

Minimising the time to a decision

Saul Jacka, Jon Warren, Peter Windridge

Department of Statistics, University of Warwick, Coventry, CV4 7AL.

January 8, 2010

Abstract

Suppose we have three independent copies of a regular diffusion on $[0,1]$ with absorbing boundaries. Of these diffusions, either at least two are absorbed at the upper boundary or at least two at the lower boundary. In this way, they determine a majority decision between 0 and 1. We show that the strategy that always runs the diffusion whose value is currently between the other two reveals the majority decision whilst minimising the total time spent running the processes.

Primary MSC classification: 93E20 Optimal stochastic control, Secondary MSC classification: 60J70 Applications of Brownian motions and diffusion theory (population genetics, absorption problems, etc.),

Keywords: optimal stochastic control, dynamic resource allocation, multiparameter processes, continuous bandits, recursive majority of three.

1 Introduction

Let X_1, X_2 and X_3 be three independent copies of a regular diffusion on $[0, 1]$ with absorbing boundaries. Eventually, either at least two of the diffusions are absorbed at the upper boundary of the interval or at least two are absorbed at the lower boundary. In this way, the diffusions determine a *majority decision* between 0 and 1.

In order to identify this decision we run the three processes –not simultaneously, but switching from one to another– until we observe at least two of them reaching a common boundary point. Our aim is to switch between the processes in a way that minimises the total time required to find the majority decision.

More precisely, we allocate our time between the three processes according to a suitably adapted $[0, \infty)^3$ -valued increasing process \mathcal{C} with $\sum_{i=1}^3 \mathcal{C}_i(t) = t$. Such a process is called a *strategy* and $\mathcal{C}_i(t)$ represents the amount of time spent observing X_i after $t \geq 0$ units of calendar time have elapsed. Accordingly, the process we observe is

$$X^{\mathcal{C}} \stackrel{\text{def}}{=} (X_1(\mathcal{C}_1(t)), X_2(\mathcal{C}_2(t)), X_3(\mathcal{C}_3(t)); t \geq 0),$$

and the *decision time* $\tau^{\mathcal{C}}$ for the strategy \mathcal{C} is the first time that two components of $X^{\mathcal{C}}$ are absorbed at the same end point of $[0, 1]$, i.e.

$$\tau^{\mathcal{C}} \stackrel{\text{def}}{=} \inf\{t \geq 0 : X_i^{\mathcal{C}}(t) = X_j^{\mathcal{C}}(t) \in \{0, 1\} \text{ for distinct } i, j\}.$$

In this paper we find a strategy \mathcal{C}^* that minimises this time. Roughly speaking, \mathcal{C}^* runs whichever diffusion is currently observed to have “middle value” (see Lemma 1.4 for a precise description). Our main theorem is that the decision time $\tau^{\mathcal{C}^*}$ of this strategy is the *stochastic minimum* of all possible decision times, i.e.

Theorem 1.1. *The decision time $\tau^{\mathcal{C}^*}$ of the “run the middle” strategy \mathcal{C}^* given in lemma 1.4 satisfies*

$$\mathbb{P}(\tau^{\mathcal{C}^*} > t) = \inf_{\mathcal{C}} \mathbb{P}(\tau^{\mathcal{C}} > t), \text{ for every } t \geq 0.$$

where the infimum is taken over all strategies and $\tau^{\mathcal{C}}$ is the corresponding decision time.

This model fits into the existing literature on optimal dynamic resource allocation (see 1.1 below for a brief review) but our original motivation for studying it was to gain an understanding of the problem of evaluating the “recursive majority of three” function on random input. The latter can be described as follows – take the complete ternary tree on n levels, place independent Bernoulli(p) ($0 < p < 1$) random variables on each of the 3^n leaves and recursively define the internal nodes to take the majority value of their children. We wish to determine the value of the root node, but may only accrue knowledge about the tree by sequentially observing leaves, paying £1 each time for the privilege. It remains an open problem to determine the strategy with least expected cost, r_n . However, r is sub-multiplicative (i.e. $r_{n+m} \leq r_n r_m$ for any $n, m \in \mathbb{N}$) and so, by Fekete’s lemma

$$\gamma \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} r_n^{1/n}$$

exists with the trivial bounds $2 \leq \gamma \leq 3$. The value of γ , despite attracting the attention of several investigators, is not known (see section 6.2). Our idea was to study it by considering a continuous approximation to the large n tree. It was this continuous approximation that inspired the diffusive model introduced in this paper, but the reader is warned that the results we present here do not shed light on the value of γ .

However, the problem of switching between diffusions is worthwhile in its own right. It has a similar flavour to the continuous multi-armed bandit problem but does not seem to have the same mathematical anatomy. Nevertheless there is an interesting structure to be revealed – in particular we make use of the heuristic equation (2.17) in order to evaluate the value function for the discounted problem, and the same equation plays a central role in proving the much stronger stochastic minimality property.

1.1 Dynamic resource allocation

The problem we have described is one of optimal dynamic resource allocation in continuous time. The most widely studied example of this is the continuous multi-armed bandit problem (see, for example, El Karoui and Karatzas [6], Mandelbaum and Kaspi [12]). Here, a gambler chooses the rates at which he will pull the arms on different slot machines. Each slot machine rewards the gambler at rates which follow a stochastic process independent of the reward processes for the other machines. These general bandit problems find application in several fields where agents must choose between exploration and exploitation, typified in economics and clinical trials. An optimal strategy is surprisingly easy to describe. Associated to each machine is a process known as the Gittins index, which may be interpreted as the equitable surrender value. It is a celebrated theorem

that at each instant, we should play whichever machine currently has the largest Gittins index. This is in direct analogy to the discrete time result of Gittins and Jones [8].

There is no optimal strategy of index type for our problem. This reflects the fact that the reward processes associated to running each of the diffusions are not independent – once two of the diffusions are absorbed, it may be pointless to run the third.

In [16], a different dynamic allocation problem is considered. It has a similar flavour in that one must choose the rates at which to run two Brownian motions on $[0, 1]$, and we stop once *one* of the processes hits an endpoint. The rates are chosen to maximise a terminal payoff, as specified by a function defined on the boundary of the square (the generalisation of this problem to several Brownian motions is considered in [21]). An optimal strategy is determined by a partition of the square into regions of indifference, preference for the first Brownian motion and preference for the second. However, there is no notion of a reward (cost) being *accrued* as in our problem. So, our problem, in which time is costly *and* there is a terminal cost of minus infinity for finishing on a part of $\partial\mathcal{S}$ which does not give a majority decision could be seen as lying between continuous bandits and the Brownian switching in [16].

1.2 Overview of paper

The rest of the paper is laid out as follows. Section 1.3 contains a precise statement of the problem and our assumptions and a clarification of Theorem 1.1. The proof of this theorem begins in Section 2, where we show that the Laplace transform of the distribution of the decision time $\tau^{\mathcal{C}^*}$ solves certain differential equations. This fact is then used in Section 3 to show that the tail of $\tau^{\mathcal{C}^*}$ solves, in a certain sense, the appropriate Hamilton-Jacobi-Bellman PDE. From here, martingale optimality arguments complete the proof. Section 4 shows the existence and uniqueness of the strategy \mathcal{C}^* and in Section 5 we explain the connection between the controlled process and doubly perturbed diffusions. In the final section, we make a conjecture about an extension to the model and then, to close, we ask a few questions relating to the discrete recursive majority of three problem that motivated us originally.

1.3 Problem statement and solution

We are given a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ supporting three independent Itô diffusions $(X_i(t), t \geq 0)$, $i \in V = \{1, 2, 3\}$, each of which is started in the unit interval $[0, 1]$ and absorbed at the endpoints. The diffusions all satisfy the same stochastic differential equation

$$dX_i(t) = \sigma(X_i(t))dB_i(t) + \mu(X_i(t))dt, \quad t \geq 0, \quad (1.1)$$

where $\sigma : [0, 1] \rightarrow (0, \infty)$ is continuous, $\mu : [0, 1] \rightarrow \mathbb{R}$ is Borel and $(B_i(t), t \geq 0)$, $i \in V$ are independent Brownian motions.

We denote by \mathcal{S} the unit cube $[0, 1]^3$, by \mathbb{R}_+ the set of non-negative real numbers $[0, \infty)$ and \preceq its usual partial order on \mathbb{R}_+^3 . It is assumed that we have a standard Markovian setup, i.e. there is a family of probability measures $(\mathbb{P}_x; x \in \mathcal{S})$ under which $X(0) = x$ almost surely and the filtration $\mathcal{F}_i = (\mathcal{F}_i(t); t \geq 0)$ generated by X_i is augmented to satisfy the usual conditions.

From here, we adopt the framework for continuous dynamic allocation models proposed by Mandelbaum in [15]. This approach relies on the theory of multiparameter time changes; the reader may consult Appendix A for a short summary of this.

For $\eta \in \mathbb{R}_+^3$ we define the σ -algebra

$$\mathcal{F}(\eta) \stackrel{\text{def}}{=} \sigma(\mathcal{F}_1(\eta_1), \mathcal{F}_2(\eta_2), \mathcal{F}_3(\eta_3)),$$

which corresponds to the information revealed by running X_i for η_i units of time. The family $(\mathcal{F}(\eta); \eta \in \mathbb{R}_+^3)$ is called a *multiparameter filtration* and satisfies the “usual conditions” of right continuity, completeness and property (F4) of Cairoli and Walsh [2]. It is in terms of this filtration that we define the sense in which our strategies must be adapted.

A *strategy* is an \mathbb{R}_+^3 -valued stochastic process

$$\mathcal{C} = (\mathcal{C}_1(t), \mathcal{C}_2(t), \mathcal{C}_3(t); t \geq 0)$$

such that

(C1) for $i = 1, 2, 3$, $\mathcal{C}_i(0) = 0$ and $\mathcal{C}_i(\cdot)$ is nondecreasing,

(C2) for every $t \geq 0$, $\mathcal{C}_1(t) + \mathcal{C}_2(t) + \mathcal{C}_3(t) = t$ and

(C3) $\mathcal{C}(t)$ is a stopping *point* of the multiparameter filtration $(\mathcal{F}(\eta); \eta \in \mathbb{R}_+^3)$, i.e.

$$\{\mathcal{C}(t) \preceq \eta\} \in \mathcal{F}(\eta) \text{ for every } \eta \in \mathbb{R}_+^3.$$

Remark 1.2. In the language of multiparameter processes, \mathcal{C} is an *optional increasing path* after Walsh [22].

Remark 1.3. Conditions (C1) and (C2) together imply that for any $s \leq t$, $|\mathcal{C}_i(t) - \mathcal{C}_i(s)| \leq t - s$. It follows that the measure $d\mathcal{C}_i$ is absolutely continuous and so it makes sense to talk about the *rate* $\dot{\mathcal{C}}_i(t) = d\mathcal{C}_i(t)/dt$, $t \geq 0$, at which X_i is to be run.

The interpretation is that $\mathcal{C}_i(t)$ models the total amount of time spent running X_i by calendar time t , and accordingly, the *controlled process* $X^{\mathcal{C}}$ is defined by

$$X^{\mathcal{C}}(t) \stackrel{\text{def}}{=} (X_1(\mathcal{C}_1(t)), X_2(\mathcal{C}_2(t)), X_3(\mathcal{C}_3(t))), t \geq 0.$$

Continuity of \mathcal{C} implies that $X^{\mathcal{C}}$ is a continuous process in \mathcal{S} that is adapted to the (one parameter) filtration $\mathcal{F}^{\mathcal{C}}$ defined by

$$\mathcal{F}^{\mathcal{C}}(t) \stackrel{\text{def}}{=} \{F \in \mathcal{F} : F \cap \{\mathcal{C}(t) \preceq \eta\} \in \mathcal{F}(\eta) \text{ for every } \eta \in \mathbb{R}_+^3\}, t \geq 0.$$

The *decision time* $\tau^{\mathcal{C}}$ for a time allocation strategy \mathcal{C} is the first time that $X^{\mathcal{C}}$ hits the *decision set*

$$D \stackrel{\text{def}}{=} \{(x_1, x_2, x_3) \in \mathcal{S} : x_i = x_j \in \{0, 1\} \text{ for some } 1 \leq i < j \leq 3\}$$

The objective is to find a strategy whose associated decision time is a stochastic minimum. Clearly, it is possible to do very badly by only ever running one of the processes

as a decision may never be reached (these strategies do not need to be ruled out in our model). A more sensible thing to do is to pick two of the processes, and run them until they are absorbed. Only if they disagree do we run the third. This strategy is much better than the pathological one (the decision time is almost surely finite!) but we can do better.

We do not think it is obvious what the best strategy is. In the situation that $X_1(0)$ is close to zero and $X_3(0)$ close to one, it is probable that X_1 and X_3 will be absorbed at different end points of $[0, 1]$. So, if $X_2(0)$ is close to 0.5 say, it seems likely that X_2 will be pivotal and so we initially run it, even though X_1 and X_3 might be absorbed much more quickly. Our guess is to run the diffusion whose value lies between that of the other two processes. But if all the processes are near one, it is not at all clear this is the best thing to do. For example, one could be tempted to run the process with largest value in the hope that it will give a decision very quickly.

It turns out that we must always “run the middle”. That is, if, at any moment $t \geq 0$, we have $X_1^C(t) < X_2^C(t) < X_3^C(t)$, then we should run X_2 exclusively until it hits $X_1^C(t)$ or $X_3^C(t)$. We need not concern ourselves with what happens when the processes are equal. This is because there is, almost surely, only one strategy that runs the middle of the three diffusions when they are separated. To state this result, let us say that for a strategy \mathcal{C} , component \mathcal{C}_i increases at time $t \geq 0$ if $\mathcal{C}_i(u) > \mathcal{C}_i(t)$ for every $u > t$.

Lemma 1.4. *There exists a unique time allocation strategy \mathcal{C}^* such that \mathcal{C}_i^* increases only at times $t \geq 0$ such that $X_j^{C^*}(t) \leq X_i^{C^*}(t) \leq X_k^{C^*}(t)$ under some labelling $\{i, j, k\} = V$ of the processes.*

If \mathcal{C} is any other strategy with this property, then $\mathcal{C}(t) = \mathcal{C}^(t)$ for all $t \geq 0$ almost surely (with respect to any of the measures \mathbb{P}_x).*

This lemma is proved in section 4 and Theorem 1.1 states that \mathcal{C}^* gives a stochastic minimum for the decision time.

In the sequel, the drift term μ is assumed to vanish. This is not a restriction, for if a drift is present we may eliminate it by rewriting the problem in natural scale.

2 The Laplace transform of the distribution of $\tau^{\mathcal{C}^*}$

The proof of Theorem 1.1 begins by computing the Laplace transform

$$\hat{v}_r(x) \stackrel{\text{def}}{=} \mathbb{E}_x(\exp(-r\tau^{\mathcal{C}^*})),$$

of the distribution of the decision time. This non-trivial task is carried out using the “guess and verify” method. Loosely, the guess is inspired by comparing the payoffs of doing something optimal against doing something nearly optimal. This leads to a surprisingly tractable heuristic equation from which \hat{v}_r can be recovered.

The argument which motivates the heuristic proceeds as follows. From any strategy \mathcal{C} it is possible to construct (but we omit the details) another strategy, $\hat{\mathcal{C}}$, that begins by running X_1 for some small time $h > 0$ (i.e. $\hat{\mathcal{C}}(t) = (t, 0, 0)$ for $0 \leq t \leq h$) and then does not run X_1 again until \mathcal{C}_1 exceeds h , if ever. In the meantime, $\hat{\mathcal{C}}_2$ and $\hat{\mathcal{C}}_3$ essentially follow \mathcal{C}_2 and \mathcal{C}_3 with the effect that once \mathcal{C}_1 exceeds h , \mathcal{C} and $\hat{\mathcal{C}}$ coincide.

This means that if the amount of time, $\mathcal{C}_1(\tau^{\mathcal{C}})$, that \mathcal{C} spends running X_1 is at least h , then $\tau^{\hat{\mathcal{C}}}$ and $\tau^{\mathcal{C}}$ are identical. On the other hand, if $\mathcal{C}_1(\tau^{\mathcal{C}}) < h$, then $\hat{\mathcal{C}}$ runs X_1 for longer than \mathcal{C} , with some of the time $\hat{\mathcal{C}}$ spends running X_1 being wasted. In fact, outside a set with probability $o(h)$ we have

$$\tau^{\hat{\mathcal{C}}} = \tau^{\mathcal{C}} + (h - T_1)^+, \quad (2.2)$$

where $T_i = \mathcal{C}_i(\tau^{\mathcal{C}})$ is the amount of time that \mathcal{C} spends running X_i while determining the decision.

We compare $\hat{\mathcal{C}}$ with the strategy that runs X_1 for time h and then behaves *optimally*. If we suppose that \mathcal{C}^* itself is optimal and recall that \hat{v}_r is the corresponding payoff, this yields the inequality

$$\mathbb{E}_x \left(\exp(-r\tau^{\hat{\mathcal{C}}}) \right) \leq \mathbb{E}_x \left(\exp(-rh) \hat{v}_r(X_1(h), X_2(0), X_3(0)) \right). \quad (2.3)$$

Now, we take $\mathcal{C} = \mathcal{C}^*$ and use (2.2) to see that the left hand side of (2.3) is equal to

$$\mathbb{E}_x \left(\exp(-r(\tau^{\mathcal{C}^*} + (h - T_1)^+)) \right) + o(h),$$

which, in turn, may be written as

$$\hat{v}_r(x) + \mathbb{E}_x \left((\exp(-r(\tau^{\mathcal{C}^*} + h)) - \exp(-r\tau^{\mathcal{C}^*})) \mathbb{1}_{[T_i=0]} \right) + o(h). \quad (2.4)$$

On the other hand, if we assume \hat{v}_r is suitably smooth, the right hand side of (2.3) is

$$\hat{v}_r(x) + h (\mathcal{G}^1 - r) \hat{v}_r(x) + o(h), \quad x_1 \in (0, 1). \quad (2.5)$$

where we have introduced the differential operator \mathcal{G}^i defined by

$$\mathcal{G}^i f(x) \stackrel{\text{def}}{=} \frac{1}{2} \sigma^2(x_i) \frac{\partial^2}{\partial x_i^2} f(x), \quad x_i \in (0, 1).$$

After substituting these expressions back into (2.3) and noticing that there was nothing special about choosing X_1 to be the process that we moved first, we see that

$$\mathbb{E}_x \left(\exp(-r(\tau^{\mathcal{C}^*} + h)) - \exp(-r\tau^{\mathcal{C}^*}); T_i = 0 \right) \leq h (\mathcal{G}^i - r) \hat{v}_r(x) + o(h), \quad (2.6)$$

for each $x_i \in (0, 1)$ and $i \in V$.

Dividing both sides by h , and taking the limit $h \rightarrow 0$ yields the inequality

$$(\mathcal{G}^i - r) \hat{v}_r(x) \leq -r \mathbb{E}_x \left(\exp(-r\tau^{\mathcal{C}^*}); T_i = 0 \right). \quad (2.7)$$

Now, in some simpler, but nevertheless related problems, we can show that (2.7) is true with an *equality* replacing the inequality. This prompts us to *construct* a function satisfying (2.7) with equality. Our effort culminates in

Lemma 2.1. *There exists a continuous function $h_r : \mathcal{S} \rightarrow \mathbb{R}$ such that*

- $h_r(x) = 1$ for $x \in D$,

- the partial derivatives $\frac{\partial^2 \hat{h}_r}{\partial x_i \partial x_j}$ exist and are continuous on $\{x \in \mathcal{S} \setminus D : x_i, x_j \in (0, 1)\}$ (for any $i, j \in V$ not necessarily distinct) and
- furthermore, for each $i \in V$ and $x \notin D$ with $x_i \in (0, 1)$,

$$(\mathcal{G}^i - r) h_r(x) = -r \hat{f}_r^i(x),$$

where $\hat{f}_r^i(x) \stackrel{\text{def}}{=} \mathbb{E}_x (\exp(-r\tau^{C^*}) \mathbb{1}_{[T_i=0]})$.

Proof. We begin by factorising $\hat{f}_r^i(x)$ into a product of Laplace transforms of diffusion exit time distributions. This factorisation is useful as it allows us to construct h by solving a series of ordinary differential equations. Note that in this proof, we will typically suppress the r dependence for notational convenience.

The diffusions all obey the same stochastic differential equation and so we lose nothing by assuming that the components of x satisfy $0 \leq x_1 \leq x_2 \leq x_3 \leq 1$. Further, we suppose that $x \notin D$ because otherwise $T_i = 0$ \mathbb{P}_x -almost-surely.

In this case, $T_2 > 0$ \mathbb{P}_x -almost surely, because for any $t > 0$, there exist times $t_1, t_3 < t/2$ at which $X_1(t_1) < x_1 \leq x_2 \leq x_3 < X_3(t_3)$ and so it is certain our strategy allocates time to X_2 . It follows that $\hat{f}^2(x)$ vanishes.

Now consider \hat{f}^1 . There is a \mathbb{P}_x -negligible set off which $T_1 = 0$ occurs if, and only if, both of the independent diffusions X_2 and X_3 exit the interval $(X_1(0), 1)$ at the upper boundary. Furthermore, τ^{C^*} is just the sum of the exit times. That is, if

$$\mathbf{m}_a^{(i)} \stackrel{\text{def}}{=} \inf\{t > 0 : X_i(t) = a\}, \quad a \in I, i \in V, \quad (2.8)$$

then

$$\hat{f}^1(x) = \mathbb{E}_x \left(\exp(-r(\mathbf{m}_1^{(2)} + \mathbf{m}_1^{(3)})) \mathbb{1}_{[\mathbf{m}_1^{(2)} < \mathbf{m}_{x_1}^{(2)}, \mathbf{m}_1^{(3)} < \mathbf{m}_{x_1}^{(3)}]} \right).$$

Using independence of X_2 and X_3 , we have the factorisation

$$\hat{f}^1(x) = \prod_{i=2}^3 \mathbb{E}_x \left(\exp(-r\mathbf{m}_1^{(i)}) \mathbb{1}_{[\mathbf{m}_1^{(i)} < \mathbf{m}_{x_1}^{(i)}]} \right).$$

Note that our assumption $x \notin D$ guarantees that $x_1 < 1$.

To write this more cleanly, let us introduce, for $0 \leq a < b \leq 1$, the functions

$$h_{a,b}^+(u) \stackrel{\text{def}}{=} \mathbb{E}_u \left(\exp(-r\mathbf{m}_b^{(1)}); \mathbf{m}_b^{(1)} < \mathbf{m}_a^{(1)} \right),$$

where the expectation operator \mathbb{E}_u corresponds to the (marginal) law of X_1 when it begins at $u \in [0, 1]$.

The diffusions obey the same SDE, and so

$$\hat{f}^1(x) = h_{x_1,1}^+(x_2) h_{x_1,1}^+(x_3),$$

Similarly,

$$\hat{f}^3(x) = h_{0,x_3}^-(x_1) h_{0,x_3}^-(x_2)$$

where

$$h_{a,b}^-(u) \stackrel{\text{def}}{=} \mathbb{E}_u \left(\exp(-r\mathbf{m}_a^{(i)}); \mathbf{m}_a^{(i)} < \mathbf{m}_b^{(i)} \right).$$

We take, as building blocks for the construction of h , the functions $h_{0,1}^\pm$, abbreviated to h^\pm in the sequel. The regularity of each of our (non-singular) diffusions together with the Markov property shows that if $a < b$, $u \in [a, b]$ then

$$h^+(u) = h_{a,b}^+(u)h^+(b) + h_{a,b}^-(u)h^+(a)$$

and

$$h^-(u) = h_{a,b}^+(u)h^-(b) + h_{a,b}^-(u)h^-(a).$$

Solving these equations gives

$$h_{a,b}^+(u) = \frac{h^-(a)h^+(u) - h^-(u)h^+(a)}{h^-(a)h^+(b) - h^-(b)h^+(a)} \quad (2.9)$$

and

$$h_{a,b}^-(u) = \frac{h^-(u)h^+(b) - h^-(b)h^+(u)}{h^-(a)h^+(b) - h^-(b)h^+(a)}. \quad (2.10)$$

The functions h^+ and h^- are C^2 on $(0, 1)$ and continuous on $[0, 1]$. Furthermore, they solve $\mathcal{G}f = rf$ where $\mathcal{G}f \stackrel{\text{def}}{=} \frac{1}{2}\sigma^2(\cdot)f''$. In light of this, and remembering our assumption that the components of x are ordered, we will look for functions λ^+ and λ^- of x_1 and x_3 such that

$$h(x) = \lambda^-(x_1, x_3)h^-(x_2) + \lambda^+(x_1, x_3)h^+(x_2) \quad (2.11)$$

has the desired properties. For other values of $x \notin D$ we will define h by symmetry.

To get started, use (2.9) and (2.10) so see that $\hat{f}^i(x)$ has a linear dependence on $h^+(x_2)$ and $h^-(x_2)$, that is, there are functions ψ_\pm^i such that

$$\hat{f}^i(x) = \psi_-^i(x_1, x_3)h^-(x_2) + \psi_+^i(x_1, x_3)h^+(x_2).$$

For example,

$$\begin{aligned} \psi_+^1(x_1, x_3) &= \frac{h^-(x_1)h^+(x_3) - h^-(x_3)h^+(x_1)}{h^-(x_1)} \\ \psi_-^1(x_1, x_3) &= -\frac{h^+(x_1)}{h^-(x_1)}\psi_+^1(x_1, x_3) \end{aligned}$$

Linearity of the operator $(\mathcal{G}^i - r)$ and linear independence of h^- and h^+ then show the requirement that $(\mathcal{G}^i - r)h = -r\hat{f}^i$ boils down to requiring

$$(\mathcal{G}^i - r)\lambda_\pm = -r\psi_\pm^i.$$

Of course, the corresponding homogenous (eigenfunction) problems are solved with linear combinations of h^+ and h^- – what remains is the essentially computational task of finding particular integrals and some constants. This endeavour begins with repeated application of Lagrange’s variation of parameters method, determining constants using the boundary conditions $h(x) = 1$ for $x \in D$ where possible. Eventually we are left wanting only for real constants, an unknown function of x_1 and a function of x_3 . At this point we appeal to the “smooth pasting” conditions

$$\left(\frac{\partial}{\partial x_i} - \frac{\partial}{\partial x_j} \right) h \Big|_{x_i=x_j} = 0, \quad i, j \in V. \quad (2.12)$$

After some manipulation, we are furnished with differential equations for our unknown functions and equations for the constants. These we solve with little difficulty and, in doing so, determine that

$$\begin{aligned}\lambda_-(x_1, x_3) &= h^-(x_1) - h^+(x_1)h^+(x_3) \int_{x_3}^1 \frac{\frac{\partial}{\partial u} h^-(u)}{h^+(u)^2} du \\ &\quad + h^-(x_1)h^+(x_3) \int_0^{x_1} \frac{\frac{\partial}{\partial u} h^+(u)}{h^-(u)^2} du \\ &\quad + \frac{2rh^-(x_3)}{\phi} \int_0^{x_1} \left(\frac{h^+(u)}{\sigma(u)h^-(u)} \right)^2 (h^-(x_1)h^+(u) - h^-(u)h^+(x_1)) du,\end{aligned}$$

and

$$\begin{aligned}\lambda_+(x_1, x_3) &= h^+(x_3) + h^-(x_1)h^-(x_3) \int_0^{x_1} \frac{\frac{\partial}{\partial u} h^+(u)}{h^-(u)^2} du \\ &\quad - h^-(x_1)h^+(x_3) \int_{x_3}^1 \frac{\frac{\partial}{\partial u} h^-(u)}{h^+(u)^2} du \\ &\quad + \frac{2rh^+(x_1)}{\phi} \int_{x_3}^1 \left(\frac{h^-(u)}{\sigma(u)h^+(u)} \right)^2 (h^-(u)h^+(x_3) - h^-(x_3)h^+(u)) du,\end{aligned}$$

where ϕ denotes the constant value of $h^-(u)\frac{\partial}{\partial u}h^+(u) - h^+(u)\frac{\partial}{\partial u}h^-(u)$.

These expressions for λ^\pm are valid for any x not lying in D with weakly ordered components; so h is defined outside of D via (2.11). Naturally, we define h to be equal to one on D .

Having defined h , we now show that it is continuous and has the required partial derivatives. Continuity is inherited from h^+ and h^- on the whole of \mathcal{S} apart from at the exceptional points $(0, 0, 0)$ and $(1, 1, 1)$ in D . For these two points, a few lines of justification is needed. We shall demonstrate continuity at the origin, continuity at the upper right hand corner $(1, 1, 1)$ follows by the same argument. Let x^n be a sequence of points in S that converge to $(0, 0, 0)$; we must show $h(x^n) \rightarrow h(0, 0, 0) = 1$. Without loss of generality assume that the components of x^n are ordered $x_1^n \leq x_2^n \leq x_3^n$ and that x^n is not in D (if $x^n \in D$, then $h(x^n) = 1$ and it may be discarded from the sequence). By examining the expressions for λ^\pm , we see that it is sufficient to check that

$$(i) \lambda^-(x_1^n, x_3^n) \rightarrow 1 \text{ and } (ii) h^+(x_2^n)\lambda^+(x_1^n, x_3^n) \rightarrow 0.$$

For (i), the only doubt is that the term involving the first integral in the expression for λ^- does not vanish in the limit. The fact that it does can be proved by the Dominated Convergence Theorem. The term is

$$h^+(x_1^n)h^+(x_3^n) \int_{x_3^n}^1 \frac{\frac{\partial}{\partial u} h^-(u)}{h^+(u)^2} du = \int_0^1 \mathbb{1}_{[u > x_3^n]} \frac{h^+(x_1^n)h^+(x_3^n)}{h^+(u)^2} \frac{\partial}{\partial u} h^-(u) du.$$

The ratio $\frac{h^+(x_1^n)h^+(x_3^n)}{h^+(u)^2}$ is bounded above by one when $u > x_3^n \geq x_1^n$ since h^+ is increasing. Further, the derivative of h^- is integrable and so the integrand is dominated by an integrable function, and converges to zero.

For the second limit (ii), there are two terms to check. Firstly, that

$$h^+(x_2^n)h^-(x_1^n)h^+(x_3^n) \int_{x_3^n}^1 \frac{\frac{\partial}{\partial u}(u)}{h^+(u)^2} du \rightarrow 0$$

follows from essentially the same argument as before. The second term of concern is

$$h^+(x_1^n) \int_{x_3^n}^1 \left(\frac{h^-(u)}{\sigma(u)h^+(u)} \right)^2 (h^-(u)h^+(x_3^n) - h^-(x_3^n)h^+(u)) du.$$

Again, one may write this as the integral of a dominated function (recalling that σ is bounded away from zero) that converges to zero. Thus, the integral above converges to zero as required.

Now that we have established continuity of h , we can begin tackling the partial derivatives. When the components of x are distinct, differentiability comes from that of our building blocks h^+ and h^- . It is at the switching boundaries, when two or more components are equal, where we have to be careful. The key here is to remember that we constructed h to satisfy the smooth pasting property (2.12) – this allows us to show that the one-sided partial derivatives are equal in $(0, 1)$. For example, provided the limit exists,

$$\left. \frac{\partial}{\partial x_1} h(x_1, x_2, x_3) \right|_{x_1=x_2=x_3} = \lim_{h \rightarrow 0} \frac{1}{h} (h(x_1 + h, x_1, x_1) - h(x_1 - h, x_1, x_1)).$$

Using (2.11) and the differentiability of λ , the limit from above is

$$\left. \frac{\partial}{\partial x_3} (\lambda^-(x_1, x_3)h^-(x_2) + \lambda^+(x_1, x_3)h^+(x_2)) \right|_{x_1=x_2=x_3}.$$

This is equal to the limit from below,

$$\left. \frac{\partial}{\partial x_1} (\lambda^-(x_1, x_3)h^-(x_2) + \lambda^+(x_1, x_3)h^+(x_2)) \right|_{x_1=x_2=x_3},$$

by the smooth pasting property. The other first order partial derivatives exist by similar arguments. Note that we do not include in our hypothesis the requirement that these first order partial derivatives exist at the boundary points of I .

The second order derivatives are only slightly more laborious to check. As before it is at switching boundaries where we must take care in checking that the limits from above and below agree. For the partial derivatives $\frac{\partial^2}{\partial x_i^2} h$ at a point x not in D with $x_i \in (0, 1)$, it is continuity of \hat{f}^i at x that allows us to equate the limits and show that the result is continuous, rather than smooth pasting. For the mixed partial derivatives, a priori, we don't have this helping hand. Instead, when two components are equal, we can always assume that one is the "middle" component that enters through the terms h^+ and h^- in (2.11) while the other is an "upper" or "lower" term that enters through λ^+ and λ^- . This makes it easy to check that the partial derivations of h commute at x . Now we can use the smooth pasting condition (2.12) to show that, for example,

$$\left. \frac{\partial^2}{\partial x_3 \partial x_2} h(x_1, x_2, x_3) \right|_{x_1=x_2=x_3} = \left. \frac{\partial^2}{\partial x_1 \partial x_2} h(x_1, x_2, x_3) \right|_{x_1=x_2=x_3}.$$

Thus, h satisfies all the properties we required. □

From here, we need a verification lemma to check that the function we constructed really is equal to \hat{v}_r . The following result does just that, and, as a corollary, shows that \hat{v}_r is maximal among Laplace transforms of decision time distributions (note that this is weaker than the stochastic minimality claimed in Theorem 1.1). The result is essentially that Bellman's principle of optimality holds (specialists in optimal control will notice that the function we constructed in Lemma 2.1 satisfies the Hamilton-Jacobi-Bellman PDE).

Lemma 2.2. *Suppose that $h_r : \mathcal{S} \rightarrow \mathbb{R}$ satisfies*

- h_r is continuous on \mathcal{S}
- for $i, j \in V$, $\frac{\partial^2 h_r}{\partial x_i \partial x_j}$ exists and is continuous on $\{x \in \mathcal{S} \setminus D : x_i, x_j \in (0, 1)\}$.
- $h_r(x) = 1$ for $x \in D$.
- $(\mathcal{G}^i - r) h_r(x) \leq 0$

Then,

$$h_r(x) \geq \sup_{\mathcal{C}} \mathbb{E}_x (\exp(-r\tau^{\mathcal{C}})).$$

Furthermore, if $(\mathcal{G}^i - r) h_r(x)$ vanishes whenever $x_j \leq x_i \leq x_k$ (under some labelling) then,

$$h_r(x) = \hat{v}_r(x) = \mathbb{E}_x (\exp(-r\tau^{\mathcal{C}^*})).$$

Proof. Let \mathcal{C} be an arbitrary strategy and define the function $g : \mathcal{S} \times [0, \infty) \rightarrow \mathbb{R}$ by $g(x, t) \stackrel{\text{def}}{=} \exp(-rt)h_r(x)$. Then, by hypothesis, g is $C^{2,1}$ on $(0, 1)^3 \times [0, \infty)$. Thus, if dist denotes Euclidean distance and $\rho_n \stackrel{\text{def}}{=} \inf\{t \geq 0 : \text{dist}(X^{\mathcal{C}}(t), \partial\mathcal{S}) < n^{-1}\}$, Ito's formula shows that,

$$\begin{aligned} g(X^{\mathcal{C}}(\rho_n), \rho_n) - g(X^{\mathcal{C}}(0), 0) &= \sum_i \int_0^{\rho_n} \frac{\partial}{\partial x_i} g(X^{\mathcal{C}}(s), s) dX_i^{\mathcal{C}}(s) \\ &+ \int_0^{\rho_n} \frac{\partial}{\partial t} g(X^{\mathcal{C}}(s), s) ds \\ &+ \frac{1}{2} \sum_{i,j} \int_0^{\rho_n} \frac{\partial^2}{\partial x_i \partial x_j} g(X^{\mathcal{C}}(s), s) d[X_i^{\mathcal{C}}, X_j^{\mathcal{C}}]_s. \end{aligned}$$

Theorem A.1 implies $[X_i^{\mathcal{C}}]_s = [X_i]_{\mathcal{C}_i(s)}$ and that $X_i^{\mathcal{C}}$ and $X_j^{\mathcal{C}}$ are orthogonal martingales. Hence, using absolute continuity of \mathcal{C} and Proposition 1.5, Chapter V of [19],

$$\begin{aligned} g(X^{\mathcal{C}}(\rho_n), \rho_n) - g(X^{\mathcal{C}}(0), 0) &= \sum_i \int_0^{\rho_n} \frac{\partial}{\partial x_i} g(X^{\mathcal{C}}(s), s) dX_i^{\mathcal{C}}(s) \\ &+ \sum_i \int_0^{\rho_n} \exp(-rs) (\mathcal{G}^i - r) h(X^{\mathcal{C}}(s)) \dot{\mathcal{C}}_i(s) ds. \end{aligned}$$

The integrand of the stochastic integral against the square integrable martingale $X_i^{\mathcal{C}}$ is continuous and hence bounded on each compact subset of $(0, 1)^3$. Thus, the integral's expectation vanishes, i.e.

$$\mathbb{E}_x \left(\int_0^{\rho_n} \frac{\partial}{\partial x_i} g(X^{\mathcal{C}}(s), s) dX_i^{\mathcal{C}}(s) \right) = 0$$

Next, the fact that $(\mathcal{G}^i - r)h$ is non-positive gives

$$\mathbb{E}_x \left(\int_0^{\rho_n} \exp(-rs) (\mathcal{G}^i - r) h(X^{\mathcal{C}}(s)) \dot{\mathcal{C}}_i(s) ds \right) \leq 0,$$

and so

$$\mathbb{E}_x (\exp(-r\rho_n) h(X^{\mathcal{C}}(\rho_n))) - h(x) \leq 0. \quad (2.13)$$

Now, the times ρ_n taken for $X^{\mathcal{C}}$ to come within distance n^{-1} of the boundary of \mathcal{S} converge to $\rho \stackrel{\text{def}}{=} \inf\{t \geq 0 : X^{\mathcal{C}}(t) \in \partial\mathcal{S}\}$ as $n \rightarrow \infty$. So, the continuity of h and the Dominated Convergence Theorem together imply

$$\mathbb{E}_x (\exp(-r\rho) h(X^{\mathcal{C}}(\rho))) \leq h(x). \quad (2.14)$$

In summary, inequality (2.14) arises by applying the three dimensional Ito formula to g composed with the controlled process stopped inside $(0, 1)^3$ and then using continuity of h . But, from time ρ onwards, our controlled process runs on a face or an edge of the cube and Ito's formula in three dimensions does not apply. This is not a problem though – a similar argument with Ito's formula in one (or two) dimensions does the trick. That is, if ρ' denotes the first time that $X^{\mathcal{C}}$ hits an edge of \mathcal{S} (so $0 \leq \rho \leq \rho' \leq \tau^{\mathcal{C}}$), then both

$$\mathbb{E}_x (\exp(-r\rho') h(X^{\mathcal{C}}(\rho')) - \exp(-r\rho) h(X^{\mathcal{C}}(\rho))) \leq 0, \quad (2.15)$$

and

$$\mathbb{E}_x (\exp(-r\tau^{\mathcal{C}}) h(X^{\mathcal{C}}(\tau^{\mathcal{C}})) - \exp(-r\rho') h(X^{\mathcal{C}}(\rho'))) \leq 0. \quad (2.16)$$

Summing these differences and using the boundary condition $h(x) = 1$ for $x \in D$ yields

$$\mathbb{E}_x (\exp(-r\tau^{\mathcal{C}})) = \mathbb{E}_x (\exp(-r\tau^{\mathcal{C}}) h(X^{\mathcal{C}}(\tau^{\mathcal{C}}))) \leq h(x).$$

Thus h is an upper bound for the Laplace transform of the distribution of the decision time arising from any strategy. It remains to prove that h is equal to the Laplace transform \hat{v}_r .

Suppose that \mathcal{C} is the strategy \mathcal{C}^* from Lemma 1.4, then for almost every $s \geq 0$, $\dot{\mathcal{C}}_i(s)$ is positive only when $X_j^{\mathcal{C}}(s) \leq X_i^{\mathcal{C}}(s) \leq X_k^{\mathcal{C}}(s)$ under some labelling. So, $(\mathcal{G}^i - r)h(X^{\mathcal{C}}(s))\dot{\mathcal{C}}_i(s)$ vanishes for almost every $s \geq 0$ and (2.13) is an equality. Taking limits shows that (2.14) – (2.16) are also equalities. \square

So, \hat{v}_r is twice differentiable in each component and satisfies the heuristic equation

$$(\mathcal{G}^i - r) \hat{v}_r(x) = -r \hat{f}_r^i(x), \quad x \notin D, \quad x_i \in (0, 1). \quad (2.17)$$

In the next section we will show that $\mathbb{P}_x(\tau^{\mathcal{C}^*} > t)$ is the probabilistic solution to certain parabolic partial differential equations. To do this, we need to rewrite \hat{v}_r in a more convenient form.

It is convenient to introduce the notation $X^{(1)}(t) = (X_1(t), X_2(0), X_3(0))$, $X^{(2)}(t) = (X_1(0), X_2(t), X_3(0))$ and $X^{(3)}(t) = (X_1(0), X_2(0), X_3(t))$ for each $t \geq 0$. We define $\rho^{(i)}$ to be the absorption time of X_i , i.e.

$$\rho^{(i)} \stackrel{\text{def}}{=} \inf\{t \geq 0 : X_i(t) \notin (0, 1)\}.$$

Lemma 2.3. *For any $x \notin D$, \hat{v}_r can be written as*

$$\hat{v}_r(x) = \mathbb{E}_x \left(\exp(-r\rho^{(i)})\hat{v}_r(X^{(i)}(\rho^{(i)})) + r \int_0^{\rho^{(i)}} \hat{f}_r^i(X^{(i)}(s)) \exp(-rs) ds \right).$$

Proof. Fix $x \notin D$, then the function $x_i \mapsto \hat{v}_r(x)$ is C^2 on $(0, 1)$ and C^0 on $[0, 1]$. Introduce the a.s. finite \mathcal{F}_i stopping time $\rho_n^{(i)} \stackrel{\text{def}}{=} \inf\{t \geq 0 : X_i(t) \notin (n^{-1}, 1 - n^{-1})\}$, so Ito's formula (in one dimension) gives

$$\begin{aligned} \exp(-r\rho_n^{(i)})\hat{v}_r(X^{(i)}(\rho_n^{(i)})) - \hat{v}_r(X(0)) &= \int_0^{\rho_n^{(i)}} \exp(-rs) \frac{\partial}{\partial x_i} \hat{v}_r(X^{(i)}(s)) dX_i(s) \\ &\quad + \int_0^{\rho_n^{(i)}} \exp(-rs) (\mathcal{G}^i - r) \hat{v}_r(X^{(i)}(s)) ds. \end{aligned}$$

The function $\frac{\partial}{\partial x_i} \hat{v}_r$ is continuous on $(0, 1)$ and hence bounded on the compact subsets $[n^{-1}, 1 - n^{-1}]$. It follows that the expectation of the stochastic integral against dX_i vanishes. So, using equation (2.17),

$$\begin{aligned} \hat{v}_r(x) &= \mathbb{E}_x \left(\exp(-r\rho_n^{(i)})\hat{v}_r(X^{(i)}(\rho_n^{(i)})) \right) \\ &\quad + r \mathbb{E}_x \left(\int_0^{\rho_n^{(i)}} \exp(-rs) \hat{f}_r^i(X^{(i)}(s)) ds \right). \end{aligned}$$

The stopping times $\rho_n^{(i)}$ converge to $\rho^{(i)}$ as $n \rightarrow \infty$ and so by continuity of X_i , \hat{v}_r , the exponential function and the integral,

$$\begin{aligned} \exp(-r\rho_n^{(i)})\hat{v}_r(X^{(i)}(\rho_n^{(i)})) &\rightarrow \exp(-r\rho^{(i)})\hat{v}_r(X^{(i)}(\rho^{(i)})) \text{ and} \\ \int_0^{\rho_n^{(i)}} \exp(-rs) \hat{f}_r^i(X^{(i)}(s)) ds &\rightarrow \int_0^{\rho^{(i)}} \exp(-rs) \hat{f}_r^i(X^{(i)}(s)) ds. \end{aligned}$$

To finish the proof, use the Dominated Convergence Theorem to exchange the limit and expectation. \square

Remark 2.4. We can generalise our heuristic to value functions of the form

$$J(x, t) \stackrel{\text{def}}{=} \mathbb{E}_x(g(\tau^{C^*} + t)), \quad x \in \mathcal{S}, \quad t \geq 0,$$

for differentiable g . It reads

$$\left(\mathcal{G}^i + \frac{\partial}{\partial t} \right) J(x, t) = \mathbb{E}_x(g'(\tau^{C^*} + t); T_i = 0). \quad (2.18)$$

Equation (2.17) is the specialisation $g(t) = \exp(-rt)$. Such a choice of g is helpful because it effectively removes the time dependence in (2.18), making it easier to solve. The benefit is the same if g is linear and it is not difficult to construct and verify (compare Lemmas 2.1 and 2.2) an explicit expression for $J(x) \stackrel{\text{def}}{=} \mathbb{E}_x(\tau^{\mathcal{C}^*})$. In terms of the integrals

$$I_k(x_1) \stackrel{\text{def}}{=} \int_0^{x_3} \frac{G(u)}{(1-u)^k} du \text{ and } J_k(x_3) \stackrel{\text{def}}{=} \int_{x_3}^1 \frac{G(u)}{u^k} du, \quad k \in \mathbb{N},$$

the expression for J reads,

$$\begin{aligned} J(x) = & G(x_2) + (1-x_1)^{-2}G(x_1) \left((1-x_2)((1-x_1) - (1-x_3)) + (1-x_1)(1-x_3) \right) \\ & - 2I_3(x_1) \left((1-x_2)((1-x_1) + (1-x_3)) + (1-x_1)(1-x_3) \right) + \\ & 6I_4(x_1)(1-x_2)(1-x_1)(1-x_3) + x_3^{-2}G(x_3) (x_2(x_3-x_1) + x_1x_3) \\ & - 2J_3(x_3) (x_2(x_3+x_1) + x_1x_3) + 6J_4(x_3)x_1x_2x_3. \end{aligned}$$

3 A representation for $\mathbb{P}_x(\tau^{\mathcal{C}^*} > T)$

The aim of this section is to connect the tail probability $v : \mathcal{S} \times [0, \infty) \rightarrow [0, 1]$ defined by

$$v(x, t) \stackrel{\text{def}}{=} \mathbb{P}_x(\tau^{\mathcal{C}^*} > t),$$

to the formula for \hat{v}_r from Lemma 2.3. Before continuing, let us explain the key idea. Just for a moment, suppose that v is smooth and consider the Laplace transform of $(\mathcal{G}^i - \frac{\partial}{\partial t})v(x, \cdot)$. It is straightforward to show that the Laplace transform of v satisfies (see (3.22)),

$$\int_0^\infty v(t, x) \exp(-rt) dt = r^{-1} (1 - \hat{v}_r(x)).$$

Bringing \mathcal{G}^i through the integral and integrating by parts in t ,

$$\int_0^\infty \exp(-rt) \left(\mathcal{G}^i - \frac{\partial}{\partial t} \right) v(x, t) dt = -r^{-1} (\mathcal{G}^i - r) \hat{v}_r(x).$$

Combining this with the heuristic equation (2.17) gives

$$\int_0^\infty \exp(-rt) \left(\mathcal{G}^i - \frac{\partial}{\partial t} \right) v(x, t) dt = \hat{f}_r^i(x). \quad (3.19)$$

This shows that $(\mathcal{G}^i - \frac{\partial}{\partial t})v$ is non-negative, (i.e. v satisfies the associated Hamilton-Jacobi-Bellman equation). From here, one could use Ito's formula (c.f. the proof of Lemma 2.2) to see that $(v(X^{\mathcal{C}}(t), T-t), 0 \leq t \leq T)$ is a sub-martingale for any strategy \mathcal{C} . In particular,

$$\mathbb{P}_x(\tau^{\mathcal{C}} > T) = \mathbb{E}_x(v(X^{\mathcal{C}}(T), 0)) \geq v(x, T).$$

So, ideally, to prove Theorem 1.1, we would establish that v is smooth enough to apply Ito's formula. We are given some hope, by noticing that if we can show that \hat{f}_r^i is the Laplace transform of a function f_i say, then (3.19) implies v solves

$$\int_0^\infty \exp(-rt) \left(\mathcal{G}^i - \frac{\partial}{\partial t} \right) v = f_i. \quad (3.20)$$

We can show such a density f_i exists (Lemma 3.1 below) but (surprisingly) not that it is Hölder continuous. Unfortunately, without the latter we cannot show that (3.20) has a classical solution. Nevertheless, we can deduce the sub-martingale inequality by showing merely that v solves (3.20) in a weaker sense (Lemma 3.2).

To commence, let us first verify the claim that \hat{f}_r^i is the Laplace transform of a function.

Lemma 3.1. *For each $x \notin D$ and $i \in V$, the Borel measure $B \mapsto \mathbb{P}_x(\tau^{C^*} \in B, T_i = 0)$ has a (defective) density $f_i : \mathcal{S} \times [0, \infty) \rightarrow [0, \infty)$, i.e.*

$$\mathbb{P}_x(\tau^{C^*} \in dt, T_i = 0) = f_i(x, t)dt, \quad t \geq 0.$$

Proof. This is essentially a corollary of the decomposition of τ^{C^*} on $\{T_i = 0\}$ that was discussed in the proof of Lemma 2.1.

Recall that if $\mathbf{m}_a^{(i)}$ is the first hitting time of level a by X_i (defined in (2.8)), then for $x_1 \leq x_2 \leq x_3$,

$$\mathbb{P}_x(\tau^{C^*} \in B, T_1 = 0) = \mathbb{P}_x(\mathbf{m}_1^{(2)} + \mathbf{m}_1^{(3)} \in B, \mathbf{m}_1^{(2)} < \mathbf{m}_{x_1}^{(2)}, \mathbf{m}_1^{(3)} < \mathbf{m}_{x_1}^{(3)}).$$

This is the convolution of the sub-probability measures

$$\mathbb{P}_x(\mathbf{m}_1^{(i)} \in \cdot, \mathbf{m}_1^{(i)} < \mathbf{m}_{x_1}^{(i)}), \quad i = 1, 2.$$

If $x_1 = x_2$, then $T_1 > 0$ almost surely under \mathbb{P}_x , and when $x \notin D$, $x_2 < 1$. So, we may assume that x_2 is in the interval $(x_1, 1)$. In this case, $\{\mathbf{m}_1^{(2)} < \mathbf{m}_{x_1}^{(2)}\}$ is not null and X_2 can be conditioned, via a Doob h -transform, to exit $(x_1, 1)$ at the upper boundary. The conditioned process is again a diffusion and so the arguments of §4.11 of [10] show that $\mathbb{P}_x(\mathbf{m}_1^{(2)} \in \cdot, \mathbf{m}_1^{(2)} < \mathbf{m}_{x_1}^{(2)})$ is absolutely continuous. Hence, $\mathbb{P}_x(\tau^{C^*} \in \cdot, T_1 = 0)$ is the convolution of two measures, at least one of which has a density.

The other cases are treated with essentially identical arguments. \square

The next step is to show that v solves (3.20) in a probabilistic sense.

Lemma 3.2. *Fix $i \in V$ and define the function $u : \mathcal{S} \times [0, \infty) \rightarrow \mathbb{R}$ by*

$$u(x, t) \stackrel{\text{def}}{=} \mathbb{E}_x \left(v(X^{(i)}(t \wedge \rho^{(i)}), (t - \rho^{(i)})^+) - \int_0^{t \wedge \rho^{(i)}} f_i(X^{(i)}(s), t - s) ds \right), \quad (3.21)$$

where $\rho^{(i)} = \inf\{t \geq 0 : X_i(t) \notin (0, 1)\}$ and f_i is the density from Lemma 3.1. Then,

- (a) for each $x \notin D$, $u(x, \cdot)$ has the same Laplace transform as $v(x, \cdot)$,
- (b) both $u(x, \cdot)$ and $v(x, \cdot)$ are right continuous, and as a result
- (c) the tail probability v is equal to u and so has the representation given in (3.21).

Proof. (a) The Laplace transform of the tail probability is, for $x \notin D$,

$$\begin{aligned} \int_0^\infty v(t, x) \exp(-rt) dt &= \mathbb{E}_x \int_0^\infty \mathbb{1}_{[\tau^{c^*} > t]} \exp(-rt) dt \\ &= \mathbb{E}_x \int_0^{\tau^{c^*}} \exp(-rt) dt \\ &= r^{-1} (1 - \hat{v}_r(x)), \end{aligned}$$

by Fubini's Theorem since the integrand is non-negative. Furthermore, for $x \in D$, both $v(t, x)$ and $1 - \hat{v}_r(x)$ vanish and so in fact, for *any* $x \in \mathcal{S}$ we have

$$\int_0^\infty v(t, x) \exp(-rt) dt = r^{-1} (1 - \hat{v}_r(x)). \quad (3.22)$$

Now, we consider the Laplace transform of u . By linearity of the expectation operator,

$$u(x, t) = \mathbb{E}_x (v(X^{(i)}(t \wedge \rho^{(i)}), (t - \rho^{(i)})^+)) - \mathbb{E}_x \left(\int_0^{t \wedge \rho^{(i)}} f_i(X^{(i)}(s), t - s) ds \right).$$

First consider the Laplace transform of the first member of the right hand side,

$$\int_0^\infty \mathbb{E}_x (v(X^{(i)}(t \wedge \rho^{(i)}), (t - \rho^{(i)})^+)) \exp(-rt) dt.$$

Applying Fubini's theorem, the preceding expression becomes

$$\mathbb{E}_x \left(\int_0^\infty v(X^{(i)}(t \wedge \rho^{(i)}), (t - \rho^{(i)})^+) \exp(-rt) dt \right),$$

which can be decomposed into the sum

$$\mathbb{E}_x \left(\int_0^{\rho^{(i)}} v(X^{(i)}(t), 0) \exp(-rt) dt \right) + \mathbb{E}_x \left(\int_{\rho^{(i)}}^\infty v(X^{(i)}(\rho^{(i)}), t - \rho^{(i)}) \exp(-rt) dt \right).$$

The first term is

$$\mathbb{E}_x \int_0^{\rho^{(i)}} v(X^{(i)}(t), 0) \exp(-rt) dt = r^{-1} \mathbb{E}_x (1 - \exp(-r\rho^{(i)})), \quad (3.23)$$

because when $x \notin D$, \mathbb{P}_x -almost surely we have $X^{(i)}(t) \notin D$ for $t < \rho^{(i)}$. The second term,

$$\mathbb{E}_x \int_{\rho^{(i)}}^\infty v(X^{(i)}(\rho^{(i)}), t - \rho^{(i)}) \exp(-rt) dt.$$

If we shift the variable of integration to $u = t - \rho^{(i)}$ and then use (3.22), the last expression becomes

$$r^{-1} \mathbb{E}_x (\exp(-r\rho^{(i)})(1 - \hat{v}_r(X^{(i)}(\rho^{(i)}))). \quad (3.24)$$

The treatment of

$$\int_0^\infty \mathbb{E}_x \left(\int_0^{t \wedge \rho^{(i)}} f_i(X^{(i)}(s), t-s) ds \right) \exp(-rt) dt \quad (3.25)$$

proceeds in a similar fashion – exchange the expectation and outer integral and then decompose the integrals into $t < \rho^{(i)}$ and $t \geq \rho^{(i)}$. The integral over $t < \rho^{(i)}$ is

$$\mathbb{E}_x \int_0^{\rho^{(i)}} \int_0^t f_i(X^{(i)}(s), t-s) ds \exp(-rt) dt.$$

Exchanging the integrals in t and s gives

$$\mathbb{E}_x \int_0^{\rho^{(i)}} \int_s^{\rho^{(i)}} f_i(X^{(i)}(s), t-s) \exp(-rt) dt ds.$$

For the integral over $t \geq \rho^{(i)}$, we again exchange the integrals in t and s to give

$$\mathbb{E}_x \int_0^{\rho^{(i)}} \int_{\rho^{(i)}}^\infty f_i(X^{(i)}(s), t-s) \exp(-rt) dt ds.$$

Summing these final two expressions and substituting $u = t - s$ shows that (3.25) is equal to

$$\mathbb{E}_x \int_0^{\rho^{(i)}} \int_0^\infty f_i(X^{(i)}(s), u) \exp(-rt) du \exp(-rs) ds.$$

The Laplace transform is a linear operator, and so we may sum (3.23), (3.24) and (3.25) to show that the Laplace transform of u is equal to

$$r^{-1} \mathbb{E}_x (1 - \exp(-r\rho^{(i)}) \hat{v}_r(X^{(i)}(\rho^{(i)}))) + \mathbb{E}_x \left(\int_0^{\rho^{(i)}} \hat{f}_r^i(X^{(i)}(s)) \exp(-rs) ds \right), \quad (3.26)$$

where we have used

$$\int_0^\infty f_i(x, u) \exp(-rt) du = \hat{f}_r^i(x)$$

for $x \notin D$.

But, (3.26) is exactly what we get by substituting the representation for \hat{v}_r from Lemma (2.3) into (3.22), and so we're done.

(b) Right-continuity of v in t follows from the Monotone Convergence Theorem. A little more work is required to see that u is right-continuous. We begin by observing that if $\rho^{(i)} > t$ then X_i has not been absorbed by time t and so, if $x \notin D$, there is a \mathbb{P}_x -negligible set outside of which $X^{(i)}(t) \notin D$.

It follows that $\{X^{(i)}(t) \in D, \rho^{(i)} > t\} = \{\rho^{(i)} > t\}$ almost surely. Combining this with the fact that $v(\cdot, 0) = \mathbb{1}_{[\cdot \notin D]}$ shows

$$\mathbb{E}_x (v(X^{(i)}(t \wedge \rho^{(i)}), (t - \rho^{(i)})^+); \rho^{(i)} > t) = \mathbb{P}_x (\rho^{(i)} > t) \text{ for } x \notin D.$$

The latter is right-continuous in t by the Monotone Convergence Theorem. The complementary expectation

$$\mathbb{E}_x (v(X^{(i)}(t \wedge \rho^{(i)}), (t - \rho^{(i)})^+); \rho^{(i)} \leq t)$$

is equal to

$$\mathbb{E}_x (v(X^{(i)}(\rho^{(i)}), t - \rho^{(i)}); \rho^{(i)} \leq t),$$

the right continuity of which follows from that of v and the indicator $\mathbb{1}_{[\rho^{(i)} \leq t]}$, together with the Dominated Convergence Theorem.

We now consider the expectation of the integral,

$$\mathbb{E}_x \left(\int_0^{t \wedge \rho^{(i)}} f_i(X^{(i)}(s), t - s) ds \right).$$

Using Fubini's theorem we may exchange the integral and expectation to get

$$\int_0^t \mathbb{E}_x (f_i(X^{(i)}(s), t - s); \rho^{(i)} > s) ds. \quad (3.27)$$

This suggests the introduction of $(p_s^\dagger; s \geq 0)$, the transition kernel of X_i killed (and sent to a cemetery state) on leaving $(0, 1)$. Such a density exists by the arguments of §4.11 of [10].

For notational ease, let us assume $i = 1$, then (3.27) can be written

$$\int_0^t \int_0^1 p_s^\dagger(x_1, y) f_1((y, x_2, x_3), t - s) dy ds.$$

Finally, changing the variable of integration from s to $s' = t - s$ gives

$$\int_0^t \int_0^1 p_{t-s'}^\dagger(x_1, y) f_1((y, x_2, x_3), s') dy ds',$$

and so regularity of (3.27) in t is inherited from p^\dagger . This is sufficient because p_t^\dagger is continuous in $t > 0$ (again see [10]).

(c) It follows from (a) that for each $x \notin D$, $u(x, t)$ and $v(x, t)$ are equal for almost every $t \geq 0$. Hence, right continuity is enough to show $v(x, t) = u(x, t)$ for every $t \geq 0$. \square

From the probabilistic representation for v , we need to deduce some sub-martingale type inequalities for $v(X^C(t), T - t)$, $0 \leq t \leq T$. As we will see later, it is enough to consider strategies that, for some $\epsilon > 0$, run only process during the interval $(k\epsilon, (k+1)\epsilon)$, for integers $k \geq 0$. In other words, the rates for each process are either zero or one and are constant over $(k\epsilon, (k+1)\epsilon)$. More specifically,

Definition 3.3 (ϵ -strategy). For $\epsilon > 0$ we let Π_ϵ denote the set of strategies \mathcal{C}^ϵ such that for any integer $k \geq 0$,

$$\mathcal{C}^\epsilon(t) = \mathcal{C}^\epsilon(k\epsilon) + (t - k\epsilon)\xi_k, \quad k\epsilon \leq t \leq (k+1)\epsilon,$$

where ξ_k takes values in the set of standard basis elements $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$.

Lemma 3.4. *Suppose $x \in \mathcal{S}$ and $0 \leq t \leq T$, then the following sub-martingale inequalities hold.*

(a) For $i \in V$,

$$\mathbb{E}_x (v(X^{(i)}(t), T - t)) \geq v(x, T).$$

(b) If $\mathcal{C}^\epsilon \in \Pi_\epsilon$ then

$$\mathbb{E}_x (v(X^{\mathcal{C}^\epsilon}(t), T - t)) \geq v(x, T).$$

Proof. Consider first the quantity

$$\mathbb{E}_x \left(\mathbb{E}_{X^{(i)}(t)} (v(X^{(i)}((T - t) \wedge \rho^{(i)}), (T - t - \rho^{(i)})^+)) \right). \quad (3.28)$$

Our Markovian setup comes with a shift operator $\theta = \theta^{(i)}$ for $X^{(i)}$ defined by $X^{(i)} \circ \theta_s(\omega, t) = X^{(i)}(\theta_s \omega, t) = X^{(i)}(\omega, s + t)$ for each $\omega \in \Omega$. In terms of this operator, (3.28) becomes

$$\mathbb{E}_x \left(\mathbb{E}_x (v(X^{(i)}((T - t) \wedge \rho^{(i)}), (T - t - \rho^{(i)})^+) \circ \theta_t | \mathcal{F}_i(t)) \right).$$

From here, use the Tower Property and the fact that $\rho^{(i)} \circ \theta_t = (\rho^{(i)} - t) \vee 0$ to find that (3.28) equals

$$\mathbb{E}_x (v(X^{(i)}(T \wedge \rho^{(i)}), (T - \rho^{(i)})^+)). \quad (3.29)$$

We can give a similar treatment for

$$\mathbb{E}_x \left(\mathbb{E}_{X^{(i)}(t)} \left(\int_0^{(T-t) \wedge \rho^{(i)}} f_i(X^{(i)}(s), T - t - s) ds \right) \right). \quad (3.30)$$

Using the Markov property of $X^{(i)}$, (3.30) becomes

$$\mathbb{E}_x \left(\mathbb{E}_x \left(\int_0^{(T-t) \wedge \rho^{(i)}} f_i(X^{(i)}(s), T - t - s) ds \circ \theta_t | \mathcal{F}_i(t) \right) \right).$$

Substituting in for $X^{(i)} \circ \theta_t$ and $\rho^{(i)} \circ \theta_t$ and using the Tower Property, the latter expectation is seen to be

$$\mathbb{E}_x \left(\int_0^{(T-t) \wedge (\rho^{(i)} - t) \vee 0} f_i(X^{(i)}(s + t), T - t - s) ds \right).$$

Now make the substitution $u = s + t$ in the integral and use the fact that f_i is non-negative to show that (3.30) is less than or equal to

$$\mathbb{E}_x \left(\int_0^{T \wedge \rho^{(i)}} f_i(X^{(i)}(u), T - u) du \right). \quad (3.31)$$

The final step is to note that, by Lemma 3.2,

$$\begin{aligned} v(x, T - t) &= \mathbb{E}_x (v(X^{(i)}(T - t \wedge \rho^{(i)}), (T - t - \rho^{(i)})^+)) \\ &\quad - \mathbb{E}_x \left(\int_0^{(T-t) \wedge \rho^{(i)}} f(X^{(i)}(s), T - t - s) ds \right), \end{aligned}$$

and so $\mathbb{E}_x(v(X^{(i)}(t), T - t))$ is equal to (3.28) minus (3.30), which by the argument above is greater than or equal to

$$\mathbb{E}_x(v(X^{(i)}(T \wedge \rho^{(i)}), (T - \rho^{(i)})^+)) - \mathbb{E}_x\left(\int_0^{T \wedge \rho^{(i)}} f_i(X^{(i)}(u), T - u) du\right).$$

Again appealing to Lemma 3.2 shows that the latter is exactly $v(x, T)$.

(b) It is sufficient to prove that for $k\epsilon \leq t \leq (k+1)\epsilon$ we have

$$\mathbb{E}_x(v(X^{C^\epsilon}(t), T - t) | \mathcal{F}^{C^\epsilon}(k\epsilon)) \geq v(X^{C^\epsilon}(k\epsilon), T - k\epsilon). \quad (3.32)$$

The desired result then follows by applying the Tower Property of conditional expectation and iterating this inequality.

Let us take $\nu \stackrel{\text{def}}{=} C^\epsilon(k\epsilon)$ and $\mathcal{H} \stackrel{\text{def}}{=} \mathcal{F}^{C^\epsilon}(k\epsilon)$. Then ν takes values in the grid $\mathcal{Z} \stackrel{\text{def}}{=} \{0, \epsilon, 2\epsilon, \dots\}^3$ and $\Lambda \in \mathcal{H}$ implies that $\Lambda \cap \{\nu = z\}$ is an element of the σ -field $\mathcal{F}(z) = \sigma(\mathcal{F}_1(z_1), \dots, \mathcal{F}_3(z_3))$ for $z \in \mathcal{Z}$. It follows from the definition of conditional expectation that \mathbb{P}_x almost surely we have

$$\mathbb{E}_x(\cdot | \mathcal{H}) = \mathbb{E}_x(\cdot | \mathcal{F}(z)) \text{ on } \{\nu = z\}. \quad (3.33)$$

Now, by continuity of C_i^ϵ and right-continuity of \mathcal{F}^{C^ϵ} , ξ_k must be \mathcal{H} -measurable. So, if $A \stackrel{\text{def}}{=} A_1 \times A_2 \times A_3$ with A_i Borel measurable for each $i \in V$, (3.33) gives the equality

$$\mathbb{E}_x(\mathbb{1}_{[\nu=z, X^{C^\epsilon}(t) \in A, \xi_k=e_i]} | \mathcal{H}) = \mathbb{1}_{[\nu=z, \xi_k=e_i]} \mathbb{E}_x(\mathbb{1}_{[X(z+(t-k\epsilon)e_i) \in A]} | \mathcal{F}(z)),$$

where, as before, $X(z) = (X_1(z_1), X_2(z_2), X_3(z_3))$.

Next we use the facts that $\mathbb{1}_{[X_j(z_j) \in A_j]}$ is $\mathcal{F}(z)$ measurable for each j and that the filtration \mathcal{F}_i of X_i is independent of \mathcal{F}_j for $j \neq i$, to show that the preceding expression is equal to

$$\mathbb{1}_{[\nu=z, \xi_k=e_i, X_j(z_j) \in A_j, j \neq i]} \mathbb{E}_x(\mathbb{1}_{[X_i(z_i+(t-k\epsilon)) \in A_i]} | \mathcal{F}_i(z_i)).$$

Finally, the Markov property of X_i allows us to write this as

$$\mathbb{1}_{[\nu=z, \xi_k=e_i]} \mathbb{E}_{X(z)}(\mathbb{1}_{[X^{(i)}(t-k\epsilon) \in A]}).$$

As $\mathbb{E}_x(v(X^{(i)}(t), s))$ is Borel measurable for any $s, t \geq 0$, this is enough to conclude that in our original notation, on $\{\xi_k = e_i\}$,

$$\mathbb{E}_x(v(X^{C^\epsilon}(t), T - t) | \mathcal{F}^{C^\epsilon}(k\epsilon)) = \mathbb{E}_{X^{C^\epsilon}(k\epsilon)}(v(X^{(i)}(t - k\epsilon), T - t)). \quad (3.34)$$

But part (a) shows that

$$\mathbb{E}_x(v(X^{(i)}(t - k\epsilon), (T - k\epsilon) - (t - k\epsilon))) \geq v(x, T - k\epsilon),$$

and so the right hand side of (3.34) is greater than or equal to $v(X^{C^\epsilon}(k\epsilon), T - k\epsilon)$. \square

It is now relatively painless to combine the ingredients above. We take an arbitrary strategy \mathcal{C} , use Lemma A.2 to approximate it by the family \mathcal{C}^ϵ , $\epsilon > 0$ and then use Lemma 3.4 part (b) with $t = T \geq 0$ to show that

$$\mathbb{P}_x(\tau^{\mathcal{C}^\epsilon} > T) = \mathbb{E}_x v(X^{\mathcal{C}^\epsilon}(T), 0) \geq v(x, T)$$

for any $x \notin D$ (equality holds trivially for $x \in D$).

The approximations are such that $\mathcal{C}(t) \preceq \mathcal{C}^\epsilon(t + M\epsilon)$ for some constant $M > 0$. Thus, $\tau^{\mathcal{C}} \leq t$ implies that $\tau^{\mathcal{C}^\epsilon} \leq t + M\epsilon$. More usefully, the contrapositive is that $\tau^{\mathcal{C}^\epsilon} > t + M\epsilon$ implies $\tau^{\mathcal{C}} > t$ and so monotonicity of the probability measure \mathbb{P}_x then ensures

$$\mathbb{P}_x(\tau^{\mathcal{C}} > t) \geq \mathbb{P}_x(\tau^{\mathcal{C}^\epsilon} > t + M\epsilon) \geq v(x, t + M\epsilon).$$

Taking the limit $\epsilon \rightarrow 0$ and using right continuity of $v(x, t)$ in t completes the proof.

4 Existence and almost sure uniqueness of \mathcal{C}^\star

In this section we give a proof for Lemma 1.4. Recall that we wish to study strategies \mathcal{C} that satisfy the property

(RTM) \mathcal{C}_i increases at time $t \geq 0$ (i.e. for every $s > t$, $\mathcal{C}_i(s) > \mathcal{C}_i(t)$) only if, under some labelling of the processes,

$$X_j^{\mathcal{C}}(t) \leq X_i^{\mathcal{C}}(t) \leq X_k^{\mathcal{C}}(t).$$

Our idea is to reduce the existence and uniqueness of our strategy to a one-sided problem. Then, we can use the following result, taken from Proposition 5 and Corollary 13 in [15] (alternatively §5.1 of [11]).

Lemma 4.1. *Suppose that $(Y_i(t); t \geq 0)$, $i = 1, 2$ are independent and identically distributed regular Ito diffusions on \mathbb{R} , beginning at the origin and with complete, right continuous filtrations $(\mathcal{H}_i(t); t \geq 0)$. Then*

- (a) *there exists a strategy $\gamma = (\gamma_1(t), \gamma_2(t); t \geq 0)$ (with respect to the multiparameter filtration $\mathcal{H} = (\sigma(\mathcal{H}_1(z_1), \mathcal{H}_2(z_2)); z \in \mathbb{R}_+^2)$) such that γ_i increases only at times $t \geq 0$ with*

$$Y_i^\gamma(t) = Y_1^\gamma(t) \wedge Y_2^\gamma(t),$$

i.e. “ γ follows the minimum of Y_1 and Y_2 ”.

- (b) *If γ' is another strategy with this property, then, almost surely, $\gamma'(t) = \gamma(t)$ for every $t \geq 0$. That is, γ is a.s. unique.*

- (c) *the maximum $Y_1^\gamma(t) \vee Y_2^\gamma(t)$ increases with t .*

We first consider the question of uniqueness, it will then be obvious how \mathcal{C}^\star must be defined. Suppose that \mathcal{C} is a strategy satisfying (RTM).

If $X_1(0) < X_2(0) = X_3(0)$, then \mathcal{C} cannot run X_1 (i.e. \mathcal{C}_1 does not increase) before the first time ν that either $X_2^{\mathcal{C}}$ or $X_3^{\mathcal{C}}$ hit $X_1(0)$. Until then (or until a decision is made, whichever comes first), \mathcal{C}_2 may increase only at times $t \geq 0$ when $X_2^{\mathcal{C}}(t) \leq X_3^{\mathcal{C}}(t)$ and \mathcal{C}_3

only when $X_3^{\mathcal{C}}(t) \leq X_2^{\mathcal{C}}(t)$. Hence, on $\tau^{\mathcal{C}} \wedge \nu \geq t$, the value of $\mathcal{C}(t)$ is determined by the strategy in lemma 4.1. Now, $X_2^{\mathcal{C}} \vee X_3^{\mathcal{C}}$ increases during this time, and so if $\nu < \tau^{\mathcal{C}}$, we have

$$X_1(0) = X_1^{\mathcal{C}}(\nu) = X_2^{\mathcal{C}}(\nu) \wedge X_3^{\mathcal{C}}(\nu) < X_2^{\mathcal{C}}(\nu) \vee X_3^{\mathcal{C}}(\nu).$$

So again, we are in a position to apply the argument above, and can do so repeatedly until a decision is made. In fact, it takes only a finite number of iterations of the argument to determine $\mathcal{C}(t)$ for each $t \geq 0$ (on $\tau^{\mathcal{C}} \geq t$) because each diffusion X_i is continuous, the minimum $X_1^{\mathcal{C}} \wedge X_2^{\mathcal{C}} \wedge X_3^{\mathcal{C}}$ is decreasing and the maximum $X_1^{\mathcal{C}} \vee X_2^{\mathcal{C}} \vee X_3^{\mathcal{C}}$ increasing. If $X_1(0) < X_2(0) < X_3(0)$ then \mathcal{C} must run X_2 exclusively until it hits either $X_1(0)$ or $X_3(0)$. From then on, the arguments of the previous case apply.

The remaining possibility is that $X_1(0) = X_2(0) = X_3(0) = a \in (0, 1)$. We shall define random times ν_ϵ , $0 < \epsilon < (1 - a) \wedge a$ such that

- $\mathcal{C}(\nu_\epsilon)$ is determined by the property (RTM),
- under some labelling,

$$a - \epsilon < X_1^{\mathcal{C}}(\nu_\epsilon) < a < X_2^{\mathcal{C}}(\nu_\epsilon) = X_3^{\mathcal{C}}(\nu_\epsilon) = a + \epsilon,$$

and

- $\nu_\epsilon \rightarrow 0$ as $\epsilon \rightarrow 0$.

Again, we may then use the one-sided argument to see that, almost surely, on $\nu_\epsilon \leq t \leq \tau^{\mathcal{C}}$, $\mathcal{C}(t)$ is determined by (RTM). This is sufficient because $\nu_\epsilon \rightarrow 0$ as $\epsilon \rightarrow 0$.

To construct ν_ϵ , suppose, without loss of generality, that X_1 and X_2 both exit $(a - \epsilon, a + \epsilon)$ at the upper boundary. We denote by α_i the finite time taken for this to happen, i.e.

$$\alpha_i \stackrel{\text{def}}{=} \inf\{t > 0 : X_i(t) \notin (a - \epsilon, a + \epsilon)\}.$$

Define

$$l_i \stackrel{\text{def}}{=} \inf_{0 \leq s \leq \alpha_i} X_i(s)$$

to be the lowest value attained by X_i before it exits $(a - \epsilon, a + \epsilon)$. By Proposition 5 of [15], it is almost sure that the l_i are not equal and so, we may assume that $l_3 < l_2 < l_1$ (by relabelling if necessary).

Intuitively, (RTM) means that $X_1^{\mathcal{C}}$ and $X_2^{\mathcal{C}}$ should hit $a + \epsilon$ together while $X_3^{\mathcal{C}}$ gets left down at l_2 . We already know it takes time α_i for X_i to hit $a + \epsilon$ ($i = 1, 2$) and X_3 takes time

$$\beta_3 \stackrel{\text{def}}{=} \inf\{t > 0 : X_3(t) = l_2\}.$$

to reach l_2 . So, we set $\nu_\epsilon = \alpha_1 + \alpha_2 + \beta_3$, and claim that

$$\mathcal{C}(\nu_\epsilon) = (\alpha_1, \alpha_2, \beta_3).$$

The proof proceeds by examining the various cases. Firstly, if $\mathcal{C}_1(\nu_\epsilon) > \alpha_1$ and $\mathcal{C}_1(\nu_\epsilon) \geq \alpha_1$, then necessarily $\mathcal{C}_3(\nu_\epsilon) < \beta_3$ and $X_3(z_3) > l_2$ for any $z_3 \leq \mathcal{C}_3(\nu_\epsilon)$. But, then there exist times $\alpha'_i < \mathcal{C}_i(\nu_\epsilon)$ ($i = 1, 2$) with

$$l_2 = X_2(\alpha'_2) < X_3(z_3) < X_1(\alpha'_1) = a + \epsilon$$

for any $z_3 \leq \mathcal{C}_3(\nu_\epsilon)$, contradicting (RTM).

The second case is that $\mathcal{C}_1(\nu_\epsilon) < \alpha_1$ and $\mathcal{C}_2(\nu_\epsilon) \leq \alpha_2$. Necessarily we then have $\mathcal{C}_3(\nu_\epsilon) > \beta_3$. Now, $X_i(z_i) \geq l_2$ for $z_i \leq \alpha_i$, $i = 1, 2$ and so (RTM) implies that $X_3(z_3) \geq l_2$ as well for $z_3 \leq \mathcal{C}_3(\nu_\epsilon)$. In addition, (RTM) and $\mathcal{C}_3(\nu_\epsilon) > \beta_3$ imply that

$$\mathcal{C}_2(\nu_\epsilon) \geq \inf\{t > 0 : X_2(t) = l_2\}$$

(otherwise $X_3(\beta_3) < X_i(z_i)$ for $z_i \leq \mathcal{C}_i(\nu_\epsilon)$, $i = 1, 2$). So, both X_2 and X_3 have attained l_2 and then stayed above it for a positive amount of time. But, by Proposition 5 in [15], this event that “the lower envelopes of X_2 and X_3 are simultaneously flat” has probability zero.

The final case $\mathcal{C}_1(\nu_\epsilon) > \alpha_1$ and $\mathcal{C}_2(\nu_\epsilon) \geq \alpha_2$ has two subcases, $\mathcal{C}_3(\nu_\epsilon) \leq \beta_3$ and $\mathcal{C}_3(\nu_\epsilon) > \beta_3$ – both can be eliminated by the methods above. The only remaining possibility is that $\mathcal{C}_i(\nu_\epsilon) = \alpha_i$ for $i = 1, 2$ and $\mathcal{C}_3(\nu_\epsilon) = \beta_3$.

The discussion above tells us how to define \mathcal{C}^* – if $X_1(0) < X_2(0) \leq X_3(0)$ under some labelling, then we just alternate the one-sided construction from lemma 4.1 repeatedly to give a strategy satisfying (C1) – (C3). If $X_1(0) = X_2(0) = X_3(0) = a \in (0, 1)$, take $0 < \epsilon < a \wedge (1 - a)$ and define $\mathcal{C}^*(\nu_u)$, $0 < u \leq \epsilon$ via the construction above. Now, ν_u is only left continuous, so we have yet to define \mathcal{C}^* on the stochastic intervals $(\nu_u, \nu_{u+}]$, $u \leq \epsilon$. But, this is easily done because $X^{\mathcal{C}^*}(\nu_u)$ has exactly two components equal and so we can again use the one-sided construction. We define \mathcal{C}^* on $(\nu_\epsilon, \tau^{\mathcal{C}^*}]$ similarly. The properties (C1) and (C2) are readily verified. To confirm (C3), we first observe that \mathcal{C}^* satisfies (RTM). But (RTM) gives us almost sure uniqueness of the paths of \mathcal{C}^* . It follows that our definition of \mathcal{C}^* does not depend on ϵ . The second observation is that $\nu_u \rightarrow 0$ as $u \rightarrow 0$. As a consequence, for $\eta \in \mathbb{R}_+^3$ and $\delta > 0$,

$$\begin{aligned} \{\mathcal{C}^*(t) \preceq \eta\} &= \{\mathcal{C}^*(t) \preceq \eta, \nu_u < \delta \text{ some } u < \epsilon\} \\ &= \bigcup_q \{\mathcal{C}^*(t) \preceq \eta, \nu_q < \delta\}, \end{aligned}$$

where the union is over rational numbers $0 < q < \epsilon$. Using the fact that \mathcal{F} is complete,

$$\{\mathcal{C}^*(t) \preceq \eta, \nu_q < \delta\} \in \mathcal{F}(\eta_1 + \delta, \eta_2 + \delta, \eta_3 + \delta).$$

From this we conclude that $\{\mathcal{C}^*(t) \preceq \eta\} \in \mathcal{F}(\eta)$ because \mathcal{F} is right continuous. This confirms (C3).

5 $X^{\mathcal{C}^*}$ as a doubly perturbed diffusion

We now turn our attention to the optimally controlled process $X^{\mathcal{C}^*}$. For convenience, we will work with the minimum

$$I_t \stackrel{\text{def}}{=} X_1^{\mathcal{C}^*}(t) \wedge X_2^{\mathcal{C}^*}(t) \wedge X_3^{\mathcal{C}^*}(t),$$

maximum

$$S_t \stackrel{\text{def}}{=} X_1^{\mathcal{C}^*}(t) \vee X_2^{\mathcal{C}^*}(t) \vee X_3^{\mathcal{C}^*}(t),$$

and middle value

$$M_t \stackrel{\text{def}}{=} (X_1^{C^*}(t) \vee X_2^{C^*}(t)) \wedge (X_1^{C^*}(t) \vee X_3^{C^*}(t)) \wedge (X_2^{C^*}(t) \vee X_3^{C^*}(t)), t \geq 0$$

of the components of X^{C^*} (so, if $X_1^{C^*}(t) \leq X_2^{C^*}(t) \leq X_3^{C^*}(t)$, then $I_t = X_1^{C^*}(t)$, $M_t = X_2^{C^*}(t)$, $S_t = X_3^{C^*}(t)$). There is no ambiguity when the values of the components are equal since we are not formally *identifying* I_t , M_t and S_t with a particular component of X^{C^*}

Clearly, M behaves as an Ito diffusion solving (1.1) away from the extrema I and S , while at the extrema it experiences a perturbation. This behaviour is reminiscent of *doubly perturbed Brownian motion*, which is defined as the (pathwise unique) solution $(X'_t; t \geq 0)$ of the equation

$$X'_t = B'_t + \alpha \sup_{s \leq t} X'_s + \beta \inf_{s \leq t} X'_s,$$

where $\alpha, \beta < 1$ and $(B'_t; t \geq 0)$ is a Brownian motion starting from the origin. This process was introduced by Le Gall and Yor in [13]; the reader may consult the survey [18] and introduction of [4] for further details. In §2 of [4], this definition is generalised to accommodate non-zero initial values for the maximum and minimum processes in the obvious way – if $i_0, s_0 \geq 0$, we take

$$X'_t = B'_t + \alpha \left(\sup_{s \leq t} X'_s - s_0 \right)^+ - \beta \left(\inf_{s \leq t} X'_s + i_0 \right)^-,$$

i.e. X' hits $-i_0$ or s_0 before the perturbations begin. As usual $a^+ = \max(a, 0)$ and $a^- = \max(-a, 0)$.

Our suspicion that M should solve this equation if the underlying processes are Brownian motions is confirmed in the following

Lemma 5.1. *Suppose that $0 \leq i_0 \leq m_0 \leq s_0 \leq 1$ and $\sigma = 1$. Then, under $\mathbb{P}_{(i_0, m_0, s_0)}$, there is a standard Brownian motion $(B'_t; t \geq 0)$ (adapted to \mathcal{F}^{C^*}) for which the process $M' = M_t - m_0$, $t \geq 0$ satisfies*

$$M'_t = B'_t - \left(\sup_{s \leq t} M'_s - s'_0 \right)^+ + \left(\inf_{s \leq t} M'_s + i'_0 \right)^-,$$

where $i'_0 = m_0 - i_0$ and $s'_0 = s_0 - m_0$. In other words, M is a doubly perturbed Brownian motion with parameters $\alpha = \beta = -1$.

Proof. The multiparameter martingale $(X_1(z_1) + X_2(z_2) + X_3(z_3); z \in \mathbb{R}_+^3)$ is bounded and right continuous. Hence, Theorem A.1 implies that

$$\xi_t \stackrel{\text{def}}{=} X_1^{C^*}(t) + X_2^{C^*}(t) + X_3^{C^*}(t), t \geq 0$$

is a continuous (single parameter) martingale with respect to the filtration \mathcal{F}^{C^*} . But, the X_i are independent Brownian motions and so the same argument applies to the multiparameter martingale

$$\left((X_1(z_1) + X_2(z_2) + X_3(z_3))^2 - (z_1 + z_2 + z_3); z \in \mathbb{R}_+^3 \right),$$

i.e. $\xi_t^2 - t$ is a martingale. It follows that $(\xi_t; t \geq 0)$ is a Brownian motion with $\xi_0 = i_0 + m_0 + s_0$ and we can take $B' = \xi - (i_0 + m_0 + s_0)$.

Now, \mathcal{C}^* always “runs M ” away from the extrema I and S of $X^{\mathcal{C}^*}$ and so it is no surprise that

$$I_t = \inf_{s \leq t} M_s \wedge i_0, \quad S_t = \sup_{s \leq t} M_s \vee s_0,$$

relationships which can be proved using the arguments of section 4. It follows that

$$M'_t = M_t - m_0 = \xi_t - m_0 - S_t - I_t = B'_t - \sup_{s \leq t} M_s \vee s_0 + s_0 - \inf_{s \leq t} M_s \wedge i_0 + i_0$$

The result now follows by noting that for real a and b we have $a \wedge b - b = -(a - b)^-$ and $a \vee b - b = (a - b)^+$. \square

Lemma 5.1 is relevant because $\tau^{\mathcal{C}^*}$ is precisely the time taken for the doubly perturbed Brownian motion M to exit the interval $(0, 1)$. In particular, the expression we find for the Laplace transform $\hat{v}_r(x)$ can be recovered from Theorems 4 and 5 in Chaumont and Doney [3].

We have so far assumed that $\sigma = 1$ and are yet to say anything about more general “perturbed diffusion processes”. There are several papers that consider this problem. Doney and Zhang [5] consider the existence and uniqueness of diffusions perturbed at their maximum. More recently, Luo [14] has shown that solutions to

$$X'_t = \int_0^t \mu(s, X'_s) ds + \int_0^t \sigma(s, X'_s) dB'_s + \alpha \sup_{s \leq t} X'_s + \beta \inf_{s \leq t} X'_s, \quad (5.35)$$

exist and are unique, but only in the case that $|\alpha| + |\beta| < 1$. A more general perturbed process is considered in [9] but similar restrictions on α and β apply.

That is, there are no existence and uniqueness results for doubly perturbed diffusions which cover our choice of α and β , and less still for the Laplace transform of the distribution of the time taken to exit an interval.

This is where our results seem to contribute something new. Lemma 5.1 easily generalises to continuous $\sigma > 0$, and this combined with the other results in this paper, lets us see that if μ is bounded and Borel measurable and $\sigma > 0$ is continuous, then there is a solution to

$$M'_t = \int_0^t \mu(M'_s) dB'_s + \int_0^t \sigma(M'_s) dB'_s - \sup_{s \leq t} M_s - \inf_{s \leq t} M_s.$$

Furthermore, we can compute the Laplace transform of the distribution of the time taken for any solution of this equation to exit any interval $(-a, b)$ when μ is zero.

6 Concluding remarks and future work

6.1 Majority decisions of $2k + 1$ diffusions and veto voting

The problem that we have solved has natural generalisations in which there are m diffusions instead of the three that we have considered.

In particular, one might ask for the majority decision of an odd number of ‘diffusive voters’ ($X_i(t); t \geq 0$), $i = 1, \dots, m$. Again, we believe that the optimal strategy is to “run the middle”. In other words, if $m = 2k + 1$, and

$$X_1^{C^*}(t) \leq \dots \leq X_k^{C^*}(t) < X_{k+1}^{C^*}(t) < X_{k+2}^{C^*}(t) \leq \dots \leq X_m^{C^*}(t)$$

then C_k^* should increase at unit rate until $X_{k+1}^{C^*}$ hits either $X_k^{C^*}(t)$ or $X_{k+2}^{C^*}(t)$. This prescribes that until then, all other components of X^{C^*} are constant.

A special case of majority voting is ‘veto voting’, where we have an arbitrary number $m' > 0$ of diffusions, and declare a negative decision if at least $k \leq m'$ of them get absorbed at the lower boundary (otherwise no veto occurs and a positive decision is made). To see that this is a majority voting problem, suppose that there is no veto if the majority of voters return positive decisions (i.e. $2k < m'$). This is equivalent to asking for a majority of $m = 2(m' - k) + 1$ diffusive voters, with $m + 1 - 2k$ of them beginning in a state of absorption at the origin. The case $2k \geq m'$ admits a similar majority voting description and in particular, our conjecture for veto voting is that if

$$X_1^{C^*}(t) \leq \dots \leq X_{k-1}^{C^*}(t) < X_k^{C^*}(t) < X_{k+1}^{C^*}(t) \leq \dots \leq X_{m'}^{C^*}(t)$$

then C_k^* should increase at unit rate until $X_k^{C^*}$ hits either $X_{k-1}^{C^*}(t)$ or $X_{k+1}^{C^*}(t)$. In other words, we “run the component with k^{th} order statistic”. The extreme of this is true veto voting in which a single diffusion being absorbed at zero will veto the others. This is the case $k = 1$, and the conjecture is that we should always “run the minimum” of the diffusions.

One might also consider diffusions which obey different stochastic differential equations. We have found an implicit equation for the switching boundaries in the optimal strategy for $m' = 2, k = 1$ ‘veto voting’ problem by solving a free boundary problem. However, we have no conjecture for the general solution.

6.2 Recursive majority revisited

To close, we return to the discrete recursive majority model that motivated us originally (see discussion in the introduction). Recall that r_n denotes the expected cost of the optimal strategy for the n layer tree. The best lower bound for the limit $\gamma = \lim_{n \rightarrow \infty} r_n^{1/n}$ in the literature¹ is $\frac{9}{4}$, which one arrives at by computing the Fourier coefficients of the recursive majority function and applying either an equality due to O’Donnell and Servedio (see §3 of [17]) or Theorem 1.8 of [20]. But, numerics suggest $\gamma \approx 2.45$, leaving big a gap. The best upper bound known to us is $\gamma \leq 2.472$.

In the introduction, we hinted at a continuous approximation to the discrete tree. What we had in mind was to replace each of the Bernoulli random variables on the leaves with a Brownian motion starting from p . These Brownian motions are absorbed at the endpoints of $(0, 1)$ and scaled so that the expected absorption time is one. As with the diffusion model treated in this paper, the observer is billed for the time they spend running each Brownian motion. Let R_n denote the least expect (time) cost for the Brownian tree. In

¹It has been communicated to us that Oded Schramm and Mark Braverman improved this bound to 2.28 but did not publish details.

this paper (see Remark 2.4), we have shown

$$R_1(p) = -\frac{6}{p(1-p)} (p(1-p) + p^2 \ln(p) + (1-p)^2 \ln(1-p)),$$

so $R_1 \leq r_1 = 2(1 + p(1-p))$.

Now, any strategy for the discrete model can also be used in this Brownian case and so R_n is not greater than r_n . It follows that

$$\gamma \geq \limsup_{n \rightarrow \infty} R_n^{1/n},$$

from which we conclude that studying the Brownian tree may help give a lower bound for γ .

Often, the Brownian version of a difficult discrete problem is easier to solve because we have the heavy machinery of stochastic calculus at our disposal. But, we concede that there is no particular reason to think that the n layer Brownian model may be more tractable than the discrete counterpart. Indeed, while we have treated the $n = 1$ case in this paper, we are unable even to give a conjecture on the $n = 2$ optimal strategy.

Still, we might ask, even if it is not possible to determine the optimal strategy, can we say anything about the asymptotics of the expected cost R_n , and in doing so sharpen the bound on γ ? We have not, for example, been able to prove that $R_n^{1/n}$ is eventually monotone in n . Nor do we have the sub-multiplicative structure to guarantee that $\Gamma = \lim_{n \rightarrow \infty} R_n^{1/n}$ even exists.

If the limit Γ does exist, is it equal to γ ? One is tempted to guess affirmatively but it is possible that the optimal strategy runs an exponentially growing number of leaf Brownian motions for very short time, leading to $\Gamma < \gamma$. To us at least, this seems a tough question to answer.

A A result from the theory of multiparameter processes

The proofs of Lemmas 2.2 and 5.1 appealed to the fact that a multiparameter martingales composed with a strategy is again a martingale. Moreover, it was asserted that we can approximate an arbitrary strategy with a discrete one. This appendix contains a precise statement of these results, together with basic definitions (adopted from §4 of [7]).

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space, \mathbb{R}_+ denote the set of non-negative reals $[0, \infty)$ and $d \geq 2$. A family $(\mathcal{F}(\eta); \eta \in \mathbb{R}^d)$ of σ -algebras contained in \mathcal{F} is called a multiparameter filtration if, for every $\eta, \nu \in \mathbb{R}_+^d$ with $\eta \preceq \nu$,

$$\mathcal{F}(\eta) \subseteq \mathcal{F}(\nu).$$

It is assumed that the following ‘‘usual conditions’’ hold;

- $\mathcal{F}(\eta) = \bigcap_{\eta \preceq \nu} \mathcal{F}(\nu)$ for every $\eta \in \mathbb{R}_+^d$ (right continuity)
- $\mathcal{F}(0)$ contains all null sets (completeness)

- for any $\eta, \nu \in \mathbb{R}_+^d$, the σ -algebras $\mathcal{F}(\eta)$ and $\mathcal{F}(\nu)$ are conditionally independent given $\mathcal{F}(\eta \wedge \nu)$.

The final condition (usually referred to as “assumption (F4)” after [2]) is trivially satisfied in our case since \mathcal{F} is generated from independent filtrations.

A real valued process $(Z(\eta); \eta \in \mathbb{R}_+^d)$ is called a multiparameter super-martingale with respect to $(\mathcal{F}(\eta); \eta \in \mathbb{R}_+^d)$ if for every η ,

- $\mathbb{E} |Z(\eta)| < \infty$, i.e. Z is integrable,
- $Z(\eta)$ is $\mathcal{F}(\eta)$ measurable and
- $\mathbb{E}(Z(\eta)|\mathcal{F}(\nu)) \leq Z(\nu)$ for every $\eta \preceq \nu$.

Recall that a strategy \mathcal{C} is a \mathbb{R}_+^d valued process such that \mathcal{C}_i increases from the origin, $\sum_i \mathcal{C}_i(t) = t$ and $\{\mathcal{C}(t) \preceq \eta\} \in \mathcal{F}(\eta)$ for every $t \geq 0$ and $\eta \in \mathbb{R}_+^d$ (conditions (C1)–(C3) from section 1.3). Then, the filtration $(\mathcal{F}^{\mathcal{C}}(t); t \geq 0)$ defined by

$$\mathcal{F}^{\mathcal{C}}(t) \stackrel{\text{def}}{=} \{F \in \mathcal{F} : F \cap \{\mathcal{C}(t) \preceq \eta\} \in \mathcal{F}(\eta) \forall \eta \in \mathbb{R}_+^d\}, t \geq 0$$

satisfies the usual conditions. The process

$$Z^{\mathcal{C}} \stackrel{\text{def}}{=} (Z_1(\mathcal{C}_1(t)), \dots, Z_d(\mathcal{C}_d(t)); t \geq 0)$$

is adapted to this filtration.

The idea is that $Z^{\mathcal{C}}$ should be a super-martingale with respect to $\mathcal{F}^{\mathcal{C}}$. Indeed, Proposition 4.3 in [7] reads

Theorem A.1. *Suppose that Z is a right continuous multi-parameter super-martingale and that \mathcal{C} is a strategy. Then $Z^{\mathcal{C}}$ is a (local) $\mathcal{F}^{\mathcal{C}}$ -super-martingale.*

This theorem appears in various guises throughout the literature (a good reference for the discrete case is chapter 1 of [1]), we do not give the proof. Merely, we will mention one of its stepping stones – approximation of an arbitrary strategy with a discrete one.

Recall from definition 3.3 that for any $\epsilon > 0$, Π_ϵ denotes the set of strategies which only increase in one component over each interval $[k\epsilon, (k+1)\epsilon)$, $k = 0, 1, \dots$, i.e. \mathcal{C}^ϵ is in Π_ϵ if $\dot{\mathcal{C}}_i$ a.e. takes only values 0 or 1 and is constant on each interval $(k\epsilon, (k+1)\epsilon)$. The promised approximation result is

Lemma A.2. *For any strategy \mathcal{C} , there exist a family of strategies $\mathcal{C}^\epsilon \in \Pi_\epsilon$, $\epsilon > 0$ that converge to \mathcal{C} in the sense that*

$$\lim_{\epsilon \rightarrow 0} \sup_{t \geq 0} |\mathcal{C}(t) - \mathcal{C}^\epsilon| = 0,$$

where $|\cdot|$ is any norm on \mathbb{R}^d .

Moreover, there is a positive constant $M > 0$ for which $\mathcal{C}(t) \preceq \mathcal{C}^\epsilon(t + M\epsilon)$ for every $t \geq 0$.

The existence and uniform convergence part of this lemma is exactly Theorem 7 of Mandelbaum [15] and the second part is a corollary to the author’s constructive proof. The details are omitted.

Corresponding author: Peter Windridge p.windridge@warwick.ac.uk

References

- [1] R. Cairoli and R. C. Dalang. *Sequential stochastic optimization*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, 1996. A Wiley-Interscience Publication.
- [2] R. Cairoli and J. B. Walsh. Stochastic integrals in the plane. *Acta Math.*, 134:111–183, 1975.
- [3] L. Chaumont and R. A. Doney. Some calculations for doubly perturbed Brownian motion. *Stochastic Process. Appl.*, 85(1):61–74, 2000.
- [4] L. Chaumont, R. A. Doney, and Y. Hu. Upper and lower limits of doubly perturbed brownian motion. *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, 36(2):219 – 249, 2000.
- [5] R. A. Doney and T. Zhang. Perturbed Skorohod equations and perturbed reflected diffusion processes. *Ann. Inst. H. Poincaré Probab. Statist.*, 41(1):107–121, 2005.
- [6] N. El Karoui and I. Karatzas. Dynamic allocation problems in continuous time. *Annals of Applied Probability*, 4(2):255–286, 1994.
- [7] N. El Karoui and I. Karatzas. Synchronization and optimality for multi-armed bandit problems in continuous time. *Mat. Apl. Comput.*, 16(2):117–151, 1997.
- [8] J. Gittins and D. Jones. A dynamic allocation index for the sequential design of experiments. In J. Gani, editor, *Progress in Statistics*, pages 241–266. North-Holland, Amsterdam, NL, 1974.
- [9] L. Hu and Y. Ren. Doubly perturbed neutral stochastic functional equations. *J. Comput. Appl. Math.*, 231(1):319–326, 2009.
- [10] K. Itô and H. P. McKean, Jr. *Diffusion processes and their sample paths*. Springer-Verlag, Berlin, 1974. Die Grundlehren der mathematischen Wissenschaften, Band 125.
- [11] H. Kaspi and A. Mandelbaum. Lévy bandits: multi-armed bandits driven by Lévy processes. *Ann. Appl. Probab.*, 5(2):541–565, 1995.
- [12] H. Kaspi and A. Mandelbaum. Multi-armed bandits in discrete and continuous time. *The Annals of Applied Probability*, 8:1270–1290, 1998.
- [13] J. Le Gall and M. Yor. Excursions browniennes et carrés de processus de Bessel. *Comptes rendus de l’Académie des sciences. Série 1, Mathématique*, 303(3):73–76, 1986.
- [14] J. Luo. Doubly perturbed jump-diffusion processes. *J. Math. Anal. Appl.*, 351(1):147–151, 2009.
- [15] A. Mandelbaum. Continuous multi-armed bandits and multiparameter processes. *Annals of Probability*, 15(4):1527–1556, 1987.

- [16] A. Mandelbaum, L. Shepp, and R. Vanderbei. Optimal switching between a pair of Brownian motions. *Annals of Probability*, 18(3):1010–1033, 1990.
- [17] Y. Peres, O. Schramm, S. Sheffield, and D. Wilson. Random-turn Hex and other selection games. *American Mathematical Monthly*, 114(5):373–387, 2007.
- [18] M. Perman and W. Werner. Perturbed Brownian motions. *Probability Theory and Related Fields*, 108(3):357–383, 1997.
- [19] D. Revuz and M. Yor. *Continuous martingales and Brownian motion*, volume 293 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, third edition, 1999.
- [20] O. Schramm and J. Steif. Quantitative noise sensitivity and exceptional times for percolation. *Ann. Math. (to appear)*, 2009.
- [21] R. J. Vanderbei. Optimal switching among several brownian motions. *SIAM Journal on Control and Optimization*, 30(5):1150–1162, 1992.
- [22] J. B. Walsh. Optional increasing paths. In *Processus Aléatoires à Deux Indices (Proceedings, Paris 1980)*, pages 172–201. Springer, Berlin, 1981.