

Probabilistic compactification methods for stochastic optimal control and mean field games

Daniel Lacker

June 2018

Contents

| | | |
|----------|--|-----------|
| 1 | Overview | 1 |
| 2 | Weak convergence of probability measures | 5 |
| 2.1 | $\mathcal{P}(\mathcal{X})$ as a metric space | 7 |
| 2.2 | Compactness in $\mathcal{P}(\mathcal{X})$ | 8 |
| 2.3 | Useful continuity results | 9 |
| 3 | Deterministic optimal control theory | 9 |
| 3.1 | An example of the chattering phenomenon | 14 |
| 4 | Stochastic optimal control | 15 |
| 4.1 | Existence of optimal relaxed controls | 18 |
| 4.2 | Existence of optimal strict controls | 21 |
| 4.3 | Extensions | 24 |
| 5 | Mean field games | 25 |
| 5.1 | Relaxed mean field equilibrium | 27 |
| 5.2 | Markovian equilibria | 32 |

1 Overview

These notes are the companion for a four-lecture series given in June 2018 at the IPAM Graduate Summer School on Mean Field Games and Applications. The goal of the course is to explain a methodology for the theory of mean field games coming from a series of papers of the author [7, 29–31]. Most of the material is drawn from [29]. In fact, the first half of the course is

solely about optimal control theory, the goal being to expose the powerful but somewhat old-fashioned methodology of generalized or *relaxed controls*. Only in the second half of the course will we adapt these tools to the study of mean field games.

The key ideas for this methodology dates back to the work of L.C. Young on generalized solutions of problems of calculus of variations [41, 42], for which he developed the theory of what is known today as *Young measures* or, in a control-theoretic context, relaxed controls. Essentially, when an optimization problem involves the choice of a measurable function $\alpha : [0, T] \rightarrow A$, where $T > 0$ and A is some metric space, it is very often convenient to relax the problem by viewing α as inducing a positive measure on $[0, T] \times A$, namely

$$dt\delta_{\alpha_t}(da).$$

The closure of the set of such measures in the topology of weak convergence (i.e., the dual topology induced by bounded continuous functions on $[0, T] \times A$), where α ranges over measurable functions from $[0, T] \rightarrow A$, is precisely the set \mathcal{V} of all measures on $[0, T] \times A$ whose first marginal is Lebesgue measure. This (metrizable) space has two extremely convenient properties:

- (1) Compact subsets of \mathcal{V} are plentiful and easy to identify using Prokhorov's theorem. In particular, if A is compact, then so is \mathcal{V} .
- (2) The set \mathcal{V} is convex.

With property (1) in mind, a “relaxed” optimal control problem can typically be formulated as the optimization of a continuous function over a compact set, and this immediately gives existence of an optimizer. Once existence of a relaxed optimizer is established, with a bit more work one can often then show that the optimizer is in fact a “strict” control, i.e., induced by a measurable function in the above sense.

This rather abstract approach to optimal control was developed mainly in the 60s-80s for optimal control problems, both deterministic [36, 42] and stochastic [2, 15, 16]. It was primarily used to prove merely the existence of optimal controls, which was a popular problem at the time, with similar analysis avoiding explicit use of relaxed controls appearing in [4, 13, 26, 40]. This capstone of this line of research was arguably the work of El Karoui et al. [9], Hausmann-Lepeltier [19], and later Kurtz-Stockbridge [25], which proved in great generality the existence of optimal *Markovian* controls. The methods are powerful enough to adapt to mixed control/stopping problems [19], singular control [20, 21], and partial information [10]. Beyond an

existence theory, this approach facilitates a general proof of the dynamic programming principle [9, 11, 12], the precursor to the derivation of the HJB equation. Finally, we mention the work of Kushner on approximations of and numerical methods for stochastic control problems, which takes full advantage of the pleasant topological setting afforded by relaxed controls; see [27] and also the book of Kushner-Dupuis [28].

There is a price for the great generality and cleanliness afforded by this powerful theory: It is completely abstract, yielding essentially no information about the nature of the optimal control which is proven to exist. This is in stark contrast with the two more standard approaches to optimal control, which reduce the problem to a Hamilton-Jacobi-Bellman equation or a forward-backward stochastic differential equation arising from an application of the Pontryagin maximum principle. Both of these approaches are appealing in that these equations, once solved, provide a way to *construct* the optimal control. Solving these equations explicitly, however, is rarely possible, and even proving well-posedness theorems requires much more restrictive assumptions than the relaxed control approach.

At this point, we can finally discuss why this powerful but somewhat old-fashioned theory is worth reviving for the analysis of mean field games (MFGs). To solve a MFG requires that one solve a family of optimal control problems and then resolve a fixed point problem built on these optimal control problems. Occasionally, one is lucky enough to solve this fixed point problem explicitly, for instance by solving the PDE system introduced in the seminal work of Lasry-Lions [33–35]. More often, one must be content to merely prove existence (and occasionally uniqueness) of a solution. These notes will show how the relaxed control methodology provides a powerful framework for proving quite general theorems on the existence of (Markovian) solutions of this fixed point problem. Moreover, the convenient compactness property (1) above enables a detailed analysis of the convergence of n -player Nash equilibria to the MFG limit, though it is unlikely we will have time to discuss this much.

Let us dwell a bit longer, and more philosophically, on the advantages and disadvantages of this approach to MFG analysis. To prove existence of a fixed point, the vast majority of MFG literature proceeds as follows: First, solve (in the loosest sense of the word) each of the optimal control problems using one’s favorite method for studying optimal control, and then apply one of two fixed point theorems:

- (A) Banach’s theorem, which requires the fixed point map to be a contraction on some complete metric space. This has the advantage of

providing a way to constructively find the fixed point, by iterating the fixed point map.

- (B) Compactness-based fixed point theorems à la Schauder or Kakutani, which requires a continuous map from a convex compact set into itself. This is non-constructive in the sense that it does not provide an algorithm for finding the fixed point; in particular, iterating the fixed point map does not converge to a solution in general.

The original work of Huang-Malhamé-Caines [22] on MFGs showed that for small time horizon one can find a contraction. In general, however, contractions are unavailable, and one must rely on the compactness-based fixed point theorems. In doing so, one essentially gives up hope of *constructing* a fixed point! Even if one starts with an extremely constructive method of solving each of the control problems, this concreteness is lost in the application of an abstract fixed point theorem. By working with relaxed controls one embraces this abstractness from the start, and doing so allows for a remarkable level of generality. Moreover, this approach streamlines what is typically the crux of a fixed point analysis, which is to show that the optimal control depends continuously on the model inputs (i.e., the measure flow on which one is performing a fixed point analysis). To find this continuity by PDE methods, on the other hand, typically requires much stronger assumptions on the structure, smoothness, or convexity of various coefficients.

If the reader will forgive one more paragraph of a philosophical nature: In a game-theoretic context, relaxed controls have a natural interpretation in terms of *mixed strategies*. A relaxed control is a measure $q(dt, da)$ on $[0, T] \times A$ with first marginal equal to Lebesgue measure, and the disintegration theorem permits us to write $q(dt, da) = dtq_t(da)$ for some measurable function $[0, T] \ni t \mapsto q_t \in \mathcal{P}(A)$, where $\mathcal{P}(A)$ is the space of probability measures on A . We may interpret this to mean that, at each time t , the controller chooses not a single action in A but rather a probability measure over the action space from which a specific action is randomly sampled. This is completely in line with the notion of a mixed strategy in game theory, in which players choose probability measures over the action set rather than single actions (pure strategies). In fact, once the appropriate spaces are set up, our proof of existence of a solution of the MFG fixed point problem is remarkably parallel to Nash's original proof of existence of mixed strategy equilibria in finite games [37].

These notes will mostly focus on existence theory for MFGs, following the framework put forth in [29]. Subsequently, other authors have extended

the methodology to cover MFGs built on jump-diffusion dynamics [3] and singular controls [17]. But it is worth stressing that the scope of the methodology is not limited to existence proofs, although this is largely the focus of these notes. An equally exciting advantage of the framework developed here is its amenability to proving limit theorems, relating n -player equilibria to the MFGs. Most of the key ingredients for proving these limit theorems are already present in some form in the existence theory presented here, but the details are much more involved. The curious reader should refer to [30] or [14] for the convergence of open-loop n -player equilibria or to the forthcoming [32] for the closed-loop case, whenever the paper is finished (hopefully in 2018).

The notes are organized as follows. First, Section 2 contains a brief summary of the most pertinent facts about weak convergence of probability measures on metric spaces; this material will not be covered in the lectures and should be bread and butter for the working probabilist, but hopefully it is a useful reference for other readers. Section 3 introduces relaxed controls in the simpler setting of deterministic optimal control, with complete proofs of existence of optimal relaxed and (under additional assumptions) strict optimal controls. Section 4 then turns to stochastic optimal control, extending the relaxed approach from the deterministic case from Section 3. Finally, these ideas from optimal control theory are applied to mean field games in Section 5, which proves existence of equilibria.

2 Weak convergence of probability measures

This short section collects some basic and standard facts about weak convergence of probability measures on metric spaces. Again, this material will be reviewed only *very* briefly in the lectures and is included in the notes as a reference for the less probabilistically oriented reader. This machinery forms the foundation for many arguments of these notes, and, in order to get to the meat of the course, we will cover this material far too quickly. For more details, refer to the classic textbook of Billingsley [5]. Weak convergence of processes is treated concisely in Kallenberg's tome [23, Chapter 14], and the book of Parthasarathy [38] is a nice reference for a more topological perspective.

Throughout the section, let (\mathcal{X}, d) denote a metric space. We always equip \mathcal{X} with the Borel σ -field, meaning the σ -field generated by the open sets of \mathcal{X} . We will write $\mathcal{P}(\mathcal{X})$ for the set of (Borel) probability measures on X . For a measurable function φ from \mathcal{X} to another metric space \mathcal{Y} , we

define the image measure $\mu \circ \varphi^{-1} \in \mathcal{P}(\mathcal{Y})$ by setting $\mu \circ \varphi^{-1}(A) := \mu(\varphi \in A) = \mu\{x \in \mathcal{X} : \varphi(x) \in A\}$.

Let $C_b(\mathcal{X})$ denote the set of bounded continuous real-valued functions on X . The fundamental definition is the following, which we state in two equivalent forms, one measure-theoretic and one probabilistic:

Definition 2.1. Given $\mu, \nu_n \in \mathcal{P}(\mathcal{X})$, we say that μ_n converges weakly to μ , or $\mu_n \rightarrow \mu$, if

$$\lim_{n \rightarrow \infty} \int_{\mathcal{X}} f d\mu_n = \int_{\mathcal{X}} f d\mu, \text{ for every } f \in C_b(\mathcal{X}).$$

Definition 2.2. Given a sequence of \mathcal{X} -valued random variables (X_n) , we say that X_n converges weakly (or in distribution) to another \mathcal{X} -valued random variable X (often denoted $X_n \Rightarrow X$) if

$$\lim_{n \rightarrow \infty} \mathbb{E}[f(X_n)] = \mathbb{E}[f(X)], \text{ for every } f \in C_b(\mathcal{X}).$$

When we say X is a \mathcal{X} -valued random variable, we mean the following: Behind the scenes, there is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a function $X : \Omega \rightarrow \mathcal{X}$, measurable with respect to the Borel σ -field on \mathcal{X} . We will not always be explicit about the choice of probability space.

It is important to notice that the weak convergence $\mu_n \rightarrow \mu$ does not imply that $\mu_n(A) \rightarrow \mu(A)$ for every measurable set $A \subset \mathcal{X}$ (unless the metric space \mathcal{X} is finite, in which case $\mathcal{P}(\mathcal{X})$ can be identified with a compact subset of $\mathbb{R}^{|\mathcal{X}|}$). Nonetheless, the following famous theorem clarifies what weak convergence does tell us about setwise convergence:

Theorem 2.3 (Portmanteau theorem). *Let $\mu, \mu_n \in \mathcal{P}(\mathcal{X})$. The following are equivalent:*

- (i) $\mu_n \rightarrow \mu$.
- (ii) $\liminf_{n \rightarrow \infty} \mu_n(U) \geq \mu(U)$ for every open set $U \subset \mathcal{X}$.
- (iii) $\limsup_{n \rightarrow \infty} \mu_n(C) \leq \mu(C)$ for every closed set $C \subset \mathcal{X}$.
- (iv) $\lim_{n \rightarrow \infty} \mu_n(A) = \mu(A)$ for every Borel set $A \subset \mathcal{X}$ with $\mu(A^\circ) = \mu(\bar{A})$.¹
- (v) $\int f d\mu_n \rightarrow \int f d\mu$ for every bounded uniformly continuous function f on \mathcal{X} .

¹Here A° denotes the interior of the set A and \bar{A} the closure.

We omit the proof of Theorem 2.3, as it is entirely classical (see [5, Theorem 2.1] or [23, Theorem 4.25]). The following theorem will often be used implicitly and is completely trivial to prove:

Theorem 2.4 (Continuous mapping theorem). *Suppose \mathcal{X} and \mathcal{Y} are metric spaces, and (X_n) is a sequence of \mathcal{X} -valued random variables converging in distribution to another \mathcal{X} -valued random variable X . Suppose $g : \mathcal{X} \rightarrow \mathcal{Y}$ is a continuous function. Then $g(X_n) \Rightarrow g(X)$.*

2.1 $\mathcal{P}(\mathcal{X})$ as a metric space

It is an important fact that weak convergence of probability measures corresponds to a metric topology on $\mathcal{P}(\mathcal{X})$, at least when the underlying metric space (\mathcal{X}, d) is separable. The weak convergence topology on $\mathcal{P}(\mathcal{X})$ is of course generated by the basic open sets of the form

$$\left\{ \nu \in \mathcal{P}(\mathcal{X}) : \left| \int f_i d\nu - \int f_i d\mu \right| < \epsilon_i, \quad i = 1, \dots, k \right\},$$

where $k \in \mathbb{N}$, $\mu \in \mathcal{P}(\mathcal{X})$, and $f_1, \dots, f_k \in C_b(\mathcal{X})$. The following theorem is well known and can be found, for instance, in [38, Section II.6].

Theorem 2.5. *If (\mathcal{X}, d) is a separable metric space, then $\mathcal{P}(\mathcal{X})$ can be metrized as a separable metric space. If (\mathcal{X}, d) is a complete and separable metric space, then $\mathcal{P}(\mathcal{X})$ can be metrized as a complete and separable metric space. If (\mathcal{X}, d) is a compact metric space, then $\mathcal{P}(\mathcal{X})$ can be metrized as a compact metric space.*

There are several popular metrics, including the Levy-Prokhorov metric, the family of Wasserstein metrics (most of which metrize a somewhat stronger topology than that of weak convergence), or the bounded-Lipschitz metric. While we will rarely need to explicitly use such a metric, the bounded-Lipschitz metric is often quite convenient to use in practice. For $\mu, \nu \in \mathcal{P}(\mathcal{X})$, define

$$d_{\text{BL}}(\mu, \nu) = \sup_f \left(\int_{\mathcal{X}} f d\mu - \int_{\mathcal{X}} f d\nu \right),$$

where the supremum is over all functions $f : \mathcal{X} \rightarrow \mathbb{R}$ satisfying

$$\sup_{x \in \mathcal{X}} |f(x)| \leq 1, \quad |f(x) - f(y)| \leq d(x, y) \quad \forall x, y \in \mathcal{X}.$$

See [8, Theorem 11.3.3] for proof of the following:

Proposition 2.6. *If (\mathcal{X}, d) is a separable metric space and $\mu_n, \mu \in \mathcal{P}(\mathcal{X})$, then $\mu_n \rightarrow \mu$ weakly if and only if $d_{\text{BL}}(\mu_n, \mu) \rightarrow 0$.*

2.2 Compactness in $\mathcal{P}(\mathcal{X})$

Given the title of the course, it should come as no surprise that characterizing compact sets of $\mathcal{P}(\mathcal{X})$ will be extremely useful. The classical theorem of Prokhorov accomplishes this. Given a set $S \subset \mathcal{P}(\mathcal{X})$, we say that the family S of probability measures is *tight* if for all ϵ there exists a compact set $K \subset X$ such that

$$\sup_{\mu \in S} \mu(K^c) \leq \epsilon.$$

The importance of this definition lies in the following theorem, the proof of which can be found in [5, Theorem 6.1, 6.2] and [23, Theorem 16.3]

Theorem 2.7 (Prokhorov's theorem). *Suppose $S \subset \mathcal{P}(\mathcal{X})$. If S is tight, then it is pre-compact in $\mathcal{P}(\mathcal{X})$. Conversely, if S is pre-compact, and if the metric space (\mathcal{X}, d) is separable and complete, then S is tight.*

As a first application, we see how easily tightness lends itself to working on product spaces:

Lemma 2.8. *Suppose \mathcal{X}_1 and \mathcal{X}_2 are complete, separable metric spaces, and endow $\mathcal{X}_1 \times \mathcal{X}_2$ with any metric compatible with the product topology. Define the projections $\pi_i : \mathcal{X}_1 \times \mathcal{X}_2 \rightarrow \mathcal{X}_i$, for $i = 1, 2$. Then a set $S \subset \mathcal{P}(\mathcal{X}_1 \times \mathcal{X}_2)$ is tight if and only if the sets $S_1 = \{\mu \circ \pi_1^{-1} : \mu \in S\}$ and $S_2 = \{\mu \circ \pi_2^{-1} : \mu \in S\}$ are tight in $\mathcal{P}(\mathcal{X}_1)$ and $\mathcal{P}(\mathcal{X}_2)$, respectively.*

Proof. Suppose first that S_1 and S_2 are tight. Let $\epsilon > 0$. for $i = 1, 2$, find a compact set $K_i \subset \mathcal{X}_i$ such that $\sup_{\mu \in S} \mu \circ \pi_i^{-1}(K_i^c) \leq \epsilon/2$. Then $K_1 \times K_2$ is compact, and for each $\mu \in S$ we have

$$\begin{aligned} \mu((K_1 \times K_2)^c) &= \mu(S_1 \times K_2^c) + \mu(K_1^c \times K_2) \\ &\leq \mu \circ \pi_2^{-1}(K_2^c) + \mu \circ \pi_1^{-1}(K_1^c) \\ &\leq \epsilon. \end{aligned}$$

On the other hand, if S is tight, then there exists a compact set $K \subset \mathcal{X}_1 \times \mathcal{X}_2$ such that $\sup_{\mu \in S} \mu(K^c) \leq \epsilon$. Then $\pi_i(K) \subset \mathcal{X}_i$ is compact for each $i = 1, 2$, and for each $\mu \in S$ we have

$$\begin{aligned} \mu \circ \pi_i^{-1}(\pi_i(K)) &= \mu \{(x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2 : \pi_i(x_1, x_2) \in \pi_i(K)\} \\ &\geq \mu \{(x_1, x_2) \in K : \pi_i(x_1, x_2) \in \pi_i(K)\} \\ &= \mu(K) \geq 1 - \epsilon. \end{aligned}$$

This shows that S_1 and S_2 are tight. □

2.3 Useful continuity results

In the analysis in these notes, we will frequently need continuity properties of the following form:

Lemma 2.9. *Suppose \mathcal{X} and \mathcal{Y} are complete, separable metric spaces, and let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ be bounded and jointly continuous. Then the map*

$$\mathcal{P}(\mathcal{X}) \times \mathcal{Y} \ni (\mu, y) \mapsto \int_{\mathcal{X}} f(x, y) \mu(dy) \in \mathbb{R}$$

is jointly continuous.

Closely related is the following:

Lemma 2.10. *Suppose \mathcal{X} , \mathcal{Y} , and \mathcal{Z} are complete, separable metric spaces, and let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ be jointly continuous. Then the map*

$$\mathcal{P}(\mathcal{X}) \times \mathcal{Y} \ni (\mu, y) \mapsto \mu \circ f(\cdot, y)^{-1} \in \mathcal{P}(\mathcal{Z})$$

is jointly continuous.

Exercise 2.11. Prove Lemmas 2.9 and 2.10.

3 Deterministic optimal control theory

As a warm-up, we consider the deterministic optimal control problem described as follows:

$$(P_s) \begin{cases} \sup_{\alpha} & \left[\int_0^T f(t, x_t^\alpha, \alpha_t) dt + g(x_T) \right], \\ \text{s.t. } x_t^\alpha & = x_0 + \int_0^t b(s, x_s^\alpha, \alpha_s) ds, \quad t \in [0, T]. \end{cases}$$

Here, $x_0 \in \mathbb{R}^d$ is a given initial state, $T > 0$ is a time horizon, and the control α takes values in a set A . For simplicity, we make the following assumptions

Assumption A. The functions $b : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}^d$, $f : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}$, and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ are bounded and continuous. The action space A is a compact metric space.

One needs to be a bit careful about specifying what controls are admissible. Most convenient is to work with open-loop controls, where $\alpha \in L^0([0, T]; A)$ is simply a measurable function of time. Often one wishes to work with feedback controls, where $\alpha_s = \hat{\alpha}(s, x_s)$ for some measurable

function $\hat{\alpha} : [0, T] \times \mathbb{R}^d \rightarrow A$. Neither of these sets lend themselves to compactness arguments. While $L^0([0, T]; A)$ can be topologized by convergence in measure and becomes separable and completely metrizable, compact sets are hard to come by in this space. The idea of *relaxed controls* is to compactify the set of admissible controls in order to turn the problem (P_s) into a problem of maximizing a continuous function over a compact set.

To do this, we embed our controls into a space of measures, for which compactness criteria are straightforward. Precisely, let \mathcal{V} denote the set of positive Borel measures on $[0, T] \times A$ with first marginal equal to Lebesgue measure. That is, $q \in \mathcal{V}$ satisfies $q([s, t] \times A) = t - s$ for $0 \leq s < t \leq T$. Note that every $q \in \mathcal{V}$ thus has total mass T . We may endow \mathcal{V} with the topology of weak convergence, which means that $q_n \rightarrow q$ if and only if $\int \varphi dq_n \rightarrow \int \varphi dq$ for every bounded continuous function $\varphi : [0, T] \times A \rightarrow \mathbb{R}$. Two incredibly useful lemmas are the following:

Lemma 3.1. *Suppose $q, q_n \in \mathcal{V}$ with $q_n \rightarrow q$. Then, for any bounded measurable function $\varphi : [0, T] \times A \rightarrow \mathbb{R}$ such that $\varphi(t, \cdot)$ is continuous on A for each $t \in [0, T]$, we have $\int \varphi dq_n \rightarrow \int \varphi dq$.*

Exercise 3.2. Prove Lemma 3.1. **Hint:** Show that the map $\Phi : [0, T] \rightarrow C(A)$ defined by $\Phi(t) = \varphi(t, \cdot)$ is Borel measurable. Then, by Lusin's theorem, there exists for each $\delta > 0$ a *continuous* function $\Phi_\delta : [0, T] \rightarrow C(A)$ such that $\Phi_\delta = \Phi$ except on a set of Lebesgue measure less than δ .

Lemma 3.3. *The space \mathcal{V} is compact and metrizable.*

Proof. Assume $T = 1$ for simplicity, as the general case amounts to no more than a rescaling. Because $[0, 1] \times A$ is a compact metric space, the space $\mathcal{P}([0, 1] \times A)$ is compact and metrizable when equipped with the topology of weak convergence (see Theorem 2.5). We also note that \mathcal{V} is a closed subset of $\mathcal{P}([0, 1] \times A)$: A measure $\mu \in \mathcal{P}([0, 1] \times A)$ belongs to \mathcal{V} if and only if its first marginal is the uniform (Lebesgue) measure. The map $[0, 1] \times A \ni (t, a) \mapsto t \in [0, 1]$ is continuous, and so if $\mu_n \rightarrow \mu$ in $\mathcal{P}([0, 1] \times A)$ then the first marginal must converge, thanks to the continuous mapping Theorem 2.4.

Alternatively, instead of referring to Theorem 2.5, we can note that the Banach space $C([0, T] \times A)$ of continuous functions on $[0, T] \times A$ is separable. Hence, by Banach-Alaoglu, closed bounded sets in the dual $C([0, T] \times A)^*$ are weak*-compact and metrizable. Clearly \mathcal{V} is bounded, and closedness is argued as above. \square

By the disintegration theorem, note that every $q \in \mathcal{V}$ can be written as

$$q(dt, da) = dtq_t(da),$$

for some measurable family of probability measures $(q_t)_{t \in [0, T]}$ on A , uniquely defined up to almost-everywhere equality. Any (open-loop) control $\alpha \in L^0([0, T]; A)$,² now called a *strict control*, can be embedded in \mathcal{V} by identifying it with the measure

$$q(dt, da) = dt\delta_{\alpha_t}(da).$$

We may write \mathcal{V}_0 for the set of $q \in \mathcal{V}$ of the above form. Elements of \mathcal{V}_0 are called *strict controls*, while general elements of \mathcal{V} are called *relaxed controls*. The following remarkable fact is quite important, but we leave it as a (tricky) exercise, given that we will not use it in this course:

Exercise 3.4. Show that \mathcal{V}_0 is dense in \mathcal{V} . This is sometimes known in the literature as the *chattering lemma*, and the example of Section 3.1 illustrates the meaning of the term “chattering” in this context. **Hint:** You may use the following implication of the Borel Isomorphism Theorem: Suppose \mathcal{X} and \mathcal{Y} are complete, separable metric spaces. Suppose μ and ν are nonnegative measures on \mathcal{X} and \mathcal{Y} , respectively, with the same total mass, i.e., $\mu(\mathcal{X}) = \nu(\mathcal{Y})$. If μ is nonatomic, then there exists a Borel-measurable function $\varphi : \mathcal{X} \rightarrow \mathcal{Y}$ such that $\mu \circ \varphi^{-1} = \nu$.

We can now define a relaxed analogue of (P_s) . To be precise, we should notice that the ODE defining x^q may not be well-posed. Hence, we make the following definitions: Let

$$\mathcal{C}^d = C([0, T]; \mathbb{R}^d)$$

denote the continuous path space, endowed with supremum norm. Let \mathcal{R} denote the set of $(x, q) \in \mathcal{C}^d \times \mathcal{V}$ such that

$$x_t = x_0 + \int_{[0, t] \times A} b(s, x_s, a)q(ds, da), \quad t \in [0, T].$$

We may now define:

$$(P_r) \quad \sup_{(x, q) \in \mathcal{R}} \left[\int_{[0, T] \times A} f(t, x_t, a)q(dt, da) + g(x_T) \right].$$

Using the above two lemmas, we quickly prove existence of a solution to (P_r) .

²We write $L^0([0, T]; A)$ for the set of (equivalence classes of a.e. equal) measurable functions from $[0, T]$ to A .

Theorem 3.5. *Suppose \mathcal{R} is nonempty. Then the supremum in (P_r) is attained.*

Proof. The goal is to check that \mathcal{R} is compact and that the map $\Gamma : \mathcal{C}^d \times \mathcal{V} \rightarrow \mathbb{R}$ defined by

$$\Gamma(x, q) := \int_{[0, T] \times A} f(t, x_t, a) q(dt, da) + g(x_T) \quad (3.1)$$

is continuous. In fact, continuity of Γ is a consequence of Lemma 2.9. Hence, we focus on compactness. Fix a sequence $(x^n, q^n) \in \mathcal{R}$. Because \mathcal{V} is compact, we may assume that there exists $q \in \mathcal{V}$ such that $q^n \rightarrow q$. By boundedness of b , we have

$$|x_t^n - x_s^n| = \left| \int_{[s, t] \times A} b(r, x_r^n, a) q^n(dr, da) \right| \leq \|b\|_\infty (t - s),$$

for $s < t$. Hence, by Arzela-Ascoli, the sequence (x^n) is precompact in \mathcal{C}^d . By passing to a further subsequence, we may assume $x^n \rightarrow x$ for some $x \in \mathcal{C}^d$, endowed with the supremum norm. We then find, for each $t \in [0, T]$

$$\begin{aligned} x_t &= \lim_n x_t^n = \lim_n \left(x_0 + \int_{[0, t] \times A} b(s, x_s^n, a) q^n(ds, da) \right) \\ &= x_0 + \int_{[0, t] \times A} b(s, x_s, a) q(ds, da), \end{aligned}$$

which shows that $(x, q) \in \mathcal{R}$. Indeed, the last limit follows from Lemma 2.9. This completes the proof. \square

Now that we have proven existence of a relaxed solution, we want to see when we can build from it a *strict* optimal control. This can be done under a suitable convexity assumption, often known as Roxin's condition:

Theorem 3.6. *Suppose that $\mathcal{R} \neq \emptyset$, and suppose also that for each $(t, x) \in [0, T] \times \mathbb{R}^d$ the set*

$$K(t, x) = \{(b(t, x, a), z) : a \in A, f(t, x, a) \geq z\} \subset \mathbb{R}^d \times \mathbb{R}$$

*is convex. (As usual, we implicitly grant Assumption **A** as well.) Then, for every $(x, q) \in \mathcal{R}$, there exists $q^0 \in \mathcal{V}_0$ such that $(x, q^0) \in \mathcal{R}$ and $\Gamma(x, q^0) \geq \Gamma(x, q)$.*

Proof. Disintegrate $q(dt, da) = q_t(da)dt$. By convexity of $K(t, x)$, we first observe that

$$\int_A \left(b(t, x_t, a), f(t, x_t, a) \right) q_t(da) \in K(t, x_t),$$

for each $t \in [0, T]$. Hence, for each $t \in [0, T]$, we may find $\alpha_t \in A$ and $z_t \in \mathbb{R}$ such that $z_t \leq f(t, x_t, \alpha_t)$ and

$$\int_A \left(b(t, x_t, a), f(t, x_t, a) \right) q_t(da) = \left(b(t, x_t, \alpha_t), z_t \right).$$

There is an important subtle point here, which is that we must choose (α_t, z_t) to be measurable functions t . There are established measurable selection theorems that ensure this is possible, and we will take this for granted; see, for example, Filippov's implicit function theorem [1, Theorem 18.17] or [19, Lemma A.9]. With measurable choices of α and z in hand, we define

$$q^0(dt, da) := dt\delta_{\alpha_t}(da),$$

and we complete the proof by two simple observations. First, for each $t \in [0, T]$,

$$x_t = x_0 + \int_0^t \int_A b(s, x_s, a) q_s(da) ds = x_0 + \int_0^t b(s, x_s, \alpha_s) ds.$$

Second, we have

$$\begin{aligned} \int_{[0, T] \times A} f(t, x_t, a) q^0(dt, da) &= \int_0^T f(t, x_t, \alpha_t) dt \\ &\geq \int_0^T z_t dt \\ &= \int_0^T \int_A f(t, x_t, a) q_t(da) dt. \end{aligned}$$

Hence,

$$\begin{aligned} \Gamma(x, q^0) &= \int_{[0, T] \times A} f(t, x_t, a) q^0(dt, da) + g(x_T) \\ &\geq \int_{[0, T] \times A} f(t, x_t, a) q(dt, da) + g(x_T) \\ &= \Gamma(x, q). \end{aligned}$$

□

Combining Theorems 3.5 and 3.6 provides a pretty satisfactory existence result:

Corollary 3.7. *Suppose that \mathcal{R} is nonempty and that the convexity assumption of Theorem 3.6 holds. Then the relaxed and strict control problems (P_r) and (P_s) have the same value, and there exists a strict optimal control.*

Example 3.8. The most common example of the assumption of Theorem 3.6 is the following. Assume the control space A is a convex subset of \mathbb{R}^k for some k . Assume b is affine in a , meaning it is of the form $b(t, x, a) = b_1(t, x) + b_2(t, x)a$, where b_1 and b_2 take values in \mathbb{R}^d and $\mathbb{R}^{d \times k}$, respectively. Lastly, assume that $a \mapsto f(t, x, a)$ is concave for each (t, x) . We may then write $K(t, x) = b_1(t, x) + b_2(t, x)K'(t, x)$, where $K'(t, x)$ is the set

$$K'(t, x) = \{(a, z) : a \in A, z \leq f(t, x, a)\} \subset A \times \mathbb{R} \subset \mathbb{R}^k \times \mathbb{R}.$$

The set $K'(t, x)$ is precisely the *hypograph* of the function $a \mapsto f(t, x, a)$, or the set of coordinates lying below the graph of the function. For a concave function the hypograph is a concave set, and so $K'(t, x)$ is convex. Hence, the linear transformation $K(t, x)$ of this set is also convex.

Remark 3.9. The first statement of Corollary (3.7), that the relaxed and strict control problems (P_r) and (P_s) have the same value, is true even without the convexity assumption of Theorem 3.6. Indeed, this follows from continuity of the functional Γ defined in (3.1) and from Exercise 3.4, which says that the set of strict controls is dense in the set of relaxed controls.

3.1 An example of the chattering phenomenon

It is instructive to see an example of when a relaxed optimal control exists while there is no strict optimal control. Consider the one-dimensional problem describe by the data

$$T = 1, \quad A = \{-1, 1\}, \quad b(t, x, a) = a, \quad f(t, x, a) = -|x|, \quad g(x) = 0.$$

That is, the controller tries to solve the following optimization problem:

$$V := \inf_{\alpha \in L^0([0,1];\{-1,1\})} \int_0^1 |x_t^\alpha| dt, \quad \text{where} \quad x_t^\alpha := \int_0^t \alpha_s ds.$$

Clearly $V \geq 0$. On the other hand, consider the control α^n which alternates between ± 1 on intervals of length $1/n$. Precisely, define $\alpha_t^n = +1$ if $t \in$

$[\frac{k}{n}, \frac{k+1}{n})$ for some odd integer k , and set $\alpha_t^n = -1$ otherwise. Then, for n even, the function $t \mapsto x_t^{\alpha^n}$ is periodic, with period $2/n$, and we have

$$x_t^{\alpha^n} = \begin{cases} t & \text{if } t \in [0, 1/n), \\ \frac{2}{n} - t & \text{if } t \in [1/n, 2/n). \end{cases}$$

The integral of this “triangular” function

$$\int_0^1 |x_t^{\alpha^n}| dt = n \int_0^{1/n} t dt = \frac{1}{2n}.$$

This shows that $V \leq 1/2n$ for each n , and so in fact $V = 0$. But there is no control achieving this value! Indeed, if $\alpha \in L^0([0, 1]; \{-1, 1\})$ had $\int_0^1 |x_t^\alpha| dt = 0$, then we would have $x_t^\alpha = 0$ for a.e. t , and thus $\alpha_t = 0$ for a.e. t . But this is not allowed, as controls must take values in $\{-1, 1\}$.

We see that the controller can approach the optimal value by using the rapidly oscillating or *chattering* controls α^n . As relaxed controls these controls simply approximate $q(dt, da) := dt (\frac{1}{2}\delta_1 + \frac{1}{2}\delta_{-1})(da)$. That is, $dt\delta_{\alpha_t^n}(da) \rightarrow q$. Indeed, if we allow the controller to choose among relaxed controls, then this choice q gives $x_t = 0$ for all t and thus achieves the optimal value.

In a sense, the problem here is that the control space A is not convex. Using the above *chattering* controls effectively convexifies the control space.

4 Stochastic optimal control

We now develop the analogous story for stochastic control problems, in which the state process is governed not by an ODE but by an SDE driven by Brownian motion. This methodology matured in the 80s with the papers [9, 19], but the essential ideas go back to the work of Fleming and Nisio [15, 16]. The compactification method is powerful enough to handle quite general setups, in which the controlled state process follows dynamics of the form

$$dX_t = b(t, X_t, \alpha_t)dt + \sigma(t, X_t, \alpha_t)dW_t. \quad (4.1)$$

For the sake of the reader less familiar with stochastic calculus, and to avoid a significant increase in technical difficulty, we choose to restrict our attention to the case where the diffusion coefficient σ is constant, and in fact we take it to be the identity matrix. That is, our state process will follow the dynamics

$$dX_t = b(t, X_t, \alpha_t)dt + dW_t.$$

Here $X = (X_t)_{t \in [0, T]}$ is a d -dimensional process, and $W = (W_t)_{t \in [0, T]}$ is a d -dimensional Brownian motion.

To define the control problem precisely, we work under the same Assumption **A** defined in the previous section, and we work with a highly convenient *weak formulation*, in which the controller gets to choose the probability space. Throughout, we will fix a probability measure λ_0 on \mathbb{R}^d to represent a given initial distribution for the state process.

Definition 4.1. A *control rule* is a tuple $\mathcal{A} = (\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha, X)$, where:

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$.
2. W is an \mathbb{F} -Brownian motion.
3. X is a continuous \mathbb{F} -adapted process, and $X_0 \sim \lambda_0$.
4. α is a progressively measurable A -valued process.
5. The state equation holds,

$$X_t = X_0 + \int_0^t b(s, X_s, \alpha_s) ds + W_t, \quad t \in [0, T].$$

For a control rule $\mathcal{A} = (\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha, X)$, we define the *value of \mathcal{A}* as

$$J(\mathcal{A}) := \mathbb{E}^{\mathbb{P}} \left[\int_0^T f(t, X_t, \alpha_t) dt + g(X_T) \right].$$

An *optimal control rule* optimizes J over all control rules.

Remark 4.2. Stochastic control problems are often posed in the *strong formulation*, as opposed to the above *weak formulation*. In the strong formulation, the probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is fixed, typically with \mathbb{F} as the filtration generated by the Brownian motion. The controller must choose α to be progressively measurable with respect to this given filtration, as a response to the noise of the Brownian motion. One must then be careful to restrict attention to those controls for which the state equation is well-posed. When b is Lipschitz in x , uniformly in (t, a) , this is not much of a restriction, and it can typically be shown that the optimal value is the same in both the strong and weak formulations. See [9, Section 4] or the more recent [12, Section 4.4] for details, or even [31, Theorem 2.4] for the case of controlled McKean-Vlasov dynamics.

Remark 4.3. It is worth noting also that an alternative weak formulation is available in this special case in which the diffusion coefficient is uncontrolled and the drift is bounded. This approach uses Girsanov's theorem to write the value function as the solution of a BSDE [39, Section 6.4], characterizing an optimal but generally path-dependent control. It is not easy with these techniques to find optimal *Markovian* controls, nor is it possible to adapt the framework to cover controlled diffusion coefficient.

Remark 4.4. A nice feature of stochastic control in the weak formulation is that it is quite easy to check that the set of control rules is nonempty. Indeed, as we have assumed our drift function b to be bounded, for any constant control $a \in A$ the SDE

$$dX_t = b(t, X_t, a)dt + dW_t$$

has a weak solution which is unique in law. This is a well known consequence of Girsanov's theorem [24, Propositions 5.3.6 and 5.3.10].

Definition 4.5. A *relaxed control rule* is a tuple $\mathcal{A} = (\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda, X)$, where:

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$.
2. W is an \mathbb{F} -Brownian motion.
3. X is a continuous \mathbb{F} -adapted process, and $X_0 \sim \lambda_0$.
4. Λ is a \mathcal{V} -valued random variable, and $\Lambda([0, s] \times S)$ is \mathcal{F}_t -measurable for all $s \leq t$ and all Borel sets $S \subset A$.
5. The state equation holds,

$$X_t = X_0 + \int_{[0, t] \times A} b(s, X_s, a) \Lambda(ds, da) + W_t, \quad t \in [0, T].$$

For a relaxed control rule $\mathcal{A} = (\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda, X)$, we define the *value of \mathcal{A}* as

$$J(\mathcal{A}) := \mathbb{E}^{\mathbb{P}} \left[\int_{[0, T] \times A} f(t, X_t, \alpha_t) \Lambda(dt, da) + g(X_T) \right].$$

An *optimal relaxed control rule* optimizes J over all control rules.

4.1 Existence of optimal relaxed controls

The first main theorem is the following, which proves existence of optimal relaxed controls in a rather general setting:

Theorem 4.6. *Under Assumption **A**, there exists an optimal relaxed control rule.*

To prove Theorem 4.6, we reformulate the problem slightly, by redefining a “control rule” in terms of the probability law of the state and control processes. In this section, we work on the canonical measurable space $\Omega = \mathcal{C}^d \times \mathcal{V}$, equipped with its Borel σ -field \mathcal{F} . We define *canonical processes* X and Λ , simply by setting $X(x, q) = x$ and $\Lambda(x, q) = q$ for $(x, q) \in \Omega$, and similarly $X_t(x, q) = x_t$ for $t \in [0, T]$. We endow Ω with the natural filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$, where \mathcal{F}_t is the σ -field generated by the following the random variables X_s and $\Lambda([0, s] \times B)$, where $s \leq t$ and $B \subset A$ is a Borel set.

Recall in the following that the initial state distribution $\lambda_0 \in \mathcal{P}(\mathbb{R}^d)$ is fixed throughout.

Definition 4.7. Let \mathcal{R} denote the set of probability measures P on $\mathcal{C}^d \times \mathcal{V}$ such that:

1. $P \circ X_0^{-1} = \lambda_0$.
2. Under P , the process $W = (W_t)_{t \in [0, T]}$ defined by

$$W_t = X_t - X_0 - \int_{[0, t] \times A} b(s, X_s, a) \Lambda(ds, da) \quad (4.2)$$

is a \mathbb{F} -Brownian motion. Equivalently, $P \circ W^{-1}$ equals Wiener measure, and it holds that $W_s - W_t$ is independent of \mathcal{F}_t under P whenever $0 \leq t \leq s \leq T$.

The point of this definition is to pose the relaxed optimal control problem as an optimization over the set \mathcal{R} , which the following lemma formalizes:

Lemma 4.8. *For each relaxed control rule $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}}, \tilde{\mathbb{P}}, \tilde{W}, \tilde{\Lambda}, \tilde{X})$, the probability measure $\tilde{\mathbb{P}} \circ (\tilde{X}, \tilde{\Lambda})^{-1}$ belongs to \mathcal{R} . Conversely, for every $P \in \mathcal{R}$, the tuple $(\Omega, \mathcal{F}, \mathbb{F}, P, W, \Lambda, X)$ is a relaxed control rule, where W is defined as in (4.2).*

Exercise 4.9. Prove Lemma 4.8.

Remark 4.10. For problems with controlled diffusion coefficient, as in (4.1), Definition 4.7 should be refined. In particular, it is not immediately obvious how to adapt property (2). The right way to proceed is to formulate the controlled SDEs in terms of *martingale problems*, and the curious reader is referred to [9] for details. Once the right formulation is ironed out, the essential structure of the arguments is the same as presented here.

This lemma shows that finding a relaxed optimal control rule is equivalent to finding an optimizer in the control problem

$$\sup_{P \in \mathcal{R}} \int_{\Omega} \Gamma dP = \sup_{P \in \mathcal{R}} \mathbb{E}^P \left[\int_{[0, T] \times A} f(t, X_t, a) \Lambda(dt, da) + g(X_T) \right]$$

where we define the functional $\Gamma : \Omega \rightarrow \mathbb{R}$ as in (3.1) by

$$\Gamma(x, q) := \int_{[0, T] \times A} f(t, x_t, a) q(dt, da) + g(x_T).$$

Theorem 4.11. *There exists $P^* \in \mathcal{R}$ such that $\sup_{P \in \mathcal{R}} \int_{\Omega} \Gamma dP = \int_{\Omega} \Gamma dP^*$.*

Note that Theorem 4.6 follows immediately from Theorem 4.11 and Lemma 4.8. To prove Theorem 4.11, we will again argue that the problem is nothing but the maximization of a continuous function over a compact set. Recall that Γ is continuous by Lemma 2.9. As it is bounded, we conclude that the map $\mathcal{P}(\Omega) \ni P \mapsto \int_{\Omega} \Gamma dP \in \mathbb{R}$ is continuous.

It suffices to prove that \mathcal{R} is compact. Do accomplish this, we will use a nice tightness criterion due to Aldous [23, Theorem 16.11]:

Theorem 4.12 (Aldous' criterion for tightness). *Consider a family (X^n) of \mathcal{C}^d -valued random variables (i.e., continuous stochastic processes with values in \mathbb{R}^d). Suppose that for any $\delta_n > 0$ with $\delta_n \downarrow 0$ and any stopping times τ_n with values in $[0, T - \delta_n]$ we have³*

$$\lim_{n \rightarrow \infty} \sup_n \mathbb{E} [\min\{1, |X_{\tau_n + \delta_n}^n - X_{\tau_n}^n|\}] = 0.$$

Then (X^n) is tight.

Proof of Theorem 4.11. We first show that \mathcal{R} is tight. It suffices (see Lemma 2.8) to show that the families $\{P \circ X^{-1} : P \in \mathcal{R}\} \subset \mathcal{P}(\mathcal{C}^d)$ and $\{P \circ \Lambda^{-1} :$

³We say that τ_n is a *stopping time* to mean relative to the filtration of X^n . That is, τ_n is a $[0, T]$ -valued random variable with the property that the set $\{\tau_n \leq t\}$ is measurable with respect to $\sigma(X_s^n : s \leq t)$, for each $t \in [0, T]$.

$P \in \mathcal{R}\} \subset \mathcal{P}(\mathcal{V})$ are tight. The latter set is tight by Prokhorov's Theorem 2.7, because the space \mathcal{V} is compact. To show that the former set is tight, we use Theorem 4.12. For any $\delta > 0$ and any \mathbb{F} -stopping time τ on Ω with values in $[0, T - \delta]$, we have

$$\begin{aligned} |X_{\tau+\delta} - X_\tau| &= \left| \int_{[\tau, \tau+\delta] \times A} b(t, X_t, a) \Lambda(dt, da) + W_{\tau+\delta} - W_\tau \right| \\ &\leq \|b\|_\infty \Lambda([\tau, \tau+\delta] \times A) + |W_{\tau+\delta} - W_\tau| \\ &= \delta \|b\|_\infty + |W_{\tau+\delta} - W_\tau|. \end{aligned}$$

Take expectations to get

$$\mathbb{E}|X_{\tau+\delta} - X_\tau| \leq \delta \|b\|_\infty + \sqrt{\delta}, \quad (4.3)$$

where we used the strong Markov property and Jensen's inequality to get

$$\mathbb{E}|W_{\tau+\delta} - W_\tau| = \mathbb{E}|N(0, \delta)| \leq \mathbb{E}[|N(0, \delta)|^2]^{1/2} = \sqrt{\delta},$$

where $N(0, \delta)$ denotes a generic normal random variable with mean zero and variance δ . From (4.3) and Theorem 4.12, we conclude that $\{P \circ X^{-1} : P \in \mathcal{R}\} \subset \mathcal{P}(\mathcal{C}^d)$ is tight and thus that \mathcal{R} is tight.

To complete the proof of Theorem 4.11, it suffices now to show that \mathcal{R} is closed. To do this, let us fix $P_n \in \mathcal{R}$ and $P \in \mathcal{P}(\Omega)$ with $P_n \rightarrow P$. We show that $P \in \mathcal{R}$ by checking the two properties of Definition 4.7. By the continuous mapping Theorem 2.4,

$$P \circ X_0^{-1} = \lim_n P_n \circ X_0^{-1} = \lambda_0.$$

Similarly, $P \circ W^{-1} = \lim_n P_n \circ W^{-1}$ equals Wiener measure. It remains to show that $W_s - W_t$ is independent of \mathcal{F}_t , for $0 \leq t \leq s \leq T$. To do this, fix $t < s$, a bounded continuous function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$, and a bounded, \mathcal{F}_t -measurable, and continuous function $h : \Omega \rightarrow \mathbb{R}$. By continuity and the fact that $P_n \in \mathcal{R}$ for each n , we have

$$\begin{aligned} \mathbb{E}^P [\varphi(W_s - W_t)h(X, \Lambda)] &= \lim_n \mathbb{E}^{P_n} [\varphi(W_s - W_t)h(X, \Lambda)] \\ &= \lim_n \mathbb{E}^{P_n} [\varphi(W_s - W_t)] \mathbb{E}^{P_n} [h(X, \Lambda)] \\ &= \mathbb{E}^P [\varphi(W_s - W_t)] \mathbb{E}^P [h(X, \Lambda)]. \end{aligned}$$

In fact, it follows then that the identity

$$\mathbb{E}^P [\varphi(W_s - W_t)h(X, \Lambda)] = \mathbb{E}^P [\varphi(W_s - W_t)] \mathbb{E}^{P_n} [h(X, \Lambda)]$$

holds not only for *continuous* functions φ and h but also for merely measurable functions; indeed, a bounded measurable function of $W_s - W_t$ can be approximated in $L^1(P)$ by bounded continuous functions of $W_s - W_t$, and, with a bit more thought, a bounded \mathcal{F}_t -measurable function of (X, Λ) can be approximated in $L^1(P)$ by bounded, continuous, \mathcal{F}_t -measurable functions of (X, Λ) . \square

4.2 Existence of optimal strict controls

Adapting the idea of Theorem 3.6, under the same convexity assumption we may find an optimal strict control.

Theorem 4.13. *Suppose that for each $(t, x) \in [0, T] \times \mathbb{R}^d$ the set*

$$K(t, x) = \{(b(t, x, a), z) : a \in A, f(t, x, a) \geq z\} \subset \mathbb{R}^d \times \mathbb{R}$$

is convex. Then, for every $P \in \mathcal{R}$, there exists a strict Markovian control $P^0 \in \mathcal{R}$ such that $\int_{\Omega} \Gamma dP^0 \geq \int_{\Omega} \Gamma dP$ and $P^0 \circ X_t^{-1} = P \circ X_t^{-1}$ for all $t \in [0, T]$. By strict Markovian control we mean that there exists a measurable function $\alpha : [0, T] \times \mathbb{R}^d \rightarrow A$ such that $P^0(\Lambda = dt\delta_{\alpha(t, X_t)}(da)) = 1$.

The second part of the conclusion, that $P^0 \circ X_t^{-1} = P \circ X_t^{-1}$ for all $t \in [0, T]$, is not important in this section, but it will be absolutely essential when we study mean field games in Section 5.

The proof of Theorem 4.13 makes use of a beautiful theorem, sometimes referred to as a *mimicking theorem* or *Markovian projection*. It is originally due to Gyöngy [18], later extended by Brunick and Shreve [6].

Theorem 4.14 (Markovian projection). *Suppose that we are given three \mathbb{R}^d -valued stochastic process $(X_t)_{t \in [0, T]}$, $(W_t)_{t \in [0, T]}$, and $(b_t)_{t \in [0, T]}$, with $\mathbb{E} \int_0^T |b_t|^2 dt < \infty$, and with*

$$dX_t = b_t dt + dW_t.$$

All three processes are defined on a common probability space and are adapted with respect to a common filtration with respect to which W is a Brownian motion. Then there exists a measurable function $\hat{b} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that:

(i) $\hat{b}(t, X_t) = \mathbb{E}[b_t | X_t]$ almost surely for each $t \in [0, T]$,

(ii) *There exists a weak solution of the SDE*

$$d\tilde{X}_t = \hat{b}(t, \tilde{X}_t)dt + d\tilde{W}_t, \tag{4.4}$$

such that X_t and \tilde{X}_t have the same law, for each $t \in [0, T]$. To be precise, we can find a filtered probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}}, \tilde{\mathbb{P}})$ supporting d -dimensional $\tilde{\mathbb{F}}$ -adapted processes \tilde{X} and \tilde{W} , where \tilde{W} is an $\tilde{\mathbb{F}}$ -Brownian motion and the above SDE is satisfied, with X_t and \tilde{X}_t having the same law for each t .

To clarify the meaning of this remarkable theorem: Given any Itô process with random coefficients, the time- t marginal distributions can be matched (or *mimicked*) by the solution of a Markovian SDE, with no additional randomness. Note that X and \tilde{X} do not have the same law when viewed as \mathcal{C}^d -valued random variables; the joint distributions of (X_s, X_t) and $(\tilde{X}_s, \tilde{X}_t)$ can be different for $s \neq t$.

Proof sketch for Theorem 4.14. Apply Itô's formula to a smooth function φ of compact support to get

$$\begin{aligned} & \mathbb{E}[\varphi(t, X_t)] - \mathbb{E}[\varphi(0, X_0)] \\ &= \int_0^t \mathbb{E} \left[\partial_t \varphi(s, X_s) + b_s \cdot \nabla \varphi(s, X_s) + \frac{1}{2} \Delta \varphi(s, X_s) \right] ds. \end{aligned}$$

By the tower property of conditional expectation, this implies

$$\begin{aligned} & \mathbb{E}[\varphi(t, X_t)] - \mathbb{E}[\varphi(0, X_0)] \\ &= \int_0^t \mathbb{E} \left[\partial_t \varphi(s, X_s) + \hat{b}(t, X_s) \cdot \nabla \varphi(s, X_s) + \frac{1}{2} \Delta \varphi(s, X_s) \right] ds. \end{aligned}$$

It follows quickly from Girsanov's theorem that the law of X_t admits a density $p(t, \cdot)$ with respect to Lebesgue measure, and the above equation implies that it satisfies

$$\begin{aligned} & \int_{\mathbb{R}^d} \varphi(x) (p(t, x) - p(0, x)) dx \\ &= \int_0^t \int_{\mathbb{R}^d} \left(\partial_t \varphi(s, x) + \hat{b}(s, x) \cdot \nabla \varphi(s, x) + \frac{1}{2} \Delta \varphi(s, x) \right) p(s, x) dx ds. \end{aligned}$$

In other words, p is a weak solution of the Fokker-Planck equation

$$\partial_t p(t, x) = -\operatorname{div}(\hat{b}(t, x)p(t, x)) + \frac{1}{2} \Delta p(t, x).$$

On the other hand, the same argument shows that, if $\tilde{p}(t, \cdot)$ denotes the law of \tilde{X}_t , then \tilde{p} is a weak solution of the same PDE. Uniqueness for weak solutions of this PDE then yields the desired result. Notably, however, the result is valid even for SDEs which correspond to non-unique Fokker-Planck equations, and the arguments of [6] get around this issue. \square

Remark 4.15. Recall that may disintegrate any $q \in \mathcal{V}$ by $q(dt, da) = dtq_t(da)$, where the measurable map $[0, T] \ni t \mapsto q_t \in \mathcal{P}(A)$ is uniquely determined up to (Lebesgue-) almost everywhere equality. In a sense, we think equivalently either about measures on $[0, T] \times A$ or about $\mathcal{P}(A)$ -valued functions. The same is true for random elements of \mathcal{V} , but a bit more care is required to make this precise. That is, for our canonical \mathcal{V} -valued random variable Λ defined on the canonical space Ω , we would like to be able to write $\Lambda(dt, da) = \Lambda_t(da)dt$, where $(\Lambda_t)_{t \in [0, T]}$ is a $\mathcal{P}(A)$ -valued process *with nice measurability properties*. By *nice measurability properties* we mean that the map $[0, T] \times \Omega \ni (t, x, q) \mapsto \Lambda(t, x, q)_t \in \mathcal{P}(A)$ should be predictable with respect to the canonical filtration \mathbb{F} on Ω . This is purely technical and can be accomplished without too much trouble; see [29, Lemma 3.2]. We will henceforth take this “nice disintegration” for granted.

Proof of Theorem 4.13. This follows a similar argument to Theorem 3.6. Fix $P \in \mathcal{R}$. By convexity of $K(t, x)$, we first observe that

$$\int_A \left(b(t, X_t, a), f(t, X_t, a) \right) \Lambda_t(da) \in K(t, X_t),$$

for each $t \in [0, T]$. Hence, taking conditional expectations, we still have

$$\mathbb{E}^P \left[\int_A \left(b(t, X_t, a), f(t, X_t, a) \right) \Lambda_t(da) \middle| X_t \right] \in K(t, X_t).$$

Hence, for each $t \in [0, T]$ and $x \in \mathbb{R}^d$, we may find $\alpha(t, x) \in A$ and $z(t, x) \in \mathbb{R}$ such that $z(t, x) \leq f(t, x, \alpha(t, x))$ and

$$\mathbb{E}^P \left[\int_A \left(b(t, X_t, a), f(t, X_t, a) \right) \Lambda_t(da) \middle| X_t \right] = \left(b(t, X_t, \alpha(t, X_t)), z(t, X_t) \right).$$

We will again ignore the important issue that we must choose $(\alpha(t, x), z(t, x))$ to be jointly measurable functions of (t, x) ; see again [19, Lemma A.9] for details.

Notice that, under P , we have the SDE

$$dX_t = \int_A b(t, X_t, a) \Lambda_t(da) dt + dW_t.$$

We have just found an expression for the conditional expectation of the drift, namely

$$b(t, X_t, \alpha(t, X_t)) = \mathbb{E}^P \left[\int_A b(t, X_t, a) \Lambda_t(da) \middle| X_t \right].$$

Hence, by Theorem 4.14, we can find $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}}, \tilde{\mathbb{P}}, Y, \tilde{B})$ as in Theorem 4.14 such that $\tilde{\mathbb{P}} \circ Y_t^{-1} = P \circ X_t^{-1}$ for each $t \in [0, T]$ and also

$$dY_t = b(t, Y_t, \alpha(t, Y_t))dt + d\tilde{B}_t.$$

Define

$$P^0 = \tilde{\mathbb{P}} \circ (Y, dt\delta_{\alpha(t, Y_t)}(da))^{-1}.$$

Then clearly P^0 belongs to \mathcal{R} , and we use Fubini's theorem to get

$$\begin{aligned} \int_{\Omega} \Gamma dP^0 &= \mathbb{E}^{\tilde{\mathbb{P}}} \left[\int_0^T f(t, Y_t, \alpha(t, Y_t))dt + g(Y_T) \right] \\ &= \int_0^T \mathbb{E}^{\tilde{\mathbb{P}}} [f(t, Y_t, \alpha(t, Y_t))] dt + \mathbb{E}^{\tilde{\mathbb{P}}} [g(Y_T)] \\ &= \int_0^T \mathbb{E}^P [f(t, X_t, \alpha(t, X_t))] dt + \mathbb{E}^{\tilde{\mathbb{P}}} [g(X_T)] \\ &= \mathbb{E}^P \left[\int_0^T f(t, X_t, \alpha(t, X_t))dt + g(X_T) \right] \\ &\geq \mathbb{E}^P \left[\int_0^T z(t, X_t)dt + g(X_T) \right] \\ &= \mathbb{E}^P \left[\int_0^T \int_A f(t, X_t, a)\Lambda_t(da)dt + g(X_T) \right] \\ &= \int_{\Omega} \Gamma dP. \end{aligned}$$

□

4.3 Extensions

Some of these points may have been obvious to the astute reader, but it is worth mentioning some of the many ways in which the above results can be generalized.

First, the coefficients (b, f, g) need not be bounded, and the action space A need not be compact. Instead, one should impose suitable growth assumptions for (b, f, g) , along with a corresponding integrability assumption on the initial distribution λ_0 . Also, quite importantly, the running cost should be *coercive* in a sense; if A is unbounded, then $f(t, x, a)$ should grow “quickly enough” with a in order to ensure that sub-level sets of the objective function $\mathcal{R} \ni P \mapsto \int \Gamma dP$ are precompact. See [19] for details.

Our existence theorems come from the fact that a continuous function achieves its maximum over a compact set. But it is enough that the function

is merely upper semicontinuous. Hence, f and g need only be upper semicontinuous. Moreover, by choosing objective functions which are allowed to equal $-\infty$, one can incorporate hard constraints on the state and control process. See again [19] for details.

A controlled diffusion coefficient $\sigma(t, X_t, \alpha_t)dW_t$ is more difficult but certainly possible to handle in this approach. See [9, 19]. Moreover, one can handle discontinuous coefficients (b, σ) (see [9, Section 7]) and even jump-diffusion state processes (see [9, Section 8]). In fact, this approach has been extended to extremely general contexts of state processes taking values in locally compact metric spaces by Kurtz and Stockbridge [25].

5 Mean field games

We now turn our attention to mean field games, with the prototypical problem described loosely but concisely as follows:

$$(MFG) \begin{cases} \alpha^* \in \arg \max_{\alpha} & \mathbb{E} \left[\int_0^T f(t, X_t^{\mu, \alpha}, \mu_t, \alpha_t) dt + g(X_T^{\mu, \alpha}, \mu_T) \right], \\ \text{s.t. } dX_t^{\mu, \alpha} & = b(t, X_t^{\mu, \alpha}, \mu_t, \alpha_t) dt + dW_t, \quad t \in [0, T], \\ \mu_t & = \mathcal{L}(X_t^{\mu, \alpha}), \quad \forall t \in [0, T]. \end{cases}$$

Here $\mathcal{L}(Z)$ denotes the law of a random variable Z . We make the following assumptions on the coefficients:

Assumption B. The functions $b : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \rightarrow \mathbb{R}^d$, $f : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \rightarrow \mathbb{R}$, and $g : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ are bounded and jointly continuous, where $\mathcal{P}(\mathbb{R}^d)$ is endowed with the topology of weak convergence. The action space A is a compact metric space. Lastly, assume that b is Lipschitz in x , uniformly in the other variables; that is, there exists $L > 0$ such that, for all $(t, m, a) \in [0, T] \times \mathcal{P}(\mathbb{R}^d) \times A$ and all $x, y \in \mathbb{R}^d$, we have

$$|b(t, x, m, a) - b(t, y, m, a)| \leq L|x - y|.$$

Finally, the initial distribution $\lambda_0 \in \mathcal{P}(\mathbb{R}^d)$ is arbitrary.

Actually, all of the results of this section hold true without the Lipschitz assumption, but we include it in order to simplify one key proof step. This point is discussed briefly in Remark 5.12.

More completely, an equilibrium is defined as follows. For a measure $\mu \in \mathcal{P}(\mathcal{C}^d)$, we define $\mu_t \in \mathcal{P}(\mathbb{R}^d)$ for each $t \in [0, T]$ as the time- t marginal. That is, μ_t is the image of μ through the map $\mathcal{C}^d \ni x \mapsto x_t \in \mathbb{R}^d$.

Definition 5.1. A *Markovian mean field equilibrium (MFE)* is a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, \mu, W, \alpha, X)$ such that

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$.
2. $\mu \in \mathcal{P}(\mathcal{C}^d)$
3. W is an \mathbb{F} -Brownian motion.
4. X is a continuous \mathbb{F} -adapted process, and $X_0 \sim \lambda_0$.
5. $\alpha : [0, T] \times \mathbb{R}^d \rightarrow A$ is a measurable function.
6. The state equation holds,

$$dX_t = b(t, X_t, \mu_t, \alpha(t, X_t))dt + dW_t.$$

7. The control is optimal, in the sense that, for any other tuple of the form $(\Omega', \mathcal{F}', \mathbb{F}', \mathbb{P}', \mu, W', \alpha', X')$ satisfying (1-6) with the same μ , we have

$$\begin{aligned} & \mathbb{E}^P \left[\int_0^T f(t, X_t, \mu_t, \alpha(t, X_t))dt + g(X_T, \mu_T) \right] \\ & \geq \mathbb{E}^{P'} \left[\int_0^T f(t, X'_t, \mu_t, \alpha'(t, X'_t))dt + g(X'_T, \mu_T) \right]. \end{aligned}$$

8. The consistency condition holds, $\mu = \mathbb{P} \circ X^{-1}$.

Our goal in this section is to prove the following:

Theorem 5.2. *Suppose Assumption B holds. Assume that for each $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$ the set*

$$K(t, x, m) = \{(b(t, x, m, a), z) : a \in A, f(t, x, m, a) \geq z\} \subset \mathbb{R}^d \times \mathbb{R}$$

is convex. Then there exists a Markovian MFE.

The way we will prove this, once again, is by first proving existence of a *relaxed* MFE, and then arguing that we may massage any given equilibrium into a Markovian one.

5.1 Relaxed mean field equilibrium

A workable relaxed notion of equilibrium is defined as follows:

Definition 5.3. A *relaxed mean field equilibrium (MFE)* is a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, \mu, W, \Lambda, X)$ such that

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$.
2. $\mu \in \mathcal{P}(\mathcal{C}^d)$
3. W is an \mathbb{F} -Brownian motion.
4. X is a continuous \mathbb{F} -adapted process, and $X_0 \sim \lambda_0$.
5. Λ is a \mathcal{V} -valued random variable, and $\Lambda([0, s] \times S)$ is \mathcal{F}_t -measurable for all $s \leq t$ and all Borel sets $S \subset A$.
6. The state equation holds,

$$X_t = X_0 + \int_{[0, t] \times A} b(s, X_s, \mu_s, a) \Lambda(ds, da) + W_t, \quad t \in [0, T].$$

7. The control is optimal, in the sense that, for any other tuple of the form $(\Omega', \mathcal{F}', \mathbb{F}', \mathbb{P}', \mu, W', \Lambda', X')$ satisfying (1-6) with the same μ , we have

$$\begin{aligned} & \mathbb{E}^P \left[\int_{[0, T] \times A} f(t, X_t, \mu_t, a) \Lambda(dt, da) + g(X_T, \mu_T) \right] \\ & \geq \mathbb{E}^{P'} \left[\int_{[0, T] \times A} f(t, X'_t, \mu_t, a) \Lambda'(dt, da) + g(X'_T, \mu_T) \right]. \end{aligned}$$

8. The consistency condition holds, $\mu = \mathbb{P} \circ X^{-1}$.

The first main theorem is the following, which proves existence of optimal relaxed controls in a rather general setting:

Theorem 5.4. *Under Assumption **B**, there exists a relaxed MFE.*

To prove Theorem **B**, we first reformulate the problem as in Section 4.1, in terms of the probability law of the state and control processes. Again, we work on the canonical measurable space $\Omega = \mathcal{C}^d \times \mathcal{V}$, equipped with its Borel σ -field \mathcal{F} . We define *canonical processes* X and Λ , simply by setting

$X(x, q) = x$ and $\Lambda(x, q) = q$ for $(x, q) \in \Omega$, and similarly $X_t(x, q) = x_t$ for $t \in [0, T]$. We endow Ω with the natural filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$, where \mathcal{F}_t is the σ -field generated by the following the random variables X_s and $\Lambda([0, s] \times B)$, where $s \leq t$ and $B \subset A$ is a Borel set.

Recall that the initial state distribution $\lambda_0 \in \mathcal{P}(\mathbb{R}^d)$ is fixed throughout.

Definition 5.5. For $\mu \in \mathcal{P}(\mathcal{C}^d)$, let $\mathcal{R}(\mu)$ denote the set of probability measures P on $\Omega = \mathcal{C}^d \times \mathcal{V}$ such that:

1. $P \circ X_0^{-1} = \lambda_0$.
2. Under P , the process $W = (W_t)_{t \in [0, T]}$ defined by

$$W_t^\mu = X_t - X_0 - \int_{[0, t] \times A} b(s, X_s, \mu_s, a) \Lambda(ds, da) \quad (5.1)$$

is a \mathbb{F} -Brownian motion. Equivalently, $P \circ W^{-1}$ equals Wiener measure, and it holds that $W_s - W_t$ is independent of \mathcal{F}_t under P whenever $0 \leq t \leq s \leq T$.

Define the function $\Gamma : \mathcal{P}(\mathcal{C}^d) \times \mathcal{C}^d \times \mathcal{V} \rightarrow \mathbb{R}$ by

$$\Gamma(\mu, x, q) := \int_{[0, T] \times A} f(t, x_t, \mu_t, a) q(dt, da) + g(x_T).$$

Finally, define $\mathcal{R}^*(\mu)$ as the set of $P^* \in \mathcal{R}(\mu)$ such that

$$\int_{\Omega} \Gamma(\mu, \cdot) dP^* = \sup_{P \in \mathcal{R}(\mu)} \int_{\Omega} \Gamma(\mu, \cdot) dP,$$

or more briefly $\mathcal{R}^*(\mu) := \arg \max_{P \in \mathcal{R}(\mu)} \int_{\Omega} \Gamma(\mu, \cdot) dP$.

The point of these definitions is that a mean field equilibrium is nothing but a fixed point of the set-valued map $\mathcal{P}(\mathcal{C}^d) \ni \mu \mapsto \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\} \subset \mathcal{P}(\mathcal{C}^d)$.

Lemma 5.6. *If $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}}, \tilde{\mathbb{P}}, \tilde{W}, \mu, \tilde{\Lambda}, \tilde{X})$ is a relaxed MFE if and only if $\mu \in \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\}$. More specifically, if $\mu = P \circ X^{-1}$ for some $P \in \mathcal{R}^*(\mu)$, then the tuple $(\Omega, \mathcal{F}, \mathbb{F}, P, W^\mu, \mu, \Lambda, X)$ is a relaxed MFE, where W^μ is defined as in (5.1).*

Exercise 5.7. Prove Lemma 5.6.

The goal now is to show the following:

Theorem 5.8. *There exists $\mu \in \mathcal{P}(\mathcal{C}^d)$ such that $\mu \in \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\}$.*

Notice that Theorem 5.4 follows immediately from Theorem 5.8 and Lemma 5.6.

The proof of Theorem 5.8 follows a well known line of argument in game theory, indeed dating back to the original existence proof of Nash [37]. The idea is to apply *Kakutani's fixed point theorem*, which we quote without proof from [1, Corollary 17.55].

Theorem 5.9 (Kakutani's fixed point theorem). *Let K be a nonempty compact convex subset of a locally convex Hausdorff topological vector space. Suppose a map $\Psi : K \rightarrow 2^K$ has the following properties:*

- (i) $\Psi(x)$ is nonempty and convex for each $x \in K$.
- (ii) The graph $\{(x, y) : x \in K, y \in \Psi(x)\}$ is closed.

Then Ψ has a fixed point. That is, there exists $x \in K$ such that $x \in \Psi(x)$.

We wish to apply this to the map

$$\Psi(\mu) := \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\}, \quad \Psi : \mathcal{P}(\mathcal{C}^d) \rightarrow 2^{\mathcal{P}(\mathcal{C}^d)}. \quad (5.2)$$

Our ambient topological vector space will of course be the set of bounded signed measures on \mathcal{C}^d , of which $\mathcal{P}(\mathcal{C}^d)$ is a convex subset. The following lemma designs an appropriate compact convex set $K \subset \mathcal{P}(\mathcal{C}^d)$.

Lemma 5.10. *Let K_0 denote the set of probability measures $\mu \in \mathcal{P}(\mathcal{C}^d)$ such that $\mu_0 = \lambda_0$ and*

$$\mathbb{E}^\mu |X_{\tau+\delta} - X_\tau| \leq \delta \|b\|_\infty + \sqrt{\delta},$$

for any $\delta > 0$ and any stopping time τ with values in $[0, T - \delta]$. Then the following hold:

- (i) $\bigcup_{\mu \in \mathcal{P}(\mathcal{C}^d)} \mathcal{R}(\mu) \subset K_0$.
- (ii) K_0 is tight.
- (iii) K_0 is convex.

In particular, the closure $K = \overline{K_0}$ is compact and convex.

Proof. Claim (i) follows from the same argument which led to (4.3) in the proof of Theorem 4.11. Claim (ii) follows from Aldous' criterion for tightness, Theorem 4.12. Claim (iii) is straightforward to check, as K_0 is defined in terms of constraints which are linear in the measure μ . \square

The next step is to show that the map Ψ defined in (5.2) satisfies the two key hypotheses of Kakutani's theorem. This is done with the help of another famous theorem, widely used in game theory, and typically attributed to Berge. We quote the result from [1, Theorem 17.31]:

Theorem 5.11 (Berge's Maximum Theorem). *Suppose \mathcal{X} and \mathcal{Y} are metric spaces, with \mathcal{Y} compact. Let $F : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$ and $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ be given. Suppose that f is jointly continuous, and F satisfies the following two properties:*

- (i) *The graph $\text{Gr}(F) := \{(x, y) : x \in \mathcal{X}, y \in F(x)\}$ is closed.*
- (ii) *F is lower hemicontinuous, which means the following: Given $x_n \rightarrow x$ in \mathcal{X} , and given $y \in F(x)$, we may find $n_1 < n_2 < \dots$ and $y_{n_k} \in F(x_{n_k})$ such that $y_{n_k} \rightarrow y$.*

Define $v : \mathcal{X} \rightarrow \mathbb{R}$ and $F^* : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$ by setting

$$v(x) = \sup_{y \in F(x)} f(x, y), \quad F^*(x) = \arg \max_{y \in F(x)} f(x, y).$$

Then v is continuous, $F^*(x) \neq \emptyset$ for all $x \in \mathcal{X}$, and the graph $\text{Gr}(F^*) := \{(x, y) : x \in \mathcal{X}, y \in F^*(x)\}$ is closed.

We are now ready for the main proof:

Proof of Theorem 5.8. Recall the definition of the set-valued map $\mathcal{R} : \mathcal{P}(\mathcal{C}^d) \rightarrow \mathcal{P}(\Omega)$, given by

$$\mathcal{R}^*(\mu) = \arg \max_{P \in \mathcal{R}(\mu)} \int_{\Omega} \Gamma(\mu, \cdot) dP.$$

We would like to apply Berge's Theorem 5.11 to conclude that $\mathcal{R}^*(\mu) \neq \emptyset$ for each μ and that the graph $\text{Gr}(\mathcal{R}^*) := \{(\mu, P) : \mu \in \mathcal{X}, P \in \mathcal{R}^*(\mu)\}$ is closed. Knowing this, we complete the proof as follows. Because $\int_{\Omega} \Gamma(\mu, \cdot) dP$ is linear in P , the set of maximizers $\mathcal{R}^*(\mu)$ must be convex for each μ . If we then define Ψ as in (5.2), it is readily checked using the fact that $P \mapsto P \circ X^{-1}$ is linear and continuous that $\Psi(\mu)$ is convex and nonempty for each μ and that the graph $\text{Gr}(\Psi)$ is closed. Hence, Kakutani's Theorem 5.9 yields the existence of a fixed point.

It remains to check the hypotheses of Berge's Theorem 5.11. It follows from Lemma 2.9 that $\int_{\Omega} \Gamma(\mu, \cdot) dP$ is jointly continuous in (μ, P) . It remains to check that \mathcal{R} satisfies the hypotheses (i) and (ii) of Berge's Theorem hold.

- (i) Suppose $(\mu^n, P^n) \rightarrow (\mu, P)$, with $(\mu^n, P^n) \in \text{Gr}(\mathcal{R})$ and (μ, P) in $\mathcal{P}(\mathcal{C}^d) \times \mathcal{P}(\Omega)$. Note that $(\mu^n, P^n) \in \text{Gr}(\mathcal{R})$ simply means $\mu^n \in \mathcal{P}(\mathcal{C}^d)$ and $P^n \in \mathcal{R}(\mu^n)$. Thus $P^n \circ X_0^{-1} = \lambda_0$ for each n , and we conclude $P \circ X_0^{-1} = \lambda_0$. Finally, we must check that W^μ is an \mathbb{F} -Brownian motion under P . First, notice that the map $\widehat{W} : \mathcal{P}(\mathcal{C}^d) \times \Omega \rightarrow \mathcal{C}^d$ defined by

$$\widehat{W}(\nu, x, q)_t := x_t - x_0 - \int_{[0,t] \times A} b(s, x_s, \nu_s, a) q(ds, da) \quad (5.3)$$

is jointly continuous (see Lemma 2.9). It then follows from Lemma 2.10 that

$$P \circ (W^\mu)^{-1} = P \circ \widehat{W}(\mu, \cdot)^{-1} = \lim_n P^n \circ \widehat{W}(\mu^n, \cdot)^{-1}$$

equals Wiener measure. It remains to show that $W_s^\mu - W_t^\mu$ is independent of \mathcal{F}_t , for $0 \leq t \leq s \leq T$. To do this, fix $t < s$, a bounded continuous function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$, and a bounded, \mathcal{F}_t -measurable, and continuous function $h : \Omega \rightarrow \mathbb{R}$. By continuity and the fact that $P_n \in \mathcal{R}$ for each n , we have (implicitly using Lemma 2.9)

$$\begin{aligned} \mathbb{E}^P [\varphi(W_s^\mu - W_t^\mu) h(X, \Lambda)] &= \lim_n \mathbb{E}^{P_n} [\varphi(W_s^{\mu^n} - W_t^{\mu^n}) h(X, \Lambda)] \\ &= \lim_n \mathbb{E}^{P_n} [\varphi(W_s^{\mu^n} - W_t^{\mu^n})] \mathbb{E}^{P_n} [h(X, \Lambda)] \\ &= \mathbb{E}^P [\varphi(W_s^\mu - W_t^\mu)] \mathbb{E}^P [h(X, \Lambda)]. \end{aligned}$$

This is enough to conclude the desired independence.

- (ii) Fix $\mu_n \rightarrow \mu$ in $\mathcal{P}(\mathcal{C}^d)$ and $P \in \mathcal{R}(\mu)$. Note that under P we have

$$dX_t = \int_A b(t, X_t, \mu_t, a) \Lambda_t(da) dt + dW_t,$$

where $W := W^\mu$ is a Brownian motion. Now, Assumption **B** ensures that b is Lipschitz in x . Hence, we may find a unique strong solution of the SDE

$$dX_t^n = \int_A b(t, X_t^n, \mu_t^n, a) \Lambda_t(da) dt + dW_t, \quad X_0^n = X_0,$$

constructed on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$. It is then clear that $P^n := P \circ (X^n, \Lambda)^{-1}$ belongs to $\mathcal{R}(\mu^n)$ for each n . We complete

the proof by checking that $P^n \rightarrow P$ weakly. A standard argument shows $X^n \rightarrow X$ almost surely in \mathcal{C}^d . Indeed, note that

$$\begin{aligned}
|X_t^n - X_t| &\leq \int_{[0,t] \times A} |b(s, X_s^n, \mu_s^n, a) - b(s, X_s, \mu_s, a)| \Lambda(ds, da) \\
&\leq L \int_{[0,t] \times A} |X_s^n - X_s| \Lambda(ds, da) \\
&\quad + \int_{[0,t] \times A} |b(s, X_s, \mu_s^n, a) - b(s, X_s, \mu_s, a)| \Lambda(ds, da) \\
&= L \int_0^t |X_s^n - X_s| ds \\
&\quad + \int_{[0,t] \times A} |b(s, X_s, \mu_s^n, a) - b(s, X_s, \mu_s, a)| \Lambda(ds, da),
\end{aligned}$$

and then use Gronwall's inequality to get

$$\begin{aligned}
&\mathbb{E}^P \sup_{t \in [0, T]} |X_t^n - X_t| \\
&\leq L e^{LT} \mathbb{E}^P \int_{[0, T] \times A} |b(s, X_s, \mu_s^n, a) - b(s, X_s, \mu_s, a)| \Lambda(ds, da).
\end{aligned}$$

By dominated convergence and continuity of b in the measure argument, the right-hand side converges to zero. We conclude that $P^n = P \circ (X^n, \Lambda)^{-1}$ converges to $P \circ (X, \Lambda)^{-1} = P$.

□

Remark 5.12. The only place we used the Lipschitz part of Assumption **B** is in this last step, to prove lower hemicontinuity. Reflecting on the proof, what is really needed is some form of *weak uniqueness*. That is, for any choice of μ and $P \in \mathcal{R}(\mu)$, we want to be able to exhibit P as the *unique* solution of some SDE, for which suitable stability properties let us prove convergence. Annoyingly, the controlled SDEs we encounter in this way have random coefficients, and the weak solution theory for such equations is not incredibly well-developed. But there are other ways to make this work, and we Remark that Theorems 5.2, 5.8, and 5.4 all hold without the Lipschitz assumption, albeit with a trickier proof. See [29, Theorem 6.2].

5.2 Markovian equilibria

Using a relaxed equilibrium, the existence of which we now know from Theorem 5.4, we can use it to construct a Markovian equilibrium, thus proving Theorem 5.2. We do this by adapting the argument of Section 4.2.

Proof of Theorem 5.2. By Theorem 5.8, there exist $(\mu, P) \in \mathcal{P}(\mathcal{C}^d) \times \mathcal{P}(\mathcal{C}^d \times \mathcal{V})$ such that $P \circ X^{-1} = \mu$ and $P \in \mathcal{R}^*(\mu)$. By Theorem 4.13, we can find a Markovian strict optimal control $P^0 \in \mathcal{R}^*(\mu)$ such that $P^0 \circ X_t^{-1} = P \circ X_t^{-1} = \mu_t$ for each $t \in [0, T]$. Next note that all of the coefficients (b, f, g) depend on μ only through the marginal flow $(\mu_t)_{t \in [0, T]}$. Hence, if we define $\mu^0 := P^0 \circ X^{-1}$, then we have $\mu_t^0 = \mu_t$ for all $t \in [0, T]$, and thus $\mathcal{R}(\mu) = \mathcal{R}(\mu^0)$ and $\mathcal{R}^*(\mu) = \mathcal{R}^*(\mu^0)$. Hence, $P^0 \in \mathcal{R}^*(\mu^0)$, and we conclude that μ^0 is a mean field equilibrium. \square

References

- [1] C. Aliprantis and K. Border, *Infinite dimensional analysis: A hitchhiker's guide*, 3rd ed., Springer, 2007.
- [2] H. Becker and V. Mandrekar, *On the existence of optimal random controls*, Journal of mathematics and mechanics **18** (1969), no. 12, 1151–1166.
- [3] C. Benazzoli, L. Campi, and L. Di Persio, *Mean-field games with controlled jumps*, arXiv preprint arXiv:1703.01919 (2017).
- [4] V.E. Beneš, *Existence of optimal stochastic control laws*, SIAM Journal on Control **9** (1971), no. 3, 446–472.
- [5] P. Billingsley, *Convergence of probability measures*, John Wiley & Sons, 2013.
- [6] G. Brunick and S. Shreve, *Mimicking an Itô process by a solution of a stochastic differential equation*, The Annals of Applied Probability **23** (2013), no. 4, 1584–1628.
- [7] R. Carmona, F. Delarue, and D. Lacker, *Mean field games with common noise*, The Annals of Probability **44** (2016), no. 6, 3740–3803.
- [8] R.M. Dudley, *Real analysis and probability*, Chapman and Hall/CRC, 2018.
- [9] N. El Karoui, D. Nguyen, and M. Jeanblanc-Picqué, *Compactification methods in the control of degenerate diffusions: existence of an optimal control*, Stochastics: an international journal of probability and stochastic processes **20** (1987), no. 3, 169–219.
- [10] N. El Karoui, D.H. Nguyen, and M. Jeanblanc-Picqué, *Existence of an optimal Markovian filter for the control under partial observations*, SIAM journal on control and optimization **26** (1988), no. 5, 1025–1061.
- [11] N. El Karoui and X. Tan, *Capacities, measurable selection and dynamic programming part I: abstract framework*, arXiv preprint arXiv:1310.3363 (2013).
- [12] ———, *Capacities, measurable selection and dynamic programming part II: application in stochastic control problems*, arXiv preprint arXiv:1310.3364 (2013).
- [13] A.F. Filippov, *On certain questions in the theory of optimal control*, Journal of the Society for Industrial & Applied Mathematics, Series A: Control **1** (1962), no. 1, 76–84.
- [14] M. Fischer, *On the connection between symmetric n -player games and mean field games*, The Annals of Applied Probability **27** (2017), no. 2, 757–810.
- [15] W.H. Fleming, *Generalized solutions in optimal stochastic control*, 1976.

- [16] W.H. Fleming and M. Nisio, *On the existence of optimal stochastic controls*, Journal of Mathematics and Mechanics **15** (1966), no. 5, 777–794.
- [17] G. Fu and U. Horst, *Mean field games with singular controls*, SIAM Journal on Control and Optimization **55** (2017), no. 6, 3833–3868.
- [18] I. Gyöngy, *Mimicking the one-dimensional marginal distributions of processes having an Itô differential*, Probability theory and related fields **71** (1986), no. 4, 501–516.
- [19] U.G. Haussmann and J.P. Lepeltier, *On the existence of optimal controls*, SIAM Journal on Control and Optimization **28** (1990), no. 4, 851–902.
- [20] U.G. Haussmann and W. Suo, *Singular optimal stochastic controls I: Existence*, SIAM Journal on Control and Optimization **33** (1995), no. 3, 916–936.
- [21] ———, *Singular optimal stochastic controls II: Dynamic programming*, SIAM Journal on Control and Optimization **33** (1995), no. 3, 937–959.
- [22] M. Huang, R. Malhamé, and P. Caines, *Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle*, Communications in Information & Systems **6** (2006), no. 3, 221–252.
- [23] O. Kallenberg, *Foundations of modern probability*, Springer Science & Business Media, 2006.
- [24] I. Karatzas and S.E. Shreve, *Brownian motion and stochastic calculus*, Graduate Texts in Mathematics, Springer New York, 1991.
- [25] T.G. Kurtz and R.H. Stockbridge, *Existence of Markov controls and characterization of optimal Markov controls*, SIAM Journal on Control and Optimization **36** (1998), no. 2, 609–653.
- [26] H.J. Kushner, *Existence results for optimal stochastic controls*, Journal of Optimization Theory and Applications **15** (1975), no. 4, 347–359.
- [27] ———, *Numerical methods for stochastic control problems in continuous time*, SIAM Journal on Control and Optimization **28** (1990), no. 5, 999–1048.
- [28] H.J. Kushner and P. Dupuis, *Numerical methods for stochastic control problems in continuous time*, Vol. 24, Springer Science & Business Media, 2013.
- [29] D. Lacker, *Mean field games via controlled martingale problems: existence of Markovian equilibria*, Stochastic Processes and their Applications **125** (2015), no. 7, 2856–2894.
- [30] ———, *A general characterization of the mean field limit for stochastic differential games*, Probability Theory and Related Fields **165** (2016), no. 3-4, 581–648.
- [31] ———, *Limit theory for controlled McKean–Vlasov dynamics*, SIAM Journal on Control and Optimization **55** (2017), no. 3, 1641–1672.
- [32] ———, *On the convergence of closed-loop Nash equilibria to the mean field game limit*, Work in progress (2018).
- [33] J.-M. Lasry and P.-L. Lions, *Jeux à champ moyen. i—le cas stationnaire*, Comptes Rendus Mathématique **343** (2006), no. 9, 619–625.
- [34] ———, *Jeux à champ moyen. ii—horizon fini et contrôle optimal*, Comptes Rendus Mathématique **343** (2006), no. 10, 679–684.
- [35] ———, *Mean field games*, Japanese journal of mathematics **2** (2007), no. 1, 229–260.

- [36] E.J. McShane, *Relaxed controls and variational problems*, SIAM Journal on Control **5** (1967), no. 3, 438–485.
- [37] J.F. Nash, *Equilibrium points in n -person games*, Proceedings of the national academy of sciences **36** (1950), no. 1, 48–49.
- [38] K.R. Parthasarathy, *Probability measures on metric spaces*, Vol. 352, American Mathematical Soc., 2005.
- [39] H. Pham, *Continuous-time stochastic control and optimization with financial applications*, Vol. 61, Springer Science & Business Media, 2009.
- [40] E. Roxin, *The existence of optimal controls.*, The Michigan Mathematical Journal **9** (1962), no. 2, 109–119.
- [41] L.C. Young, *Generalized curves and the existence of an attained absolute minimum in the calculus of variations*, Comptes Rendus de la Societe des Sci. et des Lettres de Varsovie **30** (1937), 212–234.
- [42] ———, *Lecture on the calculus of variations and optimal control theory*, Vol. 304, American Mathematical Soc., 1980.