Random relaxed controls and partially observed stochastic systems

by

Nigel J. Cutland
Department of Pure Mathematics
University of Hull
Hull HU6 7RX
England

Tom Lindstrøm
Department of Mathematics
University of Oslo
P.O. Box 1053, Blindern
N–0316 Oslo 3
Norway

1 Introduction

Consider the partially observed controlled stochastic system (for $t \leq 1$)

$$x_u(t) = \int_0^t a(x_u, y_u, s, u(y_u, s))ds + \int_0^t b(x_u, y_u, s, u(y_u, s))dw_1(s) \tag{1.1}$$

$$y_u(t) = \int_0^t c(x_u, y_u, s)ds + \int_0^t d(y_u, s)dw_2(s) \tag{1.2}$$

where w_1 and w_2 are independent Brownian motions and where the coefficients a, b, c, d are bounded, continuous functions which may depend on the past of x_u and y_u. The control u takes values in a compact, separable metric space K and is allowed to depend on the past of the observation process y_u. The cost of u is defined to be

$$j(u) = E(h(x_u)), \tag{1.3}$$

where h is a bounded, continuous function on path space. We consider the functions a, b, c, d and h as fixed and wish to find a control u which minimizes the cost. The interpretation is the usual one; y is a series of noisy and partial observations of the process x, and we are seeking the best strategy for controlling x on the basis of these observations.

An important question is what kinds of controls we should allow. The natural choice is the class of *ordinary controls*, which is just the set of all measurable functions

$$u : C([0,1], \mathbf{R}^m) \times [0,1] \to K$$

which are nonanticipating in the sense that

$$u(y, s) = u(z, s)$$

if $y(r) = z(r)$ for all $r \leq s$. But this class has bad closure properties, and the infimum

$$\alpha = \inf\{j(u) : u \text{ is an ordinary control}\}$$

is in general not attained; i.e. an optimal ordinary control does not exist.

One way of getting around this problem is to introduce classes of generalized controls with better closure properties. Two examples from the literature are relaxed controls and wide sense admissible controls. A *relaxed control* is a nonanticipating, measurable function

$$u : C([0,1], \mathbf{R}^m) \times [0,1] \to \mathcal{M}(K),$$

where $\mathcal{M}(K)$ is the set of Radon probability measures on K, while a *wide sense admissible control* to each $y \in C([0,1], \mathbf{R}^m)$ associates not a single path $u(y, \cdot)$, but a whole probability distribution on the set of all such paths (in, of course, a nonanticipating way). Relaxed controls have a history of about thirty years going back to papers by Filippov [15], McShane [23] and Warga [25], but wide sense admissible controls are much more recent; slightly different formulations have been discussed by, among others, Fleming and Pardoux [17], Bismut [4], Haussmann [19], Borkar [5], and Fleming and Nisio [16] (see also [2], [6], [14], and [20] for later developments).

The purpose of the present paper is to introduce a new class of generalized controls called *random relaxed controls*, and to show that under quite general conditions an optimal random relaxed control u exists and satisfies

$$j(u) = \inf\{j(v) : v \text{ is an ordinary control}\}$$

Random relaxed controls are natural amalgamations of relaxed and wide sense admissible controls; to each $y \in C([0,1], \mathbf{R}^m)$ they associate in a nonanticipating way a distribution on the set of measurable functions $\mu : [0,1] \to \mathcal{M}(K)$ (see Section 3 for the technical details).

Although random relaxed controls are entirely standard objects, our approach to them is based on nonstandard analysis and the following very simple idea: Consider the nonstandard version

$$X_U(t) = \int_0^t {}^*a(X_U, Y_U, s, U(Y_U, s))ds + \int_0^t {}^*b(X_U, Y_U, s, U(Y_U, s))d\,{}^*w_1(s) \tag{1.4}$$

$$Y_U(t) = \int_0^t {}^*c(X_U, Y_U, s)ds + \int_0^t {}^*d(Y_U, s)d\,{}^*w_2(s) \tag{1.5}$$

$$J(U) = E({}^*h(X_U)) \tag{1.6}$$

of the system (1.4)–(1.6), and note that by the transfer principle of nonstandard analysis

$$\alpha \equiv \inf\{j(u) : u \text{ is a standard control}\} = \inf\{J(U) : U \text{ is a nonstandard control}\}.$$

Pick a nonstandard control U such that $J(U)$ is infinitely close to α. The idea is that the standard control induced by U will be an optimal control for (1.1)–(1.3).

If we are thinking in terms of ordinary controls, this argument breaks down at the very last step; there just isn't any reasonable way of getting a general nonstandard control to induce an ordinary control. However, we shall show that the random relaxed controls are in a natural sense exactly the standard objects induced by the set of nonstandard controls, and hence the argument above proves the existence of an optimal random relaxed control of cost α. This type of argument is not new, the first author has used it before to study optimal relaxed controls for various kinds of deterministic and stochastic systems, see [8], [9], [10], [11].

The paper is organized as follows. In the next section, we introduce the spaces of measures we shall be working with and give a brief description of their topological properties.

The random relaxed controls are introduced in Section 3, and the relationship between standard and nonstandard controls is studied in Section 4 and 5 – the main result in this part of the paper is Theorem 5.4 which shows that (under certain technical conditions) any random relaxed control can be represented by a nonstandard *ordinary* control. It's not entirely obvious how to obtain a solution of (1.1)–(1.2) when u is a random relaxed control, and we explain our approach to this problem in Section 7 – it requires some knowledge of how the solution of the equation

$$x_{y,\mu}(t) = \int_0^t a(x_{y,\mu}, y, s, \mu(s))ds + \int_0^t b(x_{y,\mu}, y, s, \mu(s))dw(s)$$

(where $y \in C([0,1], \mathbf{R}^m)$ and $\mu : [0,1] \to \mathcal{M}(K)$) depends on y and μ, and these rather technical results are presented in Section 6. In Section 8 we combine results from Sections 2, 6, and 7 to show that the costs induced by corresponding standard and nonstandard controls are equal, and in Section 9 we put all the pieces together and prove the existence of an optimal random relaxed control. We also show that the minimal cost can be approximated arbitrarily well by very simple, finitary controls, and we end the paper by a brief discussion of the conditions we have had to impose.

The paper makes substantial use of nonstandard measure and probability theory, and the reader can find the necessary background in [1] or a combination of [12] and [7].

2 Measure theoretic preliminaries

This section is something of a nuisance; it presents a few facts from measure theory which are important to our later arguments. Since these facts and arguments show up in two different settings and are needed for the formulation of our problems as well as for their solution, we have chosen to give an abstract treatment of them at the outset.

If X is a Hausdroff space, let $\mathcal{M}(X)$ be the space of all Radon probability measures on X endowed with the weak topology. It is known that if X is a metric space, then $\mathcal{M}(X)$ is metrizable by the Prohorov metric (see Appendix III in Billingsley [3] for an exposition; the rather annoying conditions concerning measurable cardinals can be removed using results of Fremlin [18], see also [22].) Two much simpler results are that if X is either compact or metric and separable, then $\mathcal{M}(X)$ has the same properties. Recall that a topological space is *Polish* if it is separable and admits a complete metric.

Let us now fix two Polish spaces X and C, and a Radon probability measure Q on X. We shall assume that C is compact. Define $\mathcal{R}(X, C)$ to be the set of all measurable (w.r.t. Q) functions

$$\mu : X \to \mathcal{M}(C), \tag{2.1}$$

and identify two elements μ_1 and μ_2 of $\mathcal{R}(X, C)$ if $\mu_1(x) = \mu_2(x)$ for Q-almost all x. Define also a subset $\mathcal{M}_Q(X \times C)$ of $\mathcal{M}(X \times C)$ by letting $P \in \mathcal{M}_Q(X \times C)$ if and only if

$$P(A \times C) = Q(A) \tag{2.2}$$

for all Borel sets $A \subset X$.

Given a $\mu \in \mathcal{R}(X \times C)$, we can construct an element $\hat{\mu}$ in $\mathcal{M}(X \times C)$ by letting

$$\hat{\mu}(A \times B) = \int_A \mu(x)(B) dQ(x) \tag{2.3}$$

for all Borel sets A and B, and then extending to a Radon measure on $X \times C$. That such an extension exists and is unique is standard measure theory (it follows, for example, immediately from theorem 3.5.1 in [1]). Using conditional probabilities we can reverse the construction:

2.1 Lemma. The map $\mu \mapsto \hat{\mu}$ is a bijection from $\mathcal{R}(X, C)$ to $\mathcal{M}_Q(X \times C)$.

Proof: Let $\hat{\mu} \in \mathcal{M}_Q(X \times C)$, and let $\mathcal{A}$ denote the σ-algebra consisting of $\hat{\mu}$-measurable sets of the form $A \times C$. Since $X \times C$ is Polish, Theorem 1.1.6 in Stroock-Varadhan [24] tells us that the conditional probability $\mu(x)(\cdot)$ of $\hat{\mu}$ with respect to $\mathcal{A}$ exists, and that for each x, $\mu(x)(\cdot)$ is an element of $\mathcal{M}(C)$ satisfying (2.3)

The next lemma is a natural and useful extension of formula (2.3):

2.2 Lemma If $f : X \times C \to \mathbf{R}$ is a bounded Borel function and $\mu \in \mathcal{R}(X, C)$, then

$$\int f(x, c) d\hat{\mu}(x, c) = \iint f(x, c) d\mu(x)(c) dQ(x) \tag{2.4}$$

Proof: Let $\mathcal{A}$ be the class of all subsets D of $X \times C$ satisfying

$$\hat{\mu}(D) = \iint 1_D(x, c) d\mu(x)(c) dQ(x)\,. \tag{2.5}$$

Clearly, $\mathcal{A}$ contains the family $\mathcal{A}_0$ of all finite unions $\bigcup_{i=1}^{n} (A_i \times B_i)$ where A_i, B_i are Borel sets. Moreover, by the Monotone Convergence Theorem $\mathcal{A}$ is closed under increasing, countable unions. Since $\mathcal{A}_0$ is an algebra, the Monotone Class Theorem tells us that $\mathcal{A}$ contains the σ-algebra $\sigma(\mathcal{A}_0)$ generated by $\mathcal{A}_0$. But since X and C are separable metric spaces, $\sigma(\mathcal{A}_0)$ is exactly the Borel algebra on $X \times C$, and hence (2.5) holds for all Borel sets. Approximating f by simple functions, the lemma follows.

The natural topology on $\mathcal{M}_Q(X \times C)$ is the one inherited from $\mathcal{M}(X \times C)$. It turns out to have very nice properties:

2.3 Lemma $\mathcal{M}_Q(X \times C)$ is a compact Polish space.

Proof: Since $X \times C$ is a separable metric space, $\mathcal{M}(X \times C)$ is separable and metrizable. It thus suffices to show that $\mathcal{M}_Q(X \times C)$ is compact.

We shall use the nonstandard characterization of compactness; given $\mu \in {}^*\mathcal{M}_Q(X \times C)$, we must show that μ is nearstandard and that its standard part belongs to $\mathcal{M}_Q(X \times C)$. This is almost trivial; since C is compact and $\mu(A \times C) = {}^*Q(A)$, the "pushed down" Loeb-measure $L(\mu) \circ st^{-1}$ is the standard part of μ and it clearly belongs to $\mathcal{M}_Q(X \times C)$ (see section 3.4 of [1] for the necessary background).

On $\mathcal{R}(X, C)$ we put the topology generated by the basic open sets

$$\mathcal{O}_{f,\epsilon,\mu_0} = \{\mu : \left| \int\int f(x,c) d\mu(x)(c) dQ(x) - \int\int f(x,c) d\mu_0(x)(c) dQ(x) \right| < \epsilon\}$$

where $\mu_0 \in \mathcal{R}(X, C)$, $\epsilon \in \mathbf{R}_+$, and $f : X \times C \to \mathbf{R}$ is a bounded continuous function.

2.4 Corollary $\mathcal{R}(X, C)$ is a compact Polish space.

Proof: According to lemmas 2.1 and 2.2, the map $\mu \mapsto \hat{\mu}$ is an homeomorphism, and the result thus follows from Lemma 2.3.

Our last result in this section concerns the interplay between $\mathcal{M}_Q(X, C)$ and its nonstandard version ${}^*\mathcal{M}_Q(X, C)$. In Section 8 it will be used to establish the relationship between standard and nonstandard costs.

2.5 Lemma Let $\hat{U} \in {}^*\mathcal{M}_Q(X \times C)$ and define $\hat{u} = L(\hat{U}) \circ st^{-1}$. If

$$\theta : X \times C \to \mathbf{R}$$

is a bounded, measurable function which is continuous in the second variable, then

$$\int \theta(x,c) d\hat{u}(x,c) =^\circ \int {}^*\theta(x,c) d\hat{U}(x,c) \,.$$

Proof: Define $\tilde{\theta} : X \to C(C, \mathbf{R})$ by

$$\tilde{\theta}(x) = \theta(x, \cdot) \,.$$

By Anderson's Lusin Theorem (see, e.g., Corollary 3.4.9 in [1]), there is a set $X_0 \subset {}^*X$ of $L({}^*Q)$-measure one such that ${}^{\circ *}\tilde{\theta}(x) = \tilde{\theta}(\circ x)$ for all $x \in X_0$. Hence ${}^{\circ *}\theta(x, c) = \tilde{\theta}(\circ x, \circ c)$ for all $x \in X_0$ and all $c \in C$, and since this means that ${}^*\theta$ is a lifting of θ with respect to $\hat{U}$, the lemma follows.

Before we end this section, let us observe that any measurable function $g : X \to C$ may be considered as an element of $\mathcal{R}(X, C)$; just identify $g(x)$ with the unit point mass $\delta_{g(x)}$ at $g(x)$. We shall denote this subspace of $\mathcal{R}(X, C)$ by $\mathcal{R}_0(X, C)$, i.e.,

$$\mathcal{R}_0(X, C) = \{\delta_g \in \mathcal{R}(X, C) | g : X \to C \text{ is measurable}\} \tag{2.6}$$

In the sequel we shall apply the results of this section in two different settings. In the first, X is the interval [0,1], C is the control space K, and Q is the Lebesgue measure. In the second, X is the path space of the observation process, C is $\mathcal{R}([0,1], K)$, and Q is a reference measure on X such that the measure induced by the observation process is always absolutely continuous with respect to Q.

3 Ordinary, relaxed, and random relaxed controls

The spaces $\mathcal{R}_0([0,1], K)$ and $\mathcal{R}([0,1], K)$ (where K is a fixed compact space, the control space) will play important parts in this paper, and it is convenient to introduce the abbreviations

$$\mathcal{R}_0 = \mathcal{R}_0([0,1], K), \qquad \mathcal{R} = \mathcal{R}([0,1]), K). \tag{3.1}$$

The underlying measure on [0,1] will always be the Lebesgue measure. An element in $\mathcal{R}_0$ is called a *response*, while an element in $\mathcal{R}$ is a *relaxed response.*

In what follows, we shall think of

$$\mathcal{Y} = C([0,1], \mathbf{R}^m)$$

as the space of all possible observations. If $y \in \mathcal{Y}$ and $t \in [0,1]$, we shall write $y \upharpoonright t$ for the restriction of y to $[0,t]$. We shall also fix a Radon measure Q on $\mathcal{Y}$, and think of it as the measure induced by the observation process (or, more correctly, as a fixed reference measure on $\mathcal{Y}$ such that the measure induced by the observation process is always absolutely continuous w.r.t. Q). In this setting, an *ordinary control* (or simply a *control*) is just a measurable function

$$u : \mathcal{Y} \to \mathcal{R}_0 \tag{3.2}$$

which is *nonanticipating* in the sense that if $y \upharpoonright t = y' \upharpoonright t$, then $u(y)(t) = u(y')(t)$. A *relaxed control* is a measurable function

$$u : \mathcal{Y} \to \mathcal{R} \tag{3.3}$$

satisfying the same nonanticipation condition. Hence a (relaxed) control is a nonanticipating function which to each observation associates a (relaxed) response. Roughly speaking, a random relaxed control is a nonanticipating function which to each observation assigns a probability distribution on the set of relaxed responses. To make this precise, we must first agree on what it should mean for such a function to be nonanticipating.

A subset A of $\mathcal{R}$ is *determined at time t* if it is Borel and has the property that if $\mu \in A$ and $\mu \restriction t = \mu' \restriction t$, then $\mu' \in A$. In other words, $A \in \mathcal{F}_t$ where $\mathcal{F}_t = \sigma\{\mu(s) | s \leq t\}$ is the natural filtration on $\mathcal{R}$.

3.1 Definition A *random relaxed control* u is a measurable function $u : \mathcal{Y} \to \mathcal{M}(\mathcal{R})$ with the following property: If $y \restriction t = y' \restriction t$ and A is determined at time t, then

$$u(y)(A) = u(y')(A) \tag{3.4}$$

Equivalently, we could say that for each $A \in \mathcal{F}_t$, the map $y \to u(y)(A)$ is $\mathcal{C}_t$-measurable, where $\mathcal{C}_t = \sigma\{y(s) : s \leq t\}$ is the natural filtration on $\mathcal{Y}$.

Note that we can also think of a random relaxed control as an element of the space $\mathcal{R}(\mathcal{Y}, \mathcal{R})$.

Since all our topologies are defined in terms of continuous functions, it will be useful to have a characterization of random relaxed controls in terms of such functions rather than sets. A function $f : \mathcal{R} \to \mathbf{R}$ is *determined at time t* if

$$f(\mu) = f(\mu') \tag{3.5}$$

whenever $\mu \restriction t = \mu' \restriction t$. If $k : \mathcal{Y} \to \mathbf{R}$, let

$$k_t = E_Q(k|\mathcal{C}_t) \tag{3.6}$$

be the conditional expectation of k with respect to the measure Q and the filtration $\mathcal{C}_t$ generated up to time t. The following lemma is a straightforward exercise in measure theory which we shall leave to the reader.

3.2 Lemma A measurable function $u : \mathcal{Y} \to \mathcal{M}(\mathcal{R})$ is a random relaxed control if and only if the following holds: For all bounded, continuous functions $f : \mathcal{R} \to \mathbf{R}$, $k : \mathcal{Y} \to \mathbf{R}$ such that f is determined at time t,

$$\int\int [k(y) - k_t(y)] \cdot f(\mu) du(y)(\mu) dQ(y) = 0 \tag{3.7}$$

(Let us make it quite clear what the left hand side of (3.7) means. For each $y \in \mathcal{Y}$, $u(y)(\cdot)$ is a measure on the space $\mathcal{R}$ of relaxed responses, and we first integrate the function

$\mu \mapsto [k(y) - k_t(y)]f(\mu)$ against this measure. The result is a function of y, which we then integrate against the measure Q on $\mathcal{Y}$.)

Observe that given a relaxed control u, we can construct a random relaxed control u' by

$$u'(y) = \delta_{u(y)}\,,$$

where $\delta_{u(y)}$ is the unit mass at $u(y)$. Hence we can always consider the relaxed controls as a subset of the random relaxed controls. Since the ordinary controls are special kinds of relaxed controls, an ordinary control can also be considered as a random relaxed control in the obvious way.

4 Standard parts of nonstandard controls

Random relaxed controls are quite complicated, abstract objects, and the reader may well wonder where they come from and what they are good for. In this section, we shall give a partial answer to these questions by showing that random relaxed controls arise naturally as the standard parts of nonstandard ordinary controls.

Let us first try to explain this informally. Assume that U is a nonstandard control; i.e. U is a nonanticipating, internal function

$$U : {}^*\mathcal{Y} \to {}^*\mathcal{R}_0\,, \tag{4.1}$$

and let us try to find U's standard part u. There are two aspects of U we cannot capture if we insist that u should be an ordinary control. To see the first, let $y \in \mathcal{Y}$ and $t \in [0,1]$. If $s, r \in {}^*[0,1]$ are both infinitely close to t, there is no reason why $U({}^*y)(s)$ and $U({}^*y)(r)$ should be infinitely close. Hence there is no single, natural value to assign to $u(y)$ at time t; all we can prescribe is the distribution of ${}^\circ U({}^*y)(s)$ as s ranges over the monad of t. This explains why, in general, the standard part of U will have to be a relaxed control. The other difficulty is of a similar nature. Assume as before that $y \in \mathcal{Y}$, and let y_1, y_2 be two elements in ${}^*\mathcal{Y}$ infinitely close to y. Again there is no reason why $U(y_1)$ and $U(y_2)$ should be infinitely close, and thus there is no canonical way of assigning a single relaxed response to y. What is naturally given is the distribution of ${}^\circ U$ over the monad of y, and this leads us to the notion of a random relaxed control.

When we next try to make this argument rigorous, it will be useful to work with a more general problem. Starting with a nonstandard, random relaxed control

$$U : {}^*\mathcal{Y} \to {}^*\mathcal{M}(\mathcal{R})\,, \tag{4.2}$$

we shall see how it can be turned into a standard random relaxed control. Assume that $\tilde{Q}$ is an internal, Borel probability measure on ${}^*\mathcal{Y}$ supported on the nearstandard elements, and let

$$Q = L(\tilde{Q}) \circ st^{-1} \tag{4.3}$$

be its standard part. Let

$$\hat{U} \in {}^*\mathcal{M}_{\tilde{Q}}({}^*\mathcal{Y} \times {}^*\mathcal{M}(\mathcal{R})) \tag{4.4}$$

be the measure induced by U and $\tilde{Q}$ as defined in (2.3). Since $\mathcal{M}(\mathcal{R})$ is compact and Q is the standard part of $\tilde{Q}$, it is easy to check that

$$\hat{u} = L(\hat{U}) \circ st^{-1} \tag{4.5}$$

is an element of $\mathcal{M}_Q(\mathcal{Y} \times \mathcal{M}(\mathcal{R}))$. Using the bijection in lemma 2.1, we obtain an element

$$u : \mathcal{Y} \to \mathcal{M}(\mathcal{R}) \tag{4.6}$$

which we shall refer to as the *standard part* of U.

4.1 Lemma The standard part of a nonstandard, random relaxed control is a random relaxed control.

Proof: Let U be the nonstandard, random relaxed control and u its standard part. Assume that $f : \mathcal{R} \to \mathbf{R}$ and $k : \mathcal{Y} \to \mathbf{R}$ are bounded continuous functions, and that f is determined at time t. We have

$$\begin{aligned} &{}^\circ\!\int\!\!\int [{}^*k(y) - {}^*k_t(y)]\, {}^*f(\mu) dU(y)(\mu) d\tilde{Q}(y) \\ &\quad = {}^\circ\!\int\!\!\int [{}^*k(y) - {}^*k_t(y)]\, {}^*f(\mu) d\hat{U}(y,\mu) \\ &\quad = \int\!\!\int [k(y) - k_t(y)] f(\mu) d\hat{u}(y,\mu) \\ &\quad = \int\!\!\int [k(y) - k_t(y)] f(\mu) du(y)(\mu) dQ(y)\,, \end{aligned} \tag{4.7}$$

where the first and last equality hold by lemma 2.2, and the second one by the definition of $\hat{u}$. By the *-version of lemma 3.2, the first integral in (4.7) is zero. Hence the last integral is also zero, and thus u is a random relaxed control by lemma 3.2.

5 Liftings of random relaxed controls

Assume that $\tilde{Q}$ is an internal, nearstandardly concentrated probability measure on ${}^*\mathcal{Y}$ as in the last section, and let $Q = L(\tilde{Q}) \circ st^{-1}$ be its standard part. A *lifting* of a random relaxed control u is a nonstandard random relaxed control U whose standard part is u. If U is a nonstandard ordinary control, we call it an *ordinary lifting* of u. The key observation behind this paper is that (under some technical conditions) all random relaxed controls have ordinary liftings. To prove this, we shall need the following lemma which is just a nonstandard way of saying that $\mathcal{R}_0$ is dense in $\mathcal{R}$.

5.1 Lemma There is an internal map $\pi : {}^*\mathcal{R} \to {}^*\mathcal{R}_0$ such that $\pi(\mu) \approx \mu$, and $\pi(\mu)(t) = \pi(\mu')(t)$ if $\mu{\restriction}t = \mu'{\restriction}t$.

Proof: Partition ${}^*K$ into a hyperfinite family of *Borel subsets $K_1, K_2, \ldots, K_H$ in such a way that each K_j is contained in a monad. Pick one element k_j from each equivalence class K_j in an internal way. For each $\mu \in {}^*\mathcal{R}$ and $i < H$, let μ_i be the measure on ${}^*K$ defined by

$$\mu_i(A) = H \int_{i/H}^{(i+1)/H} \mu_t(A)dt \tag{5.1}$$

The function $\bar{\mu} = \pi(\mu) : {}^*[0,1] \to {}^*K$ will be constant on each interval of the form $\left[\frac{j}{H^3}, \frac{(j+1)}{H^3}\right)$, and will take no other values than $k_1, k_2, \ldots, k_H$, so it is sufficient to define $\bar{\mu}(s)$ for $s = \frac{j}{H^3}$.

On the interval $\left[0, \frac{1}{H}\right)$ let $\bar{\mu}$ take some arbitrary value – say $\bar{\mu}(s) = k_1$. Now assume that $\bar{\mu}(s)$ has been defined for all $s < \frac{j}{H^3}$ and that $\frac{(i+1)}{H} \leq \frac{j}{H^3} < \frac{(i+2)}{H}$. Define $\bar{\mu}\left(\frac{j}{H^3}\right) = k_r$ where r is the smallest number maximizing

$$H^2\mu_i(K_r) - \left|\{s : \frac{i+1}{H} \leq s < \frac{j}{H^3} \text{ and } \bar{\mu}(s) = k_r\}\right|. \tag{5.2}$$

If we put

$$a_r = \left|\{s : \frac{i+1}{H} \leq s < \frac{i+2}{H} \text{ and } \bar{\mu}(s) = k_r\}\right|,$$

then clearly $\sum a_r = H^2$ and $a_r \leq H^2\mu_i(K) + 1$. Thus for any $D \in^* \sigma\{K_1, \ldots, K_H\}$

$$\frac{1}{H^2} \sum_{k_r \in D} a_r \leq \sum_{k_r \in D} \mu_i(K_r) + \frac{1}{H} \approx \mu_i(D),$$

and so, in fact,

$$\frac{1}{H^2} \sum_{k_r \in D} a_r \approx \mu_i(D).$$

From this it is easy to see that $\bar{\mu} \approx \mu$ (as in [11]). Since we have used μ_i to define $\bar{\mu}$ on $\left[\frac{(i+1)}{H^3}, \frac{(i+2)}{H^3}\right]$, it is immediately clear that $\pi(\mu)(t) = \pi(\mu')(t)$ if $\mu \restriction t = \mu' \restriction t$. This completes the proof.

Given an element ν in ${}^*\mathcal{M}(\mathcal{R})$, we can turn it into an element $\pi(\nu)$ in ${}^*\mathcal{M}(\mathcal{R}_0)$ by

$$\pi(\nu)(A) = \nu\{\mu : \pi(\mu) \in A\}. \tag{5.3}$$

Since $\pi(\mu) \approx \mu$, it is easy to check that $\pi(\nu) \approx \nu$. We can now prove:

5.2 Lemma Any random relaxed control u has a random ordinary lifting U; i.e., there exists a nonanticipating function $U : {}^*\mathcal{Y} \to {}^*\mathcal{M}(\mathcal{R}_0)$ whose standard part is u.

Proof: Since $\mathcal{M}(\mathcal{R})$ is separable, there is an internal map $\tilde{U} : {}^*\mathcal{Y} \to {}^*\mathcal{M}(\mathcal{R})$ such that ${}^\circ\tilde{U}(y) = u({}^\circ y)$ $L(\tilde{Q})$-almost everywhere. Define $\tilde{U}_0 : {}^*\mathcal{Y} \to {}^*\mathcal{M}(\mathcal{R}_0)$ by $\tilde{U}_0(y) = \pi(\tilde{U}(y))$. By construction of π, ${}^\circ\tilde{U}_0(y) = u({}^\circ y)$ $L(\tilde{Q})$-almost everywhere. It remains to turn $\tilde{U}_0$ into a nonanticipating function U.

Let $\mathcal{C}_t$ be the internal σ-algebra on ${}^*\mathcal{Y}$ generated up to time t; i.e., $\mathcal{C}_t$ consists of all *Borel sets A with the property that if $y \in A$ and $y(s) = z(s)$ for all $s \leq t$, then $z \in A$. The idea is to make $U \restriction t$ the conditional expectation of $\tilde{U}_0$ with respect to $\mathcal{C}_t$. To make this precise, observe that by the proof of Lemma 5.1, $\tilde{U}_0$ is supported on the set of responses which are constant on each interval $\left[\frac{j}{H^3}, \frac{j+1}{H^3}\right)$ and take values in the set $\hat{K} = \{k_1, k_2, \ldots, k_H\}$. Since U will be supported on the same set, it suffices to specify

$$U(y)(\langle c_0, c_1, \ldots, c_j \rangle) \tag{5.4}$$

for each internal sequence $c_0, c_1, \ldots, c_j \in \hat{K}$, $j < H^3$, where

$$\langle c_0, c_1, \ldots, c_j \rangle = \left\{ \mu \in {}^*\mathcal{R}_0 : \mu(s) = c_i \text{ for all } s \in \left[\frac{i}{H^3}, \frac{i+1}{H^3}\right) \right\}. \tag{5.5}$$

Let

$$U(y)(\langle c_0, c_1, \ldots, c_j \rangle) = E_{\tilde{Q}}(\tilde{U}_0(\cdot)(\langle c_0, c_1, \ldots, c_j \rangle) \big| \mathcal{C}_{j/H^3})(y), \tag{5.6}$$

then for $\tilde{Q}$-almost all y,

$$U(y)(\langle c_0, c_1, \ldots, c_j \rangle) = \sum_{c \in \hat{K}} U(y)(\langle c_0, c_1, \ldots, c_j, c \rangle). \tag{5.7}$$

By modifying U appropriately on the remaining null set, we can make (5.7) hold for all y. But then each $U(y)(\cdot)$ can be extended to an internal measure on ${}^*\mathcal{R}_0$ in an obvious way, and it's easy (but a bit tedious) to check that U is a random ordinary lifting of u.

Remark. If $\tilde{Q} = {}^*Q$ is the nonstandard version of a standard measure Q, there is a very simple proof of Lemma 5.2 – we can simply let $U(y) = \pi({}^*U(y))$.

We need to go one step further and turn the random ordinary liftings in Lemma 5.2 into ordinary liftings. In doing so, we shall find the following notation and terminology helpful.

If M is an infinite integer, then a set of the form

$$\left\{(\xi_1,\ldots,\xi_m)\in {}^*\mathbf{R}^m \middle| \frac{j_i}{M} \leq \xi_i < \frac{j_i+1}{M} \text{ for } i=1,\ldots,m\right\},$$

where $j_1,\ldots,j_m \in {}^*\mathbf{Z}$, is called an *M-set.* For each $t \in {}^*\mathbf{R}_+$ define an equivalence relation $\sim_{t,M}$ on ${}^*\mathcal{Y}$ by

$$y_1 \sim_{t,M} y_2 \Leftrightarrow y_1\left(\tfrac{i}{H^3}\right) \text{ and } y_2\left(\tfrac{i}{H^3}\right) \text{ belong to the same } M\text{-set for all } \tfrac{i}{H^3} \leq t$$

Let $[y]_{t,M}$ denote the equivalence class of y with respect to $\sim_{t,M}$.

5.3 Definition We shall call the internal measure $\tilde{Q}$ *smooth* if the following two conditions are satisfied for all $y \in {}^*\mathcal{Y}$ and all $k \in {}^*\mathbf{Z}$, $0 \leq k < H^3$:

(i) There is an infinite integer M and a positive infinitesimal ε such that for all $\mathcal{C}_{k/H^3}$-measurable subsets B of $[y]_{k/H^3,M}$ and all M-sets A

$$\left(1+\frac{\varepsilon}{H^3}\right)^{-1} \leq \frac{\tilde{Q}\{z_{(k+1)/H^3} \in A | z \in [y]_{k/H^3,M}\}}{\tilde{Q}\{z_{(k+1)/H^3} \in A | z \in B\}} \leq 1+\frac{\varepsilon}{H^3}.$$

(ii) If $B \subset [y]_{k/H^3,M}$ is $\mathcal{C}_{k/H^3}$-measurable with $\tilde{Q}(B) > 0$, then the measure

$$A \mapsto \tilde{Q}\{z \in A | z \in B\}$$

is nonatomic.

In most examples, the integer M will be infinite compared to the (already infinite) integer H^3.

5.4 Theorem Let $\tilde{Q}$ be an internal, nearstandardly concentrated and smooth probability measure on ${}^*\mathcal{Y}$. Then any random relaxed control u is the standard part of a nonstandard ordinary control U.

Proof: Given a random relaxed control u, let $\tilde{U}$ be the random ordinary lifting constructed in the proof of Lemma 5.2. We shall construct U by modifying $\tilde{U}$, and just as for $\tilde{U}$, all the responses of U will be constant on the intervals $\left[\frac{i}{H^3}, \frac{i+1}{H^3}\right)$ and take values in the hyperfinite set $\hat{K} = \{k_1, \ldots, k_H\}$.

Since $\tilde{Q}$ is smooth, we can fix an infinite integer M satisfying Definition 5.3.1. To simplify the notation, we shall write $\sim_k$ and $[y]_k$ for $\sim_{k/H^3,M}$ and $[y]_{k/H^3,M}$, respectively. We shall also find it convenient to write

$$\overline{U(z)}(k) = \bar{c}_k$$

as an abbreviation of "$U(z)\left(\frac{i}{H^3}\right) = c_i$ for each $i \le k$".

We are now ready to construct U, but before we begin, let us admit that our U will be slightly flawed in one respect – instead of depending on the behaviour of z up to time $\frac{k}{H^3}$ as it should in order to be nonanticipating, $U(z)\left(\frac{k}{H^3}\right)$ will, in fact, depend on z all the way up to $\frac{k+1}{H^3}$. This flaw is easily fixed; if we just delay the execution of the strategy by $\frac{1}{H^3}$, we get a new strategy which is nonanticipating and which has the same standard part as the old one. (We could, of course, have avoided this problem by defining the delayed control directly, but this would have made our formulas much less intuitive.)

Assume that we can define a nonstandard ordinary control U such that for all $y \in {}^*\mathcal{Y}$ and all $n < H^3$, we have

$$\tilde{Q}\{z \in [y]_{n+1} \wedge \overline{U(z)}(n) = \bar{c}_n\} = \alpha_n(y) \int_{[y]_{n+1}} \tilde{U}(z)(c_0, \ldots, c_n) d\tilde{Q}(z) \tag{5.8}$$

where $\alpha_n(y) \approx 1$. It is then an easy exercise in nonstandard measure theory to show that U and $\tilde{U}$ have the same standard part, and hence that U is an ordinary lifting of u. We shall leave this exercise to the reader, and concentrate on proving (5.8).

Let us assume that we have defined U up to time $\frac{n-1}{H^3}$, and that we now want to define $U(z)\left(\frac{n}{H^3}\right)$ for all $y \in {}^*\mathcal{Y}$ in such a way that (5.8) holds. Observe first that if we sum both sides of (5.8) for all possible choices of $c_n \in \hat{K}$, we get

$$\tilde{Q}\{z \in [y]_{n+1} \wedge \overline{U(z)}(n-1) = \bar{c}_{n-1}\} = \alpha_n(y) \int_{[y]_{n+1}} \tilde{U}(z)(c_0, \ldots, c_{n-1}) d\tilde{Q}(z) \,,$$

which means that $\alpha_n(y)$ has to satisfy the consistency condition

$$\alpha_n(y) = \frac{\tilde{Q}\{z \in [y]_{n+1} \wedge \overline{U(z)}(n-1) = \bar{c}_{n-1}\}}{\int_{[y]_{n+1}} \tilde{U}(z)(c_0, \ldots, c_{n-1}) d\tilde{Q}(z)} \,. \tag{5.9}$$

On the other hand, the nonatomicity condition 5.3.(ii) guarantees that once (5.9) is satisfied, it is possible to choose $U(z)\left(\frac{n}{H^3}\right)$ in such a way that (5.8) holds and $U(z)\left(\frac{n}{H^3}\right)$ only depends on the behaviour of z up to time $\frac{n+1}{H^3}$.

We now have an inductive procedure for defining U, and it only remains to show that the $\alpha_n(y)$ in (5.9) is infinitely close to one. Starting with the numerator, we see that

$$\begin{aligned}
&\tilde{Q}\{z\in[y]_{n+1} \wedge \overline{U(z)}(n-1) = \bar{c}_{n-1}\} = \\
&\quad = \tilde{Q}\{z\in[y]_{n+1} | \overline{U(z)}(n-1) = \bar{c}_{n-1} \wedge z\in[y]_n\} \cdot \tilde{Q}\{\overline{U(z)}(n-1) = \bar{c}_{n-1} \wedge z\in[y]_n\} \\
&\quad = \beta\tilde{Q}\{z\in[y]_{n+1} | z\in[y]_n\} \cdot \alpha_{n-1}(y) \int_{[y]_n} \tilde{U}(z)(c_0,\dots,c_{n-1})d\tilde{Q}(z)\,,
\end{aligned}$$

where $\left(1+\frac{\varepsilon}{H^3}\right)^{-1} \leq \beta \leq \left(1+\frac{\varepsilon}{H^3}\right)$ (we have used condition 5.3(i) in the first factor and the definition of $\alpha_{n-1}(y)$ in the second). Similarly, we get for the denominator

$$\begin{aligned}
&\int_{[y]_{n+1}} \tilde{U}(z)(c_0,\dots,c_{n-1})d\tilde{Q}(z) = \\
&\quad = \gamma \cdot \int_{[y]_n} \tilde{U}(z)(c_0,\dots,c_{n-1})d\tilde{Q}(z) \cdot \tilde{Q}\{z\in[y]_{n+1} | z\in[y]_n\}
\end{aligned}$$

where $\left(1+\frac{\varepsilon}{H^3}\right)^{-1} \leq \gamma \leq \left(1+\frac{\varepsilon}{H^3}\right)$. Hence $\alpha_n(y) = \frac{\beta}{\gamma}\alpha_{n-1}(y)$, which means that

$$\left(1+\frac{\varepsilon}{H^3}\right)^{-2}\alpha_{n-1}(y) \leq \alpha_n(y) \leq \left(1+\frac{\varepsilon}{H^3}\right)^{2}\alpha_{n-1}(y)\,.$$

By induction,

$$\left(1+\frac{\varepsilon}{H^3}\right)^{-2n} \leq \alpha_n(y) \leq \left(1+\frac{\varepsilon}{H^3}\right)^{2n}.$$

and since ε is infinitesimal and $n \leq H^3$, it follows that $\alpha_n(y) \approx 1$, and the proof is complete.

Remark: In many applications the smoothness condition is difficult to verify, but in Lemma 9.1 we shall indicate a way around this problem.

6 Dependence on observations and controls

So far we have only studied the relationship between various kinds of controls, but we have now reached the stage where we can begin to approach our stochastic system (1.1)–(1.3). Obviously, the performance part x_u of this system depends on the control u and the observations y_u, and in this section we want to study this dependence in an abstract setting.

Let $\mathcal{X} = C([0,1], \mathbf{R}^n)$ and assume that a and b are bounded, continuous functions

$$\begin{aligned}
&a : \mathcal{X} \times \mathcal{Y} \times [0,1] \times K \to \mathbf{R}^n \\
&b : \mathcal{X} \times \mathcal{Y} \times [0,1] \times K \to S(n)
\end{aligned}$$

($S(n)$ is the set of symmetric $n\times n$-matrices) which are *nonanticipating* in the sense that if $x\restriction t=x'\restriction t$ and $y\restriction t=y'\restriction t$, then $a(x,y,t,k)=a(x',y',t,k)$ for all k, and similarly for b.

Given a path $y\in\mathcal{Y}$ and a relaxed response $\mu\in\mathcal{R}$, we shall study the Itô-equation

$$x_{y,\mu}(t)=\int_0^t a(x_{y,\mu},y,s,\mu(s))ds+\int_0^t b(x_{y,\mu},y,s,\mu(s))dw(s) \tag{6.1}$$

As yet, this equation only makes sense when μ is an ordinary response, $\mu\in\mathcal{R}_0$, but for general μ we shall simply interpret it as

$$x_{y,\mu}(t)=\int_0^t a_\mu(x_{y,\mu},y,s)ds+\int_0^t b_\mu(x_{y,\mu},y,s)dw(s) \tag{6.2}$$

where

$$a_\mu(x,y,s)=\int a(x,y,s,k)d\mu(s)(k) \tag{6.3}$$

and

$$b_\mu(x,y,s)=\Big[\int b(x,y,s,k)^2d\mu(s)(k)\Big]^{1/2}. \tag{6.4}$$

The square in the definition of b_μ is natural since it is b^2 rather than b itself which determines the dynamics of the process (see Cutland [11] for further comments).

6.1 Proposition Fix a $\mu\in\mathcal{R}$, and assume that (6.1) has a pathwise unique solution for each $y\in\mathcal{Y}$. For each Radon probability measure Q on $\mathcal{Y}$, we can choose versions of these solutions such that the map $(\omega,y)\mapsto x_{y,\mu}(\omega,\cdot)$ is $P\times Q$-measurable.

Proof: The proof falls naturally into two parts. In the first we show that the map $y\mapsto x_{y,\mu}$ is continuous with respect to the norm

$$\|x_{y,\mu}\|=\Big(\int\sup_{t\leq 1}|x_{y,\mu}(\omega,t)|^2dP(\omega)\Big)^{1/2}.$$

Pick $\hat{y}\in{}^*\mathcal{Y}$ infinitely close to $y\in\mathcal{Y}$, and let

$$X_{\hat{y},{}^*\mu}(t)=\int_0^t {}^*a(X_{\hat{y},{}^*\mu},\hat{y},s,{}^*\mu)ds+\int_0^t {}^*b(X_{\hat{y},{}^*\mu},\hat{y},s,{}^*\mu)d^*w(s) \tag{6.5}$$

Clearly, the nonstandard version ${}^*x_{y,\mu}$ of the solution of (6.2) satisfies

$${}^*x_{y,\mu}(t)=\int_0^t {}^*a({}^*x_{y,\mu},{}^*y,s,{}^*\mu)ds+\int_0^t {}^*b({}^*x_{y,\mu},{}^*y,s,{}^*\mu)d^*w(s) \tag{6.6}$$

and taking standard parts on both sides of (6.5) and (6.6), we see that the standard parts $^\circ X_{\hat{y},^*\mu}$ and $^{\circ *}x_{y,\mu}$ both satisfy (6.2). Since we have assumed that (6.2) has only one solution, this means that $\|X_{\hat{y},^*\mu} - {}^*x_{y,\mu}\| \approx 0$, and hence $y \mapsto x_{y,\mu}$ is continuous.

We are now ready for the second part of the proof, in which the continuity of $y \mapsto x_{y,\mu}$ is used to approximate $(\omega, y) \mapsto x_{y,\mu}(\omega)$ by $P{\times}Q$-measurable simple functions. For each $n \in \mathrm{N}$, choose a compact set K_n with $Q(K_n) > 1 - \frac{1}{n}$ in such a way that $\{K_n\}$ is an increasing sequence whose union is $\mathcal{Y}$. Partition K_n into a finite number of sets $A_1^{(n)}, \ldots, A_{m_n}^{(n)}$ such that if $y_1, y_2 \in A_i^{(n)}$ for some i, then

$$\|x_{y_1,\mu} - x_{y_2,\mu}\| < \frac{1}{n} \tag{6.7}$$

Pick an element $y_i^{(n)}$ in each partition class $A_i^{(n)}$, and define

$$x^{(n)} : \Omega \times \mathcal{Y} \to \mathcal{X}$$

by

$$x^{(n)}(\omega, y) = x_{y_i^{(n)}}(\omega) \quad \text{if } \ y \in A_i^{(n)}$$

and

$$x^{(n)}(\omega, y) = 0 \quad \text{if } \ y \notin K_n\,.$$

Obviously, each $x^{(n)}$ is $P{\times}Q$-measurable. By (6.7)

$$\int \Big[\int \sup_t |x^{(n)}(\omega, y)(t) - x_{y,\mu}(\omega, t)|^2 dP(\omega)\Big] dQ(y) \to 0\,, \tag{6.8}$$

and thus $\{x^{(n)}\}$ is a Cauchy-sequence in $L^2(\Omega \times \mathcal{Y}, \mathcal{X})$ converging to some $P{\times}Q$-measurable function $\hat{x}$. There is a subsequence $\{x^{(n_k)}\}$ such that

$$\sup_{t \le 1} |x^{(n_k)}(\omega, y)(t) - \hat{x}(\omega, y)(t)| \to 0 \tag{6.9}$$

for $P \times Q$-almost all (ω, y). Comparing (6.8) and (6.9), we see that for Q-almost all $y, x_{y,\mu}(\omega, \cdot) = \hat{x}(\omega, y)(\cdot)$ for almost all ω, and that $\hat{x}(\cdot, y)(\cdot)$ hence is a solution of (6.1) for all these y's. Modifying $\hat{x}$ on the remaining y's if necessary, we get the version of $x_{y,\mu}$ required by the proposition.

We shall also need the following result.

6.2 Proposition Assume that for each $y \in \mathcal{Y}$ and $\mu \in \mathcal{R}$, the solution of (6.1) is unique in distribution. Then the distribution of $x_{y,\mu}$ depends continuously on (y, μ).

Proof: Let $(\tilde{y},\tilde{\mu}) \in {}^*\mathcal{Y} \times {}^*\mathcal{R}$ be nearstandard with standard part (y,μ), and let $X_{\tilde{y},\tilde{\mu}}$ be the solution of the nonstandard version

$$X_{\tilde{y},\tilde{\mu}}(t) = \int_0^t {}^*a(X_{\tilde{y},\tilde{\mu}},\tilde{y},s,\tilde{\mu}(s))ds + \int_0^t {}^*b(X_{\tilde{y},\tilde{\mu}},\tilde{y},s,\tilde{\mu}(s))d^*w(s) \tag{6.10}$$

of (6.1). The idea is to show that the standard part ${}^\circ X_{\tilde{y},\tilde{\mu}}$ of $X_{y,\mu}$ is a solution of (6.1), and that the distribution of $X_{\tilde{y},\tilde{\mu}}$ thus is infinitely close to the distribution of $x_{y,\mu}$.

It clearly suffices to show that

$${}^\circ\int_0^t {}^*a(X_{\tilde{y},\tilde{\mu}},\tilde{y},s,\tilde{\mu})ds = \int_0^t a({}^\circ X_{\tilde{y},\tilde{\mu}},y,s,\mu)ds \tag{6.11}$$

for all $t \in [0,1]$, and that there is a Brownian motion $\tilde{w}$ such that

$${}^\circ\int_0^t {}^*b(X_{\tilde{y},\tilde{\mu}},\tilde{y},s,\tilde{\mu})d^*w(s) = \int_0^t b({}^\circ X_{\tilde{y},\tilde{\mu}},y,s,\mu)d\tilde{w}(s) \tag{6.12}$$

for all $t \in [0,1]$.

The first of these equalities is an immediate consequence of the continuity of a and the choice of topology on $\mathcal{R}$. To prove (6.12), note that

$${}^\circ\int_0^t {}^*b^2(X_{\tilde{y},\tilde{\mu}},\tilde{y},s,\tilde{\mu})ds = \int_0^t b^2({}^\circ X_{\tilde{y},\tilde{\mu}},y,s,\mu)ds\,. \tag{6.13}$$

If we assume for a moment that b is invertible and define (in the notation of (6.4))

$$M(t) = {}^\circ\int_0^t {}^*b_{\tilde{\mu}}(X_{\tilde{y},\tilde{\mu}},\tilde{y},s)d^*w(s)\,, \tag{6.14}$$

then

$$\tilde{w}(t) \equiv \int_0^t b_\mu({}^\circ X_{\tilde{y},\tilde{\mu}},y,s)^{-1}dM(s) \tag{6.15}$$

is a Brownian motion (simply because the quadratic variation $[\tilde{w}](t) = t$). Inverting (6.15), we get

$$M(t) = \int_0^t b_\mu({}^\circ X_{\tilde{y},\tilde{\mu}},y,s,\mu)d\tilde{w}(s)\,, \tag{6.16}$$

and comparing this to (6.14), we see that M equals both sides of (6.12).

If b isn't invertible, the proof of Theorem 5.3 in Doob [13] shows that we can still find a Brownian motion $\tilde{w}$ such that (6.16) holds, and thus our argument goes through also in this case.

7 The probabilistic setting

Let us now return to our partially observed stochastic system

$$x_u(t) = \int_0^t a(x_u, y_u, s, u(y_u, s))ds + \int_0^t b(x_u, y_u, s, u(y_u, s)dw_1(s) \tag{7.1}$$

$$y_u(t) = \int_0^t c(x_u, y_u, s)ds + \int_0^t d(y_u, s)dw_2(s) \tag{7.2}$$

We shall assume that the coefficients a, b, c and d are bounded, continuous functions

$$a : \mathcal{X} \times \mathcal{Y} \times [0,1] \times K \rightarrow \mathbf{R}^n$$
$$b : \mathcal{X} \times \mathcal{Y} \times [0,1] \times K \rightarrow S(n)$$
$$c : \mathcal{X} \times \mathcal{Y} \times [0,1] \rightarrow \mathbf{R}^m$$
$$d : \mathcal{Y} \times [0,1] \rightarrow S(m)$$

which are nonanticipating in the sense explained before, i.e. if $x \restriction t = x' \restriction t$ and $y \restriction t = y' \restriction t$, then $a(x, y, t, k) = a(x', y', t, k)$ for all $k \in K$, etc. In addition we have to put some conditions on these coefficients to guarantee the necessary regularity:

7.1 Conditions Assume that

(i) The functions a, b, c, and d are bounded, continuous, and nonanticipating. Moreover, $d(y, s)$ is nonsingular for all y and s, and $d^{-1}(y, s)$ is bounded.

(ii) Given a relaxed response $\mu \in \mathcal{R}$, a Brownian motion w_1, and a function $y \in \mathcal{Y}$, the equation

$$x_{y,\mu}(t) = \int_0^t a(x_{y,\mu}, y, s, \mu(s))ds + \int_0^t b(x_{y,\mu}, y, s, \mu(s))dw_1(s) \tag{7.3}$$

has at most one solution.

(iii) For each Brownian motion w_2, the equation

$$y(t) = \int_0^t d(y, s)dw_2(s) \tag{7.4}$$

has exactly one solution.

7.2 Remark: Since a and b are continuous and bounded, a straightforward extension of Theorem 5.2 in [21] guarantees that there is a Brownian motion w_1 such that (7.3) has a solution for all y and μ. By the "homogeneity" and "universality" results in the same paper, condition (ii) also implies that the solutions of (7.3) are unique in distribution, i.e. solutions of (7.3) with respect to different Brownian motions w_1 induce the same measure on $\mathcal{X}$.

We mentioned in the introduction that the two Brownian motions w_1 and w_2 are supposed to be independent, and it will be convenient to work with a special realization of this independence. First choose a probability space (Ω_1, P_1) carrying a Brownian motion W_1 such that

$$x_{y,\mu}(t) = \int_0^t a(x_{y,\mu}, y, s, \mu(s))ds + \int_0^t b(x_{y,\mu}, y, s, \mu(s))dW_1(s) \tag{7.5}$$

has a solution for all y and μ. Next let $(\Omega_2, P_2) \equiv (\mathcal{Y}, P_2)$ be ordinary Wiener space, let W_2 be the canonical Brownian motion on Ω_2, and let y be the solution of

$$y(t) = \int_0^t d(y, s)dW_2(s)\,. \tag{7.6}$$

Q is the Radon measure induced by y on $\mathcal{Y}$. Finally, let

$$(\Omega, P) = (\Omega_1, P_1) \times (\Omega_2, P_2)$$

be the completed product. In an obvious way, W_1 and W_2 may be thought of as independent Brownian motions on (Ω, P).

We shall choose solutions $x_{y,\mu}$ of (7.5) such that $(\omega_1, y) \mapsto x_{y,\mu}(\omega_1, \cdot)$ is $P_1{\times}Q$-measurable for each $\mu \in \mathcal{R}$ (this is possible by Proposition 6.1). Using Girsanov's formula, we shall now turn $(x_{y,\mu}, y)$ into a solution of (7.1)–(7.2). For each $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ define

$$g(x, y) = \exp\left\{\int_0^t (d^{-2}c)^T(x, y, s)dy(s) - \frac{1}{2}\int_0^t (d^{-1}c)^2(x, y, s)ds\right\} \tag{7.7}$$

where the first term in the exponent is a stochastic integral. If P_x is the measure on $\mathcal{Y}$ given by

$$dP_x(y) = g(x, y)dQ(y)\,, \tag{7.8}$$

then Girsanov's theorem tells us that there is a Brownian motion W_x on (Ω_2, P_x) such that

$$y(t) = \int_0^t c(x, y, s)ds + \int_0^t d(y, s)dW_x(s) \tag{7.9}$$

We are now ready to describe our solution of (7.1)–(7.2). Let u be a random relaxed control, and keep in mind that for each $y \in \mathcal{Y}$, $u(y)$ will be a measure on the space $\mathcal{R}$ of relaxed responses. Define a probability measure P_u on $\Omega \times \mathcal{R}$ by

$$\int f(\omega_1, y, \mu) dP_u(\omega_1, y, \mu) = \\ = \int \left[\int \left[\int f(\omega_1, y, \mu) g(x_{y,\mu}(\omega_1), y) dP(\omega_1) \right] du(y)(\mu) \right] dQ(y) \tag{7.10}$$

for bounded, product measurable f, and let

$$x_u : \Omega \times \mathcal{R} \times [0,1] \to \mathbf{R}^n$$
$$y_u : \Omega \times \mathcal{R} \times [0,1] \to \mathbf{R}^m$$

be defined by

$$x_u(\omega_1, y, \mu, t) = x_{y,\mu}(\omega_1, t)$$
$$y_u(\omega_1, y, \mu, t) = y(t) .$$

It follows from the construction that as processes on $(\Omega \times \mathcal{R}, P_u)$, the pair (x_u, y_u) is in a natural sense a solution of (7.1)–(7.2), and it is this solution we shall work with in the sequel.

8 Standard and nonstandard costs

The *cost* of a random relaxed control u is given by

$$j(u) = E_u(h(x_u)), \tag{8.1}$$

where $h : \mathcal{X} \to \mathbf{R}$ is a (given) bounded, continuous function, and E_u denotes expectation with respect to the measure P_u defined at the end of the preceding section. More explicitly, $j(u)$ is given by

$$j(u) = \int \left[\int \left[\int h(x_{y,\mu}(\omega_1)) g(x_{y,\mu}(\omega_1), y) dP_1(\omega_1) \right] du(y)(\mu) \right] dQ(y) \tag{8.2}$$

It is easy to check that if u happens to be an ordinary control, this expression coincides with the usual definition.

If U is a nonstandard random relaxed control, we can carry through the contruction in the last section in a nonstandard setting. Abusing conventional notation slightly, we shall refer to the resulting processes as ${}^*x_U, {}^*y_U,$ and the corresponding probability measure as ${}^*P_U$. The *nonstandard cost* of U is then defined as

$$J(U) = {}^*E_U({}^*h({}^*x_U)) , \tag{8.3}$$

where ${}^*E_U$ is expectation with respect to ${}^*P_U$.

Our aim in this section is to show that if u is the standard part of U, then $j(u)$ is the standard part of $J(U)$. To do this we must impose one extra condition on our system; we need the Girsanov density $g(x, y)$ introduced in (7.7) to be continuous in the first variable. Due to the stochastic integral in the exponent, this is not entirely obvious, but the following simple lemma shows that it is enough to require that $\frac{\partial}{\partial s}\left(d^{-2}c\right)(x, y, s)$ is bounded and continuous:

8.1 Lemma Assume that

$$k : \mathcal{X} \times \mathcal{Y} \to \mathbf{R}$$

is a bounded, continuous and nonanticipating function whose first derivative $\frac{\partial}{\partial t}k(x, y, t)$ is also bounded and continuous. Then there is a measurable function

$$K : \mathcal{X} \times \mathcal{Y} \to \mathbf{R}$$

which is continuous in the first variable and satisfies

$$K(x, y) = \int_0^1 k(x, y, s)dy(s)$$

Proof: Integration by parts yields

$$\int_0^1 k(x, y, s)dy(s) = k(x, y, 1)y(1) - \int_0^1 y(s)\frac{\partial}{\partial s}k(x, y, s)ds\,,$$

from which the lemma follows immediately.

We can now prove the result announced above:

8.2 Proposition Assume that Condition 7.1 is satisfied and that the Girsanov density $g(x, y)$ is continuous in the first variable. If u is the standard part of the nonstandard random relaxed control U, then

$$j(u) = {}^\circ J(U)\,. \tag{8.4}$$

Proof: By Proposition 6.1 the map

$$\theta : (y, \mu) \mapsto \int h(x_{y,\mu}(\omega_1))g(x_{y,\mu}, y)dP_1(\omega_1)$$

is measurable, and by Proposition 6.2 and the continuity assumptions on h and g, it is continuous in the second variable. Hence by Lemma 2.5

$$\int \theta(y, \mu)d\hat{u}(y, \mu) = {}^\circ\!\int {}^*\theta(y, \mu)d\hat{u}(y, \mu)\,,$$

and according to Lemma 2.2 this is exactly what we want.

9 Optimal random relaxed controls

We are now almost ready to piece everything together and show that under the conditions we have been working with in the last two sections, an optimal random relaxed control always exists and that its cost is equal to the infimum of the costs of all ordinary controls. But there is one small problem we have to deal with first.

9.1. Lemma Let Q be the measure on $\mathcal{Y}$ induced by the solution of

$$\mathcal{Y}(t) = \int_0^t d(y, s)dW(s)\,,$$

(W is a Brownian motion) and let ${}^*Q$ be its nonstandard version. Then any random relaxed control u has an ordinary lifting U with respect to ${}^*Q$.

Proof: We would have liked to appeal to Theorem 5.4, but the problem is that there is no obvious reason why ${}^*Q$ should satisfy the smoothness condition of that theorem. To circumvent this problem, we shall first replace ${}^*Q$ by a measure $\tilde{Q}$ which *is* smooth, then we shall lift u with respect to $\tilde{Q}$, and then show that this lifting can easily be modified into a lifting of u with respect to ${}^*Q$.

We begin by observing that according to Lemma 5.2 (and its proof), u has a random ordinary lifting U (w.r.t. ${}^*Q$) which is constant on intervals of the form $\left[\frac{j}{H^3}, \frac{j+1}{H^3}\right)$ and which only takes values in a hyperfinite set $\hat{K} = \{k_1, k_2, \ldots, k_H\}$. Let $M = \frac{1}{H^7}$ and define

$$\bar{d}(y, t) = {}^*Q([y]_{t,M})^{-1} \int_{[y]_{t,M}} d(z, s)d^*Q(s)$$

to be the average value of d over the equivalence class of y (recall the definitions preceding Definition 5.3). Fix an infinite integer K which is infinitesimal compared to H, and let

$$\bar{y}(t) = \int_0^{t\wedge\tau_K} \bar{d}(\bar{y}, s)ds$$

where $\tau_k = \inf\{t : \|\bar{y}(t)\| \geq K\}$. Finally, let $\tilde{Q}$ be the measure $\bar{y}$ induces on ${}^*\mathcal{Y}$.

Observe that given the equivalence class $[y]_{j/H^3,M}$, the diffusion coefficient $\bar{d}(y,s)$ is independent of y and constant on the interval $\left[\frac{j}{H^3}, \frac{j+1}{H^3}\right)$. Hence an easy calculation with Gaussian integrals is enough to check that $\tilde{Q}$ is smooth (we shall leave this to the reader; observe that the truncation at K is necessary in order not to get in trouble far out at infinity). By Theorem 5.4, u has an ordinary lifting U' with respect to $\tilde{Q}$.

In order to modify U' into a lifting of u w.r.t. ${}^*Q$, we first observe that since both d and $\bar{d}$ are nonsingular, we can establish a one-to-one correspondence between ${}^*\mathcal{Y}$ and itself by

$$y(\omega, t) \mapsto \bar{y}(\omega, t)$$

(there's a slight nuisance caused by the truncation at K which we shall simply overlook). Note that with probability one, the two paths $y(\omega, \cdot)$ and $\bar{y}(\omega, \cdot)$ are infinitely close. We now define U by

$$U(y,t) = U'(\bar{y},t).$$

It is easy to check that since U' is nonanticipating, so is U. Moreover, since y and $\bar{y}$ are infinitely close with probability one, the standard part of U w.r.t. ${}^*Q$ must equal the standard part of U' w.r.t. $\tilde{Q}$; i.e. it equals u. Hence U is an ordinary lifting of u, and the lemma is proved.

We are now ready for the main theorem.

9.2 Theorem Assume that Conditions 7.1 are satisfied and that the Girsanov density $g(x,y)$ in (7.7) is continuous in the first variable. Then there exists a random relaxed control u which is optimal in the following sense

$$\begin{aligned} j(u) &= \inf\{j(v) : v \text{ is a random relaxed control}\} = \\ &= \inf\{j(v) : v \text{ is an ordinary control}\} \end{aligned} \tag{9.1}$$

Proof: If

$$\alpha = \inf\{j(v) : v \text{ is an ordinary control}\},$$

then by transfer

$$\alpha = \inf\{J(V) : V \text{ is a nonstandard ordinary control}\}.$$

Given a random relaxed control v, we can find an ordinary lifting V by Lemma 9.1. By Proposition 8.2

$$j(v) = {}^\circ J(V) \geq \alpha,$$

which shows that the two infimums in (9.1) are equal.

On the other hand, there must be a nonstandard control U such that $J(U) \approx \alpha$. If u is its standard part, then

$$j(u) = {}^\circ J(U) = \alpha,$$

and the theorem is proved.

9.3 Remark: A trivial modification of the proof shows that

$$\{j(v) : v \text{ is a random relaxed control}\}$$

is the closure of

$$\{j(v) : v \text{ is an ordinary control}\}.$$

As an immediate consequence of our construction, we can show that very simple, ordinary controls can bring us arbitrarily close to the minimal cost. Call an ordinary control u *finitary* if there is a finite set $\hat{K} \subset K$ and an integer $M \in \mathbf{N}$ such that for each $y \in \mathcal{Y}$, the path $u(y)(\cdot)$ is constant on intervals of the form $\left[\frac{i}{M}, \frac{i+1}{M}\right)$ and only takes values in $\hat{K}$.

9.4 Corollary $\inf\{j(u) : u \text{ is an ordinary control}\} = \inf\{j(u) : u \text{ is an finitary control}\}$.

Proof: For each infinitely large integer $H \in {}^*\mathbf{N}$, the theory developed above tells us that there is a hyperfinite set $\hat{K} = \{k_1, \ldots, k_H\}$ and a nonstandard ordinary control V_H which takes values in $\hat{K}$, is constant on intervals of the form $\left[\frac{i}{H^3}, \frac{i+1}{H^3}\right)$, and has a cost $J(V_H) \approx \alpha$. In particular, $J(V_H) < \alpha + \epsilon$ for any given $\epsilon > 0$. By the "underspill" principle of nonstandard analysis, there must be a finite H and a corresponding control V_H such that $J(V_H) < \alpha + \epsilon$. We now take v to be the standard part of V_H.

Let us end the paper with a brief and informal discussion of the conditions we have imposed on our system (1.1)–(1.3). There are no "metaphysical" reasons why we have allowed the functions a and b to depend on x_u, y_u, s and u, while c depends on x_u, y_u and s, and d only on y_u and s — we have simply chosen the most general conditions our technical machinery will allow. It is quite possible that we could extend our methods to the case where c also depends on the control u, but it is vital that d does not depend on x_u and u as we needed the measures y_u induced on $\mathcal{Y}$ for different controls u to be mutually absolutely continuous. It is also important that the diffusion coefficient d of the observation process doesn't degenerate too much; if it does, we do not have sufficient inherent randomness to approximate random relaxed controls by ordinary controls. An interesting problem for future research is to construct an example where an optimal relaxed control does not exist (but – of course – where an optimal *random* relaxed control does exist); since there are several existence results for optimal relaxed controls of different kinds of systems in the literature, such an example would probably have to be quite complicated.

References

[1] Albeverio, S., Fenstad, J. E., Høegh-Krohn, R., Lindstrøm, T.: *Nonstandard methods in stochastic analysis and mathematical physics*, Academic Press, New York, 1986.

[2] Bensoussan, A. and Nisio, M.: Nonlinear semigroup arising in the control of diffusions with partial observation, Stochastics and Stochastics Report 30 (1990), 1–45.

[3] Billingsley, P.: *Convergence of probability measures*, Wiley, New York, 1968.

[4] Bismut, J.-M.: Partially observed diffusions and their control, Siam J. Control and Optimization 20 (1982), 302–309.

[5] Borkar, V. S.: Existence of optimal controls for partially observed diffusions, Stochastics 11 (1983), 103–141.

[6] Cohen, R. and Mazziotto, G.: Stochastic continuous control of partially observed systems via impulse control problems, Stochastics and Stochastics Reports 26 (1989), 101–127.

[7] Cutland, N. J.: Nonstandard measure theory and its applications, Bull. London Math. Soc. 15 (1983), 529–589.

[8] Cutland, N. J.: Internal controls and relaxed controls, J. London Math. Soc. 27 (1983), 130–140.

[9] Cutland, N. J.: Optimal controls for partially observed stochastic systems: an infinitesimal approach, Stochastics 15 (1983), 239–257.

[10] Cutland, N. J.: Partially observed stochastic controls based on cumulative readouts of the observations, in *Proc. 4th IFIP Work. Conf. Stochastic Differential Systems, Marseille-Luminy*, Springer-Verlag, Berlin, 1984, 261–269.

[11] Cutland, N. J.: Infinitesimal methods in control theory: Deterministic and stochastic, Acta Appl. Math. 5 (1986), 105–136.

[12] Cutland, N. J. (ed.): *Nonstandard analysis and its applications*, Cambridge University Press, Cambridge, 1988.

[13] Doob, J. L.: *Stochastic processes*, Wiley, New York, 1953.

[14] El Karoui, N., Nguyen, D. H., and Jeanblanc-Picque, M.: Existence of an optimal Markovian filter for the control under partial observations, SIAM J. Control and Optimization 26 (1988), 1025–1061.

[15] Filippov, A. F.: On certain questions in the theory of optimal control, J. Soc. Indust. Appl. Math., Ser. A, Control 1, (1962), 76–84.

[16] Fleming, W. H. and Nisio, M.: On stochastic relaxed control for partially observed diffusions, Nagoya Math. J. 93 (1984), 71–108.

[17] Fleming, W. H. and Pardoux, E.: Optimal control for partially observed diffusions, SIAM J. Control and Optimization 20 (1982), 261–285.

[18] Fremlin, D. H.: Measurable functions and almost continuous functions, Manuscripta Math. 33 (1981), 387–405.

[19] Haussmann, U. G.: On the existence of optimal controls for partially observed diffusions, SIAM J. Control and Optimization 20 (1982), 385–407.

[20] Hijab, O.: Partially observed control of Markov processes I, Stochastics and Stochastics Reports 28 (1989), 123–144.

[21] Keisler, H. J.: An infinitesimal approach to stochastic analysis, Memoirs Amer. Math. Soc. 297 (1984).

[22] Koumoullis, G., and Prikry, K.: Perfect measurable spaces, Ann. Pure Appl. Logic 30 (1986), 219–248.

[23] McShane, E. J.: Relaxed controls and variational problems, SIAM J. Control and Optimization 5 (1967), 438–485.

[24] Stroock, D. W., and Varadhan, S. R. S.: *Multidimensional diffusion processes*, Springer-Verlag, New York, 1979.

[25] Warga, J.: Functions of relaxed controls, SIAM J. Control and Optimization 5 (1967), 628–641.