

Explicit solution of relative entropy weighted control

Joris Bierkens*, Hilbert J. Kappen†

July 30, 2014

Abstract

We consider the minimization over probability measures of the expected value of a random variable, regularized by relative entropy with respect to a given probability distribution. In the general setting we provide a complete characterization of the situations in which a finite optimal value exists and the situations in which a minimizing probability distribution exists. Specializing to the case where the underlying probability distribution is Wiener measure, we characterize finite relative entropy changes of measure in terms of square integrability of the corresponding change of drift. For the optimal change of measure for the relative entropy weighted optimization, an expression involving the Malliavin derivative of the cost random variable is derived. The theory is illustrated by its application to several examples, including the case where the cost variable is the maximum of a standard Brownian motion over a finite time horizon. For this example we obtain an exact optimal drift, as well as an approximation of the optimal drift through a Monte-Carlo algorithm.

Keywords and phrases: stochastic optimal control, Itô calculus, Brownian motion, Malliavin calculus, relative entropy, path integral control, Monte-Carlo sampling

MSC2010 subject classifications: primary 93E20; secondary 60H07, 94A17.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 270327

1 Introduction

In certain situations in stochastic optimal control theory, the dynamic programming or Hamilton-Jacobi-Bellman equations may be transformed, through the Hopf-transform, into linear equations [Fle82], [FS09, Chapter VI]. In the past years, within the applied control and machine learning community, there has been a significant amount of interest in this class of ‘path integral control problems’ (see e.g. [Kap05, Tod06, BH13]). This class of problems also occurs in risk sensitive control theory (see [FS09]) and the theory of large deviations (see [BD98]), and it occurs in modified form (constrained to equivalent martingale measures) in mathematical finance, in particular as the dual problem for a portfolio optimization problem [Mon13]. It is the goal of this paper to review and extend the mathematical underpinning of this optimization problem, as well as showcase some new results within this context.

The problem we consider is a minimization problem over probability measures that are absolutely continuous with respect to a given probability measure (referred to as the ‘uncontrolled measure’). The functional we wish to minimize is the sum of (i) the expectation of a given random variable with respect to any probability measure, and (ii) the relative entropy of that probability measure with respect to the uncontrolled measure. The density of the optimal probability measure with respect to the uncontrolled distribution is readily available through an explicit expression in

*Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, The Netherlands, j.bierkens@science.ru.nl

†Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, The Netherlands, b.kappen@science.ru.nl

terms of the cost random variable. The challenge is then to understand this probability measure within the context of the underlying problem.

In particular, in the special case in which we are interested in this paper, the uncontrolled distribution is Wiener measure on the space of continuous sample paths. Absolutely continuous change of distribution then corresponds, by the Girsanov theorem, to a change of drift, which we will interpret as control process. The regularizing relative entropy corresponds to squared control cost. The questions we wish to answer in this paper are: (i) under what conditions does there exist an optimal change of drift corresponding to a given cost functional and probability measure, and (ii) how can it be computed?

There is a close relation to existing theory within the field of large deviations theory and stochastic optimal control [BD98, DE97]. We should also mention the work of Föllmer [F85, F86]. For an excellent self-contained review of these results, see [Leh13]. The main aim of this paper is to review the application of the mentioned results in a mathematical control context. We also illustrate the use of Monte-Carlo methods; for recent work on the use of Monte-Carlo methods in relative entropy weighted control, see [TK14].

Furthermore, for the reader who is familiar with the available literature, theoretical contributions of our paper include:

- (i) Relaxing the usual boundedness assumptions to a condition that guarantees finite relative entropy of the optimal change of measure (condition (FE) of Section 2);
- (ii) Solution of the problem where the cost random variable is the maximum of a standard Brownian motion with controlled drift over a finite time horizon.

1.1 Outline

In Section 2, we consider the general relative entropy weighted optimization problem, and completely characterize the different situations that may arise. The situation with finite relative entropy is most useful, and problems for which the optimal change of measure has finite relative entropy are easily characterized in terms of conditions on the cost functional and the probability measure. Then in Section 3, the finite relative entropy case is further investigated within the context of a Wiener process. It is shown that a change of measure with finite relative entropy corresponds to a square integrable drift, which in particular is the case for the optimal density. In Section 3.2 we show how the optimal drift may be computed through the Clark-Ocone formula. To illustrate the use of this approach, and as an interesting result in its own right, we compute the optimal drift for the case where the cost functional is the maximum of a one dimensional Wiener process with controlled drift on a finite time horizon (Section 3.3). We also provide a Monte-Carlo algorithm for the approximation of such a solution, which is easily extended to other problems.

1.2 Notation

As is common in probability theory, we will allow random variables to assume their values within the extended reals $[-\infty, \infty]$. Resulting formal expressions may be interpreted as follows: $\log 0 = -\infty$, $\log \infty = \infty$, $\exp(-\infty) = 0$, $\exp(\infty) = \infty$ and $\infty \exp(-\infty) = 0$. For any $a \in \mathbb{R}$, we write $(a)^+ := a \vee 0$ and $(a)^- := -a \vee 0$ for the positive and negative parts of a . The euclidean norm of $x \in \mathbb{R}^d$ is denoted by $|x|$. For an adapted process θ and a continuous local martingale M , both with values in \mathbb{R}^d , we write $\int_0^t \langle \theta_s, dM_s \rangle$ to indicate $\sum_{i=1}^d \int_0^t \theta_s^i dM_s^i$. If $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space we write $\mathbb{E}^{\mathbb{P}}$ for expectation with respect to the probability measure \mathbb{P} . Lebesgue measure will be denoted by Leb .

2 Relative entropy weighted optimization

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. The probability measure \mathbb{P} will be referred to as the *uncontrolled (probability) measure*. Let C be a random variable assuming values in $[-\infty, \infty]$. The

random variable C indicates a cost we wish to minimize, as explained below.

Let \mathcal{P} denote the set of probability measures on (Ω, \mathcal{F}) . We wish to find a probability measure $\mathbb{Q} \in \mathcal{P}$ that

- (i) is absolutely continuous with respect to \mathbb{P} (denoted by $\mathbb{Q} \ll \mathbb{P}$),
- (ii) reduces the expected cost $\mathbb{E}^{\mathbb{Q}}C$, but
- (iii) has small deviation from \mathbb{P} . We take the relative entropy

$$\mathcal{H}(\mathbb{Q}; \mathbb{P}) = \int_{\Omega} \log \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right) d\mathbb{Q} = \mathbb{E}^{\mathbb{Q}} \left[\log \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right) \right]$$

as a measure of this deviation (see e.g. [DE97, Section 1.4] for preliminary results on relative entropy). Recall that $\mathcal{H}(\mathbb{Q}; \mathbb{P}) \geq 0$ for any $\mathbb{Q}, \mathbb{P} \in \mathcal{P}$, and $\mathcal{H}(\mathbb{Q}; \mathbb{P}) = 0$ if and only if $\mathbb{Q} = \mathbb{P}$.

Note that (i) is a constraint and (ii) and (iii) are conflicting optimization targets.

Let $\mathcal{P}_0 := \{\mathbb{Q} \in \mathcal{P} : \mathbb{E}^{\mathbb{Q}}[(C)^+] < \infty \text{ and } \mathcal{H}(\mathbb{Q}; \mathbb{P}) < \infty\}$ denote the set of admissible probability measures, and note that \mathcal{P}_0 is convex. Define the cost functional $J : \mathcal{P} \rightarrow \mathbb{R}$ by

$$J(\mathbb{Q}) := \begin{cases} \mathbb{E}^{\mathbb{Q}}C + \mathcal{H}(\mathbb{Q}; \mathbb{P}) = \mathbb{E}^{\mathbb{Q}} \left[C + \log \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right) \right] & \text{if } \mathbb{Q} \in \mathcal{P}_0, \\ \infty & \text{otherwise.} \end{cases} \quad (1)$$

The definition of \mathcal{P}_0 is as non-restrictive as possible, so that for $\mathbb{Q} \in \mathcal{P}_0$ the value $J(\mathbb{Q})$ is well defined within the interval $[-\infty, \infty)$.

We arrive at the following problem:

Problem 2.1 (Relative entropy weighted optimization). Compute $J^* = \inf_{\mathbb{Q} \in \mathcal{P}_0} J(\mathbb{Q})$, and if it exists, a minimizer $\mathbb{Q}^* \in \mathcal{P}_0$ such that $J(\mathbb{Q}^*) = J^*$.

The solution of this problem is well known for the case in which $\mathbb{P}(|C| < K) = 1$ for some $K > 0$, see e.g. [DE97, Proposition 1.4.2] or [BD98, Proposition 2.5]. The purpose of this section is to provide a complete characterization of the existence of solutions of Problem 2.1 in terms of conditions on \mathbb{P} and C . To achieve this goal, we will consider the following further conditions on C and \mathbb{P} .

$$\textit{finite (relative) entropy: } \mathbb{P}(C < +\infty) > 0 \text{ and } \mathbb{E}^{\mathbb{P}}[\exp(-C)|C|] < \infty. \quad (\text{FE})$$

$$\textit{integrability: } 0 < \mathbb{E}^{\mathbb{P}}[\exp(-C)] < \infty. \quad (\text{I})$$

Condition (FE) will earn its name ('finite relative entropy') below: as we will see, it is a necessary and sufficient condition for the 'optimal' probability distribution to have finite relative entropy with respect to \mathbb{P} . The implication

$$(\text{FE}) \implies (\text{I})$$

holds as a result of the estimates $\exp(-x)\mathbb{1}_{\{x \leq -1\}} \leq \exp(-x)|x|$ and $\exp(-x)\mathbb{1}_{\{x > -1\}} \leq \frac{1}{e}$.

Example 2.2. The following examples may serve to illustrate the conditions (FE) and (I).

- (i) $\Omega = [0, \infty)$, with \mathbb{P} having density $f(\omega) = \exp(-\omega)$ with respect to Lebesgue measure; $C(\omega) = -\omega$: $\mathbb{P}(C < +\infty) > 0$ but (I) does not hold.
- (ii) $\Omega = (-\infty, \infty)$, $\frac{d\mathbb{P}}{d\text{Leb}}(\omega) = \frac{\exp(-|\omega|)}{k(1+\omega^2)}$, with k a normalization constant, and $C(\omega) = -|\omega|$. Then \mathbb{P} is a probability distribution,

$$\mathbb{E}^{\mathbb{P}}[\exp(-C)] = \frac{1}{k} \int_{-\infty}^{\infty} \frac{1}{1+x^2} dx = \frac{\pi}{k},$$

and $\mathbb{E}^{\mathbb{P}}[|C|\exp(-C)] = \infty$. So (I) holds but (FE) does not.

We have the following observation.

Lemma 2.3. *If $\mathbb{P}(C < +\infty) > 0$ then \mathcal{P}_0 is non-empty.*

Proof. Under the assumption, there exists an $M > 0$ such that $\mathbb{P}(C < M) > 0$. Let $Z = \mathbb{1}_{\{C \leq M\}} / \mathbb{P}(C \leq M)$. Then $\mathbb{E}^{\mathbb{P}} Z = 1$, so that $d\mathbb{Q}/d\mathbb{P} = Z$ defines a valid probability measure. Furthermore $\mathbb{E}^{\mathbb{Q}}(C)^+ \leq M < \infty$ and $\mathbb{E}^{\mathbb{Q}}[|\log(Z)|] < \infty$. We may conclude that $\mathbb{Q} \in \mathcal{P}_0$. \square

If (I) holds, then

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} = Z^* := \frac{\exp(-C)}{\mathbb{E}^{\mathbb{P}}[\exp(-C)]}. \quad (2)$$

defines a probability measure \mathbb{Q}^* that is absolutely continuous with respect to \mathbb{P} .

Lemma 2.4. *Suppose (I) holds, and let \mathbb{Q}^* and Z^* be given by (2).*

- (i) $\mathbb{Q}^* \in \mathcal{P}_0$ if and only if (FE) holds.
- (ii) For any $\mathbb{Q} \in \mathcal{P}_0$, we have that $\mathbb{Q} \ll \mathbb{Q}^*$, and

$$J(\mathbb{Q}) = \mathcal{H}(\mathbb{Q}; \mathbb{Q}^*) - \log \mathbb{E}^{\mathbb{P}} \exp(-C), \quad (3)$$

(so that in particular $J(\mathbb{Q}) = \infty$ if $\mathcal{H}(\mathbb{Q}; \mathbb{Q}^*) = \infty$);

- (ii) J is strictly convex over $\{\mathbb{Q} \in \mathcal{P}_0 : \mathcal{H}(\mathbb{Q}; \mathbb{Q}^*) < \infty\}$ and has unique minimizer \mathbb{Q}^* , provided $\mathbb{Q}^* \in \mathcal{P}_0$.

Proof. (i) Since $\mathbb{E}^{\mathbb{P}}[\exp(-C)C\mathbb{1}_{\{C \geq 0\}}] < \infty$ for any random variable C , it remains to check that $\mathbb{E}^{\mathbb{P}}[\exp(-C)(-C)\mathbb{1}_{\{C < 0\}}] < \infty$, which is the case if and only if (FE) holds.

- (ii) If \mathbb{Q} is not absolutely continuous with respect to \mathbb{Q}^* , then there exists a set E of \mathbb{Q}^* -measure zero, for which $\mathbb{Q}(E) > 0$. We have $Z^* = 0$, \mathbb{P} -almost surely on E , so that $C = +\infty$, \mathbb{P} -almost surely, on E . Therefore $\mathbb{E}^{\mathbb{Q}}(C)^+ \geq \mathbb{E}^{\mathbb{Q}}C\mathbb{1}_E = \infty$. So $\mathbb{Q} \notin \mathcal{P}_0$.

Now let $\mathbb{Q} \in \mathcal{P}_0$. We have just seen that $\mathbb{Q} \ll \mathbb{Q}^*$, say with density $Y = \frac{d\mathbb{Q}}{d\mathbb{Q}^*}$. We may choose a version of Y such that $Y = 1$ on $\{Z^* = 0\}$. Then $\mathbb{Q} \ll \mathbb{P}$ with density $Z = \frac{d\mathbb{Q}}{d\mathbb{P}} = YZ^*$. Write $K = \mathbb{E}^{\mathbb{P}} \exp(-C)$. Then

$$\begin{aligned} \mathcal{H}(\mathbb{Q}; \mathbb{Q}^*) &= \int_{\Omega} Z \log Y \, d\mathbb{P} = \int_{\Omega \cap \{Z^* > 0\}} Z \log Y \, d\mathbb{P} \\ &= \int_{\Omega \cap \{Z^* > 0\}} Z (\log Z - \log Z^*) \, d\mathbb{P} = \mathcal{H}(\mathbb{Q}; \mathbb{P}) - \int_{\Omega \cap \{Z^* > 0\}} Z \log Z^* \, d\mathbb{P} \\ &= \mathcal{H}(\mathbb{Q}; \mathbb{P}) + \int_{\Omega \cap \{Z^* > 0\}} Z (\log K + C) \, d\mathbb{P} = \mathcal{H}(\mathbb{Q}; \mathbb{P}) + \log K + \mathbb{E}^{\mathbb{P}}[ZC]. \end{aligned}$$

In particular

$$J(\mathbb{Q}) = \mathbb{E}^{\mathbb{P}} ZC + \mathcal{H}(\mathbb{Q}; \mathbb{P}) = \mathcal{H}(\mathbb{Q}; \mathbb{Q}^*) - \log K.$$

- (iii) Using (3), this follows from convexity of relative entropy and the fact that $\mathcal{H}(\mathbb{Q}; \mathbb{Q}^*) = 0$ if and only if $\mathbb{Q} = \mathbb{Q}^*$. \square

The following proposition gives a sufficient condition for the existence of a minimizer.

Proposition 2.5. *Suppose Hypothesis (FE) holds. Then \mathbb{Q}^* , given by (2), is the unique minimizer for Problem 2.1, and $J(\mathbb{Q}^*) = -\log \mathbb{E}^{\mathbb{P}} \exp(-C)$.*

In particular, under (FE), $\mathcal{H}(\mathbb{Q}^*; \mathbb{P}) < \infty$ which explains the name of this condition.

Proof. Note that $\mathbb{Q}^* \in \mathcal{P}_0$ if (FE) holds. The stated result is now an immediate consequence of Lemma 2.4. \square

The following proposition holds with minimal restrictions on \mathbb{P} and C . Since we are primarily interested in the case where Hypothesis (FE) holds (so as to have the existence result of Proposition 2.5), the proof is delegated to the appendix.

Proposition 2.6. *Suppose $\mathbb{P}(C < +\infty) > 0$. Then*

$$\inf_{\mathbb{Q} \in \mathcal{P}_0} J(\mathbb{Q}) = -\log \mathbb{E}^{\mathbb{P}} \exp(-C). \quad (4)$$

To conclude this section, we summarize by distinguishing the following cases, as immediate consequences of Lemma 2.4, Proposition 2.5 and Proposition 2.6.

Corollary 2.7. *(i) If Hypothesis (FE) holds, then there exists a minimizer for Problem 2.1;*

(ii) If Hypothesis (I) holds, but Hypothesis (FE) does not hold, then Problem 2.1 has optimal value $-\infty < J^ < \infty$, but there does not exist a probability measure $\mathbb{Q} \in \mathcal{P}_0$ at which this value is attained;*

(iii) If $\mathbb{P}(C < +\infty) > 0$ but Hypothesis (I) does not hold, then Problem 2.1 has optimal value $J^ = -\infty$.*

2.1 Notes and remarks

Remark 2.8. One may wish to include a factor $\beta > 0$ in the problem formulation, to indicate the relative importance of minimizing $\mathbb{E}^{\mathbb{Q}}C$ compared to minimizing $\mathcal{H}(\mathbb{Q}; \mathbb{P})$ to obtain the form $J(\mathbb{Q}) = \mathbb{E}^{\mathbb{Q}}C + \frac{1}{\beta} \mathcal{H}(\mathbb{Q}; \mathbb{P})$. In this case the optimal value becomes $J^* = -\frac{1}{\beta} \log \mathbb{E}^{\mathbb{P}} \exp(-\beta C)$, so that β admits the interpretation of inverse temperature, in the context of statistical physics.

Remark 2.9. A sufficient condition for Hypothesis (FE) to hold is that for some $\gamma > 1$,

$$0 < \mathbb{E}^{\mathbb{P}} \exp(-\gamma C) < \infty. \quad (5)$$

Indeed, if this is the case, then, since $x \leq \frac{1}{\varepsilon e} \exp(\varepsilon x)$ for $x \geq 0$,

$$\mathbb{E}^{\mathbb{P}} [\exp(-C)(-C) \mathbb{1}_{\{C \leq 0\}}] \leq \frac{1}{(\gamma - 1)e} \mathbb{E}^{\mathbb{P}} [\exp(-\gamma C) \mathbb{1}_{\{C \leq 0\}}] < \infty.$$

(And, as always, $\mathbb{E}^{\mathbb{P}} [\exp(-C)|C| \mathbb{1}_{\{C > 0\}}] \leq \frac{1}{e}$.) In turn (5), and therefore (FE), are implied by the conditions that $\mathbb{P}(C > K) = 1$ for some $K \in \mathbb{R}$, and $\mathbb{P}(C = \infty) < 1$.

3 Relative entropy weighted optimization with respect to Wiener measure

Our primary interest lies in the theory discussed in Section 2 applied to the special case where all randomness is generated by a d -dimensional Wiener process.

Equivalent changes of measure (satisfying mild conditions) may in this case be expressed as a Girsanov type transformation. The corresponding change of drift will constitute the ‘control process’. A crucial observation is that the relative entropy of such a transformation is given by the squared control costs. However, quite importantly, we will not restrict ourselves to the case in which the optimal measure is equivalent to the Wiener measure, and will only require absolute continuity, with finite relative entropy. This allows us to consider cost functionals that may have value $+\infty$ with positive probability.

In Section 3.1 we will discuss the general setting. In Section 3.2 we specialize to the case in which the cost functional has a Malliavin derivative, which provides a convenient expression for the optimal drift.

3.1 General results

Definition 3.1 (Canonical Wiener process). Let $\Omega = C([0, \infty); \mathbb{R}^d)$, i.e. the space of continuous functions mapping $[0, \infty)$ into \mathbb{R}^d . For $t \geq 0$, let $\mathcal{F}_t^o = \sigma(\{\omega(s) : 0 \leq s \leq t\})$ for $t \geq 0$. Let \mathbb{P} denote Wiener measure, and for $t \geq 0$, let (\mathcal{F}_t) be the right-continuous completion of (\mathcal{F}_t^o) with respect to \mathbb{P} (see [RW94a, Section II.67]). Let $X_t(\omega) := \omega(t)$ for $\omega \in \Omega$ and $t \geq 0$, so that X is a standard Brownian motion in \mathbb{R}^d under the probability measure \mathbb{P} . Let $\mathcal{F} := \mathcal{F}_\infty := \sigma(\cup_{t \geq 0} \mathcal{F}_t)$. The collection $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathcal{F}, \mathbb{P}, X)$ will be referred to as a *canonical d -dimensional Wiener process*.

Throughout this section let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathcal{F}, \mathbb{P}, X)$ denote a canonical d -dimensional Wiener process. Furthermore let $C : \Omega \rightarrow \mathbb{R}$ satisfy (FE), i.e.

$$\mathbb{P}(C < +\infty) > 0 \quad \text{and} \quad \mathbb{E}^{\mathbb{P}}[\exp(-C)|C|] < \infty.$$

The following result and its practical application (Remark 3.3) are of fundamental importance in the explicit computation of the optimal drift in general cases.

Proposition 3.2 (Föllmer's drift). *Let \mathbb{Q} be a probability measure on (Ω, \mathcal{F}) that is absolutely continuous with respect to \mathbb{P} . There exists a adapted process (U_t) such that under \mathbb{Q} the following hold:*

- (i) *The process $X_t^U = X_t - \int_0^t U_s ds$ is a \mathbb{Q} -Brownian motion;*
- (ii) *$\mathcal{H}(\mathbb{Q}; \mathbb{P}) = \frac{1}{2} \mathbb{E}^{\mathbb{Q}} \int_0^\infty |U_t|^2 dt$.*

Proof. See [F86, Section 2] or [Leh13, Theorem 2]. □

Remark 3.3. Since our interest is in obtaining explicit expressions for the optimal feedback control, we also state here the construction of the Föllmer drift (U_t) , and hence of the feedback control u . Let $Z = \frac{d\mathbb{Q}}{d\mathbb{P}}$, and write $Z_t = \mathbb{E}^{\mathbb{P}}[Z | \mathcal{F}_t]$, $t \geq 0$. Since (Z_t) is a continuous martingale, it allows the martingale representation $Z_t = 1 + \int_0^t \langle V_s, dX_s \rangle$ [RW94b, Theorem IV.36.1]. The Föllmer drift is given by $U_t := (V_t/Z_t) \mathbb{1}_{Z_t > 0}$. See [F86, Section 2] or [Leh13, Theorem 2] for details.

Definition 3.4 (Feedback control). A mapping $u : [0, \infty) \times C([0, \infty); \mathbb{R}^d) \rightarrow \mathbb{R}^d$ is called a *feedback control* if the process $U_t := u(t, X)$ is (\mathcal{F}_t) -adapted.

Suppose u is a feedback control and consider the stochastic differential equation

$$Y_t = \int_0^t u(s, Y) ds + B_t \quad \text{for } t \geq 0. \tag{6}$$

Recall that a *weak solution* of (6) consists of a filtered probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, (\tilde{\mathcal{F}}_t), \tilde{\mathbb{P}})$ satisfying the usual conditions, on which is defined a d -dimensional standard Brownian motion along with a $(\tilde{\mathcal{F}}_t)$ -adapted process Y with values in \mathbb{R}^d such that (6) is satisfied. Note that (Y_t) is not necessarily adapted to the filtration generated by (B_t) .

The following result shows how changes of measure may be associated with feedback controls and vice versa.

Proposition 3.5. (i) *Let u be a feedback control and suppose $(\tilde{\Omega}, \tilde{\mathcal{F}}, (\tilde{\mathcal{F}}_t), \tilde{\mathbb{P}}, Y, B)$ is a weak solution to (6), such that $\mathbb{E}^{\tilde{\mathbb{P}}}[\int_0^\infty |u(t, Y)|^2 dt] < \infty$. Then there exists a probability measure $\mathbb{Q} \ll \mathbb{P}$ on (Ω, \mathcal{F}) , such that under \mathbb{Q} , the coordinate process has the same distribution as Y , and such that $\mathcal{H}(\mathbb{Q}; \mathbb{P}) = \frac{1}{2} \mathbb{E}^{\tilde{\mathbb{P}}} \int_0^\infty |u(t, Y)|^2 dt$.*

(ii) *Suppose $\mathbb{Q} \ll \mathbb{P}$ with $\mathcal{H}(\mathbb{Q}; \mathbb{P}) < \infty$. Then there exists a feedback control u and adapted processes Y and B on $(\Omega, \mathcal{F}, (\mathcal{F}_t))$ satisfying (6) such that \mathbb{Q} is the law of Y , B is a standard Brownian motion under \mathbb{Q} , and $\mathcal{H}(\mathbb{Q}; \mathbb{P}) = \frac{1}{2} \mathbb{E}^{\mathbb{P}} \int_0^\infty |u(t, Y)|^2 dt$.*

Proof. (i) See e.g. [Leh13, Proposition 1].

- (ii) Let U be the Föllmer drift corresponding to \mathbb{Q} , and note that $u(t, X) := U_t(\omega)$ is a feedback control. Furthermore $B_t := X_t - \int_0^t u(s, X) ds$ is a standard Brownian motion under \mathbb{Q} , so that $Y := X$ satisfies (6).

In either case Proposition 3.2(ii) gives the stated expression for the entropy. \square

The following result is now a simple combination of Proposition 2.5 and Proposition 3.5(ii).

Corollary 3.6. *Suppose $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P}, X)$ is a d -dimensional canonical Wiener process. Let $C : \Omega \rightarrow \mathbb{R}$ be \mathcal{F} -measurable such that (FE) holds. Then there exists a probability measure $\mathbb{Q}^* \in \mathcal{P}_0$ that solves Problem 2.1. Furthermore there exists a feedback control u , such that*

- (i) $U_t = u(t, X)$ is the Föllmer drift (of Proposition 3.2) associated with \mathbb{Q}^* ;
- (ii) Under \mathbb{Q}^* , the law of the coordinate process X coincides with the law of the solution Y of (6); and
- (iii) $\frac{1}{2} \mathbb{E}^{\mathbb{Q}^*} \int_0^\infty |u(s, X)|^2 ds = \mathcal{H}(\mathbb{Q}^*; \mathbb{P}) < \infty$.

Remark 3.7. The practical interpretation of Corollary 3.6 is that we solve the control problem

$$\inf_u \mathbb{E}^{\tilde{\mathbb{P}}} \left[C(Y^u) + \frac{1}{2} \int_0^\infty |u(s, Y^u)|^2 ds \right],$$

with the infimum taken over all feedback controls u , and where Y^u satisfies (6) with respect to a $\tilde{\mathbb{P}}$ -Brownian motion B on some fixed probability space $(\tilde{\Omega}, \tilde{\mathbb{P}})$. The problem with making this interpretation rigorous, is that one would have to impose conditions on u such that a strong solution to (6) always exists. As illustrated by Corollary 3.6, we are able to obtain, in general, a weak solution to the mentioned problem. In [Leh13, Section 2.3] conditions are discussed such that a strong solution exists.

Remark 3.8. The result of Corollary 3.6 is relatively straightforward if $\mathbb{P}(C = +\infty) = 0$. In this case the optimal measure \mathbb{Q}^* with density $Z^* = \frac{d\mathbb{Q}^*}{d\mathbb{P}} = \frac{\exp(-C)}{\mathbb{E}^{\mathbb{P}} \exp(-C)}$ is equivalent to \mathbb{P} , so that standard Girsanov theory (e.g. [RW94b, Theorem IV.38.5(i)]) can be applied to find the representation

$$Z_t^* = \exp \left(\int_0^t \langle u(s, X), dX_s \rangle - \frac{1}{2} \int_0^t |u(s, X)|^2 ds \right) \quad (7)$$

for the density process $Z_t^* := \mathbb{E}^{\mathbb{P}}[Z^* | \mathcal{F}_t]$; it only remains to verify the stated expression for the relative entropy, which boils down to a simple computation. We have to be more careful if $\mathbb{P}(C = +\infty) > 0$. In this case $\mathbb{P}(Z^* = 0) > 0$ so that the measures \mathbb{Q}^* and \mathbb{P} are not equivalent. The representation (7) is still valid but only \mathbb{Q}^* -almost surely; see [F85, Section 2] or [Leh13, Theorem 2].

The following proposition provides a method to compute the optimal drift through Monte Carlo sampling, by employing the law of large numbers to approximate the (conditional) expectation. In Section 3.3.2 this idea is used in the approximate minimization of the running maximum of a Brownian motion with drift.

Proposition 3.9. *Let $\mathbb{Q} \ll \mathbb{P}$ with $\mathcal{H}(\mathbb{Q}; \mathbb{P}) < \infty$ with density $Z = \frac{d\mathbb{Q}}{d\mathbb{P}}$ and let u be the feedback control of Proposition 3.5. Suppose that Z is \mathcal{F}_T -measurable for some $T \geq 0$. Then the corresponding Föllmer drift $U_t = u(t, X)$ at time t satisfies*

$$U_t = \lim_{h \downarrow 0} \frac{1}{h} \mathbb{E}^{\mathbb{Q}}[(X_{t+h} - X_t) | \mathcal{F}_t] \quad \text{in } L^2(\mathbb{Q}) \text{ for } 0 \leq t \leq T. \quad (8)$$

Proof. This is [F86, Proposition 2.5]. \square

3.1.1 Examples and remarks

In all the examples below, $Z^* = \frac{d\mathbb{Q}^*}{d\mathbb{P}}$, with \mathbb{Q}^* as in Corollary 3.6. Also, $Z_t^* = \mathbb{E}^{\mathbb{P}}[Z^* | \mathcal{F}_t]$, $t \geq 0$.

Remark 3.10. Finite time horizon problems fit well within the theory of this section: In this case one just considers a cost random variable C that is \mathcal{F}_T -measurable for some deterministic $T > 0$. The optimal drift (U_t) will be \mathcal{F}_T -measurable for all $t \geq 0$, by its construction.

As an illustration of this remark, consider the following example.

Example 3.11 (Constrained problem, obstacle avoidance). Let $A \subset \mathbb{R}^d$ be open and let $T > 0$. Then the event $\{\omega : X_T \in A\}$ is \mathcal{F}_T -measurable. Let

$$C = \begin{cases} +\infty & \text{if } X_T \in A, \\ 0 & \text{otherwise.} \end{cases}$$

Then (FE) is satisfied as long as $\mathbb{P}(C < +\infty) = \mathbb{P}(X_T \notin A) > 0$. Optimally controlled paths will avoid the set A at time T . For example, in the one-dimensional case, let us take $A = (a, b)$, with $a < b$, possibly with $a = -\infty$ or $b = +\infty$, but not both. We compute for $t < T$,

$$N_t := \mathbb{E}^{\mathbb{P}}[\exp(-C) | \mathcal{F}_t] = \mathbb{E}^{\mathbb{P}}[\mathbb{1}_{X_T \notin A} | \mathcal{F}_t] = 1 - \frac{1}{\sqrt{2\pi(T-t)}} \int_a^b \exp(-(x - X_t)^2/2(T-t)) dx.$$

By Itô calculus, differentiating under the integral, for $t < T$,

$$\begin{aligned} dN_t &= - \left(\frac{1}{\sqrt{2\pi(T-t)}} \int_a^b \left\{ \exp(-(x - X_t)^2/2(T-t)) \left(\frac{x - X_t}{T-t} \right) \right\} dx \right) dX_t \\ &= \left(\frac{1}{\sqrt{2\pi(T-t)}} \left\{ \exp(-(b - X_t)^2/2(T-t)) - \exp(-(a - X_t)^2/2(T-t)) \right\} \right) dX_t \end{aligned}$$

Since $Z_t^* = N_t/N_0$, and by Remark 3.3, we find that

$$U_t = \frac{\exp(-(b - X_t)^2/2(T-t)) - \exp(-(a - X_t)^2/2(T-t))}{\sqrt{2\pi(T-t)} - \int_a^b \exp(-(x - X_t)^2/2(T-t)) dx} \mathbb{1}_{t < T}$$

is the optimal drift (U_t) .

Note that we have no difficulty in dealing with the fact that $\mathbb{P}(C = +\infty) > 0$. This is also illustrated by the following example.

Example 3.12 (Whittle's 'flypaper' example [Whi82, Section 20.3],[RW94b, Sections V.7, V.15]). Consider the case in which X is a one-dimensional Wiener process, i.e. $d = 1$. Let $\tau = \inf\{t > 0 : X_t \in \{a, b\}\}$, with $a < 0 < b$. As is well known, $\mathbb{E}\tau < \infty$. Consider the cost functional $C = k_a \mathbb{1}_{X_\tau = a} + k_b \mathbb{1}_{X_\tau = b} + \gamma\tau$, with k_a and k_b given scalar constants in $(-\infty, \infty]$, not both equal to $+\infty$, and $\gamma \geq 0$. Then (FE) is satisfied for C . Suppose that φ is the unique twice continuously differentiable solution on $[a, b]$ to the linear differential equation

$$\begin{cases} \frac{1}{2}\varphi''(x) - \gamma\varphi(x) = 0, & (a < x < b), \\ \varphi(a) = \exp(-k_a), & \varphi(b) = \exp(-k_b). \end{cases}$$

Define $N_t := \exp(-\gamma(t \wedge \tau))\varphi(X_{t \wedge \tau})$. Then by Itô's formula, for $0 \leq t \leq \tau$,

$$dN_t = -\gamma N_t dt + \frac{1}{2} \exp(-\gamma t) \varphi''(X_t) + \exp(-\gamma t) \varphi'(X_t) dX_t = \exp(-\gamma t) \varphi'(X_t) dX_t,$$

and $dN_t = 0$ for $t > \tau$, so that (N_t) is a local martingale. In fact, since φ has bounded derivative, (N_t) is a martingale. Since $N_\infty = \exp(-C)$, we have $Z_t^* = N_t/N_0$. By Remark 3.3, it follows that

$$u(t, X) = \varphi'(X_t)/\varphi(X_t) \mathbb{1}_{t < \tau}$$

is the optimal feedback control.

Remark 3.13. Example 3.12 illustrates the basic approach to more general problems where C is of the form

$$C = \int_0^\tau v(s, \xi_s) ds + k(\tau, \xi_\tau),$$

with (ξ_t) a process in \mathbb{R}^n satisfying an SDE of the form

$$\xi_t = \xi_0 + \int_0^t b(s, \xi_s) ds + \int_0^t \sigma(s, \xi_s) dX_s, \quad (9)$$

and $\tau = \inf\{t > 0 : \xi_t \notin U \text{ or } t > T\}$ denoting the exit time of a region $U \times [0, T]$ with U open and $T \in (0, \infty]$. Assume that $\mathbb{P}(\tau < \infty) = 1$. The optimal solution of Corollary 3.6 may in principle be obtained by finding a solution to the PDE

$$\begin{cases} L\varphi(t, x) - v(t, x)\varphi(t, y) = 0, & x \in U, 0 < t < T, \\ \varphi(t, x) = \exp(-k(t, x)), & x \in \partial U \text{ or } t = T, \end{cases}$$

where L is the infinitesimal generator of (9), and verifying that $N_t := \exp\left(-\int_0^{t \wedge \tau} v(s, \xi_s) ds\right) \varphi(t \wedge \tau, \xi_{t \wedge \tau})$ is a martingale. Then $Z_t^* = N_t/N_0$, and an application of Itô's formula provides (by Remark 3.3) the optimal drift

$$U_t = \frac{[\sigma(t, \xi_t)]^T \nabla \varphi(t, \xi_t)}{\varphi(t, \xi_t)} \mathbb{1}_{t \leq \tau} \mathbb{1}_{\varphi(t, \xi_t) > 0}.$$

The optimal drift is then a Markov control for (ξ_t) (i.e. of the form $u(t, \xi_t)$), and under the optimal change of measure, ξ satisfies (in law) the ‘controlled’ SDE

$$\xi_t = \xi_0 + \int_0^t \{b(s, \xi_s) + \sigma(s, \xi_s)u(s, \xi_s)\} ds + \int_0^t \sigma(s, \xi_s) dB_s,$$

where (B_t) is a standard Brownian motion.

Remark 3.14. In [BD98], it is proven that for C bounded from above, and \mathcal{F}_1 -measurable (i.e. in the finite time horizon case), we have

$$\inf_U J(U) = -\log \mathbb{E}^\mathbb{P} \exp(-C),$$

where the minimization is over a suitable class of processes that are progressively measurable with respect to (\mathcal{F}_t) . This result is used in the field of large deviations theory [BD98, DE97]. See also [Leh13, Theorems 4, 9] for an elegant alternative approach, extending the above Boué-Dupuis formula to infinite time horizon. Our focus is on explicit computations in optimal control, so we have taken care so that our approach (i) is not restricted to finite time horizon problems; (ii) establishes explicitly the existence of a minimizer; and (iii) does not impose the restriction on the cost functional to be bounded from above.

Remark 3.15. Our set up, using Wiener measure on the time interval $[0, \infty)$ allows for control over an infinite time horizon. At the same time, by Corollary 3.6, condition (FE) excludes e.g. the case in which the control process is \mathbb{Q} -almost surely equal to a non-zero constant on an infinite time horizon (since then $\mathcal{H}(\mathbb{Q}; \mathbb{P}) = \infty$). However, condition (FE) allows us to consider problems whose cost function is \mathcal{F}_τ -measurable, with τ a stopping time that is unbounded but \mathbb{P} -almost surely finite, as illustrated by Example 3.12.

Remark 3.16. An alternative method for finding an expression for the optimal drift in the case of the control of diffusion processes with cost function as in Remark 3.13 is through the dynamic programming principle, i.e. the Hamilton-Jacobi-Bellman PDE. For the class of problems considered in this paper, the HJB equation may be transformed into the linear PDE (3.13) through the Hopf- or logarithmic transform [FS09, Kap05].

3.2 Computation of optimal drift by means of Clark-Ocone representation

In this section we will apply the Clark-Ocone theorem of Malliavin calculus to obtain an explicit representation of the optimal drift in terms of the Malliavin derivative of the cost random variable. This representation also occurs e.g. in [Leh13, Lemma 3]. Section 3.3 provides an illustration of the use of Corollary 3.18.

As usual, let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathcal{F}, \mathbb{P}, X)$ denote a canonical d -dimensional Wiener process. Recall the following definitions and notations (see [Nua06, Chapter 1] for details). Let H denote the Hilbert space $L^2([0, \infty); \mathbb{R}^d)$. For suitable random variables $F : \Omega \rightarrow \mathbb{R}$, the Malliavin derivative $DF \in L^2((\Omega, \mathbb{P}); H)$ may be defined. The domain of D in $L^2(\Omega, \mathcal{F}, \mathbb{P})$ is denoted by \mathbb{D}^2 . If $F \in \mathbb{D}^2$, so that $DF \in L^2(\Omega; H) \cong L^2(\Omega \times [0, \infty); \mathbb{R}^d)$, then we may identify DF with a stochastic process, denoted $D_t F$, as usual:

$$(D_t F)(\omega) = DF(\omega, t), \quad \omega \in \Omega, \quad t \geq 0.$$

Note that $(D_t F)_{t \geq 0}$ is not necessarily adapted, and the value of $D_t F(\omega)$ is defined $\mathbb{P} \otimes \text{Leb}$ -almost everywhere. We can now quote the following result, which is essentially an application of the Clark-Ocone formula [Nua06, Proposition 1.3.14].

Lemma 3.17. *Suppose $Z \in \mathbb{D}^2$ and let (U_t) be the Föllmer process associated to the measure \mathbb{Q} with density Z with respect to \mathbb{P} . Let $Z_t := \mathbb{E}^{\mathbb{P}}[Z \mid \mathcal{F}_t]$ for $t \geq 0$. Then $\mathbb{P} \otimes \text{Leb}$ -almost everywhere,*

$$U_t = \frac{\mathbb{E}^{\mathbb{P}}[D_t Z \mid \mathcal{F}_t]}{Z_t} \mathbb{1}_{\{Z_t > 0\}}.$$

Proof. See [Leh13, Lemma 3]. □

Corollary 3.18. *Suppose $\mathbb{P}(0 \leq C < \infty) = 1$, and $C \in \mathbb{D}^2$. Then the optimal drift $(U_t)_{t \geq 0}$ of Corollary 3.6 satisfies, $\mathbb{P} \otimes \text{Leb}$ -almost everywhere,*

$$U_t = -\frac{\mathbb{E}^{\mathbb{P}}[\exp(-C)D_t C \mid \mathcal{F}_t]}{\mathbb{E}^{\mathbb{P}}[\exp(-C) \mid \mathcal{F}_t]}.$$

Proof. Hypothesis (FE) is satisfied. By Proposition 2.5, the optimal density in Problem 2.1 has density $Z^* = \frac{\exp(-C)}{\mathbb{E}^{\mathbb{P}}[\exp(-C)]}$. Since $\mathbb{P}(C < +\infty) = 1$, $Z > 0$, \mathbb{P} -almost surely and therefore $Z_t > 0$, \mathbb{P} -almost surely for all $t \geq 0$ [RW94a, Theorem II.78.1(ii)]. Furthermore $x \mapsto \exp(-x)$ has bounded derivative for x bounded from below. Therefore we may apply the chain rule of Malliavin calculus to $\exp(-C)$ (see [Nua06, Proposition 1.2.3]), to conclude that $\exp(-C) \in \mathbb{D}^2$ with $D_t \exp(-C) = -\exp(-C)D_t C$. The result now follows from Lemma 3.17. □

Remark 3.19. In [OK91], the Clark-Ocone formula is used in a different optimization context, namely for portfolio optimization in mathematical finance.

3.3 Minimization of the maximum of a Brownian motion with drift over a finite time horizon

In this section we illustrate the theory by obtaining the optimal drift for the minimization of the maximum of a standard Brownian motion with drift over a finite time horizon. As it turns out, the value of optimal drift can be explicitly computed as a function of the difference between the running maximum and the current value of the Wiener process. For a different optimization problem related to the maximum of a Wiener process see [HS91].

First, in Section 3.3.1, we apply the result of Section 3.2 to obtain an explicit expression for the optimal drift through the Clark-Ocone formula. Next, in Section 3.3.2, we show how to apply Proposition 3.9 to obtain a Monte-Carlo estimate of the optimal drift.

Let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathcal{F}, \mathbb{P}, X)$ denote a one-dimensional canonical Wiener process. Define $M_t := \max_{0 \leq s \leq t} X_s$ and take $C := M_T$ for some $T > 0$. Hypothesis (FE) is satisfied by Remark 2.9 and the observation that $\mathbb{P}(M_T = \infty) = 0$.

3.3.1 Exact solution

For the distribution of M_t we have by virtue of the reflection principle [KS91, Section 2.8.A]

$$\mathbb{P}(M_t \geq a) = \left(\frac{2}{\pi}\right)^{1/2} \int_{a/\sqrt{t}}^{\infty} \exp(-\xi^2/2) d\xi, \quad t \geq 0, a \geq 0.$$

We will make use of the error function (erf) and complimentary error function (erfc), defined by

$$\operatorname{erf}(x) := \frac{2}{\sqrt{\pi}} \int_0^x \exp(-\eta^2) d\eta, \quad \operatorname{erfc}(x) := 1 - \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} \exp(-\eta^2) d\eta, \quad x \geq 0.$$

We compute the optimal drift corresponding to the minimization of C , using Corollary 3.18. This means we need to compute $\mathbb{E}^{\mathbb{P}}[\exp(-M_T)D_tM_T | \mathcal{F}_t]$ and $\mathbb{E}^{\mathbb{P}}[\exp(-M_T) | \mathcal{F}_t]$. We start with the latter. Conditional on \mathcal{F}_t , the event $M_T = M_t$ occurs when the maximum over $[t, T]$ does not exceed $y := M_t$. This has the same probability as the event that the maximum over $[0, T-t]$ does not exceed $y - X_t$, so

$$\mathbb{P}(M_T = M_t | X_t = x, M_t = y) = \mathbb{P}(M_{T-t} \leq y - x) = \left(\frac{2}{\pi}\right)^{1/2} \int_0^{\frac{M_t - X_t}{\sqrt{T-t}}} \exp(-\xi^2/2) d\xi.$$

For $0 \leq x \leq y < z$ we compute

$$\begin{aligned} \mathbb{P}(M_T \geq z | X_t = x, M_t = y) &= \mathbb{P}(M_T = M_t | X_t = x, M_t = y) + \mathbb{P}(M_{T-t} \geq z - x) \\ &= \left(\frac{2}{\pi}\right)^{1/2} \int_0^{\frac{M_t - X_t}{\sqrt{T-t}}} \exp(-\xi^2/2) d\xi + \left(\frac{2}{\pi}\right)^{1/2} \int_{\frac{z - X_t}{\sqrt{T-t}}}^{\infty} \exp(-\xi^2/2) d\xi. \end{aligned}$$

Therefore the density function of M_T conditional on \mathcal{F}_t is equal to

$$f_{M_T | \mathcal{F}_t}(\xi) = \left(\frac{2}{\pi(T-t)}\right)^{1/2} \exp\left(-\frac{(\xi - X_t)^2}{2(T-t)}\right), \quad \text{for } \xi > M_t \geq X_t.$$

We compute

$$\begin{aligned} N_t &:= \mathbb{E}^{\mathbb{P}}[\exp(-M_T) | \mathcal{F}_t] = \mathbb{E}^{\mathbb{P}}[\exp(-M_T) \mathbb{1}_{M_T = M_t} | \mathcal{F}_t] + \mathbb{E}^{\mathbb{P}}[\exp(-M_T) \mathbb{1}_{M_T > M_t} | \mathcal{F}_t] \\ &= \exp(-M_t) \mathbb{P}(M_T = M_t | \mathcal{F}_t) + \mathbb{E}^{\mathbb{P}}[\exp(-M_T) \mathbb{1}_{M_T > M_t} | \mathcal{F}_t] \\ &= \exp(-M_t) \left(\frac{2}{\pi}\right)^{1/2} \int_0^{\frac{M_t - X_t}{\sqrt{T-t}}} \exp(-\xi^2/2) d\xi \\ &\quad + \left(\frac{2}{\pi(T-t)}\right)^{1/2} \int_{M_t}^{\infty} \exp(-\xi) \exp\left(-\frac{(\xi - X_t)^2}{2(T-t)}\right) d\xi \\ &= \exp(-M_t) \operatorname{erf}\left(\frac{M_t - X_t}{\sqrt{2(T-t)}}\right) + \exp\left(-X_t + \frac{1}{2}(T-t)\right) \operatorname{erfc}\left(\frac{M_t - X_t + (T-t)}{\sqrt{2(T-t)}}\right) \quad (10) \end{aligned}$$

The Malliavin derivative of M_T is given by $D_tM_T = \mathbb{1}_{[0, \tau]}(t) = \mathbb{1}_{M_t < M_T}$, where τ is the a.s. unique point where X attains its maximum [Nua06, Exercise 1.2.11]. Therefore

$$\begin{aligned} V_t &:= -\mathbb{E}^{\mathbb{P}}[\exp(-M_T)D_tM_T | \mathcal{F}_t] = -\mathbb{E}^{\mathbb{P}}[\exp(-M_T) \mathbb{1}_{M_T > M_t} | \mathcal{F}_t] \\ &= -\left(\frac{2}{\pi(T-t)}\right)^{1/2} \int_{M_t}^{\infty} \exp\left(-\xi - \frac{(\xi - X_t)^2}{2(T-t)}\right) d\xi \\ &= -\exp\left(-X_t + \frac{1}{2}(T-t)\right) \operatorname{erfc}\left(\frac{M_t - X_t + (T-t)}{(2(T-t))^{1/2}}\right). \end{aligned}$$

We finally compute the optimal feedback control

$$u(t, X) := V_t/N_t = u(t, X_t, M_t), \quad (11)$$

where (with some abuse of notation), for $0 \leq t < T$, $x \in \mathbb{R}$, and $m \geq x$,

$$u(t, x, m) = \frac{-\exp(-x + \frac{1}{2}(T-t)) \operatorname{erfc}\left(\frac{m-x+T-t}{\sqrt{2(T-t)}}\right)}{\exp(-m) \operatorname{erf}\left(\frac{m-x}{\sqrt{2(T-t)}}\right) + \exp(-x + \frac{1}{2}(T-t)) \operatorname{erfc}\left(\frac{m-x+T-t}{\sqrt{2(T-t)}}\right)}. \quad (12)$$

Remark 3.20. This example illustrates how the theory developed in this paper may be applied to obtain non-Markov controls, and therefore provides a method that applies where a dynamic programming (i.e. the HJB equation) can not be used. One could in principle obtain a Markov control for the extended process (X_t, M_t) . However solution by means of dynamic programming is still far from straightforward.

Remark 3.21. Alternatively, one can find the required martingale representation by applying Itô's formula to the expression for N_t obtained in (10), as in the examples of Section 3.1.1. Since the semimartingale decomposition of $Z_t^* := N_t/N_0$ is unique [RW94b, Section IV.31], the decomposition obtained using Itô's formula is necessarily equal to the required martingale representation of (Z_t^*) as in Remark 3.3.

3.3.2 Approximate solution through Monte-Carlo sampling

Algorithm 1 provides pseudocode for the approximate computation of the optimal drift (U_t) , employing a Monte-Carlo approximation of the expression (8) for (U_t) .¹ See [TK14] for recent results in this direction. Parameters of the algorithm are:

- δ : small time step for numerical simulation;
- Δ : small time step determining the interval length over which the approximate controls will be constant; Δ should be an integer multiple of δ ;
- N : the number of samples paths generated at every iteration.

In the algorithm, $\mathcal{N}(\mu, \sigma^2)$ denotes a (pseudo) random number generator that draws from a normal distribution with mean μ and variance σ^2 . In the experiment we set $T = 1$, $N = 10^2$, $\Delta = 0.1$, $\delta = 0.01$. An illustration of the exact and approximate computational methods of this section is provided in Figure 1.

A Appendix

A.1 Proof of Proposition 2.6

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and C be a \mathcal{F} -measurable random variable assuming its values within the extended reals. Write \mathcal{P} for the set of probability measures on (Ω, \mathcal{F}) , $\mathcal{P}_0 := \{\mathbb{Q} \in \mathcal{P} : \mathbb{E}^{\mathbb{Q}}[(C)^+] < \infty \text{ and } \mathcal{H}(\mathbb{Q}; \mathbb{P}) < \infty\}$.

Proof of Proposition 2.6. We will distinguish two cases.

- (i) $\mathbb{E}^{\mathbb{P}}[(C)^-] = \infty$. In this case, essentially just using Jensen's inequality,

$$\begin{aligned} \mathbb{E}^{\mathbb{P}}[\exp(-C)] &\geq \mathbb{E}^{\mathbb{P}}[\exp((C)^-)\mathbb{1}_{\{C \leq 0\}}] = \mathbb{E}^{\mathbb{P}}[\exp((C)^-)(1 - \mathbb{1}_{\{C > 0\}})] \\ &\geq \mathbb{E}^{\mathbb{P}}[\exp((C)^-)] - \mathbb{P}(C > 0) \geq \exp(\mathbb{E}^{\mathbb{P}}[(C)^-]) - \mathbb{P}(C > 0) = \infty. \end{aligned}$$

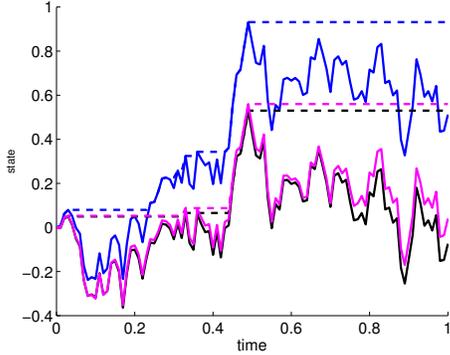
¹We wish to acknowledge drs. S. Thijssen (Radboud Universiteit Nijmegen) for providing us with this pseudocode and an implementation of the algorithm.

Algorithm 1 Monte-Carlo algorithm for relative entropy weighted minimization of the maximum of a Wiener process over a finite time interval

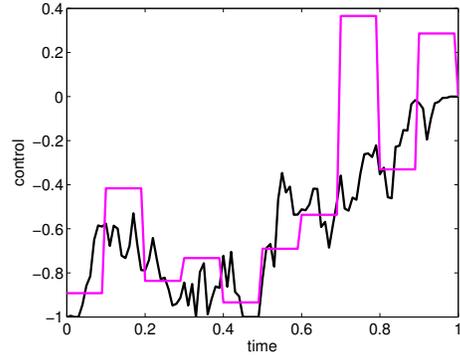
```

 $X_0 \leftarrow 0, Y_0 \leftarrow 0, M_0 \leftarrow 0$ 
for  $t = 0 : \Delta : T - \Delta$  do
  for  $n = 1, 2, \dots, N$  do                                     Perform  $N$  ‘rollouts’ until end time  $T$ 
     $X_t^n \leftarrow Y_t, M_t^n \leftarrow M_t$ 
    for  $s = t, t + \delta, \dots, T - \delta$  do
       $X_{s+\delta}^n \leftarrow X_s^n + \mathcal{N}(0, \delta)$ 
       $M_{s+\delta}^n \leftarrow \max(M_s^n, X_{s+\delta}^n)$                                      Compute the value of the cost for every rollout
    end for
  end for
   $\hat{U}_t \leftarrow \frac{\sum_{n=1}^N (X_{t+\Delta}^n - X_t^n) e^{-M_T^n}}{\sum_{n=1}^N e^{-M_T^n}}$                                      By law of large numbers, this is an approximation to (8)
  for  $s = t, t + \delta, \dots, t + \Delta - \delta$  do
     $X_{s+\delta} \leftarrow X_s + \mathcal{N}(0, \delta)$                                      Simulate a time step  $\Delta$  using control  $\hat{U}_t$ 
     $Y_{s+\delta} \leftarrow Y_s + \hat{U}_t \delta + X_{s+\delta} - X_s$                                       $X$  is uncontrolled Wiener process
     $M_s \leftarrow \max(M_s, Y_{s+\delta})$                                       $Y$  is (approximately) optimally controlled process
  end for
end for

```



(a) Sample path of Brownian motion with drift (solid line) along with its maximum process (dashed line), uncontrolled (blue), controlled with exact optimal drift (black) and controlled with approximately optimal drift, obtained by Algorithm 1 (magenta). It is seen how the controlled paths reduce the value of the maximum process at time $t = 1$, where the exact control outperforms the Monte-Carlo solution (but possibly with higher control cost).



(b) The corresponding exact (black) optimal drift ($u(t, Y)$), given by (11), (12), and approximate (magenta) optimal drift (\hat{U}_t), as obtained by Algorithm 1. The approximate control is piecewise constant with interval size $\Delta = 0.1$. We see how the approximation of the optimal drift may deviate from the exact solution due to discretization and statistical variation in the Monte-Carlo sampling.

Figure 1: Minimization of the maximum of a Brownian motion with controlled drift, subject to squared control costs.

Therefore the lower bound,

$$\inf_{\mathbb{Q} \in \mathcal{P}_0} J(\mathbb{Q}) \geq -\log \mathbb{E}^{\mathbb{P}}[\exp(-C)]$$

is immediate.

In the simple case $\mathbb{P}(C = -\infty) > 0$, we may choose $d\mathbb{Q}/d\mathbb{P} = \mathbb{1}_{\{C=-\infty\}}$, which establishes the required upper bound.

Assume now $\mathbb{P}(C = -\infty) = 0$ and (still) $\mathbb{E}^{\mathbb{P}}[(C)^-] = \infty$. Define measures \mathbb{Q}_n with densities $d\mathbb{Q}_n/d\mathbb{P} = \mathbb{1}_{[-n,0]}(C)/\mathbb{P}(C \in [-n,0])$. Since $\mathbb{P}(C \in [-n,0]) \leq \mathbb{P}(C \in (-\infty,0])$, we have by monotone convergence

$$\mathbb{E}^{\mathbb{Q}_n}[C] = \frac{\mathbb{E}[C\mathbb{1}_{[-n,0]}(C)]}{\mathbb{P}(C \in [-n,0])} \leq \frac{\mathbb{E}[C\mathbb{1}_{[-n,0]}(C)]}{\mathbb{P}(C \in (-\infty,0])} \downarrow -\frac{\mathbb{E}[(C)^-]}{\mathbb{P}(C \in (-\infty,0])} = -\infty$$

Furthermore $\mathcal{H}(\mathbb{Q}_n; \mathbb{P}) \downarrow -\log \mathbb{P}(C \in [0, \infty))$ as $n \rightarrow \infty$. Therefore $J(\mathbb{Q}_n) \downarrow -\infty$. This minimizing sequence establishes the required upper bound.

- (ii) $\mathbb{E}[(C)^-] < \infty$: Define $C_n := C \vee (-n)$. Since C_n is bounded from below, $\mathbb{E}^{\mathbb{P}}[\exp(-C_n)|C_n] < \infty$, and $\mathbb{P}(C_n < +\infty) = \mathbb{P}(C < +\infty) > 0$, so (FE) holds for C_n . Define $\mathcal{P}_0(C_n) = \{\mathbb{Q} \in \mathcal{P} : \mathbb{E}^{\mathbb{Q}}(C_n)^+ < \infty \text{ and } \mathcal{H}(\mathbb{Q}; \mathbb{P}) < \infty\}$. Using Proposition 2.5,

$$\inf_{\mathbb{Q} \in \mathcal{P}_0(C_n)} J(\mathbb{Q}) = -\log \mathbb{E}^{\mathbb{P}} \exp(-C_n).$$

We have trivially $\mathcal{P}_0 = \mathcal{P}_0(C_n)$. Furthermore by monotone convergence $-\log \mathbb{E}^{\mathbb{P}} \exp(-C_n) \downarrow -\log \mathbb{E}^{\mathbb{P}} \exp(-C)$ and, for $\mathbb{Q} \in \mathcal{P}_0$, $\mathbb{E}^{\mathbb{Q}} C_n \downarrow \mathbb{E}^{\mathbb{Q}} C$. Therefore, exchanging two infima,

$$\begin{aligned} \inf_{\mathbb{Q} \in \mathcal{P}_0} J(\mathbb{Q}) &= \inf_{\mathbb{Q} \in \mathcal{P}_0(C_n)} \left(\inf_{n \in \mathbb{N}} \mathbb{E}^{\mathbb{Q}} C_n + \mathcal{H}(\mathbb{Q}; \mathbb{P}) \right) = \inf_{n \in \mathbb{N}} (-\log \mathbb{E}^{\mathbb{P}} \exp(-C_n)) \\ &= -\log \mathbb{E}^{\mathbb{P}} \exp(-C). \end{aligned}$$

□

Acknowledgements

We are grateful to dr. ir. O. van Gaans (Mathematical Institute, Leiden University), for his comments on our work, and drs. S. Thijssen (Radboud University, Nijmegen) for providing us with his implementation of the Monte-Carlo scheme. We also wish to acknowledge dr. J. Lehec (Université Paris Dauphine) for providing us with a some details concerning his work [Leh13]. Also we wish to acknowledge the important suggestions of the associate editor and several anonymous reviewers. These comments have had significant impact upon this paper, for which we are very much indebted.

References

- [BD98] M Boué and P Dupuis. A variational representation for certain functionals of Brownian motion. *The Annals of Probability*, pages 1641–1659, 1998.
- [BH13] Ralf Banisch and Carsten Hartmann. Meshless discretization of LQ-type stochastic control problems. <http://arxiv.org/abs/1309.7497>, September 2013.
- [DE97] Paul Dupuis and Richard S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations (Wiley Series in Probability and Statistics)*. Wiley-Interscience, 1997.
- [F85] H Föllmer. An entropy approach to the time reversal of diffusion processes. *Stochastic Differential Systems Filtering and Control*, 8092, 1985.

- [F86] H Föllmer. *Time reversal on Wiener space*, volume 1158 of *Lecture Notes in Math.* Springer, Berlin, 1986.
- [Fle82] Wendell H Fleming. Logarithmic transformations and stochastic control. In *Advances in filtering and optimal stochastic control (Cocoyoc, 1982)*, volume 42 of *Lecture Notes in Control and Inform. Sci.*, pages 131–141. Springer, Berlin, 1982.
- [FS09] Wendell H. Fleming and Halil Mete Soner. *Controlled Markov Processes and Viscosity Solutions (Stochastic Modelling and Applied Probability)*. Springer, 2009.
- [HS91] AC Heinricher and RH Stockbridge. Optimal control of the running max. *SIAM journal on control and optimization*, 29(4):936–953, 1991.
- [Kap05] Hilbert J. Kappen. Linear Theory for Control of Nonlinear Stochastic Systems. *Physical Review Letters*, 95(20):200201, November 2005.
- [KS91] Ioannis Karatzas and Steven Shreve. *Brownian Motion and Stochastic Calculus (Graduate Texts in Mathematics)*. Springer, 1991.
- [Leh13] J Lehec. Representation formula for the entropy and functional inequalities. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, pages 1–18, 2013.
- [Mon13] Michael Monoyios. Malliavin calculus method for asymptotic expansion of dual control problems. *SIAM J. Financial Math.*, 4(1):884–915, 2013.
- [Nua06] David Nualart. *The Malliavin Calculus and Related Topics*. 2006.
- [OK91] Daniel L Ocone and Ioannis Karatzas. A generalized clark representation formula, with application to optimal portfolios. *Stochastics and Stochastic Reports*, 34(3-4):187–220, 1991.
- [RW94a] L C G Rogers and David Williams. *Diffusions, Markov processes, and martingales. Vol. 1.* Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester, second edition, 1994.
- [RW94b] L C G Rogers and David Williams. *Diffusions, Markov processes, and martingales. Vol. 2.* Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., New York, 2nd edition, 1994.
- [TK14] Sep Thijssen and H. J. Kappen. Path Integral Control and State Dependent Feedback. <http://arxiv.org/abs/1406.4026>, June 2014.
- [Tod06] Emanuel Todorov. Linearly-solvable Markov decision problems. In B Schölkopf, J Platt, and T Hoffman, editors, *Advances in Neural Information Processing Systems 19*, number 1, pages 1369–1376. MIT Press, Cambridge, MA, 2006.
- [Whi82] Peter Whittle. *Optimization over time*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, Ltd., Chichester, 1982.