The Dissertation Committee for Roman Fayvisovich
certifies that this is the approved version of the following dissertation:

Martingale-Generated Control Structures and a Framework for the Dynamic Programming Principle

Committee:

Gordan Žitković, Supervisor

Mihai Sîrbu

Thaleia Zariphopoulou

Kasper Larsen

Martingale-Generated Control Structures and a Framework for the Dynamic Programming Principle

by

Roman Fayvisovich

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2017

Dedicated to my parents, Aleksandr and Larisa Fayvisovich.

Acknowledgments

I wish to thank the members of my PhD committee: Gordan Žitković, Mihai Sîrbu, Thaleia Zariphopoulou, and Kasper Larsen, for teaching and advising me in the area of Stochastic Control. I am especially grateful to my adviser Gordan Žitković for guiding, supporting, and encouraging me throughout my work on this research.

Martingale-Generated Control Structures and a Framework for the Dynamic Programming Principle

Publication No. ______________

Roman Fayvisovich, Ph.D.
The University of Texas at Austin, 2017

Supervisor: Gordan Žitković

This thesis constructs an abstract framework in which the dynamic programming principle (DPP) can be proven for a broad range of stochastic control problems. Using a distributional formulation of stochastic control, we prove the DPP for problems that optimize over sets of martingale measures. As an application, we use the classical martingale problem to prove the DPP for weak solutions of controlled diffusions, and use it show that the value function is a viscosity solution of the associated Hamilton-Jacobi-Bellman equation.

Table of Contents

Chapter 1

Introduction

There are two main goals in this thesis: create a framework, sufficiently abstract to encompass many stochastic control problems, in which the dynamic programming principle can be shown to hold, and then apply this framework to weak solutions of controlled diffusion SDEs.

The representation of stochastic control problems used in the abstract framework follows the distributional formulation similar to [11]. The system being modeled has an implicit state, which evolves over time according to a probability distribution that is chosen by the controller from an admissible set of measures. The set of admissible measures depends on the current state, and is analogous in spirit to the set of admissible controls typically used in strong formulations. Ultimately, the controller wants to choose the best measure to maximize an objective function at the infinite time horizon.

The dynamic programming principle (abbreviated as DPP from here on) is an old idea in which, intuitively speaking, an optimization problem is broken up into smaller pieces, where optimizing the overall problem is equivalent to optimizing the sub-problems. In our abstract framework for stochastic control, these sub-problems involve choosing one measure up to a stopping time, and then a selector of measures to be used after that time.

One of the major difficulties in building a framework where the DPP holds for general stochastic control problems is the conflict between the simple measurability structure needed to prove the DPP, and the messy completed filtration needed to define weak solutions of SDEs. To make this problem tractable, we use the classic martingale problem go between measures in the distributional framework and weak solutions of controlled diffusions.

Previous work on the DPP for stochastic control problems includes a rigorous and complete treatment of the discrete time case by Bertsekas and Shreve in [2]. Several authors have also worked on the continuous time case, including Bouchard and Touzi in [3], and Karoui and Tan in [5]. A different approach that skips the DPP in favor of directly getting viscocity solutions, is the stochastic Perron method by Bayraktar and Sîrbu in [1].

The approach taken in this thesis is to first build a minimal framework (T-space, Definition 2.1.1), which is designed to have an easy-to-use filtration that extends well to stopping times. This framework is then augmented with concatenation operators for both paths and measures (TC-space, Definition 2.5.1), which allows for a rigorous proof of an abstract DPP (Theorem 2.10.1). Using this robust foundation, a martingale oriented control structure (Section 3.1) is shown to satisfy the three main assumptions for the DPP: analyticity, concatenability, and disintegrability. This results in a martingale version of the DPP (Theorem 3.5.1). Finally, the classical martingale problem is used to bridge the gap between weak solutions and the martingale control structure, which gives a DPP for weak solutions (Theorem 4.1.3).

1.1 Notation and Conventions

1.1.1 Integration.

Both probabilistic $\mathbb{E}^{\mathbb{P}}[X]$ and analytic $\int G\, d\mu$ notation for integration will be used. The former will appear mostly in examples, and the latter in the abstract part.

1.1.2 Polish spaces.

Many of our probability spaces come with Polish (completely metrizable, separable) sample spaces and Borel probability measures or their completions. When the Polish structure is present, measurability will always refer to the associated Borel σ-alebra, denoted by $\mathsf{Borel}(\Omega)$.

1.1.3 Analytic sets.

A subset A of a Polish space Ω is called **analytic** if it can be realized as a projection of a Borel subset of $\Omega \times \mathbb{R}$ onto Ω. We remind the reader that analytic subsets of Polish spaces are closed under countable unions, intersections and products, but not necessarily under complements. It will be important for us that each analytic set is in the **universal** σ-algebra - denoted by $\mathsf{Univ}(\Omega)$ - i.e., the family of all sets which belong to the completion $(\mathsf{Borel}(\Omega))^*_\mu$ for each $\mu \in \mathsf{Prob}(\Omega)$. We refer the reader to [9] for all the necessary details concerning descriptive set theory.

1.1.4 Probability measures and kernels

The set of all probability measures on $\mathsf{Borel}(\Omega)$ is denoted by $\mathsf{Prob}(\Omega)$. We topologize $\mathsf{Prob}(\Omega)$ with the topology of (probabilist's) weak convergence. This way, $\mathsf{Prob}(\Omega)$ becomes a Polish space. The following well-known fact, proved in a standard way via the monotone-class theorem, will be used throughout without mention: Let U and V be Polish spaces and let $f : U \times V \to [0, \infty]$ be a Borel-measurable function. The map

$$U \times \mathsf{Prob}(V) \ni (x, \mu) \mapsto \mathbb{E}^{\mu}[f(x, \cdot)] = \int_V f(x, y)\, \mu(dy)$$

is Borel measurable.

A probability measure defined on $\mathsf{Borel}(\Omega)$ admits a natural extension to $\mathsf{Univ}(\Omega)$. Similarly, our kernels *will always be universally measurable.* More precisely, for Polish spaces $\Omega, \tilde{\Omega}$, a map $\nu : \Omega \times \mathsf{Borel}(\tilde{\Omega}) \to [0, 1]$ is called a **kernel** if $\nu(\omega, \cdot) \in \mathsf{Prob}(\tilde{\Omega})$ for each $\omega \in \Omega$ and $\nu(\cdot, B)$ is a universally-measurable map on Ω, for each $B \in \mathsf{Borel}(\tilde{\Omega})$. Depending on the situation we use both notations $\nu(\omega, \cdot)$ and ν_ω for the probability measure associated by ν to ω.

1.1.5 Standard Borel spaces

A standard Borel space is, by definition, a measurable space which admits a measurable bijection to a Borel subset of some $\mathbb{R}^n$, whose inverse is also measurable (a **bimesurable isomorphism**). All standard Borel spaces of the same cardinality are bimesurable isomorphic, and, so, each standard

Borel space can be given a complete and separable (Polish) metric so that the induced measurable structure matches the original one. With this in mind, we talk standard Borel spaces when only the measurable structure is relevant, and about Polish spaces when topological properties are required.

Chapter 2

An Abstract Framework for the DPP

Let the **time set** Time be either $[0, \infty)$ or $\mathbb{N}_0$. An overwhelming majority of applications will only use these two time sets, so we do not aim for greater generality. We do note that the results of this section will hold for more general time structures (such as intersections with $[0, \infty)$ of Borel-measurable additive subgroups of $\mathbb{R}$).

2.1 T-spaces (truncated spaces)

Definition 2.1.1 (T-spaces)**.** A filtered measurable space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \in \mathsf{Time}})$ is called a **T-space** (or a **truncated space**) if

1. Ω is a standard Borel space and $\mathcal{F} = \bigvee_{t \in \mathsf{Time}} \mathcal{F}_t$.
2. there exists a family $\{T_t\}_{t \in \mathsf{Time}}$ of maps $T_t : \Omega \to \Omega$ - called a **truncation** - such that
   (a) $(t, \omega) \mapsto T_t(\omega)$ is (jointly) measurable,
   (b) $T_t \circ T_s = T_{s \wedge t}$ for all $s, t \in \mathsf{Time}$, and
   (c) $\mathcal{F}_t = \sigma(T_t)$ for each $t \in \mathsf{Time}$.

2.2 First examples of T-spaces

As T-spaces are necessarily countably generated, not every filtered probability space can be endowed with the structure of a T-space. Nevertheless, as our examples below aim to show, many spaces used in stochastic analysis and optimal stochastic control are natural T-spaces.

In all of the examples below we take $\mathsf{Time} = [0, \infty)$. We leave it to the reader to make the necessary minor adjustments needed to translate all of the examples below to $\mathsf{Time} = \mathbb{N}_0$. Once we describe various natural constructions involving T-space in subsection 2.4 below, the reader will be able to produce many more examples.

2.2.1 The path space D_E

Let E be a Polish space, and let D_E denote the family of all càdlàg functions from Time to E. For $t \in \mathsf{Time}$, we define the truncation map $T_t : D_E \to D_E$ by

$$T_t(\omega)(s) = \omega(t \wedge s) \text{ for } s \in \mathsf{Time},$$

so that (2b) of Definition 2.1.1 holds. It is well-known that D_E is a Polish space under the Skorokhod topology. The map T_t is Skorokhod-continuous, and therefore, measurable. Therefore, as a Caratheodory function, $T : \mathsf{Time} \times \Omega \to \Omega$ is (jointly) measurable. The filtration $\mathcal{F}_t = \sigma(T_t), t \in \mathsf{Time}$ clearly conincides with the (raw) filtration generated by the coordinate maps $\omega \mapsto \omega(t)$.

2.2.2 The path space G_E and the related spaces

Analogous constructions can be performed on the space G_E of left-continuous and right limited paths (with the suitable variation of the Skorokhod topology) from Time to E. In the case $E = \mathbb{R}$, all three $D_{\mathbb{R}}$, $C_{\mathbb{R}}$ and $G_{\mathbb{R}}$ are subspaces of the (generalized) Skorokhod space of paths whose left and right limits exist at each point, but do not necessarily match the value. This space is also a Polish space (see [8, Section VII.6., p. 231] for details), and can be given a truncation structure (in several ways).

We will also have use for the space $\mathrm{Lip}_{\mathbb{R}}^{L,x_0}$ consisting of all functions $x : [0,\infty) \to \mathbb{R}$ such that $x(0) = x_0$ and $|x(t) - x(s)| \leq L\,|t-s|$ for all $s,t \in [0,\infty)$.

2.2.3 The space $\mathbb{L}^0_A$

Let A be a standard Borel space, let λ be the Lebesgue measure (or any other Radon measure) on $[0,\infty)$, and let $\hat{\lambda}$ denote an equivalent probability measure on $[0,\infty)$ (e.g., $\hat{\lambda}(dt) = e^{-t}\,\lambda(dt)$, when λ is the Lebesgue measure). We define $\mathbb{L}^0_A$ as the set of all λ-a.e.-equivalence classes of Borel functions $\alpha : [0,\infty) \to A$. Given a bimeasurable isomorphism $\phi : A \to [-1,1]$ (which exists thanks to the standard Borel property of A) we metrize $\mathbb{L}^0_A$ by

$$d(\alpha,\beta) = ||\phi(\alpha) - \phi(\beta)||_{\mathbb{L}^1(\hat{\lambda})}.$$

This way, $\Omega = \mathbb{L}^0_A$ becomes a Polish space and a natural truncation on it is defined by

$$T_t(\alpha) = \begin{cases} \alpha_u, & u < t \\ \phi^{-1}(0), & u \geq t. \end{cases}$$

We note that the equivalence class of the right-hand side depends on α only through its equivalence class, and that, while d and the induced Polish topology depend on the choice of ϕ and $\hat{\lambda}$, the resulting standard Borel structure does not. The choice of the particular ϕ makes it easy to show that T_t is jointly measurable; indeed, it will be continuous under d in both of its arguments.

2.2.4 Spaces of measures

For a metrized Polish space U, let $\mathcal{M}^{\#}(U)$ be the family of all boundedly-finite Borel measures on U, i.e,. those measures μ such that $\mu(B) < \infty$, as soon as B is a bounded Borel set. There exists a metric on $\mathcal{M}^{\#}(U)$, whose topology coincides with the topology of weak convergence when restricted to measures supported by a fixed bounded set (see [4, Section A2.6, p. 402] for the proof of this and other statements about the space $\mathcal{M}^{\#}(U)$ we make below). Under the full topology induced by this metric, called the $w^{\#}$-topology, $\mathcal{M}^{\#}(U)$ becomes a Polish space. Moreover, a sequence $\{\mu_n\}_{n\in\mathbb{N}}$ in $\mathcal{M}^{\#}(U)$ converges if and only if $\int f\, d\mu_n \to \int f \mu$ for each bounded and continuous function $f : \Omega \to \mathbb{R}$ which vanishes outside a bounded set. The Borel σ-algebra on $\mathcal{M}^{\#}(U)$ is generated by the evaluation maps $\mu \mapsto \mu(A)$, where A ranges over a family of all bounded Borel subsets of U. The subsets $\mathcal{M}^f(U)$ and $\mathcal{M}^p(U) = \mathsf{Prob}(U)$ of $\mathcal{M}^{\#}(U)$, consisting only of finite or probability measures (respectively), are easily seen to be Borel subsets of $\mathcal{M}^{\#}(\Omega)$, and, therefore, standard Borel spaces themselves.

For a Polish space E, we set $\Omega = \mathcal{M}^*(U)$, where $U = [0,\infty) \times E$ and $* \in \{\#, f, p\}$. The truncation maps are given by

$$\mu_{\leq t}(A) = \mu\Big(([0,t] \times E) \cap A\Big), \text{ for } t \in [0,\infty), A \in [0,\infty) \times E.$$

With the filtration generated by the maps T_t, it is clear that $\vee_t \mathcal{F}_t$ is the Borel σ-algebra on Ω. The only remaining property from Definition 2.1.1 is (2a), for which it is sufficient to note that for any boundedly supported function f we have $\int f \, d\mu_{\leq t} = \int f \mathbf{1}_{[0,t]\times E} \, d\mu$. Indeed, it follows that $(t,\mu) \mapsto \mu_{\leq t}$ is a Caratheodory functions as it is right continuous in t and measurable in μ.

2.2.5 Predictable truncations

In many the examples above, it is possible to define several different truncations on the same underlying Polish space. For example, in the case of the canonical space D_E, we may set

$$T'_t(\omega)(s) = \begin{cases} \omega_s, & s < t \\ \omega_{t-}, & s \geq t \end{cases}.$$

It is easily checked that T'_t is indeed, a truncation on D_E.

2.3 Truncating at stopping times

With the set of all stopping times is denoted by Stop, the index set for the family of truncation operators on a T-space can be extended to Stop by setting

$$T_\tau(\omega) = T_{\tau(\omega)}(\omega) \text{ for } \tau \in \mathsf{Stop} \text{ and } \omega \in \Omega,$$

where the convention that T_∞ is the identity map is used. The notation $T_\tau(\omega)$ will often be replaced by the less cumbersome (and more suggestive) $\omega_{\le\tau}$.

Proposition 2.3.1. *For all $t \in \mathsf{Time}$, $\omega \in \Omega$, $\tau, \kappa \in \mathsf{Stop}$ and we have*

1. *T_τ and T_κ are measurable maps on Ω and $T_\tau \circ T_\kappa = T_{\tau\wedge\kappa}$*
2. *$\sigma(T_\tau) = \{A \in \mathcal{F} : T_\tau^{-1}(A) = A\}$, and*

   $$\text{“}A \in \sigma(T_\tau)\text{”} \quad \textit{is equivalent to} \quad \text{“}\omega \in A \;\Leftrightarrow\; \omega_{\le\tau} \in A''$$

3. *$\tau(\omega) = \tau(T_\tau(\omega))$, and hence τ is $\sigma(T_\tau)$-measurable*
4. *$\sigma(T_\tau) = \mathcal{F}_\tau$, where $\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \le t\} \in \mathcal{F}_t, \forall t \in \mathsf{Time}\}$*
5. *Let $(S, \mathcal{S})$ be a standard Borel space. An $(\mathcal{F}, \mathcal{S})$-measurable map $Z : \Omega \to S$ is $(\mathcal{F}_\tau, \mathcal{S})$-measurable if and only if $Z \circ T_\tau = Z$.*

Proof.

1. Measurability of T_τ follows directly from the measurability of stopping times and the joint measurability of $(t, \omega) \mapsto T_t(\omega)$ on $(\mathsf{Time} \cup \{\infty\}) \times \Omega$. Applying Definition 2.1.1, part (2b) pointwise for $t = \tau(\omega)$ and $s = \kappa(\omega)$ gives $T_\tau \circ T_\kappa = T_{\tau\wedge\kappa}$.

2. By part (1) we have $T_\tau = T_\tau \circ T_\tau$ for each $\tau \in \mathsf{Stop}$, and so for any $A \in \mathcal{F}$, we have

$$A = T_\tau^{-1}(B) \text{ for some } B \in \mathcal{F} \quad \Leftrightarrow \quad A = T_\tau^{-1}(A).$$

Furthermore the condition $A = T_\tau^{-1}(A)$ is equivalent to:

$$\omega \in A \quad \Leftrightarrow \quad \omega_{\leq \tau} \in A$$

3. Fix $\omega \in \Omega$, let $t = \tau(\omega)$, and let $A = \{\tau = t\}$. Since $\tau \in \mathsf{Stop}$, then $A \in \mathcal{F}_t = \sigma(T_t)$. Combining part (2) with the fact that $\omega \in A$ implies $T_t(\omega) \in A$. Therefore:

$$\tau(T_{\tau(\omega)}(\omega)) = \tau(T_t(\omega)) = t = \tau(\omega)$$

4. For the forward inclusion, let $A \in \sigma(T_\tau)$. Thanks to (2) above, we have $A = T_\tau^{-1}(A)$. Therefore for all $t \in \mathsf{Time}$ we have:

$$\begin{aligned} A \cap \{\tau \leq t\} &= \{\omega \in \Omega \,:\, T_{\tau(\omega)}(\omega) \in A, \tau(\omega) \leq t\} \\ &= \{\omega \in \Omega \,:\, T_{\tau(\omega) \wedge t}(\omega) \in A, \tau(\omega) \leq t\} \\ &= T_{\tau \wedge t}^{-1}(A) \cap \{\tau \leq t\} \in \mathcal{F}_t, \end{aligned}$$

where we used the fact that $T_{\tau \wedge t} = T_{\tau \wedge t} \circ T_t$ is $\mathcal{F}_t$-measurable. Therefore $A \in \mathcal{F}_\tau$, and hence $\sigma(T_\tau) \subset \mathcal{F}_\tau$.

For the backward inclusion, let $A \in \mathcal{F}_\tau$. By part (2), it suffices to show:

$$\omega \in A \quad \Leftrightarrow \quad \omega_{\leq \tau} \in A$$

First suppose $\omega \in A$ and let $t = \tau(\omega)$. Since $A \in \mathcal{F}_\tau$, then $\omega \in A \cap \{\tau \leq t\} \in \mathcal{F}_t$. Applying (2) to $A \cap \{\tau \leq t\}$ gives $\omega_{\leq \tau} \in A \cap \{\tau \leq t\} \subset A$.

For the other direction, suppose $\omega_{\leq \tau} \in A$. By part (3) we have $\tau(\omega_{\leq \tau}) = \tau(\omega)$ and hence $\omega_{\leq \tau} \in A \cap \{\tau \leq t\} \in \mathcal{F}_t$. Applying (2) to $A \cap \{\tau \leq t\}$ gives $\omega \in A \cap \{\tau \leq t\} \subset A$.

5. If $Z = Z \circ T_\tau$, then Z is $\mathcal{F}_\tau$-measurable as a measurable transformation of the $\mathcal{F}_\tau$-measurable map T_τ. Conversely, if Z is $\mathcal{F}_\tau$-measurable, the standard Borel propery and the Doob-Dynkin lemma guarantee the existence of a measurable map $\zeta : \Omega \to S$ such that $Z = \zeta \circ T_\tau$. A composition with T_τ yields that

$$Z \circ T_\tau = \zeta \circ T_\tau \circ T_\tau = \zeta \circ T_\tau = Z. \qquad \square$$

2.4 Constructions on T-spaces

Next, we describe several natural notions and constructions on T-spaces, as well as various operations that produce new T-spaces from the old ones. For the remainder of this subsection, let $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \in \mathsf{Time}})$ and $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}} = \{\tilde{\mathcal{F}}_t\}_{t \in \mathsf{Time}})$ be two T-spaces, with truncations $\{T_t\}_{t \in \mathsf{Time}}$ and $\{\tilde{T}_t\}_{t \in \mathsf{Time}}$, respectively.

2.4.1 Structure-preserving maps

A measurable map $F : \Omega \to \tilde{\Omega}$ is said to be **non-anticipating** if it is $(\mathcal{F}_t, \tilde{\mathcal{F}}_t)$-measurable, i.e. $F^{-1}(\tilde{\mathcal{F}}_t) \subseteq \mathcal{F}_t$ for each $t \in \mathsf{Time}$. We have the following characterization using the truncation maps:

Proposition 2.4.1. *A measurable map* $F : (\Omega, \mathcal{F}) \to (\tilde{\Omega}, \tilde{\mathcal{F}})$ *is non-anticipating if and only if* $\tilde{T}_t \circ F \circ T_t = \tilde{T}_t \circ F$*, for all* $t \in \mathsf{Time}$*.*

Proof. By Proposition 2.3.1 part (2) we have $\tilde{\mathcal{F}}_t = \sigma(\tilde{T}_t) = \tilde{T}_t^{-1}(\tilde{\mathcal{F}})$, and by part (5) we have $\tilde{T}_t \circ F$ is $\mathcal{F}_t$-measurable if and only if $\tilde{T}_t \circ F \circ T_t = \tilde{T}_t \circ F$. Therefore for all $t \in \mathsf{Time}$:

$$\begin{aligned} F^{-1}(\tilde{\mathcal{F}}_t) \subset \mathcal{F}_t \quad &\Leftrightarrow \quad F^{-1}(\tilde{T}_t^{-1}(\tilde{\mathcal{F}})) \subset \mathcal{F}_t \\ &\Leftrightarrow \quad \tilde{T}_t \circ F \text{ is } \mathcal{F}_t\text{-measurable} \\ &\Leftrightarrow \quad \tilde{T}_t \circ F \circ T_t = \tilde{T}_t \circ F \end{aligned}$$

□

2.4.2 T-subspaces

We say that a T-space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}} = \{\tilde{\mathcal{F}}_t\}_{t\in\mathsf{Time}})$ is a T**-subspace** of $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}})$ if $\tilde{\Omega} \subseteq \Omega$ and $\tilde{\mathcal{F}}_t \subseteq \mathcal{F}_t$, for all $t \in \mathsf{Time}$. As the following result show, subsets preserved by truncation inherit a structure of a T-space:

Proposition 2.4.2. *Let* $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}})$ *be a T-space, and let* Ω' *be a measurable subset of* Ω *with the property that* $T_t(\Omega') \subseteq \Omega'$*, for all* $t \in \mathsf{Time}$*. Then the family* $\{T'_t\}_{t\in\mathsf{Time}}$ *given by* $T'_t = T_t|_{\Omega'}$*, is a truncation, and the filtered space* $(\Omega', \mathcal{F}', \{\mathcal{F}'_t\}_{t\in\mathsf{Time}})$*, given by* $\mathcal{F}' = \{B \in \mathcal{F} : B \subseteq \Omega'\}$*,* $\mathcal{F}_t = \sigma(T'_t)$*,* $t \in \mathsf{Time}$*, is a T-space and a subspace of* $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}})$*.*

Proof. Clearly $(\Omega', \mathcal{F}')$ is a subspace of $(\Omega, \mathcal{F})$. To satisfy Definition 2.5.1 of T-spaces, note that part (1) follows from the construction of Ω' and $\mathcal{F}'$, and the properties of part (2) are passed down from T to T'. $\square$

Example 2.4.3. Truncation operators on D_E leave invariant several important measurable subsets of D_E. Among the examples are

1. C_E, the family of all everywhere continuous elements of D_E
2. $D_E^{E_0}$, the family of paths in D_E which start from a point in E_0
3. D_{E^F}, the family of paths in D_E stopped once they hit the closed subset F of E
4. $FV\,(FV^+, FV^-)$, the family of all paths in $D_{\mathbb{R}}$ all of whose components are of finite variation (nondecreasing, nonincreasing)
5. $\mathrm{Lip}_{\mathbb{R}}^L$, the family of all Lipschitz continuous maps from $[0, \infty)$ to $\mathbb{R}$, with the Lipschitz constant at most L

2.4.3 Products

T-spaces behave well under products, too. Indeed, the standard Borel space $\hat{\Omega} = \Omega \times \tilde{\Omega}$ admits a natural truncation given by the family $\{\hat{T}_t\}_{t \in \mathsf{Time}}$ of maps on $\hat{\Omega}$ defined by

$$\tilde{T}_t(\omega, \tilde{\omega}) = (T_t(\omega), \tilde{T}_t(\tilde{\omega})). \tag{2.4.1}$$

The resulting T-space $\hat{\Omega}$, together with the natural filtration generated by $\{\hat{T}_t\}_{t\in\mathsf{Time}}$, is called the **product** of the truncated spaces Ω and $\tilde{\Omega}$. It is not difficult to see that the same construction can be applied to countable products of truncated spaces.

2.4.4 State maps

A measurable map $X : \Omega \to E$, where E is a Polish space is called a **state map**. Such maps define a class of progressively measurable E-valued stochastic processes on Ω via

$$X_t(\omega) = X(T_t(\omega)), t \in \mathsf{Time} \cup \{\infty\}, \omega \in \Omega$$

(where the convention $T_\infty(\omega) = \omega$ is used). We can also write X_τ for $X \circ T_\tau$ when $\tau \in \mathsf{Stop}$.

2.4.5 Actions on measures and kernels

For a probability measure $\mu \in \mathsf{Prob}(\Omega)$, and a stopping time $\tau \in \mathsf{Stop}$ we define the **truncated measure** $\mu_{\leq\tau}$ as the push-forward of μ via the truncation map T_τ.

Two analogous operations can be applied to kernels ν from Ω to Ω. We can truncate the second argument, leading to the **truncated kernel** $\nu_{\leq\tau}$, where, for each $\omega \in \Omega$, $\nu_{\leq\tau}(\omega, \cdot)$ is the truncation of the measure $\nu(\omega, \cdot)$, as above. On the other hand, we can define define the **restricted kernel** $\nu^{\leq\tau}$ by truncating in the first argument, i.e., by setting

$$\nu^{\le\tau}(\omega, B) = \nu(T_\tau(\omega), B).$$

That $\nu^{\le\tau}$ is, indeed, a kernel follows from the fact that a Borel measurable function (like T_τ) between to Polish spaces remains measurable under the pair of universal σ-algebras (see [2, Proposition 7.44, p. 172]).

2.5 TC-spaces (truncation-concatenation spaces)

Definition 2.5.1. A **truncation-concatenation space** (or a **TC-space**) is a truncation space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}})$ together with a jointly measurable map $* : \mathcal{C} \to \Omega$, called the **concatenation operator**, defined on a measurable subset $\mathcal{C} \subseteq \Omega\times\mathsf{Time}\times\Omega$, such that the following compatibility conditions hold:

1. for all $\omega, \omega' \in \Omega$ and $s, t \in \mathsf{Time}$ we have

$$(\omega, t, \omega') \in \mathcal{C} \Leftrightarrow (\omega_{\le t}, t, \omega') \in \mathcal{C} \Leftrightarrow (\omega, t, \omega'_{\le s}) \in \mathcal{C}. \tag{2.5.1}$$

2. if $(\omega, t, \omega') \in \mathcal{C}$, then, for all $s \in \mathsf{Time}$ we have

$$\omega *_t \omega' = \omega_{\le t} *_t \omega', \text{ as well as} \tag{2.5.2}$$

$$(\omega *_t \omega')_{\le s} = \begin{cases} \omega_{\le s}, & s \le t \\ \omega *_t \omega_{\le s-t}, & s > t \end{cases} \tag{2.5.3}$$

The action of the concatenation operator on the triplet $(\omega, t, \omega') \in \mathcal{C}$ is denoted by $\omega *_t \omega'$ and is usually interpreted as an element of Ω "obtained by following ω until time t, with ω' attached afterwards". The set $\mathcal{C}$ - the domain of $*$ - encodes a possible compatibility relation necessary for the concatenation to be possible. The set of all $\omega' \in \Omega$ such that $(\omega, t, \omega') \in \mathcal{C}$ is denoted by $\mathcal{C}_{\omega,t}$, and we say that ω' **is compatible with** ω **at** t if $\omega' \in \mathcal{C}_{\omega,t}$.

In many examples compatibility is established via a state map (as defined in subsection 2.4.4 above):

Definition 2.5.2. Given a TC-space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}})$ and a state map X, we say that the concatenation operator $*$

1. **factors through** X if

$$X_t(\omega) = X_0(\omega') \\ \Rightarrow (\omega, t, \omega') \in \mathcal{C}, \text{ and}$$

2. **is a factor of** X if

$$(\omega, t, \omega') \in \mathcal{C} \ \Rightarrow X_t(\omega) = X_0(\omega').$$

When needed, we also define $\omega *_\infty \omega' = \omega$, declaring, implicitly, any two elements of Ω compatible at $t = \infty$, so that $\mathcal{C}_{\omega,\infty} = \Omega$. This way, as in the case of the truncation spaces, the time-set Time can be extended to the set of all stopping times by setting:

$$\omega *_\tau \omega' = \omega *_{\tau(\omega)} \omega' \text{ for } \omega' \in \mathcal{C}_{\omega,\tau(\omega)}. \tag{2.5.4}$$

By Proposition 2.3.1, part (3), $\tau(\omega_{\leq\tau}) = \tau(\omega)$, and, so, the stopping-time analogue of (2.5.2) holds in TC spaces:

$$\omega *_\tau \omega' = \omega *_{\tau(\omega)} \omega' = \omega_{\leq\tau(\omega)} *_{\tau(\omega)} \omega' = \omega_{\leq\tau} *_{\tau(\omega_{\leq\tau})} \omega' = \omega_{\leq\tau} *_\tau \omega'.$$

2.6 Examples of TC-spaces

We go through the list of examples of T-spaces from subsection 2.2 and describe how a natural concatenation operator can be introduced.

2.6.1 Strict concatenation on path spaces D_E and C_E

We consider the space D_E with the truncation $\omega_{\leq t}(s) = \omega(s \wedge t)$. The **strict concatenation** operation $\bullet$ is given by

$$(\omega \bullet_t \omega')_s = \begin{cases} \omega(s), & s \leq t \\ \omega'(s-t), & s > t, \end{cases}$$

for $\omega, \omega' \in D_E$, where ω and ω' are considered t-compatible if and only if $\omega(t) = \omega'(0)$. To check that $\bullet$ is, indeed, a concatenation is straightforward, and we only remark that the joint measurability of $\bullet$ (in all three of its arguments) follows from the observation that, as a function of the inner argument t, it is right-continuous in the Skorokhod topology. When applied on its compatibility set $\mathcal{C}$, the operation $\bullet$ preserves continuity, so it can be used to define a concatenation operator on C_E, as well. Finally, it is straightforward that

$$X(\omega) := \liminf_{t \to \infty} \omega(t)$$

defines an $E = \bar{\mathbb{R}}$-valued state map with the property $X_t(\omega) = \omega(t)$ for $t \in \mathsf{Time}$ and such that the concatenation operator $\bullet$ factors through it.

Remark 2.6.1. We note that many subspaces of D_E are closed under the strict concatenation. The reader will easily check that all the spaces in Example 2.4.3 have this property, making them into TC-spaces themselves.

2.6.2 Adjusted concatenation on D_E and C_E

When E admits the structure of a linear space, we can define another concatenation operator on it, namely the **adjusted concatenation** operator $\star$. It is given for $\omega, \omega' \in D_E$ by

$$(\omega \star_t \omega')_s = \begin{cases} \omega(s), & s \leq t \\ \omega(t) + \omega'(s-t) - \omega'(0), & s > t, \end{cases}$$

with no restrictions on compatibility, i.e., with $\mathcal{C} = \Omega \times \mathsf{Time} \times \Omega$. It is clear that the strict and the adjusted concatenation operators agree on the compatibility set of $\bullet$, and that $*$ can be restricted to C_E without loosing any properties required of a concatenation.

2.6.3 Spaces of measures

We define the concatenation operator $*$ on the space $\Omega = \mathcal{M}^{\#}([0,\infty) \times E)$, described in subsection 2.2 as follows. For $\mu, \mu' \in \Omega$, we set

$$(\mu *_t \mu')(A) = \mu\Big(([0,t) \times E) \cap A\Big) + \mu'\Big(\big(([t,\infty) \times E) \cap A\big) - t\Big),$$

where $B - t = \{(x, s-t) \,:\, (x,s) \in B\}$, for $B \subseteq [t,\infty) \times E$. No compatibility restrictions are imposed. There should be no difficulty in checking that $*$ satisfies all defining properties of a concatenation, without. We also note that the same construction applies when $\mathcal{M}^{\#}$ is replaced by $\mathcal{M}^f$.

In the case when $\mathcal{M}^p$ is considered, the above operation does not preserve total mass. This cannot be fixed by restricting compatibility, but can be overcome by defining another concatenation operation as follows:

$$(\mu \tilde{*}_t \mu')(A) = \mu\Big(([0,t) \times E) \cap A\Big) + \Big(1 - \mu([0,t) \times E)\Big)\mu'\Big(\big(([t,\infty) \times E) \cap A\big) - t\Big),$$

2.6.4 $\mathbb{L}^0_A$ spaces

When the underlying measure λ is the Lebesgue measure, we usually concatenate $\mathbb{L}^0_A$ functions as follows:

$$(f *_t g)_u = \begin{cases} f_u, & u \le t \\ g_{u-t}, & u > t \end{cases},$$

with no compatibility restriction.

2.7 Concatenation of measures in TC-spaces

The ability to concatenate elements of Ω extends to probability measures and kernels on Ω. We say that a measure $\mu \in \mathsf{Prob}(\Omega)$ and a kernel $\nu \in \mathsf{Kern}(\Omega)$ on a TC-space are **compatible** at the stopping time τ if

$$\nu^{\le\tau}_{\omega}(\mathcal{C}_{\omega,\tau(\omega)}) = 1, \text{ for } \mu\text{-almost all } \omega.$$

When $*$ factors through a state map X, a sufficient condition for compatibility of $\mu \in \mathsf{Prob}(\Omega)$ and $\nu \in \mathsf{Kern}(\Omega)$ at τ is that

$$\nu^{\le\tau}_{\omega}\Big(X_0 = X_\tau(\omega)\Big) = 1, \text{ for } \mu\text{-almost all } \omega \text{ with } \tau(\omega) < \infty. \tag{2.7.1}$$

Using the convention, as above, that $\Omega \times \{\infty\} \times \Omega' \subseteq \mathcal{C}$, we also note that, given a stopping time τ, the set $\mathcal{C}_\tau = \{(\omega, \omega') : (\omega, \tau(\omega), \omega') \in \mathcal{C}\}$ is a pullback of the Borel set $\mathcal{C}$ via the measurable map $(\omega, \omega') \mapsto (\omega, \tau(\omega), \omega')$, and, therefore, itself measurable.

For $\mu \in \mathsf{Prob}(\Omega)$ and a τ-compatible kernel $\nu \in \mathsf{Kern}(\Omega)$ let $\mu \otimes \nu^{\le\tau} \in \mathsf{Prob}(\Omega \times \Omega)$ denote the product of μ and the τ-restriction of ν. The **concatenation** $\mu *_\tau \nu$ is then defined as the push-forward of this product via the measurable map $C_\tau \ni (\omega,\omega') \mapsto \omega *_{\tau(\omega)} \omega'$. We note that the compatibility relation introduced above implies that $\mu \otimes \nu^{\leq\tau}(C_\tau) = 1$, so that $\mu *_\tau \nu$ is, indeed, a probability measure. Moreover, we have

$$\begin{aligned}\int G(\omega)\,(\mu *_\tau \nu)(d\omega) &= \int G(\omega *_\tau \omega')\,(\mu \otimes \nu^{\leq\tau})(d\omega, d\omega') \\ &= \iint G(\omega *_\tau \omega')\,\nu_\omega^{\leq\tau}(d\omega')\,\mu(d\omega),\end{aligned}$$

for any sufficiently integrable random variable G on Ω. The compatibility condition (2.5.2) implies further that

$$\int G\,d(\mu *_\tau \nu) = \iint G(\omega_{\leq\tau} *_\tau \omega')\,\nu_\omega^{\leq\tau}(d\omega')\mu(d\omega) \tag{2.7.2}$$

$$= \iint G(\tilde{\omega} *_\tau \omega')\,\nu_\omega^{\leq\tau}(d\omega')\mu_{\leq\tau}(d\tilde{\omega}), \tag{2.7.3}$$

where $\mu_{\leq\tau}$ is the push forward of μ via T_τ.

2.7.1 Tail maps

Tail maps on TC-spaces will play an important role in the dynamic programming principle and will model payoffs associated to controlled processes.

Definition 2.7.1. A measurable map G from a TC-space to a measurable space S is called a **tail map** if $G(\omega *_t \omega') = G(\omega')$ for all $t \in \mathsf{Time}$, all $\omega \in \Omega$ and all $\omega' \in \mathcal{C}_{\omega,t}$. When $S = \mathbb{R}$ ($S = \bar{\mathbb{R}}$), a tail map is called a **tail random variable** (**extended tail random variable**).

The tail property of random variables extends readily to stopping times in the following form:

$$G(\omega *_\tau \omega') = \begin{cases} G(\omega'), & \tau(\omega) < \infty \\ G(\omega), & \tau(\omega) = \infty, \end{cases}$$

as long as ω' is compatible with ω at τ. Combining this expression with (2.7.2) we obtain the following equality, valid for each stopping time τ, probability $\mu \in \mathsf{Prob}(\Omega)$, a τ-compatible kernel $\nu \in \mathsf{Kern}(\Omega)$, and a sufficiently integrable tail random variable G:

$$\int G\, d(\mu *_\tau \nu) = \int \tilde{G}(\omega_{\leq \tau})\, \mu(d\omega), \tag{2.7.4}$$

where

$$\tilde{G}(\omega) = G(\omega)\mathbf{1}_{\{\tau(\omega)=\infty\}} + \int G(\omega')\, \nu_\omega(d\omega')\mathbf{1}_{\{\tau(\omega)<\infty\}}.$$

2.8 Control Structures

A map $\mathcal{P} : A \to 2^B$, where 2^B denotes the power-set of B is called a **correspondence** from A to B, denoted by $f : A \twoheadrightarrow B$. Its **graph** $\Gamma(f) \subseteq A \times B$ is given by $\Gamma(f) = \{(a, b) : a \in A, b \in f(a)\}$, and its **image** by $\mathrm{Im}(f) = \cup_{a \in A} f(a)$. A correspondence is said to be **non-empty-valued** if $f(a) \neq \emptyset$ for all $a \in A$.

Definition 2.8.1. A non-empty-valued correspondence $\mathcal{P} : \Omega \twoheadrightarrow \mathsf{Prob}(\Omega)$, on a measurable space Ω is called a **control structure**.

Given a control structure $\mathcal{P}$, a universally measurable random variable G is said to be $\mathcal{P}$**-lower semi-integrable**, denoted by $G \in \mathcal{L}^{0-1}(\mathcal{P})$, if $G^- \in$

$\mathcal{L}^1(\mu)$ for each $\mu \in \mathrm{Im}\mathcal{P}$. To each control structure $\mathcal{P}$ and each $G \in \mathcal{L}^{0-1}(\mathcal{P})$ we associate the **value function** $v : \Omega \to (-\infty, \infty]$, given by

$$v(\omega) = \sup_{\mu \in \mathcal{P}(\omega)} \int G \, d\mu. \tag{2.8.1}$$

2.9 Three key properties

As we will see below, there are three key properties that control structures must satisfy in order for our main results to apply. one:

Definition 2.9.1. A control structure $\mathcal{P}$ on standard Borel space Ω is called

1. **analytic** if its graph $\Gamma(\mathcal{P})$ is an analytic subset of the (standard Borel) space $\Omega \times \mathsf{Prob}(\Omega)$.

A control structure $\mathcal{P}$ defined on a TC space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \in \mathsf{Time}})$ is said to be

2. **concatenable** if for each $\omega \in \Omega$, $\mu \in \mathcal{P}(\omega)$, $\nu \in \mathcal{S}(\mathcal{P})$, and each stopping time τ, ν is τ-compatible with μ and

$$\mu *_\tau \nu \in \mathcal{P}(\omega).$$

3. **disintegrable** if for each $\omega \in \Omega$, $\mu \in \mathcal{P}(\omega)$ and a stopping time τ there exists $\nu \in \mathcal{S}(\mathcal{P})$ such that ν is μ-compatible at τ and

$$\mu = \mu *_\tau \nu.$$

We state for completeness the following result which will be used in the sequel, and the proof of which follows almost verbatim the argument in [11, Theorem 2.4, part 1., p. 1605], which, in turn, is a reformulation of the standard argument available, for example, in [2]. We recall that a **universally measurable selector** (or, simply, a **selector**) is a (universally measurable) kernel form Ω to $\mathsf{Prob}(\Omega)$ with the property that $\nu(\omega) \in \mathcal{P}(\omega)$, for each ω; the family of all selectors is denoted by $\mathcal{S}(\mathcal{P})$. We also remind the reader of the convention $+\infty - \varepsilon = 1/\varepsilon$, for $\varepsilon > 0$.

Proposition 2.9.2 (Universal measurability of value functions)**.** *Suppose that Ω is a standard Borel space, $\mathcal{P}$ an analytic control structure, $G : \Omega \to [-\infty, \infty]$ a Borel measurable function, and v the associated value function, given by* (2.8.1)*. Then v is universally measurable and for each $\varepsilon > 0$ there exists a (universally measurable) selector $\nu^\varepsilon \in \mathcal{S}(\mathcal{P})$ such that*

$$v(\omega) - \varepsilon \leq \int G \, d\nu^\varepsilon_\omega, \quad \textit{for all } \omega \in \Omega.$$

2.10 An abstract version of the dynamic programming principle

We are ready to state the most abstract version of the DPP that holds in our setting. A more directly applicable - and more familiar-looking - version, based on the notion of a state map will be given below. The ideas in the proof are entirely standard. In fact, our setting is constructed as the most flexible one where this proof can be applied. We provide the details in our setting for the reader's convenience.

Theorem 2.10.1 (DPP)**.** *Let $\mathcal{P}$ be an analytic control structure on a TC space Ω, $G \in \mathcal{L}^{0-1}(\mathcal{P})$ a tail random variable, and v the associated value function, given by* (2.8.1)*. Then,*

1. *If $\mathcal{P}$ is concatenable, then, for each $\omega \in \Omega$ and each stopping time τ we have*

$$v(\omega) \geq \sup_{\mu \in \mathcal{P}(\omega)} \int v \circ T_\tau \mathbf{1}_{\{\tau<\infty\}} + G\mathbf{1}_{\{\tau=\infty\}}\, d\mu \tag{2.10.1}$$

2. *If $\mathcal{P}$ is disintegrable, then, for each $\omega \in \Omega$ and each stopping time τ we have*

$$v(\omega) \leq \sup_{\mu \in \mathcal{P}(\omega)} \int v \circ T_\tau \mathbf{1}_{\{\tau<\infty\}} + G\mathbf{1}_{\{\tau=\infty\}}\, d\mu \tag{2.10.2}$$

Proof. Suppose, first, that $\mathcal{P}$ is concatenable and pick $\omega \in \Omega$, $\mu \in \mathcal{P}(\omega)$ and a stopping time τ. Given $\varepsilon > 0$, Proposition 2.9.2 guarantees the existence of an ε-optimizing selector ν^ε, i.e., such that $v^\varepsilon(\omega) := \int G\, d\nu^\varepsilon_\omega \geq v(\omega) - \varepsilon$, for each $\omega \in \Omega$. We construct the measure μ' by concatenating μ and ν^ε at τ. The assumption of concatenability implies that they are compatible and that $\mu' \in \mathcal{P}(\omega)$. Therefore,

$$\begin{aligned} v(\omega) &\geq \int G\, d\mu' = \int G\, d(\mu *_\tau \nu^\varepsilon) = \iint G(\omega *_\tau \omega')\, (\nu^\varepsilon)^{\leq\tau}_\omega(d\omega')\, \mu(d\omega) \\ &= \iint G(\omega)\mathbf{1}_{\{\tau(\omega)=\infty\}} + G(\omega')\mathbf{1}_{\{\tau(\omega)<\infty\}}(\nu^\varepsilon)^{\leq\tau}_\omega(d\omega')\mu(d\omega) \\ &\geq \int G(\omega)\mathbf{1}_{\{\tau(\omega)=\infty\}} + (v(\omega_{\leq\tau}) - \varepsilon)\mathbf{1}_{\{\tau(\omega)<\infty\}}\, \mu(d\omega), \end{aligned}$$

which implies (2.10.1).

In the disintegrable case, we pick $\varepsilon > 0$, $\omega \in \Omega$, $\tau \in \mathsf{Stop}$ and choose $\mu^\varepsilon \in \mathcal{P}(\omega)$ such that $v(\omega) - \varepsilon \leq \int G\, d\mu^\varepsilon$. By disintegrability, we can write $\mu^\varepsilon = \mu^\varepsilon *_\tau \nu$ for some $\nu \in \mathcal{S}(\mathcal{P})$, and so

$$\begin{aligned} v(\omega) - \varepsilon &\leq \int G\, d(\mu^\varepsilon *_\tau \nu) \\ &= \int G(\omega)\mathbf{1}_{\{\tau=\infty\}} + \mathbf{1}_{\{\tau<\infty\}}\Big(\int G(\omega')\nu_\omega^{\leq\tau}(d\omega')\Big)\, \mu(d\omega) \\ &\leq \int G(\omega)\mathbf{1}_{\{\tau=\infty\}} + v(\omega_{\leq\tau})\mathbf{1}_{\{\tau<\infty\}}\, \mu(d\omega). \end{aligned}$$ □

2.10.1 State maps and factoring

We remind the reader that, as defined in subsection 2.4.4, a state map $X : \Omega \to E$ is simply a measurable map from a T-space to a Polish space E, and that X_τ is a shortcut for $X \circ T_\tau$, for $\tau \in \mathsf{Stop}$. Just like (concatenation) compatibility may factor through X, so can a control structure:

Definition 2.10.2. A control structure $\mathcal{P}$ on Ω is said to **factor through** a state map X if there exists a correspondence $\bar{\mathcal{P}} : E \twoheadrightarrow \mathsf{Prob}(\Omega)$ such that $\mathcal{P}(\omega) = \bar{\mathcal{P}}(X(\omega)) \subseteq \mathsf{Prob}(\Omega)$, i.e., the following diagram commutes:

$$\begin{array}{ccc} \Omega & \xrightarrow{\;X\;} & E \\ \downarrow{\scriptstyle \mathcal{P}} & \swarrow{\scriptstyle \bar{\mathcal{P}}} & \\ \mathsf{Prob}(\Omega) & & \end{array} \tag{2.10.3}$$

A very simple, but important, consequence of the existence of a state map through which the control structure $\mathcal{P}$ factors is that, in that case, v factors through it, as well. Indeed, the function $\bar{v} : E \to [-\infty, \infty]$, given by $\bar{v}(x) = \sup_{\mu \in \bar{\mathcal{P}}(x)} \int G \, d\mu$, then has the property that $\bar{v}(X(\omega)) = v(\omega)$ and, under the conditions of Theorem 2.10.1, satisfies

$$\bar{v}(x) \leq (\geq) \sup_{\mu \in \mathcal{P}(x)} \int \left(\bar{v}(X_\tau) \mathbf{1}_{\{\tau < \infty\}} + G \mathbf{1}_{\{\tau = \infty\}} \right) d\mu$$

for all $x \in \mathrm{Im} X$, and all stopping times $\tau \in \mathsf{Stop}$.

Chapter 3

Martingale-Generated Control Structures

The next task is so take the abstraction level down a notch and study a class of control structures defined via a family of martingale conditions. These structures generalize the standard martingale formulation in the theory of stochastic optimal control and are defined via a family of structure-preserving maps into the model space space $D^0_{\mathbb{R}}$ of $\mathbb{R}$-valued càdlàg paths $x : \mathsf{Time} \to \mathbb{R}$ with $x(0) = 0$.

3.1 Canonical local martingale measures

With the T-space structure of $D_{\mathbb{R}}$ described in subsection 2.2, each non-anticipating map F from a T-space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \in \mathsf{Time}})$ into $D_{\mathbb{R}}$ induces a sequence $\{F^n\}$ of non-anticipating maps

$$F^n_t = F_{\tau^F_n \wedge t} \text{ where } \tau^F_n(\omega) = \inf\{t \geq 0 : |F_t(\omega)| \geq n\} \wedge n. \qquad (3.1.1)$$

When the choice of F is evident from context, we may drop the superscript and write $\tau_n = \tau^F_n$.

Definition 3.1.1. A probability measure $\mu \in \mathsf{Prob}(\Omega)$ is said to be a **canonical local-martingale probability for** F if the stochastic process $\{F^n_t(\cdot)\}_{t \in \mathsf{Time}}$

is a martingale under $(\mu, \mathbb{F})$ for each $n \in \mathbb{N}$. The set of all canonical local martingale probabilities for F is denoted by $\mathcal{M}^{F,loc}$.

Remark 3.1.2. The notion of a *canonical* local martingale differs from the standard notion of a local martingale in that it requires that the reducing sequence takes a particular form, namely that of the sequence of space-time exit times. This requirement is nontrivial, as it is known that there are local martingales that cannot be reduced by this particular sequence (see [10, Lemme 2.1., p. 57]). On the other hand, this notion suffices for applications, as we will be dealing with continuous processes or processes with bounded jumps; in those classes all local martingales are canonical in our sense.

With the notion of a canonical local martingale probability under our belt, we can define a large class of control structures. Housed on T-spaces, they need two ingredients to be specified: 1) a family of $\mathcal{D}$ of non-anticipating maps from $\Omega \to D_{\mathbb{R}}$, and 2) a state map X from Ω to a Polish space E. Once these are specified, for $x \in E$ we define

$$\boldsymbol{P}(x) = \bigcap_{F \in \mathcal{D}} \mathcal{M}^{F,loc} \cap \Big\{ \mu \in \mathsf{Prob}(\Omega) \, : \, X_0 = x, \; \mu\text{-a.s.} \Big\}, \tag{3.1.2}$$

where, as usual, X_0 is the shortcut for $X \circ T_0$. The $(\mathcal{D}, X)$-**generated control structure** $\mathcal{P} = \mathcal{P}(\mathcal{D}, X) : \Omega \twoheadrightarrow \mathsf{Prob}(\Omega)$ is then defined by

$$\mathcal{P}(\omega) = \boldsymbol{P}(X(\omega)) \text{ for } \omega \in \Omega,$$

so that it naturally factors through X.

3.2 Sufficient conditions for analyticity

The ubiquitous Polish-space structure woven into all the ingredients of our setup makes it possible to give widely met sufficient conditions on the family $\mathcal{D}$ such that the resulting $(\mathcal{D}, X)$-structure becomes analytic. The countability condition we impose on $\mathcal{D}$ is not the weakest possible, but since it holds in most relevant examples, we only comment on some possible routes towards establishing weaker versions in Remark 3.2.3 below.

Proposition 3.2.1. *Let $\mathcal{D}$ be a countable family of nonanticipating maps from a T-space Ω to $D_{\mathbb{R}}$ and let $X : \Omega \to E$ be a state map. Then the $(\mathcal{D}, X)$-generated control structure $\mathcal{P}$ is analytic.*

The proof is based on a modification of [11, Lemma 3.6, p. 1611], where

$$\mathsf{QStop} = \left\{ q\mathbf{1}_A + r\mathbf{1}_{A^c} \; : \; q \leq r \in \mathsf{QTime}, A \in \Pi_q \right\}$$

with QTime denoting a countable dense set in Time, and $\{\Pi_q\}_{q \in \mathsf{QTime}}$ a collection of countable π-systems such that $\sigma(\Pi_q) = \mathcal{F}_q$ for all $q \in \mathsf{QTime}$. The exact choice of QTime or $\{\Pi_q\}_{q \in \mathsf{QTime}}$ is unimportant, as long as it is fixed throughout.

Lemma 3.2.2. *For each non-anticipative map F, we have*

$$\mathcal{M}^{F,loc} = \bigcap \left\{ \mu \in \mathsf{Prob}(\Omega) \; : \; F^n_q, F^n_r \in \mathbb{L}^1(\mu) \text{ and } \mathbb{E}^\mu[F^n_r \mathbf{1}_A] = \mathbb{E}^\mu[F^n_q \mathbf{1}_A] \right\} \tag{3.2.1}$$

where the intersection is taken over all $n \in \mathbb{N}$, $q < r \in \mathsf{QTime}$ and $A \in \Pi_q$.

Proof. The inclusion $\mathcal{M}^{F,loc} \subseteq \ldots$ is straightforward. Conversely, let $\mu \in \mathsf{Prob}(\Omega)$ be an element of the right-hand side of (3.2.1). We first show that $\mu \in \mathcal{M}^{F^n}_{\mathsf{QTime}}$, where $\mathcal{M}^{F^n}_{\mathsf{QTime}}$ denotes the set of all $\mu \in \mathsf{Prob}(\Omega)$ with the property that $\{F^n_t\}_{t \in \mathsf{QTime}}$ is a μ-martingale with respect to $\{\mathcal{F}_t\}_{t \in \mathsf{QTime}}$. That is an immediate consequence of the equalities of expectations under μ on the right-hand-side of (3.2.1). Considered over all $A \in \Pi_q$, with $q < r \in \mathsf{QTime}$, they amount to $\mathbb{E}^\mu[F^n_r|\mathcal{F}_q] = F^n_q$, a.s., by π-λ-theorem.

It remains to argue that F^n is a μ-martingale on entire Time. Assuming, without loss of generality, that $\mathsf{Time} = [0, \infty)$, we start by picking $s \in \mathsf{Time} \setminus \mathsf{QTime}$ and $r \in \mathsf{QTime}$ with $r > s$. The backward martingale convergence theorem implies that

$$\mathbb{E}^\mu[F^n_r|\mathcal{F}_{s+}] = F^n_s, \ \mu\text{-a.s.}$$

Since F^n is non-anticipative, F^n_s is $\mathcal{F}_s$-measurable and we may replace $\mathcal{F}_{s+}$ by $\mathcal{F}_s$ in the equality above. Finally, for $t \in \mathsf{Time}$ with $t > s$, we approximate F^n_t by a sequence $\{F^n_{r_m}\}_{m \in \mathbb{N}}$ with $r_m \searrow t$ and $r_m \in \mathsf{QTime}$, to conclude that F^n is, indeed, a martingale under μ. □

Proof of Proposition 3.2.1. For each $r \in \mathsf{Time}$, the coordinate maps are Borel measurable on $D_\mathbb{R}$ and, so, $\mu \mapsto \mathbb{E}^\mu[F_r \mathbf{1}_A]$ is Borel on Ω. It is easy to see that the family of probability measure under which a given real-valued Borel map is integrable is also a Borel set, so it follows that $\mathcal{M}^{F,loc}$ is Borel for each F. The countability of $\mathcal{D}$ guarantees that $\cap_{F \in \mathcal{D}} \mathcal{M}^{F,loc}$, as well. Finally, the graph of $\mathcal{P}$ is analytic (in fact Borel) as it is given as an intersection of a Borel sets

$$\Gamma(\mathcal{P}) = \Big\{(\omega,\mu)\ :\ \mu\Big(X_0 = X_0(\omega)\Big) = 1\Big\} \cap \Big(\Omega \times \bigcap_{F\in\mathcal{D}} \mathcal{M}^F\Big). \qquad \square$$

Remark 3.2.3. When $\mathcal{D}$ is not countable, the set $\cap_{F\in\mathcal{D}}\mathcal{M}^{F,loc}$ is not necessarily Borel measurable (or even analytic) in general. The situation is somewhat more pleasant when $\mathcal{D}$ admits a structure of a Borel space with the property that the maps

$$\mathcal{D} \ni F \mapsto \mathbb{E}^{\mu}[F_r], r \in \mathsf{Time},$$

are measurable for each probability measure $\mu \in \mathsf{Prob}(\Omega)$. In that case, the intersection $\cap_{F\in\mathcal{D}}\mathcal{M}^{F,loc}$ can be represented as a *co-projection*

$$\{\mu \in \mathsf{Prob}(\Omega)\ :\ \forall\, F \in \mathcal{D},\ (F,\mu) \in \mathcal{M}\}$$

of the Borel set $\mathcal{M} = \{(F,\mu) \in \mathcal{D} \times \mathsf{Prob}(\Omega)\ :\ \mu \in \mathcal{M}^{F,loc}\}$. Unlike projections, the images of co-projections are co-analytic, but not necessarily analytic sets. Not everything is lost, however, as we usually know a great deal more about the set $\mathcal{M}$, other than the fact that it is a Borel set. Indeed, the countable case of Proposition (3.2.1) corresponds to the measurable-selection theorem of Lusin for sets with countable sections (see [9, Theorem 5.7.2, p. 205]. On the other side of the spectrum are measurable selection theorems with large sections (see Section 5.8 in [9]), which can be used for certain uncountable $\mathcal{D}$.

3.3 Sufficient conditions for concatenability

Having discussed analyticity, we turn to the second major assumption of our abstract DPP theorem, namely concatenability. It is not hard to see that without additional requirements on $\mathcal{D}$, no $(\mathcal{D}, X)$-generated control structure should be expected to be concatenable. A natural requirement, as we will see below, is that the maps F preserve the structure of TC-spaces.

Definition 3.3.1. A measurable map $F : \Omega \to \tilde{\Omega}$ between two TC-spaces, with concatenation operators $*$ and $\tilde{*}$ (and compatibility sets $\mathcal{C}$ and $\tilde{\mathcal{C}}$) is called a **TC-morphism** if

1. F is non-anticipating, and
2. for all $t \in \mathsf{Time}$, and all $\omega, \omega' \in \Omega$ with $\omega' \in \mathcal{C}_{\omega,t}$ we have $F(\omega') \in \tilde{\mathcal{C}}_{F(\omega),t}$ and
$$F(\omega *_t \omega') = F(\omega) \,\tilde{*}_t\, F(\omega').$$

TC-morphisms into $D^0_{\mathbb{R}}$ are especially important for martingale-generated structures. We remind the reader that $D_{\mathbb{R}}$ comes with two different, natural, concatenations, namely, the strict one $(\bullet)$ and the adjusted one $(\star)$. We will only work with the adjusted one in this section, but, in order to avoid any confusion, we will write $(D_{\mathbb{R}}, \star)$ and $(D^0_{\mathbb{R}}, \star)$ throughout.

Definition 3.3.2. A map $F : \Omega \to D_{\mathbb{R}}$ is said to be **canonically locally bounded** if there exists a sequence $\{M_n\}_{n\in\mathbb{N}}$ of positive constants so that

$$|F^n(\omega)_t| \leq M_n \text{ for all } \omega \in \Omega, t \in \mathsf{Time}. \tag{3.3.1}$$

A simple sufficient condition for canonical local boundedness is that the jumps of F (when seen as a stochastic process on Ω) are uniformly bounded.

Proposition 3.3.3. *Let $\mathcal{D}$ a family of canonically locally bounded TC-morphisms into $(D^0_{\mathbb{R}}, \star)$, and let X be a state map. Then the $(\mathcal{D}, X)$-generated control structure $\mathcal{P}$ is closed under concatenation.*

The proof is based on the several lemmas. We omit the straightforward proof of the first one.

Lemma 3.3.4. *Suppose that F is a TC-morphism into $(D_{\mathbb{R}}, \star)$. For all stopping times κ we have*

$$F_{\kappa+t}(\omega *_\kappa \omega') - F_{\kappa+s}(\omega *_\kappa \omega') = F_t(\omega') - F_s(\omega')$$

for all $\omega \in \Omega$ with $\kappa(\omega) < \infty$, $\omega' \in \mathcal{C}_{\omega,\kappa(\omega)}$ and all $s, t \in \mathsf{Time}$.

Our second lemma gives a convenient characterization of canonical local martingales. We use Stop, as in the case of T-spaces, to denote the set of all Time-valued (raw) stopping times. We also write $Y^n = Y^{\tau_n}$, where $\tau_n = \inf\{t \geq 0 : |Y_t| \geq n\} \wedge n$, and note that all sampled values of Y in the statement are well-defined thanks to the fact that each Y^n is constant after $t = n$.

Lemma 3.3.5. *Let $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}}, \mathbb{P})$ be a filtered probability space, $\{Y_t\}_{t\in\mathsf{Time}}$ a càdlàg and adapted process, and κ a stopping time with $Y^n_\kappa \in \mathbb{L}^1$ for each $n \in \mathbb{N}$. Then, the following two statements are equivalent*

1. *Y is a canonical local martingale.*

2. $G \in \mathbb{L}^1$ *and* $\mathbb{E}[G] = 0$ *for all*

$$G \in \bigcup_{n\in\mathbb{N}} \mathcal{X}_n^{\leq\kappa}(Y) \cup \mathcal{X}_n^{\geq\kappa}(Y),$$

where the countable sets $\mathcal{X}_n^{\leq\kappa}$ *and* $\mathcal{X}_n^{\geq\kappa}$ *are given by*

$$\mathcal{X}_n^{\leq\kappa}(Y) = \Big\{ Y^n_{\tau\wedge\kappa} - Y^n_\kappa \,:\, \tau \in \mathsf{QStop} \Big\},$$
$$\mathcal{X}_n^{\geq\kappa}(Y) = \Big\{ Y^n_{\tau\vee\kappa} - Y^n_\kappa \,:\, \tau \in \mathsf{QStop} \Big\}.$$

Proof. (1) $\Rightarrow$ (2) Assuming that Y is a canonical local martingale, each Y^n is martingale constant after $t = n$, and, so, a uniformly-integrable martingale. Stopping times in QStop are bounded, so, by the optional sampling theorem, (2) holds.

(2) $\Rightarrow$ (1) Suppose that (2) holds and that $n \in \mathbb{N}$ is fixed. We take the advantage of the fact that Y is càdlàg to conclude (as in the proof of Lemma 3.2.2) that it suffices to show that Y^n is a martingale on QTime. For that, in turn, we choose $\tau \in \mathsf{QStop}$, so that $\tau = p\mathbf{1}_A + q\mathbf{1}_{A^c}$ for some $p \leq q \in \mathsf{QTime}$ and $A \in \Pi_p$ and note that

$$Y^n_\tau - Y^n_\kappa = \Big(Y^n_{\tau\wedge\kappa} - Y^n_\kappa \Big) + \Big(Y^n_{\tau\vee\kappa} - Y^n_\kappa \Big).$$

Since $Y^n_{\tau\wedge\kappa} - Y^n_\kappa \in \mathcal{X}^{\leq\kappa}$ $Y^n_{\tau\vee\kappa} - Y^n_\kappa \in \mathcal{X}^{\geq\kappa}$ and $Y^n_\kappa \in \mathbb{L}^1$, we conclude that $Y^n_\tau \in \mathbb{L}^1$ and that $\mathbb{E}[Y^n_\tau] = \mathbb{E}[Y^n_\kappa]$. It follows that the value of $\mathbb{E}[Y^n_\tau]$ does not depend on the choice of τ, making Y^n into a martingale. □

Lemma 3.3.6. *Let Ω be a TC-space and $\kappa, \tau \in \mathsf{Stop}$ such that $\kappa \leq \tau$. For $\omega \in \Omega$ we define τ'_ω by*

$$\tau'_\omega(\omega') = \begin{cases} \tau(\omega *_\kappa \omega') - \kappa(\omega), & \kappa(\omega) < \infty \text{ and } \omega' \in \mathcal{C}_{\omega,\kappa(\omega)} \\ +\infty, & \text{otherwise,} \end{cases}$$

Then the map $(\omega, \omega') \mapsto \tau'_\omega(\omega')$ is jointly measurable, $\tau'_\omega \in \mathsf{Stop}$ for any fixed $\omega \in \Omega$, and $\tau(\omega *_\kappa \omega') = \kappa(\omega) + \tau'_\omega(\omega')$.

Proof. By construction, we clearly have $\tau(\omega *_\kappa \omega') = \kappa(\omega) + \tau'_\omega(\omega')$. With the convention that $\tau(\omega *_\kappa \omega') - \kappa(\omega) = \infty$ when $\kappa(\omega) = \infty$, we note that τ' can be expressed as:

$$\tau'_\omega(\omega') = (+\infty)\mathbf{1}_{\mathcal{C}^c}(\omega, \kappa(\omega), \omega') + (\tau(\omega *_\kappa \omega') - \kappa(\omega))\mathbf{1}_{\mathcal{C}}(\omega, \kappa(\omega), \omega')$$

and is hence jointly measurable. It remains to argue that τ'_ω is a stopping time. We fix $\omega \in \Omega$ with $k = \kappa(\omega) < \infty$, and for $s \in \mathsf{Time}$ define

$$A = \{\omega' \in \Omega \, : \, \tau'(\omega') \leq s\} = \{\omega' \in \mathcal{C}_{\omega,k} \, : \, \tau(\omega *_k \omega') \leq s + k\}.$$

By Proposition 2.3.1, part (1), it will suffice to show that $T_s^{-1}(A) = A$, i.e., for $\omega' \in \Omega$ we have $(a) \Leftrightarrow (b)$, where

(a) $\omega' \in \mathcal{C}_{\omega,k}$ and $\tau(\omega *_k \omega') \leq s + k$, and

(b) $(\omega')_{\leq s} \in \mathcal{C}_{\omega,k}$ and $\tau(\omega *_k (\omega'_{\leq s})) \leq s + k$.

The first, compatibility-related, parts of statements of (a) and (b) are equivalent to each other by the assumptions in (2.5.1) of Definition 2.5.1. To deal with the inequalities involving τ we use Proposition 2.3.1, part (2), as well as the assumption 2.5.3 of Definition 2.5.1 to conclude that

$$\begin{aligned}\tau\Big(\omega *_k (\omega'_{\leq s})\Big) \leq s+k &\Leftrightarrow \tau\Big(\big(\omega *_k (\omega'_{\leq s})\big)_{\leq s+k}\Big) \leq s+k \\ &\Leftrightarrow \tau\Big((\omega *_k \omega')_{\leq s+k}\Big) \leq s+k \\ &\Leftrightarrow \tau\Big(\omega *_k \omega'\Big) \leq s+k. \qquad \square\end{aligned}$$

Proof of Proposition 3.3.3. Let $\mathcal{P}$ be the $(\mathcal{D}, X)$-generated control structure as in the statement, and let $\omega_0 \in \Omega$, $\mu \in \mathcal{P}(\omega_0)$, a kernel $\nu \in \mathcal{S}(\mathcal{P})$ and a stopping time κ be given.

First, we argue that ν is κ-compatible with μ. By the definition of $\mathcal{P}$, we have $\nu_\omega(X_0 = X(\omega)) = 1$ for each $\omega \in \Omega$. After a composition with T_κ, we get $\nu_\omega^{\leq\kappa}(X_0 = X_\kappa(\omega)) = 1$ for each $\omega \in \Omega$, which implies compatibility, according to the criterion of (2.7.1).

Next, we show that $\mu' = \mu *_\kappa \nu \in \mathcal{P}(\omega_0)$. Part (2) of Definition 2.5.1 makes it clear that for $x = X_0(\omega_0)$ we have $\mu'(X_0 = x) = 1$. Therefore, we need to argue that $\mu' \in \mathcal{M}^{F,loc}$, for each $F \in \mathcal{D}$. By Lemma 3.3.5, this is equivalent to checking $\int G\, d(\mu *_\kappa \nu) = 0$ for all $G \in \cup_{n\in\mathbb{N}} \mathcal{X}_n^{\leq\kappa}(F) \cup \mathcal{X}_n^{\geq\kappa}(F)$. We fix $n \in \mathbb{N}$ and treat the two cases separately:

1. $G \in \mathcal{X}_n^{\leq\kappa}(F)$: In this case there exists $\tau \in \mathsf{QStop}$, such that $G(\omega) = F^n_{(\tau\wedge\kappa)(\omega)}(\omega) - F^n_{\kappa(\omega)}(\omega)$. By Definition 2.5.1, part (2), we have $(\tau\wedge\kappa)(\omega *_\kappa \omega') = (\tau\wedge\kappa)(\omega)$ and $\kappa(\omega *_\kappa \omega') = \kappa(\omega)$, so that, by the non-anticipativity of F^n (which

follows from the non-anticipativity of F), we have

$$G(\omega *_\kappa \omega') = F^n_{(\tau\wedge\kappa)(\omega)}(\omega *_\kappa \omega') - F^n_{\kappa(\omega)}(\omega *_\kappa \omega')$$
$$= F^n_{(\tau\wedge\kappa)(\omega)}(\omega) - F^n_{\kappa(\omega)}(\omega) = G(\omega).$$

Since G is bounded (since so is F^n) we have

$$\int G\, d\mu' = \iint G(\omega *_\kappa \omega')\, \nu^{\leq\kappa}_{\omega}(d\omega')\, \mu(d\omega) = \int G(\omega)\, \mu(d\omega) = 0,$$

where the last equality follows from the fact that $\mu \in \mathcal{M}^{F,loc}$.

2. $G \in \mathcal{X}^{\geq\kappa}_n(F)$: Let $\tau \in \mathsf{QStop}$ be such that $G = F^n_{\tau\vee\kappa} - F^n_\kappa$. Then

$$\int F^n_{\tau\vee\kappa}(\omega) - F^n_\kappa(\omega)\, \mu'(d\omega) = \int \mathbf{1}_{\{\tau_n>\kappa\}}(\omega)(F^n_{\tau\vee\kappa}(\omega) - F^n_\kappa(\omega))\, \mu'(d\omega)$$
$$= \int \mathbf{1}_{\{\tau_n>\kappa\}}(\omega)(F_{(\tau\wedge\tau_n)\vee\kappa}(\omega) - F_\kappa(\omega))\, \mu'(d\omega)$$

Note that $(\tau \wedge \tau_n) \vee \kappa \geq \kappa$, and let τ' be as in Lemma 3.3.6 (applied to $(\tau \wedge \tau_n) \vee \kappa$). Also note that by Proposition 2.3.1, $\{\tau_n > \kappa\} \in \mathcal{F}_\kappa = \sigma(T_\kappa)$. Therefore $\mathbf{1}_{\{\tau_n>\kappa\}}$ is $\sigma(T_\kappa)$-measurable and so $\mathbf{1}_{\{\tau_n>\kappa\}}(\omega *_\kappa \omega') = \mathbf{1}_{\{\tau_n>\kappa\}}(\omega_{\leq\kappa}) = \mathbf{1}_{\{\tau_n>\kappa\}}(\omega)$. Continuing with the equalities from above, we have

$$\int F^n_{\tau\vee\kappa}(\omega) - F^n_\kappa(\omega)\, \mu'(d\omega) =$$
$$= \iint \mathbf{1}_{\{\tau_n>\kappa\}}(\omega)(F_{\kappa(\omega)+\tau'_\omega(\omega')}(\omega *_\kappa \omega') - F_\kappa(\omega *_\kappa \omega'))\, \nu_\omega(d\omega')\mu(d\omega)$$
$$= \iint \mathbf{1}_{\{\tau_n>\kappa\}}(\omega)(F_{\tau'_\omega}(\omega') - F_0(\omega'))\, \nu_\omega(d\omega')\mu(d\omega)$$
$$= \iint \mathbf{1}_{\{\tau_n>\kappa\}}(\omega)F_{\tau'_\omega}(\omega')\, \nu_\omega(d\omega')\mu(d\omega).$$

where the last equality used the TC-morphism assumption together with Lemma 3.3.4. With M_n given by (3.3.1), $|F|$ is bounded on $[0, \tau'_\omega]$ by $2M_{2n}$

when $\omega \in \{\kappa < \tau_n\}$, By the canonical local martingale property, we have $\int F_{\tau'_\omega}(\omega')\,\nu_\omega(d\omega') = 0$ for each $\omega \in \{\kappa < \tau_n\}$. Thanks to boundedness, again, the integral of G can be computed as an iterated integral:

$$\int G\,d\mu' = \int \mathbf{1}_{\{\tau_n > \kappa\}}(\omega) \left(\int F_{\tau'_\omega}(\omega')\,\nu_\omega(d\omega') \right) \mu(d\omega)$$

and therefore $\int G\,d\mu' = 0$. □

3.4 Sufficient conditions for disintegrability

The key to disintegrability for martingale-generated control structures is the existence of a shift operator, as described below. It plays the role of a partial inverse of the concatenation operator in the second argument.

Definition 3.4.1. A measurable map $\theta : \mathsf{Time} \times \Omega \to \Omega$ is said to be a **shift operator** if for all $\omega \in \Omega$, $t, s \in \mathsf{Time}$ and $\omega' \in \mathcal{C}_{\omega,t}$:

1. $\theta_t(\omega) \in \mathcal{C}_{\omega,t}$ and $\omega *_t \theta_t(\omega) = \omega$,
2. $(\theta_t(\omega))_{\leq t+s} = (\theta_t(\omega_{\leq s}))_{\leq t+s}$

Remark 3.4.2. Since $\mathcal{F}_t = \sigma(T_t)$ on Ω, then part (2) of Definition 3.4.1 is equivalent to:

$$\forall\, t, s \in \mathsf{Time} : \quad \theta_t^{-1}(\mathcal{F}_{t+s}) \subset \mathcal{F}_s$$

The stopping-time version of a shift operator θ is defined in the natural way

$$\theta_\tau(\omega) = \theta_{\tau(\omega)}(\omega),$$

where, for definiteness, we set $\theta_\infty(\omega) = \omega$, for all ω. This way, $\theta_\tau : \Omega \to \Omega$ is Borel measurable and retains the property that $\omega *_\tau \theta_\tau(\omega) = \omega$, for all $\omega \in \Omega$ and $\tau \in \mathsf{Stop}$.

Lemma 3.4.3. *For any $\kappa, \sigma \in \mathsf{Stop}$, the following is also a stopping time:*

$$\tau(\omega) := \kappa(\omega) + \sigma(\theta_\kappa(\omega))$$

Proof. Fix any $t \in \mathsf{Time}$ and $\omega \in \Omega$. In order to show $\{\tau \leq t\} \in \mathcal{F}_t$, it is enough to show that $\tau(\omega) \leq t$ if and only if $\tau(\omega_{\leq t}) \leq t$. Applying Proposition 2.3.1 to σ and using part (2) of the definition of θ gives the following equivalence:

$$\begin{aligned}
\tau(\omega) \leq t &\Leftrightarrow \sigma(\theta_{\kappa(\omega)}(\omega)) \leq t - \kappa(\omega) \\
&\Leftrightarrow \sigma((\theta_{\kappa(\omega)}(\omega))_{\leq t-\kappa(\omega)}) \leq t - \kappa(\omega) \\
&\Leftrightarrow \sigma((\theta_{\kappa(\omega)}(\omega_{\leq t}))_{\leq t-\kappa(\omega)}) \leq t - \kappa(\omega) \\
&\Leftrightarrow \sigma(\theta_{\kappa(\omega)}(\omega_{\leq t})) \leq t - \kappa(\omega)
\end{aligned}$$

First suppose $\tau(\omega) \leq t$. Since κ is a stopping time and $\kappa(\omega) \leq \tau(\omega) \leq t$, then $\kappa(\omega) = \kappa(\omega_{\leq t})$. Together with the above equivalence, this implies:

$$\begin{aligned}
\tau(\omega_{\leq t}) &= \kappa(\omega_{\leq t}) + \sigma(\theta_{\kappa(\omega_{\leq t})}(\omega_{\leq t})) \\
&= \kappa(\omega) + \sigma(\theta_{\kappa(\omega)}(\omega_{\leq t})) \leq t
\end{aligned}$$

For the other direction, suppose $\tau(\omega_{\leq t}) \leq t$. Since κ is a stopping time and $\kappa(\omega_{\leq t}) \leq \tau(\omega_{\leq t}) \leq t$, then $\kappa(\omega_{\leq t}) = \kappa(\omega)$. Therefore:

$$\kappa(\omega) + \sigma(\theta_{\kappa(\omega)}(\omega_{\leq t})) = \kappa(\omega_{\leq t}) + \sigma(\theta_{\kappa(\omega_{\leq t})}(\omega_{\leq t})) = \tau(\omega_{\leq t}) \leq t,$$

which implies $\tau(\omega) \leq t$ by the equivalence above. □

Proposition 3.4.4. *Let $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \in \mathsf{Time}})$ be a TC-space with concatenation operator $*$, on which a shift operator θ is defined. Suppose each $F \in \mathcal{D}$ is a canonically locally bounded TC-morphism into $(D^0_{\mathbb{R}}, \star)$, and that $*$ is a factor of X. Then $\mathcal{P}(\mathcal{D}, X)$ is disintegrable.*

Proof. Having fixed a shift operator θ, we pick $\omega_0 \in \Omega$, $\mu \in \mathcal{P}(\omega_0)$ and $\kappa \in \mathsf{Stop}$. For a stopping time $\sigma \in \mathsf{QStop}$ and define

$$\sigma_n(\omega) = (\sigma \wedge \tau_n)(\omega)$$
$$\tau(\omega) = \kappa(\omega) + \sigma(\theta_\kappa(\omega))$$

so that τ is a stopping time by Lemma 3.4.3. Since F is a TC-morphism into $(D^0_{\mathbb{R}}, \star)$ Lemma 3.3.4 implies that

$$F_\tau(\omega) - F_\kappa(\omega) = F_{\kappa+\sigma_n(\theta_\kappa)}(\omega *_\kappa \theta_\kappa \omega) - F_\kappa(\omega) = F_{\sigma_n}(\theta_\kappa \omega) = F^n_\sigma(\theta_\kappa \omega).$$

The same Lemma implies that $|F|$ is bounded by $|F_\kappa| + M_n$ on the entire stochastic interval $[0, \tau]$. In particular, for $A_m = \{|F_\kappa| \leq m\}$ we have

$$\mathbf{1}_{A_m} F_{\sigma_n} \circ \theta_\kappa = \mathbf{1}_{A_m}\Big(F_\tau - F_\kappa\Big) = \mathbf{1}_{A_m}\Big(F^{m+M_n}_\tau - F^{m+M_n}_\kappa\Big).$$

Since F^{m+M_n} is a bounded martingale under μ, for any bounded measurable function H on E we have $\int H(X(\omega_{\leq \kappa}))\mathbf{1}_{A_m}(\omega) F^n_\sigma(\theta_\kappa \omega)\, \mu(d\omega) = 0$, and, given that F^n is bounded, we can pass to the limit $m \to \infty$ by the dominated convergence theorem to obtain

$$\int H(X(\omega_{\leq \kappa})) F^n_\sigma(\theta_\kappa \omega)\, \mu(d\omega) = 0, \tag{3.4.1}$$

for all bounded and measurable H. With ν_x denoting a version of the regular conditional distribution of θ_κ given $X_\kappa = x$, we then have

$$0 = \int H(X(\omega_{\leq\kappa}))F^n_\sigma(\theta_\kappa\omega)\mu(d\omega) = \iint H(x)F^n_\sigma(\omega')\,\nu_x(d\omega')\,\mu_{X_\kappa}(dx),$$

where μ_{X_κ} is the μ-distribution of X_κ. Since H is arbitrary, it follows that

$$\int F^n_\sigma\,d\nu_x = 0 \text{ for } \mu_{X_\kappa}\text{-almost all } x \in E, \tag{3.4.2}$$

for all $\sigma \in \mathsf{QStop}$ and all $n \in \mathbb{N}$. Since QStop is countable, there exists a set $\mathcal{N}_1 \in \mathsf{Borel}(E)$ such that $\mu_{X_\kappa}(\mathcal{N}_1) = 0$, and the equality in (3.4.2) holds for all $x \in E \setminus \mathcal{N}_1$ and $\sigma \in \mathsf{QStop}$. Therefore $\nu_x \in \mathcal{M}^{F,loc}$ for all $x \in E \setminus \mathcal{N}_1$.

Since $*$ is a factor of X, we have $X(T_\kappa(\omega)) = X_0(\theta_\kappa(\omega))$ for all ω, and so

$$1 = \int \mathbf{1}_{\{X_\kappa(\omega)=X_0(\omega_{\geq\kappa})\}}\,\mu(d\omega) = \iint \mathbf{1}_{\{x=X_0(\omega')\}}\nu_x(d\omega')\mu_{X_\kappa}(dx),$$

This implies that there exists another zero set $\mathcal{N}_2 \in \mathsf{Borel}(E)$ such that $\mu_{X_\kappa}(\mathcal{N}_2) = 0$ and $X_0 = x$, ν_x-a.s. for all $x \in E \setminus \mathcal{N}_2$. Hence, $\nu_x \in \bar{\mathcal{P}}(x)$ for all $x \notin \mathcal{N}_1 \cup \mathcal{N}_2$. By picking a selector ν' of $\bar{\mathcal{P}}$ (which is nonempty by Proposition 2.9.2) and using it to set the values of ν_x on $\mathcal{N}_1 \cup \mathcal{N}_2$, we can arrange that $\nu_x \in \bar{\mathcal{P}}(x)$, for all $x \in E$. □

3.5 Main Result

Theorem 3.5.1 (DPP for Martingales)**.** *Let* $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathsf{Time}})$ *be a TC-space with concatenation operator* $*$ *and a shift operator* θ*. Suppose that* X *is* *state map from Ω to a Polish space E such that $*$ is a factor of X, and that $\mathcal{D}$ is a countable collection of canonically locally bounded TC-morphisms from $(\Omega, *)$ into $(D^0_{\mathbb{R}}, \star)$. Define the control structure as:*

$$\boldsymbol{P}(x) = \bigcap_{F \in \mathcal{D}} \mathcal{M}^{F,loc} \cap \Big\{\mu \in \mathsf{Prob}(\Omega) \ : \ X_0 = x, \ \mu\text{-}a.s.\Big\}$$
$$\mathcal{P}(\omega) = \boldsymbol{P}(X(\omega)) \text{ for } \omega \in \Omega$$

Let $G \in \mathcal{L}^{0-1}(\mathcal{P})$ be a tail random variable, and define the value function as

$$\bar{v}(x) = \sup_{\mu \in \boldsymbol{P}(x)} \int G \, d\mu$$

Then for all $\omega \in \Omega$, $x \in E$, and $\tau \in \mathsf{Stop}$ we have:

$$\bar{v}(x) = \sup_{\mu \in \boldsymbol{P}(x)} \int \bar{v}(X_\tau) \mathbf{1}_{\{\tau < \infty\}} + G \mathbf{1}_{\{\tau = \infty\}} \, d\mu$$

Proof. Use Propositions 3.2.1, 3.3.3, and 3.4.4 to get the analyticity, concatenability, and disintegrability (respectively) of the control structure $(\mathcal{D}, X)$. Then apply Theorem 2.10.1. □

Chapter 4

Controlled Diffusions: the Weak Formulation

4.1 Problem formulation and the main result

Throughout this section we fix the following:

1. a nonempty open set $\mathcal{O}$ in $\mathbb{R}^n$ and set $E = \mathrm{Cl}\mathcal{O}$ (the **state space**),

2. a nonempty standard Borel space A, (the **control space**),

3. Borel measurable functions $\beta : E \times A \to \mathbb{R}^n$ and $\sigma : E \times A \to \mathbb{R}^{n\times n}$ (the **coefficients**),

4. a Borel measurable function $g : E \to [-\infty, \infty]$ (the **objective function**).

We remind the reader that $C_{E^{\partial\mathcal{O}}}$ denotes the set of all continuous trajectories with values in E that get absorbed once they hit the boundary $\partial\mathcal{O}$.

4.1.1 Weak solutions to controlled SDEs

With Einstein's summation convention used throughout, we start by making precise what we mean by a controlled diffusion.

Definition 4.1.1 (Weak solutions to controlled SDEs)**.** A probability measure μ on $C_{E^{\partial\mathcal{O}}}$ is said to be a **weak solution of the controlled SDE**

$$d\xi_t^i = \beta^i(\xi_t, \alpha_t)\, dt + \sigma_k^i(\xi_t, \alpha_t)\, dW_t^k, \ \xi_0 = x, \tag{4.1.1}$$

with absorption in $\partial\mathcal{O}$ - denoted by $\mu \in \mathcal{L}^x(\beta, \sigma)$ - if there exists filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\in[0,\infty)}, \mathbb{P})$ on which three stochastic process $\{W_t\}_{t\in[0,\infty)}$, $\{\xi_t\}_{t\in[0,\infty)}$ and $\{\alpha_t\}_{t\in[0,\infty)}$ are defined, such that:

1. W is an $\mathbb{R}^n$ valued $\{\mathcal{F}_t\}_{t\in[0,\infty)}$-Brownian motion,
2. ξ is adapted and $\xi(\omega) \in C_{E^{\partial\mathcal{O}}}$ for all ω,
3. α is A-valued and progressively measurable,
4. $\int_0^t |\beta^i(\xi_u, \alpha_u)|\, du + \int_0^t (\sigma_k^i(\xi_u, \alpha_u))^2\, du < \infty$, a.s. for all i, k and $t \geq 0$,
5. $\xi_t = x + \int_0^t \beta^i(\xi_u, \alpha_u)\, du + \int_0^t \sigma_k^i(\xi_u, \alpha_u)\, dW_t^k$, a.s., for all $t \in [0, \tau_{\partial\mathcal{O}}]$, where
$$\tau_{\partial\mathcal{O}} = \inf\{t \geq 0 \,:\, \xi_t \in \partial\mathcal{O}\}, \text{ and}$$
6. μ is the law of $\xi_\cdot$ on $C_{E^{\partial\mathcal{O}}}$.

4.1.2 The stochastic optimal control problem

Given $x \in E$ and $\mu \in \mathcal{L}^x(\beta, \sigma)$, we set

$$J(\mu) = \mathbb{E}^\mu[G(\xi)] \text{ where } G(\xi) = \liminf_{t\to\infty} g(\xi_t), \tag{4.1.2}$$

with ξ denoting the coordinate map on $C_{E^{\partial \Theta}}$, where we use the convention that $\mathbb{E}[Y] = -\infty$ as soon as $\mathbb{E}[Y^-] = \infty$. The **value function** of the associated control problem is then given by

$$v(x) = \sup_{\mu \in \mathcal{L}^x(\beta,\sigma)} J(\mu), \ x \in E. \tag{4.1.3}$$

Remark 4.1.2. By choosing the state process ξ appropriately, this setup includes various common formulations of optimal stochasting control, including problems on a finite horizon (when $E = E_0 \times [0, T]$ and the last component plays the role of time) with terminal and/or running costs, discounted problems and stationary problems.

4.1.3 DPP for controlled diffusions

Theorem 4.1.3 (A dynamic programming principle for controlled diffusions - the weak formulation)**.** *Suppose that,*

1. *there exist locally bounded real functions $\hat{\beta} : E \to \mathbb{R}$ and $\hat{\sigma} : E \to \mathbb{R}$ such that*

$$|\beta^i(x, \alpha)| \leq \hat{\beta}(x) \text{ and } |\sigma^i_k(x, \alpha)| \leq \hat{\sigma}(x) \text{ for all } \alpha \in A,$$

2. *for each $x \in E$ we have $\mathcal{L}^x(\beta, \sigma) \neq \emptyset$, and*

3. *$J(\mu) > -\infty$ for each $\mu \in \mathcal{L}^x(\beta, \sigma)$.*

Then, the value function $v : E \to (-\infty, \infty]$ is universally measurable and satisfies the dynamic programming principle:

$$v(x) = \sup_{\mu \in \mathcal{L}^x(\beta,\sigma)} \mathbb{E}^\mu[v(\xi_\tau)\mathbf{1}_{\{\tau<\infty\}} + G(\xi.)\mathbf{1}_{\{\tau=\infty\}}], \text{ for all } x \in R,$$

for each (raw) stopping time τ on $C_{E^{\partial\heartsuit}}$.

Remark 4.1.4. Condition (1) in Theorem 4.1.3 is far from necessary and is placed mostly for convenience. It can be replaced by a different condition or relaxed by choosing a different control part Ω^{α} of the universal space $\Omega^{\alpha\xi}$ in the proof below.

4.2 Proof of Theorem 4.1.3

Our proof of Theorem 4.1.3 consists of two steps. In the first one, we observe that the family $\mathcal{L}^x(\beta,\sigma)$ can be manufactured by varying admissible controls on a single, universal, filtered probability space, and that it admits a martingale characterization there. In the second one we show that this, equivalent, setup fits our abstract framework of Section 3 so that Theorem 3.5.1 can be applied.

4.2.1 Construction of a universal setup

Let $\Omega^{\alpha} = \mathbb{L}^0_A$ be the space of all Lebesgue-a.e equivalence classes of A-valued Borel functions from $[0,\infty)$ to A, and let Ω^{ξ} be the subspace $C_{E^{\partial\heartsuit}}$ of the canonical space $C_{\mathbb{R}^n}$. Both can be given the structure of a filtered measurable space, namely $(\Omega^{\alpha},\mathcal{F}^{\alpha},\mathbb{F}^{\alpha}=\{\mathcal{F}^{\alpha}_t\}_{t\in\mathsf{Time}})$, $(\Omega^{\xi},\mathcal{F}^{\xi},\mathbb{F}^{\xi}=\{\mathcal{F}^{\xi}_t\}_{t\in\mathsf{Time}})$, as described in more detail in subsection 2.2 and in Example 2.4.3. We define the (universal) filtered measurable space $(\Omega^{\alpha\xi},\mathcal{F}^{\alpha\xi},\mathbb{F}^{\alpha\xi}=\{\mathcal{F}^{\alpha\xi}_t\}_{t\in\mathsf{Time}})$ simply as their product. In particular $\mathcal{F}^{\alpha\xi}_t=\mathcal{F}^{\alpha}_t\otimes\mathcal{F}^{\xi}_t$. It will be used in the second step that $\Omega^{\alpha\xi}$ is, in fact, a T-space - the product of T-spaces Ω^{α} and Ω^{ξ}.

Let $\mathsf{Coord} = \{x_i, x_i x_j \;:\; 1 \leq i,j \leq n\}$ be the family of coordinate functions and their products on $\mathbb{R}^n$, and let QCoord denote an arbitrary, but fixed throughout, countable family of bounded C^2-functions on $\mathbb{R}^n$ such that for each $f \in \mathsf{Coord}$ and each compact set $K \subseteq \mathbb{R}^n$ there exists $\tilde{f} \in \mathsf{QCoord}$ such that $f = \tilde{f}$ on K. Also, for $f \in C^2$ and $a \in A$ we define the $\mathcal{G}^a f$ by

$$(\mathcal{G}^a f)(x) = \beta^i(x,a)\partial_i f(x) + \tfrac{1}{2}\gamma^{ij}(a,x)\partial_{ij} f(x), \text{ with } \gamma^{ij} = \textstyle\sum_k \sigma^i_k \sigma^j_k,$$

Proposition 4.2.1 (A martingale characterization of weak solutions to controlled SDEs)**.** *The following two statements are equivalent for a probability measure μ on $C_{E^{\partial\mathcal{O}}}$:*

1. *μ is a weak solution to the controlled SDE* (4.1.1) *with absorption at $\partial\mathcal{O}$ starting at x, and*
2. *there exists a probability measure $\bar{\mu}$ on $\Omega^{\alpha\xi}$ whose Ω^{ξ}-marginal is μ such that*
   - *(a) $\xi_0 = x$, $\bar{\mu}$-a.s.,*
   - *(b) $\int_0^t |\beta^i(\xi_u, \alpha_u)|\, du + \int_0^t (\sigma^i_k(\xi_u, \alpha_u))^2\, du < \infty$ for all i,k and $t \in [0, \tau_{\partial\mathcal{O}}]$, $\bar{\mu}$-a.s., and*
   - *(c) for each $f \in \mathsf{QCoord}$,*

$$f(\xi_t) - f(\xi_0) - \int_0^{t\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u)\, du$$

is a $(\{\mathcal{F}^{\alpha\xi}_t\}_{t\in[0,\infty)}, \bar{\mu})$-local martingale .

If (1) *holds, then* (2b) *is true for all* $f \in C^2(E)$.

The proof follows, almost verbatim, the steps in the standard proof of the equivalence in the non-controlled case (see, e.g., Proposition 4.6, p. 315, [7]) so we omit the details. The only observation that needs to be made is that α is not a stochastic process in the classical sense. This difficulty can be circumvented by considering appropriate versions as in the following lemma. We remind the reader that an A-valued process $\{\hat{\alpha}_t\}_{t\in[0,\infty)}$ is considered progressively measurable if $\{\phi(\hat{\alpha}_t)\}_{t\in[0,\infty)}$ is progressively measurable for each Borel measurable $\phi : A \to [-1, 1]$.

Lemma 4.2.2. *There exists an* $\{\mathcal{F}_t^{\alpha\xi}\}_{t\in[0,\infty)}$*-progressively measurable process* $\{\hat{\alpha}_t\}_{t\in[0,\infty)}$ *with values in* A *such that* $\{\hat{\alpha}_t(\omega)\}_{t\geq 0}$ *is a Leb-a.e.-representative of the coordinate map* $\alpha(\omega)$ *for each* ω.

Conversely, let (ξ, α) *be a pair consisting of a continuous process* ξ *with values in* $\mathbb{R}^n$ *and an* A*-valued progressive process* α *defined on some filtered probability space* $(\Omega, \{\mathcal{F}_t\}_{t\in[0,\infty)}, \mathcal{F}, \mathbb{P})$. *Then* (ξ, α) *admits an* $\Omega^{\alpha\xi}$*-distribution, i.e., a probability measure* $\bar{\mu}$ *on* $\Omega^{\alpha\xi}$ *such that the* $\mathbb{P}$*-distribution of* $\int_0^t \varphi(u, \xi_u, \alpha_u)\, du$ *coincides with the* $\bar{\mu}$*-distribution of* $\int_{[0,t]} \varphi(u, \alpha, \xi)\, d\lambda$, *for each bounded and measurable* φ *and all* $t \geq 0$.

Proof. Let ϕ be an isomorphism (a bimeasurable bijection) between A and the closed interval $[-1, 1]$. Given $\alpha(\omega) \in \mathbb{L}^0([0,\infty), A)$, we define $\hat{\alpha}$ by

$$\hat{\alpha}(t) = \phi^{-1}\Big(\liminf_{n\to\infty} \Phi_t^n(\omega)\Big) \text{ where } \Phi_t^n(\omega) = \tfrac{1}{n} \int_{(t-1/n)^+}^t \phi\big(\alpha_u(\omega)\big)\, du.$$

It is straightforward to check that $\hat{\alpha}(\omega)$ is a representative of $\alpha(\omega)$ for each ω. Moreover $\phi(\hat{\alpha})$ (and, therefore, α) is a progressively-measurable process, as a pointwise limit of continuous adapted processes.

For the converse, and under the assumptions of the second part of the Lemma, let $\bar{\mu}$ be the pushforward of $\mathbb{P}$ via the map $\Phi : \Omega \to \Omega^{\alpha\xi}$ defined as follows:

$$\Phi(\omega) = \Big((\xi_t(\omega))_{t\geq 0}, \alpha(\omega) \Big),$$

where $\alpha(\omega)$ is the Leb-a.e.-equivalence class of $(\alpha_t(\omega))_{t\geq 0}$. (Progressive) measurability of α guarantees that Φ is a measurable map. The equality of the distributions of two integrals in the statement is then a simple consequence of the monotone-class theorem. $\square$

4.2.2 An application of the abstract DPP

Proposition 4.2.1 allows us to reformulate our control problem so as to fit the setting of the first part of our paper. Indeed, it states that the value function $v(x)$ can be represented as

$$v(x) = \sup_{\bar{\mu} \in \bar{\mathcal{P}}^x} \mathbb{E}^{\mu}[G(\xi)]$$

where $\bar{\mathcal{P}}^x$ is the family of all probability measures on $\Omega^{\alpha\xi}$ such that (2a), (2b) and (2c) hold, and our job is to show that it is, in fact, a martingale generated control structure which satisfies all the requirements of the abstract Theorem 3.5.1.

Thanks to the discussion and examples in subsections 2.4 and 2.6, the space $\Omega^{\alpha\xi}$ admits a natural structure of a TC-space, with the strict concatenation used for the ξ component. The map $X : \Omega^{\alpha\xi} \to E$, given by $X(\xi,\alpha) = \liminf_{t\to\infty} \xi_t$ computed componentwise, and suitably measurably altered to take values in E and when the limits inferior take infinite values, so that $X_t(\xi,\alpha) = \xi_t$. Given that the concatenation operator in α requires no compatibility conditions, and the one in ξ is strict, the product concatenation operator $*$ factors through X (and is a factor of X). Also, there is a naturally-defined shift operator θ on $\Omega^{\alpha\xi}$.

Condition (1) of Theorem 4.1.3 takes care of the integrability condition (2b) of Proposition 4.2.1, so we can conclude that we are, indeed, dealing with a martingale-generated control structure with the state map X, generated by the family $\mathcal{D}$ which consist of (well-defined) maps of the form

$$F(\alpha,\xi)_t = f(\xi_t) - f(\xi_0) - \int_0^{t\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u)\, du$$

with f ranging through the countable set QCoord. The last thing we need to check, before we can apply Theorem 3.5.1, is that each such F is a TC-morphism into $(D^0_{\mathbb{R}}, \star)$. We fix $f \in \mathsf{QCoord}$, and note that the corresponding functional F clearly takes values in $D^0_{\mathbb{R}}$ and that it is non-anticipating. To establish the TC-morphism property let us fix $s, t \in \mathsf{Time}$ and $\omega, \omega' \in \Omega^{\alpha\xi}$ such that ω is compatible with ω' at t. The case of $s \leq t$ is straightforward, so suppose $s > t$. Since the ξ component uses the strict concatenation operator, then $\xi_t(\omega) = \xi_0(\omega')$, and furthermore:

$$\tau_{\partial\mathcal{O}}(\omega) \leq t \quad \Leftrightarrow \quad \xi_t(\omega) \in \partial\mathcal{O} \quad \Leftrightarrow \quad \xi_0(\omega') \in \partial\mathcal{O} \quad \Leftrightarrow \quad \tau_{\partial\mathcal{O}}(\omega') = 0$$

Combining this with the properties of concatenation gives:

$$\begin{aligned}
\int_{t\wedge\tau_{\partial\mathcal{O}}}^{s\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega *_t \omega'))\, du &= \mathbf{1}_{\{\tau_{\partial\mathcal{O}}>t\}}(\omega) \int_{t\wedge\tau_{\partial\mathcal{O}}}^{s\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega *_t \omega'))\, du \\
&= \mathbf{1}_{\{\tau_{\partial\mathcal{O}}>0\}}(\omega') \int_0^{(s-t)\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega'))\, du \\
&= \int_0^{(s-t)\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega'))\, du
\end{aligned}$$

Putting everything together gives:

$$\begin{aligned}
F(\omega *_t \omega')_s &= f(\xi_s(\omega *_t \omega')) - f(\xi_0(\omega *_t \omega')) - \int_0^{s\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega *_t \omega'))\, du \\
&= \left(f(\xi_t(\omega *_t \omega')) - f(\xi_0(\omega *_t \omega')) - \int_0^{t\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega *_t \omega'))\, du \right) \\
&\quad + \left(f(\xi_s(\omega *_t \omega')) - f(\xi_t(\omega *_t \omega')) - \int_{t\wedge\tau_{\partial\mathcal{O}}}^{s\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega *_t \omega'))\, du \right) \\
&= \left(f(\xi_t(\omega)) - f(\xi_0(\omega)) - \int_0^{t\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega))\, du \right) \\
&\quad + \left(f(\xi_{s-t}(\omega')) - f(\xi_0(\omega')) - \int_0^{(s-t)\wedge\tau_{\partial\mathcal{O}}} \mathcal{G}^{\alpha_u} f(\xi_u(\omega'))\, du \right) \\
&= F(\omega)_t + F(\omega')_{t-s} = (F(\omega) \star_t F(\omega'))_s
\end{aligned}$$

4.3 Viscosity solutions

We conclude this example by showing how our result can be applied to show that value functions of stochastic control problems are viscosity solutions to the associated Hamilton-Jacobi-Bellman equations under weak conditions.

In particular, we do not require that the equation itself admit an a-priori solution, or that any solution is smooth or unique (i.e, that the comparison principle hold). Our results, in particular, imply some of the results in [1] and [3] under weaker assumptions. We note that the lack of any strong ellipticity allow us keep assuming, without loss of generality, that the problem is time-independent; time can be incorporated as just another (space) variable with linear dynamics and the terminal condition imposed as part of the boundary condition.

For a C^2 function $\varphi : \mathcal{O} \to \mathbb{R}$ we define the **Hamiltonian** $H\varphi : \mathcal{O} \to (-\infty, \infty]$ by

$$H\varphi(x) = \sup_{a\in A} \mathcal{G}^a\varphi(x) = \sup_{a\in A} \Big(\beta^i(x,a)\partial_{x_i}\varphi(x) + \tfrac{1}{2}\gamma^{ij}(x,a)\partial_{x_i x_j}\varphi(x)\Big).$$

4.3.1 The viscosity property of the value function

Definition 4.3.1. Let v be a real-valued function defined in a neighborhood $\mathcal{V}$ of a point $\bar{x} \in \mathcal{O}$, and let v_* and v^* denote its lower and upper semicontinuous envelopes, respectively. We say that v is a

1. **viscosity supersolution** of the equation $Hv = 0$ at $\bar{x}$ if $H\varphi(\bar{x}) \leq 0$ for each $\varphi \in C^2(\mathcal{V})$ with the property that $\varphi(\bar{x}) = v_*(\bar{x})$ and $\varphi(x) < v_*(x)$ for $x \in \mathcal{V} \setminus \{\bar{x}\}$, and

2. **viscosity subsolution** of the equation $Hv = 0$ at $\bar{x}$ if $H\varphi(\bar{x}) \leq 0$ for each $\varphi \in C^2(\mathcal{V})$ with the property that $\varphi(\bar{x}) = v^*(\bar{x})$ and $\varphi(x) > v^*(x)$ for $x \in \mathcal{V} \setminus \{\bar{x}\}$.

A function which is both a viscosity supersolution and a viscosity subsolution is called a **viscosity solution** to $Hv = 0$ at $\bar{x}$.

For $x \in \mathbb{R}^n$ and $r > 0$ we define

$$\tau^{r,x} = \inf\{t \geq 0 \,:\, d(x, \xi_t) \geq r\} \wedge r,$$

where d denotes the Euclidean distance on $\mathbb{R}^n$, so that $\tau^{r,x}$ is a raw stopping times on $\Omega^{\alpha\xi}$.

Theorem 4.3.2. *Given $\bar{x} \in \mathcal{O}$, suppose that there exists a neighboorhood $\mathcal{V}$ of $\bar{x}$ in $\mathcal{O}$ such that*

1. ***(availability of DPP)*** *the assumptions of Theorem 4.1.3 hold and v is finite on $\mathcal{V}$,*
2. ***(continuity of coefficients)*** *$x \mapsto \beta^i(x, a)$ and $x \mapsto \sigma^i_k(x, a)$ are continuous functions on $\mathcal{V}$ for all $a \in A$,*
3. ***(admissibility of locally constant controls)*** *there exists a constant $r > 0$ such that for each $x \in \mathcal{V}$ and $a \in A$ there exists a control process $\{\alpha_t\}_{t\in[0,\infty)}$ and an associated weak solution $\{\xi_t\}_{t\in[0,\infty)}$ of the controlled SDE* (4.1.1) *with $\xi_0 = x$ (defined on some filtered probability space) such that*

$$\alpha_t = a \quad \text{for } t \in [0, \tau] \text{ a.s.}, \quad \text{where } \tau = \inf\{t \geq 0 \,:\, d(\xi_t, \bar{x}) \geq r\} \wedge r.$$

Then the value function v is a viscosity solution to $Hv = 0$ at x_0.

Proof. We split the proof into two parts, in which we establish the supersolution and the subsolution property of v separately.

The supersolution property. We take $\varphi \in C^2$ which touches v_* at $\bar{x}$ from below, i.e. $v_*(\bar{x}) = \varphi(\bar{x})$ and $\varphi(x) < v_*(x)$ for $x \neq \bar{x}$ This implies that there exists a sequence $\{x_m\}_{m\in\mathbb{N}}$ such that

$$v(x_m) \leq \varphi(x_m) + \tfrac{1}{m} \text{ and } d(x_m, \bar{x}) \leq \tfrac{1}{m}. \tag{4.3.1}$$

Suppose, for contradiction, that $H\varphi(\bar{x}) > 0$. Then there exists $a \in A$ such that $(\mathcal{G}^a\varphi)(\bar{x}) > 0$. Since $\mathcal{G}^a\varphi$ is continuous in x, there exist constants $\varepsilon > 0$ and $r > 0$ such that $(\mathcal{G}^a\varphi)(x) \geq \varepsilon$ when $d(x, \bar{x}) \leq r$. Using the fact that $\varphi(x) < v_*(x)$ as soon as $x \neq \bar{x}$ and that the function $v_* - \varphi$ is lower semicontinuous, we find that

$$\delta = \min\{v_*(x) - \varphi(x) \,:\, d(x, \bar{x}) = r\} > 0.$$

For each $m \in \mathbb{N}$, let μ_m be the law of the weak solution $\{\xi_t\}_{t\in[0,\infty)}$ described in part 3 of the statement, where we assume, without loss of generality, that the same constant $r > 0$, as above, can be used. Proposition 4.2.1 and the local nonnegativity of $\mathcal{G}^a\varphi - \varepsilon$ imply that $\varphi(\xi_t) - \varepsilon t$ is a bounded μ_m-submartingale under μ_m on $[0, \tau^{r,\bar{x}}]$. Therefore, with $\tau = \tau^{r,\bar{x}}$ and for $m > 1/r$, we get

$$\begin{aligned}
\varphi(x_m) &\leq \mathbb{E}^{\mu_m}[\varphi(\xi_\tau) - \varepsilon\tau] \leq \mathbb{E}^{\mu_m}[\varphi(\xi_\tau)\mathbf{1}_{\{\tau<r\}}] + \mathbb{E}^{\mu_m}[(\varphi(\xi_\tau) - \varepsilon r)\mathbf{1}_{\{\tau=r\}}] \\
&\leq \mathbb{E}^{\mu_m}[(v_*(\xi_\tau) - \delta)\mathbf{1}_{\{\tau<r\}}] + \mathbb{E}^{\mu_m}[(v_*(\xi_\tau) - \varepsilon r)\mathbf{1}_{\{\tau=r\}}] \\
&\leq \mathbb{E}^{\mu_m}[v_*(\xi_\tau)] - \min(\delta, \varepsilon r).
\end{aligned}$$

Using the dynamic programming principle of Theorem 4.1.3 and the relation (4.3.1) above, we finally obtain

$$\begin{aligned} v(x_m) - \tfrac{1}{m} + \min(\delta, \varepsilon r) &\leq \mathbb{E}^{\mu_m}[v_*(\xi_\tau)] \leq \mathbb{E}^{\mu_m}[v(\xi_\tau)] \\ &\leq \sup_{\mu \in \mathcal{L}^{x_m}(\beta,\sigma)} \mathbb{E}^{\mu}[v(\xi_\tau)] = v(x_m), \end{aligned}$$

and reach a contradiction by taking m large enough.

The subsolution property. We pick $\varphi \in C^2$ which touches v^* at $\bar{x}$ from above, i.e. $v^*(\bar{x}) = \varphi(\bar{x})$ and $\varphi(x) > v_*(x)$ for $x \neq \bar{x}$. As in the first part of the proof, this implies that there exists a sequence $\{x_m\}_{m \in \mathbb{N}}$ such that

$$v(x_m) \geq \varphi(x_m) - \tfrac{1}{m} \text{ and } d(x_m, \bar{x}) \leq \tfrac{1}{m}. \tag{4.3.2}$$

Suppose, for contradiction, that $H\varphi(\bar{x}) < 0$. Being representable as a supremum of continuous functions, $H\varphi$ is upper semicontinuous, and, so, there exist constants $r > 0$ and $\varepsilon > 0$ such that $H\varphi(x) \leq -\varepsilon$ for all x with $d(x, \bar{x}) \leq r$. Using the fact that $\varphi(x) > v^*(x)$ as soon as $x \neq \bar{x}$ and that the function $\varphi - v^*$ is lower semicontinuous, we find, as above, that

$$\delta = \min\{\varphi(x) - v^*(x) \,:\, d(x, \bar{x}) = r\} > 0.$$

Let the laws $(\mu_m)_{m \in \mathbb{N}}$ be defined as in the first part of the proof, so that under each μ_m the process $\varphi(\xi_t) + \varepsilon t$ is supermartingale on $[0, \tau^{r,\bar{x}}]$. It follows that, with $\tau = \tau^{r,x}$, we have

$$\begin{aligned}
\varphi(x_m) &\geq \mathbb{E}^{\mu}[\varphi(\xi_\tau) + \varepsilon\tau] \\
&= \mathbb{E}^{\mu}[(\varphi(\xi_\tau) + \varepsilon\tau)\mathbf{1}_{\{\tau=r\}}] + \mathbb{E}^{\mu}[(\varphi(\xi_\tau) + \varepsilon\tau)\mathbf{1}_{\{\tau<r\}}] \\
&\geq \mathbb{E}^{\mu}[(v^*(\xi_\tau) + \delta)\mathbf{1}_{\{\tau=r\}}] + \mathbb{E}^{\mu}[(\varphi(\xi_\tau) + \varepsilon r)\mathbf{1}_{\{\tau<r\}}] \\
&\geq \mathbb{E}^{\mu}[v(\xi_\tau)] + \min(\delta, \varepsilon r)
\end{aligned}$$

We take a supremum over all $\mu \in \mathcal{P}^{x_m}$ on the right hand side and use the DPP to conclude that $\varphi(x_m) \geq v(x_m) + \min(\delta, \varepsilon r)$ for all m - a contradiction with (4.3.2). □

Bibliography

[1] Erhan Bayraktar and Mihai Sîrbu. Stochastic Perron's method for Hamilton-Jacobi-Bellman equations. *SIAM J. Control Optim.*, 51(6):4274–4294, 2013.

[2] Dimitri P. Bertsekas and Steven E. Shreve. *Stochastic optimal control*, volume 139 of *Mathematics in Science and Engineering*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York, 1978. The discrete time case.

[3] Bruno Bouchard and Nizar Touzi. Weak dynamic programming principle for viscosity solutions. *SIAM J. Control Optim.*, 49(3):948–962, 2011.

[4] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes. Vol. I*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2003. Elementary theory and methods.

[5] N. El Karoui and X. Tan. Capacities, Measurable Selection and Dynamic Programming Part I: Abstract Framework. *ArXiv e-prints*, 2013. 1310.3363.

[6] N. El Karoui and X. Tan. Capacities, Measurable Selection and Dynamic Programming Part II: Application in Stochastic Control Problems. *ArXiv e-prints*, 2013. 1310.3364.

[7] Ioannis Karatzas and Steven E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991.

[8] K. R. Parthasarathy. *Probability measures on metric spaces*. AMS Chelsea Publishing, Providence, RI, 2005. Reprint of the 1967 original.

[9] S. M Srivastava. *A course on Borel sets*, volume 180. Springer, New York, 1998.

[10] Christophe Stricker. Quasimartingales, martingales locales, semimartingales et filtration naturelle. *Probability Theory and Related Fields*, 39:55–63, 1977.

[11] Gordan Žitković. Dynamic programming for controlled Markov families: abstractly and over martingale measures. *SIAM Journal of Control and Optimization*, 52(3):1597–1621, 2014.