A quasi-sure approach to the control of non-Markovian stochastic differential equations[*]

Marcel Nutz[†]

Abstract

We study stochastic differential equations (SDEs) whose drift and diffusion coefficients are path-dependent and controlled. We construct a value process on the canonical path space, considered simultaneously under a family of singular measures, rather than the usual family of processes indexed by the controls. This value process is characterized by a second order backward SDE, which can be seen as a non-Markovian analogue of the Hamilton-Jacobi-Bellman partial differential equation. Moreover, our value process yields a generalization of the G-expectation to the context of SDEs.

1 Introduction

We consider a controlled stochastic differential equation (SDE) of the form

$$X_t = x + \int_0^t \mu(r, X, \nu_r)\, dr + \int_0^t \sigma(r, X, \nu_r)\, dW_r, \quad 0 \le t \le T, \tag{1.1}$$

where ν is an adapted control process, W is a Brownian motion and the Lipschitz-continuous coefficients $\mu(r, X, \nu_r)$ and $\sigma(r, X, \nu_r)$ may depend on the past trajectory $\{X_s, 0 \le s \le r\}$ of the solution. Denoting by X^ν the solution corresponding to ν, we are interested in the stochastic optimal control problem

$$V_0 := \sup_\nu E[\xi(X^\nu)], \tag{1.2}$$

where ξ is a given functional. The standard approach to such a non-Markovian control problem (cf. [9, 10]) is to consider for each control ν the associated value process

$$J_t^\nu = \operatorname*{ess\,sup}_{\tilde\nu:\ \tilde\nu = \nu \text{ on } [0,t]} E[\xi(X^{\tilde\nu})|\mathcal{F}_t], \tag{1.3}$$

[†]Department of Mathematics, Columbia University, New York, USA. E-mail: mnutz@math.columbia.edu

where $(\mathcal{F}_t)$ is the given filtration. The dependence on ν reflects the presence of a *forward* component in the optimization problem.

The situation is quite different in Markovian optimal control (cf. [12]), where one uses a single value function which depends on certain state variables but not on a control. This is essential to describe the value function by a differential equation, such as the Hamilton-Jacobi-Bellman PDE, which is the main merit of the dynamic programming approach. It is worth noting that this equation is always *backward* in time. An analogous description for (1.3) via backward SDEs (BSDEs, cf. [20]) is available for certain popular problems such as utility maximization with power or exponential utility functions (e.g., [13, 17]) or drift control (e.g., [10]). However, this relies on a very particular algebraic structure which allows for a separation of J^ν into a backward part independent of ν and a forward part depending on ν.

In this paper, we consider the problem (1.2) on the canonical space by recasting it as

$$V_0 = \sup_\nu E^{P^\nu}[\xi(B)], \tag{1.4}$$

where P^ν is the distribution of X^ν and B is the canonical process, and we describe its dynamic value by a single value process $V = \{V_t(\omega)\}$. Formally, V corresponds to a value function in the Markovian sense if we see the whole trajectory of the controlled system as a state variable. Even though (1.2) has features of coupled forward-backward type, the value process is defined in a purely backward manner: one may say that by constructing V on the whole canonical space, we essentially calculate the value for all possible outcomes of the forward part. An important ingredient in the same vein is that V is defined "quasi-surely" under the family of mutually singular measures $\{P^\nu\}$. Rather than forming a family of processes as in (1.3), the necessary information is stored in a single process which is defined on a "large" part of the probability space; indeed, the process V "seen under P^ν" should be thought of as an analogue of J^ν. Clearly, this is a necessary step to obtain a (second order) backward SDE. We remark that [22] considered the same control problem (1.2) and also made a connection to nonlinear expectations. However, in [22], the value process was considered only under the given probability measure.

We first consider a fairly regular functional ξ and define $V_t(\omega)$ as a conditional version of (1.4). Applying and advancing ideas from [28] and [18], regular conditional probability distributions are used to define $V_t(\omega)$ for *every* ω and prove a pathwise dynamic programming principle (Theorem 3.2). In a second step, we enlarge the class of functionals ξ to an L^1-type space and prove that the value process admits a (quasi-sure) càdlàg modification (Theorem 5.1).

We also show that the value process falls into the class of sublinear expectations studied in [19]. Indeed, if ξ is considered as a random variable on the canonical space, the mapping $\xi \mapsto V_t$ can be seen as a generalization of the G-expectation [23, 24], which, by [7], corresponds to the case $\mu \equiv 0$ and $\sigma(r, X, \nu_r) = \nu_r$, where the SDE (1.1) degenerates to a stochastic integral. Moreover, V_t can be seen as a variant of the random G-expectation [18]; cf. Remark 6.5.

Finally, we characterize V by a second order backward SDE (2BSDE) in the spirit of [19]; cf. Theorem 6.4. The second order is clearly necessary since the Hamilton-Jacobi-Bellman PDE for the Markovian case is fully nonlinear, while ordinary BSDEs correspond to semilinear equations. 2BSDEs were introduced in [5], and in [27] for the non-Markovian case. We refer to [27] for the precise relation between 2BSDEs in the quasi-sure formulation and fully nonlinear parabolic PDEs.

We remark that our approach is quite different from the (backward) stochastic *partial* differential equations studied in [21] for a similar control problem (mainly for un-

controlled volatility) and in [14, 15, 2, 3, 4] for so-called pathwise stochastic control problems. The relation to the path-dependent PDEs, introduced very recently in [25], is yet to be explored.

The remainder of the paper is organized as follows. In Section 2 we detail the controlled SDE and its conditional versions, define the value process for the case when ξ is uniformly continuous and establish its regularity. The pathwise dynamic programming principle is proved in Section 3. In Section 4 we extend the value process to a more general class of functionals ξ and state its quasi-sure representation. The càdlàg modification is constructed in Section 5. In the concluding Section 6, we provide the Hamilton-Jacobi-Bellman 2BSDE and interpret the value process as a variant of the random G-expectation.

2 Construction of the Value Function

In this section, we first introduce the setting and notation. Then, we define the value function $V_t(\omega)$ for a uniformly continuous reward functional ξ and examine the regularity of V_t.

2.1 Notation

We fix a constant $T > 0$ and let $\Omega := C([0, T]; \mathbb{R}^d)$ be the canonical space of continuous paths equipped with the uniform norm $\|\omega\|_T := \sup_{0 \le s \le T} |\omega_s|$, where $|\cdot|$ is the Euclidean norm. We denote by B the canonical process $B_t(\omega) = \omega_t$, by P_0 the Wiener measure, and by $\mathbb{F} = \{\mathcal{F}_t\}_{0 \le t \le T}$ the (raw) filtration generated by B. Unless otherwise stated, probabilistic notions requiring a filtration (such as adaptedness) refer to $\mathbb{F}$.

For any probability measure P on Ω and any $(t, \omega) \in [0, T] \times \Omega$, we can construct the corresponding regular conditional probability distribution P_t^ω; cf. [29, Theorem 1.3.4]. We recall that P_t^ω is a probability kernel on $\mathcal{F}_t \times \mathcal{F}_T$; i.e., P_t^ω is a probability measure on $(\Omega, \mathcal{F}_T)$ for fixed ω and $\omega \mapsto P_t^\omega(A)$ is $\mathcal{F}_t$-measurable for each $A \in \mathcal{F}_T$. Moreover, the expectation under P_t^ω is the conditional expectation under P:

$$E^{P_t^\omega}[\xi] = E^P[\xi | \mathcal{F}_t](\omega) \quad P\text{-a.s.}$$

whenever ξ is $\mathcal{F}_T$-measurable and bounded. Finally, P_t^ω is concentrated on the set of paths that coincide with ω up to t,

$$P_t^\omega \{\omega' \in \Omega : \omega' = \omega \text{ on } [0, t]\} = 1. \tag{2.1}$$

While P_t^ω is not defined uniquely by these properties, we choose and fix one version for each triplet (t, ω, P).

Let $t \in [0, T]$. We denote by $\Omega^t := \{\omega \in C([t, T]; \mathbb{R}^d) : \omega_t = 0\}$ the shifted canonical space of paths starting at the origin. For $\omega \in \Omega$, the shifted path $\omega^t \in \Omega^t$ is defined by $\omega_r^t := \omega_r - \omega_t$ for $t \le r \le T$, so that $\Omega^t = \{\omega^t : \omega \in \Omega\}$. Moreover, we denote by P_0^t the Wiener measure on Ω^t and by $\mathbb{F}^t = \{\mathcal{F}_r^t\}_{t \le r \le T}$ the (raw) filtration generated by B^t, which can be identified with the canonical process on Ω^t.

Given two paths ω and $\tilde{\omega}$, their concatenation at t is the (continuous) path defined by

$$(\omega \otimes_t \tilde{\omega})_r := \omega_r \mathbf{1}_{[0,t)}(r) + (\omega_t + \tilde{\omega}_r^t) \mathbf{1}_{[t,T]}(r), \quad 0 \le r \le T.$$

Given an $\mathcal{F}_T$-measurable random variable ξ on Ω and $\omega \in \Omega$, we define the conditioned random variable $\xi^{t,\omega}$ on Ω by

$$\xi^{t,\omega}(\tilde{\omega}) := \xi(\omega \otimes_t \tilde{\omega}), \quad \tilde{\omega} \in \Omega.$$

Note that $\xi^{t,\omega}(\tilde{\omega}) = \xi^{t,\omega}(\tilde{\omega}^t)$; in particular, $\xi^{t,\omega}$ can also be seen as a random variable on Ω^t. Then $\tilde{\omega} \mapsto \xi^{t,\omega}(\tilde{\omega})$ is $\mathcal{F}_T^t$-measurable and moreover, $\xi^{t,\omega}$ depends only on the restriction of ω to $[0, t]$. We note that for an $\mathbb{F}$-progressively measurable process $\{X_r, \, r \in [s, T]\}$, the conditioned process $\{X_r^{t,\omega}, \, r \in [t, T]\}$ is $\mathbb{F}^t$-progressively measurable. If P is a probability on Ω, the measure $P^{t,\omega}$ on $\mathcal{F}_T^t$ defined by

$$P^{t,\omega}(A) := P_t^\omega(\omega \otimes_t A), \quad A \in \mathcal{F}_T^t, \quad \text{where } \omega \otimes_t A := \{\omega \otimes_t \tilde{\omega} : \tilde{\omega} \in A\},$$

is again a probability by (2.1). We then have

$$E^{P^{t,\omega}}[\xi^{t,\omega}] = E^{P_t^\omega}[\xi] = E^P[\xi|\mathcal{F}_t](\omega) \quad P\text{-a.s.}$$

Analogous notation will be used when ξ is a random variable on Ω^s and $\omega \in \Omega^s$, where $0 \le s \le t \le T$. We denote by $\Omega_t^s := \{\omega|_{[s,t]} : \omega \in \Omega^s\}$ the restriction of Ω^s to $[s, t]$, equipped with $\|\omega\|_{[s,t]} := \sup_{r \in [s,t]} |\omega_r|$. Note that Ω_t^s can be identified with the set $\{\omega \in \Omega^s : \omega_r = \omega_t \text{ for } r \in [t, T]\}$. The meaning of Ω_t is analogous.

2.2 The Controlled SDE

Let U be a nonempty Borel subset of $\mathbb{R}^m$ for some $m \in \mathbb{N}$. We consider two given functions

$$\mu : [0, T] \times \Omega \times U \to \mathbb{R}^d \quad \text{and} \quad \sigma : [0, T] \times \Omega \times U \to \mathbb{R}^{d \times d},$$

the drift and diffusion coefficients, such that $(t, \omega) \mapsto \mu(t, X(\omega), \nu_t(\omega))$ and $(t, \omega) \mapsto \sigma(t, X(\omega), \nu_t(\omega))$ are progressively measurable for any continuous adapted process X and any U-valued progressively measurable process ν. In particular, $\mu(t, \omega, u)$ and $\sigma(t, \omega, u)$ depend only on the past trajectory $\{\omega_r, \, r \in [0, t]\}$, for any $u \in U$. Moreover, we assume that there exists a constant $K > 0$ such that

$$|\mu(t, \omega, u) - \mu(t, \omega', u)| + |\sigma(t, \omega, u) - \sigma(t, \omega', u)| \le K\|\omega - \omega'\|_t \qquad (2.2)$$

for all $(t, \omega, \omega', u) \in [0, T] \times \Omega \times \Omega \times U$. We denote by $\mathcal{U}$ the set of all U-valued progressively measurable processes ν such that

$$\int_0^T |\mu(r, X, \nu_r)| \, dr < \infty \quad \text{and} \quad \int_0^T |\sigma(r, X, \nu_r)|^2 \, dr < \infty \qquad (2.3)$$

hold path-by-path for any continuous adapted process X. Given $\nu \in \mathcal{U}$, the stochastic differential equation

$$X_t = x + \int_0^t \mu(r, X, \nu_r) \, dr + \int_0^t \sigma(r, X, \nu_r) \, dB_r, \quad 0 \le t \le T \quad \text{under } P_0$$

has a P_0-a.s. unique strong solution for any initial condition $x \in \mathbb{R}^d$, which we denote by $X(0, x, \nu)$. We shall denote by

$$\bar{P}(0, x, \nu) := P_0 \circ X(0, x, \nu)^{-1} \qquad (2.4)$$

the distribution of $X(0, x, \nu)$ on Ω and by

$$P(0, x, \nu) := P_0 \circ \left(X(0, x, \nu)^0 \right)^{-1}$$

the distribution of $X(0, x, \nu)^0 \equiv X(0, x, \nu) - x$; i.e., the solution which is translated to start at the origin. Note that $P(0, x, \nu)$ is concentrated on Ω^0 and can therefore be seen as a probability measure on Ω^0.

We shall work under the following nondegeneracy condition.

Assumption 2.1. *Throughout this paper, we assume that*

$$\overline{\mathbb{F}^X}^{P_0} \supseteq \mathbb{F} \quad \textit{for all} \quad X = X(0, x, \nu), \tag{2.5}$$

where $\overline{\mathbb{F}^X}^{P_0}$ is the P_0-augmentation of the filtration generated by X and (x, ν) varies over $\mathbb{R}^d \times \mathcal{U}$.

One can construct situations where Assumption 2.1 fails. For example, if $x = 0$, $\sigma \equiv 1$ and $\mu(r, X, \nu_r) = \nu_r$, then (2.5) fails for a suitable choice of ν; see, e.g., [11]. The following is a positive result which covers many applications.

Remark 2.2. *Let σ be strictly positive definite and assume that $\mu(r, X, \nu_r)$ is a progressively measurable functional of X and $\sigma(r, X, \nu_r)$. Then (2.5) holds true.*

Note that the latter assumption is satisfied in particular when μ is uncontrolled; i.e., $\mu(r, \omega, u) = \mu(r, \omega)$.

Proof. Let $X = X(0, x, \nu)$. As the quadratic variation of X, the process $\int \sigma\sigma^\top(r, X, \nu_r)\, dr$ is adapted to the filtration generated by X. In view of our assumptions, it follows that

$$M := \int \sigma(r, X, \nu_r)\, dB_r = X - x - \int \mu(r, X, \nu_t)\, dr$$

has the same property. Hence $B = \int \sigma(r, X, \nu_r)^{-1}\, dM_r$ is again adapted to the filtration generated by X. $\qquad\square$

Remark 2.3. *For some applications, in particular when the SDE is of geometric form, requiring (2.5) to hold for all $x \in \mathbb{R}^d$ is too strong. One can instead fix the initial condition x throughout the paper, then it suffices to require (2.5) only for that x.*

Next, we introduce for fixed $t \in [0, T]$ an SDE on $[t, T] \times \Omega^t$ induced by μ and σ. Of course, the second argument of μ and σ requires a path on $[0, T]$, so that it is necessary to specify a "history" for the SDE on $[0, t]$. This role is played by an arbitrary path $\eta \in \Omega$. Given η, we define the conditioned coefficients

$$\mu^{t,\eta} : [0, T] \times \Omega^t \times U \to \mathbb{R}^d, \qquad \mu^{t,\eta}(r, \omega, u) := \mu(r, \eta \otimes_t \omega, u),$$
$$\sigma^{t,\eta} : [0, T] \times \Omega^t \times U \to \mathbb{R}^{d \times d}, \qquad \sigma^{t,\eta}(r, \omega, u) := \sigma(r, \eta \otimes_t \omega, u).$$

(More precisely, these functions are defined also when ω is a path not necessarily starting at the origin, but clearly their value at (r, ω, u) depends only on ω^t.) We observe that the Lipschitz condition (2.2) is inherited; indeed,

$$|\mu^{t,\eta}(r, \omega, u) - \mu^{t,\eta}(r, \omega', u)| + |\sigma^{t,\eta}(r, \omega, u) - \sigma^{t,\eta}(r, \omega', u)|$$
$$\leq K\|\eta \otimes_t \omega - \eta \otimes_t \omega'\|_{[t,r]}$$
$$= K\|\omega^t - \omega'^t\|_{[t,r]}$$
$$\leq 2K\|\omega - \omega'\|_{[t,r]}.$$

We denote by $\mathcal{U}^t$ the set of all $\mathbb{F}^t$-progressively measurable, U-valued processes ν such that $\int_t^T |\mu(r, X, \nu_r)|\, dr < \infty$ and $\int_t^T |\sigma(r, X, \nu_r)|^2\, dr < \infty$ for any continuous $\mathbb{F}^t$-adapted process $X = \{X_r, r \in [0, T]\}$. For $\nu \in \mathcal{U}^t$, the SDE

$$X_s = \eta_t + \int_t^s \mu^{t,\eta}(r, X, \nu_r)\, dr + \int_t^s \sigma^{t,\eta}(r, X, \nu_r)\, dB_r^t, \quad t \leq s \leq T \quad \text{under } P_0^t \tag{2.6}$$

has a unique solution $X(t, \eta, \nu)$ on $[t, T]$. Similarly as above, we define

$$P(t, \eta, \nu) := P_0^t \circ \left(X(t, \eta, \nu)^t \right)^{-1}$$

to be the distribution of $X(t, \eta, \nu)^t \equiv X(t, \eta, \nu) - \eta_t$ on Ω^t. Note that this is consistent with the notation $P(0, x, \nu)$ if x is seen as a constant path.

2.3 The Value Function

We can now define the value function for the case when the reward functional ξ is an element of $\mathrm{UC}_b(\Omega)$, the space of bounded uniformly continuous functions on Ω.

Definition 2.4. *Given $t \in [0, T]$ and $\xi \in \mathrm{UC}_b(\Omega)$, we define the value function*

$$V_t(\omega) = V_t(\xi; \omega) = \sup_{\nu \in \mathcal{U}^t} E^{P(t, \omega, \nu)}[\xi^{t, \omega}], \quad (t, \omega) \in [0, T] \times \Omega. \tag{2.7}$$

The function ξ is fixed throughout Sections 2 and 3 and hence often suppressed in the notation. In view of the double dependence on ω in (2.7), the measurability of V_t is not obvious. We have the following regularity result.

Proposition 2.5. *Let $t \in [0, T]$ and $\xi \in \mathrm{UC}_b(\Omega)$. Then $V_t \in \mathrm{UC}_b(\Omega_t)$ and in particular V_t is $\mathcal{F}_t$-measurable. More precisely,*

$$|V_t(\omega) - V_t(\omega')| \leq \rho(\|\omega - \omega'\|_t) \quad \text{for all} \quad \omega, \omega' \in \Omega$$

with a modulus of continuity ρ depending only on ξ, the Lipschitz constant K and the time horizon T.

The first source of regularity for V_t is our assumption that ξ is uniformly continuous; the second one is the Lipschitz property of the SDE. Before stating the proof of the proposition, we examine the latter aspect in detail.

Lemma 2.6. *Let $\psi \in \mathrm{UC}_b(\Omega^t)$. There exists a modulus of continuity $\rho_{K, T, \psi}$, depending only on K, T and the minimal modulus of continuity of ψ, such that*

$$\left| E^{P(t, \omega, \nu)}[\psi] - E^{P(t, \omega', \nu)}[\psi] \right| \leq \rho_{K, T, \psi}\left(\|\omega - \omega'\|_t \right)$$

for all $t \in [0, T]$, $\nu \in \mathcal{U}^t$ and $\omega, \omega' \in \Omega$.

Proof. We set $E[\,\cdot\,] := E^{P_0^t}[\,\cdot\,]$ to alleviate the notation. Let $\omega, \bar{\omega} \in \Omega$, set $X := X(t, \omega, \nu)$ and $\bar{X} := X(t, \bar{\omega}, \nu)$, and recall that $X^t = X - X_t = X - \omega_t$ and similarly $\bar{X}^t = \bar{X} - \bar{\omega}_t$.

(i) We begin with a standard SDE estimate. Let $X^t = M^t + A^t$ and $\bar{X}^t = \bar{M}^t + \bar{A}^t$ be the semimartingale decompositions and $t \leq \tau \leq T$ be a stopping time such that $M^t, \bar{M}^t, A^t, \bar{A}^t$ are bounded on $[t, \tau]$. Then Itô's formula and the Lipschitz property (2.2) of σ yield that

$$E[|M_\tau^t - \bar{M}_\tau^t|^2] \leq E \int_t^\tau |\sigma(r, \omega \otimes_t X, \nu_r) - \sigma(r, \bar{\omega} \otimes_t \bar{X}, \nu_r)|^2 \, dr$$

$$\leq K^2 E \int_t^\tau \|\omega \otimes_t X - \bar{\omega} \otimes_t \bar{X}\|_r^2 \, dr$$

$$\leq K^2 E \int_t^\tau \left(\|\omega - \bar{\omega}\|_t + \|X^t - \bar{X}^t\|_{[t, r]} \right)^2 \, dr$$

$$\leq 2K^2 T \|\omega - \bar{\omega}\|_t^2 + 2K^2 \int_t^T E\left[\|X^t - \bar{X}^t\|_{[t, r \wedge \tau]}^2 \right] dr.$$

Hence, Doob's maximal inequality implies that

$$
E\big[\|M^t - \bar{M}^t\|_{[t,\tau]}^2\big] \leq 4E[|M_\tau^t - \bar{M}_\tau^t|^2]
$$
$$
\leq 8K^2 T\|\omega - \bar{\omega}\|_t^2 + 8K^2 \int_t^T E\big[\|X^t - \bar{X}^t\|_{[t, r \wedge \tau]}^2\big]\, dr.
$$

Moreover, using the Lipschitz property (2.2) of μ, we also have that

$$
|A_s^t - \bar{A}_s^t| \leq \int_t^s |\mu(r, \omega \otimes_t X, \nu_r) - \mu(r, \bar{\omega} \otimes_t \bar{X}, \nu_r)|\, dr
$$
$$
\leq K \int_t^s \big(\|\omega - \bar{\omega}\|_t + \|X^t - \bar{X}^t\|_{[t,r]}\big)\, dr
$$

for all $t \leq s \leq T$ and then Jensen's inequality yields that

$$
E\big[\|A^t - \bar{A}^t\|_{[t,\tau]}^2\big] \leq 2K^2 T^2 \|\omega - \bar{\omega}\|_t^2 + 2K^2 T \int_t^T E\big[\|X^t - \bar{X}^t\|_{[t, r \wedge \tau]}^2\big]\, dr.
$$

Hence, we have shown that

$$
E\big[\|X^t - \bar{X}^t\|_{[t,\tau]}^2\big] \leq C_0 \|\omega - \bar{\omega}\|_t^2 + C_0 \int_t^T E\big[\|X^t - \bar{X}^t\|_{[t, r \wedge \tau]}^2\big]\, dr,
$$

where C_0 depends only on K and T, and we conclude by Gronwall's lemma that

$$
E\big[\|X^t - \bar{X}^t\|_{[t,\tau]}^2\big] \leq C\|\omega - \bar{\omega}\|_t^2, \quad C := C_0 e^{C_0 T}.
$$

By the continuity of their sample paths, there exists a localizing sequence $(\tau_n)_{n \geq 1}$ of stopping times such that $M^t, \bar{M}^t, A^t, \bar{A}^t$ are bounded on $[t, \tau_n]$ for each n. Therefore, monotone convergence and the previous inequality yield that

$$
E\big[\|X^t - \bar{X}^t\|_{[t,T]}^2\big] \leq C\|\omega - \bar{\omega}\|_t^2. \tag{2.8}
$$

(ii) Let $\tilde{\rho}$ be the minimal (nondecreasing) modulus of continuity for ψ,

$$
\tilde{\rho}(z) := \sup \big\{|\psi(\tilde{\omega}) - \psi(\tilde{\omega}')| : \; \tilde{\omega}, \tilde{\omega}' \in \Omega^t, \|\tilde{\omega} - \tilde{\omega}'\|_{[t,T]} \leq z\big\},
$$

and let ρ be the concave hull of $\tilde{\rho}$. Then ρ is a bounded continuous function satisfying $\rho(0) = 0$ and $\rho \geq \tilde{\rho}$. Let $P := P(t, \omega, \nu)$ and $\bar{P} := P(t, \bar{\omega}, \nu)$, then P and $\bar{P}$ are the distributions of X^t and $\bar{X}^t$, respectively; therefore,

$$
\big|E^P[\psi] - E^{\bar{P}}[\psi]\big| = \big|E[\psi(X^t) - \psi(\bar{X}^t)]\big| \leq E\big[\rho\big(\|X^t - \bar{X}^t\|_{[t,T]}\big)\big]. \tag{2.9}
$$

Moreover, Jensen's inequality and (2.8) yield that

$$
E\big[\rho\big(\|X^t - \bar{X}^t\|_{[t,T]}\big)\big] \leq \rho\big(E\big[\|X^t - \bar{X}^t\|_{[t,T]}\big]\big)
$$
$$
\leq \rho\big(E\big[\|X^t - \bar{X}^t\|_{[t,T]}^2\big]^{1/2}\big)
$$
$$
\leq \rho\big(\sqrt{C}\|\omega - \bar{\omega}\|_t\big)
$$

for every n. In view of (2.9), we have $|E^P[\psi] - E^{\bar{P}}[\psi]| \leq \rho(\sqrt{C}\|\omega - \bar{\omega}\|_t)$; i.e., the result holds for $\rho_{K,T,\psi}(z) := \rho(\sqrt{C}z)$. $\qquad \square$

After these preparations, we can prove the continuity of V_t.

Proof of Proposition 2.5. To disentangle the double dependence on ω in (2.7), we first consider the function

$$(\eta, \omega) \mapsto E^{P(t,\eta,\nu)}[\xi^{t,\omega}], \quad (\eta, \omega) \in \Omega \times \Omega.$$

Since $\xi \in \mathrm{UC}_b(\Omega)$, there exists a modulus of continuity $\rho^{(\xi)}$ for ξ; i.e.,

$$|\xi(\omega) - \xi(\omega')| \le \rho^{(\xi)}(\|\omega - \omega'\|_T), \quad \omega, \omega' \in \Omega.$$

Therefore, we have for all $\tilde{\omega} \in \Omega^t$ that

$$\begin{aligned}
|\xi^{t,\omega}(\tilde{\omega}) - \xi^{t,\omega'}(\tilde{\omega})| &= |\xi(\omega \otimes_t \tilde{\omega}) - \xi(\omega' \otimes_t \tilde{\omega})| \\
&\le \rho^{(\xi)}(\|\omega \otimes_t \tilde{\omega} - \omega' \otimes_t \tilde{\omega}\|_T) \\
&= \rho^{(\xi)}(\|\omega - \omega'\|_t).
\end{aligned} \tag{2.10}$$

For fixed $\eta \in \Omega$, it follows that

$$\left| E^{P(t,\eta,\nu)}[\xi^{t,\omega}] - E^{P(t,\eta,\nu)}[\xi^{t,\omega'}] \right| \le \rho^{(\xi)}(\|\omega - \omega'\|_t). \tag{2.11}$$

Fix $\omega \in \Omega$ and let $\psi_\omega := \xi^{t,\omega}$. Then $\rho^{(\xi)}$ yields a modulus of continuity for ψ_ω; in particular, this modulus of continuity is uniform in ω. Thus Lemma 2.6 implies that the mapping $\eta \mapsto E^{P(t,\eta,\nu)}[\psi_\omega]$ admits a modulus of continuity $\rho_{T,K,\xi}$ depending only on T, K, ξ. In view of (2.11), we conclude that

$$\left| E^{P(t,\eta,\nu)}[\xi^{t,\omega}] - E^{P(t,\eta',\nu)}[\xi^{t,\omega'}] \right| \le \rho^{(\xi)}(\|\omega - \omega'\|_t) + \rho_{T,K,\xi}(\|\eta - \eta'\|_t)$$

for all $\eta, \eta', \omega, \omega' \in \Omega$ and in particular that

$$\left| E^{P(t,\omega,\nu)}[\xi^{t,\omega}] - E^{P(t,\omega',\nu)}[\xi^{t,\omega'}] \right| \le \rho(\|\omega - \omega'\|_t), \quad \rho := \rho^{(\xi)} + \rho_{T,K,\xi} \tag{2.12}$$

for all $\omega, \omega' \in \Omega$. Passing to the supremum over $\nu \in \mathcal{U}^t$, we obtain that $|V_t(\omega) - V_t(\omega')| \le \rho(\|\omega - \omega'\|_t)$ for all $\omega, \omega' \in \Omega$, which was the claim. $\qquad\square$

3 Pathwise Dynamic Programming

In this section, we provide a pathwise dynamic programming principle which is fundamental for the subsequent sections. As we are working in the weak formulation (1.4), the arguments used here are similar to, e.g., [26], while [22] gives a related construction in the strong formulation (i.e., working only under P_0).

We assume in this section that the following conditional version of Assumption 2.1 holds true; however, we shall see later (Lemma 4.4) that this extended assumption holds automatically outside certain nullsets.

Assumption 3.1. *Throughout Section 3, we assume that*

$$\overline{\mathbb{F}^X}^{P_0^t} \supseteq \mathbb{F}^t \quad \text{for} \quad X := X(t, \eta, \nu), \tag{3.1}$$

for all $(t, \eta, \nu) \in [0, T] \times \Omega \times \mathcal{U}^t$.

The main result of this section is the following dynamic programming principle. We shall also provide more general, quasi-sure versions of this result later (the final form being Theorem 5.2).

Theorem 3.2. *Let $0 \leq s \leq t \leq T$, $\xi \in \mathrm{UC}_b(\Omega)$ and set $V_r(\cdot) = V_r(\xi; \cdot)$. Then*

$$V_s(\omega) = \sup_{\nu \in \mathcal{U}^s} E^{P(s,\omega,\nu)}\big[(V_t)^{s,\omega}\big] \quad \text{for all} \quad \omega \in \Omega. \tag{3.2}$$

We remark that in view of Proposition 2.5, we may see $\xi \mapsto V_r(\xi; \cdot)$ as a mapping $\mathrm{UC}_b(\Omega) \to \mathrm{UC}_b(\Omega)$ and recast (3.2) as the semigroup property

$$V_s = V_s \circ V_t \quad \text{on} \quad \mathrm{UC}_b(\Omega) \quad \text{for all} \quad 0 \leq s \leq t \leq T. \tag{3.3}$$

Some auxiliary results are needed for the proof of Theorem 3.2, which is stated at the end of this section. We start with the (well known) observation that conditioning the solution of an SDE yields the solution of a suitably conditioned SDE.

Lemma 3.3. *Let $0 \leq s \leq t \leq T$, $\nu \in \mathcal{U}^s$ and $\bar{\omega} \in \Omega$. If $\bar{X} := X(s, \bar{\omega}, \nu)$, then*

$$\bar{X}^{t,\omega} = X\big(t, \bar{\omega} \otimes_s \bar{X}(\omega), \nu^{t,\omega}\big) \quad P_0^t\text{-a.s.}$$

for all $\omega \in \Omega^s$.

Proof. Let $\omega \in \Omega^s$. Using the definition and the flow property of $\bar{X}$, we have

$$\bar{X}_r = \bar{\omega}_s + \int_s^r \mu^{s,\bar{\omega}}(u, \bar{X}, \nu_u)\, du + \int_s^r \sigma^{s,\bar{\omega}}(u, \bar{X}, \nu_u)\, dB_u^s$$

$$= \bar{X}_t + \int_t^r \mu(u, \bar{\omega} \otimes_s \bar{X}, \nu_u)\, du + \int_t^r \sigma(u, \bar{\omega} \otimes_s \bar{X}, \nu_u)\, dB_u^s \quad P_0^s\text{-a.s.}$$

for all $r \in [t, T]$. Hence, using that $(P_0^s)^{t,\omega} = P_0^t$ by the P_0^s-independence of the increments of B^s,

$$\bar{X}_r^{t,\omega} = \bar{X}_t^{t,\omega} + \int_t^r \mu(u, \bar{\omega} \otimes_s \bar{X}^{t,\omega}, \nu_u^{t,\omega})\, du + \int_t^r \sigma(u, \bar{\omega} \otimes_s \bar{X}^{t,\omega}, \nu_u^{t,\omega})\, dB_u^t \quad P_0^t\text{-a.s.} \tag{3.4}$$

Since $\bar{X}$ is adapted, we have $\bar{X}^{t,\omega}(\cdot) = \bar{X}(\omega \otimes_t \cdot) = \bar{X}(\omega)$ on $[s, t]$ and in particular

$$\bar{\omega} \otimes_s \bar{X}^{t,\omega} = \bar{\omega} \otimes_s \bar{X}(\omega) \otimes_t \bar{X}^{t,\omega} = \eta \otimes_t \bar{X}^{t,\omega}, \quad \text{for} \quad \eta := \bar{\omega} \otimes_s \bar{X}(\omega).$$

Therefore, recalling that $\bar{X}_s = \bar{\omega}_s$, (3.4) can be stated as

$$\bar{X}_r^{t,\omega} = \eta_t + \int_t^r \mu^{t,\eta}(u, \bar{X}^{t,\omega}, \nu^{t,\omega})\, du + \int_t^r \sigma^{t,\eta}(u, \bar{X}^{t,\omega}, \nu^{t,\omega})\, dB_u^t \quad P_0^t\text{-a.s.};$$

i.e., $\bar{X}^{t,\omega}$ solves the SDE (2.6) for the parameters $(t, \eta, \nu^{t,\omega})$. Now the result follows by the uniqueness of the solution to this SDE. $\qquad\square$

Given $t \in [0, T]$ and $\omega \in \Omega$, we define

$$\mathcal{P}(t, \omega) = \big\{ P(t, \omega, \nu) : \nu \in \mathcal{U}^t \big\}. \tag{3.5}$$

These sets have the following invariance property.

Lemma 3.4. *Let $0 \leq s \leq t \leq T$ and $\bar{\omega} \in \Omega$. If $P \in \mathcal{P}(s, \bar{\omega})$, then*

$$P^{t,\omega} \in \mathcal{P}(t, \bar{\omega} \otimes_s \omega) \quad \text{for} \quad P\text{-a.e. } \omega \in \Omega^s.$$

Proof. Since $P \in \mathcal{P}(s, \bar{\omega})$, we have $P = P(s, \bar{\omega}, \nu)$ for some $\nu \in \mathcal{U}^s$; that is, setting

$X := X(s, \bar{\omega}, \nu)$, P is the distribution of

$$X^s = \int_s^\cdot \mu(r, \bar{\omega} \otimes_s X, \nu_r)\, dr + \int_s^\cdot \sigma(r, \bar{\omega} \otimes_s X, \nu_r)\, dB_r^s \quad \text{under} \quad P_0^s.$$

We set $\hat{\mu}_r := \mu(r, \bar{\omega} \otimes_s X, \nu_r)$ and $\hat{\sigma}_r := \sigma(r, \bar{\omega} \otimes_s X, \nu_r)$ and see the above as the integral $\int_s^\cdot \hat{\mu}_r\, dr + \int_s^\cdot \hat{\sigma}_r\, dB_r^s$ rather than an SDE. As in [26, Lemma 2.2], the nondegeneracy assumption (3.1) implies the existence of a progressively measurable transformation $\beta_\nu : \Omega^s \to \Omega^s$ (depending on $s, \bar{\omega}, \nu$) such that

$$\beta_\nu(X^s) = B^s \quad P_0^s\text{-a.s.} \tag{3.6}$$

Furthermore, a rather tedious calculation as in the proof of [26, Lemma 4.1] shows that

$$P^{t,\omega} = P_0^t \circ \left(\int_s^\cdot \hat{\mu}_r^{t,\beta_\nu(\omega)}\, dr + \int_s^\cdot \hat{\sigma}_r^{t,\beta_\nu(\omega)}\, dB_r^t \right)^{-1} \quad \text{for} \quad P\text{-a.e. } \omega \in \Omega^s.$$

Note that, abbreviating $\check{\omega} := \beta_\nu(\omega)$, we have

$$\hat{\mu}_r^{t,\beta_\nu(\omega)} = \mu\big(r, \bar{\omega} \otimes_s X^{t,\check{\omega}}, \nu_r^{t,\check{\omega}}\big) = \mu\big(r, \bar{\omega} \otimes_s X(\check{\omega}) \otimes_t X^{t,\check{\omega}}, \nu_r^{t,\check{\omega}}\big)$$

and similarly for $\hat{\sigma}^{t,\beta_\nu(\omega)}$. Hence, we deduce by Lemma 3.3 that

$$P^{t,\omega} = P(t, \bar{\omega} \otimes_s X(\check{\omega}), \nu^{t,\check{\omega}}) \quad \text{for} \quad P\text{-a.e. } \omega \in \Omega^s.$$

In view of (3.6), we have

$$X^s(\beta_\nu(B^s)) = B^s \quad P\text{-a.s.;} \tag{3.7}$$

i.e., $X(\check{\omega})^s = X^s(\check{\omega}) = \omega$ for P-a.e. $\omega \in \Omega^s$, and we conclude that

$$P^{t,\omega} = P(t, \bar{\omega} \otimes_s \omega, \nu^{t,\check{\omega}}) \quad \text{for} \quad P\text{-a.e. } \omega \in \Omega^s. \tag{3.8}$$

In particular, $P^{t,\omega} \in \mathcal{P}(t, \bar{\omega} \otimes_s \omega)$. $\qquad\qquad\square$

Lemma 3.5 (Pasting). *Let $0 \leq s \leq t \leq T$, $\bar{\omega} \in \Omega$, $\nu \in \mathcal{U}^s$ and set $P := P(s, \bar{\omega}, \nu)$, $X := X(s, \bar{\omega}, \nu)$. Let $(E^i)_{0 \leq i \leq N}$ be a finite $\mathcal{F}_t^s$-measurable partition of Ω^s, $\nu^i \in \mathcal{U}^t$ for $1 \leq i \leq N$ and define $\bar{\nu} \in \mathcal{U}^s$ by*

$$\bar{\nu}(\omega) := \mathbf{1}_{[s,t)}\nu(\omega) + \mathbf{1}_{[t,T]}\left[\nu(\omega)\mathbf{1}_{E^0}(X(\omega)^s) + \sum_{i=1}^N \nu^i(\omega^t)\mathbf{1}_{E^i}(X(\omega)^s) \right].$$

Then $\bar{P} := P(s, \bar{\omega}, \bar{\nu})$ satisfies $\bar{P} = P$ on $\mathcal{F}_t^s$ and

$$\bar{P}^{t,\omega} = P(t, \bar{\omega} \otimes_s \omega, \nu^i) \quad \text{for} \quad P\text{-a.e. } \omega \in E^i, \quad 1 \leq i \leq N.$$

Proof. As $\bar{\nu} = \nu$ on $[s,t)$, we have $X(s, \bar{\omega}, \bar{\nu}) = X$ on $[s,t]$ and in particular $\bar{P} = P$ on $\mathcal{F}_t^s$. Let $1 \leq i \leq N$, we show that $\bar{P}^{t,\omega} = P(t, \bar{\omega} \otimes_s \omega, \nu^i)$ for P-a.e. $\omega \in E^i$. Recall from (3.8) that

$$\bar{P}^{t,\omega} = P(t, \bar{\omega} \otimes_s \omega, \bar{\nu}^{t,\beta_{\bar{\nu}}(\omega)}) \quad \text{for} \quad \bar{P}\text{-a.e. } \omega \in \Omega^s,$$

where $\beta_{\bar{\nu}}$ is defined as in (3.6). Since both sides of this equality depend only on the restriction of ω to $[s,t]$ and $X(s, \bar{\omega}, \bar{\nu}) = X$ on $[s,t]$, we also have that

$$\bar{P}^{t,\omega} = P(t, \bar{\omega} \otimes_s \omega, \bar{\nu}^{t,\check{\omega}}) \quad \text{for} \quad P\text{-a.e. } \omega \in \Omega^s, \tag{3.9}$$

where $\breve{\omega} = \beta_\nu(\omega)$ is defined as below (3.6). Note that by (3.7), $\omega \in E^i$ implies $X(\breve{\omega})^s \in E^i$ under P. (More precisely, if $A \subseteq \Omega^s$ is a set such that $A \subseteq E^i$ P-a.s., then $\{X(\breve{\omega})^s : \omega \in A\} \subseteq E^i$ P-a.s.) In fact, since X is adapted and $E^i \in \mathcal{F}_t^s$, we even have that

$$X(\breve{\omega} \otimes_t \tilde{\omega})^s \in E^i \quad \text{for all} \quad \tilde{\omega} \in \Omega^t, \quad \text{for} \quad P\text{-a.e.} \ \omega \in E^i.$$

By the definition of $\bar{\nu}$, we conclude that

$$\bar{\nu}^{t,\breve{\omega}}(\tilde{\omega}) = \bar{\nu}(\breve{\omega} \otimes_t \tilde{\omega}) = \nu^i((\breve{\omega} \otimes_t \tilde{\omega})^t) = \nu^i(\tilde{\omega}), \quad \tilde{\omega} \in \Omega^t, \quad \text{for} \quad P\text{-a.e.} \ \omega \in E^i.$$

In view of (3.9), this yields the claim. $\qquad \square$

Remark 3.6. *In [26] and [18], it was possible to use a pasting of measures as follows: in the notation of Lemma 3.5, it was possible to specify measures P^i on Ω^t, corresponding to certain admissible controls, and use the pasting $\hat{P}(A) := P(A \cap E^0) + \sum_{i=1}^N E^P[P^i(A^{t,\omega})\mathbf{1}_{E^i}(\omega)]$ to obtain a measure in Ω^s which again corresponded to some admissible control and satisfied*

$$\hat{P}^{t,\omega} = P^i \quad \text{for} \quad \hat{P}\text{-a.e. } \omega \in E^i, \tag{3.10}$$

which was then used in the proof of the dynamic programming principle.

This is not possible in our SDE-driven setting. Indeed, suppose that $\hat{P}$ is of the form $P(s, \bar{\omega}, \nu)$ for some $\bar{\omega}$ and ν, then we see from (3.8) that, when σ is general, $\hat{P}^{t,\omega}$ will depend explicitly on ω, which contradicts (3.10). Therefore, the subsequent proof uses an argument where (3.10) holds only at one specific $\omega \in E^i$; on the rest of E^i, we confine ourselves to controlling the error.

We can now show the dynamic programming principle. Apart from the difference remarked above, the basic pattern of the proof is the same as in [26, Proposition 4.7].

Proof of Theorem 3.2. Using the notation (3.5), our claim (3.2) can be stated as

$$\sup_{P \in \mathcal{P}(s,\bar{\omega})} E^P[\xi^{s,\bar{\omega}}] = \sup_{P \in \mathcal{P}(s,\bar{\omega})} E^P[(V_t)^{s,\bar{\omega}}] \quad \text{for all} \quad \bar{\omega} \in \Omega. \tag{3.11}$$

(i) We first show the inequality "$\leq$" in (3.11). Fix $\bar{\omega} \in \Omega$ and $P \in \mathcal{P}(s,\bar{\omega})$. Lemma 3.4 shows that $P^{t,\omega} \in \mathcal{P}(t, \bar{\omega} \otimes_s \omega)$ for P-a.e. $\omega \in \Omega^s$ and hence that

$$\begin{aligned}
E^{P^{t,\omega}}[(\xi^{s,\bar{\omega}})^{t,\omega}] &= E^{P^{t,\omega}}[\xi^{t,\bar{\omega}\otimes_s\omega}] \\
&\leq \sup_{P' \in \mathcal{P}(t,\bar{\omega}\otimes_s\omega)} E^{P'}[\xi^{t,\bar{\omega}\otimes_s\omega}] \\
&= V_t(\bar{\omega} \otimes_s \omega) \\
&= V_t^{s,\bar{\omega}}(\omega) \quad \text{for } P\text{-a.e. } \omega \in \Omega^s.
\end{aligned}$$

Since V_t is measurable by Proposition 2.5, we can take $P(d\omega)$-expectations on both sides to obtain that

$$E^P[\xi^{s,\bar{\omega}}] = E^P\Big[E^{P^{t,\omega}}[(\xi^{s,\bar{\omega}})^{t,\omega}]\Big] \leq E^P[V_t^{s,\bar{\omega}}].$$

We take the supremum over $P \in \mathcal{P}(s,\bar{\omega})$ on both sides and obtain the claim.

(ii) We now show the inequality "$\geq$" in (3.11). Fix $\bar{\omega} \in \Omega$, $\nu \in \mathcal{U}^s$ and let $P = P(s,\bar{\omega},\nu)$. We fix $\varepsilon > 0$ and construct a countable cover of the state space as follows.

Let $\hat{\omega} \in \Omega^s$. By the definition of $V_t(\bar{\omega} \otimes_s \hat{\omega})$, there exists a control $\nu^{(\hat{\omega})} \in \mathcal{U}^t$ such that

$P^{(\hat{\omega})} := P(t, \bar{\omega} \otimes_s \hat{\omega}, \nu^{(\hat{\omega})})$ satisfies

$$V_t(\bar{\omega} \otimes_s \hat{\omega}) \leq E^{P^{(\hat{\omega})}}[\xi^{t,\bar{\omega} \otimes_s \hat{\omega}}] + \varepsilon. \tag{3.12}$$

Let $B(\varepsilon, \hat{\omega}) \subseteq \Omega^s$ denote the open $\|\cdot\|_{[s,t]}$-ball of radius ε around $\hat{\omega}$. Since $(\Omega^s, \|\cdot\|_{[s,t]})$ is a separable (quasi-)metric space and therefore Lindelöf, there exists a sequence $(\hat{\omega}^i)_{i \geq 1}$ in Ω^s such that the balls $B^i := B(\varepsilon, \hat{\omega}^i)$ form a cover of Ω^s. As an $\|\cdot\|_{[s,t]}$-open set, each B^i is $\mathcal{F}_t^s$-measurable and hence

$$E^1 := B^1, \quad E^{i+1} := B^{i+1} \setminus (E^1 \cup \cdots \cup E^i), \quad i \geq 1$$

defines a partition $(E^i)_{i \geq 1}$ of Ω^s. Replacing E^i by

$$\left(E^i \cup \{\hat{\omega}^i\}\right) \setminus \{\hat{\omega}^j : j \geq 1, j \neq i\}$$

if necessary, we may assume that $\hat{\omega}^i \in E^i$ for $i \geq 1$. To alleviate the notation, we set $\nu^i := \nu^{(\hat{\omega}^i)}$ and $P^i := P(t, \bar{\omega} \otimes_s \hat{\omega}^i, \nu^i)$.

Next, we paste the controls ν^i. Fix $N \in \mathbb{N}$ and let $A_N := E^1 \cup \cdots \cup E^N$, then $\{A_N^c, E^1, \ldots, E^N\}$ is a finite partition of Ω^s. Let $X := X(s, \bar{\omega}, \nu)$, define

$$\bar{\nu}(\omega) := \mathbf{1}_{[s,t)}\nu(\omega) + \mathbf{1}_{[t,T]}\left[\nu(\omega)\mathbf{1}_{A_N^c}(X(\omega)^s) + \sum_{i=1}^{N} \nu^i(\omega^t)\mathbf{1}_{E^i}(X(\omega)^s)\right]$$

and let $\bar{P} := P(s, \bar{\omega}, \bar{\nu})$. By Lemma 3.5, we have $\bar{P} = P$ on $\mathcal{F}_t^s$ and $\bar{P}^{t,\omega} = P(t, \bar{\omega} \otimes_s \omega, \nu^i)$ for all $\omega \in \tilde{E}^i$, for some subset $\tilde{E}^i \subseteq E^i$ of full measure P. Let us assume for the moment that

$$\hat{\omega}^i \in \tilde{E}^i \quad \text{for} \quad 1 \leq i \leq N, \tag{3.13}$$

then we may conclude that

$$\bar{P}^{t,\hat{\omega}^i} = P^i \quad \text{for} \quad 1 \leq i \leq N. \tag{3.14}$$

Recall from Proposition 2.5 that V_t admits a modulus of continuity $\rho^{(V_t)}$. Moreover, we obtain similarly as in (2.10) that there exists a modulus of continuity $\rho^{(\xi)}$ such that

$$|\xi^{t,\bar{\omega} \otimes_s \omega} - \xi^{t,\bar{\omega} \otimes_s \omega'}| \leq \rho^{(\xi)}(\|\omega - \omega'\|_{[s,t]}).$$

Let $\omega \in E^i \subseteq \Omega^s$ for some $1 \leq i \leq N$, then $\|\omega - \hat{\omega}^i\|_{[s,t]} < \varepsilon$. Together with (3.12) and (3.14), we obtain that

$$\begin{aligned} V_t^{s,\bar{\omega}}(\omega) &\leq V_t^{s,\bar{\omega}}(\hat{\omega}^i) + \rho^{(V_t)}(\varepsilon) \\ &\leq E^{P^i}[\xi^{t,\bar{\omega} \otimes_s \hat{\omega}^i}] + \varepsilon + \rho^{(V_t)}(\varepsilon) \\ &= E^{\bar{P}^{t,\hat{\omega}^i}}[\xi^{t,\bar{\omega} \otimes_s \hat{\omega}^i}] + \varepsilon + \rho^{(V_t)}(\varepsilon). \end{aligned} \tag{3.15}$$

Recall from (2.12) that the mapping

$$\omega' \mapsto E^{P(t,\omega',\nu^i)}[\xi^{t,\omega'}]$$

is uniformly continuous with a modulus $\tilde{\rho}$ independent of i and N. Since $\omega \in E^i$, it

follows that

$$
\begin{aligned}
E^{\bar{P}^{t,\hat{\omega}^i}}&[\xi^{t,\bar{\omega}\otimes_s\hat{\omega}^i}] - E^{\bar{P}^{t,\omega}}[\xi^{t,\bar{\omega}\otimes_s\omega}] \\
&= E^{P(t,\bar{\omega}\otimes_s\hat{\omega}^i,\nu^i)}[\xi^{t,\bar{\omega}\otimes_s\hat{\omega}^i}] - E^{P(t,\bar{\omega}\otimes_s\omega,\nu^i)}[\xi^{t,\bar{\omega}\otimes_s\omega}] \\
&\leq \tilde{\rho}(\varepsilon) \quad \text{for} \quad \bar{P}\text{-a.e. } \omega \in E^i.
\end{aligned} \tag{3.16}
$$

Setting $\rho(\varepsilon) := \tilde{\rho}(\varepsilon) + \varepsilon + \rho^{(V_t)}(\varepsilon)$ and noting that

$$
E^{\bar{P}^{t,\omega}}[\xi^{t,\bar{\omega}\otimes_s\omega}] = E^{\bar{P}^{t,\omega}}[(\xi^{s,\bar{\omega}})^{t,\omega}] = E^{\bar{P}}[\xi^{s,\bar{\omega}}|\mathcal{F}_t^s](\omega),
$$

the inequalities (3.15) and (3.16) imply that

$$
V_t^{s,\bar{\omega}}(\omega) \leq E^{\bar{P}}\big[\xi^{s,\bar{\omega}}\big|\mathcal{F}_t^s\big](\omega) + \rho(\varepsilon) \tag{3.17}
$$

for $\bar{P}$-a.e. (and thus P-a.e.) $\omega \in E^i$. This holds for all $1 \leq i \leq N$. As $P = \bar{P}$ on $\mathcal{F}_t^s$, taking P-expectations yields

$$
E^P[V_t^{s,\bar{\omega}}\mathbf{1}_{A_N}] \leq E^{\bar{P}_N}[\xi^{s,\bar{\omega}}\mathbf{1}_{A_N}] + \rho(\varepsilon), \tag{3.18}
$$

where we write $\bar{P}_N = \bar{P}$ to recall the dependence on N. Since $A_N \uparrow \Omega^s$, we have $\bar{P}_N(A_N^c) = P(A_N^c) \to 0$ as $N \to \infty$. In view of

$$
E^{\bar{P}_N}[\xi^{s,\bar{\omega}}\mathbf{1}_{A_N}] = E^{\bar{P}_N}[\xi^{s,\bar{\omega}}] - E^{\bar{P}_N}[\xi^{s,\bar{\omega}}\mathbf{1}_{A_N^c}] \leq E^{\bar{P}_N}[\xi^{s,\bar{\omega}}] + \|\xi\|_\infty P_N(A_N^c),
$$

we conclude from (3.18) that

$$
E^P[V_t^{s,\bar{\omega}}] \leq \limsup_{N\to\infty} E^{\bar{P}_N}[\xi^{s,\bar{\omega}}] + \rho(\varepsilon) \leq \sup_{P'\in\mathcal{P}(s,\bar{\omega})} E^{P'}[\xi^{s,\bar{\omega}}] + \rho(\varepsilon).
$$

Since $P \in \mathcal{P}(s,\bar{\omega})$ was arbitrary, letting $\varepsilon \to 0$ completes the proof of (3.11).

It remains to argue that our assumption (3.13) does not entail a loss of generality. Indeed, assume that $\hat{\omega}^i \notin \tilde{E}^i$ for some i. Then there are two possible cases. The case $P(E^i) = 0$ is easily seen to be harmless; recall that the measure P was fixed throughout the proof. In the case $P(E^i) > 0$, we also have $P(\tilde{E}^i) > 0$ and in particular $\tilde{E}^i \neq \emptyset$. Thus we can replace $\hat{\omega}^i$ by an arbitrary element of $\tilde{E}^i$ (which can be chosen independently of N). Using the continuity of the value function (Proposition 2.5) and of the reward function (2.12), we see that the above arguments still apply if we add an additional modulus of continuity in (3.15). $\qquad\square$

4 Extension of the Value Function

In this section, we extend the value function $\xi \mapsto V_t(\xi;\cdot)$ to an L^1-type space of random variables ξ, in the spirit of, e.g., [8]. While the construction of V_t in the previous section required a precise analysis "ω by ω", we can now move towards a more probabilistic presentation. In particular, we shall often write $V_t(\xi)$ for the random variable $\omega \mapsto V_t(\xi;\omega)$.

For reasons explained in Remark 4.2 below, we fix from now on an initial condition $x \in \mathbb{R}^d$ and let

$$
\mathcal{P}_x := \{P(0,x,\nu) : \nu \in \mathcal{U}\}
$$

be the corresponding set of measures at time $s = 0$. Given a random variable ψ on Ω, we write ψ^x as a shorthand for $\psi^{0,x} \equiv \psi(x \otimes_0 \cdot)$. We also write $V_t^x(\xi)$ for $(V_t(\xi))^x$.

Given $p \in [1,\infty)$, we define $L_{\mathcal{P}_x}^p$ to be the space of $\mathcal{F}_T$-measurable random variables

X satisfying

$$\|X\|_{L^p_{\mathcal{P}_x}} := \sup_{P \in \mathcal{P}_x} \|X\|_{L^p(P)} < \infty,$$

where $\|X\|^p_{L^p(P)} := E^P[|X|^p]$. More precisely, we identify functions which are equal $\mathcal{P}_x$-quasi-surely, so that $L^p_{\mathcal{P}_x}$ becomes a Banach space. (Two functions are equal $\mathcal{P}_x$-quasi-surely, $\mathcal{P}_x$-q.s. for short, if they are equal P-a.s. for all $P \in \mathcal{P}_x$.) Furthermore, given $t \in [0, T]$,

$$\mathbb{L}^p_{\mathcal{P}_x}(\mathcal{F}_t) \text{ is defined as the } \| \cdot \|_{L^p_{\mathcal{P}_x}}\text{-closure of } \mathrm{UC}_b(\Omega_t) \subseteq L^p_{\mathcal{P}_x}.$$

Since any $L^p_{\mathcal{P}_x}$-convergent sequence has a $\mathcal{P}_x$-q.s. convergent subsequence, any element of $\mathbb{L}^p_{\mathcal{P}_x}(\mathcal{F}_t)$ has an $\mathcal{F}_t$-measurable representative. For brevity, we shall often write $\mathbb{L}^p_{\mathcal{P}_x}$ for $\mathbb{L}^p_{\mathcal{P}_x}(\mathcal{F}_T)$.

Remark 4.1. *The space $\mathbb{L}^p_{\mathcal{P}_x}$ can be described as follows. We shall say that $\xi \in L^p_{\mathcal{P}_x}$ is $\mathcal{P}_x$-quasi uniformly continuous if ξ has a representative ξ' with the property that for all $\varepsilon > 0$ there exists an open set $G \subseteq \Omega$ such that $P(G) < \varepsilon$ for all $P \in \mathcal{P}$ and such that the restriction $\xi'|_{\Omega \setminus G}$ is uniformly continuous. Then $\mathbb{L}^p_{\mathcal{P}_x}$ consists of all $\xi \in L^p_{\mathcal{P}_x}$ such that ξ is $\mathcal{P}_x$-quasi uniformly continuous and $\lim_{n \to \infty} \|\xi \mathbf{1}_{\{|\xi| \geq n\}}\|_{L^p_{\mathcal{P}_x}} = 0$. Moreover, If $\mathcal{P}_x$ is weakly relatively compact, then $\mathbb{L}^p_{\mathcal{P}_x}$ contains all bounded continuous functions on Ω.*

The proof is the same as in [18, Proposition 5.2], which, in turn, followed an argument of [7].

Before extending the value function to $\mathbb{L}^1_{\mathcal{P}_x}$, let us explain why we are working under a fixed initial condition $x \in \mathbb{R}^d$.

Remark 4.2. *There is no fundamental obstruction to writing the theory without fixing the initial condition x; in fact, most of the results would be more elegant if stated using $\bar{\mathcal{P}}$ instead of $\mathcal{P}_x$, where $\bar{\mathcal{P}}$ is the set of all distributions of the form (2.4), with arbitrary initial condition. However, the set $\bar{\mathcal{P}}$ is very large and therefore the corresponding space $\mathbb{L}^1_{\bar{\mathcal{P}}}$ is very small, which is undesirable for the domain of our extended value function. As an illustration, consider a random variable of the form $\xi(\omega) := f(\omega_0)$ on Ω, where $f : \mathbb{R}^d \to \mathbb{R}$ is a measurable function. Then*

$$\|\xi\|_{L^1_{\bar{\mathcal{P}}}} = \sup_{x \in \mathbb{R}^d} |f(x)|;$$

i.e., ξ is in $L^1_{\bar{\mathcal{P}}}$ only when f is uniformly bounded. As a second issue in the same vein, it follows from the Arzelà-Ascoli theorem that the set $\bar{\mathcal{P}}$ is never weakly relatively compact. The latter property, which is satisfied by $\mathcal{P}_x$ for example when μ and σ are bounded, is sometimes useful in the context of quasi-sure analysis.

Lemma 4.3. *Let $p \in [1, \infty)$. The mapping V^x_t on $\mathrm{UC}_b(\Omega)$ is 1-Lipschitz,*

$$\|V^x_t(\xi) - V^x_t(\psi)\|_{L^p_{\mathcal{P}_x}} \leq \|\xi - \psi\|_{L^p_{\mathcal{P}_x}} \quad \text{for all} \quad \xi, \psi \in \mathrm{UC}_b(\Omega).$$

As a consequence, V^x_t uniquely extends to a Lipschitz-continuous mapping

$$V^x_t : \mathbb{L}^p_{\mathcal{P}_x}(\mathcal{F}_T) \to \mathbb{L}^p_{\mathcal{P}_x}(\mathcal{F}_t).$$

Proof. The argument is standard and included only for completeness. Note that $|\xi - \psi|^p$ is again in $\mathrm{UC}_b(\Omega)$. Moreover, the definition of V^x_t and Jensen's inequality imply that $|V^x_t(\xi) - V^x_t(\psi)|^p \leq V^x_t(|\xi - \psi|)^p \leq V^x_t(|\xi - \psi|^p)$. Therefore,

$$\|V^x_t(\xi) - V^x_t(\psi)\|_{L^p_{\mathcal{P}_x}} \leq \sup_{P \in \mathcal{P}_x} E^P\left[V^x_t(|\xi - \psi|^p)\right]^{1/p} = \sup_{P \in \mathcal{P}_x} E^P[|\xi - \psi|^p]^{1/p},$$

where the equality is due to (3.2) applied with $s = 0$. (For the case $s = 0$, the additional Assumption 3.1 was not used in the previous section.) Recalling from Proposition 2.5 that V_t^x maps $\mathrm{UC}_b(\Omega)$ to $\mathrm{UC}_b(\Omega_t)$, it follows that the extension maps $\mathbb{L}_{\mathcal{P}_x}^p$ to $\mathbb{L}_{\mathcal{P}_x}^p(\mathcal{F}_t)$. $\qquad\square$

4.1 Quasi-Sure Properties of the Extension

In this section, we provide some auxiliary results of technical nature. The first one will (quasi-surely) allow us to appeal to the results in the previous section without imposing Assumption 3.1. This is desirable since we would like to end up with quasi-sure theorems whose statements do not involve regular conditional probability distributions.

Lemma 4.4. *Assumption 2.1 implies that Assumption 3.1 holds for $\mathcal{P}_x$-quasi-every $\eta \in \Omega$ satisfying $\eta_0 = x$.*

For the proof of this lemma, we shall use the following result.

Lemma 4.5. *Let Y and Z be continuous adapted processes, $t \in [0, T]$ and let P be a probability measure on Ω. Then $\overline{\mathbb{F}^Y}^P \supseteq \mathbb{F}^Z$ implies that*

$$\overline{\mathbb{F}^{Y^{t,\omega}}}^{P^{t,\omega}} \supseteq \mathbb{F}^{Z^{t,\omega}} \quad for \quad P\text{-}a.e.\ \omega \in \Omega.$$

Proof. The assumption implies that there exists a progressively measurable transformation $\beta : \Omega \to \Omega$ such that $Z = \beta(Y)$ P-a.s. For P-a.e. $\omega \in \Omega$, it follows that $Z(\omega \otimes_t \cdot) = \beta(Y(\omega \otimes_t \cdot))$ $P^{t,\omega}$-a.s., which, in turn, yields the result. $\qquad\square$

Proof of Lemma 4.4. Let $X := X(t, \eta, \nu)$ with $\eta_0 = x$, we have to show that $\overline{\mathbb{F}^X}^{P_0^t} \supseteq \mathbb{F}^t$ whenever η^0 is outside some $\mathcal{P}_x$-polar set. Hence we shall fix an arbitrary $\hat{P} \in \mathcal{P}_x$ and show that the result holds on a set of full measure $\hat{P}$.

Let $\hat{P} \in \mathcal{P}_x$, then $\hat{P} = P(0, x, \hat{\nu})$ for some $\hat{\nu} \in \mathcal{U}$ and $\hat{P}$ is concentrated on the image of $\hat{X}^0$, where $\hat{X} := X(0, x, \hat{\nu})$. That is, recalling that $\eta_0 = \hat{X}_0 = x$, we may assume that $\eta = \hat{X}(\omega)$ for some $\omega \in \Omega$. Let

$$\bar{\nu}(\tilde{\omega}) := \mathbf{1}_{[0,t)}\hat{\nu}(\tilde{\omega}) + \mathbf{1}_{[t,T]}\nu(\tilde{\omega}^t). \tag{4.1}$$

Then $\bar{X} := X(0, x, \bar{\nu})$ satisfies $\bar{X} = \hat{X}$ on $[0, t]$ and hence we may assume that $\eta = \bar{X}(\omega)$ on $[0, t]$. Using Lemma 3.3 and (4.1), we deduce that

$$\bar{X}^{t,\omega} = X\big(t, x \otimes \bar{X}(\omega), \bar{\nu}^{t,\omega}\big) = X(t, \eta, \nu) = X.$$

Since $\mathbb{F} \subseteq \overline{\mathbb{F}^{\bar{X}}}^{P_0}$ by Assumption 2.1, we conclude that

$$\overline{\mathbb{F}^t}^{P_0^t} \subseteq \overline{\mathbb{F}^{\bar{X}^{t,\omega}}}^{P_0^t} = \overline{\mathbb{F}^X}^{P_0^t}$$

by using Lemma 4.5 with Z being the canonical process. $\qquad\square$

The next two results show that (for $\mu \equiv 0$ and σ positive definite) the mapping $\xi \mapsto V_t^x(\xi)$ on $\mathbb{L}_{\mathcal{P}_x}^1$ falls into the general class of sublinear expectations considered in [19], whose techniques we shall apply in the subsequent section. More precisely, the two lemmas below yield the validity of its main condition [19, Assumption 4.1].

The following property is known as *stability under pasting* and well known to be important in non-Markovian control. It should not be confused with the pasting discussed in Remark 3.6, where the considered measures correspond to different points in time.

Lemma 4.6. *Let τ be an $\mathbb{F}$-stopping time and let $\Lambda \in \mathcal{F}_\tau$. Let $P, P^1, P^2 \in \mathcal{P}_x$ satisfy $P^1 = P^2 = P$ on $\mathcal{F}_\tau$. Then*

$$\bar{P}(A) := E^P\big[P^1(A|\mathcal{F}_\tau)\mathbf{1}_\Lambda + P^2(A|\mathcal{F}_\tau)\mathbf{1}_{\Lambda^c}\big], \quad A \in \mathcal{F}_T$$

defines an element of $\mathcal{P}_x$.

Proof. It follows from the definition of the conditional expectation that $\bar{P}$ is a probability measure which is characterized by the properties

$$\bar{P} = P \text{ on } \mathcal{F}_\tau \quad \text{and} \quad \bar{P}^{\tau(\omega),\omega} = \begin{cases} (P^1)^{\tau(\omega),\omega} & \text{for} \quad P\text{-a.e. } \omega \in \Lambda, \\ (P^2)^{\tau(\omega),\omega} & \text{for} \quad P\text{-a.e. } \omega \in \Lambda^c. \end{cases} \tag{4.2}$$

Let $\nu, \nu^1, \nu^2 \in \mathcal{U}$ be such that $P = P(0, x, \nu)$ and $P^i = P(0, x, \nu^i)$ for $i = 1, 2$. Moreover, let $X := X(0, x, \nu)$, define $\bar{\nu} \in \mathcal{U}$ by

$$\bar{\nu}_r(\omega) := \mathbf{1}_{[0,\tau(X(\omega)^0))}(r)\nu_r(\omega)$$
$$+ \mathbf{1}_{[\tau(X(\omega)^0),T]}(r)\Big[\nu_r^1(\omega)\mathbf{1}_\Lambda(X(\omega)^0) + \nu_r^2(\omega)\mathbf{1}_{\Lambda^c}(X(\omega)^0)\Big]$$

and let $P_* := P(0, x, \bar{\nu}) \in \mathcal{P}_x$. We show that P_* satisfies the three properties from (4.2). Indeed, $\nu = \bar{\nu}$ on $[0, \tau(X^0))$ implies that $P_* = P$ on $\mathcal{F}_\tau$. Moreover, as in (3.8),

$$P_*^{\tau(\omega),\omega} = P\big(\tau(\omega), x \otimes_0 \omega, \bar{\nu}^{\tau(X(\omega)^0),\breve{\omega}}\big) \quad \text{for} \quad P\text{-a.e. } \omega \in \Omega.$$

Similarly as below (3.9), we also have that

$$\bar{\nu}^{\tau(X(\omega)^0),\breve{\omega}} = \nu^1\big(\breve{\omega} \otimes_{\tau(X(\omega)^0)} \cdot\big) = (\nu^1)^{\tau(X(\omega)^0),\breve{\omega}} \quad \text{for} \quad P\text{-a.e. } \omega \in \Lambda.$$

Therefore,

$$P_*^{\tau(\omega),\omega} = P\big(\tau(\omega), x \otimes_0 \omega, (\nu^1)^{\tau(X(\omega)^0),\breve{\omega}}\big) = (P^1)^{\tau(\omega),\omega} \quad \text{for} \quad P\text{-a.e. } \omega \in \Lambda.$$

An analogous argument establishes the third property from (4.2) and we conclude that $\bar{P} = P_* \in \mathcal{P}_x$. $\qquad\square$

The second property is the quasi-sure representation of $V_t^x(\xi)$ on $\mathbb{L}^1_{\mathcal{P}_x}$, a result which will be generalized in Theorem 5.2 below.

Lemma 4.7. *Let $t \in [0, T]$ and $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$. Then*

$$V_t^x(\xi) = \operatorname*{ess\,sup}_{P' \in \mathcal{P}_x(\mathcal{F}_t, P)}^P E^{P'}[\xi^x|\mathcal{F}_t] \quad P\text{-a.s.} \quad \text{for all} \quad P \in \mathcal{P}_x, \tag{4.3}$$

where $\mathcal{P}_x(\mathcal{F}_t, P) := \{P' \in \mathcal{P}_x : P' = P \text{ on } \mathcal{F}_t\}$.

Proof. Recall that Lemma 4.4 allows us to appeal to the results of Section 3.

(i) We first prove the inequality "$\leq$" for $\xi \in \mathrm{UC}_b(\Omega)$. Fix $P \in \mathcal{P}_x$. We use Step (ii) of the proof of Theorem 3.2, in particular (3.17), for the special case $s = 0$ and obtain that for given $\varepsilon > 0$ and $N \geq 1$ there exists a measure $\bar{P}_N \in \mathcal{P}_x(\mathcal{F}_t, P)$ such that

$$V_t^x(\omega) \leq E^{\bar{P}_N}[\xi^x|\mathcal{F}_t](\omega) + \rho(\varepsilon) \quad \text{for} \quad P\text{-a.s. } \omega \in E^1 \cup \cdots \cup E^N,$$

where $V_t^x(\omega) := V_t^x(\xi; \omega)$. Since $\bigcup_{i \geq 1} E^i = \Omega^0 = \Omega$ P-a.s., we deduce that

$$V_t^x(\omega) \leq \sup_{N \geq 1} E^{\bar{P}_N}[\xi^x|\mathcal{F}_t](\omega) + \rho(\varepsilon) \quad \text{for} \quad P\text{-a.s. } \omega \in \Omega.$$

The claim follows by letting $\varepsilon \to 0$.

(ii) Next, we show the inequality "$\geq$" in (4.3) for $\xi \in \mathrm{UC}_b(\Omega)$. Fix $P, P' \in \mathcal{P}_x$ and recall that $(P')^{t,\omega} \in \mathcal{P}(t, x \otimes_0 \omega)$ for P'-a.s. $\omega \in \Omega$ by Lemma 3.4. Therefore, (3.2) applied with $s := t$ and $t := T$ yields that

$$V_t^x(\omega) = V_t(x \otimes_0 \omega) \geq E^{(P')^{t,\omega}}[\xi^{t, x \otimes_0 \omega}] = E^{(P')^{t,\omega}}[(\xi^x)^{t,\omega}] = E^{P'}[\xi^x | \mathcal{F}_s](\omega)$$

P'-a.s. on $\mathcal{F}_t$. If $P' \in \mathcal{P}_x(\mathcal{F}_t, P)$, then $P' = P$ on $\mathcal{F}_t$ and the inequality holds also P-a.s. The claim follows as $P' \in \mathcal{P}_x(\mathcal{F}_t, P)$ was arbitrary.

(iii) So far, we have proved the result for $\xi \in \mathrm{UC}_b(\Omega)$. The general case $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$ can be derived by an approximation argument exploiting the stability under pasting (Lemma 4.6). We omit the details since the proof is exactly the same as in [18, Theorem 5.4]. $\qquad\square$

5 Path Regularity for the Value Process

In this section, we construct a càdlàg $\mathcal{P}_x$-modification for $V^x(\xi)$; that is, a càdlàg process Y^x such that $Y_t^x = V_t^x(\xi)$ $\mathcal{P}_x$-q.s. for all $t \in [0, T]$. (Recall that the initial condition $x \in \mathbb{R}^d$ has been fixed.) To this end, we extend the raw filtration $\mathbb{F}$ as in [19]: we let $\mathbb{F}^+ = \{\mathcal{F}_{t+}\}_{0 \leq t \leq T}$ be the minimal right-continuous filtration containing $\mathbb{F}$ and we augment $\mathbb{F}^+$ by the collection $\mathcal{N}^{\mathcal{P}_x}$ of $(\mathcal{P}_x, \mathcal{F}_T)$-polar sets to obtain the filtration

$$\mathbb{G} = \{\mathcal{G}_t\}_{0 \leq t \leq T}, \quad \mathcal{G}_t := \mathcal{F}_{t+} \vee \mathcal{N}^{\mathcal{P}_x}.$$

We note that $\mathbb{G}$ depends on $x \in \mathbb{R}^d$ since $\mathcal{N}^{\mathcal{P}_x}$ does, but for brevity, we shall not indicate this in the notation. In fact, the dependence on x is not crucial: we could also work with $\mathbb{F}^+$, at the expense of obtaining a modification which is $\mathcal{P}_x$-q.s. equal to a càdlàg process rather than being càdlàg itself.

We recall that in the quasi-sure setting, value processes similar to the one under consideration do not admit càdlàg modifications in general; indeed, while the right limit exists quasi-surely, it need not be a modification (cf. [19]). Both the regularity of $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$ and the regularity induced by the SDE are crucial for the following result.

Theorem 5.1. Let $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$. There exists a $(\mathcal{P}_x$-q.s. unique$)$ $\mathbb{G}$-adapted càdlàg $\mathcal{P}_x$-modification $\mathcal{E}^x(\xi) = \{\mathcal{E}_t^x(\xi)\}_{t \in [0,T]}$ of $\{V_t^x(\xi)\}_{t \in [0,T]}$. Moreover,

$$\mathcal{E}_t^x(\xi) = \operatorname*{ess\,sup}_{P' \in \mathcal{P}_x(\mathcal{G}_t, P)}{}^P E^{P'}[\xi^x | \mathcal{G}_t] \quad P\text{-a.s.} \quad \text{for all} \quad P \in \mathcal{P}_x, \tag{5.1}$$

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

Proof. In view of Lemmata 4.6 and 4.7, we obtain exactly as in [19, Proposition 4.5] that there exists a $\mathcal{P}_x$-q.s. unique $\mathbb{G}$-adapted càdlàg process $\mathcal{E}^x(\xi)$ satisfying (5.1) and

$$\mathcal{E}_t^x(\xi) = V_{t+}^x(\xi) := \lim_{r \downarrow t} V_r^x(\xi) \quad \mathcal{P}_x\text{-q.s.} \quad \text{for all} \quad 0 \leq t < T. \tag{5.2}$$

The observation made there is that (4.3) implies that $V^x(\xi)$ is a $(P, \mathbb{F})$-supermartingale for all $P \in \mathcal{P}_x$, so that one can use the standard modification argument for supermartingales under each P. This argument, cf. [6, Theorem VI.2], also yields that $E^P[\mathcal{E}_t^x(\xi) | \mathcal{F}_{t+}] \leq V_t^x(\xi)$ P-a.s. and in particular

$$E^P[\mathcal{E}_t^x(\xi)] \leq E^P[V_t^x(\xi)] \quad \text{for all} \quad P \in \mathcal{P}_x.$$

Hence, it remains to show that

$$\mathcal{E}_t^x(\xi) \geq V_t^x(\xi) \quad \mathcal{P}_x\text{-q.s.} \tag{5.3}$$

for $t \in [0, T)$, which is the part that is known to fail in a more general setting. We give the proof in several steps.

(i) We first show that $\mathcal{E}_t^x$ maps $\mathrm{UC}_b(\Omega)$ to $\mathbb{L}^1_{\mathcal{P}_x}(\mathcal{F}_t)$, and in fact even to $\mathrm{UC}_b(\Omega_t)$ if a suitable representative is chosen. Let $\xi \in \mathrm{UC}_b(\Omega)$, $r \in (t, T]$ and set $V_r^x := V_r^x(\xi)$. By Proposition 2.5, there exists a modulus of continuity ρ independent of r such that

$$|V_r^x(\omega) - V_r^x(\omega')| \leq \rho(\|\omega - \omega'\|_r).$$

Hence the $\mathcal{P}_x$-q.s. limit from (5.2) satisfies

$$|V_{t+}(\omega) - V_{t+}(\omega')| \leq \rho(\|\omega - \omega'\|_r) \quad \text{for all} \quad r \in (t, T] \cap \mathbb{Q}, \quad \mathcal{P}_x\text{-q.s.}$$

Since $\mathcal{E}_t^x(\xi) = V_{t+}(\xi)$, taking the limit $r \downarrow t$ yields that

$$|\mathcal{E}_t^x(\xi)(\omega) - \mathcal{E}_t^x(\xi)(\omega')| \leq \rho(\|\omega - \omega'\|_t) \quad \mathcal{P}_x\text{-q.s.}$$

By a variant of Tietze's extension theorem, cf. [16], this implies that $\mathcal{E}_t^x(\xi)$ coincides $\mathcal{P}_x$-q.s. with an element of $\mathrm{UC}_b(\Omega_t)$. In particular, $\mathcal{E}_t^x(\xi) \in \mathbb{L}^1_{\mathcal{P}_x}(\mathcal{F}_t)$.

(ii) Next, we show that $\mathcal{E}_t^x$ is Lipschitz-continuous. Let $\xi, \psi \in \mathbb{L}^1_{\mathcal{P}_x}$ and $t_n \downarrow t$. Using (5.2), Fatou's lemma and Lemma 4.3, we obtain that

$$\begin{aligned}
\|\mathcal{E}_t^x(\xi) - \mathcal{E}_t^x(\psi)\|_{L^1_{\mathcal{P}_x}} &= \left\|\lim_n |V_{t_n}^x(\xi) - V_{t_n}^x(\psi)|\right\|_{L^1_{\mathcal{P}_x}} \\
&= \sup_{P \in \mathcal{P}_x} E^P\left[\lim_n |V_{t_n}^x(\xi) - V_{t_n}^x(\psi)|\right] \\
&\leq \sup_{P \in \mathcal{P}_x} \liminf_n E^P\left[|V_{t_n}^x(\xi) - V_{t_n}^x(\psi)|\right] \\
&\leq \|\xi - \psi\|_{L^1_{\mathcal{P}_x}}.
\end{aligned}$$

(iii) Let $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$. Then there exist $\xi^n \in \mathrm{UC}_b(\Omega)$ such that $\xi^n \to \xi$ in $L^1_{\mathcal{P}_x}$ and in thus $\mathcal{E}_t^x(\xi^n) \to \mathcal{E}_t^x(\xi)$ in $L^1_{\mathcal{P}_x}$ by Step (ii). Since $\mathcal{E}_t^x(\xi^n) \in \mathbb{L}^1_{\mathcal{P}_x}(\mathcal{F}_t)$ by Step (i) and since $\mathbb{L}^1_{\mathcal{P}_x}(\mathcal{F}_t)$ is closed in $L^1_{\mathcal{P}_x}$, we conclude that

$$\mathcal{E}_t^x(\xi) \in \mathbb{L}^1_{\mathcal{P}_x}(\mathcal{F}_t) \quad \text{for all} \quad \xi \in \mathbb{L}^1_{\mathcal{P}_x}.$$

(iv) Let $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$. Since V_t^x is the identity on $\mathbb{L}^1_{\mathcal{P}_x}(\mathcal{F}_t)$, Step (iii) implies that $\mathcal{E}_t^x(\xi) = V_t^x(\mathcal{E}_t^x(\xi))$. Moreover, the representations (4.3) and (5.1) yield

$$\begin{aligned}
V_t^x(\mathcal{E}_t^x(\xi)) &= \operatorname*{ess\,sup}^P_{P' \in \mathcal{P}_x(\mathcal{F}_t, P)} E^{P'}[\mathcal{E}_t^x(\xi)|\mathcal{F}_t] \\
&\geq \operatorname*{ess\,sup}^P_{P' \in \mathcal{P}_x(\mathcal{F}_t, P)} E^{P'}\left[E^{P'}[\xi^x|\mathcal{G}_t]\big|\mathcal{F}_t\right] \\
&= \operatorname*{ess\,sup}^P_{P' \in \mathcal{P}_x(\mathcal{F}_t, P)} E^{P'}[\xi^x|\mathcal{F}_t] \\
&= V_t^x(\xi) \quad P\text{-a.s.} \quad \text{for all} \quad P \in \mathcal{P}_x.
\end{aligned}$$

We conclude that (5.3) holds true. $\qquad\square$

Since $\mathcal{E}^x(\xi)$ is a càdlàg process, its value $\mathcal{E}_\tau^x(\xi)$ at a stopping time τ is well defined. The following result states the quasi-sure representation of $\mathcal{E}_\tau^x(\xi)$ and the quasi-sure

version of the dynamic programming principle in its final form.

Theorem 5.2. *Let $0 \leq \varrho \leq \tau \leq T$ be $\mathbb{G}$-stopping times and $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$. Then*

$$\mathcal{E}^x_\varrho(\xi) = \operatorname*{ess\,sup}_{P' \in \mathcal{P}_x(\mathcal{G}_\varrho, P)}^P E^{P'}[\mathcal{E}^x_\tau(\xi)|\mathcal{G}_\varrho] \quad \textit{P-a.s.} \quad \textit{for all} \quad P \in \mathcal{P}_x \tag{5.4}$$

and in particular

$$\mathcal{E}^x_\varrho(\xi) = \operatorname*{ess\,sup}_{P' \in \mathcal{P}_x(\mathcal{G}_\varrho, P)}^P E^{P'}[\xi^x|\mathcal{G}_\varrho] \quad \textit{P-a.s.} \quad \textit{for all} \quad P \in \mathcal{P}_x.$$

Moreover, there exists for each $P \in \mathcal{P}_x$ a sequence $P_n \in \mathcal{P}_x(\mathcal{G}_\varrho, P)$ such that

$$\mathcal{E}^x_\varrho(\xi) = \lim_{n \to \infty} E^{P_n}[\xi^x|\mathcal{G}_\varrho] \quad \textit{P-a.s.}$$

with a P-a.s. increasing limit.

Proof. In view of Lemmata 4.6 and 4.7, the result is derived exactly as in [19, Theorem 4.9]. $\qquad\square$

As in (3.3), the relation (5.4) can be seen as a semigroup property

$$\mathcal{E}^x_\varrho(\xi) = \mathcal{E}^x_\varrho(\mathcal{E}^x_\tau(\xi)),$$

at least when $\mathcal{E}^x_\tau(\xi)$ is in the domain $\mathbb{L}^1_{\mathcal{P}_x}$ of $\mathcal{E}^x_\varrho$. The latter is guaranteed by Lemma 4.3 when τ is a deterministic time. However, one cannot expect $\mathcal{E}^x_\tau(\xi)$ to be quasi uniformly continuous (cf. Remark 4.1) for a general stopping time, for which reason we prefer to express the right hand side as in (5.4).

6 Hamilton-Jacobi-Bellman 2BSDE

In this section, we characterize the value process $\mathcal{E}^x(\xi)$ as the solution of a 2BSDE. To this end, we first examine the properties of B under a fixed $P \in \mathcal{P}_x$. The following result is in the spirit of [28, Section 8].

Proposition 6.1. *Let $x \in \mathbb{R}^d$, $\nu \in \mathcal{U}$ and $P := P(0, x, \nu)$. There exists a progressively measurable transformation $\beta : \Omega \to \Omega$ (depending on x, ν) such that $W := \beta(B)$ is a P-Brownian motion and*

$$\overline{\mathbb{F}}^P = \overline{\mathbb{F}^W}^P. \tag{6.1}$$

Moreover, B is the P-a.s. unique strong solution of the SDE

$$B = \int_0^\cdot \mu(t, x + B, \nu_t(W)) \, dt + \int_0^\cdot \sigma(t, x + B, \nu_t(W)) \, dW_t \quad \textit{under} \quad P.$$

Proof. Let $X := X(0, x, \nu)$. As in Lemma 3.4, Assumption 2.1 implies the the existence of a progressively measurable transformation $\beta : \Omega \to \Omega$ such that

$$\beta(X^0) = B \quad P_0\text{-a.s.} \tag{6.2}$$

Let $W := \beta(B)$. Then

$$(B, X^0)_{P_0} = (\beta(X^0), X^0)_{P_0} = (\beta(B), B)_P = (W, B)_P;$$

i.e., the distribution of (B, X^0) under P_0 coincides with the distribution of (W, B) under P. In particular, W is a P-Brownian motion. Moreover, we have $\overline{\mathbb{F}^{X^0}}^{P_0} = \overline{\mathbb{F}^{\beta(X^0)}}^{P_0}$ by

Assumption 2.1 and therefore

$$\overline{\mathbb{F}^B}^P = \overline{\mathbb{F}^{\beta(B)}}^P,$$

which is (6.1). Note that

$$X^0 = X^0(B) = \int \mu(t, x + X^0, \nu_t(B)) \, dt + \int \sigma(t, x + X^0, \nu_t(B)) \, dB_t$$

under P_0. Let Y be the (unique, strong) solution of the analogous SDE

$$Y = \int \mu(t, x + Y, \nu_t(W)) \, dt + \int \sigma(t, x + Y, \nu_t(W)) \, dW_t \quad \text{under} \quad P.$$

Using the definition of P and (6.2), we have that

$$
\begin{aligned}
(Y, W)_P &= (X^0(W), W)_P \\
&= (X^0(B), B)_{P_0} \\
&= (X^0, \beta(X^0))_{P_0} \\
&= (B, \beta(B))_P \\
&= (B, W)_P.
\end{aligned}
$$

In view of (6.1), it follows that $Y = B$ holds P-a.s. $\qquad\square$

In the sequel, we denote by $M^{B,P}$ the local martingale part in the canonical semimartingale decomposition of B under P.

Corollary 6.2. *Let $P \in \mathcal{P}_x$. Then the filtration $\overline{\mathbb{F}}^P$ is right-continuous. If, in addition, σ is invertible, then $(M^{B,P}, P)$ has the predictable representation property.*

The latter statement means that any right-continuous $(\overline{\mathbb{F}}^P, P)$-local martingale N has a representation $N = N_0 + \int Z \, dM^{B,P}$ under P, for some $\overline{\mathbb{F}}^P$-predictable process Z.

Proof. We have seen in Proposition 6.1 that $\overline{\mathbb{F}}^P$ is generated by a Brownian motion W, hence right-continuous, and that $M^{B,P} = \int_0^{\cdot} \hat{\sigma}_t \, dW_t$ for $\hat{\sigma}_t := \sigma(t, x + B, \nu_t(W))$, where $\nu \in \mathcal{U}$. By changing $\hat{\sigma}$ on a $dt \times P$-nullset, we may assume that $\hat{\sigma}$ is $\overline{\mathbb{F}}^P$-predictable. Using the Brownian representation theorem and $W = \int \hat{\sigma}^{-1} \, dM^{B,P}$, we deduce that $M^{B,P}$ has the representation property. $\qquad\square$

The following formulation of 2BSDE is, of course, inspired by [27].

Definition 6.3. *Let $\xi \in L^1_{\mathcal{P}_x}$ and consider a pair (Y, Z) of processes with values in $\mathbb{R} \times \mathbb{R}^d$ such that Y is càdlàg $\mathbb{G}$-adapted while Z is $\mathbb{G}$-predictable and $\int_0^T |Z_s|^2 \, d\langle B \rangle_s < \infty$ $\mathcal{P}_x$-q.s. Then (Y, Z) is called a solution of the 2BSDE (6.3) if there exists a family $(K^P)_{P \in \mathcal{P}_x}$ of $\overline{\mathbb{F}}^P$-adapted increasing processes satisfying $E^P[|K_T^P|] < \infty$ such that*

$$Y_t = \xi - \int_t^T Z_s \, dM_s^{B,P} + K_T^P - K_t^P, \quad 0 \le t \le T, \quad P\text{-a.s.} \quad \text{for all } P \in \mathcal{P}_x \qquad (6.3)$$

and such that the following minimality condition holds for all $0 \le t \le T$:

$$\operatorname*{ess\,inf}_{P' \in \mathcal{P}_x(\mathcal{G}_t, P)}^P E^{P'}\big[K_T^{P'} - K_t^{P'} \big| \mathcal{G}_t\big] = 0 \quad P\text{-a.s.} \quad \text{for all } P \in \mathcal{P}_x. \qquad (6.4)$$

Moreover, a càdlàg process Y is said to be *of class* $(D,\mathcal{P}_x)$ if the family $\{Y_\tau\}_\tau$ is uniformly integrable under P for all $P \in \mathcal{P}_x$, where τ runs through all $\mathbb{G}$-stopping times. The following is our main result.

Theorem 6.4. *Assume that σ is invertible and let $\xi \in \mathbb{L}^1_{\mathcal{P}_x}$.*

  (i) *There exists a $(dt \times \mathcal{P}_x$-q.s. unique$)$ $\mathbb{G}$-predictable process Z^ξ such that*

$$Z^\xi = \left(d\langle B, B\rangle^P\right)^{-1} d\langle \mathcal{E}^x(\xi), B\rangle^P \quad P\text{-a.s.} \quad \text{for all} \quad P \in \mathcal{P}_x. \tag{6.5}$$

  (ii) *The pair $(\mathcal{E}^x(\xi), Z^\xi)$ is the minimal solution of the 2BSDE (6.3); i.e., if (Y, Z) is another solution, then $\mathcal{E}^x(\xi) \leq Y$ $\mathcal{P}_x$-q.s.*

  (iii) *If (Y, Z) is a solution of (6.3) such that Y is of class $(D,\mathcal{P}_x)$, then $(Y, Z) = (\mathcal{E}^x(\xi), Z^\xi)$.*

In particular, if $\xi \in \mathbb{L}^p_{\mathcal{P}_x}$ for some $p > 1$, then $(\mathcal{E}^x(\xi), Z^\xi)$ is the unique solution of (6.3) in the class $(D,\mathcal{P}_x)$.

Proof. Given two processes which are (càdlàg) semimartingales under all $P \in \mathcal{P}_x$, their quadratic covariation can be defined $\mathcal{P}_x$-q.s. by using the integration-by-parts formula and Bichteler's pathwise stochastic integration [1, Theorem 7.14]; therefore, the right hand side of (6.5) can be used as a definition of Z^ξ. The details of the argument are as in [19, Proposition 4.10].

Let $P \in \mathcal{P}_x$. By Proposition 6.1, B is an Itô process under P; in particular, we have $\langle B, S\rangle^P = \langle M^{B,P}, S\rangle^P$ P-a.s. for any P-semimartingale S. The Doob-Meyer theorem under P and Corollary 6.2 then yield the decomposition

$$\mathcal{E}^x(\xi) = \mathcal{E}^x_0(\xi) + \int Z^\xi \, dM^{B,P} - K^P \quad P\text{-a.s.}$$

and we obtain (ii) and (iii) by following the arguments in [19, Theorem 4.15]. If $\xi \in \mathbb{L}^p_{\mathcal{P}_x}$ for some $p \in (1, \infty)$, then $\mathcal{E}^x(\xi)$ is of class $(D,\mathcal{P}_x)$ as a consequence of Jensen's inequality (cf. [19, Lemma 4.14]). Therefore, the last assertion follows from the above. $\square$

We conclude by interpreting the canonical process B, seen under the "set of scenarios" $\mathcal{P}_x$, as a model for drift and volatility uncertainty in the Knightian sense.

Remark 6.5. *Consider the set-valued process*

$$\mathbf{D}_t(\omega) := \left\{\left(\mu(t, \omega, u), \sigma(t, \omega, u)\right) : u \in U\right\} \subseteq \mathbb{R}^d \times \mathbb{R}^{d \times d}.$$

In view of Proposition 6.1, each $P \in \mathcal{P}_x$ can be seen as a scenario in which the drift and the volatility (of B) take values in $\mathbf{D}$, P-a.s. Then, the upper expectation $\mathcal{E}^x(\xi)$ is the corresponding worst-case expectation (see [19] for a connection to superhedging in finance). Note that $\mathbf{D}$ is a random process although the coefficients of our controlled SDE are non-random. Indeed, the path-dependence of the SDE translates to an ω-dependence in the weak formulation that we are considering.

In particular, for $\mu \equiv 0$, we have constructed a sublinear expectation similar to the random G-expectation of [18]. While the latter is defined by specifying a set-valued process like $\mathbf{D}$ in the first place, we have started here from a controlled SDE under P_0. It seems that the present construction is somewhat less technical that the one in [18]; in particular, we did not work with the process $\hat{a} = d\langle B\rangle_t/dt$ which played an important role in [26] and [18]. However, it seems that the Lipschitz conditions on μ and σ are essential, while [18] merely used a notion of uniform continuity.

References

[1] K. Bichteler, *Stochastic integration and L^p-theory of semimartingales*, Ann. Probab. **9** (1981), no. 1, 49–89.MR-0606798

[2] R. Buckdahn and J. Ma, *Stochastic viscosity solutions for nonlinear stochastic partial differential equations. I.*, Stochastic Process. Appl. **93** (2001), no. 2, 181–204. MR-1828772

[3] _____, *Stochastic viscosity solutions for nonlinear stochastic partial differential equations. II.*, Stochastic Process. Appl. **93** (2001), no. 2, 205–228. MR-1831830

[4] _____, *Pathwise stochastic control problems and stochastic HJB equations*, SIAM J. Control Optim. **45** (2007), no. 6, 2224–2256. MR-2285722

[5] P. Cheridito, H. M. Soner, N. Touzi, and N. Victoir, *Second-order backward stochastic differential equations and fully nonlinear parabolic PDEs*, Comm. Pure Appl. Math. **60** (2007), no. 7, 1081–1110. MR-2319056

[6] C. Dellacherie and P. A. Meyer, *Probabilities and potential B*, North Holland, Amsterdam, 1982. MR-0745449

[7] L. Denis, M. Hu, and S. Peng, *Function spaces and capacity related to a sublinear expectation: application to G-Brownian motion paths*, Potential Anal. **34** (2011), no. 2, 139–161. MR-2754968

[8] L. Denis and C. Martini, *A theoretical framework for the pricing of contingent claims in the presence of model uncertainty*, Ann. Appl. Probab. **16** (2006), no. 2, 827–852. MR-2244434

[9] N. El Karoui, *Les aspects probabilistes du contrôle stochastique*, Ecole d'été de probabilitées de Saint-Flour (Springer, Berlin), Lecture Notes in Math., vol. 876, 1981, pp. 73–238. MR-0637471

[10] R. J. Elliott, *Stochastic calculus and applications*, Springer, New York, 1982. MR-0678919

[11] J. Feldman and M. Smorodinsky, *Simple examples of non-generating Girsanov processes*, Séminaire de Probabilités XXXI, Lecture Notes in Math., vol. 1655, Springer, Berlin, 1997, pp. 247–251. MR-1478733

[12] W. H. Fleming and H. M. Soner, *Controlled Markov processes and viscosity solutions*, 2nd ed., Springer, New York, 2006. MR-2179357

[13] Y. Hu, P. Imkeller, and M. Müller, *Utility maximization in incomplete markets*, Ann. Appl. Probab. **15** (2005), no. 3, 1691–1712. MR-2152241

[14] P.-L. Lions and P. Souganidis, *Fully nonlinear stochastic partial differential equations*, C. R. Acad. Sci. Paris Sér. I Math. **326** (1998), no. 9, 1085–1092. MR-1647162

[15] _____, *Fully nonlinear stochastic partial differential equations: non-smooth equations and applications*, C. R. Acad. Sci. Paris Sér. I Math. **327** (1998), no. 8, 735–741. MR-1659958

[16] M. Mandelkern, *On the uniform continuity of Tietze extensions*, Arch. Math. **55** (1990), no. 4, 387–388. MR-1076066

[17] M. Nutz, *The Bellman equation for power utility maximization with semimartingales*, Ann. Appl. Probab. **22** (2012), no. 1, 363–406.

[18] _____, *Random G-expectations*, Preprint arXiv:1009.2168v1 (2010).

[19] M. Nutz and H. M. Soner, *Superhedging and dynamic risk measures under volatility uncertainty*, Preprint arXiv:1011.2958v1 (2010).

[20] E. Pardoux and S. Peng, *Adapted solution of a backward stochastic differential equation*, Systems Control Lett. **14** (1990), no. 1, 55–61. MR-1037747

[21] S. Peng, *Stochastic Hamilton-Jacobi-Bellman equations*, SIAM J. Control Optim. **30** (1992), no. 2, 284–304. MR-1149069

[22] _____, *Filtration consistent nonlinear expectations and evaluations of contingent claims*, Acta Math. Appl. Sin. Engl. Ser. **20** (2004), no. 2, 191–214. MR-2064000

[23] _____, *G-expectation, G-Brownian motion and related stochastic calculus of Itô type*, Stochastic Analysis and Applications (Springer, Berlin), Abel Symp., vol. 2, 2007, pp. 541–567.

[24] _____, *Multi-dimensional G-Brownian motion and related stochastic calculus under G-expectation*, Stochastic Process. Appl. **118** (2008), no. 12, 2223–2253. MR-2474349

[25] _____, *Note on viscosity solution of path-dependent PDE and G-martingales*, Preprint arXiv:1106.1144v1 (2011).

[26] H. M. Soner, N. Touzi, and J. Zhang, *Dual formulation of second order target problems*, To appear in Ann. Appl. Probab. (2010).

[27] _____, *Wellposedness of second order backward SDEs*, To appear in Probab. Theory Related Fields (2010).

[28] _____, *Quasi-sure stochastic analysis through aggregation*, Electron. J. Probab. **16** (2011), no. 2, 1844–1879. MR-2842089

[29] D. Stroock and S. R. S. Varadhan, *Multidimensional diffusion processes*, Springer, New York, 1979. MR-0532498