The Annals of Applied Probability
2022, Vol. 32, No. 3, 1902–1928
https://doi.org/10.1214/21-AAP1724

A DYNAMIC PROGRAMMING APPROACH TO DISTRIBUTION-CONSTRAINED OPTIMAL STOPPING

BY SIGRID KÄLLBLAD[a]

Department of Mathematics, Royal Institute of Technology, [a]*sigrid.kallblad@math.kth.se*

We consider an optimal stopping problem where a constraint is placed on the distribution of the stopping time. Reformulating the problem in terms of so-called measure-valued martingales enables us to transform the distributional constraint into an initial condition and view the problem as a stochastic control problem; we establish the corresponding dynamic programming principle. The method offers a systematic approach for solving the problem for general constraints and under weak assumptions on the cost function. In addition, we provide certain continuity results for the value of the problem viewed as a function of its distributional constraint.

1. Introduction. Our main problem of study is an optimal stopping problem where a distributional constraint is placed on the stopping time. Specifically, given a nonanticipative cost function, say c, we are interested in the problem of finding

$$\sup_{\tau} \mathbb{E}\big[c(B_{\cdot\wedge\tau}, \tau)\big],$$

where $(B_t)_{t\geq 0}$ is a Brownian motion and the optimisation is taken over stopping times satisfying the constraint $\tau \sim \mu$ for an a priori specified probability distribution μ. We consider the weak formulation of this problem where the optimisation is performed over filtered probability spaces supporting the Brownian motion and the constrained stopping times; the problem can therefore also be understood as a particular optimal transport problem. While of interest in its own right, the problem is also related to the so-called inverse first-passage-time problem which has a long history. It has also attracted recent attention; see, for example, Beiglböck et al. [8], where existence and characterisation of an optimiser is provided for certain cost functions, or, for example, Ekström and Janson [16] to which we refer for further motivation, references and an exposition of its role within financial and actuarial mathematics. Among the related literature we also single out the articles by Bayraktar and Miller [5], who considered an optimal stopping problem equipped with a distributional constraint of atomic form, as well as Ankirchner et al. [2] and Miller [30], where optimal stopping problems featuring constraints on the expected value of the stopping time were studied.

The idea herein is to address this problem—for general distributional constraints and under weak assumptions on the cost function—via a reformulation in terms of so-called measure-valued martingales. The reformulation transforms the distributional constraint on the stopping time into an initial condition for the associated measure-valued martingale enabling addressing the problem by use of dynamic programming arguments. The method thus unleashes the full machinery of dynamic programming and stochastic control and opens up for a systematic approach for solving distribution-constrained optimal stopping problems.

Received September 2019; revised April 2021.

MSC2020 subject classifications. Primary 60G40, 60G57, 49L20, 62L15; secondary 49Q20, 93E20, 58E25, 91A15, 28B20, 60J25, 06A06.

Key words and phrases. Distribution-constrained optimal stopping, measure-valued martingales, dynamic programming principle, optimal transport, measurable selection.

In a nutshell, each stopping time τ is identified with the process specifying its distribution conditional on current information; that is,

$$\xi_t \mathrel{\widehat{=}} \mathcal{L}(\tau|\mathcal{F}_t), \quad t \geq 0.$$

Albeit, in order to account for the fact that we consider a weak formulation, we will rather consider conditional distributions of so-called randomised stopping times. The thus defined process belongs to a specific class of measure-valued martingales which we refer to as adapted; given that $\tau \sim \mu$, this process will also satisfy the initial condition $\xi_0 = \mu$. More precisely, a measure-valued martingale (MVM) is a process, say $(\xi_t)_{t\geq 0}$, which takes its values in the space of probability-measures and for which $(\xi_t(\varphi))_{t\geq 0}$ is a martingale for nice enough test functions φ. The key property satisfied by stopping times—namely that the information available at each time is enough to determine whether or not to stop—implies that the corresponding conditional distribution process in addition need to satisfy a certain adaptedness property. Crucially, the processes satisfying these properties also characterise all distribution-constrained stopping times. In effect, the distribution-constrained optimal stopping problem admits an equivalent formulation as an optimisation problem over adapted MVMs starting off in μ—our first result formalises this equivalence.

Our main result then establishes the dynamic programming principle for this reformulated problem; the result ultimately relies on the fact that the distributional constraint now is incorporated as an initial condition for the measure-valued state process. To establish the DPP, we consider its weak formulation on the associated canonical path space; the canonical framework has previously successfully been used for the study of stochastic control problems by, for example, El Karoui and Tan [18, 19], see also Neufeld and Nutz [31], Nutz and van Handel [32] and Žitković [38]. We note that since we are dealing with MVMs, a component of our canonical space will consist of functions from the positive reals into the space of probability measures which we equip with the topology induced by the first Wasserstein distance rendering it a Polish space. We then establish the DPP by proving analyticity and stability under concatenation and disintegration of certain sets of measures, and, in turn, applying Jankov-von Neumann's measurable selection theorem. The approach enables us to establish the DPP for general distributional constraints and under weak assumptions on the cost function; while imposing certain regularity, we notably do not require the cost to be of any specific "structural form". In parallel, a priori assuming certain continuity of the value function and restricting to Markovian cost functions, we provide an alternative proof of the DPP; we choose to report this complementary argument due to its simplicity. Following Bouchard and Touzi [11], see also Bouchard and Nutz [10], the idea is to exploit continuity properties of the value function in order to explicitly construct an approximately optimal kernel and thus circumvent the need for measurable selection arguments. To deal with the measure-valued argument we here borrow ideas from optimal transport and the argument makes use of the structure of the first Wasserstein distance. The fact that we are dealing with adapted MVMs implies that the ordering on the set of probability measures becomes of key importance; we note that these straightforward covering arguments are difficult to generalise precisely due to such ordering considerations.

Finally, we consider the stability of the distribution-constrained optimal stopping problem in that we establish continuity properties of the value of the problem as a function of its distributional constraint. While upper semicontinuity holds in great generality, continuity is more involved. Again, the question turns out to be linked to the ordering on the underlying space of probability distributions. Equipping this space with the usual stochastic ordering—while still making use of its properties as a metric space under the Wasserstein distance—we establish "right-continuity" of the value function for a specific class of cost functions. While we do not exclude that continuity might hold under appropriate assumptions, we note that the present result suffices to ensure that a stopping problem equipped with a general constraint may be approximated by problems featuring atomic constraints which are easier to handle by use of numerical methods.

We now comment on the related literature: Measure-valued martingales have been around for a long time; see, for example, Horowitz [25] and the lecture notes by Dawson [14] and Walsh [37]. They were recently introduced to the study of financially motivated robust pricing problems by Cox and the present author in [13]; that is, to optimisation problems over martingales satisfying marginal constraints (guaranteeing fit to market data). While the methods used in the present article and in [13] are related, the distinct features of the two problems calls for decisive differences: First and foremost, the present connection between the MVMs and the a priori given optimisation objects requires the use of adapted MVMs; in consequence, the ordering on the set of probability measures plays a central role. Moreover, to account for the fact that we consider a weak formulation of the optimal stopping problem, our MVMs correspond to conditional distributions of disintegrated randomised stopping times; in effect, they will not necessarily terminate. Bayraktar and Miller [5] notably study a distribution-constrained optimal stopping problem using a similar perspective. They, however, restrict to a more specific class of cost functions and atomic constraints only; in contrast to us, they also consider the strong formulation of the problem. Combining their result with ours, we obtain, however, as a corollary that for the class of cost functions and constraints considered in [5], the weak and strong formulations of the problem coincide. Our method of proof is notably distinct from theirs; indeed, a key motivation for the present article has been to establish the DPP for general distributional constraints and path-dependent cost functions which called for a different set of arguments. Ankirchner et al. [2] and Miller [30] (cf. also Bouchard and Nutz [10]) consider an optimal stopping problem with a constraint placed on the expected value of the stopping time; the conditional expected value of the terminal constraint is then incorporated as an additional state process and the problem is addressed using BSDEs. At an abstract level, our approach bears resemblance to theirs for our MVMs—corresponding to conditional laws of stopping times—may be viewed as infinite-dimensional additional state processes rendering the problem dynamically consistent. Finally, we note that distribution-constrained optimal stopping problems are closely related to so-called causal optimal transport problems; see, for example, Beiglböck and Lacker [9]. Our approach may thus be extended to such problems in a natural way. In the context of (discrete) causal transport, a recursive scheme was notably introduced by Backhoff et al. [3], where the corresponding dynamic programming principle was also established. It is interesting to compare their scheme to ours; since the causality condition is nonsymmetric, the approaches are, however, conceptually different.

The remainder of the article is organised as follows: In Section 2, we introduce our problem of study and the notion of adapted measure-valued martingales; we then establish the equivalence between the original problem formulation and the corresponding optimisation problem featuring MVMs. In Section 3, we introduce the conditional version of the problem and prove the dynamic programming principle. In Section 4, we consider continuity properties of the value function and provide the direct proof of the DPP using covering arguments. Appendix A collects results on the continuity and convergence of MVMs which are of independent interest; various auxiliary arguments are deferred to Appendix B.

Notation: $\diamond$ We write $\mathcal{B}$ for the Borel algebra and $b\mathcal{B}$, C and C_b for real-valued functions being, respectively, bounded and measurable, continuous, and continuous and bounded; if no other domain is explicitly mentioned, it is understood to be $\mathbb{R}$. $\diamond$ $C_0(\mathbb{R}_+)$ denotes the space of continuous functions from $\mathbb{R}_+$ to $\mathbb{R}$ which equal zero at the origin and are equipped with the topology of uniform convergence on compact sets. $\diamond$ $\mathcal{P}$ denotes the set of probability measures on the reals which are concentrated on $(0, \infty)$ and have finite first moment;

occasionally, we equip $\mathcal{P}$ with the weak topology but most frequently with the (subspace) topology induced by the first Wasserstein distance denoted by $\mathcal{W}_1$. ⋄ Given a μ-integrable function φ, we write $\mu(\varphi) = \int_0^\infty \varphi(s)\mu(\mathrm{d}s)$; in particular, given a $\mathcal{P}$-valued stochastic process $\xi = (\xi_t)_{t\geq 0}$ and $\varphi \in b\mathcal{B}$, $\xi(\varphi)$ denotes a real-valued stochastic process. ⋄ We employ the convention $\mathbb{E}[\Xi] = \mathbb{E}[\Xi^+] - \mathbb{E}[\Xi^-]$ with $\infty - \infty = -\infty$. ⋄ $\mathbb{Q}$ denotes the rational numbers.

2. Distribution-constrained optimal stopping and MVMs. For our main problem of study, we consider a fixed filtered probability space $(\Omega, \mathcal{G}, \mathbb{G}, \mathbb{P})$ with $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0}$. We suppose that it is rich enough to support a Brownian motion $(W_t)_{t\geq 0}$ and an independent $\mathcal{G}_0$-measurable random variable which is uniformly distributed on $[0, 1]$. For $\mu \in \mathcal{P}$, the set of stopping times which satisfy the constraint $\tau \sim \mu$ is denoted by $\mathcal{T}(\mu)$. We take $c : C_0(\mathbb{R}_+) \times \mathbb{R}_+ \to \mathbb{R}$ to be a given measurable cost function and suppose that it is nonanticipative in the sense that $c(\omega, t) = c(\omega_{\cdot\wedge t}, t)$. Given $\mu \in \mathcal{P}$, we then consider the distribution-constrained optimal stopping problem of maximising

$$\mathbb{E}\big[c(W_{\cdot\wedge\tau}, \tau)\big], \quad \text{over } \tau \in \mathcal{T}(\mu). \tag{2.1}$$

Throughout, we impose the following nondegeneracy assumption on the cost function; in consequence, $v(\mu)$ is well defined for $\mu \in \mathcal{P}$ with values in $(-\infty, \infty]$, where

$$v(\mu) := \sup_{\tau\in\mathcal{T}(\mu)} \mathbb{E}\big[c(W_{\cdot\wedge\tau}, \tau)\big].$$

ASSUMPTION 2.1. For each $\mu \in \mathcal{P}$, $\mathbb{E}[c(W_{\cdot\wedge\tau}, \tau)]$ is well defined in $[-\infty, \infty)$ for all $\tau \in \mathcal{T}(\mu)$, and there exists some $\tau \in \mathcal{T}(\mu)$ with $\mathbb{E}[c(W_{\cdot\wedge\tau}, \tau)] > -\infty$.

We note that the required richness of the probability space implies that we effectively are dealing with a weak formulation of the optimal stopping problem. The remainder of this section is devoted to demonstrating how this problem can be reformulated in terms of measure-valued martingales so as to obtain a problem which is natural to address by use of dynamic programming arguments.

2.1. *Adapted measure-valued martingales and alternative problem formulation.* To motivate our alternative problem formulation, consider first a stopping time $\tau \sim \mu$ in the filtered probability space $(\Omega, \mathcal{G}, \mathbb{F}, \mathbb{P})$, where $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ denotes the augmented filtration generated by the Brownian motion alone. Defining the process (this will be made precise below)

$$\xi_t \mathrel{\widehat{=}} \mathcal{L}(\tau|\mathcal{F}_t), \quad t \geq 0,$$

we then obtain a process taking its values in $\mathcal{P}$ and which is a measure-valued martingale in the sense that $\xi(\varphi)$ is a martingale for any $\varphi \in C_b$. Morally, the process ξ gradually reveals the full structure of the stopping time τ upon arrival of new information: $\xi_0 = \mu$ and as t tends to infinity, $\xi_t \to \delta_\tau$ weakly, a.s. The fact that an MVM eventually collapses into a random dirac measure is a feature referred to as termination. More pertinently, denoting the almost sure weak limit of ξ by ξ_∞, the fact that τ is a stopping time implies that the following (stronger) property must hold:

$$\inf\big\{t > 0 : \xi_t = \delta_x, x \in (0, \infty)\big\} \leq \inf\{x : x \in \operatorname{supp} \xi_\infty\}, \quad \text{a.s.}; \tag{2.2}$$

that is, almost surely, ξ must eventually terminate and for each $t \geq 0$, ξ_t is concentrated either on a singleton or on (t, ∞). Crucially, also the converse holds true; that is, every process satisfying the above properties defines a stopping time with distribution μ. In particular, if ξ is sufficiently nice, the latter can be recovered from

$$\tau = \arg\{t \geq 0 : \xi_t = \delta_t\}; \tag{2.3}$$

indeed, the fact that ξ is terminating together with (2.2) implies that τ is a well-defined stopping time, and the fact that ξ is an MVM satisfying $\xi_0 = \mu$ ensures that $\tau \sim \mu$. Putting the above together, we see that there is a one-to-one correspondence between stopping times $\tau \sim \mu$ in the Brownian filtration, and MVMs satisfying (2.2) and $\xi_0 = \mu$. The associated distribution-constrained optimal stopping problem may thus, equivalently, be formulated as an optimisation problem over such MVMs. Since these processes are constructed dynamically forwards in time, the distributional constraint is now encoded by an initial condition and the formulation is the appropriate one for addressing the problem via dynamic programming arguments.

While the above discussion highlights the ideas underpinning our approach, we note that in order to handle the optimal stopping problem in (2.1), we need to consider a larger class of stopping times than those discussed above. The way we choose to formalise this is by *relaxing the condition of termination*. There are two ways to understand this approach: From the perspective of the problem formulation in (2.1), if $\mathbb{G}$ is a Brownian filtration initially enlarged by an independent uniformly distributed random variable, this corresponds to identifying a $\mathbb{G}$-stopping time with the process yielding its conditional law given the Brownian filtration alone. Alternatively, since problem (2.1) is given in a weak formulation, it can be argued to be equivalent to optimizing the objective over so-called randomised stopping times assigning to each path not a specific time to stop but a 'stopping distribution'; viewed from this perspective, an MVM effectively corresponds to the *conditional distribution* of a *randomised stopping time*.

To formalise the above, we here introduce a class of adapted but not necessarily terminating MVMs. For a general account on measure-valued martingales we refer to, for example, [14, 25, 37]; see also [7, 13] for recent applications. In Appendix A, we provide various auxiliary results on the convergence and continuity of MVMs that will be used in the subsequent analysis.

DEFINITION 2.2. Given a filtered probability space supporting a $\mathcal{P}$-valued adapted process $\xi = (\xi_t)_{t\geq 0}$, we say that:

- ξ is a *measure-valued martingale* (MVM) if $\xi(\varphi)$ is a martingale for $\varphi \in C_b$;
- ξ is *continuous* (resp. *càdlàg*) if it is almost surely continuous (resp. càdlàg) w.r.t. $\mathcal{W}_1$;
- ξ is an *adapted* MVM if an MVM and $\xi_t([0,s]) = \xi_s([0,s])$ a.s., for all $t \geq s \geq 0$.

We denote by $\mathsf{MVM}(\mu)$ the set of continuous adapted MVMs ξ with $\xi_0 = \mu$; we emphasise that an MVM ξ being adapted refers to a stronger property than ξ being an adapted process. Moreover, for any given MVM ξ, we define the associated real-valued process

$$A_t^\xi(\omega) := \inf\{\xi_q([0,q])(\omega) : t < q \in \mathbb{Q}\}, \quad t \geq 0;$$

the process A^ξ is notably nondecreasing and right-continuous by definition and therefore specifies a (random) cumulative distribution function.

REMARK 2.3 (Adapted MVMs). (i) By the monotone class theorem, a $\mathcal{P}$-valued process is an MVM, if and only if, $\xi(\varphi)$ is a martingale for any $\varphi \in b\mathcal{B}$. More pertinently, given an MVM, $\xi(\varphi)$ is a martingale for every measurable function with $\xi_0(|\varphi|) < \infty$; see Remark 2 in [13]. In particular, $\xi(\varphi)$ is a martingale for any $\varphi \in C$ asymptotically of at most linear growth. (ii) Our notion of continuity is notably different from, for example, [13]; it is, however, natural since when the filtration satisfies the usual conditions, according to Lemma A.1, every MVM admits a version which is right-continuous in the topology induced by $\mathcal{W}_1$. (iii) For an adapted MVM, by the monotone class theorem, $\xi_t(B) = \xi_s(B)$ a.s., $t \geq s$, for any $B \in \mathcal{B}([0,s])$. (iv) For an adapted MVM, $\xi_t([0,s]) = A_s^\xi$ a.s., $t \geq s$; cf. the proof of Lemma 3.9 below.

For the measure-valued martingale formulation of the problem, we consider a Brownian setup. To this end, let $\Omega = C_0(\mathbb{R}_+)$ and denote the canonical process on this space by $(B_t)_{t\geq 0}$ and the Wiener measure by $\mathbb{W}$. Further, let the filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ be the usual augmentation of the canonical filtration and set $\mathcal{F} = \bigvee_{t\geq 0} \mathcal{F}_t$; whenever using the notation $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{W})$, we implicitly refer to this setup. The MVM-formulation of the optimal stopping problem then takes the following form.

PROBLEM 2.4. Given $\mu \in \mathcal{P}$ and the space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{W})$, maximise

$$\mathbb{E}\left[\int_0^\infty c(B_{\cdot\wedge s}, s)\,\mathrm{d}A_s^\xi\right] \quad \text{over } \xi \in \mathsf{MVM}(\mu). \tag{2.4}$$

The integral inside the expectation is to be understood in the usual sense of pathwise Lebesgue–Stieltjes integration. In the next subsection (see Proposition 2.5 below), we establish the equivalence between this problem and our original problem of study.

2.2. *The equivalence between the two formulations.* Recall that for our optimal stopping problem we require the filtration to be rich enough to support not only a Brownian motion but also an independent uniformly distributed random variable. As already pointed out, problem (2.1) is therefore equivalent to its weakly formulated counterpart. We first recall some facts on equivalent weak formulations before proving the equivalence between the original problem and the measure-valued martingale version provided in Problem 2.4.

Preliminaries on the weak problem formulation. For the study of weak formulations of optimal stopping problems, one often makes use of the notion of randomised stopping times. The notion has a long history; see, for example, [4, 15, 21, 29, 34], and notably recently appeared in [6] and [8] (see also [18, 19, 24, 27] for related formulations).

To specify the setup, recall the definition of the space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{W})$. Moreover, given the product space $C_0(\mathbb{R}_+) \times \mathbb{R}_+$, for $\mu \in \mathcal{P}$, denote by $\mathsf{Cpl}(\mathbb{W}, \mu)$ the set of couplings between $\mathbb{W}$ and μ; that is, the set of probability measures on the product space $C_0(\mathbb{R}_+) \times \mathbb{R}_+$ with marginal laws $\mathbb{W}$ and μ, respectively. The set of randomised stopping times with pre-specified law μ, is then given by[1]

$$\mathsf{RST}(\mu) := \{\gamma \in \mathsf{Cpl}(\mathbb{W}, \mu) : \omega \mapsto \gamma_\omega([0, t]) \text{ is } \mathcal{F}_t\text{-measurable, for } t \geq 0\},$$

where $(\gamma_\omega)_{\omega\in C_0(\mathbb{R}_+)}$ denotes the disintegration of γ in the first coordinate.

The distribution-constrained optimal stopping problem can then be formulated as an optimisation problem over such randomised stopping times; specifically, denoting by (B, T) the canonical pair on $C_0(\mathbb{R}_+) \times \mathbb{R}_+$, problem (2.1) is equivalent to the problem of optimising

$$\mathbb{E}^\gamma[c(B, T)] \quad \text{over } \gamma \in \mathsf{RST}(\mu). \tag{2.5}$$

Indeed, any given stopping time induces a measure in $\mathsf{RST}(\mu)$. Conversely, given a randomised stopping time, since $(\Omega, \mathcal{G}, \mathbb{G}, \mathbb{P})$ supports not only a Brownian motion but also an independent uniformly distributed $\mathcal{G}_0$-measurable random variable, we may construct a stopping time τ in this space such that (B, τ) has precisely that given law; see Lemma 3.11 in [6]. Since randomised stopping times are characterised by their disintegration kernels, working

[1] In [6], sub-probability measures γ with $\mathrm{proj}_{C_0(\mathbb{R}_+)}\gamma \leq \mathbb{W}$ are considered. However, since we here require $\mathrm{proj}_{\mathbb{R}_+}\gamma = \mu$, where $\mu \in \mathcal{P}$ has mass 1, it follows that so has γ and it suffices to restrict to so-called 'finite RST's' for which $\mathrm{proj}_{C_0(\mathbb{R}_+)}\gamma = \mathbb{W}$; see also p. 10 in [8].

on the fixed probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{W})$ and still denoting by B the canonical process on $\Omega = C_0(\mathbb{R}_+)$, the problem can also be formulated as optimising

$$\mathbb{E}\left[\int_0^\infty c(B_{\cdot\wedge s}, s)\,\mathrm{d}A_s^\gamma\right] \quad \text{over } \gamma \in \mathsf{RST}(\mu), \tag{2.6}$$

where $A^\gamma(\omega)$ denotes the cumulative distribution function associated with γ_ω for each $\omega \in C_0(\mathbb{R}_+)$; that is, $A_t^\gamma(\omega) := \gamma_\omega([0,t])$, $t \geq 0$. This formulation emphasises the intuition behind randomised stopping times: while a standard stopping time assigns to each path a single time to stop, a stopping time depending on some external randomisation may be understood as assigning to each path a distribution specifying the probability to stop at various times.

The above implies in particular that the specific choice of the probability space $(\Omega, \mathcal{G}, \mathbb{G}, \mathbb{P})$ has in fact no bearing on the problem. Hence, writing $\mathcal{A}(\mu)$ for the class of tuples $\alpha = (\Omega^\alpha, \mathcal{G}^\alpha, \mathbb{G}^\alpha, \mathbb{P}^\alpha, W^\alpha, \tau^\alpha)$, with $(\Omega^\alpha, \mathcal{G}^\alpha, \mathbb{G}^\alpha, \mathbb{P}^\alpha)$ a probability space in which W^α is a Brownian motion and τ^α is a stopping time with $\tau^\alpha \sim \mu$, the value of the problem remains the same if optimising

$$\mathbb{E}^\alpha\big[c\big(W^\alpha_{\cdot\wedge\tau^\alpha}, \tau^\alpha\big)\big] \quad \text{over } \alpha \in \mathcal{A}(\mu);$$

this illustrates that we are indeed dealing with a weak problem formulation.

Since the above formulations are all equivalent, we are free to switch between them and in the sequel we will consider the formulation most convenient in each particular situation.

The equivalence to the MVM formulation. We are now ready to establish the equivalence between our original problem of study and the measure-valued martingale formulation; the proof relies on identifying disintegrated randomised stopping times with measure-valued martingales specifying their conditional law.

PROPOSITION 2.5. *Problem 2.4 and the optimisation problem in (2.1) are equivalent in the sense that their values coincide and from any optimiser to the former problem we may construct an optimiser to the latter and vice versa.*

PROOF. Recall that problem (2.1) is equivalent to problem (2.6). Let $\gamma \in \mathsf{RST}(\mu)$; we write γ also for its disintegration kernel in the first variable $(\gamma_\omega)_{\omega \in C_0(\mathbb{R}_+)}$. Invoking a Getoor-type result, cf. [20], we may then construct an MVM returning the conditional distribution of this random measure on the space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{W})$. More precisely, applying Theorem 1.3 in [25] together with its proof, since $\mu \in \mathcal{P}$, we deduce the existence of a $\mathcal{P}$-valued process $\xi = (\xi_t)_{t\geq 0}$ such that $\xi(\varphi)$ is a version of the optional projection of $\gamma(\varphi)$, for every $\varphi \in b\mathcal{B}$, where $\gamma(\varphi)$ denotes the random variable $\omega \mapsto \gamma_\omega(\varphi)$; since the filtration is right-continuous, this implies that

$$\xi_t(\varphi) = \mathbb{E}\big[\gamma(\varphi)|\mathcal{F}_t\big] \quad \text{a.s.}, t \geq 0, \varphi \in b\mathcal{B},$$

and the process ξ is thus an MVM. In consequence, since the filtration is the augmented one generated by a Brownian motion, according to Lemma A.1, there exists a version of ξ such that $\xi(\varphi)$ is continuous a.s. for every $\varphi \in b\mathcal{B}$ and every $\varphi \in C$ asymptotically of at most linear growth; we choose this version. In particular, for any such φ, the process $\xi_t(\varphi)$, $t \geq 0$, is then indistinguishable from the martingale $\mathbb{E}[\gamma(\varphi)|\mathcal{F}_t]$, $t \geq 0$. Moreover, ξ is a.s. continuous in the topology induced by $\mathcal{W}_1$; cf. the proof of Lemma A.1. Since $\gamma \in \mathsf{RST}(\mu)$, we also have that $\gamma([0,t])$ is $\mathcal{F}_t$-measurable, which implies that

$$\xi_u\big([0,t]\big) = \mathbb{E}\big[\gamma\big([0,t]\big)|\mathcal{F}_u\big] = \gamma\big([0,t]\big) \quad \text{a.s.}, u \geq t.$$

The thus defined process ξ is therefore a continuous adapted MVM. Moreover, the fact that $\gamma \in \mathsf{CPL}(\mathbb{W}, \mu)$ implies that

$$\mu(B) = \gamma\big(C_0(\mathbb{R}_+) \times B\big) = \int_{C_0(\mathbb{R}_+)} \gamma_\omega(B)\, \mathrm{d}\mathbb{W}(\omega) = \mathbb{E}\big[\gamma(B)\big], \quad B \in \mathcal{B};$$

hence $\xi_0 = \mu$ by Blumenthal's zero-one law. It remains to argue that the objective functions in (2.4) and (2.6) coincide when evaluated at ξ and γ, respectively. To this end, note that there exists a null set $\mathcal{N}$ such that, for $\omega \notin \mathcal{N}$, $\xi_q([0, q])(\omega) = \gamma_\omega([0, q])$ for every $q \in \mathbb{Q}$; without loss of generality, suppose that $q \mapsto \gamma_\omega([0, q])$ is nondecreasing off $\mathcal{N}$. Then,

$$A_t^\xi(\omega) = \inf_{t<q\in\mathbb{Q}} \xi_q\big([0, q]\big)(\omega) = \inf_{t<q\in\mathbb{Q}} \gamma_\omega\big([0, q]\big) = \gamma_\omega\big([0, t]\big), \quad t \geq 0, \omega \notin \mathcal{N},$$

which completes the proof of the claim.

Conversely, consider an arbitrary $\xi \in \mathsf{MVM}(\mu)$. Invoking again a Getoor-type result we may construct a limiting object. More precisely, according to Lemma A.2, there exists a $\mathcal{P}$-valued random measure ξ_∞ such that $\lim_{t\to\infty} \xi_t(\varphi) = \xi_\infty(\varphi)$ a.s., for any $\varphi \in b\mathcal{B}$. We define a measure γ on $C_0(\mathbb{R}_+) \times \mathbb{R}_+$ via the following specification for $c \in b\mathcal{B}(C_0(\mathbb{R}_+) \times \mathbb{R}_+)$:

$$\gamma(c) = \iint c(\omega, x)\xi_\infty(\mathrm{d}x)(\omega)\, \mathrm{d}\mathbb{W}(\omega).$$

Since $\gamma(C_0(\mathbb{R}_+) \times B) = \mathbb{E}[\xi_\infty(B)] = \xi_0(B)$, for all $B \in \mathcal{B}$, we have that $\gamma \in \mathsf{CPL}(\mathbb{W}, \mu)$. Further, recall that $\xi_u([0, t]) = \xi_t([0, t])$ a.s. for $u \geq t$. Hence, there exists some null set, say again $\mathcal{N}$, off which $\xi_\infty([0, q])(\omega) = \xi_q([0, q])(\omega)$ for every $q \in \mathbb{Q}$; in consequence

$$A_t^\xi(\omega) = \inf_{t<q\in\mathbb{Q}} \xi_\infty\big([0, q]\big)(\omega) = \xi_\infty\big([0, t]\big)(\omega), \quad t \geq 0, \omega \notin \mathcal{N}. \tag{2.7}$$

Since the disintegration of γ is $\mathbb{W}$-a.e. unique, this implies that the objective functions in (2.4) and (2.6) coincide when evaluated at ξ and γ, respectively. Moreover, to ensure that $\gamma \in \mathsf{RST}(\mu)$, it now suffices to argue that A_t^ξ is $\mathcal{F}_t$-measurable for $t \geq 0$. To see this, note that off $\mathcal{N}$, we have $\xi_q([0, q]) \leq \xi_r([0, r])$ for any $q \leq r$ with $q, r \in \mathbb{Q}$; hence, for any $\varepsilon > 0$, it holds (up to a null set) that

$$\big\{A_t^\xi \geq a\big\} = \bigcap_{\substack{t<q\\ q\in\mathbb{Q}}} \big\{\xi_q\big([0, q]\big) \geq a\big\} = \bigcap_{\substack{t<q<t+\varepsilon\\ q\in\mathbb{Q}}} \big\{\xi_q\big([0, q]\big) \geq a\big\} \in \mathcal{F}_{t+\varepsilon}, \tag{2.8}$$

which completes the proof. □

REMARK 2.6. Consider a filtered probability space $(\Omega, \mathcal{H}, \mathbb{H}, \mathbb{P})$ satisfying the usual conditions and supporting a Brownian motion W; in turn, let ξ be a càdlàg adapted MVM in this space with $\xi_0 = \mu$. By use of analogous arguments to those employed above, in the space $(\overline{\Omega}, \overline{\mathcal{H}}, \overline{\mathbb{H}}, \overline{\mathbb{P}})$ obtained by initially enlarging the former space by an independent uniformly distributed random variable, we may then construct a stopping time $\tau \sim \mu$ such that $\overline{\mathbb{E}}[c(W_{\cdot\wedge\tau}, \tau)] = \mathbb{E}[\int c(W_{\cdot\wedge s}, s)\, \mathrm{d}A_s^\xi]$. In consequence, from Proposition 2.5 and the independence of the specific choice of probability space for problem (2.1), we obtain that

$$v(\mu) = \sup_{\mathcal{K}(\mu)} \mathbb{E}^\kappa\bigg[\int c\big(W_{\cdot\wedge s}^\kappa, s\big)\, \mathrm{d}A_s^{\xi^\kappa}\bigg],$$

where we denote by $\mathcal{K}(\mu)$ the class of tuples $\kappa = (\Omega^\kappa, \mathcal{H}^\kappa, \mathbb{H}^\kappa, \mathbb{P}^\kappa, W^\kappa, \xi^\kappa)$ such that $(\Omega^\kappa, \mathcal{H}^\kappa, \mathbb{H}^\kappa, \mathbb{P}^\kappa)$ is a filtered probability space in which W^κ is a Brownian motion and ξ^κ is a càdlàg adapted MVM with $\xi_0^\kappa = \mu$.

3. The dynamic programming principle. The reformulation established in the previous section transformed the distributional constraint into an initial condition for an additional state process and we are now facing an optimisation problem over adapted MVMs starting off in μ. We may therefore address the problem by the use of dynamic programming arguments. Such methods offer a powerful tool for characterising the solution as well as for enabling the use of numerical methods. We here formally establish this link by proving that the dynamic programming principle holds for this problem.

3.1. *The conditional problem and the DPP.*

The conditional problem and the value function. The dynamic programming approach relies on embedding the problem of original interest into a family of conditional problems. To this end, focusing on the distribution-constrained optimal stopping problem in its equivalent form (2.4), we define the associated conditional problem and its value function as follows: Continuing to work on the space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{W})$, let $v : \mathbb{R}_+ \times C_0(\mathbb{R}_+) \times \mathcal{P} \to \mathbb{R}$ be given by

$$v(t, \mathbf{w}, \mu) := \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\left[\int_t^\infty c\big(W^{t,\mathbf{w}}_{\cdot\wedge u}, u\big)\,\mathrm{d}A^\xi_u\right],$$

where $W^{t,\mathbf{w}}_s$, $s \geq 0$, denotes the solution to the SDE

$$\mathrm{d}W^{t,\mathbf{w}}_s = \mathrm{d}B_s, \quad s \in [t, \infty), \qquad W^{t,\mathbf{w}}_s = \mathbf{w}_s, \quad \text{a.s.}, s \in [0, t],$$

and $\mathsf{MVM}^t(\mu)$ denotes the set of continuous adapted MVMs with $\xi_t = \mu$ a.s.

The value function notably depends on $\mathbf{w}$ up to time t only in the sense that $v(t, \mathbf{w}, \mu) = v(t, \mathbf{w}_{\cdot\wedge t}, \mu)$. In particular, according to Proposition 2.5, we have that $v(\mu) = v(0, \mathbf{w}, \mu)$, for any $\mathbf{w} \in C_0(\mathbb{R}_+)$. Moreover, if $\mu \in \mathcal{P}$ has an atom at time t, it will notably not contribute to the value of $v(t, \mathbf{w}, \mu)$; put differently, $v(t, \mathbf{w}, \mu)$ depends on the restriction of μ to (t, ∞) only. For this reason we introduce the notation $\mu^{(t}$ for the (re-weighted) restriction of μ to (t, ∞) given by

$$\mu^{(t}(A) := \frac{\mu(A \cap (t, \infty))}{\mu((t, \infty))}, \quad A \in \mathcal{B}; \tag{3.1}$$

we use the convention $\mu^{(t} = \mu$ if $\mu((t, \infty)) = 0$. From the definition of v, it then follows that $v(t, \mathbf{w}, \mu) = \mu((t, \infty))v(t, \mathbf{w}, \mu^{(t})$.

The dynamic programming principle. We are now ready to present our main result which establishes the dynamic programming principle for the MVM-formulation of the optimal stopping problem; the proof is provided in Section 3.3 below.

THEOREM 3.1 (Dynamic programming principle). *Let $(t, \mathbf{w}, \mu) \in \mathbb{R}_+ \times C_0(\mathbb{R}_+) \times \mathcal{P}$. For any $\mathbb{F}$-stopping time θ with values in (t, ∞), it then holds that*

$$v(t, \mathbf{w}, \mu) = \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\left[\int_t^\theta c\big(W^{t,\mathbf{w}}_{\cdot\wedge u}, u\big)\,\mathrm{d}A^\xi_u + v\Big(\theta, W^{t,\mathbf{w}}_{\cdot\wedge\theta}, \xi_\theta\Big)\right]. \tag{3.2}$$

REMARK 3.2. In the formulation of the DPP, for the same reasons as above, if an MVM $\xi \in \mathsf{MVM}^t(\mu)$ has an atom at θ it will contribute to the objective function via the integral from t to θ rather than via the value function evaluated at θ on the right-hand side of (3.2). Using the notation introduced in (3.1), we therefore have the following equivalent formulation of the DPP:

$$v(t, \mathbf{w}, \mu) = \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\left[\int_t^\theta c\big(W^{t,\mathbf{w}}_{\cdot\wedge u}, u\big)\,\mathrm{d}A^\xi_u + \xi_\theta\big((\theta, \infty)\big)v\Big(\theta, W^{t,\mathbf{w}}_{\cdot\wedge\theta}, \xi^{(\theta}_\theta\Big)\right].$$

As always when the dynamic programming principle holds, it follows that the problem can be solved by backward induction. When the distributional constraint is supported on a finite number of atoms, the problem thus amounts to solving a finite number of subsequent one-period problems (see Corollary 3.6 below). More generally, for an arbitrary distributional constraint, the DPP forms the basis for solving the problem via stochastic control methods. While we leave the detailed development of such arguments for future study, we note that MVMs are the canonical objects for guaranteeing preservation of a distributional constraint when letting random distributions evolve over time; they are thus indispensable for the development of such a theory. We also note that already the DPP itself ensures that the martingale optimality principle holds (see [17]): the process $(\omega, t) \mapsto v(t, \omega, \xi_t(\omega))$ is thus a supermartingale when evaluated for any MVM $\xi \in \mathsf{MVM}(\mu)$, and a martingale—if and only if—evaluated at an optimal MVM; this offers an alternative characterisation of optimisers to our original problem of study.

To understand the role played by the dynamic programming principle, it is illustrative to consider the following schematic example where the conceptual idea behind our approach is pinned down in a discrete setup.

EXAMPLE 3.3 (Schematic illustration). Consider discrete time points $n = 0, 1, 2, 3$, and let us replace the Brownian motion with a simple random walk; that is, $|\Omega| = 8$, and $(B_n(\omega_1))_{n=0}^3 = (0, 1, 2, 3)$, $(B_n(\omega_2))_{n=0}^3 = (0, 1, 2, 1)$ etc. Let $\mu = \mu^1\delta_1 + \mu^2\delta_2 + \mu^3\delta_3$, and consider a cost function c depending on the path of B and the stopping time. The analogue of our weakly formulated stopping problem then amounts to optimally associate with each ω_i a probability distribution $(\gamma^n_{\omega_i})_{n=1}^3$ such that γ^n is $\sigma(B_1, \dots, B_n)$-measurable and $\frac{1}{8}\sum_{i=1}^8 \gamma^n_{\omega_i} = \mu^n$, for $n = 1, 2, 3$. The idea is to do this in a recursive way. To this end, suppose that given (B_1, B_2), there is remaining mass ξ_2^3 to assign to the last atom. There is only one way of doing this: no matter whether the random walk goes up or down, we must let $\xi_3^3 = \xi_2^3$; the contribution to the objective is given by

$$v\big((B_1, B_2), \xi_2^3\big) = \frac{1}{2}\sum_{i=\pm 1} c\big((B_1, B_2, B_2 + i), 3\big)\xi_3^3.$$

Taking one step back, given B_1, consider having mass (ξ_1^2, ξ_1^3) to assign to the last two atoms; the problem then amounts to find nonnegative pairs $(\xi_2^2, \xi_2^3)(i)$, $i = \pm 1$, which attain the following maximum:

$$v\big(B_1, (\xi_1^2, \xi_1^3)\big) = \max_{(\xi_2^2, \xi_2^3)} \frac{1}{2}\sum_{i=\pm 1} c\big((B_1, B_1 + i), 2\big)\xi_2^2(i) + v\big((B_1, B_1 + i), \xi_2^3(i)\big), \tag{3.3}$$

subject to the constraint

$$\xi_2^2(i) + \xi_2^3(i) = \xi_1^2 + \xi_1^3, \quad i = \pm 1, \qquad \frac{1}{2}\big(\xi_2^n(+1) + \xi_2^n(-1)\big) = \xi_1^n, \quad n = 2, 3. \tag{3.4}$$

Because of the joint dependence on the path of B and the stopping time, (3.3)–(3.4) has a nontrivial solution specifying the best choice for the given constraint. In turn, given B_0 and μ, solving the analogue of (3.3)–(3.4) to obtain $(\xi_1^1, \xi_1^2, \xi_1^3)(i)$, $i = \pm 1$, and substituting this solution as input for the above problems, the optimal γ is then obtained from

$$(\gamma^n)_{n=1}^3 = (\xi_1^1, \xi_2^2, \xi_3^3).$$

In essence, in each time-step, we are facing the problem of "splitting the mass to be taken care of" between the two events corresponding to the random walk going up or down; this is done in an optimal way, taking into consideration the best value one can obtain for each of these scenarios when facing a certain constraint to embed. That we respect the "total mass that needs to be embedded" is guaranteed through the measure-valued martingale condition; cf. the second part of (3.4). Meanwhile, the adaptedness condition is satisfied since when incorporating the knowledge of, say B_2, we keep refining only how to "split the mass" assigned to the atoms at $n = 2$ and $n = 3$, while the mass reserved for $n = 1$ is "frozen" at this stage; cf. the first part of condition (3.4).

3.2. *The special cases of terminating and atomic MVMs.* We comment next on two particular cases: the case when the weak problem formulation is equivalent to the strong one and the case of atomic constraints.

The case of terminating MVMs and the strong problem formulation. Although we consider the weak problem formulation studied herein to be the natural one (in particular since it always admits a solution) we note that whenever problem (2.1) admits a strong solution—here referring to an optimal stopping time adapted to the filtration generated by the Brownian motion alone—Theorem 3.1 admits a stronger formulation. To specify the DPP in this case, we first formalise the notion of termination briefly discussed in Section 2.1. To this end, let

$$\mathcal{P}^s := \{\mu \in \mathcal{P} : \mu = \delta_y, y \in \mathbb{R}_+\}.$$

Further, recall from Lemma A.2 that for each MVM ξ there exists a limiting (random) measure ξ_∞ such that for any $\varphi \in b\mathcal{B}$, $\xi_t(\varphi) \to \xi_\infty(\varphi)$ a.s. as $t \to \infty$; notably $\xi_\infty \in \mathcal{P}$ a.s. and convergence holds also a.s. in the topology induced by $\mathcal{W}_1$. Following [13], we define the concept of termination as follows.

DEFINITION 3.4. An MVM ξ is *terminating* if $\xi_\infty \in \mathcal{P}^s$ almost surely; it is *finitely terminating* if $\inf\{t > 0 : \xi_t \in \mathcal{P}^s\}$ is finite almost surely.

A terminating MVM is adapted in the sense of Definition 2.2, if and only if, condition (2.2) holds; in consequence, any adapted terminating MVM is also finitely terminating. The associated stopping time, say τ^ξ, is given by (cf. (2.3) and (2.7)):

$$\tau^\xi = \inf\{t > 0 : A_t^\xi = 1\}.$$

We denote the set of continuous adapted and finitely terminating MVMs by $\mathsf{MVM}_{\text{term}}$, and define $\mathsf{MVM}^t_{\text{term}}(\mu)$ analogously to above.

COROLLARY 3.5 (DPP: strong formulation). *Let* $(t, \mathbf{w}, \mu) \in \mathbb{R}_+ \times C_0(\mathbb{R}_+) \times \mathcal{P}$. *Suppose that restricting in* (2.1) *to* $\tau \in \mathcal{T}(\mu)$ *which are stopping times in the filtration generated by the Brownian motion alone does not affect the value of the problem. For any* $\mathbb{F}$*-stopping time* θ *with values in* (t, ∞), *it then holds that*

$$v(t, \mathbf{w}, \mu) = \sup_{\xi \in \mathsf{MVM}^t_{\text{term}}(\mu)} \mathbb{E}\Big[c\big(W^{t,\mathbf{w}}_{\cdot\wedge\tau^\xi}, \tau^\xi\big)\mathbf{1}_{\{\tau^\xi \le \theta\}} + v\big(\theta, W^{t,\mathbf{w}}_{\cdot\wedge\theta}, \xi_\theta\big)\mathbf{1}_{\{\tau^\xi > \theta\}}\Big]. \tag{3.5}$$

We note that the assumptions of this corollary are satisfied, for example, when (2.1) admits a so-called barrier solution, which according to [8] happens whenever $c(\omega, t) = f \circ \omega(t)$ for some $f : \mathbb{R} \to \mathbb{R}$ with $f''' > 0$.

The case of atomic MVMs and atomic constraints. Consider the situation where the initial condition μ has support on a finite number of atoms $0 < t_1 < \cdots < t_r$, $r \in \mathbb{N}$. An MVM ξ with $\xi_0 = \mu$ must notably remain supported on the same atoms; the process $t \mapsto \xi_t(\{t_i\})$ is a real-valued martingale for each $i = 1, \ldots, r$. This implies that the problem reduces to a finite-dimensional one. To formalise this, let Δ^{k-1} be the standard $(k-1)$-simplex in $\mathbb{R}^k$, $k \in \mathbb{N}$; that is, $\Delta^{k-1} = \{(x_1, \ldots, x_k) \in [0,1]^k : x_1 + \cdots + x_k = 1\}$. Moreover, let $\mathcal{M}(\Delta^k)$ denote the set of martingales taking values in Δ^k. For $k \in \{1, \ldots, r\}$ and $t \in [t_{r-k}, t_{r-k+1})$, we then define $v_t : C_0(\mathbb{R}_+) \times \Delta^{k-1} \to \mathbb{R}$ by

$$v_t(\mathbf{w}, y_{(r-k+1):r}) := v(t, \mathbf{w}, \mu) \quad \text{with } \mu = \sum_{i=r-k+1}^{r} y_i \delta_{t_i}, \tag{3.6}$$

where we use the notation $y_{(r-k+1):r}$ for the vector $(y_{r-k+1}, \ldots, y_r) \in \Delta^{k-1}$ and similarly for vector-valued stochastic processes.

COROLLARY 3.6 (DPP: atomic constraints). *Let $0 < t_1 < \cdots < t_r$, $r \in \mathbb{N}$. For $k \in \{1, \ldots, r\}$, $y_{(r-k):r} \in \Delta^k$ and $\mathbf{w} \in C_0(\mathbb{R}_+)$, we then have that*

$$\begin{aligned} v_{t_{r-k-1}}(\mathbf{w}, y_{(r-k):r}) = \sup_{\substack{Y \in \mathcal{M}(\Delta^k) \\ Y_{t_{r-k-1}} = y_{(r-k):r}}} \mathbb{E}\Bigg[& Y^{r-k}_{t_{r-k}} c\big(W^{t,\mathbf{w}}_{\cdot \wedge t_{r-k}}, t_{r-k}\big) \\ & + \big(1 - Y^{r-k}_{t_{r-k}}\big) v_{t_{r-k}}\Bigg(W^{t,\mathbf{w}}_{\cdot \wedge t_{r-k}}, \frac{Y^{(r-k+1):r}_{t_{r-k}}}{1 - Y^{r-k}_{t_{r-k}}}\Bigg)\Bigg]. \end{aligned} \tag{3.7}$$

As already mentioned, in [5], the authors consider the distribution-constrained optimal stopping problem for atomic constraints and cost functions of the form $c(\omega_{\cdot \wedge s}, s) = f \circ \omega(s)$ for some Lipschitz function $f : \mathbb{R} \to \mathbb{R}$. A priori, they consider the strong formulation where the filtration is the one generated by the Brownian motion alone and for this problem they establish the DPP; that is, (3.5) which for atomic constraints takes the form (3.7) with the additional condition $Y^{r-k}_{t_{r-k}} \in \{0, 1\}$ a.s. In turn, they relax this additional condition and show that their value function also satisfies (3.7). Combining our Corollary 3.6 with their result, we obtain that for distributional constraints and cost functions of the form considered in [5], the value function resulting from the restriction to Brownian stopping times and the value function for the weak problem formulation coincide. This should come as no surprise for when considering non-path-dependent cost functions and atomic constraints, one has the possibility to let the Brownian motion run for a positive amount of time before stopping; even if a priori working with the filtration generated by the Brownian motion alone, additional randomisation can thus be obtained by conditioning on the past of the Brownian motion itself.

3.3. *Proof of the dynamic programming principle.*

Reformulation onto Mayer form. Since our objective is given in Lagrange form (with a reward function integrated over time), we first introduce an additional state variable (governing its accumulated value) in order to transform it onto Mayer form. Specifically, given $(t, \mathbf{w}, y) \in \mathbb{R}_+ \times C_0(\mathbb{R}_+) \times \mathbb{R}$ and $\xi \in \mathsf{MVM}$, we define the process

$$Y^{t,\mathbf{w},y}_u(\xi) := y + \int_t^{t \vee u} c\big(W^{t,\mathbf{w}}_{\cdot \wedge s}, s\big)\, \mathrm{d}A^{\xi}_s;$$

we note that it admits a well-defined limit which we denote by $Y_\infty^{t,\mathbf{w},y}(\xi)$. For $(t,\mathbf{w},y,\mu) \in \mathbb{R}_+ \times C_0(\mathbb{R}_+) \times \mathbb{R} \times \mathcal{P}$, we re-introduce the value function

$$v(t,\mathbf{w},y,\mu) := \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\big[Y_\infty^{t,\mathbf{w},y}(\xi)\big]. \tag{3.8}$$

To establish Theorem 3.1, we then need to show that for every $\mathbb{F}$-stopping time θ with values in (t,∞),

$$v(t,\mathbf{w},y,\mu) = \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\Big[v\Big(\theta, W_{\cdot\wedge\theta}^{t,\mathbf{w}}, Y_\theta^{t,\mathbf{w},y}(\xi), \xi_\theta\Big)\Big]. \tag{3.9}$$

The DPP formulated on the canonical path space. For our proof we consider the weak formulation of the DPP on the associated canonical path space; the canonical framework has previously successfully been used for the study of stochastic control problems in [18, 19], see also [38] and [31, 32]. We denote by $\mathbb{D}$ the set of càdlàg paths on $[0,\infty)$ taking values in $E := \mathbb{R} \times \mathbb{R} \times \mathcal{P}$, where we equip $\mathcal{P}$ with the topology induced by $\mathcal{W}_1$ and E with the product topology; this renders E a Polish space (cf. Theorem 6.18 in [36]) and using the Skorokhod topology on $\mathbb{D}$ it is a Polish space too. A generic path in $\mathbb{D}$ is denoted by ω and we use $X = (W, Y, \xi)$ for the coordinate process $X_t(\omega) = (W_t, Y_t, \xi_t)(\omega) = \omega(t)$. Further, we let $\mathbf{X} = C_0(\mathbb{R}_+) \times \mathbb{R} \times \mathcal{P}$, $\mathcal{X} = \mathcal{B}(\mathbf{X})$, denote a generic element in $\mathbf{X}$ by $\mathbf{x}$, and write $\mathbf{x}_t$ for the mapping from $\mathbb{R}_+ \times \mathbb{D}$ to $\mathbf{X}$ returning $\mathbf{x}_t(X) := (W_{\cdot\wedge t}, Y_t, \xi_t)$. Let $\mathbb{F}^0 = (\mathcal{F}_t^0)_{t\ge 0}$ denote the natural filtration generated by the coordinate process X and denote by $\mathfrak{P}$ the set of all probability measures on $\mathcal{B}(\mathbb{D})$. Given $(t,\mathbf{x}) \mathrel{\widehat{=}} (t,\mathbf{w},y,\mu) \in \mathbb{R}_+ \times \mathbf{X}$, we then introduce the set $\mathfrak{P}_{t,\mathbf{x}}$ of probability measures $\mathbb{P}$ in $\mathfrak{P}$ for which:

i. $\mathbb{P}$-a.s., $W_s = \mathbf{w}_s$ for $s \le t$, $Y_t = y$ and $\xi_t = \mu$;
ii. $(W_u - W_t)_{u\ge t}$ is an $(\mathbb{F}^0, \mathbb{P})$-Brownian motion;
iii. $(\xi_u)_{u\ge t}$ is an adapted $(\mathbb{F}^0, \mathbb{P})$-MVM;
iv. $\mathbb{P}$-a.s., $Y_u - Y_t = \int_t^u c(W_{\cdot\wedge s}, s)\,\mathrm{d}A_s^\xi$, $u > t$.

Note that there exists a measurable functional $Y_\infty : \mathbb{D} \to \mathbb{R}_+$ such that $Y_\infty(\omega) = \lim_{t\to\infty} Y_t(\omega)$ whenever the limit exists (recall that $c(\cdot,\cdot)$ is measurable and see Lemma 3.12 in [38]); in particular, for any $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$, we have $Y_\infty = \lim_{t\to\infty} Y_t$, $\mathbb{P}$-a.s. We define

$$v(t,\mathbf{x}) := \sup_{\mathbb{P}\in\mathfrak{P}_{t,\mathbf{x}}} \mathbb{E}^{\mathbb{P}}[Y_\infty], \quad (t,\mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X}; \tag{3.10}$$

with a slight abuse of notation we will also write $v(\theta, X_{\cdot\wedge\theta})$ for $v(\theta, \mathbf{x}_\theta(X))$.

LEMMA 3.7. *The value functions given by* (3.8) *and* (3.10) *coincide*; *that is*, *for all* $(t,\mathbf{x}) = (t,(\mathbf{w},y,\mu)) \in \mathbb{R}_+ \times \mathbf{X}$, $v(t,\mathbf{x}) = v(t,\mathbf{w},y,\mu)$.

PROOF. Let $(t,\mathbf{x}) = (t,(\mathbf{w},y,\mu)) \in \mathbb{R}_+ \times \mathbf{X}$. In turn, let $\mathcal{K}(t,\mathbf{x})$ denote the class of tuples $\kappa = (\Omega^\kappa, \mathcal{F}^\kappa, \mathbb{F}^\kappa, \mathbb{P}^\kappa, W^\kappa, Y^\kappa, \xi^\kappa)$ such that $(\Omega^\kappa, \mathcal{F}^\kappa, \mathbb{F}^\kappa, \mathbb{P}^\kappa)$ is a filtered probability space in which ξ^κ is a càdlàg adapted MVM with $\xi_t^\kappa = \mu$ a.s.; W^κ solves $\mathrm{d}W_s^\kappa = \mathrm{d}B_s^\kappa$, $s \in [t,\infty)$, with $W_s^\kappa = \mathbf{w}_s$ a.s., $s \in [0,t]$, for some $(\mathbb{F}^\kappa, \mathbb{P}^\kappa)$-Brownian motion B^κ; and $Y_u^\kappa = y + \int_t^u c(W_{\cdot\wedge s}^\kappa, s)\,\mathrm{d}A_s^{\xi^\kappa}$, $u \ge t$. According to Remark 2.6, we then have that

$$v(t,\mathbf{w},y,\mu) = \sup_{\mathcal{K}(t,\mathbf{x})} \mathbb{E}^\kappa\big[Y_\infty^\kappa\big].$$

Note that each tuple $\kappa \in \mathcal{K}(t,\mathbf{x})$ induces a term in $\mathfrak{P}_{t,\mathbf{x}}$. Conversely, any probability measure $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$ together with the space $(\mathbb{D}, \mathcal{B}(\mathbb{D}), \mathbb{F}^{\mathbb{P}})$ and the canonical process (W, Y, ξ) produces such a tuple for if properties (ii) and (iii) hold with respect to $\mathbb{F}^0$ they hold also with respect to the augmented filtration $\mathbb{F}^{\mathbb{P}}$. This implies that $\sup_{\mathcal{K}(t,\mathbf{x})} \mathbb{E}^{\kappa}[Y^{\kappa}_{\infty}] = v(t,\mathbf{x})$, which completes the proof. $\square$

The DPP for the problem introduced in (3.10) takes the following form.

THEOREM 3.8. *For all $(t,\mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X}$, and any $\mathbb{F}^0$-stopping time θ with values in (t,∞), it holds that*

$$v(t,\mathbf{x}) = \sup_{\mathbb{P}\in\mathfrak{P}_{t,\mathbf{x}}} \mathbb{E}^{\mathbb{P}}\big[v(\theta, X_{\cdot\wedge\theta})\big].$$

In order to obtain Theorem 3.1, it then suffices to establish Theorem 3.8. Indeed, applying Lemma 3.7 combined with the same arguments as used for its proof, it follows directly from the latter result that relation (3.9), and thus relation (3.2), holds for any θ which is a stopping time in the raw filtration generated by the Brownian motion. Theorem 3.1 then follows since any $\mathbb{F}$-stopping time is predictable and since for any $\mathbb{F}$-predictable time θ there exists a predictable time $\bar{\theta}$ in the raw filtration with $\bar{\theta} = \theta$ a.s.

The remainder of this section is devoted to the proof of Theorem 3.8.

Analyticity of $\mathfrak{P}_{t,\mathbf{x}}$ and forward and backward concatenation properties. To establish the above DPP, we use the same method of proof as employed in, for example, [32] or [19]; the key step is to establish certain analyticity and stability properties.

We first show that the graph corresponding to $\mathfrak{P}_{t,\mathbf{x}}$ is analytic:

LEMMA 3.9. *The set $\Gamma := \{(t,\mathbf{x},\mathbb{P}) : (t,\mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X}, \mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}\}$ is a Borel set.*

PROOF. For $0 < q < r$, $\Psi = \mathbf{1}_A$ with $A \in \mathcal{F}^0_q$, $\psi \in C^2_b$ and $\varphi \in C_b$, with C^2_b denoting the set of functions with continuous and bounded derivatives of second order, we consider the following subsets of $\mathbb{R}_+ \times \mathbf{X} \times \mathfrak{P}$:

- ◇ $\{(t,\mathbf{x},\mathbb{P}) : \mathbb{P}(W_{r\wedge t} = \mathbf{w}(r\wedge t), Y_t = y, \xi_t = \mu) = 1\}$;
- ◇ $\{(t,\mathbf{x},\mathbb{P}) : \mathbb{E}^{\mathbb{P}}[\Psi(\psi(W_{t\vee r}) - \psi(W_{t\vee q}) - \frac{1}{2}\int_{t\vee q}^{t\vee r} \psi''(W_u)\,\mathrm{d}u)] = 0\}$;
- ◇ $\{(t,\mathbf{x},\mathbb{P}) : \mathbb{E}^{\mathbb{P}}[\Psi(\xi_{t\vee r}(\varphi) - \xi_{t\vee q}(\varphi))] = 0\}$;
- ◇ $\{(t,\mathbf{x},\mathbb{P}) : \mathbb{P}(\xi_{t\vee r}([0,q]) = \xi_{t\vee q}([0,q]) = 1\}$;
- ◇ $\{(t,\mathbf{x},\mathbb{P}) : \mathbb{P}(Y_{t\vee r} - Y_{t\vee q} = \int_{t\vee q}^{t\vee r} c(W_{\cdot\wedge u}, u)\,\mathrm{d}A^{\xi}_u) = 1\}$.

The above sets are all Borel measurable. Furthermore, we claim that Γ is the intersection of countably many sets of the above form and thus itself a Borel set. More precisely, we claim that Γ is the intersection of the above sets when $0 < q < r$ are allowed to vary among all rational numbers; A among a countable family generating $\mathcal{F}^0_q$ for each q; and ψ and φ among countable subsets generating, respectively, C^2_b and C_b via pointwise convergence. Indeed, by use of the monotone class theorem and dominated convergence we then obtain that the second and third properties hold for any $\Psi \in \mathcal{F}^0_q$, $\psi \in C^2_b$ and $\varphi \in C_b$, for rational $0 < q < r$; since W and $\xi(\varphi)$, $\varphi \in C_b$, are càdlàg this implies that $(W_u)_{u\geq t}$ is a Brownian motion and that $(\xi_u)_{u\geq t}$ is an MVM. To argue the adaptedness property, without loss of generality, let $t = 0$ and let $\mathbb{P}$ belong to the above intersection of sets. Using the MVM-property, Remark 2.3(i) and the fourth property above, we obtain for any $u \geq q$ with $q \in \mathbb{Q}$,

$$\xi_u([0,q]) = \mathbb{E}[\xi_r([0,q])|\mathcal{F}^0_u] = \mathbb{E}[\xi_q([0,q])|\mathcal{F}^0_u] = \xi_q([0,q]), \quad \text{a.s.},$$

where we picked $r \geq u$ with $r \in \mathbb{Q}$. In turn, recalling that there exists a null set off which $q \mapsto \xi_q([0,q])$ is nondecreasing, we obtain for any $u > s$,

$$\xi_u([0,s]) = \inf_{q\in(s,u)\cap\mathbb{Q}} \xi_u([0,q]) = \inf_{q\in(s,u)\cap\mathbb{Q}} \xi_q([0,q]) = A^{\xi}_s, \quad \text{a.s.}$$

Let $\bar{\mathbb{F}} = (\bar{\mathcal{F}}_u)_{u\geq 0}$ be the augmentation of $\mathbb{F}^0$. By Lévy's downward theorem, the càdlàg $(\mathbb{F}^0, \mathbb{P})$-martingales $\xi(\varphi)$, $\varphi \in C_b$, are also $(\bar{\mathbb{F}}, \mathbb{P})$-martingales and ξ is thus an $(\bar{\mathbb{F}}, \mathbb{P})$-MVM. Using again Remark 2.3(i) and the fact that A_s^ξ is $\bar{\mathcal{F}}_s$-measurable (cf. (2.8)), we obtain

$$\xi_s([0, s]) = \mathbb{E}[\xi_u([0, s])|\bar{\mathcal{F}}_s] = \mathbb{E}[A_s^\xi|\bar{\mathcal{F}}_s] = A_s^\xi, \quad \text{a.s.;}$$

hence, $\xi_u([0, s]) = \xi_s([0, s])$ a.s., $u \geq s$, and ξ is an adapted MVM. □

Next, we establish that certain consistency properties hold when concatenating measures in $\mathfrak{P}_{t,\mathbf{x}}$. To this end, we first introduce some further notation: A map $\mathbb{Q} : \mathbb{D} \times \mathcal{B}(\mathbb{D}) \to [0, 1]$ is called a (universally) measurable kernel if (i) $\mathbb{Q}(\omega, \cdot) \in \mathfrak{P}$ for all $\omega \in \mathbb{D}$, and (ii) $\mathbb{D} \ni \omega \mapsto \mathbb{Q}(\omega, A)$ is (universally) measurable for all $A \in \mathcal{B}(\mathbb{D})$. Recall that the universal σ-algebra is the intersection of the completions of the Borel σ-algebra over all probability measures on the space, and that universally measurable functions are integrable with respect to any such probability measure; we will use the superscript U to refer to the universal completion. We write $\mathbb{Q}_\omega$ for the probability measure $\mathbb{Q}(\omega, \cdot)$ and interpret $\mathbb{Q}$ as a (universally) measurable map from $\mathbb{D}$ to $\mathfrak{P}$. Further, given a random time θ and two paths $\omega, \omega' \in \mathbb{D}$ with $X_{\theta(\omega)}(\omega) = X_{\theta(\omega)}(\omega')$, we define the concatenation $\omega \otimes_\theta \omega'$ to be the element of $\mathbb{D}$ given by

$$X_t(\omega \otimes_\theta \omega') = \mathbf{1}_{\{t<\theta(\omega)\}} X_t(\omega) + \mathbf{1}_{\{t\geq\theta(\omega)\}} X_t(\omega').$$

Given a probability measure $\mathbb{P} \in \mathfrak{P}$ and a universally measurable kernel $\mathbb{Q}$, we then define the concatenation $\mathbb{P} \otimes_\theta \mathbb{Q}$ as the probability measure in $\mathfrak{P}$ given by

$$(\mathbb{P} \otimes_\theta \mathbb{Q})(A) = \iint \mathbf{1}_A(\omega \otimes_\theta \omega')\mathbb{Q}_\omega(\mathrm{d}\omega')\mathbb{P}(\mathrm{d}\omega), \quad A \in \mathcal{B}(\mathbb{D}).$$

LEMMA 3.10. *Let $(t, \mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X}$, $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$ and let θ be a (t, ∞)-valued $\mathbb{F}^0$-stopping time. Then, (i) there exists a regular conditional probability $(\mathbb{P}_\omega)_{\omega\in\mathbb{D}}$ of $\mathbb{P}$ given $\mathcal{F}_\theta^0$ with $\mathbb{P}_\omega \in \mathfrak{P}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$ for $\mathbb{P}$-a.a. $\omega \in \mathbb{D}$; and (ii) given $(\mathbb{Q}_\omega)_{\omega\in\mathbb{D}}$ such that $\omega \mapsto \mathbb{Q}_\omega$ is $\mathcal{F}_\theta^{0,U}$-measurable and $\mathbb{Q}_\omega \in \mathfrak{P}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$ for $\mathbb{P}$-a.a. $\omega \in \mathbb{D}$, it holds that $\mathbb{P} \otimes_\theta \mathbb{Q} \in \mathfrak{P}_{t,\mathbf{x}}$.*

PROOF. (i) We note that $\mathcal{F}_\theta^0$ is generated by the map $\omega \mapsto (\omega_{\cdot\wedge\theta(\omega)}, \theta(\omega))$ where $\mathbb{D} \times \mathbb{R}_+$ is equipped with its Borel algebra; see, for example, page 10 in [19]. Since $\mathcal{B}(\mathbb{D} \times \mathbb{R}_+)$ is countably generated, hence so is $\mathcal{F}_\theta^0$. In turn, since $(\mathbb{D}, \mathcal{B}(\mathbb{D}))$ is a standard Borel space and $\mathcal{F}_\theta^0$ is countably generated, there exists a (proper) regular conditional probability $(\mathbb{P}_\omega)_{\omega\in\mathbb{D}}$ of $\mathbb{P}$ given $\mathcal{F}_\theta^0$, such that $\mathbb{P}_\omega(X_s = \omega_s, s \in [0, \theta(\omega)]) = 1$ for $\mathbb{P}$-a.a. $\omega \in \mathbb{D}$. According to Theorem 1.2.10 in [35], given a martingale M on $[t, \infty)$, there exists a $\mathbb{P}$-null set $N \in \mathcal{F}_\theta^0$ such that M is a $\mathbb{P}_\omega$-martingale after $\theta(\omega)$ for every $\omega \notin N$. Since the fact that W is a BM and ξ an MVM are characterised by the martingale property of $M^\psi_\cdot = \psi(W_{t\vee\cdot}) - \psi(W_t) - \frac{1}{2}\int_t^{t\vee\cdot} \psi''(W_u)\,\mathrm{d}u$ and $\xi_{t\vee\cdot}(\varphi)$, where ψ and φ run through some countable sets of functions (cf. the proof of Lemma 3.9), we conclude that there exists a $\mathbb{P}$-null set $N \in \mathcal{F}_\theta^0$ such that for every $\omega \notin N$, W is a BM and ξ an MVM under $\mathbb{P}_\omega$ on $[\theta(\omega), \infty)$. The remaining properties which must be checked to ensure that $\mathbb{P}_\omega \in \mathfrak{P}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$, for $\mathbb{P}$-almost every $\omega \in \mathbb{D}$, follow from the fact that any $\mathbb{P}$-null set is a $\mathbb{P}_\omega$-null set for $\mathbb{P}$-almost all $\omega \in \Omega$.

(ii) Given $(\mathbb{Q}_\omega)_{\omega\in\mathbb{D}}$ as specified in the statement of the lemma, applying once again Theorem 1.2.10 in [35] (cf. also the proof of Lemma 3.3 in [19]), we obtain that $M^\psi_{t\vee\cdot}$ and $\xi_{t\vee\cdot}(\varphi)$ defined as above are indeed martingales also under $\mathbb{P} \otimes_\theta \mathbb{Q}$. Moreover, since a set which is a null set under $\mathbb{Q}_\omega$, for $\mathbb{P}$-almost all $\omega \in \mathbb{D}$, is a null set also under $\mathbb{P} \otimes_\theta \mathbb{Q}$, we conclude that also the remaining properties hold under the latter measure; hence, $\mathbb{P} \otimes_\theta \mathbb{Q} \in \mathfrak{P}_{t,\mathbf{x}}$. □

Proof of the dynamic programming principle. We are now ready to complete the proof of our main result.

PROOF OF THEOREM 3.8. Given $(t, \mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X}$ and an $\mathbb{F}^0$-stopping time θ with values in (t, ∞), let $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$ and recall that according to Lemma 3.10, there exists a regular conditional probability of $\mathbb{P}$ given $\mathcal{F}^0_\theta$, say $(\mathbb{P}_\omega)_{\omega\in\mathbb{D}}$, such that $\mathbb{P}_\omega \in \mathfrak{P}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$ for $\mathbb{P}$-almost every $\omega \in \mathbb{D}$. In particular, $\mathbb{E}^{\mathbb{P}_\omega}[Y_\infty] \le v(\theta(\omega), \mathbf{x}_{\theta(\omega)}(\omega))$, for $\mathbb{P}$-almost every $\omega \in \mathbb{D}$, and we obtain

$$\mathbb{E}^{\mathbb{P}}[Y_\infty] = \mathbb{E}^{\mathbb{P}}\big[\mathbb{E}^{\mathbb{P}_\omega}[Y_\infty]\big] \le \mathbb{E}^{\mathbb{P}}\big[v(\theta, X_{\cdot\wedge\theta})\big].$$

Since $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$ was arbitrarily chosen, this gives us $v(t, \mathbf{x}) \le \sup_{\mathbb{P}\in\mathfrak{P}_{t,\mathbf{x}}} \mathbb{E}^{\mathbb{P}}[v(\theta, X_{\cdot\wedge\theta})]$.

Next, recall that according to Lemma 3.9, the set $\Gamma = \{(t, \mathbf{x}, \mathbb{P}) : \mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}\}$ is Borel. In particular, this implies that v is upper semianalytic. Indeed, the level sets of v are given by $\{(t, \mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X} : v(t, \mathbf{x}) > c\} = \mathrm{proj}_{\mathbb{R}_+\times\mathbf{X}} \mathrm{L}^{>c}$ with

$$\mathrm{L}^{>c} := \big\{(t, \mathbf{x}, \mathbb{P}) \in \mathbb{R}_+ \times \mathbf{X} \times \mathfrak{P} : \mathbb{E}^{\mathbb{P}}[Y_\infty] > c\big\} \cap \Gamma, \quad c \in \mathbb{R},$$

where we recall that $\mathbb{P} \mapsto \mathbb{E}^{\mathbb{P}}[Y_\infty]$ is Borel measurable. Given $\varepsilon > 0$, it follows that also the following set is analytic:

$$\Gamma_\varepsilon = \big\{(t, \mathbf{x}, \mathbb{P}) : \mathbb{E}^{\mathbb{P}}[Y_\infty] \ge v^\varepsilon(t, \mathbf{x})\big\} \cap \Gamma,$$

where $v^\varepsilon(t, \mathbf{x}) := (v(t, \mathbf{x}) - \varepsilon)\mathbf{1}_{\{v(t,\mathbf{x})<\infty\}} + \frac{1}{\varepsilon}\mathbf{1}_{\{v(t,\mathbf{x})=\infty\}}$. Hence, we may apply Jankov-von Neumann's measurable selection theorem to obtain $(\mathbb{Q}_{t,\mathbf{x}})$ such that $(t, \mathbf{x}) \mapsto \mathbb{Q}_{t,\mathbf{x}}$ is $(\mathcal{B} \otimes \mathcal{X})^U$-meas. and such that $\mathbb{Q}_{t,\mathbf{x}} \in \mathfrak{P}_{t,\mathbf{x}}$ and $\mathbb{E}^{\mathbb{Q}_{t,\mathbf{x}}}[Y_\infty] \ge v^\varepsilon(t, \mathbf{x})$, for every $(t, \mathbf{x}) \in \mathbb{R}_+ \times \mathbf{X}$. In turn, given an $\mathbb{F}^0$-stopping time θ with values in (t, ∞), we define $\mathbb{Q}_\omega := \mathbb{Q}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$ for $\omega \in \mathbb{D}$; then,

$$\mathbb{Q}_\omega \in \mathfrak{P}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)} \text{ and } \mathbb{E}^{\mathbb{Q}_\omega}[Y_\infty] \ge v^\varepsilon\big(\theta(\omega), \mathbf{x}_{\theta(\omega)}(\omega)\big), \quad \omega \in \mathbb{D}.$$

We claim that $(\mathbb{Q}_\omega)_{\omega\in\mathbb{D}}$ is also universally measurable. To see this, given $\mathbb{P} \in \mathfrak{P}$, let $\hat{\mathbb{P}}$ be a measure on $(\mathbb{R}_+ \times \mathbf{X}, \mathcal{B} \otimes \mathcal{X})$ induced by $\mathbb{P}$ via the mapping $\omega \mapsto (\theta(\omega), \mathbf{x}_{\theta(\omega)}(\omega))$; that is,

$$\hat{\mathbb{P}}\big((t, \mathbf{x}) \in A\big) := \mathbb{P}\big(\big(\theta, (W_{\cdot\wedge\theta}, Y_\theta, \xi_\theta)\big) \in A\big), \quad A \in \mathbb{R}_+ \times \mathbf{X}.$$

Let $\mathbb{Q}'_{t,\mathbf{x}}$ be such that $\mathbb{Q}'_{t,\mathbf{x}} = \mathbb{Q}_{t,\mathbf{x}}$, $\hat{\mathbb{P}}$-a.s., and $(t, \mathbf{x}) \mapsto \mathbb{Q}'_{t,\mathbf{x}}$ is $\mathcal{B} \otimes \mathcal{X}$-meas. Then $(t, \omega) \mapsto \mathbb{Q}'_{t,\mathbf{x}_t(\omega)}$ is $\mathcal{B} \times \mathcal{F}^0_\infty$-meas. and thus $\mathcal{O}$-meas. since $\mathbb{Q}'_{t,\mathbf{x}_t(\omega)}$ depends only on $(t, \omega_{\cdot\wedge t})$. In consequence, $\omega \mapsto \mathbb{Q}'_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$ is $\mathcal{F}^0_\theta$-meas. It follows that $\mathbb{Q}_{\theta(\omega),\mathbf{x}_{\theta(\omega)}(\omega)}$ is $\mathcal{F}^{0,\mathbb{P}}_\theta$-meas. and therefore $\mathcal{F}^{0,U}_\theta$-meas. since $\mathbb{P}$ was arbitrarily chosen.

We may now easily conclude: According to Lemma 3.10, for any $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$ we have $\mathbb{P} \otimes_\theta \mathbb{Q} \in \mathfrak{P}_{t,\mathbf{x}}$. In consequence,

$$v(t, \mathbf{x}) \ge \mathbb{E}^{\mathbb{P}\otimes_\theta\mathbb{Q}}[Y_\infty] = \mathbb{E}^{\mathbb{P}}\big[\mathbb{E}^{\mathbb{Q}_\omega}\big] \ge \mathbb{E}^{\mathbb{P}}\big[v^\varepsilon(\theta, X_{\cdot\wedge\theta})\big].$$

Since $\varepsilon > 0$ and $\mathbb{P} \in \mathfrak{P}_{t,\mathbf{x}}$ were both arbitrarily chosen, we therefore have $v(t, \mathbf{x}) \ge \sup_{\mathbb{P}\in\mathfrak{P}_{t,\mathbf{x}}} \mathbb{E}^{\mathbb{P}}[v(\theta, X_{\cdot\wedge\theta})]$ which completes the proof. □

4. Stability of the value function and covering arguments. In this section we study continuity properties of the optimal stopping problem with respect to its distributional constraint. Moreover, for the case of Markovian costs, we provide an alternative proof of the dynamic programming principle.

Throughout this section we work under the following assumption on the cost function.

ASSUMPTION 4.1. There exists a modulus of continuity φ such that for any two ordered stopping times, say $\tau \le \rho$ a.s., in a filtered probability space supporting a Brownian motion $(W_t)_{t\ge 0}$, it holds that

$$\big|\mathbb{E}\big[c(W_{\cdot\wedge\tau}, \tau) - c(W_{\cdot\wedge\rho}, \rho)\big]\big| \le \varphi\big(\mathbb{E}\big[|\tau - \rho|\big]\big).$$

While this assumption defines a particular class of cost functions we note that it is wide enough to include a number of relevant examples.

EXAMPLE 4.2. Assumption 4.1 holds in each of the following cases:

- $c(\omega, t)$ satisfying $|c(\omega, t) - c(\omega, s)| \le \varphi(|t - s|)$ for a concave modulus of continuity φ;
- $c(\omega, t) = f(\omega_t)$, where $|f(x) - f(y)| \le \varphi(|x - y|)$ for a concave modulus of continuity φ, since then $\widetilde{\varphi}(\cdot) = \varphi(\sqrt{\cdot})$ is a concave modulus of continuity and by Jensen's inequality

$$\mathbb{E}\big[\big|f(W_\tau) - f(W_\rho)\big|\big] \le \mathbb{E}\big[\widetilde{\varphi}\big(|W_\tau - W_\rho|^2\big)\big] \le \widetilde{\varphi}\big(\mathbb{E}\big[|W_\tau - W_\rho|^2\big]\big) = \widetilde{\varphi}\big(\mathbb{E}\big[|\tau - \rho|\big]\big);$$

- $c(\omega, t) = f(\omega_t^*)$, where $\omega_t^* = \sup_{\cdot\le t} \omega_\cdot$ and $|f(x) - f(y)| \le \varphi(|x - y|)$ for a concave modulus of continuity φ, since defining $\widetilde{\varphi}$ as above and applying Doob's inequality,

$$\mathbb{E}\big[\big|f(W_\tau^*) - f(W_\rho^*)\big|\big] \le \widetilde{\varphi}\big(\mathbb{E}\big[\big|W_\tau^* - W_\rho^*\big|^2\big]\big) \le \widetilde{\varphi}\big(4\mathbb{E}\big[|\tau - \rho|\big]\big);$$

- variations of the above cases obtained when replacing W by a local martingale M "evolving slower" than the Brownian motion in the sense that $\langle M\rangle_t \le t$.

4.1. *Stability properties of the value function.* We first study the dependence of the problem on its distributional constraint. Specifically, recalling our original problem formulation, we consider continuity properties of the mapping $\mu \mapsto v(\mu)$, $\mu \in \mathcal{P}$. We note that this value function is continuous, for example, when the problem admits a solution of barrier type; cf. [8, 23, 28, 33]. More generally, such continuity properties are easier to establish when taking the ordering on the underlying space of probability distributions into account; we therefore choose to view $\mathcal{P}$ as a partially ordered space while still making use of its properties as a metric space under $\mathcal{W}_1$. More precisely, we write $\mu \preccurlyeq \eta$ if η can be obtained from μ by "moving mass to the right"; that is, if there exists a coupling between μ and η whose disintegration in the first variable, say $m(\cdot, \mathrm{d}y)$, is such that for μ-a.a. x, $m(x, \mathrm{d}y)$ is concentrated on $[x, \infty)$. This ordering is usually referred to as the first stochastic ordering and the fact that η stochastically dominates μ can equivalently be expressed as $\mu(f) \le \eta(f)$ for any nondecreasing function f.[2]

Our first result establishes a certain "right-continuity" of the value function. The result is of relevance for it ensures that the value of an optimal stopping problem equipped with a general distributional constraint can be approximated by problems featuring atomic constraints which are easier to handle by use of numerical methods (cf. Corollary 3.6).

THEOREM 4.3. *Suppose that the cost function c is bounded with $t \mapsto c(\omega_{\cdot\wedge t}, t)$ continuous and that it satisfies Assumption 4.1. Let $\mu \in \mathcal{P}$ and take $\mu_n \in \mathcal{P}$, $n \in \mathbb{N}$, to be a sequence with $\mu \preccurlyeq \mu_n$, $n \in \mathbb{N}$, such that $\mathcal{W}_1(\mu, \mu_n) \to 0$, as $n \to \infty$. Then, $v(\mu) = \lim_{n\to\infty} v(\mu_n)$.*

[2]We note that $\mathcal{P}$ equipped with the usual stochastic ordering is a lattice (a poset such that every nonempty finite set has a supremum and infimum); if restricting to measures supported on some given compact set, it is also a complete lattice (any subset admits an infimum). It is also a continuous lattice in the sense of [22].

PROOF. For each $n \in \mathbb{N}$, recalling that $\mu \preccurlyeq \mu_n$, let Γ^n such that $\mathcal{W}_1(\mu, \mu_n) = \iint |s - t|\Gamma^n(\mathrm{d}s, \mathrm{d}t)$ and such that its disintegration kernel $m^n(s, \mathrm{d}t)$ satisfies both $\Gamma^n(\mathrm{d}s, \mathrm{d}t) = \mu(\mathrm{d}s)m^n(s, \mathrm{d}t)$ and $\operatorname{supp} m^n(s, \cdot) \subseteq [s, \infty)$, for μ-a.e. $s > 0$. Let $\varepsilon > 0$. Recalling in turn the probability space $(\Omega, \mathcal{G}, \mathbb{G}, \mathbb{P})$ from problem (2.1), take $\tau \in \mathcal{T}(\mu)$ such that $\mathbb{E}[c(W_{\cdot\wedge\tau}, \tau)] \geq v(\mu) - \varepsilon$. Consider now the enlarged probability space given by

$$(\overline{\Omega}, \overline{\mathcal{G}}, (\overline{\mathcal{G}}_t)_{t\geq 0}, \overline{\mathbb{P}}) = (\Omega \times [0, 1], \mathcal{G} \otimes \mathcal{B}([0, 1]), (\mathcal{G}_t \otimes \mathcal{B}([0, 1]))_{t\geq 0}, \mathbb{P} \times \mathrm{Leb});$$

we write $\bar{\omega} = (\omega, u)$ for $\bar{\omega} \in \overline{\Omega}$. Defining $U(\bar{\omega}) = u$, it follows that the random variable U is uniformly distributed on $[0, 1]$ and independent of $\mathcal{G}$. Denoting by $M^{n,(-1)}(s, \cdot) : [0, 1) \to [s, \infty)$ the right-continuous inverse of $M^n(s, \cdot) := \int_s^{\cdot} m^n(s, \mathrm{d}t)$, for $n \in \mathbb{N}$, and using the obvious identification $\tau(\bar{\omega}) = \tau(\omega)$, we then define

$$\tau_n := M^{n,(-1)}(\tau, U), \quad n \in \mathbb{N};$$

it follows that τ_n is a stopping time in the enlarged space. Moreover, $\tau_n \sim \mu_n$ and

$$\overline{\mathbb{E}}[|\tau - \tau_n|] = \overline{\mathbb{E}}[|\tau - M^{n,(-1)}(\tau, U)|] = \iint |s - t|\mu(\mathrm{d}s)m^n(s, \mathrm{d}t) = \mathcal{W}_1(\mu, \mu_n).$$

We can therefore choose $N \in \mathbb{N}$ such that $\overline{\mathbb{E}}[|\tau - \tau_n|] \leq \varphi^{-1}(\varepsilon)$, for all $n \geq N$, where φ is the modulus of continuity provided by Assumption 4.1; making use of this assumption we then obtain

$$v(\mu_n) \geq \overline{\mathbb{E}}[c(W_{\cdot\wedge\tau_n}, \tau_n)] \geq \overline{\mathbb{E}}[c(W_{\cdot\wedge\tau}, \tau)] - \varepsilon \geq v(\mu) - 2\varepsilon \quad n \geq N.$$

Since $\varepsilon > 0$ was arbitrarily chosen, we therefore have $\liminf_{n\to\infty} v(\mu_n) \geq v(\mu)$, which combined with Proposition 4.4 below yields the result. □

The next result establishes (general) upper semicontinuity of the value function; the argument is a variation of the existence proof given in [8] and notably does not rely on Assumption 4.1. Since the value function v is concave,[3] the result implies in particular that when restricting v to the set of finitely supported probability measures—rendering it a function from $\mathbb{R}^n$ to $\mathbb{R}$ (cf. (3.6))—continuity holds on the domain where it is finite.

PROPOSITION 4.4. *Suppose that the cost function c is bounded from above and that $t \mapsto c(\omega_{\cdot\wedge t}, t)$ is upper semicontinuous for $\mathbb{W}$-a.e. $\omega \in C_0(\mathbb{R}_+)$. Then, $\mu \mapsto v(\mu)$ is upper semicontinuous on $\mathcal{P}$ in the topology induced by $\mathcal{W}_1$.*

PROOF. Fix $\mu \in \mathcal{P}$ and let $\mu_n \in \mathcal{P}$, $n \in \mathbb{N}$, be a sequence such that $\mathcal{W}_1(\mu_n, \mu) \to 0$, as $n \to \infty$; without loss of generality, we suppose that $\mathcal{W}_1(\mu_n, \mu)$ is nonincreasing in n. In turn, recalling the equivalent problem formulation from (2.5), let $\gamma_n \in \mathsf{RST}(\mu_n)$, $n \in \mathbb{N}$, be a sequence such that

$$\lim_{n\to\infty} \mathbb{E}^{\gamma_n}[c(B, T)] = \limsup_{n\to\infty} v(\mu_n). \tag{4.1}$$

We first argue that the sequence (γ_n) is tight; for this it suffices to show that its respective projections onto $C_0(\mathbb{R}_+)$ and $\mathbb{R}_+$ are tight. The projections onto $C_0(\mathbb{R}_+)$ all coincide with

[3] Indeed, let $\mu_1, \mu_2 \in \mathcal{P}$, $\varepsilon > 0$, and let τ_1, τ_2 be corresponding ε-optimal stopping times. Given $\lambda \in (0, 1)$, let $\mu^\lambda = \lambda\mu_1 + (1 - \lambda)\mu_2$. Consider the enlarged space $(\overline{\Omega}, \overline{\mathcal{G}}, \overline{\mathbb{G}}, \overline{\mathbb{P}})$ constructed in the proof of Theorem 4.3 and define $\tau = \mathbf{1}_{\{U\leq\lambda\}}\tau_1 + \mathbf{1}_{\{U>\lambda\}}\tau_2$. Then, $\tau \in \mathcal{T}(\mu^\lambda)$; moreover,

$$v(\mu^\lambda) \geq \overline{\mathbb{E}}[\mathbf{1}_{\{U\leq\lambda\}}c(W_{\cdot\wedge\tau_1}, \tau_1) + \mathbf{1}_{\{U>\lambda\}}c(W_{\cdot\wedge\tau_2}, \tau_2)] \geq \lambda v(\mu_1) + (1 - \lambda)v(\mu_2) - 2\varepsilon.$$

$\mathbb{W}$ and are thus trivially tight. On the other hand, for any $\varepsilon > 0$, using the properties of the sequence (μ_n), we may choose $r > 0$ such that $\mu_n((r, \infty)) < \varepsilon$, for all $n \in \mathbb{N}$. Since $\gamma_n(T > r) = \mu_n((r, \infty))$, it follows that also the projections onto $\mathbb{R}_+$ are tight.

Let now γ be an accumulation point; by passing if necessary to a subsequence, we may assume that γ_n converges weakly to γ. From the continuity of $(\omega, t) \mapsto \omega$, we obtain that $\gamma|_{C_0(\mathbb{R}_+)} = \mathbb{W}$. Further, using the continuity of $(\omega, t) \mapsto t$ combined with the fact that $\mathcal{W}_1(\mu_n, \mu) \to 0$, we obtain for any $\varphi \in C_b$,

$$\gamma|_{\mathbb{R}_+}(\varphi) = \mathbb{E}^{\gamma}[\varphi(T)] = \lim_{n\to\infty} \mathbb{E}^{\gamma_n}[\varphi(T)] = \lim_{n\to\infty} \mu_n(\varphi) = \mu(\varphi).$$

Hence, $\gamma \in \mathsf{Cpl}(\mathbb{W}, \mu)$. We may now invoke Theorem 3.8 in [6] (cf. the equivalence between properties (2) and (3) therein), to deduce that $\gamma \in \mathsf{RST}(\mu)$.

Next, according to Proposition 2.4 in [26] (see also Lemma 4.2 in [24]), on $\bigcup_{\mu\in\mathcal{P}} \mathsf{RST}(\mu)$, the weak convergence topology coincides with the stable convergence topology, which is the coarsest topology under which $\gamma \to \mathbb{E}^{\gamma}[\phi]$ is continuous for all bounded measurable functions $\phi : C_0(\mathbb{R}_+) \times \mathbb{R}_+ \to \mathbb{R}$ such that $t \mapsto \phi(\omega, t)$ is continuous for all $\omega \in C_0(\mathbb{R}_+)$. Hence, by use of the imposed assumptions on c and Portmanteau's lemma, we obtain

$$v(\mu) \geq \mathbb{E}^{\gamma}[c(B, T)] \geq \limsup_{n\to\infty} \mathbb{E}^{\gamma_n}[c(B, T)],$$

which combined with (4.1) yields the required upper semicontinuity. □

4.2. *The DPP in the Markovian case.* In this section we provide an alternative proof of the DPP for the case of Markovian cost functions. While the result naturally follows from our previous one we provide this separate argument due to its simplicity. The argument relies on the fact that under some additional structural assumptions, we can construct a measurable optimiser explicitly via covering arguments; we may thus circumvent the need for abstract measurable selection arguments and provide a straightforward proof. The approach is in the spirit of Bouchard and Touzi [11] (see also [10] and [1]); it was also the method employed in [5] for the optimal stopping problem with atomic constraints.

The Markovian setting. Recalling the setup from our original problem formulation (2.1) and denoting the augmented filtration generated by the Brownian motion alone by $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$, we say that a cost function c is of *Markov type* if it admits the representation

$$c(W_{\cdot\wedge\tau}, \tau) = f(X_\tau) \quad \text{a.s.},$$

for some $\mathbb{R}^n$-valued $\mathbb{F}$-Markov process $(X_t)_{t\geq 0}$ and $f \in C(\mathbb{R}^n)$. Let $S = \{(t, x, \mu) \in \mathbb{R}_+ \times \mathbb{R}^n \times \mathcal{P} : \mu([t, \infty)) = 1\}$ and $S_0 = \{(t, x, \mu) \in S : \mu((t, \infty)) = 1\}$; with a slight abuse of notation we write S and S_0 also for their projections onto $\mathbb{R}_+ \times \mathcal{P}$. For the class of cost functions of Markov type, we then re-introduce the value function $v : S_0 \to \mathbb{R}$ by

$$v(t, x, \mu) := \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\left[\int_t^\infty f(X_s^{t,x})\, \mathrm{d}A_s^{\xi}\right],$$

where $\mathsf{MVM}^t(\mu)$ here denotes the set of continuous adapted MVMs with $\xi_t = \mu$ a.s., which are independent of $\mathcal{F}_t$, and where we now refer to an MVM as continuous if $\xi(\varphi)$ is continuous a.s. for every $\varphi \in b\mathcal{B}$.[4]

For the Markovian case, the DPP then takes the following form; the result can be deduced directly from Theorem 3.1 via the identification $v(t, X_t(\mathbf{w}), \mu) \mathrel{\hat=} v(t, \mathbf{w}, \mu)$ or proven analogously to Theorem 3.1 by applying measurable selection to deduce a measurable optimiser on the state space S_0.

[4]An inspection of the proof of Proposition 2.5 together with Lemma A.1 verifies that this modified continuity assumption will not affect the value of the problem; that the value remains unaltered if relaxing the condition of independence of $\mathcal{F}_t$ follows from (B.3) below and, for example, Proposition 2.4 in [12].

COROLLARY 4.5 (DPP: Markovian case). *Let $(t, x, \mu) \in S_0$. For any $\mathbb{F}$-stopping time θ with values in (t, ∞), it then holds that*

$$v(t, x, \mu) = \sup_{\xi \in \mathsf{MVM}^t(\mu)} \mathbb{E}\bigg[\int_t^\theta f(X_u^{t,x}) \, \mathrm{d}A_u^\xi + \xi_\theta\big((\theta, \infty)\big) v\Big(\theta, X_\theta^{t,x}, \xi_\theta^{(\theta)}\Big)\bigg].$$

Covering arguments and alternative proof. Our direct proof of the above DPP relies on the fact that given an MVM which is optimal at a point, we may in a measurable way construct a modified MVM which respects the state constraint and is approximately optimal in some neighbourhood of that point. To specify this, given $\mu, \eta \in \mathcal{P}$, let $m^{\mu \to \eta}$ denote the disintegration kernel for which

$$\mathcal{W}_1(\mu, \eta) = \iint |s - u| \mu(\mathrm{d}u) m^{\mu \to \eta}(u, \mathrm{d}s).$$

Consider $(\widetilde{t}, \widetilde{\mu}) \in S$, $\underset{\sim}{t} < \widetilde{t}$ and the MVM $\widetilde{\xi} \in \mathsf{MVM}^{\underset{\sim}{t}}(\widetilde{\mu})$ to be given. For every $(s, \eta) \in S_0$ with $s < \underset{\sim}{t}$, we then define a process ξ as follows: $\xi_{s \vee \cdot \wedge \underset{\sim}{t}} := \eta$ and

$$\xi_{\cdot \vee \underset{\sim}{t}}(A) := \eta\big(A \cap (s, \widetilde{t}]\big) + \eta\big((\widetilde{t}, \infty)\big) \int \widetilde{\xi}_\cdot(\mathrm{d}u) m^{\widetilde{\mu} \to \eta^{(\widetilde{t}}}(u, A), \quad A \in \mathcal{B}. \tag{4.2}$$

If $\widetilde{\mu} \preccurlyeq \eta^{(\widetilde{t}}$, the thus defined process ξ is an adapted MVM which lies in $\mathsf{MVM}^s(\eta)$. More pertinently, under some further assumptions, the next result verifies that $\widetilde{\xi}$ and ξ are close in a certain sense and if the former is optimally chosen, the latter will thus be approximately optimal in some neighbourhood.

ASSUMPTION 4.6. The cost function c is of Markov type with f locally bounded and satisfies Assumption 4.1. Moreover, for every $\eta \in \mathcal{P}$, there exists $\varepsilon > 0$ such that $(t, x) \mapsto \mathbb{E}[f(X_\tau^{t \wedge l_\tau, x})]$ is continuous on $\mathbb{R}_+ \times \mathbb{R}^n$ uniformly for $\tau \in \mathcal{T}(\mu)$ with $\mu \in B_\varepsilon(\eta)$, where l_τ denotes the left end-point of the support of τ.

PROPOSITION 4.7. *Suppose that Assumption 4.6 holds and that v is continuous on S_0. Given $\varepsilon > 0$, there exists $\delta : S_0 \to \mathbb{R}_+$ such that for every $(t, x, \mu) \in S_0$ the following holds: If $A \subseteq B_{\delta(t,x,\mu)}(t, x, \mu)$, $(\widetilde{t}, \widetilde{\mu}) \in S$, $\underset{\sim}{t} < \widetilde{t}$, and it holds for all $(s, z, \eta) \in A$ that*

$$(\widetilde{t} - s) \vee \eta\big((0, \widetilde{t}]\big) \vee \mathcal{W}_1\big(\eta^{(\widetilde{t}}, \widetilde{\mu}\big) < \delta(t, x, \mu), \quad s < \underset{\sim}{t} \text{ and } \widetilde{\mu} \preccurlyeq \eta^{(\widetilde{t}}; \tag{4.3}$$

and if $\widetilde{\xi} \in \mathsf{MVM}^{\underset{\sim}{t}}(\widetilde{\mu})$ is chosen such that

$$\mathbb{E}\bigg[\int f\big(X_u^{\underset{\sim}{t},x}\big) \, \mathrm{d}A_u^{\widetilde{\xi}}\bigg] \geq v(\underset{\sim}{t}, x, \widetilde{\mu}) - \delta(t, x, \mu);$$

then it holds for any $(s, z, \eta) \in A$ that the MVM ξ given by (4.2) belongs to $\mathsf{MVM}^s(\eta)$ and satisfies

$$\mathbb{E}\bigg[\int f\big(X_u^{s,z}\big) \, \mathrm{d}A_u^\xi\bigg] \geq v(s, z, \eta) - \varepsilon.$$

The second component needed for the proof is the fact that S_0 admits a measurable partition satisfying the properties required for Proposition 4.7 to apply; the following lemma provides such a partition.

LEMMA 4.8. *For $\delta : S_0 \to \mathbb{R}_+$ given, S_0 admits a measurable partition (A_i), $i \in \mathbb{N}$, with associated sequences of reference points $(t_i, x_i, \mu_i) \in S_0$, $(\widetilde{t}_i, \widetilde{\mu}_i) \in S$ and $\underset{\sim}{t_i} < \widetilde{t}_i$, such that $A_i \subseteq B_{\delta(t_i,x_i,\mu_i)}(t_i, x_i, \mu_i)$ and (4.3) holds with respect to (t_i, x_i, μ_i), $(\widetilde{t}_i, \widetilde{\mu}_i)$ and $\underset{\sim}{t_i}$, for all $(s, z, \eta) \in A_i$, $i \in \mathbb{N}$.*

Combining Lemma 4.8 and Proposition 4.7, it is now straightforward to explicitly construct a measurable and approximately optimal MVM, which enables proving the DPP without invoking any abstract measurable selection arguments. For completeness, we provide the full argument in Appendix B together with the proofs of Proposition 4.7 and Lemma 4.8. We note that this direct proof relies on continuity of v while in light of Theorem 4.3, "right-continuity" would have been a more natural assumption. It seems, however, unclear whether the argument can be extended to account for this, which illustrates the limitations of covering arguments. In conclusion, while covering arguments provide neat proofs whenever applicable, the complexity of the distribution-constrained optimal stopping problem typically calls for more subtle measurable selection arguments as employed in our main proof herein.

APPENDIX A: ON THE CONTINUITY AND CONVERGENCE OF MEASURE-VALUED MARTINGALES

In this appendix we establish some auxiliary results on measure-valued martingales; the underlying filtered probability space is here taken to be a general one satisfying the usual conditions. Moreover, we here apply Definition 2.2 with the convention that $\mathcal{P}$ is replaced by the set of all probability measures on $\mathbb{R}$; that is, we do not a priori impose any assumptions on the integrability or support of the MVMs.

First, we consider the existence of continuous versions of MVMs.

LEMMA A.1. *Given an MVM ξ, there exists a version of it such that $\xi(\varphi)$ is càdlàg a.s. for every $\varphi \in b\mathcal{B}$. Moreover, ξ is then almost surely right-continuous in the weak topology.*

If $\xi_0(|\cdot|) < \infty$, the version can be chosen such that $\xi(\varphi)$ is càdlàg a.s. for every $\varphi \in b\mathcal{B}$ and every $\varphi \in C$ asymptotically of at most linear growth. Moreover, ξ is then almost surely right-continuous in the topology induced by $\mathcal{W}_1$.

PROOF. According to Proposition 2.5 in [25] and its proof, there exists a version of ξ such that for every $\varphi \in b\mathcal{B}$, $\xi(\varphi)$ is càdlàg a.s. Let $\mathcal{A}$ be the set of finite unions of open intervals with rational endpoints. Then there exists a null set $\mathcal{N}$ such that for $\omega \notin \mathcal{N}$, $t \mapsto \xi_t(A)(\omega)$ is càdlàg for all $A \in \mathcal{A}$. Let $U \subseteq \mathbb{R}$ be an arbitrary open set; then $U = \bigcup_{n=1}^{\infty} A_n$ for some sequence (A_n) with $A_n \in \mathcal{A}$ and $A_n \subset A_{n+1}$, $n \in \mathbb{N}$. Since $A_n \subset U$, $n \in \mathbb{N}$, we obtain off $\mathcal{N}$ that for any $t \geq 0$,

$$\liminf_{u \downarrow t} \xi_u(U) \geq \sup_{n \in \mathbb{N}} \liminf_{u \downarrow t} \xi_u(A_n) = \sup_{n \in \mathbb{N}} \xi_t(A_n) = \xi_t(U);$$

by Portmanteau's lemma ξ is thus weakly right-continuous off $\mathcal{N}$.

Next, suppose in addition that $\xi_0(|\cdot|) < \infty$. For any $\varphi \in b\mathcal{B}$ it then holds that

$$\int \varphi(x)\big(1 \vee |x|\big)\xi_0(\mathrm{d}x) \leq \sup_{x \in \mathbb{R}} \varphi(x)\big(1 + \xi_0(|\cdot|)\big) < \infty,$$

and therefore $\xi_t(\varphi(\cdot)(1 \vee |\cdot|))$, $t \geq 0$, is a martingale for any $\varphi \in b\mathcal{B}$. In consequence, defining $\eta_t(\mathrm{d}x) = (1 \vee |x|)\xi_t(\mathrm{d}x)$, it holds that $\eta_t(\varphi)$, $t \geq 0$, is a martingale for any $\varphi \in b\mathcal{B}$. Moreover, $\eta_t(\mathrm{d}x)$ is a nonnegative random measure. Hence, we may again apply the results in [25] to deduce the existence of a version of η, and thus of ξ, such that $\eta(\varphi)$ is càdlàg a.s. for any $\varphi \in b\mathcal{B}$. In particular, for any $\varphi \in b\mathcal{B}$ or $\varphi \in C$ with $\lim_{|x| \to \infty} |\frac{\varphi(x)}{x}| \leq 1$, since it then holds that $\frac{\varphi(\cdot)}{1 \vee |\cdot|} \in b\mathcal{B}$ and

$$\int \varphi(x)\xi_t(\mathrm{d}x) = \int \frac{\varphi(x)}{1 \vee |x|}\eta_t(\mathrm{d}x), \quad t \geq 0,$$

we obtain that $\xi(\varphi)$ is càdlàg a.s. By use of the same arguments as above, it follows that ξ is a.s. weakly right-continuous. Since also $\xi(|\cdot|)$ is càdlàg a.s., it follows that ξ is a.s. right-continuous in the topology induced by $\mathcal{W}_1$. □

Next, we turn to the question of convergence of MVMs to a limiting object.

LEMMA A.2. *Given an MVM ξ, there exists a random measure ξ_∞ which satisfies $\lim_{t\to\infty}\xi_t(\varphi)=\xi_\infty(\varphi)$ a.s. for every $\varphi\in b\mathcal{B}$. Moreover, as $t\to\infty$, ξ_t then converges almost surely to ξ_∞ in the weak topology.*

If $\xi_0(|\cdot|)<\infty$, there exists a random measure ξ_∞ which satisfies $\lim_{t\to\infty}\xi_t(\varphi)=\xi_\infty(\varphi)$ a.s. for every $\varphi\in b\mathcal{B}$ and every $\varphi\in C$ asymptotically of at most linear growth. Moreover, as $t\to\infty$, ξ_t then converges almost surely to ξ_∞ in the topology induced by $\mathcal{W}_1$.

PROOF. According to Proposition 2.1 in [25] and its proof, there exists a random measure ξ_∞ such that $\xi_t(\varphi)\to\xi_\infty(\varphi)$ a.s., as $t\to\infty$, for any $\varphi\in b\mathcal{B}$. Let $\mathcal{A}$, U and $(A_n)_{n\in\mathbb{N}}$ be as in the proof of Lemma A.1; further, let $\mathcal{N}$ be a null set off which $\xi_t(A_n)(\omega)\to\xi_\infty(A_n)(\omega)$, for all $n\in\mathbb{N}$. Off $\mathcal{N}$ it then holds that

$$\liminf_{t\to\infty}\xi_t(U)\geq\lim_{n\to\infty}\liminf_{t\to\infty}\xi_t(A_n)=\lim_{n\to\infty}\xi_\infty(A_n)=\xi_\infty(U),$$

and the almost sure weak convergence follows from Portmanteau's lemma.

Next, suppose in addition that $\xi_0(|\cdot|)<\infty$. Define the measure-valued process η as in the proof of Lemma A.1. Applying again the results in [25], we deduce the existence of a random measure η_∞ such that $\eta_t(\varphi)\to\eta_\infty(\varphi)$ a.s., for any $\varphi\in b\mathcal{B}$. Defining $\xi_\infty(\mathrm{d}x)=\frac{1}{1\vee|x|}\eta_\infty(\mathrm{d}x)$, we obtain for any $\varphi\in b\mathcal{B}$ or $\varphi\in C$ with $\lim_{|x|\to\infty}|\frac{\varphi(x)}{x}|\leq 1$, that

$$\lim_{t\to\infty}\xi_t(\varphi)=\lim_{t\to\infty}\eta_t\left(\frac{\varphi(\cdot)}{1\vee|\cdot|}\right)=\eta_\infty\left(\frac{\varphi(\cdot)}{1\vee|\cdot|}\right)=\xi_\infty(\varphi),\quad \text{a.s.}$$

Convergence in the weak topology then follows by use of the same arguments as above. Convergence in the topology induced by $\mathcal{W}_1$ then holds off the union of the two null sets off which, respectively, the first moment converges and weak convergence holds. □

The next result refines the convergence properties for adapted MVMs.

LEMMA A.3. *Consider an adapted MVM such that $\xi(\varphi)$ is càdlàg a.s. for every $\varphi\in b\mathcal{B}$ and let ξ_∞ be a random measure such that $\lim_{t\to\infty}\xi_t(\varphi)=\xi_\infty(\varphi)$ a.s. for every $\varphi\in b\mathcal{B}$. Then, for any $s>0$, there exists a null set $\mathcal{N}^s$ such that*

$$\xi_t(B)(\omega)=\xi_s(B)(\omega)\quad \textit{for all } t\in[s,\infty] \textit{ and } B\in\mathcal{B}([0,s]),\omega\notin\mathcal{N}^s.$$

In particular, there exists a null set $\mathcal{N}$ such that for any $B\in\mathcal{B}$ bounded by some $\kappa>0$, $\xi_t(B)(\omega)=\xi_\infty(B)(\omega)$ for all $t>\kappa$, $\omega\notin\mathcal{N}$.

PROOF. For any fixed $t\geq 0$, $\xi_u([0,t])=\xi_t([0,t])$, a.s., $u\geq t$. Since $\xi_\cdot([0,t])$ is càdlàg a.s. there thus exists some null set, possibly depending on t, off which $\xi_u([0,t])(\omega)=\xi_t([0,t])(\omega)$, $u\geq t$. In consequence, also $\xi_\infty([0,t])(\omega)=\xi_t([0,t])(\omega)$ off this null set. For fixed $s>0$, it follows that there exists some null set $\mathcal{N}^s$, off which

$$\xi_t([0,q])=\xi_q([0,q]),\quad t\in[q,\infty],\ \text{for all } q\in[0,s]\cap\mathbb{Q}\cup\{s\}.$$

Note that $\mathcal{D} := \{B \in \mathcal{B}([0,s]) : \xi_t(B)(\omega) = \xi_s(B)(\omega), t \geq s, \omega \notin \mathcal{N}^s\}$ is a λ-system. Indeed, clearly $[0,s] \in \mathcal{D}$ and $A \in \mathcal{D} \implies [0,s] \setminus A \in \mathcal{D}$. Further, given $(A_i)_{i\in\mathbb{N}}$ with $A_i \in \mathcal{D}$ and $A_i \cap A_j = \varnothing$, $i \neq j$, it holds off $\mathcal{N}^s$ that

$$\xi_t\Big(\bigcup_{i\in\mathbb{N}} A_i\Big) = \sum_{i\in\mathbb{N}} \xi_t(A_i) = \sum_{i\in\mathbb{N}} \xi_s(A_i) = \xi_s\Big(\bigcup_{i\in\mathbb{N}} A_i\Big), \quad t \geq s.$$

Since $\{[0,q] : q \in [0,s] \cap \mathbb{Q} \cup \{s\}\}$ is a π-system contained in $\mathcal{D}$ this completes the proof. $\square$

APPENDIX B: PROOF OF THE DPP USING COVERING ARGUMENTS

We here provide the full argument for how to prove Theorem 3.1 using covering arguments. To this end we first provide the two missing proofs from Section 4.2.

PROOF OF PROPOSITION 4.7. Let $(t,x,\mu) \in S_0$ and $\delta > 0$. Further, let $A \subset B_\delta(t,x,\mu)$, $(\widetilde{t},\widetilde{\mu}) \in S$, $\underset{\sim}{t} < \widetilde{t}$ and $\widetilde{\xi} \in \mathsf{MVM}^{\underset{\sim}{t}}(\widetilde{\mu})$ be such that the conditions listed in the proposition are satisfied with respect to this fixed δ. In turn, let $(s,z,\eta) \in A$ and define $\xi \in \mathsf{MVM}^s(\eta)$ by (4.2). By use of Assumption 4.1 we first obtain

$$\begin{aligned}
&\bigg|\mathbb{E}\Big[\iint f(X_r^{s,z}) m^{\widetilde{\mu}\to\eta^{(\widetilde{t}}}(u,\mathrm{d}r)\widetilde{\xi}_\infty(\mathrm{d}u) - \int f(X_u^{s,z})\widetilde{\xi}_\infty(\mathrm{d}u)\Big]\bigg| \\
&\quad \leq \varphi\Big(\mathbb{E}\Big[\iint |u-s| m^{\widetilde{\mu}\to\eta^{(\widetilde{t}}}(s,\mathrm{d}u)\widetilde{\xi}_\infty(\mathrm{d}s)\Big]\Big) \\
&\quad \leq \varphi\Big(\iint |u-s| m^{\widetilde{\mu}\to\eta^{(\widetilde{t}}}(s,\mathrm{d}u)\widetilde{\mu}(\mathrm{d}s)\Big) \leq \varphi\big(\mathcal{W}_1(\widetilde{\mu},\eta^{(\widetilde{t}})\big),
\end{aligned}$$

for some modulus of continuity φ (cf. the proof of Theorem 4.3); that is,

$$\mathbb{E}\Big[\int_{\widetilde{t}}^\infty f(X_u^{s,z})\xi_\infty(\mathrm{d}u)\Big] \geq \mathbb{E}\Big[\int f(X_u^{s,z})\widetilde{\xi}_\infty(\mathrm{d}u)\Big] - \varphi\big(\mathcal{W}_1(\widetilde{\mu},\eta^{(\widetilde{t})})\big).$$

Second, using again Assumption 4.1, we also have

$$\bigg|\mathbb{E}\Big[\int_s^{\widetilde{t}} f(X_u^{s,z})\frac{\eta(\mathrm{d}u)}{\eta((s,\widetilde{t}])} - f(x)\Big]\bigg| \leq \widetilde{\varphi}\Big(\int_s^{\widetilde{t}} |u-s|\eta(\mathrm{d}u)\Big) \leq \widetilde{\varphi}\big(\eta((s,\widetilde{t}])|s-\widetilde{t}|\big),$$

for some modulus of continuity $\widetilde{\varphi}$; that is,

$$\mathbb{E}\Big[\int_s^{\widetilde{t}} f(X_u^{s,z})\xi_\infty(\mathrm{d}u)\Big] \geq f(x)\eta((s,\widetilde{t}]) - \widetilde{\varphi}\big(\eta((s,\widetilde{t}])|s-\widetilde{t}|\big).$$

Applying Assumption 4.6 we see that by requiring δ to be small enough, we may thus ensure that $\mathbb{E}[\int f(X_u^{s,z})\xi_\infty(\mathrm{d}u)] \geq \mathbb{E}[\int f(X_u^{s,z})\widetilde{\xi}_\infty(\mathrm{d}u)] - \varepsilon$, for all $(s,z,\eta) \in A$.

On the other hand, applying again Assumption 4.6 and using that v is continuous on S_0, requiring if necessary $\delta > 0$ to be even smaller, we also have for all $(s,z,\eta) \in A \subset B_\delta(t,x,\mu)$,

$$\bigg|\mathbb{E}\Big[\int f(X_u^{s,z})\widetilde{\xi}_\infty(\mathrm{d}u)\Big] - \mathbb{E}\Big[\int f(X_u^{\underset{\sim}{t},x})\widetilde{\xi}_\infty(\mathrm{d}u)\Big]\bigg| \leq \varepsilon$$

and $|v(\underset{\sim}{t},x,\widetilde{\mu}) - v(s,z,\eta)| \leq \varepsilon$, which allows us to conclude. $\square$

PROOF OF LEMMA 4.8. Let $(t,x,\mu) \in S_0$ and $0 < \delta < 1$. Let $t_0 \in (t, t+\delta)$ such that $\mu((0,t_0]) < \delta$. In turn, let $t < \underset{\sim}{t} < \widetilde{t} < t_0$ and define the set

$$A^{(t,x,\mu)} := (t-\delta, \underset{\sim}{t}) \times B_\delta(x) \times B_{\delta(t_0-\widetilde{t})}(\mu);$$

then $\widetilde{t}-s < 2\delta$ and $\eta((0,\widetilde{t}]) < 2\delta$ for all $(s,z,\eta) \in A^{(t,x,\mu)} \subseteq B_\delta(t,x,\mu)$. In particular, for all $(s,z,\eta) \in A^{(t,x,\mu)}$, $\mathcal{W}_1(\eta^{(\widetilde{t}}, \eta) < \varphi(\delta)$ for some modulus of continuity φ; hence, there exists $\widetilde{\mu} \in \mathcal{P}$ with $\operatorname{supp}\widetilde{\mu} \subseteq [\widetilde{t},\infty)$, such that $\widetilde{\mu} \preccurlyeq \eta^{(\widetilde{t}}$ and $\mathcal{W}_1(\widetilde{\mu}, \eta^{(\widetilde{t})}) < \varphi(\delta)$ for all $(s,z,\eta) \in A^{(t,x,\mu)}$, where φ was if necessary modified. By choosing $\delta > 0$ small enough, we may then ensure that the thus defined set $A^{(t,x,\mu)}$ lies in $B_{\delta(t,x,\mu)}$ and together with the points $(\widetilde{t},\widetilde{\mu})$ and $\underset{\sim}{t} < \widetilde{t}$ satisfies (4.3).

The collection of sets $\{A^{(t,x,\mu)} : (t,x,\mu) \in S_0\}$ provides an open cover of S_0; since we are dealing with a Polish space, there exists a countable subcover. We denote the latter by $(B_i)_{i\in\mathbb{N}}$, and associate with each B_i the reference points $(t_i,x_i,\mu_i) \mathrel{\widehat{=}} (t,x,\mu)$, $(\widetilde{t}_i,\widetilde{\mu}_i) \mathrel{\widehat{=}} (\widetilde{t},\widetilde{\mu})$ and $\underset{\sim}{t_i} \mathrel{\widehat{=}} \underset{\sim}{t}$. Defining

$$A_1 = B_1 \cap S_0,\, A_{i+1} = \big(B_{i+1} \setminus (B_1 \cup \cdots \cup B_i)\big) \cap S_0, \quad i \geq 1,$$

we then obtain a measurable partition of S_0 with the required properties. □

Our direct proof of the DPP relies on continuity of v; we now highlight the difficulties associated with relaxing this to right-continuity only. For simplicity, consider the case when the only state process is the MVM itself and the value function is given by $\mu \mapsto v(\mu)$. Let $\Xi = (\mu_i)_{i\in\mathbb{N}}$ be the countable family of measures in $\mathcal{P}$ whose cumulative distribution functions are "step-functions" with step-lengths and -heights given by multiples of rational numbers. Given $\varepsilon > 0$, one can then cover $\mathcal{P}$ by countably many sets of the form $B_\varepsilon^{\preccurlyeq}(\mu_i)$, $i \in \mathbb{N}$, where

$$B_\varepsilon^{\preccurlyeq}(\eta) := \{\mu \succcurlyeq \eta : \mathcal{W}_1(\mu,\eta) < \varepsilon\}. \tag{B.1}$$

Given $\delta > 0$, choosing $\varepsilon > 0$ such that $v(\mu) \leq v(\mu_i)+\delta$, for $\mu \in B_\varepsilon^{\preccurlyeq}(\mu_i)$, $i \in \mathbb{N}$, a measurable optimiser can be constructed—and the DPP proven—using similar arguments to those in Section 4.2. While this argument relies on right-continuity only as opposed to continuity, it notably requires *uniform* (semi) right-continuity of v. To generalise arguments of this type beyond uniform (semi) right-continuity seems difficult. Indeed, consider a set $S \subset \mathcal{P}$ and a cover of the form $\bigcup_{\eta\in S} B_{\varepsilon_\eta}^{\preccurlyeq}(\eta)$, where ε_η is chosen such that $v(\mu) \leq v(\eta) + \delta$, for $\mu \in B_{\varepsilon_\eta}^{\preccurlyeq}(\eta)$, $\eta \in S$. To proceed along the same lines as above, one would like to find a countable subcover; however, it is not clear that the topology defined with respect to open sets of the form (B.1) would be Lindelöf. Indeed, note first that the second countable topology generated by $\{B_q^{\preccurlyeq}(\mu) : q \in \mathbb{Q}, \mu \in \Xi\}$ is strictly coarser than the topology generated by all sets of the form (B.1) since for $\mu \in B_\varepsilon^{\preccurlyeq}(\eta)$, $\varepsilon > 0$ and $\eta \in \mathcal{P}$, there need not exist $q \in \mathbb{Q}$ and $\mu_i \in \Xi$ such that $\mu \in B_q^{\preccurlyeq}(\mu_i) \subseteq B_\varepsilon^{\preccurlyeq}(\eta)$. Moreover, note that considering sets of the form

$$B_\varepsilon^{\ll}(\eta) := \{\mu \gg \eta : \mathcal{W}_1(\mu,\eta) < \varepsilon\} \subset B_\varepsilon^{\preccurlyeq}(\eta), \tag{B.2}$$

where $\ll$ denotes the "way-below" relation associated with the stochastic ordering (see [22]), the topology generated by $\{B_q^{\ll}(\mu) : q \in \mathbb{Q}, \mu \in \Xi\}$ does coincide with the topology generated by all sets of the form (B.2); cf. the so-called Scott topology [22]. This topology is therefore second countable, hence Lindelöf, and any cover of the form $\bigcup_{\eta\in S} B_{\varepsilon_\eta}^{\ll}(\eta)$ admits a countable subcover. It is tempting to think that this could be used to argue that the topology generated by (B.1) would be Lindelöf even if not second countable—in the same way as the Sorgenfrey line is; however, since the anti-chains for the stochastic ordering are not countable, the set $S \setminus \bigcup_{\eta\in S} B_{\varepsilon_\eta}^{\ll}(\eta)$ cannot be covered in a countable way and analogous arguments do not apply.

We finally provide the details for the direct proof of the DPP using covering arguments.

PROOF OF COROLLARY 4.5. (Imposing Assumption 4.6 and continuity of v on S_0). Without loss of generality, let $t=0$. Given $\xi \in \mathsf{MVM}(\mu)$, we then introduce the modified MVM ξ^ε as follows: $\xi^\varepsilon_{\cdot\wedge\theta} := \xi_{\cdot\wedge\theta}$ and

$$\xi^\varepsilon_{\cdot\vee\theta}(A) := \xi_\theta\big(A\cap(0,\theta]\big) + \xi_\theta\big((\theta,\infty)\big)\sum_{i\in\mathbb{N}} \mathbf{1}_{\{(\theta, X^x_\theta, \xi^{(\theta}_\theta)\in A_i\}} \xi^{i,\theta,\xi^{(\theta}_\theta}_{\cdot\vee\theta}(A), \qquad A\in\mathcal{B},$$

where $(A_i)_{i\in\mathbb{N}}$ is the partition of S_0 provided by Lemma 4.8 with associated reference points, and where $\xi^{i,s,\eta}_{\cdot\vee s}$, $(s,z,\eta)\in A_i$, is the MVM given by (4.2) with $(\widetilde{t},\widetilde{\mu}) \widehat{=} (\widetilde{t}_i,\widetilde{\mu}_i)$, $\underset{\sim}{t} \widehat{=} \underset{\sim}{t_i}$ and $\widetilde{\xi} \widehat{=} \xi_i$ with $\xi_i \in \mathsf{MVM}^{\underset{\sim}{t_i}}(\widetilde{\mu}_i)$ chosen such that

$$\mathbb{E}\bigg[\int f\big(X_u^{\underset{\sim}{t_i},x_i}\big)\,\mathrm{d}A^{\xi_i}_u\bigg] \ge v(\underset{\sim}{t_i}, x_i, \widetilde{\mu}_i) - \delta(t_i,x_i,\mu_i);$$

we note that the thus defined process is a well-defined continuous and adapted MVM with $\xi^\varepsilon_0=\mu$. By use of Proposition 4.7 and the facts that X is a Markov process and ξ_i is independent of $\mathcal{F}_{\underset{\sim}{t_i}}$, we obtain the following on $\{(\theta, X^x_\theta, \xi^{(\theta)}_\theta)\in A_i\}$:

$$\mathbb{E}\bigg[\int f(X^x_s)\xi^{i,\theta,\xi^{(\theta}_\theta}_\infty(\mathrm{d}s)|\mathcal{F}_\theta\bigg] = \mathbb{E}\bigg[\int f(X^{t,x}_s)\xi^{i,t,\mu}_\infty(\mathrm{d}s)\bigg]_{(t,x,\mu)=(\theta,X^x_\theta,\xi^{(\theta)}_\theta)} \ge v(\theta, X^x_\theta, \xi^{(\theta)}_\theta) - \varepsilon.$$

Hence,

$$\begin{aligned} &v(0,x,\mu)\\ &\quad\ge \mathbb{E}\bigg[\int_0^\infty f(X^x_s)\,\mathrm{d}A^{\xi^\varepsilon}_s\bigg]\\ &\quad= \mathbb{E}\bigg[\int_0^\theta f(X^x_s)\xi_\theta(\mathrm{d}s) + \xi_\theta\big((\theta,\infty)\big)\sum_{i\in\mathbb{N}}\mathbf{1}_{\{(\theta,X^x_\theta,\xi^{(\theta)}_\theta)\in A_i\}}\mathbb{E}\bigg[\int f(X^x_s)\xi^{i,\theta,\xi^{(\theta}_\theta}_\infty(\mathrm{d}s)|\mathcal{F}_\theta\bigg]\bigg]\\ &\quad\ge \mathbb{E}\bigg[\int_0^\theta f(X^x_s)\xi_\theta(\mathrm{d}s) + \xi_\theta\big((\theta,\infty)\big)v\big(\theta, X^x_\theta,\xi^{(\theta)}_\theta\big)\bigg] - \varepsilon. \end{aligned}$$

Since $\varepsilon>0$ and $\xi\in\mathsf{MVM}(\mu)$ were both arbitrarily chosen, this yields the first inequality.

In order to show the reverse inequality, let $\xi\in\mathsf{MVM}(\mu)$ and let θ be a finite $\mathbb{F}$-stopping time. The so-called pseudo Markov property notably holds (cf. [12]), and for almost all $\omega\in\Omega$ we have that

$$\begin{aligned} \mathbb{E}\bigg[\int_\theta^\infty f(X^x_s)\xi^{(\theta}_\infty(\mathrm{d}s)|\mathcal{F}_\theta\bigg](\omega) &= \int_\Omega\int_{\theta(\omega)}^\infty f\big(X_s^{\theta(\omega),X^x_{\theta(\omega)}}(\widetilde{\omega})\big)\widetilde{\xi}^{\theta(\omega),\omega}_\infty(\widetilde{\omega};\mathrm{d}s)\,\mathrm{d}\mathbb{W}(\widetilde{\omega})\\ &\le v\big(\theta(\omega), X^x_{\theta(\omega)}(\omega), \xi^{(\theta(\omega)}_{\theta(\omega)}(\omega)\big), \end{aligned} \tag{B.3}$$

where $\widetilde{\xi}^{\theta(\omega),\omega}_u(\widetilde{\omega}) = \xi^{(\theta(\omega)}_{\theta(\omega)}(\omega)\mathbf{1}_{\{u<\theta(\omega)\}} + \xi^{(\theta(\omega)}_u(\omega *_{\theta(\omega)}\widetilde{\omega})\mathbf{1}_{\{u\ge\theta(\omega)\}}$ with the concatenated path given by $\omega *_t \widetilde{\omega}(s) = \mathbf{1}_{\{0\le s<t\}}\omega(s) + \mathbf{1}_{\{t\le s\}}(\omega(t)+\widetilde{\omega}(s)-\widetilde{\omega}(t))$; indeed, for $\omega\in\Omega$ fixed, $\widetilde{\xi}^{\theta(\omega),\omega}_\cdot$ is independent of $\mathcal{F}_\theta$ and thus lies in $\mathsf{MVM}^{\theta(\omega)}(\xi^{(\theta(\omega)}_{\theta(\omega)}(\omega))$. In consequence,

$$\begin{aligned} \mathbb{E}\bigg[\int_0^\infty f(X^x_s)\,\mathrm{d}A^\xi_s\bigg] &= \mathbb{E}\bigg[\int_0^\theta f(X^x_s)\xi_\theta(\mathrm{d}s) + \xi_\theta\big((\theta,\infty)\big)\mathbb{E}\bigg[\int_\theta^\infty f(X^x_s)\xi^{(\theta}_\infty(\mathrm{d}s)|\mathcal{F}_\theta\bigg]\bigg]\\ &\le \mathbb{E}\bigg[\int_0^\theta f(X^x_s)\xi_\theta(\mathrm{d}s) + \xi_\theta\big((\theta,\infty)\big)v\big(\theta, X^x_\theta,\xi^{(\theta)}_\theta\big)\bigg], \end{aligned}$$

which completes the proof since ξ and θ were arbitrarily chosen. □

Funding. The author gratefully acknowledges financial support from the Swedish Research Council (Grant 2020-03449).

REFERENCES

[1] AKSAMIT, A., HOU, Z. and OBŁÓJ, J. (2020). Robust framework for quantifying the value of information in pricing and hedging. *SIAM J. Financial Math.* **11** 27–59.

[2] ANKIRCHNER, S., KLEIN, M. and KRUSE, T. (2019). A verification theorem for optimal stopping problems with expectation constraints. *Appl. Math. Optim.* **79** 145–177.

[3] BACKHOFF, J., BEIGLBÖCK, M., LIN, Y. and ZALASHKO, A. (2017). Causal transport in discrete time and applications. *SIAM J. Optim.* **27** 2528–2562.

[4] BAXTER, J. R. and CHACON, R. V. (1977). Compactness of stopping times. *Z. Wahrsch. Verw. Gebiete* **40** 169–181.

[5] BAYRAKTAR, E. and MILLER, C. W. (2019). Distribution-constrained optimal stopping. *Math. Finance* **29** 368–406.

[6] BEIGLBÖCK, M., COX, A. M. G. and HUESMANN, M. (2017). Optimal transport and Skorokhod embedding. *Invent. Math.* **208** 327–400.

[7] BEIGLBÖCK, M., COX, A. M. G., HUESMANN, M. and KÄLLBLAD, S. (2017). Measure-valued martingales and optimality of Bass-type solutions to the Skorokhod Embedding Problem. ArXiv preprint. Available at arXiv:1708.07071.

[8] BEIGLBÖCK, M., EDER, M., ELGERT, C. and SCHMOCK, U. (2018). Geometry of distribution-constrained optimal stopping problems. *Probab. Theory Related Fields* **172** 71–101.

[9] BEIGLBÖCK, M. and LACKER, D. (2018). Denseness of adapted processes among causal couplings. ArXiv preprint. Available at arXiv:1805.03185.

[10] BOUCHARD, B. and NUTZ, M. (2012). Weak dynamic programming for generalized state constraints. *SIAM J. Control Optim.* **50** 3344–3373.

[11] BOUCHARD, B. and TOUZI, N. (2011). Weak dynamic programming principle for viscosity solutions. *SIAM J. Control Optim.* **49** 948–962.

[12] CLAISSE, J., TALAY, D. and TAN, X. (2016). A pseudo-Markov property for controlled diffusion processes. *SIAM J. Control Optim.* **54** 1017–1029.

[13] COX, A. M. G. and KÄLLBLAD, S. (2017). Model-independent bounds for Asian options: A dynamic programming approach. *SIAM J. Control Optim.* **55** 3409–3436.

[14] DAWSON, D. A. (1993). Measure-valued Markov processes. In *École D'Été de Probabilités de Saint-Flour XXI*—1991. *Lecture Notes in Math.* **1541** 1–260. Springer, Berlin.

[15] EDGAR, G. A., MILLET, A. and SUCHESTON, L. (1982). On compactness and optimality of stopping times. In *Martingale Theory in Harmonic Analysis and Banach Spaces* (*Cleveland*, *Ohio*, 1981). *Lecture Notes in Math.* **939** 36–61. Springer, Berlin.

[16] EKSTRÖM, E. and JANSON, S. (2016). The inverse first-passage problem and optimal stopping. *Ann. Appl. Probab.* **26** 3154–3177.

[17] EL KAROUI, N. (1981). Les aspects probabilistes du contrôle stochastique. In *Ecole D'Ete de Probabilites de Saint-Flour IX*-1979. *Lecture Notes in Math.* **876** 73–238. Springer, Berlin.

[18] EL KAROUI, N. and TAN, X. (2013). Capacities, measurable selection and dynamic programming. Part I: Abstract framework. Available at arXiv:1310.3363.

[19] EL KAROUI, N. and TAN, X. (2013). Capacities, measurable selection and dynamic programming. Part II: Application in stochastic control problems. Available at arXiv:1310.3364.

[20] GETOOR, R. K. (1975). On the construction of kernels. In *Séminaire de Probabilités*, *IX* (*Seconde Partie*, *Univ. Strasbourg*, *Strasbourg*, *Années Universitaires* 1973/1974 *et* 1974/1975). *Lecture Notes in Math.* **465** 443–463.

[21] GHOUSSOUB, N. (1982). An integral representation of randomized probabilities and its applications. In *Seminar on Probability*, *XVI*. *Lecture Notes in Math.* **920** 519–543. Springer, Berlin.

[22] GIERZ, G., HOFMANN, K. H., KEIMEL, K., LAWSON, J. D., MISLOVE, M. W. and SCOTT, D. S. (1980). *A Compendium of Continuous Lattices*. Springer, Berlin.

[23] GRASS, A.-M. (2016). Uniqueness and stability properties of barrier type Skorokhod embeddings. Masther Thesis, Univ. Vienna.

[24] GUO, G., TAN, X. and TOUZI, N. (2016). Optimal Skorokhod embedding under finitely many marginal constraints. *SIAM J. Control Optim.* **54** 2174–2201.

[25] HOROWITZ, J. (1985). Measure-valued random processes. *Z. Wahrsch. Verw. Gebiete* **70** 213–236.

[26] JACOD, J. and MÉMIN, J. (1981). Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité. In *Seminar on Probability*, *XV* (*Univ. Strasbourg*, *Strasbourg*, 1979/1980) (*French*). *Lecture Notes in Math.* **850** 529–546. Springer, Berlin.

[27] KÄLLBLAD, S., TAN, X. and TOUZI, N. (2017). Optimal Skorokhod embedding given full marginals and Azéma-Yor peacocks. *Ann. Appl. Probab.* **27** 686–719.

[28] LOYNES, R. M. (1970). Stopping times on Brownian motion: Some properties of Root's construction. *Z. Wahrsch. Verw. Gebiete* **16** 211–218.

[29] MEYER, P. A. (1978). Convergence faible et compacité des temps d'arrêt d'après Baxter et Chacon. In *Séminaire de Probabilités, XII* (*Univ. Strasbourg, Strasbourg*, 1976/1977). *Lecture Notes in Math.* **649** 411–423. Springer, Berlin.

[30] MILLER, C. W. (2017). Nonlinear PDE approach to time-inconsistent optimal stopping. *SIAM J. Control Optim.* **55** 557–573.

[31] NEUFELD, A. and NUTZ, M. (2013). Superreplication under volatility uncertainty for measurable claims. *Electron. J. Probab.* **18** no. 48.

[32] NUTZ, M. and VAN HANDEL, R. (2013). Constructing sublinear expectations on path space. *Stochastic Process. Appl.* **123** 3100–3121.

[33] ROOT, D. H. (1969). The existence of certain stopping times on Brownian motion. *Ann. Math. Stat.* **40** 715–718.

[34] ROST, H. (1971). The stopping distributions of a Markov process. *Invent. Math.* **14** 1–16.

[35] STROOCK, D. W. and VARADHAN, S. R. S. (1979). *Multidimensional Diffusion Processes. Grundlehren der Mathematischen Wissenschaften* [*Fundamental Principles of Mathematical Sciences*] **233**. Springer, Berlin.

[36] VILLANI, C. (2009). *Optimal Transport: Old and New. Grundlehren der Mathematischen Wissenschaften*; *Fundamental Principles of Mathematical Sciences* **338**. Springer, Berlin.

[37] WALSH, J. B. (1986). An introduction to stochastic partial differential equations. In *École D'été de Probabilités de Saint-Flour, XIV*—1984. *Lecture Notes in Math.* **1180** 265–439. Springer, Berlin.

[38] ŽITKOVIĆ, G. (2014). Dynamic programming for controlled Markov families: Abstractly and over martingale measures. *SIAM J. Control Optim.* **52** 1597–1621.