}, x}, Y_{t_k}^{t_{k-1}, y, \alpha}) \right] \\ &\quad \text{s.t.} \quad Y_u^{t_{k-1}, y, \alpha} \in \Delta_k \text{ for } u \geq t_{k-1} \\ &\quad Y_{t_k}^{(k), t_{k-1}, y, \alpha} \in \{0, 1\} \text{ almost-surely,} \end{aligned}$$

and

$$\begin{aligned} B &:= \sup_{\alpha \in \mathcal{A}_{t_{k-1}}} \mathbb{E} \left[Y_{t_k}^{(k), t_{k-1}, y, \alpha} f(X_{t_k}^{t_{k-1}, x}) + (1 - Y_{t_k}^{(k), t_{k-1}, y, \alpha}) v_{k+1}(X_{t_k}^{t_{k-1}, x}, P_k(Y_{t_k}^{t_{k-1}, y, \alpha})) \right] \\ &\quad \text{s.t.} \quad Y_u^{t_{k-1}, y, \alpha} \in \Delta_k \text{ for } u \geq t_{k-1}. \end{aligned}$$

By Lemma 3, we have $v_k(x, y) = A$. In the remainder of the proof we withhold the superscript t_{k-1} on X and Y for the sake of brevity.

Step 1: Let $\alpha \in \mathcal{A}_{t_{k-1}}$ be an arbitrary control for which $Y_u^{y, \alpha} \in \Delta_k$ for $u \geq t_{k-1}$ and $Y_{t_k}^{(k), y, \alpha} \in \{0, 1\}$ almost-surely. Note that $Y_{t_k}^{y, \alpha} = P_k(Y_{t_k}^{y, \alpha})$ on the set $\{Y_{t_k}^{(k), y, \alpha} = 0\}$, almost-surely. Then

$$\begin{aligned} & \mathbb{E} \left[Y_{t_k}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y, \alpha}) \right] \\ &= \mathbb{E} \left[Y_{t_k}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k}^{y, \alpha})) \right] \\ &\leq B. \end{aligned}$$

Because α was arbitrary, we conclude $A \leq B$.

Step 2: Let $\alpha \in \mathcal{A}_{t_{k-1}}$ be an arbitrary control for which $Y_u^{y, \alpha} \in \Delta_k$ for $u \geq t_{k-1}$, almost-surely. For any $0 < h < t_k - t_{k-1}$, define a random variable M as

$$M := h^{-1/2} \max_{t_k - h \leq s \leq t_k} |W_s - W_{t_k - h} - h^{-1}(s - t_k + h)(W_{t_k} - W_{t_k - h})|.$$

Then M is \mathcal{F}_{t_k} -measurable and is equal in distribution to the absolute maximum of a standard Brownian bridge on $[0, 1]$, the cumulative distribution function of which we denote by Φ_{BB} . If we define $\mathcal{G} := \sigma(\mathcal{F}_{t_k - h} \cup \sigma(W_{t_k}))$, then M is independent of \mathcal{G} .

Define a random vector \bar{Y}_{t_k} as

$$\bar{Y}_{t_k}^{(k)} := 1_{\{M \leq \Phi_{BB}^{-1}(Y_{t_k - h}^{(k), y, \alpha})\}}$$

and

$$\bar{Y}_{t_k}^{(k+1):r} := P_k(Y_{t_k - h}^{y, \alpha}) 1_{\{M > \Phi_{BB}^{-1}(Y_{t_k - h}^{(k), y, \alpha})\}},$$

where $\bar{Y}_{t_k}^{(k+1):r}$ denotes the $(k+1)$ th through r th element in the vector. Let $\bar{Y}_{t_k}^{(i)} \equiv 0$ for any $i \in \{1, \dots, k-1\}$. Then \bar{Y}_{t_k} is \mathcal{F}_{t_k} -measurable and has the key property that $\mathbb{E}[\bar{Y}_{t_k} | \mathcal{F}_{t_k - h}] = Y_{t_k - h}^{y, \alpha}$. We also note that $\mathbb{E}\left[1_{\{\bar{Y}_{t_k}^{(k)} = 1\}} | \mathcal{G}\right] = Y_{t_k - h}^{(k), y, \alpha}$.

By the Martingale Representation Theorem, there exists $\alpha_\epsilon \in \mathcal{A}_{t_{k-1}}$ such that $Y_u^{y, \alpha_\epsilon} \in \Delta_k$ for $u \geq t_{k-1}$, $Y_{t_k}^{(k), y, \alpha_\epsilon} \in \{0, 1\}$, and $Y_{t_k}^{y, \alpha_\epsilon} = \bar{Y}_{t_k}$ almost-surely. We can then compute

$$\begin{aligned} & \mathbb{E} \left[Y_{t_k}^{(k), y, \alpha_\epsilon} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha_\epsilon}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y, \alpha_\epsilon}) \right] \\ &= \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 1\}} f(X_{t_k}^x) + 1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} v_{k+1}(X_{t_k}^x, P_k(Y_{t_k - h}^{y, \alpha})) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 1\}} | \mathcal{G} \right] f(X_{t_k}^x) + \mathbb{E} \left[1_{\{\bar{Y}_{t_k}^{(k)} = 0\}} | \mathcal{G} \right] v_{k+1}(X_{t_k}^x, P_k(Y_{t_k - h}^{y, \alpha})) \right] \\ &= \mathbb{E} \left[Y_{t_k - h}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k - h}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k - h}^{y, \alpha})) \right]. \end{aligned}$$

But by the continuity and growth bounds of f and v_{k+1} , we can apply the Dominated Convergence Theorem to see

$$\begin{aligned} & \lim_{\delta \rightarrow 0^+} \mathbb{E} \left[Y_{t_k - h}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k - h}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k - h}^{y, \alpha})) \right] \\ &= \mathbb{E} \left[Y_{t_k}^{(k), y, \alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k), y, \alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k}^{y, \alpha})) \right]. \end{aligned}$$

So then for any $\epsilon > 0$, we may take $h > 0$ small enough that

$$\begin{aligned} & \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k}^{y,\alpha})) \right] \\ & \leq \mathbb{E} \left[Y_{t_k-h}^{(k),y,\alpha} f(X_{t_k}^x) + (1 - Y_{t_k-h}^{(k),y,\alpha}) v_{k+1}(X_{t_k}^x, P_k(Y_{t_k-h}^{y,\alpha})) \right] + \epsilon \\ & = \mathbb{E} \left[Y_{t_k}^{(k),y,\alpha\epsilon} f(X_{t_k}^x) + (1 - Y_{t_k}^{(k),y,\alpha\epsilon}) v_{k+1}(X_{t_k}^x, Y_{t_k}^{y,\alpha\epsilon}) \right] + \epsilon \\ & \leq A + \epsilon. \end{aligned}$$

Because ϵ and α were arbitrary, we conclude $B \leq A$. \square

C Proof of Proposition 4

Proof. We proceed in several steps, each relating the value function between nearby points. In the first three points, we consider a shift backwards in the time variable, a shift forward in time variable away from the terminal time, and lastly a jump onto the terminal time. In the fourth step, we discuss arbitrary perturbations in x . In the fifth step, we discuss a perturbation inside the interior of some face of Δ_k , including a possible jump off the face. In the sixth step, we consider a jump from an interior point onto a face of Δ_k . Lastly, in the final step, we discuss how to put these together into one coherent bound.

Step 1: Fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$ and $t' \in [t_{k-1}, t_k]$ such that $t' \leq t$. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \geq t$, almost-surely. Define a new control $\alpha' \in \mathcal{A}$ as

$$\alpha'_u := 1_{\{u \geq t\}} \alpha_u,$$

for all $u \geq t'$. We can see that $Y_u^{t',y,\alpha'} \in \Delta_k$ for all $u \in [t', t_k]$ and $Y_{t_k}^{t',y,\alpha'} = Y_{t_k}^{t,y,\alpha}$, almost-surely. Then

$$\begin{aligned} w_k(t', x, y) & \geq \mathbb{E} \left[Y_{t_k}^{(k),t',y,\alpha'} f(X_{t_k}^{t',x}) + (1 - Y_{t_k}^{(k),t',y,\alpha'}) v_{k+1}(X_{t_k}^{t',x}, P_k(Y_{t_k}^{t',y,\alpha'})) \right] \\ & = \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t',x}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{t_k}^{t',x}, P_k(Y_{t_k}^{t,y,\alpha})) \right] \\ & \geq \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y,\alpha})) \right] \\ & \quad - 2C \mathbb{E} |W_t - W_{t'}|, \end{aligned}$$

where $C > 0$ is at least as large as the Lipschitz constants in x for f and v_{k+1} . But recall that for Brownian motion we can find $C > 0$ such that

$$\mathbb{E} |W_t - W_{t'}| \leq C |t - t'|^{1/2}.$$

Using this and the fact that α was arbitrary, we then conclude

$$w_k(t, x, y) - w_k(t', x, y) \leq 2C^2 |t - t'|^{1/2}.$$

Step 2: Fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$ and $t' \in [t_{k-1}, t_k]$ such that $t \leq t'$. Define

$$\eta := \sqrt{\frac{t_k - t}{t_k - t'}} \geq 1.$$

Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t,y,\alpha'} \in \Delta_k$ for all $u \geq t$, almost-surely. Define new control $\alpha' \in \mathcal{A}$ as

$$\alpha'_u := \eta \alpha_{\tau_u},$$

where

$$\tau_u := \eta^2(u - t') + t$$

for all $u \in [t', t_k]$. Note that $\alpha'_u \in \mathcal{F}_{\tau_u}$ by definition. Because $\tau_u \leq u$, we then have $\alpha'_u \in \mathcal{F}_u$ so it is an adapted control. We can also check by the time-change properties of the Itô Integral that

$$\left(W_{t_k} - W_{t'}, Y_{t_k}^{t', y, \alpha'} \right) \stackrel{(d)}{=} \left(\eta^{-1}(W_{t_k} - W_t), Y_{t_k}^{t, y, \alpha} \right).$$

Then $Y_u^{t', y, \alpha'} \in \Delta_k$ for all $u \in [t', t_k]$, almost-surely, by the convexity of Δ_k and the martingale property of Y . Then α' is an admissible control.

We can compute

$$\begin{aligned} w_k(t', x, y) &\geq \mathbb{E} \left[Y_{t_k}^{(k), t', y, \alpha'} f(X_{t_k}^{t', x}) + (1 - Y_{t_k}^{(k), t', y, \alpha'}) v_{k+1}(X_{t_k}^{t', x}, P_k(Y_{t_k}^{t', y, \alpha'})) \right] \\ &\geq \mathbb{E} \left[Y_{t_k}^{(k), t, y, \alpha} f(X_{t_k}^{t, x}) + (1 - Y_{t_k}^{(k), t, y, \alpha}) v_{k+1}(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, y, \alpha})) \right] \\ &\quad - 2C(1 - \eta^{-1}) \mathbb{E} |W_{t_k} - W_t|. \end{aligned}$$

Now we proceed to bound the final term in this inequality. First, note that by the convexity of $x \mapsto x^{-1/2}$, we can bound

$$\eta^{-1} = \left(1 + \frac{t' - t}{t_k - t'} \right)^{-1/2} \geq 1 - \frac{t' - t}{2(t_k - t')}.$$

Furthermore, for large enough $C > 0$, depending only upon t_r , we have $\mathbb{E} |W_u| \leq C$ for all $u \in [0, t_r]$. Then we can estimate

$$(1 - \eta^{-1}) \mathbb{E} |W_{t_k} - W_t| \leq 2C \frac{t' - t}{2(t_k - t')}.$$

Putting these together and recalling that α was arbitrary, we conclude

$$w_k(t, x, y) - w_k(t', x, y) \leq 4C^2 \frac{|t - t'|}{t_k - t'}.$$

Step 3: Fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$ and let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t, y, \alpha} \in \Delta_k$ for all $u \geq t$, almost-surely. By the Lipschitz continuity of f and v_{k+1} in x we can compute

$$\begin{aligned} &\mathbb{E} \left[Y_{t_k}^{(k), t, y, \alpha} f(X_{t_k}^{t, x}) + (1 - Y_{t_k}^{(k), t, y, \alpha}) v_{k+1}(X_{t_k}^{t, x}, P_k(Y_{t_k}^{t, y, \alpha})) \right] \\ &\leq \mathbb{E} \left[Y_{t_k}^{(k), t, y, \alpha} f(x) + (1 - Y_{t_k}^{(k), t, y, \alpha}) v_{k+1}(x, P_k(Y_{t_k}^{t, y, \alpha})) \right] + C \mathbb{E} [|W_{t_k} - W_t|] \end{aligned}$$

We can bound the error term by $C |t_k - t|^{1/2}$. By viewing the term containing v_{k+1} as a perspective map applied to a concave function (See the proof of Proposition 3) and noting that the controlled process Y is a martingale, we can apply Jensen's Inequality to see

$$\begin{aligned} \mathbb{E} \left[Y_{t_k}^{(k), t, y, \alpha} f(x) + (1 - Y_{t_k}^{(k), t, y, \alpha}) v_{k+1}(x, P_k(Y_{t_k}^{t, y, \alpha})) \right] &\leq yk f(x) + (1 - yk) v_{k+1}(x, P_k(y)) \\ &\leq w_k(t_k, x, y). \end{aligned}$$

But then because α was arbitrary, we conclude that

$$w_k(t, x, y) - w(t_k, x, y) \leq C |t_k - t|^{1/2}.$$

Step 4: Fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$ and $x' \in \mathbb{R}$. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \geq t$, almost-surely. Then we immediately compute

$$\begin{aligned} w_k(t, x', y) &\geq \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t,x'}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{t_k}^{t,x'}, P_k(Y_{t_k}^{t,y,\alpha})) \right] \\ &\geq \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y,\alpha})) \right] - 2C |x - x'|, \end{aligned}$$

where $C > 0$ is at least as large as the Lipschitz constant in x of f and v_{k+1} . Because α was arbitrary we conclude

$$w_k(t, x', y) \geq w(t, x, y) - 2C |x - x'|.$$

Step 5: Fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$. Suppose for now that y is not a vertex of Δ_k , or equivalently that the value in each coordinate is less than one. Denote by

$$\mathcal{I} := \{i \in \{1, \dots, r\} \mid y_i = 0\}, \quad \mathcal{J} = \{1, \dots, r\} \setminus \mathcal{I}$$

the disjoint collections of coordinates in which y is zero and non-zero, respectively. Let

$$\delta := \min_{i \in \mathcal{J}} \{\min\{|y_i|, |1 - y_i|\}\} > 0.$$

Let $y' \in \Delta_k$ be any point satisfying $\|y - y'\|_{\ell^\infty} \leq \delta^2$. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \geq t$, almost-surely. Note that α equals zero almost surely for each coordinate in \mathcal{I} .

Define a new control $\alpha' \in \mathcal{A}$ as

$$\alpha'_u := (1 - \|y - y'\|_{\ell^\infty}^{1/2}) \alpha_u$$

and note that by construction we have

$$\min(Y_u^{(i),t,y,\alpha'}, 1 - Y_u^{(i),t,y,\alpha'}) \geq \delta \|y - y'\|_{\ell^\infty}^{1/2}$$

for each $i \in \mathcal{J}$. Similarly, we have

$$\min(Y_u^{(i),t,y',\alpha'}, 1 - Y_u^{(i),t,y',\alpha'}) \geq \delta \|y - y'\|_{\ell^2}^{1/2} - \|y - y'\|_{\ell^2} \geq (\delta - \|y - y'\|_{\ell^2}^{1/2}) \|y - y'\|_{\ell^2}^{1/2} \geq 0,$$

for all $u \geq t$, almost-surely. This together with the observation that α' equals zero in each direction in \mathcal{I} implies that $Y_u^{t,y',\alpha'} \in \Delta_k$ for all $u \geq t$ almost-surely. Furthermore, we have

$$\begin{aligned} \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty} &\leq \|y - y'\|_{\ell^\infty} + \left(1 - \sqrt{1 - \|y - y'\|_{\ell^\infty}^{1/2}}\right) \|Y_{t_k}^{t,y,\alpha} - y\|_{\ell^\infty} \\ &\leq \|y - y'\|_{\ell^\infty} + R \|y - y'\|_{\ell^\infty}^{1/2} \\ &\leq (1 + R) \|y - y'\|_{\ell^\infty}^{1/2}, \end{aligned} \tag{C.1}$$

almost-surely, where $R > 0$ is the diameter of the set Δ_k .

Before proceeding with concrete bounds, we note an estimate regarding the perspective map. For any $y, y' \in \Delta_k$ such that $y_k, y'_k \neq 1$, we have

$$\|P_k(y) - P_k(y')\|_{\ell^\infty} \leq \left(\frac{1}{1 - y_k} + \frac{1}{1 - y'_k} \right) \|y - y'\|_{\ell^\infty}.$$

That is, the perspective map fails to be Lipschitz as $y_k, y'_k \rightarrow 1$. Using this along with the Hölder continuity and bounds on v_{k+1} from Proposition 2, we can carefully bound

$$\begin{aligned} &|(1 - y_k)v_{k+1}(x, P_k(y)) - (1 - y'_k)v_{k+1}(x, P_k(y'))| \\ &\leq \min\{1 - y_k, 1 - y'_k\} |v_{k+1}(x, P_k(y)) - v_{k+1}(x, P_k(y'))| \\ &\quad + |y_k - y'_k| (|v_{k+1}(x, P_k(y))| + |v_{k+1}(x, P_k(y'))|) \\ &\leq \sqrt{2}C \|y - y'\|_{\ell^\infty}^{1/2} + 2C(1 + |x|) \|y - y'\|_{\ell^\infty}. \end{aligned}$$

This bound is easily seen in the case $y_k, y'_k \neq 1$ and may be carefully checked when either equals zero exactly.

Then using the bound above as well as the growth bounds on f , we check

$$\begin{aligned} w_k(t, x, y') &\geq \mathbb{E} \left[Y_{t_k}^{(k),t,y',\alpha'} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y',\alpha'}) v_{k+1}(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y',\alpha'})) \right] \\ &\geq \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y,\alpha})) \right] \\ &\quad - \mathbb{E} \left[\sqrt{2}C \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty}^{1/2} + 3C |X_{t_k}^{t,x}| \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty} \right]. \end{aligned}$$

Applying Hölder's Inequality and the almost-sure bound in (C.1), we bound the last term by

$$\mathbb{E} \left[|X_{t_k}^{t,x}| \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty} \right] \leq (x^2 + t_r)^{1/2} (1 + R)^{1/2} \|y - y'\|_{\ell^\infty}^{1/4}.$$

Putting these all together and recalling that α was arbitrary, we conclude

$$w_k(t, x, y) - w_k(t, x, y') \leq C(1 + |x|) \|y - y'\|_{\ell^2}^{1/4},$$

for large enough constant $C > 0$.

If y is a vertex, then the only admissible control is $\alpha \equiv 0$, so we obtain the same bound (in fact a better bound) for any nearby y' directly from Proposition 2.

Step 6: Fix some $\delta > 0$ at least small enough that $\delta < 1/(2r)$. Fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$ such that at least one element of (y_k, \dots, y_r) is in $(0, \delta^2)$. Denote by

$$\begin{aligned} \mathcal{I} &:= \{i \in \{1, \dots, r\} \mid y_i = 0\}, & \mathcal{J} &:= \{i \in \{1, \dots, r\} \mid y_i \in (0, \delta^2)\}, \\ \mathcal{K} &:= \{1, \dots, r\} \setminus (\mathcal{I} \cup \mathcal{J}) \end{aligned}$$

the disjoint collections of coordinates in which y is zero, “small”, and “large”, respectively. Let $y' \in \Delta_k$ be any point obtained from setting elements of y in \mathcal{J} to zero and adding these values to a single index $\kappa \in \mathcal{K}$. Then y' is zero in all coordinates $\mathcal{I} \cup \mathcal{J}$ and non-zero (and “large”) in all coordinates \mathcal{K} . Furthermore, $\|y - y'\|_{\ell^\infty} \geq \delta^2$, so

$$\delta \leq \|y - y'\|_{\ell^\infty}^{1/2}.$$

Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \geq t$, almost-surely. Note that α equals zero almost surely for each coordinate in \mathcal{I} . Similarly, because each component of $Y^{t,y,\alpha}$ is a martingale, we conclude

$$\delta^2 \geq y_i = \mathbb{E} \left[Y_{t_k}^{(i),t,y,\alpha} \right] \geq \delta \mathbb{P} \left[Y_{t_k}^{(i),t,y,\alpha} \right]$$

for each $i \in \mathcal{J}$. That is, $\mathbb{P} \left[Y_{t_k}^{(i),t,y,\alpha} \geq \delta \right] \leq \delta$, so the i th coordinate of $Y^{t,y,\alpha}$ stays small with high probability.

Define $\alpha' \in \mathcal{A}$ by moving the values of α at coordinates $i \in \mathcal{J}$ to the κ th coordinate. Note that, by construction,

$$Y_u^{(\kappa),t,y',\alpha'} = Y_u^{(\kappa),t,y,\alpha} + \sum_{i \in \mathcal{J}} Y_u^{(i),t,y,\alpha} \in [0, 1]$$

and $Y_u^{(\kappa),t,y',\alpha'} \in \Delta_k$ for all $u \geq t$, almost-surely. Furthermore, we have

$$Y_{t_k}^{(i),t,y',\alpha'} = Y_{t_k}^{(i),t,y,\alpha},$$

for each $i \in \mathcal{I} \cup \mathcal{K} \setminus \{\kappa\}$, and

$$\mathbb{P} \left[\left| Y_{t_k}^{(i),t,y',\alpha'} - Y_{t_k}^{(i),t,y,\alpha} \right| \geq \delta \right] \leq \delta$$

for each $i \in \mathcal{J}$. Lastly, we have

$$\mathbb{P} \left[\left| Y_{t_k}^{(\kappa),t,y',\alpha'} - Y_{t_k}^{(\kappa),t,y,\alpha} \right| \geq r\delta \right] \leq \delta.$$

In summary, we have

$$\mathbb{P} \left[\|Y_{t_k}^{t,y',\alpha'} - Y_{t_k}^{t,y,\alpha}\|_{\ell^\infty} \geq 2r\delta \right] \leq \delta.$$

Then using the same bounds as in the previous step, we can now compute

$$\begin{aligned} w_k(t, x, y') &\geq \mathbb{E} \left[Y_{t_k}^{(k),t,y',\alpha'} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y',\alpha'}) v_{k+1}(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y',\alpha})) \right] \\ &\geq \mathbb{E} \left[Y_{t_k}^{(k),t,y,\alpha} f(X_{t_k}^{t,x}) + (1 - Y_{t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{t_k}^{t,x}, P_k(Y_{t_k}^{t,y,\alpha})) \right] \\ &\quad - \mathbb{E} \left[\sqrt{2}C \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty}^{1/2} + 3C |X_{t_k}^{t,x}| \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty} \right]. \end{aligned}$$

The first term on the right-hand-side may be bounded as

$$\begin{aligned} \mathbb{E} \left[\|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty}^{1/2} \right] &\leq \sqrt{2} \mathbb{P} \left[\|Y_{t_k}^{t,y',\alpha'} - Y_{t_k}^{t,y,\alpha}\|_{\ell^\infty} \geq 2r\delta \right] + \sqrt{2r\delta} \\ &\leq \sqrt{2(r+1)}\delta^{1/2}. \end{aligned}$$

Similarly, the second term may be bounded as

$$\begin{aligned} \mathbb{E} \left[|X_{t_k}^{t,x}| \|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty} \right] &\leq \sqrt{|x|^2 + t_r} \mathbb{E} \left[\|Y_{t_k}^{t,y,\alpha} - Y_{t_k}^{t,y',\alpha'}\|_{\ell^\infty}^2 \right]^{1/2} \\ &\leq 2\sqrt{(|x|^2 + t_r)(1+r^2)}\delta. \end{aligned}$$

Putting these all together and recalling that α was arbitrary and $\delta \leq \|y - y'\|_{\ell^\infty}^{1/2}$, we conclude

$$w_k(t, x, y) - w_k(t, x, y') \leq C(1 + |x|)\|y - y'\|_{\ell^2}^{1/4},$$

for large enough constant $C > 0$.

Step 7: We now briefly remark how to put all of these estimates together. We consider the Hölder estimates in each coordinate separately as they can be combined in the end using triangle inequality. Note that the Lipschitz regularity in x has already been proven.

In the time direction, fix $(t, x, y) \in [t_{k-1}, t_k] \times \mathbb{R} \times \Delta_k$ and $t' \in [t_{k-1}, t_k]$ some small $\theta > 0$ such that $\theta < t_k - t_{k-1}$. If $t' \leq t$, then by Step 1 we have

$$w_k(t, x, y) - w_k(t', x, y) \leq 2C^2 |t - t'|^{1/2}.$$

If $t < t' = t_k$, then by Step 3 we have

$$w_k(t, x, y) - w_k(t', x, y) = w_k(t, x, y) - w_k(t_k, x, y) \leq C |t_k - t|^{1/2} = C |t - t'|^{1/2}.$$

Now suppose that $t < t' \leq t_k - |t - t'|^{1/2}$. By Step 2, we have

$$w_k(t, x, y) - w_k(t', x, y) \leq 4C^2 \frac{|t - t'|}{t_k - t'} \leq 4C^2 \frac{|t - t'|}{|t - t'|^{1/2}} = 4C^2 |t - t'|^{1/2}.$$

In the next step, we critically see where the $(1/4)$ -Hölder coefficient appears. If $t_k - |t - t'|^{1/2} \leq t < t' < t_k$ then by an application of Step 1 and Step 3, we see

$$\begin{aligned} w_k(t, x, y) - w_k(t', x, y) &\leq [w_k(t, x, y) - w_k(t_k, x, y)] + [w_k(t_k, x, y) - w_k(t', x, y)] \\ &\leq C|t_k - t| + 4C^2|t_k - t'|^{1/2} \\ &\leq C|t - t'|^{1/2} + 4C^2|t - t'|^{1/4}. \end{aligned}$$

Lastly we consider the case $t \leq t_k - |t - t'|^{1/2} \leq t' < t_k$. By an application of each of Steps 1 through 3, we see

$$\begin{aligned} w_k(t, x, y) - w_k(t', x, y) &\leq [w_k(t, x, y) - w_k(t_k - |t - t'|^{1/2}, x, y)] \\ &\quad + [w_k(t_k - |t - t'|^{1/2}, x, y) - w_k(t_k, x, y)] + [w_k(t_k, x, y) - w_k(t', x, y)] \\ &\leq 4C^2 \frac{|t_k - |t - t'|^{1/2} - t|}{|t - t'|^{1/2}} + C|t - t'|^{1/4} + 4C^2|t_k - t'|^{1/2} \\ &\leq 4C^2 \frac{|t - t'|}{|t - t'|^{1/2}} + C|t - t'|^{1/4} + 4C^2|t - t'|^{1/4} \\ &\leq (C + 8C^2)|t - t'|^{1/4}. \end{aligned}$$

Of course, by taking a large enough constant we can bound all $|t - t'|^{1/2}$ terms by $|t - t'|^{1/4}$ terms and obtain the $(1/4)$ -Hölder continuity result in t .

The $(1/4)$ -Hölder continuity result follows by a similar approach by cases as in the time perturbation case. The key idea is that Step 5 and Step 6 tell locally how to perturb in a $(1/4)$ -Hölder way, including onto and off the boundaries. Then by a covering argument and the compactness of Δ_k we can obtain a finite chain of local Hölder inequalities connecting any two points and obtain the result for sufficiently large constant. \square

D Proof of Theorem 2

This argument is essentially a time-dependent version of that given in the proof of Lemma 3. The key idea here is to use the convexity of the set Δ_k and the concavity of w_k in y to construct an ϵ -suboptimal control which satisfies the state-constraint as a convex combination of admissible controls starting from nearby points.

Proof. Fix $(t, x, y) \in [t_{k-1}, t_k) \times \mathbb{R} \times \Delta_k$ and $0 < h < t_k - t_{k-1}$. For convenience of notation, define $\theta := t_k$ and

$$\begin{aligned} A &:= \sup_{\alpha \in \mathcal{A}_t} \mathbb{E} \left[w_k \left(\tau^\alpha, X_{\tau^\alpha}^{t,x}, Y_{\tau^\alpha}^{t,y,\alpha} \right) \right] \\ &\text{s.t. } Y_u^{t,y,\alpha} \in \Delta_k \text{ for all } u \geq t. \end{aligned}$$

The inequality $w_k(t, x, y) \leq A$ is a standard result even in the case of these state-constraints. We refer the interested reader to Theorem 3.3 in [26] and instead focus on the opposite inequality.

Step 1: Fix an arbitrary $\epsilon > 0$. Choose $R > 0$ large enough that

$$\mathbb{P} \left[\sup_{t \leq u \leq t+h} |W_u - W_t| \geq R \right] \leq \epsilon^2.$$

Because w_k is continuous on the compact set $[t, t+h] \times [x-R, x+R] \times \Delta_k$, we can find $\delta > 0$ small enough that

$$|w_k(t', x', y') - w_k(t', x', y'')| \leq \epsilon$$

for all $(t', x') \in [t, t+h] \times [x-R, x+R]$ and $y', y'' \in \Delta_k$ such that

$$\|y' - y''\|_{\ell^\infty} \leq \delta.$$

Similarly, because f is Lipschitz and v_{k+1} is Lipschitz in x uniformly in y , we can find $\delta > 0$, possibly smaller than before, such that we also have

$$|f(x') - f(x'')| + |v_{k+1}(x', y') - v_{k+1}(x'', y')| \leq \epsilon$$

for all $x', x'' \in \mathbb{R}$ and $y' \in \Delta_{k+1}$ such that

$$|x' - x''| \leq \delta.$$

Finally, take $\delta > 0$ potentially even smaller so that

$$\delta^{1/2} + \delta^{1/4} \leq \epsilon.$$

Step 2: We first construct a finite mesh on $[t, t + h]$ and Δ_k which will be fine enough to take advantage the continuity of w_k . Let $\Lambda := \{t_i\}_{i=1}^M$ be a finite collection of mesh points in $[t, t + h]$ with the key property that for any $u \in [t, t + h]$, there exists $i \in \{1, \dots, M\}$ such that $u \leq t_i \leq u + \delta$.

By the compactness and convexity of Δ_k , we can obtain a finite subset of Δ_k , $\mathcal{P} := \{y_\ell\}_{\ell=1}^P$, with the property that

- The convex hull of \mathcal{P} is Δ_k , and
- Any point $y \in \Delta_k$ can be written as a convex combination of points in \mathcal{P} , each contained in a δ -neighborhood of y .

In particular, we can find a continuous function $T : \Delta_k \rightarrow [0, 1]^P$ with the properties

- $T_\ell(y) = 0$ for all $y \in \Delta_k$ such that $|y - y_\ell| > \delta$
- $\sum_{\ell=1}^P T_\ell(y) = 1$ for all $y \in \Delta_k$, and
- $\sum_{\ell=1}^P y_\ell T_\ell(y) = y$ for all $y \in \Delta_k$.

This corresponds to a continuous map from a point $y \in \Delta_k$ to a probability weighting of points in \mathcal{P} such that y is a convex combination of nearby points in \mathcal{P} .

By the same type of covering argument as in the proof of Lemma 2, we can obtain a finite and disjoint covering of $[x - R, x + R]$ by measurable sets $\{A_j\}_{j=1}^N$, each contained in a δ -ball, and controls $\alpha_{ij\ell} \in \mathcal{A}_{t_i}$ with the key properties that $Y_u^{t_i, y_\ell, \alpha_{ij\ell}} \in \Delta_k$ for all $u \geq t_i$ and

$$\mathbb{E} \left[Y_\theta^{(k), t_i, y_\ell, \alpha_{ij\ell}} f(X_\theta^{t_i, x}) + (1 - Y_\theta^{(k), t_i, y_\ell, \alpha_{ij\ell}}) v_{k+1}(X_\theta^{t_i, x}, P_k(Y_\theta^{t_i, y_\ell, \alpha_{ij\ell}})) \right] \geq w_k(t_i, x, y_\ell) - 3\epsilon$$

for each $i \in \{1, \dots, M\}$, $j \in \{1, \dots, N\}$, $\ell \in \{1, \dots, P\}$, and $x \in A_j$.

Step 3: Fix an arbitrary control $\alpha \in \mathcal{A}$ for which $Y_u^{t, y, \alpha} \in \Delta_k$ for all $u \geq t$ and let τ^α be the associated stopping time which is valued in $[t, t + h]$. We are next going to construct a new control related to the suboptimal controls $\alpha_{ij\ell}$. In words, we will follow α up to the stopping time τ^α , then set the control to zero until the first subsequent hitting time of Λ . Then we will follow an appropriate convex combination of the controls $\alpha_{ij\ell}$.

To make this precise, define a stopping time

$$\bar{\tau} := \inf\{t \geq \tau^\alpha \mid t \in \Lambda\}.$$

Define a collection of controls $\alpha_\ell \in \mathcal{A}$ as

$$\alpha_{\ell, u} := 1_{\{u \in [t, \tau^\alpha]\}} \alpha_u + 1_{\{u > \bar{\tau}\}} \sum_{i=1}^M \sum_{j=1}^N 1_{\{\bar{\tau} = t_i\}} 1_{\{X_{\bar{\tau}}^{t, x} \in A_j\}} \alpha_{ij\ell, u}$$

for each $\ell \in \{1, \dots, P\}$ and all $u \geq t$. Finally, define a control $\bar{\alpha} \in \mathcal{A}$ as

$$\bar{\alpha}_u := 1_{\{u \in [t, \tau^\alpha]\}} \alpha_u + 1_{\{u > \bar{\tau}\}} \sum_{\ell=1}^P T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) \alpha_{\ell,u}$$

for all $u \geq t$.

The control $\bar{\alpha}$ is adapted because the map T is continuous. Similarly, it can be easily seen to be square-integrable. The key property, however, is that $\bar{\alpha}$ satisfies $Y_u^{t,y,\bar{\alpha}} \in \Delta_k$ for all $u \geq t$. In words, this follows from the convexity of the set Δ_k and the fact that $\bar{\alpha}$ is a convex combination of controls, each of which satisfy the state-constraint.

Making this precise, we use the assumed properties of the map T and the dynamics of Y to compute

$$\begin{aligned} Y_\theta^{t,y,\bar{\alpha}} &= y + \int_t^{\tau^\alpha} \alpha_u dW_u + \int_{\bar{\tau}}^\theta \bar{\alpha}_u dW_u \\ &= Y_{\tau^\alpha}^{t,y,\alpha} + \sum_{i,j,\ell} 1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) \int_{\bar{\tau}}^\theta \alpha_{ij\ell,u} dW_u \\ &= 1_{\{|W_{\bar{\tau}} - W_t| \geq R\}} Y_{\tau^\alpha}^{t,y,\alpha} + \sum_{i,j} 1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} \left(Y_{\tau^\alpha}^{t,y,\alpha} + \sum_{\ell} T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) \int_{\bar{\tau}}^\theta \alpha_{ij\ell,u} dW_u \right) \\ &= 1_{\{|W_{\bar{\tau}} - W_t| \geq R\}} Y_{\tau^\alpha}^{t,y,\alpha} + \sum_{i,j} 1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} \sum_{\ell} T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) \left(y_\ell + \int_{\bar{\tau}}^\theta \alpha_{ij\ell,u} dW_u \right) \\ &= 1_{\{|W_{\bar{\tau}} - W_t| \geq R\}} Y_{\tau^\alpha}^{t,y,\alpha} + \sum_{i,j} 1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} \sum_{\ell} T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) Y_\theta^{t_i,y_\ell,\alpha_{ij\ell}} \end{aligned}$$

Recall though that $Y_\theta^{t_i,y_\ell,\alpha_{ij\ell}} \in \Delta_k$ and $Y_{\tau^\alpha}^{t,y,\alpha} \in \Delta_k$ almost-surely. Then the equality above and the convexity of Δ_k demonstrate that $Y_u^{t,y,\bar{\alpha}} \in \Delta_k$ for all $u \geq t$ almost-surely.

Step 4: We now proceed to make a very delicate series of estimates. First, we have

$$\begin{aligned} w_k(t, x, y) &\geq \mathbb{E} \left[Y_\theta^{(k),t,y,\bar{\alpha}} f(X_\theta^{t,x}) + (1 - Y_\theta^{(k),t,y,\bar{\alpha}}) v_{k+1}(X_\theta^{t,x}, P_k(Y_\theta^{t,y,\bar{\alpha}})) \right] \\ &\geq \sum_{i,j} \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} \left(Y_\theta^{(k),t,y,\bar{\alpha}} f(X_\theta^{t,x}) + (1 - Y_\theta^{(k),t,y,\bar{\alpha}}) v_{k+1}(X_\theta^{t,x}, P_k(Y_\theta^{t,y,\bar{\alpha}})) \right) \right] \\ &\quad - C \mathbb{E} \left[1_{\{|W_{\bar{\tau}} - W_t| \geq R\}} \left(1 + |X_{\bar{\tau}}^{t,x}| \right) \right], \end{aligned}$$

where the last term comes from known growth bounds on f and v_{k+1} . Of course, this term is bounded by $\epsilon \sqrt{C(1 + |x|)}$ by the choice of R and use of Hölder's Inequality. Next, rewriting each term in the sum above using the key property from the construction of $\bar{\alpha}$ in the previous step, we see

$$\begin{aligned} &\mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} \left(Y_\theta^{(k),t,y,\bar{\alpha}} f(X_\theta^{t,x}) + (1 - Y_\theta^{(k),t,y,\bar{\alpha}}) v_{k+1}(X_\theta^{t,x}, P_k(Y_\theta^{t,y,\bar{\alpha}})) \right) \right] \\ &= \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} \left(\sum_{\ell=1}^P T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) Y_\theta^{(k),t_i,y_\ell,\alpha_{ij\ell}} f(X_\theta^{t,x}) \right. \right. \\ &\quad \left. \left. + (1 - \sum_{\ell=1}^P T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) Y_\theta^{(k),t_i,y_\ell,\alpha_{ij\ell}}) v_{k+1}(X_\theta^{t,x}, P_k(\sum_{\ell=1}^P T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) Y_\theta^{t_i,y_\ell,\alpha_{ij\ell}})) \right) \right] \\ &\geq \sum_{\ell=1}^P \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} T_\ell(Y_{\tau^\alpha}^{t,y,\alpha}) \left(Y_\theta^{(k),t_i,y_\ell,\alpha_{ij\ell}} f(X_\theta^{t,x}) \right. \right. \\ &\quad \left. \left. + (1 - Y_\theta^{(k),t_i,y_\ell,\alpha_{ij\ell}}) v_{k+1}(X_\theta^{t,x}, P_k(Y_\theta^{t_i,y_\ell,\alpha_{ij\ell}})) \right) \right] \end{aligned}$$

using the concavity of v_{k+1} composed with the perspective function in y (See the proof of Proposition 3). Next, by the suboptimality conditions of α_{ijk} , we see

$$\begin{aligned} & \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} T_{\ell}(Y_{\tau^{\alpha}}^{t,y,\alpha}) \left(Y_{\theta}^{(k),t_i,y_{\ell},\alpha_{ij\ell}} f(X_{\theta}^{t,x}) + (1 - Y_{\theta}^{(k),t_i,y_{\ell},\alpha_{ij\ell}}) v_{k+1}(X_{\theta}^{t,x}, P_k(Y_{\theta}^{t_i,y_{\ell},\alpha_{ij\ell}})) \right) \right] \\ & \geq \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} T_{\ell}(Y_{\tau^{\alpha}}^{t,y,\alpha}) w_k(t_i, X_{\bar{\tau}}^{t,x}, y_{\ell}) \right] - 3\epsilon \\ & \geq \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} T_{\ell}(Y_{\tau^{\alpha}}^{t,y,\alpha}) w_k(t_i, X_{\bar{\tau}}^{t,x}, Y_{\tau^{\alpha}}^{t,y,\alpha}) \right] - 4\epsilon, \end{aligned}$$

where we used the locality property of the map T and continuity of w_k assumed in the construction of \mathcal{P} . Lastly, summing over i, j, ℓ , we see

$$\begin{aligned} & \sum_{i,j,\ell} \mathbb{E} \left[1_{\{\bar{\tau}=t_i\}} 1_{\{X_{\bar{\tau}}^{t,x} \in A_j\}} T_{\ell}(Y_{\tau^{\alpha}}^{t,y,\alpha}) w_k(t_i, X_{\bar{\tau}}^{t,x}, Y_{\tau^{\alpha}}^{t,y,\alpha}) \right] \\ & \geq \mathbb{E} \left[w_k(\bar{\tau}, X_{\bar{\tau}}^{t,x}, Y_{\tau^{\alpha}}^{t,y,\alpha}) \right] - \mathbb{E} \left[1_{\{|W_{\bar{\tau}}| \geq R\}} w_k(\bar{\tau}, X_{\bar{\tau}}^{t,x}, Y_{\tau^{\alpha}}^{t,y,\alpha}) \right] \\ & \geq \mathbb{E} \left[w_k(\tau^{\alpha}, X_{\tau^{\alpha}}^{t,x}, Y_{\tau^{\alpha}}^{t,y,\alpha}) \right] - C \mathbb{E} \left[|\bar{\tau} - \tau^{\alpha}|^{1/4} + \left| X_{\bar{\tau}}^{t,x} - X_{\tau^{\alpha}}^{t,x} \right| + 1_{\{|W_{\bar{\tau}}| \geq R\}} \left(1 + \left| X_{\bar{\tau}}^{t,x} \right| \right) \right] \\ & \geq \mathbb{E} \left[w_k(\tau^{\alpha}, X_{\tau^{\alpha}}^{t,x}, Y_{\tau^{\alpha}}^{t,y,\alpha}) \right] - C (\delta^{1/4} + \delta^{1/2} + \epsilon (1 + |x|)) \end{aligned}$$

for sufficiently large $C > 0$. In this step, we used growth bounds on w_k and the Hölder estimates from Proposition 4 together with the fact that $|\bar{\tau} - \tau^{\alpha}| \leq \delta$ by construction. By the choice of δ , however, we see this last error term is bounded by $C(2 + |x|)\epsilon$.

Putting all these computations together and recalling that α and $\epsilon > 0$ were arbitrary, we see

$$w_k(t, x, y) \geq A.$$

□

References

- [1] S. ANKIRCHNER, M. KLEIN, AND T. KRUSE, *A verification theorem for optimal stopping problems with expectation constraints*, Applied Mathematics & Optimization, (2017).
- [2] E. BAYRAKTAR AND M. SIRBU, *Stochastic perron's method for hamilton-jacobi-bellman equations*, SIAM Journal on Control and Optimization, 51 (2013), pp. 4274–4294.
- [3] E. BAYRAKTAR AND S. YAO, *A weak dynamic programming principle for zero-sum stochastic differential games with unbounded controls*, SIAM J. Control Optim., 51 (2013), pp. 2036–2080.
- [4] M. BEIGLBOECK, M. EDER, C. ELGERT, AND U. SCHMOCK, *Geometry of Distribution-Constrained Optimal Stopping Problems*, ArXiv e-prints, (2016).
- [5] J. F. BONNANS AND X. TAN, *A model-free no-arbitrage price bound for variance options*, Appl. Math. Optim., 68 (2013), pp. 43–73.
- [6] B. BOUCHARD AND M. NUTZ, *Weak dynamic programming for generalized state constraints*, SIAM J. Control Optim., 50 (2012), pp. 3344–3373.
- [7] B. BOUCHARD AND N. TOUZI, *Weak dynamic programming principle for viscosity solutions*, SIAM J. Control Optim., 49 (2011), pp. 948–962.
- [8] S. BOYD AND L. VANDENBERGHE, *Convex optimization*, Cambridge University Press, Cambridge, 2004.
- [9] R. CAPOCELLI AND L. RICCIARDI, *On the inverse of the first passage time probability problem*, Journal of Applied Probability, (1972), pp. 270–287.
- [10] A. M. G. COX AND S. KÄLLBLAD, *Model-independent bounds for Asian options: a dynamic programming approach*, To appear in SIAM Journal on Control and Optimization, (2017).
- [11] R. E. CURTO AND L. A. FIALKOW, *Recursiveness, positivity, and truncated moment problems*, Houston J. Math., 17 (1991), pp. 603–635.

- [12] L. C. EVANS, *Partial differential equations*, vol. 19 of Graduate Studies in Mathematics, American Mathematical Society, Providence, RI, second ed., 2010.
- [13] B. D. FROESE AND A. M. OBERMAN, *Convergent finite difference solvers for viscosity solutions of the elliptic Monge-Ampère equation in dimensions two and higher*, SIAM J. Numer. Anal., 49 (2011), pp. 1692–1714.
- [14] A. GALICHON, P. HENRY-LABORDÈRE, AND N. TOUZI, *A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options*, Ann. Appl. Probab., 24 (2014), pp. 312–336.
- [15] I. KARATZAS AND S. E. SHREVE, *Brownian motion and stochastic calculus*, vol. 113 of Graduate Texts in Mathematics, Springer-Verlag, New York, second ed., 1991.
- [16] M. A. KATSOUKAKIS, *Viscosity solutions of second order fully nonlinear elliptic equations with state constraints*, Indiana Univ. Math. J., 43 (1994), pp. 493–519.
- [17] J. B. LASSERRE, *Moments, positive polynomials and their applications*, vol. 1 of Imperial College Press Optimization Series, Imperial College Press, London, 2010.
- [18] C. W. MILLER, *Nonlinear PDE approach to time-inconsistent optimal stopping*, SIAM J. Control Optim., 55 (2017), pp. 557–573.
- [19] A. M. OBERMAN, *The convex envelope is the solution of a nonlinear obstacle problem*, Proc. Amer. Math. Soc., 135 (2007), pp. 1689–1694 (electronic).
- [20] ———, *Wide stencil finite difference schemes for the elliptic Monge-Ampère equation and functions of the eigenvalues of the Hessian*, Discrete Contin. Dyn. Syst. Ser. B, 10 (2008), pp. 221–238.
- [21] J. L. PEDERSEN AND G. PESKIR, *Optimal mean-variance selling strategies*, Math. Financ. Econ., 10 (2016), pp. 203–220.
- [22] G. PESKIR AND A. SHIRYAEV, *Optimal stopping and free-boundary problems*, Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 2006.
- [23] H. PHAM, *Continuous-time stochastic control and optimization with financial applications*, vol. 61, Springer Science & Business Media, 2009.
- [24] D. B. ROKHLIN, *Stochastic Perron’s method for optimal control problems with state constraints*, Electron. Commun. Probab., 19 (2014), pp. no. 73, 15.
- [25] A. N. SHIRYAEV, *Optimal stopping rules*, vol. 8 of Stochastic Modelling and Applied Probability, Springer-Verlag, Berlin, 2008. Translated from the 1976 Russian second edition by A. B. Aries, Reprint of the 1978 translation.
- [26] N. TOUZI, *Optimal stochastic control, stochastic target problems, and backward SDE*, vol. 29 of Fields Institute Monographs, Springer, New York; Fields Institute for Research in Mathematical Sciences, Toronto, ON, 2013. With Chapter 13 by Agnès Tourin.
- [27] C. ZUCCA AND L. SACERDOTE, *On the inverse first-passage-time problem for a wiener process*, The Annals of Applied Probability, (2009), pp. 1319–1346.