

Extended Mean Field Control Problems: stochastic maximum principle and transport perspective

Beatrice Acciaio* Julio Backhoff-Veraguas[†] René Carmona[‡]

August 21, 2019

Abstract

We study Mean Field stochastic control problems where the cost function and the state dynamics depend upon the joint distribution of the controlled state and the control process. We prove suitable versions of the Pontryagin stochastic maximum principle, both in necessary and in sufficient form, which extend the known conditions to this general framework. We suggest a variational approach for a weak formulation of these control problems. We show a natural connection between this weak formulation and optimal transport on path space, which inspires a novel discretization scheme.

Keywords: Controlled McKean-Vlasov SDEs; Pontryagin principle; Mean-Field interaction; Causal transport plans

MSC2010 subject classifications: 93E20, 90C08, 60H30, 60K35.

1 Introduction

The control of stochastic differential equations of Mean Field type, also known as McKean-Vlasov control, did not get much attention before the theory of Mean Field Games became a popular subject of investigation. Indeed the two topics are intimately related through the asymptotic theory of mean field stochastic systems known as *propagation of chaos*. See for example [15] for an early discussion of the similarities and the differences of the two problems. Among the earliest works on this new form of control problem, relevant to the spirit of the analysis conducted in this paper, are [10, 9, 3, 28, 8, 13]. Here, we follow the approach introduced and developed in [13]. The reader is referred to [14, Ch. 3, 4, 6] for a general overview of these problems and an extensive historical perspective. Still, most of these contributions are limited to Mean Field interactions entering the models through the statistical distribution of the state of the system alone. The goal of the present article is to investigate the control of stochastic dynamics depending upon the joint distribution of the controlled state and the control process. We refer to such problems as *extended Mean Field control problems*; see [14, Sec. 4.6].

*Department of Statistics, London School of Economics

[†]Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology

[‡]Operations Research and Financial Engineering, Princeton University, Partially supported by National Science Foundation # DMS-1716673 and Army Research Office # W911NF-17-1-0578

Our first contribution is to prove an appropriate form of the Pontryagin stochastic maximum principle, in necessary and in sufficient form, for extended Mean Field control problems. The main driver behind this search for an extension of existing tools is the importance of many practical applications, which naturally fit within the class of models for which the interactions are not only through the distribution of the state of the system, but also through the distribution of the controls. The analysis of extended Mean Field control problems had been restricted so far to the Linear Quadratic (LQ) case; see e.g. [35, 24, 6, 33]. To the best of our knowledge, the recent work [33] is the only one where more general models are considered. In that article, however, the authors restrict the analysis to closed-loop feedback controls, leading to a deterministic reformulation of the problem, which is used in order to derive the Bellman equation associated to the problem; theirs is therefore a PDE approach. In the present paper, we study the extended Mean Field control problem without any restrictions, deriving a version of the Pontryagin maximum principle via a probabilistic approach.

We apply our optimality conditions for particular classes of models, where our analysis can be pushed further. In the case of scalar interactions, in which the dynamics depend solely upon moments of the marginal distributions, we derive a more explicit form of the optimality condition. The advantage here is that the analysis can be conducted with a form of classical differential calculus, without the use of the notion of L-differentiability. The announced work [23] studies an application of such class of models in electricity markets. As a special case of scalar interaction, we study an optimal liquidation model, which we are able to solve explicitly. Finally, we consider the case of LQ models for which we easily derive explicit solutions which can be computed numerically. The results in the LQ setting are compatible with existing results in the literature.

Another contribution of the present article is the variational study of a weak formulation of the extended Mean Field control problem. Weak formulations have already been studied in the literature, without non-linear dependence in the law of the control, as in [14, Ch. 6] and [25]. In this framework, we derive an analogue of the Pontryagin principle in the form of a martingale optimality condition. Similar statements have been derived in [18, 27] under the name of Stochastic Euler-Lagrange condition for a different kind of problems. Next, we derive a natural connection between the extended Mean Field control problem and an optimal transport problem on path space. The theory of optimal transport is known to provide a set of tools and results crucial to the understanding of mean field control and mean field games. We illustrate the use of this connection by building a discretization scheme for extended Mean Field control based on transport-theoretic tools (as in [36, Ch. 3.6] for the case without Mean Field terms), and show that this scheme converges monotonically to the value of the original extended Mean Field control problem. The explosion in activity regarding numerical optimal transport gives us reason to believe that such discretization schemes might be efficiently implemented in the near future; see e.g. [19, 7, 29] for the static setting and [30, 31, 32] for the dynamic one.

The paper is organized as follows. In Section 2, we introduce the notations and basic underpinnings for extended mean field control. Section 3 provides a new form of the Pontryagin stochastic maximum principle. In Section 4, we study classes of models for which our optimality conditions lead to explicit solutions. In Section 5, we analyse the weak formulation of the problem in connection with optimal transport. In the Finally, in the Appendix, we collect some technical proofs.

2 Extended Mean Field Control Problems

The goal of this short subsection is to set the stage for the statements and proofs of the stochastic maximum principle proven in Section 3 below.

Let f , b , and σ be measurable functions on $\mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k)$ with values in \mathbb{R} , \mathbb{R}^d , and $\mathbb{R}^{d \times m}$ respectively, and g be a real valued measurable function on $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$. Here and elsewhere we denote by $\mathcal{P}(\cdot)$ (resp. $\mathcal{P}_2(\cdot)$) the set of probability measures (resp. with finite second moments) over an underlying metric space. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, $\mathcal{F}_0 \subset \mathcal{F}$ be a sub sigma-algebra, and $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ be the filtration generated by \mathcal{F}_0 and an m -dimensional Wiener process $\mathbf{W} = (W_t)_{0 \leq t \leq T}$. We denote by \mathbb{A} the set of progressively measurable processes $\alpha = (\alpha_t)_{0 \leq t \leq T}$ taking values in a given closed-convex set $A \subset \mathbb{R}^k$ and satisfying the integrability condition $\mathbb{E} \int_0^T |\alpha_t|^2 dt < \infty$.

We consider the problem of minimizing

$$J(\alpha) = \mathbb{E} \left[\int_0^T f(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)) dt + g(X_T, \mathcal{L}(X_T)) \right] \quad (2.1)$$

over the set \mathbb{A} of admissible control processes, under the dynamic constraint

$$dX_t = b(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)) dt + \sigma(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)) dW_t, \quad (2.2)$$

with X_0 a fixed \mathcal{F}_0 -measurable random variable.

The symbol \mathcal{L} stands for the law of the given random element. We shall add mild regularity conditions for the coefficients b and σ so that a solution to equation (2.2) always exists when $\alpha \in \mathbb{A}$. For the sake of simplicity, we chose to use time independent coefficients, but all the results would be the same should f , b and σ depend upon t , since time can be included as an extra state in the vector X .

The novelty of the above control problem lies in the fact that the cost functional and the controlled SDE depend on the joint distribution of state and control. For this reason, we call it *extended Mean Field control problem*. In this generality, this problem has not been studied before. We mention the works [35, 24, 6, 33] for particular cases and different approaches.

2.1 Partial L-differentiability of Functions of Measures

We introduce here the concept of L-differentiability for functions of joint probability laws (i.e. probability measures on product spaces). We refer the reader to [14, Ch. 5] for more details.

Let $u : \mathbb{R}^q \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \rightarrow \mathbb{R}$. We use the notation ξ for a generic element of $\mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k)$, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\nu \in \mathcal{P}_2(\mathbb{R}^k)$ for its marginals. We denote a generic element of \mathbb{R}^q by v .

Let $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ be a probability space and let \tilde{u} be a lifting of the function u . In other words:

$$\tilde{u} : \mathbb{R}^q \times L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d \times \mathbb{R}^k) \ni (v, \tilde{X}, \tilde{\alpha}) \mapsto \tilde{u}(v, \tilde{X}, \tilde{\alpha}) = u(v, \mathcal{L}(\tilde{X}, \tilde{\alpha})).$$

We say that u is *L-differentiable* at (v, ξ) if there exists a pair

$$(\tilde{X}, \tilde{\alpha}) \in L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d \times \mathbb{R}^k) \text{ with } \mathcal{L}(\tilde{X}, \tilde{\alpha}) = \xi,$$

such that the lifted function \tilde{u} is Fréchet differentiable at $(v, \tilde{X}, \tilde{\alpha})$; cf. [20, Ch. II.5 p.92]. When this is the case, it turns out that the Fréchet derivative depends only on the law ξ and not on the specific pair $(\tilde{X}, \tilde{\alpha})$ having distribution ξ ; see [11] or [14, Ch. 6] for details. Thanks to self-duality of L^2 spaces, the Fréchet derivative $[D\tilde{u}](v, \tilde{X}, \tilde{\alpha})$ of the lifting function \tilde{u} at $(v, \tilde{X}, \tilde{\alpha})$ can be viewed as an element $D\tilde{u}(v, \tilde{X}, \tilde{\alpha})$ of $\mathbb{R}^q \times L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d \times \mathbb{R}^k)$, in the sense that

$$[D\tilde{u}](v, \tilde{X}, \tilde{\alpha})(\tilde{Y}) = \tilde{\mathbb{E}}[D\tilde{u}(v, \tilde{X}, \tilde{\alpha}) \cdot \tilde{Y}], \quad \text{for all } \tilde{Y} \in \mathbb{R}^q \times L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d \times \mathbb{R}^k).$$

Since $\mathbb{R}^q \times L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d \times \mathbb{R}^k) \cong \mathbb{R}^q \times L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d) \times L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^k)$, as in [11] the random variable $D\tilde{u}(v, \tilde{X}, \tilde{\alpha})$ can be represented a.s. via the random vector

$$D\tilde{u}(v, \tilde{X}, \tilde{\alpha}) = \left(\partial_v u(v, \mathcal{L}(\tilde{X}, \tilde{\alpha}))(\tilde{X}, \tilde{\alpha}), \partial_\mu u(v, \mathcal{L}(\tilde{X}, \tilde{\alpha}))(\tilde{X}, \tilde{\alpha}), \partial_\nu u(v, \mathcal{L}(\tilde{X}, \tilde{\alpha}))(\tilde{X}, \tilde{\alpha}) \right),$$

for measurable functions $\partial_v u(\cdot, \mathcal{L}(\tilde{X}, \tilde{\alpha}))(\cdot, \cdot)$, $\partial_\mu u(\cdot, \mathcal{L}(\tilde{X}, \tilde{\alpha}))(\cdot, \cdot)$, $\partial_\nu u(\cdot, \mathcal{L}(\tilde{X}, \tilde{\alpha}))(\cdot, \cdot)$, all of them defined on $\mathbb{R}^q \times \mathbb{R}^d \times \mathbb{R}^k$ and valued respectively on \mathbb{R}^q , \mathbb{R}^d and \mathbb{R}^k . We call these functions the *partial L-derivatives* of u at $(v, \mathcal{L}(\tilde{X}, \tilde{\alpha}))$.

3 Stochastic Maximum Principle

Our goal is to prove a necessary and a sufficient condition for optimality in the extended class of problems considered in the paper. These are suitable extensions of the Pontryagin stochastic maximum principle conditions. We define the Hamiltonian H by:

$$H(x, \alpha, \xi, y, z) = b(x, \alpha, \xi) \cdot y + \sigma(x, \alpha, \xi) \cdot z + f(x, \alpha, \xi), \quad (3.1)$$

for $(x, \alpha, \xi, y, z) \in \mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \times \mathbb{R}^d \times \mathbb{R}^{d \times m}$. Naturally, the dot notation for matrices refers to the trace inner product. We let $\mathbb{H}^{0,n}$ stand for the collection of all \mathbb{R}^n -valued progressively measurable processes on $[0, T]$, and denote by $\mathbb{H}^{2,n}$ the collection of processes Z in $\mathbb{H}^{0,n}$ such that $\mathbb{E} \int_0^T |Z_s|^2 ds < \infty$. We shall also denote by $\mathbb{S}^{2,n}$ the space of all continuous processes $\mathbf{S} = (S_t)_{0 \leq t \leq T}$ in $\mathbb{H}^{0,n}$ such that $\mathbb{E}[\sup_{0 \leq t \leq T} |S_t|^2] < +\infty$. Here and in what follows, regularity properties, such as continuity or Lipschitz character, of functions of measures are always understood in the sense of the 2-Wasserstein distance of the respective spaces of probability measures with finite second moments; cf. [34].

Throughout this section, we assume:

- (I) The functions b , σ and f are differentiable with respect to (x, α) , for $\xi \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k)$ fixed, and the functions $(x, \alpha, \xi) \mapsto (\partial_x(b, \sigma, f)(x, \alpha, \xi), \partial_\alpha(b, \sigma, f)(x, \alpha, \xi))$ are continuous. Moreover, the functions b , σ and f are L-differentiable with respect to the variable ξ , the mapping

$$\mathbb{R}^d \times A \times L^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^d \times \mathbb{R}^k) \ni (x, \alpha, (X, \beta)) \mapsto \partial_\mu(b, \sigma, f)(x, \alpha, \mathcal{L}(X, \beta))(X, \beta)$$

being continuous. Similarly, the function g is differentiable with respect to x , the mapping $(x, \mu) \mapsto \partial_x g(x, \mu)$ being continuous. The function g is also L-differentiable with respect to the variable μ , and the following map is continuous

$$\mathbb{R}^d \times L^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^d) \ni (x, X) \mapsto \partial_\mu g(x, \mathcal{L}(X))(X) \in L^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^d).$$

(II) The derivatives $\partial_x(b, \sigma)$ and $\partial_\alpha(b, \sigma)$ are uniformly bounded, and the mapping $(x', \alpha') \mapsto \partial_\mu(b, \sigma)(x, \alpha, \xi)(x', \alpha')$ (resp. $(x', \alpha') \mapsto \partial_\nu(b, \sigma)(x, \alpha, \xi)(x', \alpha')$) has an $L^2(\mathbb{R}^d, \mu; \mathbb{R}^d \times \mathbb{R}^k)$ -norm (resp. $L^2(\mathbb{R}^k, \nu; \mathbb{R}^d \times \mathbb{R}^k)$ -norm) which is uniformly bounded in (x, α, ξ) . There exists a constant L such that, for any $R \geq 0$ and any (x, α, ξ) such that $|x|, |\alpha|, \|\xi\|_{L^2} \leq R$, it holds that

$$|\partial_x f(x, \alpha, \xi)| \vee |\partial_x g(x, \mu)| \vee |\partial_\alpha f(x, \alpha, \xi)| \leq L(1 + R),$$

and the norms in $L^2(\mathbb{R}^d \times \mathbb{R}^k, \xi; \mathbb{R}^d \times \mathbb{R}^k)$ and $L^2(\mathbb{R}^d, \xi; \mathbb{R}^d \times \mathbb{R}^k)$ of $(x', \alpha') \mapsto \partial_\mu f(x, \alpha, \xi)(x', \alpha')$, $(x', \alpha') \mapsto \partial_\nu f(x, \alpha, \xi)(x', \alpha')$ and $x' \mapsto \partial_\mu g(x, \mu)(x')$ are bounded by $L(1 + R)$.

Under these assumptions, for any admissible control $\alpha \in \mathbb{A}$, we denote by $\mathbf{X} = \mathbf{X}^\alpha$ the corresponding controlled state process satisfying (2.2). We call *adjoint processes* of \mathbf{X} (or of α), the couple (\mathbf{Y}, \mathbf{Z}) of stochastic processes $\mathbf{Y} = (Y_t)_{0 \leq t \leq T}$ and $\mathbf{Z} = (Z_t)_{0 \leq t \leq T}$ in $\mathbb{S}^{2,d} \times \mathbb{H}^{2,d \times m}$ that satisfy:

$$\begin{cases} dY_t = -\left[\partial_x H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}}[\partial_\mu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)]\right] dt + Z_t dW_t, & t \in [0, T], \\ Y_T = \partial_x g(X_T, \mathcal{L}(X_T)) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T, \mathcal{L}(X_T))(X_T)], \end{cases} \quad (3.2)$$

where $\theta_t := (X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t))$, and the tilde notation refers to an independent copy. Equation (3.2) is referred to as the *adjoint equation*. Formally, the adjoint variable Y_t reads as the derivative of the value function of the control problem with respect to the state variable. In contrast with the deterministic case, in order for the solution to be adapted to the information flow, the extra term $Z_t dW_t$ is needed. This is a standard feature of the extension of the maximum principle from deterministic control to stochastic control. As expected, it is driven by the derivative of the Hamiltonian function with respect to the state variable. In addition, since the controlled dynamics are of the McKean-Vlasov type, the state variable with respect to which we differentiate the Hamiltonian function needs to include the probability measure appearing in the state equation. This is now understood thanks to the early contributions [13] and [14, Ch. 6]. In the present case of extended Mean Field control problems, the above adjoint equation needed to account for the fact that the probability measure appearing in the state equation is in fact the joint distribution of the state X_t and the control α_t . This forces us to involve the derivative of the Hamiltonian with respect to the first marginal of this joint distribution.

Given α and as a result \mathbf{X} , θ_t appears as a (random) input in the coefficients of this equation which, except for the presence of the process copies, is a backward stochastic differential equation of the McKean-Vlasov type, which is well posed under the current assumptions. See for example the discussion in [14, Ch. 6, p.532].

3.1 A Necessary Condition

The main result of this subsection is based on the following expression of the Gâteaux derivative of the cost function $J(\alpha)$.

Lemma 3.1. *Let $\alpha \in \mathbb{A}$, \mathbf{X} be the corresponding controlled state process, and (\mathbf{Y}, \mathbf{Z}) its adjoint processes satisfying (3.2). For $\beta \in \mathbb{A}$, the Gâteaux derivative of J at α in the direction $\beta - \alpha$ is*

$$\frac{d}{d\epsilon} J(\alpha + \epsilon(\beta - \alpha)) \Big|_{\epsilon=0} = \mathbb{E} \int_0^T \left(\partial_\alpha H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}}[\partial_\nu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] \right) \cdot (\beta_t - \alpha_t) dt,$$

where $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}, \tilde{\mathbf{Z}}, \tilde{\alpha}, \tilde{\beta})$ is an independent copy of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \alpha, \beta)$ on the space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

Proof. We follow the lines of the proof of the stochastic maximum principle for the control of McKean-Vlasov equations given in [14, Sec. 6.3]. Given admissible controls α and β , for each $\epsilon > 0$ we define the admissible control $\alpha^\epsilon = (\alpha_t^\epsilon)_{0 \leq t \leq T}$ by $\alpha_t^\epsilon = \alpha_t + \epsilon(\beta_t - \alpha_t)$, and we denote by $\mathbf{X}^\epsilon = (X_t^\epsilon)_{0 \leq t \leq T}$ the solution of the state equation (2.2) for α^ϵ in lieu of α . We then consider the variation process $\mathbf{V} = (V_t)_{0 \leq t \leq T}$, defined as the solution of the linear stochastic differential equation:

$$dV_t = [\gamma_t V_t + \rho_t + \eta_t] dt + [\hat{\gamma}_t V_t + \hat{\rho}_t + \hat{\eta}_t] dW_t, \quad (3.3)$$

with $V_0 = 0$. The coefficients $\gamma_t, \hat{\gamma}_t, \eta_t$ and $\hat{\eta}_t$ are defined by

$$\gamma_t = \partial_x b(\theta_t), \quad \hat{\gamma}_t = \partial_x \sigma(\theta_t), \quad \eta_t = \partial_\alpha b(\theta_t)(\beta_t - \alpha_t), \quad \hat{\eta}_t = \partial_\alpha \sigma(\theta_t)(\beta_t - \alpha_t),$$

which are progressively measurable bounded processes with values in the spaces $\mathbb{R}^{d \times d}$, $\mathbb{R}^{(d \times d) \times d}$, \mathbb{R}^d , and $\mathbb{R}^{d \times d}$, respectively (the parentheses around $d \times d$ indicating that $\hat{\gamma}_t \cdot u$ is seen as an element of $\mathbb{R}^{d \times d}$ whenever $u \in \mathbb{R}^d$). The coefficients ρ_t and $\hat{\rho}_t$ are given by

$$\begin{aligned} \rho_t &= \tilde{\mathbb{E}}[\partial_\mu b(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)\tilde{V}_t] + \tilde{\mathbb{E}}[\partial_\nu b(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)(\tilde{\beta}_t - \tilde{\alpha}_t)], \\ \hat{\rho}_t &= \tilde{\mathbb{E}}[\partial_\mu \sigma(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)\tilde{V}_t] + \tilde{\mathbb{E}}[\partial_\nu \sigma(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)(\tilde{\beta}_t - \tilde{\alpha}_t)], \end{aligned}$$

which are progressively measurable bounded processes with values in \mathbb{R}^d and $\mathbb{R}^{d \times d}$, respectively, and where $(\tilde{X}_t, \tilde{\alpha}_t, \tilde{V}_t, \tilde{\beta}_t)$ is an independent copy of $(X_t, \alpha_t, V_t, \beta_t)$ defined on separate probability structure $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

We call $\mathbf{V} = (V_t)_{0 \leq t \leq T}$ the variation process because it is the Gâteaux derivative of the state in the direction $\beta - \alpha$, since, as detailed in [14, Lemma 6.10], it satisfies:

$$\lim_{\epsilon \searrow 0} \mathbb{E} \left[\sup_{0 \leq t \leq T} \left| \frac{X_t^\epsilon - X_t}{\epsilon} - V_t \right|^2 \right] = 0.$$

For this reason, we have:

$$\begin{aligned} \lim_{\epsilon \searrow 0} \frac{1}{\epsilon} [J(\alpha^\epsilon) - J(\alpha)] &= \mathbb{E} \int_0^T \left(\partial_x f(\theta_t) V_t + \partial_\alpha f(\theta_t)(\beta_t - \alpha_t) \right. \\ &\quad \left. + \tilde{\mathbb{E}}[\partial_\mu f(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)\tilde{V}_t] + \tilde{\mathbb{E}}[\partial_\nu f(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)(\tilde{\beta}_t - \tilde{\alpha}_t)] \right) dt \\ &\quad + \mathbb{E} \left[\partial_x g(X_T, \mathcal{L}(X_T)) V_T + \tilde{\mathbb{E}}[\partial_\mu g(X_T, \mathcal{L}(X_T))(\tilde{X}_T)\tilde{V}_T] \right] \\ &= \mathbb{E} \int_0^T \left(\partial_x f(\theta_t) V_t + \partial_\alpha f(\theta_t)(\beta_t - \alpha_t) \right. \\ &\quad \left. + \tilde{\mathbb{E}}[\partial_\mu f(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)\tilde{V}_t] + \tilde{\mathbb{E}}[\partial_\nu f(\theta_t)(\tilde{X}_t, \tilde{\alpha}_t)(\tilde{\beta}_t - \tilde{\alpha}_t)] \right) dt \\ &\quad + \mathbb{E} \left[(\partial_x g(X_T, \mathcal{L}(X_T)) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T, \mathcal{L}(X_T))(\tilde{X}_T)]) V_T \right], \end{aligned} \quad (3.4)$$

where we used Fubini's theorem to obtain the last equality. Notice that, if we introduce the adjoint processes (\mathbf{Y}, \mathbf{Z}) of $\alpha \in \mathbb{A}$ and the corresponding state process \mathbf{X} , by (3.2), we see that the last expectation above is exactly $\mathbb{E}[Y_T V_T]$. This can be computed by integration by parts, using the Itô differentials of \mathbf{Y} and \mathbf{V} , which are given respectively by (3.2) and (3.3). In this way we obtain:

$$\begin{aligned} Y_T V_T &= Y_0 V_0 + \int_0^T Y_t dV_t + \int_0^T V_t dY_t + \int_0^T d[Y, V]_t \\ &= M_T + \int_0^T \left[Y_t \partial_x b(\theta_t) V_t + Y_t \partial_\alpha b(\theta_t) (\beta_t - \alpha_t) + Y_t \tilde{\mathbb{E}} \left[\partial_\mu b(\theta_t) (\tilde{X}_t, \tilde{\alpha}_t) \tilde{V}_t \right] \right. \\ &\quad + Y_t \tilde{\mathbb{E}} \left[\partial_\nu b(\theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] - V_t \partial_x b(\theta_t) Y_t - V_t \partial_x \sigma(\theta_t) Z_t - V_t \partial_x f(\theta_t) \\ &\quad - V_t \tilde{\mathbb{E}} \left[\partial_\mu b(\tilde{\theta}_t) (X_t, \alpha_t) \tilde{Y}_t \right] - V_t \tilde{\mathbb{E}} \left[\partial_\mu \sigma(\tilde{\theta}_t) (X_t, \alpha_t) \tilde{Z}_t \right] - V_t \tilde{\mathbb{E}} \left[\partial_\mu f(\tilde{\theta}_t) (X_t, \alpha_t) \right] \\ &\quad + Z_t \partial_x \sigma(\theta_t) V_t + Z_t \partial_\alpha \sigma(\theta_t) (\beta_t - \alpha_t) + Z_t \tilde{\mathbb{E}} \left[\partial_\mu \sigma(\theta_t) (\tilde{X}_t, \tilde{\alpha}_t) \tilde{V}_t \right] \\ &\quad \left. + Z_t \tilde{\mathbb{E}} \left[\partial_\nu \sigma(t, \theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] \right] dt, \end{aligned}$$

where $(M_t)_{0 \leq t \leq T}$ is a mean zero integrable martingale which disappears when we take expectations of both sides. Applying Fubini's theorem once more, we have:

$$\begin{aligned} \mathbb{E}[Y_T V_T] &= \mathbb{E} \int_0^T \left[Y_t \partial_x b(\theta_t) V_t + Y_t \partial_\alpha b(\theta_t) (\beta_t - \alpha_t) + Y_t \tilde{\mathbb{E}} \left[\partial_\nu b(\theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] \right. \\ &\quad - V_t \partial_x b(\theta_t) Y_t - V_t \partial_x \sigma(\theta_t) Z_t - V_t \partial_x f(\theta_t) - V_t \tilde{\mathbb{E}} \left[\partial_\mu f(\tilde{\theta}_t) (X_t, \alpha_t) \right] \\ &\quad \left. + Z_t \partial_x \sigma(\theta_t) V_t + Z_t \partial_\alpha \sigma(\theta_t) (\beta_t - \alpha_t) + Z_t \tilde{\mathbb{E}} \left[\partial_\nu \sigma(t, \theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] \right] dt, \end{aligned}$$

Plugging this expression in the second equality of (3.4) we get, again by Fubini's theorem,

$$\begin{aligned} \lim_{\epsilon \searrow 0} \frac{1}{\epsilon} [J(\alpha^\epsilon) - J(\alpha)] &= \mathbb{E} \int_0^T \left(\partial_\alpha f(\theta_t) (\beta_t - \alpha_t) + \tilde{\mathbb{E}} \left[\partial_\nu f(\theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] \right. \\ &\quad + Y_t \partial_\alpha b(\theta_t) (\beta_t - \alpha_t) + Y_t \tilde{\mathbb{E}} \left[\partial_\nu b(\theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] \\ &\quad \left. + Z_t \partial_\alpha \sigma(\theta_t) (\beta_t - \alpha_t) + Z_t \tilde{\mathbb{E}} \left[\partial_\nu \sigma(t, \theta_t) (\tilde{X}_t, \tilde{\alpha}_t) (\tilde{\beta}_t - \tilde{\alpha}_t) \right] \right) dt, \end{aligned}$$

which is the desired result, by (3.1). \square

We are now ready to prove the necessary part of the Pontryagin stochastic maximum principle. In the present framework of extended Mean Field control, we obtain (3.5) below. It is not possible to improve this condition into a pointwise minimization condition as in more classical versions of the problem, when there is no non-linear dependence on the law of the control, see (6.58) in [14]. We give an example of this phenomenon in Remark 4.2.

Theorem 3.2. *Under Assumptions (I)-(II), if the admissible control $\alpha = (\alpha_t)_{0 \leq t \leq T} \in \mathbb{A}$ is optimal, $\mathbf{X} = (X_t)_{0 \leq t \leq T}$ is the associated controlled state given by (2.2), and $(\mathbf{Y}, \mathbf{Z}) = (Y_t, Z_t)_{0 \leq t \leq T}$ are the associated adjoint processes satisfying (3.2), then we have:*

$$\left(\partial_\alpha H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}} \left[\partial_\nu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t) (X_t, \alpha_t) \right] \right) \cdot (\alpha_t - a) \leq 0 \quad \forall a \in A, \quad dt \otimes d\mathbb{P} \text{ -a.s.}, \quad (3.5)$$

where $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}, \tilde{\mathbf{Z}}, \tilde{\alpha})$ is an independent copy of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \alpha)$ on $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

Proof. Given any admissible control β , we use as before the perturbation $\alpha_t^\epsilon = \alpha_t + \epsilon(\beta_t - \alpha_t)$. Since α is optimal, we have the inequality

$$\frac{d}{d\epsilon} J(\alpha + \epsilon(\beta - \alpha)) \Big|_{\epsilon=0} \geq 0.$$

Using the result of the previous lemma, we get:

$$\mathbb{E} \int_0^T \left(\partial_\alpha H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}}[\partial_\nu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] \right) \cdot (\beta_t - \alpha_t) dt \geq 0.$$

We now use the same argument as in the classical case (see e.g. [14, Theorem 6.14]). For every t and $\beta \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; A)$, we can take β_t equal to α_t except for the interval $[t, t + \varepsilon]$ where it equals β , obtaining

$$\mathbb{E} \left(\left(\partial_\alpha H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}}[\partial_\nu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] \right) \cdot (\beta - \alpha_t) \right) \geq 0. \quad (3.6)$$

Further, for any $a \in A$ we can take β to be equal to a on an arbitrary set in \mathcal{F}_t , and to coincide with α_t otherwise, establishing equation (3.5). \square

Remark 3.3. If the admissible optimal control α takes values in the interior of A , then we may replace (3.5) with the following condition (see e.g. [14, Proposition 6.15]):

$$\partial_\alpha H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}}[\partial_\nu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] = 0 \quad dt \otimes d\mathbb{P} \text{-a.s.} \quad (3.7)$$

Remark 3.4. A sharpening of (3.5) can be obtained under the convexity condition:

$$\begin{aligned} H(x, a', \xi', y, z) &\geq H(x, a, \xi, y, z) + \partial_\alpha H(x, a, \xi, y, z) \cdot (a' - a) \\ &\quad + \tilde{\mathbb{E}} \left[\partial_\nu H(x, a, \xi, y, z)(\tilde{X}_t, \tilde{\alpha}_t) \cdot (\tilde{\alpha}'_t - \tilde{\alpha}_t) \right], \end{aligned} \quad (3.8)$$

for all $x \in \mathbb{R}^d$, $a, a' \in A$, and $\tilde{\alpha}'$ a copy on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ of an admissible control α' , and where $\xi, \xi' \in \mathcal{P}_2(\mathbb{R}^d \times A)$ with $\xi = \mathcal{L}(\tilde{X}_t, \tilde{\alpha}_t)$ and $\xi' = \mathcal{L}(\tilde{X}_t, \tilde{\alpha}'_t)$. Indeed, in the framework of Theorem 3.2, if (3.8) holds, we apply it for $x = X_t(\omega)$, $a' = \beta(\omega)$, $y = Y_t(\omega)$, $z = Z_t(\omega)$, $a = \alpha_t(\omega)$ and $\alpha' = \beta$ s.t. $(\tilde{X}, \tilde{Y}, \tilde{Z}, \tilde{\alpha}, \tilde{\beta})$ is a copy of (X, Y, Z, α, β) . Passing to expectation and using (3.6), we get $\mathbb{E} [H(X_t, \beta, \mathcal{L}(X_t, \beta), Y_t, Z_t)] \geq \mathbb{E} [H(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t), Y_t, Z_t)]$, so

$$\alpha_t = \operatorname{argmin} \left\{ \mathbb{E} [H(X_t, \beta, \mathcal{L}(X_t, \beta), Y_t, Z_t)] : \beta \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; A) \right\}.$$

3.2 A Sufficient Condition

Guided by the necessary condition proven above, we derive a sufficient condition for optimality in the same spirit, though under stronger convexity assumptions. For a given pair $(\tilde{X}, \tilde{\alpha})$, these conditions read as

$$g(x', \mu') \geq g(x, \mu) + \partial_x g(x, \mu) \cdot (x' - x) + \tilde{\mathbb{E}} \left[\partial_\mu g(x, \mu)(\tilde{X}) \cdot (\tilde{X}' - \tilde{X}) \right], \quad (3.9)$$

and

$$\begin{aligned} H(x', a', \xi', y, z) &\geq H(x, a, \xi, y, z) + \partial_x H(x, a, \xi, y, z) \cdot (x' - x) + \partial_\alpha H(x, a, \xi, y, z) \cdot (a' - a) \\ &\quad + \tilde{\mathbb{E}} \left[\partial_\mu H(x, a, \xi, y, z)(\tilde{X}, \tilde{\alpha}) \cdot (\tilde{X}' - \tilde{X}) + \partial_\nu H(x, a, \xi, y, z)(\tilde{X}, \tilde{\alpha}) \cdot (\tilde{\alpha}' - \tilde{\alpha}) \right], \end{aligned} \quad (3.10)$$

for all $x, x' \in \mathbb{R}^d$, $a, a' \in A$, $y \in \mathbb{R}^d$, $z \in \mathbb{R}^{d \times m}$, and any \tilde{X}' (resp. $\tilde{\alpha}'$) copy of a process in $\mathbb{H}^{2,d}$ (resp. of an admissible control) on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, and where $\mu = \mathcal{L}(\tilde{X})$, $\mu' = \mathcal{L}(\tilde{X}')$, $\xi = \mathcal{L}(\tilde{X}, \tilde{\alpha})$ and $\xi' = \mathcal{L}(\tilde{X}', \tilde{\alpha}')$; see [14, Ch. 6].

Theorem 3.5. *Under Assumptions (I)-(II), let $\alpha = (\alpha_t)_{0 \leq t \leq T} \in \mathbb{A}$ be an admissible control, $\mathbf{X} = (X_t)_{0 \leq t \leq T}$ the corresponding controlled state process, and $(\mathbf{Y}, \mathbf{Z}) = (Y_t, Z_t)_{0 \leq t \leq T}$ the corresponding adjoint processes satisfying (3.2). Let us assume that:*

- (i) g is convex in the sense of (3.9);
- (ii) H is convex in the sense of (3.10).

Then, if (3.5) holds, α is an optimal control, i.e. $J(\alpha) = \inf_{\alpha' \in \mathbb{A}} J(\alpha')$.

As before, we use the notation $\theta_t = (X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t))$ throughout the proof.

Proof. We follow the steps of the classical proofs; see for example [14, Theorem 6.16] for the case of the control of standard McKean-Vlasov SDEs. Let $(\tilde{X}, \tilde{\alpha})$ be a copy of (X, α) on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, and let $\alpha' \in \mathbb{A}$ be any admissible control, with $\mathbf{X}' = \mathbf{X}^{\alpha'}$ the corresponding controlled state. By definition of the objective function in (2.1) and of the Hamiltonian of the control problem in (3.1), we have:

$$\begin{aligned} J(\alpha) - J(\alpha') &= \mathbb{E}[g(X_T, \mathcal{L}(X_T)) - g(X'_T, \mathcal{L}(X'_T))] + \mathbb{E} \int_0^T [f(\theta_t) - f(\theta'_t)] dt \\ &= \mathbb{E}[g(X_T, \mathcal{L}(X_T)) - g(X'_T, \mathcal{L}(X'_T))] + \mathbb{E} \int_0^T [H(\theta_t, Y_t, Z_t) - H(\theta'_t, Y_t, Z_t)] dt \\ &\quad - \mathbb{E} \int_0^T \{ [b(\theta_t) - b(\theta'_t)] \cdot Y_t + [\sigma(\theta_t) - \sigma(\theta'_t)] \cdot Z_t \} dt, \end{aligned} \quad (3.11)$$

with $\theta'_t = (X'_t, \alpha'_t, \mathcal{L}(X'_t, \alpha'_t))$. Being g convex, we have:

$$\begin{aligned} &\mathbb{E}[g(X_T, \mathcal{L}(X_T)) - g(X'_T, \mathcal{L}(X'_T))] \\ &\leq \mathbb{E}[\partial_x g(X_T, \mathcal{L}(X_T)) \cdot (X_T - X'_T) + \tilde{\mathbb{E}}[\partial_\mu g(X_T, \mathcal{L}(X_T))(\tilde{X}_T) \cdot (\tilde{X}_T - \tilde{X}'_T)]] \\ &= \mathbb{E}[(\partial_x g(X_T, \mathcal{L}(X_T)) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T, \mathcal{L}(X_T))(\tilde{X}_T)]) \cdot (X_T - X'_T)] \\ &= \mathbb{E}[(X_T - X'_T) \cdot Y_T], \end{aligned} \quad (3.12)$$

where we used Fubini and the fact that the ‘tilde random variables’ are independent copies of the ‘non-tilde’ ones. Using integration by parts and the fact that $\mathbf{Y} = (Y_t)_{0 \leq t \leq T}$ solves the adjoint equation (3.2), we get:

$$\begin{aligned} \mathbb{E}[(X_T - X'_T) \cdot Y_T] &= \mathbb{E} \left[\int_0^T (X_t - X'_t) \cdot dY_t + \int_0^T Y_t \cdot d[X_t - X'_t] + \int_0^T [\sigma(\theta_t) - \sigma(\theta'_t)] \cdot Z_t dt \right] \\ &= -\mathbb{E} \int_0^T [\partial_x H(\theta_t, Y_t, Z_t) \cdot (X_t - X'_t) + \tilde{\mathbb{E}}[\partial_\mu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] \cdot (X_t - X'_t)] dt \\ &\quad + \mathbb{E} \int_0^T [[b(\theta_t) - b(\theta'_t)] \cdot Y_t + [\sigma(\theta_t) - \sigma(\theta'_t)] \cdot Z_t] dt. \end{aligned} \quad (3.13)$$

Again by Fubini’s theorem, we get:

$$\mathbb{E} \int_0^T \tilde{\mathbb{E}}[\partial_\mu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] \cdot (X_t - X'_t) dt = \mathbb{E} \int_0^T \tilde{\mathbb{E}}[\partial_\mu H(\theta_t, Y_t, Z_t)(\tilde{X}_t, \tilde{\alpha}_t) \cdot (\tilde{X}_t - \tilde{X}'_t)] dt.$$

Together with (3.11), (3.12), and (3.13), this gives:

$$\begin{aligned} J(\alpha) - J(\alpha') &\leq \mathbb{E} \int_0^T [H(\theta_t, Y_t, Z_t) - H(\theta'_t, Y_t, Z_t)] dt \\ &\quad - \mathbb{E} \int_0^T \left[\partial_x H(\theta_t, Y_t, Z_t) \cdot (X_t - X'_t) + \tilde{\mathbb{E}}[\partial_\mu H(\theta_t, Y_t, Z_t)(\tilde{X}_t, \tilde{\alpha}_t) \cdot (\tilde{X}_t - \tilde{X}'_t)] \right] dt \\ &\leq \mathbb{E} \int_0^T \left[\partial_\alpha H(\theta_t, Y_t, Z_t) \cdot (\alpha_t - \alpha'_t) + \tilde{\mathbb{E}}[\partial_\nu H(\theta_t, Y_t, Z_t)(\tilde{X}_t, \tilde{\alpha}_t) \cdot (\tilde{\alpha}_t - \tilde{\alpha}'_t)] \right] dt \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} \int_0^T \left(\partial_\alpha H(\theta_t, Y_t, Z_t) + \tilde{\mathbb{E}}[\partial_\nu H(\tilde{\theta}_t, \tilde{Y}_t, \tilde{Z}_t)(X_t, \alpha_t)] \right) \cdot (\alpha_t - \alpha'_t) dt \\
&\leq 0,
\end{aligned}$$

because of the convexity of H , Fubini's theorem, and (3.5), showing that α is optimal. \square

4 Examples

In this section, we consider models for which the solution strategy suggested by the stochastic maximum principle proved in the previous section can be pushed further. In fact, in Sections 4.2 and 4.3, we are able to obtain explicit solutions.

4.1 The Case of Scalar Interactions

In this subsection, we state explicitly what the above forms of the Pontryagin stochastic maximum principle become in the case of scalar interactions. This is a case of particular interest because it does not need the full generality of the differential calculus on Wasserstein spaces, and can be dealt with by using standard calculus. An example of scalar interactions will be studied and explicitly solved in the next subsection; see also [23] for another application of scalar interactions.

Assume drift and cost functions to be of the form

$$b(x, \alpha, \xi) = b_0(x, \alpha, \int \varphi d\xi), \quad f(x, \alpha, \xi) = f_0(x, \alpha, \int \psi d\xi), \quad g(x, \mu) = g_0(x, \int \phi d\mu),$$

for some functions b_0, f_0 on $\mathbb{R}^d \times A \times \mathbb{R}$, g_0 on $\mathbb{R}^d \times \mathbb{R}$, φ, ψ on $\mathbb{R}^d \times A$, and ϕ on \mathbb{R}^d . In order to simplify the notation, we shall assume that the volatility is independent of the control, and actually we take $\sigma \equiv I_d$. Under these circumstances, the adjoint equation becomes:

$$\begin{aligned}
dY_t = & - \left(\partial_x b_0(X_t, \alpha_t, \mathbb{E}[\varphi(X_t, \alpha_t)]) Y_t + \partial_x f_0(X_t, \alpha_t, \mathbb{E}[\psi(X_t, \alpha_t)]) \right. \\
& + \tilde{\mathbb{E}}[\tilde{Y}_t \cdot \partial_\zeta b_0(\tilde{X}_t, \tilde{\alpha}_t, \mathbb{E}[\varphi(X_t, \alpha_t)])] \partial_x \varphi(X_t, \alpha_t) \\
& \left. + \tilde{\mathbb{E}}[\partial_\zeta f_0(\tilde{X}_t, \tilde{\alpha}_t, \mathbb{E}[\psi(X_t, \alpha_t)])] \partial_x \psi(X_t, \alpha_t) \right) dt + Z_t dW_t,
\end{aligned}$$

with terminal condition $Y_T = \partial_x g_0(X_T, \mathbb{E}[\phi(X_T)]) + \tilde{\mathbb{E}}[\partial_\zeta g_0(\tilde{X}_T, \mathbb{E}[\phi(X_T)])] \partial_x \phi(X_T)$. Accordingly, the necessary condition (3.7) for optimality will be satisfied when

$$\begin{aligned}
0 = & \partial_\alpha b_0(X_t, \alpha_t, \mathbb{E}[\varphi(X_t, \alpha_t)]) \cdot Y_t + \partial_\alpha f_0(X_t, \alpha_t, \mathbb{E}[\psi(X_t, \alpha_t)]) \\
& + \tilde{\mathbb{E}}[\tilde{Y}_t \cdot \partial_\zeta b_0(\tilde{X}_t, \tilde{\alpha}_t, \mathbb{E}[\varphi(X_t, \alpha_t)])] \partial_\alpha \varphi(X_t, \alpha_t) + \tilde{\mathbb{E}}[\partial_\zeta f_0(\tilde{X}_t, \tilde{\alpha}_t, \mathbb{E}[\psi(X_t, \alpha_t)])] \partial_\alpha \psi(X_t, \alpha_t).
\end{aligned} \tag{4.1}$$

4.2 Optimal liquidation with market impact

In this section we explicitly solve an example that lies outside the classical LQ framework, in the sense that convexity fails. This is inspired by an optimal liquidation problem with price impact, but here it is more of mathematical interest than financial one.

Consider a market where a group of investors, indexed by i , has large positions q_0^i on the same asset S . Each investor wants to liquidate her position by a fixed time $T > 0$, and

controls her trading speed α_t^i through time. Her state is then described by two variables: her inventory Q_t^i , that starts at q_0^i and changes according to α_t^i , and her wealth X_t^i , which is assumed to start at zero for all traders. Investors' speed of trading affects prices in two ways. On the one hand, it generates a permanent market impact, as the dynamics of S are assumed to linearly depend on the average trading speed of all investors. On the other hand, it produces a temporary impact, that only affects traders' own wealth process (as fees or liquidation cost), and which is assumed to be linear in their respective rate of trading. The optimality criterion is the minimization of the cost, which is composed of three factors: the wealth at time T , the final value of the inventory penalized by a terminal market impact, and a running penalty which is assumed quadratic in the inventory. The optimal trades will be a result of the trade-off between trading slowly to reduce the market impact (or execution/liquidity cost), and trading fast to reduce the risk of future uncertainty in prices; see e.g. [2, 16, 17, 12, 6].

Here we think of a continuum of investors. The initial inventories are distributed according to a measure m_0 on \mathbb{R} . We formulate the problem for a representative agent, in the case of cooperative equilibria. The inventory process then evolves as

$$dQ_t = \alpha_t dt, \quad Q_0 \sim m_0, \quad (4.2)$$

while the wealth process is given by

$$dX_t = -\alpha_t(S_t + k\alpha_t)dt, \quad X_0 = 0,$$

where $k\alpha_t$ measures the temporary market impact. The price process is modelled by

$$dS_t = \lambda \mathbb{E}[\alpha_t]dt + \sigma dW_t, \quad S_0 = s_0,$$

where $\mathbb{E}[\alpha_t]$ represents the average trading speed, hence $\lambda \mathbb{E}[\alpha_t]$ stands for the permanent market impact to which all agents contribute (naturally $\lambda \geq 0$). The cost to be minimized is given by

$$\mathbb{E} \left[-X_T - Q_T(S_T - AQ_T) + \phi \int_0^T Q_t^2 dt \right],$$

where X_T is the terminal profit due to trading in $[0, T]$, $Q_T(S_T - AQ_T)$ is the liquidation value of the remaining quantity at terminal time (with a liquidation/execution penalization), and ϕ is an ‘‘urgency’’ parameter on the running cost (the higher ϕ is, the higher is the liquidation speed at the beginning of the trading period). Using the dynamics of X , this can be rewritten as

$$\mathbb{E} \left[\int_0^T (\alpha_t S_t + k\alpha_t^2 + \phi Q_t^2) dt - Q_T(S_T - AQ_T) \right].$$

This example falls into the framework described in Section 2. We have a 2-dimensional state process (S, Q) , a 1-dimensional Wiener process W , and the control process is the trading speed α . The Hamiltonian of the system is

$$H(x_1, x_2, a, \xi, y_1, y_2) = \lambda \bar{\xi}_2 y_1 + ay_2 + \phi x_2^2 + ax_1 + ka^2,$$

where $\bar{\xi}_2 = \int v \xi(du, dv)$, and the first order condition (4.1) reads as

$$Y_t^2 + S_t + 2k\alpha_t + \lambda \mathbb{E}[Y_t^1] = 0, \quad (4.3)$$

with adjoint equations

$$dY_t^1 = -\alpha_t dt + Z_t^1 dW_t, \quad Y_T^1 = -Q_T \quad (4.4)$$

$$dY_t^2 = -2\phi Q_t dt + Z_t^2 dW_t, \quad Y_T^2 = -S_T + 2AQ_T. \quad (4.5)$$

Remark 4.1. Here the terminal cost function g reads as

$$g(x_1, x_2) = -x_1 x_2 + Ax_2^2,$$

which does not satisfy the convexity condition (3.9). However, an inspection of the proof of Theorem 3.5 reveals that this assumption was only used in order to obtain the inequality in (3.12). We are now going to show that such inequality holds in the present setting when $A \geq \lambda$ (which is satisfied for typical values of the parameters, see [17, 12]), thus guaranteeing that the first order condition (4.3) is not only necessary but also sufficient for the optimality of α . For this purpose, let $\alpha' \in \mathbb{A}$ be any admissible control, and (S', Q') the corresponding controlled state. Then

$$\begin{aligned} & \mathbb{E}[g(S_T, Q_T) - g(S'_T, Q'_T)] - \mathbb{E}[(S_T - S'_T)Y_T^1 + (Q_T - Q'_T)Y_T^2] \\ &= \lambda \left(\mathbb{E} \left[\int_0^T \alpha'_t dt - \int_0^T \alpha_t dt \right] \right)^2 - A \mathbb{E} \left[\left(\int_0^T \alpha_t dt - \int_0^T \alpha'_t dt \right)^2 \right] \\ &\leq (\lambda - A) \mathbb{E} \left[\left(\int_0^T \alpha_t dt - \int_0^T \alpha'_t dt \right)^2 \right], \end{aligned}$$

which is non-positive for $A \geq \lambda$.

An inspection of (4.4) suggests that we have $Z_t^1 = 0$, and $Y_t^1 = -Q_0 - \int_0^t \alpha_s ds = -Q_t$; Y_t^2 will be determined later. Substituting in (4.3), we have

$$Y_0^2 - 2\phi \int_0^t Q_s ds + \int_0^t Z_s^2 dW_s + s_0 + \lambda \int_0^t \mathbb{E}[\alpha_s] ds + \sigma W_t + 2k\alpha_t - \lambda(\mathbb{E}[Q_0] + \int_0^t \mathbb{E}[\alpha_s] ds) = 0,$$

that is,

$$\alpha_t = \frac{\lambda \mathbb{E}[Q_0] - Y_0^2 - s_0}{2k} + \frac{\phi}{k} \int_0^t Q_s ds - \frac{1}{2k} \int_0^t (Z_s^2 + \sigma) dW_s. \quad (4.6)$$

We now show that $Q \equiv Q^0$ and $\alpha \equiv \alpha^0$, where

$$Q_t^0 := \mathbb{E}[Q_t | Q_0], \quad \alpha_t^0 := \mathbb{E}[\alpha_t | Q_0].$$

By taking conditional expectation in (4.2) and (4.6), we get

$$Q_t^0 = Q_0 + \int_0^t \alpha_s^0 ds, \quad \alpha_t^0 = \alpha_0 + \frac{\phi}{k} \int_0^t Q_s^0 ds. \quad (4.7)$$

Setting $F(t) := Q_t^0$, we note that $F'(t) = \alpha_t^0$, and $F''(t) = \frac{\phi}{k} F(t)$. Together with the initial conditions $F(0) = Q_0$ and $F'(0) = \alpha_0$, this gives

$$F(t) = \left(\frac{Q_0}{2} - \frac{\alpha_0}{2r} \right) e^{-rt} + \left(\frac{Q_0}{2} + \frac{\alpha_0}{2r} \right) e^{rt}, \quad (4.8)$$

where $r = \sqrt{\phi/k}$. Now, by taking conditional expectation in equation (4.5), and substituting in (4.7), we obtain

$$\alpha_T^0 = \frac{\lambda \mathbb{E}[Q_0] - 2AQ_0}{2k} + \frac{\lambda}{2k} \int_0^T \mathbb{E}[\alpha_t] dt - \frac{A}{k} \int_0^T \alpha_t^0 dt$$

$$\begin{aligned}
&= \frac{\lambda \mathbb{E}[Q_0] - 2A Q_0}{2k} + \frac{\lambda}{2k} (\mathbb{E}[Q_T] - \mathbb{E}[Q_0]) - \frac{A}{k} (Q_T^0 - Q_0) \\
&= \frac{\lambda}{2k} \mathbb{E}[Q_T] - \frac{A}{k} Q_T^0,
\end{aligned} \tag{4.9}$$

that is, $F'(T) = \frac{\lambda}{2k} \mathbb{E}[F(T)] - \frac{A}{k} F(T)$. Imposing this condition, and using (4.8), we obtain:

$$\alpha_0 = Q_0 r \frac{d_1 e^{-rT} - d_2 e^{rT}}{d_1 e^{-rT} + d_2 e^{rT}} + \frac{\mathbb{E}[Q_0] 4\lambda\phi}{(d_1 e^{-rT} + d_2 e^{rT})(c_1 e^{-rT} + c_2 e^{rT})}, \tag{4.10}$$

where $d_1 = \sqrt{\phi k} - A$, $d_2 = \sqrt{\phi k} + A$, $c_1 = 2d_1 + \lambda$, $c_2 = 2d_2 - \lambda$. From (4.6), we also have an explicit expression for $Y_0^2 = \lambda \mathbb{E}[Q_0] - s_0 - 2k\alpha_0$.

Now we use the ansatz: $Z^2 \equiv -\sigma$, and show that the process

$$Y_t^2 = Y_0^2 - 2\phi \int_0^t Q_s ds - \sigma W_t \tag{4.11}$$

does satisfy the equation and terminal condition in (4.5). Only the latter needs to be shown. First note that, with this ansatz, from (4.6) and (4.2) we have:

$$\alpha_t = \alpha_0 + \frac{\phi}{k} \int_0^t Q_s ds, \quad Q_t = Q_0 + \alpha_0 t + \frac{\phi}{k} \int_0^t \int_0^s Q_u du ds,$$

thus both processes α and Q are $\sigma(Q_0)$ -measurable, that is,

$$Q_t = \mathbb{E}[Q_t | Q_0] = Q_t^0 = F(t), \quad \text{and} \quad \alpha_t = \mathbb{E}[\alpha_t | Q_0] = \alpha_t^0 = F'(t). \tag{4.12}$$

We now check that Y^2 satisfies the terminal condition in (4.5). By (4.12), (4.11) implies

$$Y_T^2 = \lambda \mathbb{E}[Q_0] - s_0 - 2k\alpha_0 - 2\phi \int_0^T Q_t^0 dt - \sigma W_T.$$

On the other hand, by (4.12), (4.9) and (4.7),

$$\begin{aligned}
-S_T + 2A Q_T &= -s_0 - \lambda(\mathbb{E}[Q_T] - \mathbb{E}[Q_0]) - \sigma W_T + 2A Q_T^0 = -s_0 + \lambda \mathbb{E}[Q_0] - 2k\alpha_T^0 - \sigma W_T \\
&= -s_0 + \lambda \mathbb{E}[Q_0] - 2k\alpha_0 - 2\phi \int_0^T Q_t^0 dt - \sigma W_T,
\end{aligned}$$

which yields $Y_T^2 = -S_T + 2A Q_T$, as wanted. This shows that the process Z^2 in the ansatz, together with Y^2 defined above, do satisfy (4.5). We have seen that this gives $Q_t = F(t)$ and $\alpha_t = F'(t)$, by (4.12), thus from (4.8) we have

$$Q_t = \left(\frac{Q_0}{2} - \frac{\alpha_0}{2r}\right) e^{-rt} + \left(\frac{Q_0}{2} + \frac{\alpha_0}{2r}\right) e^{rt}, \quad \alpha_t = \left(-\frac{Q_0 r}{2} + \frac{\alpha_0}{2}\right) e^{-rt} + \left(\frac{Q_0 r}{2} + \frac{\alpha_0}{2}\right) e^{rt}.$$

By (4.10), this gives

$$\begin{aligned}
Q_t &= Q_0 \frac{d_1 e^{-r(T-t)} + d_2 e^{r(T-t)}}{d_1 e^{-rT} + d_2 e^{rT}} + \mathbb{E}[Q_0] \frac{2\lambda \sqrt{\phi k} (-e^{-rt} + e^{rt})}{(d_1 e^{-rT} + d_2 e^{rT})(c_1 e^{-rT} + c_2 e^{rT})}, \\
\alpha_t &= Q_0 r \frac{d_1 e^{-r(T-t)} - d_2 e^{r(T-t)}}{d_1 e^{-rT} + d_2 e^{rT}} + \mathbb{E}[Q_0] \frac{2\lambda \phi (e^{-rt} + e^{rt})}{(d_1 e^{-rT} + d_2 e^{rT})(c_1 e^{-rT} + c_2 e^{rT})}.
\end{aligned}$$

4.3 The Linear-Quadratic Case

In this subsection, we use the sufficient condition derived above to solve a simple Linear Quadratic (LQ) model. Via different methods, such models have been already studied in the literature; see e.g. [35, 24, 6, 33]. For the sake of simplicity, we give the details of the computations in the scalar case $m = d = k = 1$ and with $A = \mathbb{R}$. Also, as before, we assume that the volatility is not controlled, and in fact that it is identically equal to 1. In such LQ model, the drift is of the form:

$$b(x, \alpha, \xi) = b_1 x + b_2 \alpha + \bar{b}_1 \bar{x} + \bar{b}_2 \bar{\alpha},$$

for some constants $b_1, b_2, \bar{b}_1, \bar{b}_2$, where we denote by \bar{x} and $\bar{\alpha}$ the means of the state and the control, in the sense that $\bar{x} = \int \int x \xi(dx, d\alpha)$, and $\bar{\alpha} = \int \int \alpha \xi(dx, d\alpha)$. As for the cost functions, we assume that

$$f(x, \alpha, \xi) = \frac{1}{2} [qx^2 + \bar{q}(x - s\bar{x})^2 + r\alpha^2 + \bar{r}(\alpha - \bar{s}\bar{\alpha})^2], \quad g(x, \mu) = \frac{1}{2}\gamma x^2 + \frac{\bar{\gamma}}{2}(x - \rho\bar{x})^2,$$

for some constants $q, \bar{q}, r, \bar{r}, s, \bar{s}, \gamma, \delta, \rho$ satisfying $\bar{q}, \bar{r}, \bar{\gamma} \geq 0$ and $q, r, \gamma > 0$. Under these conditions, the Hamiltonian reads:

$$H(x, \alpha, \xi, y) = (b_1 x + b_2 \alpha + \bar{b}_1 \bar{x} + \bar{b}_2 \bar{\alpha})y + \frac{1}{2} [qx^2 + \bar{q}(x - s\bar{x})^2 + r\alpha^2 + \bar{r}(\alpha - \bar{s}\bar{\alpha})^2]. \quad (4.13)$$

Accordingly, the adjoint equation reads as:

$$dY_t = -\left(b_1 Y_t + (q + \bar{q})X_t + \bar{b}_1 \mathbb{E}[Y_t] + s\bar{q}(s - 2)\mathbb{E}[X_t]\right)dt + Z_t dW_t. \quad (4.14)$$

In the present situation, conditions (i) and (ii) of Theorem 3.5 hold, and condition (3.7) of the Pontryagin stochastic maximum principle holds if

$$b_2 Y_t + \bar{b}_2 \mathbb{E}[Y_t] + (r + \bar{r})\alpha_t + \bar{r}\bar{s}(\bar{s} - 2)\mathbb{E}[\alpha_t] = 0. \quad (4.15)$$

Taking expectations, we obtain

$$\mathbb{E}[\alpha_t] = -\frac{b_2 + \bar{b}_2}{r + \bar{r}(\bar{s} - 1)^2} \mathbb{E}[Y_t]. \quad (4.16)$$

Plugging this expression into (4.15), we get:

$$\alpha_t = -\frac{1}{r + \bar{r}} \left(b_2 Y_t + \left(\bar{b}_2 - \frac{\bar{r}\bar{s}(\bar{s} - 2)(b_2 + \bar{b}_2)}{r + \bar{r}(\bar{s} - 1)^2} \right) \bar{Y}_t \right). \quad (4.17)$$

We can rewrite (4.17) and (4.16) as

$$\alpha_t = aY_t + b\mathbb{E}[Y_t] \quad \text{and} \quad \mathbb{E}[\alpha_t] = c\mathbb{E}[Y_t], \quad (4.18)$$

with

$$a = -\frac{b_2}{r + \bar{r}}, \quad b = -\frac{1}{r + \bar{r}} \left(\bar{b}_2 - \frac{\bar{r}\bar{s}(\bar{s} - 2)(b_2 + \bar{b}_2)}{r + \bar{r}(\bar{s} - 1)^2} \right), \quad \text{and} \quad c = -\frac{b_2 + \bar{b}_2}{r + \bar{r}(\bar{s} - 1)^2}. \quad (4.19)$$

With this notation, the solution of the mean field optimal control of the McKean-Vlasov SDE (2.2) reduces to the solution of the following forward-backward stochastic differential equation (FBSDE) of McKean-Vlasov type:

$$\begin{cases} dX_t = (b_1 X_t + \bar{b}_1 \mathbb{E}[X_t] + (ab_2 Y_t + (bb_2 + c\bar{b}_2)\mathbb{E}[Y_t]))dt + dW_t \\ dY_t = -\left(b_1 Y_t + (q + \bar{q})X_t + \bar{b}_1 \mathbb{E}[Y_t] + s\bar{q}(s - 2)\mathbb{E}[X_t]\right)dt + Z_t dW_t, \end{cases} \quad (4.20)$$

with terminal condition $Y_T = (\gamma + \bar{\gamma})X_T + \bar{\gamma}\rho(\rho - 2)\mathbb{E}[X_T]$. We solve this system in the usual way. First, we compute the means $\bar{x}_t = \mathbb{E}[X_t]$ and $\bar{y}_t = \mathbb{E}[Y_t]$. Taking expectations in (4.20), we obtain:

$$\begin{cases} d\bar{x}_t = \left((b_1 + \bar{b}_1)\bar{x}_t + (ab_2 + bb_2 + c\bar{b}_2)\bar{y}_t \right) dt \\ d\bar{y}_t = -\left((b_1 + \bar{b}_1)\bar{y}_t + (q + \bar{q} + s\bar{q}(s-2))\bar{x}_t \right) dt, \end{cases} \quad (4.21)$$

with terminal condition $\bar{y}_T = (\gamma + \bar{\gamma} + \bar{\gamma}\rho(\rho - 2))\bar{x}_T$. The linear system (4.21) can be solved explicitly. For instance, if we denote

$$\Delta := \sqrt{(b_1 + \bar{b}_1)^2 - (ab_2 + bb_2 + c\bar{b}_2)(q + \bar{q} + s\bar{q}(s-2))},$$

and assume that the argument of the square root is strictly positive, one can solve (4.21) via the theory of linear ODE systems in the case of real eigenvalues. We then obtain that

$$\bar{x}_t = -\frac{(b_1 + \bar{b}_1)^2 - \Delta^2}{2(q + \bar{q} + s\bar{q}(s-2))\Delta} \left\{ e^{-\Delta t} \left(y_0 + \frac{(q + \bar{q} + s\bar{q}(s-2))x_0}{b_1 + \bar{b}_1 + \Delta} \right) - e^{\Delta t} \left(y_0 + \frac{(q + \bar{q} + s\bar{q}(s-2))x_0}{b_1 + \bar{b}_1 - \Delta} \right) \right\}$$

together with

$$\bar{y}_t = -\frac{(b_1 + \bar{b}_1)^2 - \Delta^2}{2(q + \bar{q} + s\bar{q}(s-2))\Delta} \left\{ \frac{-(q + \bar{q} + s\bar{q}(s-2))e^{-\Delta t}}{b_1 + \bar{b}_1 - \Delta} \left(y_0 + \frac{(q + \bar{q} + s\bar{q}(s-2))x_0}{b_1 + \bar{b}_1 + \Delta} \right) + \frac{(q + \bar{q} + s\bar{q}(s-2))e^{\Delta t}}{b_1 + \bar{b}_1 + \Delta} \left(y_0 + \frac{(q + \bar{q} + s\bar{q}(s-2))x_0}{b_1 + \bar{b}_1 - \Delta} \right) \right\}$$

solve (4.21) for any y_0 , and choosing y_0 appropriately one can guarantee that $\bar{y}_T = (\gamma + \bar{\gamma} + \bar{\gamma}\rho(\rho - 2))\bar{x}_T$. This expression for (\bar{x}_t, \bar{y}_t) can be plugged into (4.20) in lieu of $(\mathbb{E}[X_t], \mathbb{E}[Y_t])$, reducing the latter to a standard affine FBSDE. We then make the ansatz $Y_t = \eta_t X_t + \chi_t$ for two deterministic functions $t \mapsto \eta_t$ and $t \mapsto \chi_t$, which is compatible with the terminal condition. Computing the Itô differentials of Y_t from the ansatz and from the system (4.20), and identifying the terms in the drift multiplying the unknown X_t , we find that η_t should be a solution of the scalar Riccati equation

$$\eta_t = -\frac{1}{2b_1}(q + \bar{q} + \eta'_t + ab_2\eta_t^2).$$

The latter is easily solved, and since necessarily $\bar{y}_t = \eta_t \bar{x}_t + \chi_t$, then χ_t can also be explicitly obtained. By Theorem 3.5, the control α obtained in this way is optimal. Notice that it takes the form

$$\alpha_t = a\eta_t X_t + a\chi_t + b\bar{x}_t,$$

with a and b given in (4.19).

Remark 4.2. In classical control of Mean Field type, the pointwise minimization of the Hamiltonian with respect to the control is a necessary optimality condition. Let us illustrate with the LQ example how this need not be the case in our extended framework. If we impose pointwise minimization of (4.13) with respect to α , we get $b_2 Y_t + r\alpha_t + \bar{r}(\alpha_t - \bar{s}\bar{\alpha}_t) = 0$. Integrating it, we obtain $b_2 \mathbb{E}[Y_t] + (r + \bar{r} - \bar{r}\bar{s})\bar{\alpha}_t = 0$. On the other hand, the necessary condition (3.5) implies (4.15), so we have $\bar{b}_2 \mathbb{E}[Y_t] + \bar{r}\bar{s}(\bar{s} - 1)\bar{\alpha}_t = 0$. The right choice of parameters leads to a contradiction between this and the previous equation.

5 Variational Perspective in the Weak Formulation

The goal of this section is to analyse the extended Mean Field control problem from a purely variational perspective, that is, by considering its formulation on path space. Given the intrinsic nature of Mean Field problems, it is natural to express them in terms of laws rather than controls. The main reason for exploring this point of view is that of creating a bridge with the optimal transport theory. This paves the way to the use of different sets of tools, as for example the numerical methods that are fast developing in transport theory. We start by introducing, in Section 5.1, a weak formulation of the extended Mean Field control problem, especially well-suited for variational analysis. In such formulation, the probability space is not specified a priori. We remark that a weak formulation of the Mean Field control problem has been considered in [14, Sec. 6.6] and in [25], the latter rigorously proving convergence of large systems of interacting control problems to the corresponding Mean Field control problem. However, in these works there is no non-linear dependence on the law of the control; cf. our problem (5.1) below.

We proceed in Section 5.2 to obtain what we call a *martingale optimality condition*. Such a condition can serve as a verification tool, in order to evaluate whether a given control can be optimal. It is therefore the weak-formulation analogue of the necessary Pontryagin maximum principle. This forms a bridge between the previous sections of this work, and the ensuing ones. Whenever the Pontryagin maximum principle can be used (or the martingale optimality condition in the weak formulation), it is a powerful tool to identify optimal controls and the trajectories of the state at the optimum. However, it does not say much about the optimal value of the problem. In fact, at the optimum, the adjoint process gives formally the value of the gradient of the value function when computed along the optimal trajectories. In order to study the value function of the control problem (in a situation in which PDE techniques are highly non-trivial) we recast in Section 5.3 our weak formulation in transport-theoretic terms.

Numerical optimal transport has spectacularly grown in strength over the last few years; see e.g. [19, 7, 29] and the references therein. Our connection between transport and Mean Field control is meant to lay ground for efficient numerical methods in the future. In Section 5.4 we provide, at a theoretical level, a first discretization scheme of this kind. To be specific, the optimal transport problem we obtain in the discretization has an additional *causality* constraint (see e.g. [26, 1, 4, 5]); the numerical analysis of such problems is also having a burst of activity (e.g. [30, 31, 32]).

5.1 The Weak Formulation

We present a *weak formulation* of the extended Mean Field control problem formulated in Section 2, in the sense that the probability space is not specified here. We restrict our attention to the case where the state dynamics have uncontrolled volatility, actually assuming $\sigma \equiv I_d$, $m = d$, that the drift does not depend on the law of the control, and that the initial condition X_0 is a constant x_0 . We thus consider the minimization problem

$$\inf_{\mathbb{P}, \alpha} \mathbb{E}^{\mathbb{P}} \left[\int_0^T f(X_t, \alpha_t, \mathcal{L}_{\mathbb{P}}(X_t, \alpha_t)) dt + g(X_T, \mathcal{L}_{\mathbb{P}}(X_T)) \right] \quad (5.1)$$

subject to $dX_t = b(X_t, \alpha_t, \mathcal{L}_{\mathbb{P}}(X_t)) dt + dW_t$, $X_0 = x_0$,

where the infimum is taken over filtered probability spaces $(\Omega, \mathbb{F}, \mathbb{P})$ supporting some d -dimensional Wiener process \mathbf{W} , and over control processes α which are progressively measurable on $(\Omega, \mathbb{F}, \mathbb{P})$ and \mathbb{R}^k -valued. We use $\mathcal{L}_{\mathbb{P}}$ to denote the law of the given random element under \mathbb{P} . Again, we choose time independent coefficients for simplicity, but all the results would be the same should f and b depend upon t .

We say that $(\Omega, \mathbb{F}, \mathbb{P}, \mathbf{W}, \mathbf{X}, \alpha)$ is a *feasible tuple* if it participates in the above optimization problem yielding a finite cost.

5.2 Martingale Optimality Condition

In this section, we obtain a necessary Pontryagin principle for the weak formulation (5.1). We call this the *martingale optimality condition*. Since our aim is to illustrate the method, we assume only in this part that we are dealing with a drift-control problem

$$b(x, \alpha, \mu) = \alpha, \quad m = d.$$

We start by expressing the objective function of (5.1) in canonical space, as a function of semimartingale laws. We denote by \mathcal{C}_{x_0} the space of \mathbb{R}^d -valued continuous paths started at x_0 , and by \mathbf{S} the canonical process on it. We consider the set of semimartingale laws

$$\tilde{\mathcal{P}} := \{\mu \in \mathcal{P}(\mathcal{C}_{x_0}) : dS_t = \alpha_t^\mu(\mathbf{S})dt + dW_t^\mu \text{ } \mu\text{-a.s.}\}, \quad (5.2)$$

where \mathbf{W}^μ is a μ -Brownian motion and α^μ is a progressively measurable process w.r.t. the canonical filtration, denoted by \mathcal{F} . It is then easy to see that (5.1) is equivalent to

$$\inf_{\mu \in \tilde{\mathcal{P}}} \mathbb{E}^\mu \left[\int_0^T f(S_t, \alpha_t^\mu, \mathcal{L}_\mu(S_t, \alpha_t^\mu))dt + g(S_T, \mu_T) \right]. \quad (5.3)$$

In what follows we consider perturbation of measures in $\tilde{\mathcal{P}}$ via push-forwards along absolutely continuous shifts which preserve the filtration; see the work of Cruzeiro and Lassalle [18] and the references therein. Using push-forwards instead of perturbations directly on the SDE is the main difference between the weak and the strong perspective. The main idea is to find the first order conditions for Problem (5.3) by considering perturbations of the form $\mu^{\epsilon, K} := (Id + \epsilon K)_* \mu$ around a putative optimizer μ . For this matter it is important to identify the Doob-Meyer decomposition of the canonical process under $\mu^{\epsilon, K}$, which forces an assumption on K as we now explain:

Remark 5.1. Let $\mu \in \tilde{\mathcal{P}}$. We say that an adapted process $U : \mathcal{C}_{x_0} \rightarrow \mathcal{C}_{x_0}$ is μ -invertible, if there exists $V : \mathcal{C}_{x_0} \rightarrow \mathcal{C}_{x_0}$ adapted such that $U \circ V = Id_{\mathcal{C}_{x_0}}$ holds $U(\mu)$ -a.s., and $V \circ U = Id_{\mathcal{C}_{x_0}}$ holds μ -a.s.. Now let $K = \int_0^\cdot k_t dt$ be adapted. We say that K preserves the filtration under μ , if for every U which is μ -invertible we also have that $U + K$ is μ -invertible. It follows that the set of those $K = \int_0^\cdot k_t dt$ that preserve the filtration under μ , is a linear space. It also follows that for such K we have $\mu^{\epsilon, K} := (Id + \epsilon K)_* \mu \in \tilde{\mathcal{P}}$, with $\alpha_t^{\mu^{\epsilon, K}}(\mathbf{S} + \epsilon K(\mathbf{S})) = \alpha_t^\mu(\mathbf{S}) + \epsilon k_t(\mathbf{S})$; see [18, Proposition 2.1, Lemma 3.1]. A typical case when the filtration is preserved is when K is a piecewise linear and adapted process, while an example when K does not preserve the filtration is given by Tsirelson's drift; see respectively [18, Proposition 2.4, Remark 2.1.1].

In analogy to [18, Theorem 5.1], we then obtain the following necessary condition for an optimizer in (5.3). We use here the notation $\theta_t^\mu = (S_t, \alpha_t^\mu, \mathcal{L}_\mu(S_t, \alpha_t^\mu))$.

Proposition 5.2. *Let μ be an optimizer for (5.3). Then the process N^μ given by*

$$N_t^\mu := \partial_a f(\theta_t^\mu) + \tilde{\mathbb{E}}[\partial_\nu f(\tilde{\theta}_t^\mu)(S_t, \alpha_t^\mu)] - \int_0^t \left(\partial_x f(\theta_s^\mu) + \tilde{\mathbb{E}}[\partial_\mu f(\tilde{\theta}_s^\mu)(S_s, \alpha_s^\mu)] \right) ds \quad (5.4)$$

is a μ -martingale, with terminal value equal to

$$N_T^\mu = -\partial_x g(S_T, \mu_T) - \tilde{\mathbb{E}}[\partial_\mu g(\tilde{S}_T, \mu_T)(S_T)] - \int_0^T \left(\partial_x f(\theta_s^\mu) + \tilde{\mathbb{E}}[\partial_\mu f(\tilde{\theta}_s^\mu)(S_s, \alpha_s^\mu)] \right) ds. \quad (5.5)$$

Proof. We use the notation $\mu^{\epsilon, K}$ introduced in Remark 5.1, and call $C(\mu)$ the cost function appearing in Problem (5.3). We have $\lim_{\epsilon \rightarrow 0} \frac{C(\mu^{\epsilon, K}) - C(\mu)}{\epsilon} \geq 0$ for all K . Now if K preserves the filtration under μ , then the same is true for $-K$. Therefore $\lim_{\epsilon \rightarrow 0} \frac{C(\mu^{\epsilon, K}) - C(\mu)}{\epsilon} = 0$. To conclude the proof, we use $\alpha_t^{\mu^{\epsilon, K}}(\mathbf{S} + \epsilon K(\mathbf{S})) = \alpha_t^\mu(\mathbf{S}) + \epsilon k_t(\mathbf{S})$ and similar arguments as in [18, Theorem 5.1]. \square

When (5.4)-(5.5) hold, we say that μ satisfies the *martingale optimality condition*. The interest of this condition is that it is a clear stochastic counterpart to the classical Euler-Lagrange condition in calculus of variation, except for the fact that “being equal to zero” is here replaced by “being a martingale”; see [18, 27].

Example 5.3. The martingale optimality condition is the analogue of the Pontryagin principle in the weak formulation. To wit, we verify this in a simple example. Suppose $f(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)) = \frac{1}{2}(\alpha_t - \mathbb{E}[\alpha_t])^2$ and $g(X_T, \mathcal{L}(X_T)) = \frac{1}{2}X_T^2$. The martingale optimality condition then asserts that for an optimizer μ the process $N_t^\mu := \alpha_t^\mu - \mathbb{E}[\alpha_t^\mu]$ is a martingale with $N_T^\mu = -S_T$. On the other hand the Pontryagin FBSDE states that

$$dY_t = Z_t dW_t, \quad Y_T = X_T,$$

as well as $\alpha_t - \mathbb{E}[\alpha_t] + Y_t$, by Remark 3.3. We see the compatibility of the two statements, as well as the equality in law $N_t^\mu = -Y_t$ in this particular case.

Remark 5.4. The above arguments can be adapted to the case when $b(x, \alpha, \mu) = b(x, \alpha)$. This is the case, for example, when b is a C^1 -diffeomorphism and $b(x, \mathbb{R}^k)$ is convex for each x . Indeed, in this case one may re-define the drift in the dynamics of S via $\beta_t^\mu(\mathbf{S}) := b(S_t, \alpha_t^\mu(\mathbf{S}))$, which is associated with the cost $f(S_t, b^{-1}(S_t, \beta_t^\mu(\mathbf{S})), \mathcal{L}_\mu(S_t, b^{-1}(S_t, \beta_t^\mu(\mathbf{S}))))$, where with some abuse of notation $b^{-1}(x, \cdot)$ denotes the inverse of $b(x, \cdot)$. Using this time the notation $\theta_t^\mu = (S_t, \beta_t^\mu, \mathcal{L}_\mu(S_t, \beta_t^\mu))$ one then replaces the r.h.s. of (5.4) with

$$\begin{aligned} & \partial_a f(\theta_t^\mu) \partial_a (b^{-1})(S_t, \beta_t^\mu) + \tilde{\mathbb{E}}[\partial_\nu f(\theta_t^\mu) \partial_a (b^{-1})(\tilde{S}_t, \tilde{\beta}_t)] \\ & - \int_0^t \left(\partial_x f(\theta_s^\mu) + \tilde{\mathbb{E}}[\partial_\mu f(\theta_s^\mu)(\tilde{S}_s, \tilde{\beta}_s) + \partial_\nu f(\theta_s^\mu) \partial_x (b^{-1})(\tilde{S}_s, \tilde{\beta}_s)] \right) ds, \end{aligned} \quad (5.6)$$

and the r.h.s. of (5.5) with

$$\begin{aligned} & -\partial_x g(S_T, \mu_T) - \tilde{\mathbb{E}}[\partial_\mu g(S_T, \mu_T)(\tilde{S}_T)] \\ & - \int_0^T \left(\partial_x f(\theta_s^\mu) + \tilde{\mathbb{E}}[\partial_\mu f(\theta_s^\mu)(\tilde{S}_s, \tilde{\beta}_s) + \partial_\nu f(\theta_s^\mu) \partial_x (b^{-1})(\tilde{S}_s, \tilde{\beta}_s)] \right) ds. \end{aligned} \quad (5.7)$$

5.3 Optimal Transport Reformulation

In this section we formulate a variational transport problem on $\mathcal{C} = \mathcal{C}([0, T]; \mathbb{R}^d)$, the space of \mathbb{R}^d -valued continuous paths, which is equivalent to finding the weak solutions of the extended Mean Field problem (5.1). This variational formulation is a particular type of transport problem under the so-called *causality* constraint; see [26, 1, 4, 5]. Here we recall this concept with respect to the filtrations \mathbb{F}^1 and \mathbb{F}^2 , generated by the first and by the second coordinate process on $\mathcal{C} \times \mathcal{C}$.

Definition 5.5. Given $\zeta_1, \zeta_2 \in \mathcal{P}(\mathcal{C})$, a probability measure $\pi \in \mathcal{P}(\mathcal{C} \times \mathcal{C})$ is called a causal transport plan between ζ_1 and ζ_2 if its marginals are ζ_1 and ζ_2 , and, for any $t \in [0, T]$ and any set $A \in \mathcal{F}_t^2$, the map $\mathcal{C} \ni x \mapsto \pi^x(A)$ is $\tilde{\mathcal{F}}_t^1$ -measurable, where $\pi^x(dy) := \pi(\{x\} \times dy)$ is a regular conditional kernel of π w.r.t. the first coordinate, and $\tilde{\mathbb{F}}^1$ is the completion of \mathbb{F}^1 w.r.t. ζ_1 . The set of causal transport plans between ζ_1 and ζ_2 is denoted by $\Pi_c(\zeta_1, \zeta_2)$.

The only transport plans that contribute to the variational formulation of the problem are those under which the difference of the the coordinate processes on the product space $\mathcal{C} \times \mathcal{C}$ is a.s. absolutely continuous with respect to Lebesgue measure. We denote by $(\omega, \bar{\omega})$ the generic element on $\mathcal{C} \times \mathcal{C}$, and we use $(\overleftarrow{\bar{\omega} - \omega})$ to indicate the density of the process $\bar{\omega} - \omega$ with respect to Lebesgue measure, when it exists, i.e.

$$\bar{\omega}_t - \omega_t = \bar{\omega}_0 - \omega_0 + \int_0^t (\overleftarrow{\bar{\omega} - \omega})_s ds, \quad t \in [0, T].$$

In such case, we write $\bar{\omega} - \omega \ll \mathcal{L}$. Moreover, we set

$$\gamma := \text{Wiener measure on } \mathcal{C} \text{ started at } 0,$$

and $\Pi_c^{\ll}(\gamma, \cdot) := \{\pi \in \mathcal{P}(\mathcal{C} \times \mathcal{C}) : \pi(d\omega \times \mathcal{C}) = \gamma(d\omega), \text{ and } \bar{\omega} - \omega \ll \mathcal{L}, \pi\text{-a.s.}\}$.

We present the connection between extended Mean Field control and causal transport:

Lemma 5.6. *Assume that $b(x, \cdot, \mu)$ is injective, and set*

$$u_t(\omega, \bar{\omega}, \mu) := b^{-1}(\bar{\omega}_t, \cdot, \mu)((\overleftarrow{\bar{\omega} - \omega})_t).$$

Then Problem (5.1) is equivalent to:

$$\inf \mathbb{E}^\pi \left[\int_0^T f(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t^\pi), \mathcal{L}_\pi(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t^\pi))) dt + g(\bar{\omega}_T, \mu_T^\pi) \right], \quad (5.8)$$

where the infimum is taken over transport plans $\pi \in \Pi_c^{\ll}(\gamma, \cdot)$ such that $dt \otimes d\pi$ -a.s. $(\overleftarrow{\bar{\omega} - \omega})_t \in b(\bar{\omega}_t, \mathbb{R}^d, \mu_t^\pi)$, and μ^π denotes the second marginal of π .

Proof. Fix $(\Omega, \mathbb{F}, \mathbb{P}, \mathbf{W}, \mathbf{X}, \alpha)$ feasible tuple for (5.1), if it exists, and note that $\alpha_t = u_t(\mathbf{W}, \mathbf{X}, \mathcal{L}_\mathbb{P}(X_t))$ is $\mathbb{F}^{\mathbf{X}, \mathbf{W}}$ -adapted. Then $\pi := \mathcal{L}_\mathbb{P}(\mathbf{W}, \mathbf{X})$ belongs to $\Pi_c^{\ll}(\gamma, \mathcal{L}_\mathbb{P}(\mathbf{X}))$ and generates the same cost in (5.8). Conversely, given a transport plan π participating in (5.8), the following tuple $(\Omega, \mathbb{F}, \mathbb{P}, \mathbf{W}, \mathbf{X}, \alpha)$ is feasible for (5.1): $\Omega = \mathcal{C} \times \mathcal{C}$, \mathbb{F} canonical filtration on $\mathcal{C} \times \mathcal{C}$, $\mathbb{P} = \pi$, $\mathbf{W} = \omega$, $\mathbf{X} = \bar{\omega}$, and $\alpha_t = u_t(\omega, \bar{\omega}, \mu_t^\pi)$. \square

The connection presented in the above lemma will be used in the next proposition, in order to reduce the optimization problem in (5.1) to a minimization over weak closed loop tuples, in the following sense.

Definition 5.7. We say that a feasible tuple for (5.1) is a *weak closed loop* if the control is adapted to the state (i.e. α is $\mathbb{F}^{\mathbf{X}}$ -measurable).

We will further need the following concepts of monotonicity: a function $f : \mathcal{P}(\mathbb{R}^N) \rightarrow \mathbb{R}$ is called \prec_{cm} -monotone (resp. \prec_c -monotone) if $f(m_1) \leq f(m_2)$ whenever $m_1 \prec_{cm} m_2$ (resp. $m_1 \prec_c m_2$). With the latter order of measures, we mean $\int h dm_1 \leq \int h dm_2$ for all functions h which are convex and increasing w.r.t. the usual componentwise order in \mathbb{R}^N (resp. all convex functions h) such that the integrals exist.

Proposition 5.8. *Assume*

(A1) $b(x, \cdot, \mu)$ is injective, $b(x, \mathbb{R}^k, \mu)$ is a convex set, and $b^{-1}(x, \cdot, \mu)$ is convex;

(A2) $f(x, b^{-1}(x, \cdot, \mu), \xi)$ is convex and grows at least like $\kappa_0 + \kappa_1 |\cdot|^p$ with $\kappa_1 > 0, p \geq 1$;

(A3) $f(x, \alpha, \cdot)$ is \prec_{cm} -monotone.

Then the minimization in the extended Mean Field problem (5.1) can be taken over weak closed loop tuples. Moreover, if the infimum is attained, then the optimal control α is of weak closed loop form.

The proof follows the projection arguments used in [1], which requires the above convexity assumptions. On the other hand, no regularity conditions are required here, unlike in the classical PDE or probabilistic approaches (see Assumptions (I)-(II) in Section 3). We refer to [25] for a similar statement, in a general framework, but under no non-linear dependence on the control law. This proof is postponed to Appendix A.

Remark 5.9. If b is linear with positive coefficient for α , then assumption (A3) in Proposition 5.8 can be weakened:

(A3') $f(x, \alpha, \cdot)$ is \prec_c -monotone,

as can be seen from the proof. For example, conditions (A1),(A2),(A3') are satisfied if

$$b(x, \alpha, \mu) = c_1 x + c_2 \alpha + c_3 \bar{\mu} \quad \text{and} \quad f(x, \alpha, \xi) = d_1 x + d_2 \alpha + d_3 x^2 + d_4 \alpha^2 + J(\bar{\xi}_1, \bar{\xi}_2),$$

where J is a measurable function, $\bar{\mu} = \int x \mu(dx)$, $\bar{\xi}_1 = \int \int x \xi(dx, d\alpha)$, $\bar{\xi}_2 = \int \int \alpha \xi(dx, d\alpha)$, and c_i, d_i are constants such that $c_2 \neq 0, d_4/c_2 > 0$.

5.4 A Transport-Theoretic Discretization Scheme

In this part we specialize the analysis to the following particular case of (5.1):

$$\inf_{\mathbb{P}, \alpha} \left\{ \int_0^1 f(\mathcal{L}_{\mathbb{P}}(\alpha_t)) dt + g(\mathcal{L}_{\mathbb{P}}(X_T)) : dX_t = \alpha_t dt + dW_t, X_0 = x_0 \right\}, \quad (5.9)$$

where for simplicity we took $T = 1$. Throughout this section we assume:

- (i) g is bounded from below and lower semicontinuous w.r.t. weak convergence;
- (ii) f is increasing with respect to convex order, lower semicontinuous w.r.t. weak convergence, and such that for all $\lambda \in [0, 1]$ and \mathbb{R}^k -valued random variables Z, \bar{Z} :

$$f(\mathcal{L}(\lambda Z + (1 - \lambda)\bar{Z})) \leq \lambda f(\mathcal{L}(Z)) + (1 - \lambda)f(\mathcal{L}(\bar{Z})); \quad (5.10)$$

- (iii) f satisfies the growth condition $f(\rho) \geq a + b \int |z|^p \rho(dz)$ for some $a \in \mathbb{R}, b > 0, p > 1$.

Lemma 5.6 shows the equivalence of (5.9) with the variational problem

$$\inf_{\pi \in \Pi_c^{\ll}(\gamma, \cdot)} \left\{ \int_0^1 f(\mathcal{L}_\pi(\overline{\omega - \omega})_t) dt + g(\mathcal{L}_\pi(\bar{\omega}_1)) \right\}.$$

Under the convention that $\int_0^1 f(\mathcal{L}_\pi(\overline{\omega - \omega})_t) dt = +\infty$ if $\bar{\omega} - \omega \ll \mathcal{L}$ fails under π , the latter can be expressed in the equivalent form:

$$\inf_{\mu \in \tilde{\mathcal{P}}} \inf_{\pi \in \Pi_c(\gamma, \mu)} \left\{ \int_0^1 f(\mathcal{L}_\pi(\overline{\omega - \omega})_t) dt + g(\mathcal{L}_\pi(\bar{\omega}_1)) \right\}, \quad (P)$$

where $\tilde{\mathcal{P}}$ was defined in (5.2). In the same spirit as [36, Ch. 3.6], we introduce a family of causal transport problems in finite dimension increasing to (P). For $n \in \mathbb{N}$, let $\mathbb{T}_n := \{i 2^{-n} : 0 \leq i \leq 2^n, i \in \mathbb{N}\}$ be the n -th generation dyadic grid. For measures $m \in \mathcal{P}(\mathcal{C})$ and $\pi \in \mathcal{P}(\mathcal{C} \times \mathcal{C})$, we write

$$m_n := \mathcal{L}_m(\{\omega_t\}_{t \in \mathbb{T}_n}) \in \mathcal{P}(\mathbb{R}^{(2^n+1)d}) \quad \text{and} \quad \pi_n := \mathcal{L}_\pi(\{(\omega_t, \bar{\omega}_t)\}_{t \in \mathbb{T}_n}) \in \mathcal{P}(\mathbb{R}^{(2^n+1)d} \times \mathbb{R}^{(2^n+1)d})$$

for the projections of m and π on the grid \mathbb{T}_n . We denote by $(x_0^n, x_1^n, \dots, x_{2^n}^n, y_0^n, y_1^n, \dots, y_{2^n}^n)$ a typical element of $\mathbb{R}^{(2^n+1)d} \times \mathbb{R}^{(2^n+1)d}$, and let $\Delta^n x_i := x_{i+1}^n - x_i^n$, and similarly for $\Delta^n y_i$.

We consider the auxiliary transport problems

$$\inf_{\mu \in \mathcal{P}(\mathbb{R}^{(2^n+1)d})} \inf_{\pi \in \Pi_c^n(\gamma_n, \mu)} \left\{ 2^{-n} \sum_{i=0}^{2^n-1} f\left(\mathcal{L}_\pi\left(\frac{\Delta^n y_i - \Delta^n x_i}{2^{-n}}\right)\right) + g(\mathcal{L}_\pi(y_{2^n}^n)) \right\}, \quad (P(n))$$

where, in analogy to Definition 5.5, we called

$$\Pi_c^n(\gamma_n, \mu) \subset \mathcal{P}(\mathbb{R}^{(2^n+1)d} \times \mathbb{R}^{(2^n+1)d})$$

the set of causal couplings in $\mathcal{P}(\mathbb{R}^{(2^n+1)d} \times \mathbb{R}^{(2^n+1)d})$ with marginals γ_n and μ ; see [5].

Theorem 5.10. *Suppose Problem (P) is finite, and that (i), (ii), (iii) hold. Then the value of the auxiliary problems (P(n)) increases to the value of the original problem (P), and the latter admits an optimizer.*

Remark 5.11. An example of a function satisfying Conditions (ii)-(iii) of Theorem 5.10 is $f(\rho) = R(\int h d\rho)$, for R convex and increasing, and h convex with p -power growth ($p > 1$). It also covers the case of functions of the form $f(\rho) = \int \phi(w, z) d\rho(w) d\rho(z) + \int |x|^p d\rho(x)$, with ϕ jointly convex and bounded from below, and $f(\rho) = \text{Var}(\rho) + \int |x|^p d\rho(x)$, where in both cases $p > 1$. For $p = 2$ the latter falls into the LQ case of Section 4.3.

Proof. Step 1 (Lower bound). Let $\boldsymbol{\mu} \in \tilde{\mathcal{P}}$ and $\pi \in \Pi_c(\boldsymbol{\gamma}, \boldsymbol{\mu})$ with finite cost for Problem (P). Fix $n \in \mathbb{N}$, and denote by π_n the projection of π onto the grid \mathbb{T}_n . We first observe that

$$\int_0^1 f(\mathcal{L}_{\pi}(\overline{\omega} - \omega)_t) dt + g(\mathcal{L}_{\pi}(\overline{\omega}_1)) \geq 2^{-n} \sum_{i=0}^{2^n-1} f\left(\mathcal{L}_{\pi_n}\left(\frac{\Delta^n y_i - \Delta^n x_i}{2^{-n}}\right)\right) + g(\mathcal{L}_{\pi_n}(y_{2^n}^n)). \quad (5.11)$$

Indeed, for $i \in \{0, \dots, 2^n - 1\}$ we have

$$\begin{aligned} \int_{i2^{-n}}^{(i+1)2^{-n}} f(\mathcal{L}_{\pi}(\overline{\omega} - \omega)_t) dt &\geq 2^{-n} f\left(\mathcal{L}_{\pi}\left(\int_{i2^{-n}}^{(i+1)2^{-n}} (\overline{\omega} - \omega)_t \frac{dt}{2^{-n}}\right)\right) \\ &= 2^{-n} f\left(\mathcal{L}_{\pi}\left(\frac{\overline{\omega}_{(i+1)2^{-n}} - \overline{\omega}_{i2^{-n}} - (\omega_{(i+1)2^{-n}} - \omega_{i2^{-n}})}{2^{-n}}\right)\right) \\ &= 2^{-n} f\left(\mathcal{L}_{\pi_n}\left(\frac{\Delta^n y_i - \Delta^n x_i}{2^{-n}}\right)\right), \end{aligned}$$

where for the inequality we used the convexity condition (5.10). Noticing that the first marginal of π_n is equal to $\boldsymbol{\gamma}_n$, the r.h.s. of (5.11) is bounded from below by the value of $(P(n))$. Because $\boldsymbol{\mu}, \pi$ have been chosen having finite cost for Problem (P), but otherwise arbitrary, we conclude that

$$(P) \geq (P(n)) \quad \forall n \in \mathbb{N}.$$

Step 2 (Monotonicity). For $n \in \mathbb{N}$ and $i \in \{0, \dots, 2^n - 1\}$, take k such that

$$i2^{-n} = (k-1)2^{-(n+1)} < k2^{-(n+1)} < (k+1)2^{-(n+1)} = (i+1)2^{-n}.$$

Let $\boldsymbol{\mu}_{n+1} \in \mathcal{P}(\mathbb{R}^{(2^{n+1}+1)d})$ and $\pi_{n+1} \in \Pi_c^{n+1}(\boldsymbol{\gamma}_{n+1}, \boldsymbol{\mu}_{n+1})$. By (5.10) we get

$$\begin{aligned} &2^{-(n+1)} \left\{ f\left(\mathcal{L}_{\pi_{n+1}}\left(\frac{\Delta^{n+1} y_{k-1} - \Delta^{n+1} x_{k-1}}{2^{-(n+1)}}\right)\right) + f\left(\mathcal{L}_{\pi_{n+1}}\left(\frac{\Delta^{n+1} y_k - \Delta^{n+1} x_k}{2^{-(n+1)}}\right)\right) \right\} \\ &\geq 2^{-n} f\left(\mathcal{L}_{\pi_{n+1}}\left(\frac{y_{k+1}^{n+1} - y_{k-1}^{n+1} - (x_{k+1}^{n+1} - x_{k-1}^{n+1})}{2^{-n}}\right)\right) = 2^{-n} f\left(\mathcal{L}_{\pi_n}\left(\frac{\Delta^n y_i - \Delta^n x_i}{2^{-n}}\right)\right), \end{aligned}$$

where π_n is the projection of π_{n+1} on the grid \mathbb{T}_n . Analogously to the previous step, this gives

$$(P(n+1)) \geq (P(n)) \quad \forall n \in \mathbb{N}.$$

Step 3 (Discrete to Continuous). We introduce auxiliary problems in path-space:

$$\inf_{\boldsymbol{\mu} \in \tilde{\mathcal{P}}} \inf_{\pi \in \Pi_c(\boldsymbol{\gamma}, \boldsymbol{\mu})} \left\{ 2^{-n} \sum_{i=0}^{2^n-1} f\left(\mathcal{L}_{\pi}\left(\frac{\Delta_i^n \overline{\omega} - \Delta_i^n \omega}{2^{-n}}\right)\right) + g(\mathcal{L}_{\pi}(\overline{\omega}_1)) \right\}, \quad (P^{aux}(n))$$

where $\Delta_i^n \omega := \omega_{(i+1)2^{-n}} - \omega_{i2^{-n}}$ and likewise for $\Delta_i^n \overline{\omega}$. We now prove that

$$(P^{aux}(n)) = (P(n)) \quad \forall n \in \mathbb{N}. \quad (5.12)$$

First we observe that the l.h.s. of (5.12) is larger than the r.h.s. Indeed, projecting a coupling from $\Pi_c(\boldsymbol{\gamma}, \cdot)$ onto a discretization grid gives again a causal coupling; see [36,

Lemma 3.5.1]. For the converse inequality, note that Remark 5.12 implies that, for any $\nu \in \mathcal{P}(\mathbb{R}^{(2^n+1)d})$ and $\pi \in \Pi_c^n(\gamma_n, \nu)$ with finite cost in $(P(n))$, there exist $\mu \in \tilde{\mathcal{P}}$ and $P \in \Pi_c(\gamma, \mu)$ that give the same cost in $(P^{aux}(n))$.

Step 4 (Convergence). Let us denote

$$c(\pi) := \int_0^1 f(\mathcal{L}_\pi(\overline{\omega - \omega})_t) dt \quad \text{and} \quad c^n(\pi) := 2^{-n} \sum_{i=0}^{2^n-1} f\left(\mathcal{L}_\pi\left(\frac{\Delta_i^n \overline{\omega} - \Delta_i^n \omega}{2^{-n}}\right)\right),$$

the cost functionals defining the optimization problems (P) and $(P^{aux}(n))$. Notice that Step 1 implies $c \geq c^n$, and Step 2 shows that c^n is increasing. We now show that c^n converges to c whenever the latter is finite. For this it suffices to show that

$$\liminf_n c^n(\pi) \geq c(\pi). \quad (5.13)$$

We start by representing c^n in an alternative manner, namely

$$c^n(\pi) = \int_0^1 f\left(\mathcal{L}_\pi\left(\int_{\lfloor t2^n \rfloor 2^{-n}}^{\lfloor t2^n \rfloor + 1} 2^{-n} (\overline{\omega - \omega})_s \frac{ds}{2^{-n}}\right)\right) dt.$$

By Lebesgue differentiation theorem [21, Theorem 6, Appendix E.4], for each pair $(\overline{\omega}, \omega)$ such that $\overline{\omega} - \omega$ is absolutely continuous, there exists a dt -full set of times such that

$$A(t, n) := \int_{\lfloor t2^n \rfloor 2^{-n}}^{\lfloor t2^n \rfloor + 1} 2^{-n} (\overline{\omega - \omega})_s \frac{ds}{2^{-n}} \rightarrow (\overline{\omega - \omega})_t. \quad (5.14)$$

If $c(\pi) < \infty$, the set of such pairs $(\overline{\omega}, \omega)$ is π -full. This shows that (5.14) holds $\pi(d\omega, d\overline{\omega})dt$ -a.s. By Fubini's theorem, there is a dt -full set of times $I \subset [0, 1]$ such that, for $t \in I$, the limit (5.14) holds in the π -almost sure sense (the π -null set depends on t a priori). By dominated convergence, this proves that

$$\forall t \in I : \mathcal{L}_\pi(A(t, n)) \Rightarrow \mathcal{L}_\pi((\overline{\omega - \omega})_t),$$

namely in the sense of weak convergence of measures. By lower-boundedness and lower-semicontinuity of f , together with Fatou's Lemma, we obtain

$$\liminf_n c^n(\pi) \geq \int_0^1 \liminf_n f(\mathcal{L}_\pi(A(t, n))) dt = \int_0^1 f(\mathcal{L}_\pi((\overline{\omega - \omega})_t)) dt,$$

establishing (5.13) and so that $c^n \nearrow c$.

By Steps 2 and 3, we know that the values of $(P^{aux}(n))$ are increasing and bounded from above by the value of (P) . We take π^n which is $1/n$ -optimal for $(P^{aux}(n))$. It follows then by Assumptions (i)-(iii) that $\int_0^1 [(\overline{\omega - \omega})_t]^p dt d\pi^n \leq \bar{a} + \bar{b}(P)$, for some $\bar{a}, \bar{b} \in \mathbb{R}$. By [36, Lemma 3.6.2], we obtain the tightness of $\{\pi^n\}_n$. We may thus assume that $\pi_n \Rightarrow \pi$ weakly. By [1, Lemma 5.5], the measure π is causal (and it obviously has first marginal γ). For $k \leq n$ we have

$$c^k(\pi^n) \leq c^n(\pi^n) \leq 1/n + (P^{aux}(n)) \leq 1/n + (P),$$

so, sending $n \rightarrow \infty$, we get $c^k(\pi) \leq \lim_n (P^{aux}(n)) \leq (P)$, as clearly c^k is lower semicontinuous. By letting $k \rightarrow \infty$, and using the fact that $c^k \nearrow c$, we conclude that π is optimal for (P) , and that the value of (P) is the limit of the increasing values of $(P^{aux}(n))$, which in turn equals the limit of the increasing values of $(P(n))$. \square

We complete the argument used in Step 3 with the following remark. This also shows how, from an (approximate) optimizer of the discrete-time problem ($P(n)$), an approximate optimizer of the continuous-time problem (5.9) can be built.

Remark 5.12. In Lemma B.3 we show how, given $\nu \in \mathcal{P}(\mathbb{R}^{2d})$ and $\pi \in \Pi_c^0(\gamma_0, \nu)$, there exists a weak solution (W, X) of an SDE such that $\mathcal{L}_P(W_0, W_1, X_0, X_1) = \pi$. The argument used to prove Lemma B.3 can be iterated in order to get an SDE whose unique weak solution fits any joint distribution over finitely many time-points: For any given $\nu \in \mathcal{P}(\mathbb{R}^{(2^n+1)d})$ and $\pi \in \Pi_c^n(\gamma_n, \nu)$, there exist $\mu \in \tilde{\mathcal{P}}$ and $P \in \Pi_c(\gamma, \mu)$ such that

$$\mathcal{L}_P(\omega_0, \omega_{2^{-n}}, \omega_{2^{-n+1}}, \dots, \omega_1, \bar{\omega}_0, \bar{\omega}_{2^{-n}}, \bar{\omega}_{2^{-n+1}}, \dots, \bar{\omega}_1) = \pi,$$

with P being the joint law of (W, X) , the unique weak solution of an SDE of the form

$$dX_t = \beta_t dt + dW_t.$$

Lemma B.3 covers the case $n = 0$. We now show the case $n = 1$, the general case following similarly. Fix $\nu \in \mathcal{P}(\mathbb{R}^{3d})$ and $\pi \in \Pi_c^1(\gamma_1, \nu)$. As in Lemma B.3, if U_1 is a d -dimensional uniform distribution, independent of X_0 and of the Brownian motion W , then there exists Ψ_1 such that $(0, W_{1/2}, X_0, \Psi_1(U_1, W_{1/2}, X_0)) \sim \pi_1$, where π_1 is the projection of π into the first 4 coordinates. Introducing U_2 , an independent copy of U_1 , we can apply Lemma B.1 in the Appendix, obtaining the existence of a measurable function Ψ_2 such that

$$\left(0, W_{1/2}, W_1, X_0, \Psi_1(U_1, W_{1/2}, X_0), \Psi_2(U_2, W_{1/2}, W_1, X_0, \Psi_1(U_1, W_{1/2}, X_0))\right) \sim \pi.$$

Now we define the following SDE with initial condition X_0 :

$$dX_t = \left(\frac{\Psi_1(U_1, W_t, X_0) - X_t}{1/2 - t} 1_{[0, 1/2)}(t) + \frac{\Psi_2(U_2, W_{1/2}, W_t, X_0, X_{1/2}) - X_t}{1 - t} 1_{[1/2, 1)}(t) \right) dt + dW_t.$$

This admits a unique solution in $[0, 1)$, which is given by

$$\begin{aligned} X_t &= X_0(1 - 2t)1_{[0, 1/2)}(t) + X_{1/2}(2 - 2t)1_{[1/2, 1)}(t) \\ &+ \left(\frac{1}{2} - t \wedge \frac{1}{2}\right) \int_0^{t \wedge 1/2} \frac{\Psi_1(U_1, W_s, X_0)}{(1/2 - s)^2} ds + (1 - t) \int_{t \wedge 1/2}^t \frac{\Psi_2(U_2, W_{1/2}, W_s, X_0, X_{1/2})}{(1 - s)^2} ds \\ &+ \left(\frac{1}{2} - t \wedge \frac{1}{2}\right) \int_0^{t \wedge 1/2} \frac{1}{1/2 - s} dW_s + (1 - t) \int_{t \wedge 1/2}^t \frac{1}{1 - s} dW_s. \end{aligned}$$

Noting $X_{\frac{1}{2}-} = \Psi_1(U_1, W_{1/2}, X_0)$ and $X_{1-} = \Psi_2(U_2, W_{1/2}, W_1, X_0, X_{1/2})$, we conclude.

A Proof of Proposition 5.8

Proof. Fix $(\Omega, \mathbb{F}, \mathbb{P}, \mathbf{W}, \mathbf{X}, \boldsymbol{\alpha})$ feasible tuple for (5.1), if it exists, and set $\pi := \mathcal{L}_{\mathbb{P}}(\mathbf{W}, \mathbf{X}) \in \Pi_c^{\leq}(\gamma, \cdot)$ and $\mu := \mu^\pi$. Under π we have $\bar{\omega}_t - \omega_t = x_0 + \int_0^t \beta_s ds$ for some progressive β . By (A2), the optional projection of β w.r.t. $(\pi, \{\emptyset, \mathcal{C}\} \times \mathbb{F}^2)$, which we call $\bar{\beta}$, is well defined. As in [1], one can prove that the process $M_t := \bar{\omega}_t - x_0 - \int_0^t \bar{\beta}_s(\bar{\omega}) ds$ is a (μ, \mathbb{F}^2) -martingale. Indeed, taking $0 \leq s < t \leq T$ and $h_s \in L^\infty(\mathbb{F}_s^2)$, we have

$$\mathbb{E}^\mu[(M_t - M_s)h_s(\bar{\omega})] = \mathbb{E}^\pi[(\omega_t - \omega_s)h_s(\bar{\omega})] + \mathbb{E}^\pi\left[h_s(\bar{\omega}) \int_s^t \left(\overbrace{(\bar{\omega} - \omega)}^\cdot\right)_r - \bar{\beta}_r(\bar{\omega})\right] dr$$

$$= \mathbb{E}^\pi \left[h_s(\bar{\omega}) \int_s^t \mathbb{E}^\pi \left[\left(\widehat{(\bar{\omega} - \omega)}_r - \bar{\beta}_r(\bar{\omega}) \right) | \mathbb{F}_r^2 \right] dr \right] = 0,$$

where the second equality follows since ω , which is a (γ, \mathbb{F}^1) -martingale, is consequently by causality a $(\pi, \mathbb{F}^1 \otimes \mathbb{F}^2)$ -martingale. Therefore \mathbf{M} is a (μ, \mathbb{F}^2) -martingale, as claimed.

Since $\langle \mathbf{M} \rangle_t = \langle \bar{\omega} \rangle_t = t$ under μ , then \mathbf{M} is actually a (μ, \mathbb{F}^2) -Brownian motion, by Lévy's theorem. This implies $\hat{\pi} := \mathcal{L}_\mu(\mathbf{M}, \bar{\omega}) \in \Pi_c(\gamma, \mu)$. We are next going to show that the expectation in (5.8) is smaller when considering $\hat{\pi}$ instead of π , i.e., when replacing $\beta = \widehat{\bar{\omega} - \omega}$ with $\bar{\beta}$. Then, by taking $\Omega = \mathcal{C}$, $\mathbb{P} = \mu$, $\mathbb{F} = \mathbb{F}^2$, $\mathbf{X} = \bar{\omega}$ and $\alpha = b^{-1}(\bar{\omega}_t, \cdot, \mu_t)(\bar{\beta}_t)$, we have a feasible tuple, which concludes the proof of the proposition.

Let us show our claim. Set $\bar{u}_t(\bar{\omega}, \mu) := b^{-1}(\bar{\omega}_t, \cdot, \mu_t)(\bar{\beta}_t)$ and note that, by (A2) and Jensen's inequality,

$$f(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu), \mathcal{L}_\pi(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu))) \leq \mathbb{E}^\pi \left[f(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t), \mathcal{L}_\pi(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu))) | \mathbb{F}_t^2 \right].$$

By taking expectation under π on both sides, integrating and using Fubini's theorem, we then get

$$\mathbb{E}^\mu \left[\int_0^T f(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu), \mathcal{L}_\pi(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu))) dt \right] \tag{A.1}$$

$$\leq \mathbb{E}^\pi \left[\int_0^T f(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t), \mathcal{L}_\pi(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu))) dt \right]. \tag{A.2}$$

We now establish some ordering between measures. For any measurable function $F : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$ and sigma-field σ , set $\bar{F} := \mathbb{E}^\pi[F|\sigma]$, and note that for any convex function $q : \mathbb{R} \rightarrow \mathbb{R}$, Jensen's inequality gives $\int q(x) d(\mathcal{L}_\pi(\bar{F}))(x) = \mathbb{E}^\pi[q(\bar{F})] \leq \mathbb{E}^\pi[q(F)] = \int q(x) d(\mathcal{L}_\pi(F))(x)$, i.e., $\mathcal{L}_\pi(\bar{F}) \prec_c \mathcal{L}_\pi(F)$. Analogously, for any convex function $H : \mathbb{R} \rightarrow \mathbb{R}$, we have that $\mathcal{L}_\pi(H(\bar{F})) \prec_{cm} \mathcal{L}_\pi(H(F))$. By (A1) and (A3) this implies

$$\begin{aligned} & \mathbb{E}^\pi \left[\int_0^T f(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t), \mathcal{L}_\pi(\bar{\omega}_t, \bar{u}_t(\bar{\omega}, \mu))) dt \right] \\ & \leq \mathbb{E}^\pi \left[\int_0^T f(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t), \mathcal{L}_\pi(\bar{\omega}_t, u_t(\omega, \bar{\omega}, \mu_t))) dt \right]. \end{aligned}$$

Together with (A.1), this concludes our claim, and so the proof of the proposition. \square

B Measurable selection of pushforwarding maps

The next result is obvious in dimension one. In higher dimensions it could follow easily from Brenier's theorem in optimal transport, under assumptions relating to the finiteness of second moments. We do not assume this, and therefore we need to be more careful. For the meaning of concepts such as c -cyclical monotonicity, we refer to [34].

Lemma B.1. *Let Q be a probability measure on $\mathbb{R}^r \times \mathbb{R}^\ell$, and denote by q the (joint) distribution of the first r coordinates of Q . Then there exists a Borel measurable function $F : \mathbb{R}^r \times [0, 1]^\ell \rightarrow \mathbb{R}^\ell$ such that $(I, F)(q \otimes L) = Q$, where L is the ℓ -dimensional Lebesgue measure on $[0, 1]^\ell$, and $I : \mathbb{R}^r \times [0, 1]^\ell \rightarrow \mathbb{R}^r$ is the projection map $I(x, y) = x$.*

Proof. Let $\mathbb{R}^r \ni x \mapsto Q^x$ be a regular conditional kernel of Q with respect to the first r coordinates. Consider the Borel function $x \mapsto (L, Q^x) \in (\mathcal{P}(\mathbb{R}^\ell))^2$. All assumptions of [34, Corollary 10.44] are satisfied. Thus we have, $q(dx)$ -almost surely, the existence of a unique Borel mapping $F_x(\cdot) : \mathbb{R}^\ell \rightarrow \mathbb{R}^\ell$ such that $F_x(L) = Q^x$ and such that its graph is

cyclically monotone (i.e. c -cyclically monotone for $c = \|\cdot\|^2$). By Lemma B.2 below, there exists a Borel function $F : \mathbb{R}^r \times \mathbb{R}^\ell \rightarrow \mathbb{R}^\ell$ such that $F(x, L) = Q^x$, $q(dx)$ -a.s. We finally verify that $F(q \otimes L) = Q$, which concludes the proof:

$$\begin{aligned} \int \int (h \circ (I, F)) dq \otimes dL &= \int (\int h(x, F(x, y)) L(dy)) q(dx) = \int (\int h(x, y) F(x, L)(dy)) q(dx) \\ &= \int (\int h(x, y) Q^x(dy)) q(dx) = \int h dQ. \end{aligned}$$

□

Lemma B.2. *Let (E, Σ, m) be a σ -finite measure space. Consider a measurable function $E \ni \lambda \mapsto (\mu_\lambda, \nu_\lambda) \in \mathcal{P}(\mathbb{R}^\ell) \times \mathcal{P}(\mathbb{R}^\ell)$, and a function $c : \mathbb{R}^\ell \times \mathbb{R}^\ell \rightarrow \mathbb{R}$ continuous and bounded from below. Assume that for m -a.e. λ , there exists a unique mapping $F_\lambda : \mathbb{R}^\ell \rightarrow \mathbb{R}^\ell$ satisfying: F_λ is Borel measurable with $F_\lambda(\mu_\lambda) = \nu_\lambda$, and the graph of F_λ is c -cyclically monotone. Then there exists a measurable $F : E \times \mathbb{R}^\ell \rightarrow \mathbb{R}^\ell$ such that $m(d\lambda)$ -a.s: $F(\lambda, y) = F_\lambda(y)$, $\mu_\lambda(dy)$ -a.s.*

Proof. Let $\tilde{\Pi}(\mu, \nu) := \{\pi \in \Pi(\mu, \nu) : \text{supp}(\pi) \text{ is } c\text{-cyclically monotone}\}$. We first note that the set-valued map $(\mu, \nu) \mapsto \tilde{\Pi}(\mu, \nu)$ is measurable. To wit, $\tilde{\Pi}(\mu, \nu)$ is closed and the pre-image of closed sets by $\tilde{\Pi}(\cdot, \cdot)$ are closed. The argument for the first fact is contained in the proof of Theorem 5.20 in [34, p. 77]. As for the second fact, let $\Sigma \subset \mathcal{P}(\mathbb{R}^\ell \times \mathbb{R}^\ell)$ be closed, and $(\mu_n, \nu_n) \rightarrow (\mu, \nu)$ with $(\mu_n, \nu_n) \in \tilde{\Pi}^{-1}(\Sigma)$. The latter means that there exists $\pi_n \in \Pi(\mu, \nu) \cap \Sigma$ with $\text{supp}(\pi_n)$ being c -cyclically monotone. By Prokhorov's theorem, up to selection of a subsequence, we may assume that $\pi_n \rightarrow \pi \in \Pi(\mu, \nu) \cap \Sigma$, and again reasoning as in the proof of Theorem 5.20 in [34] we also get that π has c -cyclically monotone support. This implies $(\mu, \nu) \in \tilde{\Pi}^{-1}(\Sigma)$, and all in all we get the measurability of $\tilde{\Pi}(\cdot, \cdot)$. We also remark that $\tilde{\Pi}(\mu, \nu) \neq \emptyset$, by the argument in the first paragraph of the proof of Theorem 10.42 in [34, p. 251]. We now closely follow the arguments in the proof of Theorem 1.1 in [22]. First remark that the set-valued mapping

$$(\mu, \nu) \mapsto \Phi(\mu, \nu) := \overline{\bigcup_{\pi \in \tilde{\Pi}(\mu, \nu)} \text{supp}(\pi)} \subset \mathbb{R}^\ell \times \mathbb{R}^\ell$$

is measurable. This easily follows, similarly to [22, Theorem 2.1], by the measurability of $(\mu, \nu) \mapsto \tilde{\Pi}(\mu, \nu)$. Now [22, Corollary 2.3] is valid for our Φ without any changes. Finally, the proof of Theorem 1.1 in [22] can be fully translated in our terms.

□

We provide the missing argument for Step 3 in the proof of Theorem 5.10, which is used in Remark 5.12. We use the notation adopted in that part of the article.

Lemma B.3. *Given $\nu \in \mathcal{P}(\mathbb{R}^{2d})$ and $\pi \in \Pi_c^0(\gamma_0, \nu)$, there exist $\mu \in \mathcal{P}(\mathcal{C})$ and $P \in \Pi_c(\gamma, \mu)$ such that $\mathcal{L}_P(\omega_0, \omega_1, \bar{\omega}_0, \bar{\omega}_1) = \pi$. This measure P is the joint law of the unique weak solution of an SDE of the form $dX_t = \beta_t dt + dW_t$, namely $P = \mathcal{L}(\mathbf{W}, \mathbf{X})$.*

Proof. Recall that $\gamma_0(dz_0, dz_1) = \delta_0(dz_0) \mathcal{N}(dz_1)$ where \mathcal{N} is the standard Gaussian in \mathbb{R}^d . We consider a probability space supporting a random variable U uniformly distributed in $[0, 1]^d$, a random variable X_0 distributed according to the first marginal of ν , and a standard Brownian motion \mathbf{W} , such that U, X_0, \mathbf{W} are independent. We first observe that, by Lemma B.1, there exists a Borel function $\Psi : \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$(0, W_1, X_0, \Psi(U, W_1, X_0)) \sim \pi.$$

Second, we define the following SDE, with initial condition X_0 :

$$dX_t = \frac{\Psi(U, W_t, X_0) - X_t}{1-t} dt + dW_t.$$

Note that there is at most one solution to this SDE on every interval $[0, T]$ with $T < 1$, by the theory of Lipschitz SDEs with random coefficients. This proves that the solution is unique on $[0, 1)$. Third, we observe that a solution of the above SDE is given by

$$X_t = X_0(1-t) + (1-t) \int_0^t \frac{\Psi(U, W_s, X_0)}{(1-s)^2} ds + (1-t) \int_0^t \frac{1}{1-s} dW_s,$$

and therefore this is the unique solution on $[0, 1)$. Finally, we observe that sending $t \rightarrow 1$ (by L'Hôpital rule) we have $X_1 := X_{1-} = \Psi(U, W_1, X_0)$. We now observe that $\mathcal{L}(W_0, W_1, X_0, X_1) = \pi$ as desired, and notice that $P := \mathcal{L}(\mathbf{W}, \mathbf{X})$ is causal (Definition 5.5), since \mathbf{X} is adapted to the filtration $\mathcal{G}_t := \{(U, X_0, W_s) : s \leq t\}$ and \mathbf{W} is a \mathcal{G} -Brownian motion. \square

References

- [1] Beatrice Acciaio, Julio Backhoff-Veraguas, and Anastasiia Zalashko. Causal optimal transport and its links to enlargement of filtrations and continuous-time stochastic optimization. *arXiv:1611.02610*, 2016.
- [2] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.
- [3] Daniel Andersson and Boualem Djehiche. A maximum principle for sdes of mean-field type. *Applied Mathematics & Optimization*, 63(3):341–356, 2011.
- [4] Julio Backhoff-Veraguas, Mathias Beiglböck, Manu Eder, and Alois Pichler. Fundamental properties of process distances. *arXiv:1701.03955*, 2017.
- [5] Julio Backhoff-Veraguas, Mathias Beiglböck, Yiqing Lin, and Anastasiia Zalashko. Causal transport in discrete time and applications. *SIAM Journal on Optimization*, 27(4):2528–2562, 2017.
- [6] Matteo Basei and Huyên Pham. Linear-quadratic McKean-Vlasov stochastic control problems with random coefficients on finite and infinite horizon, and applications. *arXiv:1711.09390*, 2017.
- [7] Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- [8] Alain Bensoussan, Jens Frehse, and Phillip Yam. *Mean Field Games and Mean Field Type Control Theory*, volume 101. Springer, 2013.
- [9] Rainer Buckdahn, Boualem Djehiche, and Juan Li. A general stochastic maximum principle for sdes of mean-field type. *Applied Mathematics & Optimization*, 64(2):197–216, 2011.

- [10] Rainer Buckdahn, Boualem Djehiche, Juan Li, Shige Peng, et al. Mean-field backward stochastic differential equations: a limit approach. *The Annals of Probability*, 37(4):1524–1565, 2009.
- [11] Pierre Cardaliaguet. Notes on P.L. Lions’ lectures at the College de France. Technical report, Working Paper, September, 2013.
- [12] Pierre Cardaliaguet and Charles-Albert Lehalle. Mean field game of controls and an application to trade crowding. *Mathematics and Financial Economics*, pages 1–29, 2016.
- [13] René Carmona and François Delarue. Forward–backward stochastic differential equations and controlled McKean–Vlasov dynamics. *The Annals of Probability*, 43(5):2647–2700, 2015.
- [14] René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications. Volume I: Mean Field FBSDEs, Control and Games*. Springer, 2017.
- [15] René Carmona, François Delarue, and Aimé Lachapelle. Control of McKean–Vlasov dynamics versus Mean Field Games. *Mathematics and Financial Economics*, 7(2):131–166, 2013.
- [16] René Carmona and Daniel Lacker. A probabilistic weak formulation of mean field games and applications. *The Annals of Applied Probability*, 25(3):1189–1231, 2015.
- [17] Álvaro Cartea and Sebastian Jaimungal. Incorporating order-flow into optimal execution. *Mathematics and Financial Economics*, 10(3):339–364, 2016.
- [18] Ana Bela Cruzeiro and Rémi Lassalle. Weak calculus of variations for functionals of laws of semi-martingales. *arXiv:1501.05134*, 2015.
- [19] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
- [20] Nelson Dunford and Jacob T Schwartz. *Linear Operators. Part 1: General Theory*. New York Interscience, 1957.
- [21] Lawrence Evans. *Partial differential equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 1998.
- [22] Joaquin Fontbona, Hélène Guérin, and Sylvie Méléard. Measurability of optimal transportation and strong coupling of martingale measures. *Electronic Communications in Probability*, 15:124–133, 2010.
- [23] Emmanuel Gobet and Maxime Grangereau. McKean optimal stochastic control of microgrid equipped with photo-voltaic panels and battery. *Preprint*, 2018.
- [24] P. Jameson Graber. Linear quadratic mean field type control and mean field games with common noise, with application to production of an exhaustible resource. *Applied Mathematics & Optimization*, 74(3):459–486, 2016.

- [25] Daniel Lacker. Limit theory for controlled McKean–Vlasov dynamics. *SIAM Journal on Control and Optimization*, 55(3):1641–1672, 2017.
- [26] Rémi Lassalle. Causal transport plans and their Monge–Kantorovich problems. *Stochastic Analysis and Applications*, 36(3):452–484, 2018.
- [27] Rémi Lassalle and Jean Claude Zambrini. A weak approach to the stochastic deformation of classical mechanics. *Journal of Geometric Mechanics*, 8(2), 2016.
- [28] Thilo Meyer-Brandis, Bernt Øksendal, and Xun Yu Zhou. A mean-field stochastic maximum principle via malliavin calculus. *Stochastics An International Journal of Probability and Stochastic Processes*, 84(5-6):643–666, 2012.
- [29] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [30] Georg Ch Pflug. Version-independence and nested distributions in multistage stochastic optimization. *SIAM Journal on Optimization*, 20(3):1406–1420, 2009.
- [31] Georg Ch Pflug and Alois Pichler. A distance for multistage stochastic optimization models. *SIAM Journal on Optimization*, 22(1):1–23, 2012.
- [32] Georg Ch Pflug and Alois Pichler. *Multistage stochastic optimization*. Springer, 2016.
- [33] Huyên Pham and Xiaoli Wei. Bellman equation and viscosity solutions for mean-field stochastic control problem. *ESAIM: Control, Optimisation and Calculus of Variations*, 24(1):437–461, 2018.
- [34] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [35] Jiongmin Yong. Linear-quadratic optimal control problems for mean-field stochastic differential equations. *SIAM Journal on Control and Optimization*, 51(4):2809–2838, 2013.
- [36] Anastasiia Zalashko. Causal optimal transport: theory and applications. *PhD Thesis, University of Vienna*, 2017.