



UNIVERSITÀ DEGLI STUDI DI MILANO
FACOLTÀ DI SCIENZE E TECNOLOGIE

Dipartimento di Matematica “Federigo Enriques”

Corso di Dottorato di Ricerca in Scienze Matematiche

Ciclo XXXVII

Tesi di Dottorato di Ricerca

**Coarse correlated equilibria in continuous-time
mean field games**

Settori scientifico-disciplinari: MATH-03/B, STAT-04/A

Relatore:
Prof. Luciano Campi

Coordinatore del dottorato:
Prof. Giulio Ciraolo

Candidato:
Federico Cannerozzi

A. A. 2023-2024

Abstract

Mean field games (MFGs), an active theme of research since the mid 2000's, arise as the limit formulation of symmetric N -player games with mean field interactions, and they act as a treatable approximation of such games, when the numbers of players is large enough. The usual notion of MFG solution can be regarded as the infinitely many players counterpart of the concept of Nash equilibrium in the underlying N -player game, and the connection between Nash equilibria and solutions to the MFG can be made rigorous via convergence arguments.

The goal of this thesis is to introduce more general equilibria with better properties in continuous-time MFGs. In game theory literature, coarse correlated equilibria (CCEs) have been proposed as an alternative to Nash equilibria. CCEs can be seen as generalizations of the Nash equilibrium concept, incorporating a correlation device that allows agents to adopt correlated strategies without requiring cooperation. CCEs generalize Aumann's concept of correlated equilibria as well as Nash equilibria in both pure and mixed strategies. Among the nice features of CCEs, they are shown to be able to lead to higher payoffs than Nash equilibria in standard game theory, to be computationally "easier" and more natural to be learnt by the players.

In this thesis, we introduce CCEs in continuous-time MFGs and study their properties. In each chapter a different framework is considered and different problems are addressed, each requiring specific methodologies. In particular, the first chapter presents abstract results on the existence of CCEs in MFGs and on the approximation of CCEs in the underlying N -player games by means of the former ones. Subsequent chapters focus on detailed studies of specific MFG models, in which it is possible to analytically compute CCEs (or at least some of them) and compare them with more classical solution concepts.

In Chapter 1 we consider MFGs driven by additive Wiener noise, with general drift and cost functions. The interaction term is given by a flow of probability measures, which appears both in the drift and in the cost functions. We introduce CCEs in both continuous time stochastic differential games and MFGs. The notion of coarse correlated solution to the MFG is justified by proving an approximation result. An existence result is also presented, whose proof relies on a minimax theorem.

In Chapter 2 we consider linear-quadratic MFGs. The interaction term is given by a flow of first order moments, which appears in the payoff functional only. A methodology to compute CCEs in such class of MFGs is provided and, through the study of a simple yet important example with applications in environmental economics, we show that there exist infinitely many CCEs for the MFG which both yield higher payoffs than the classical MFG solutions and are more efficient with respect to the environmental goals. Moreover, we provide instances of MFGs that do not admit any

MFG solution, whereas infinitely many CCEs exist.

In Chapter 3 we consider a simple class of stationary MFGs of singular control. The reward, of ergodic type, is given by the long-time average of an expected utility functional. The interaction term is given by the stationary mean of the distribution of the representative player, which appears only in the reward. We provide constructive existence results, as well as approximation results and comparison with both MFG solutions and mean field control solutions. Finally, we show that CCEs may exist even when MFG solutions do not.

Acknowledgements

As my PhD studies come to an end, there are many people I need to thank for their help, guidance, and support. First and foremost, I would like to express my gratitude to my supervisor, Prof. Luciano Campi. Luciano has been everything I could have asked for in a supervisor, providing invaluable guidance and help throughout my PhD studies, as well as engaging in the insightful discussions that led to the research presented in this thesis. I would also like to extend my sincere thanks to all my co-authors, whose collaboration has been both inspiring and enjoyable: Prof. Markus Fischer, Dr. Fanny Cartellier, and Prof. Giorgio Ferrari. I thank the Referees for their careful reading of this thesis and for their positive feedback, which exceeded my expectations. Special thanks go to Prof. Marco Fuhrman and Prof. Marco Maggis for their advice, which began even before my PhD. I am also grateful to Prof. Daniele Cavicchioli and Prof. Fulvia Confortola for the opportunity and trust they placed in me by allowing me to conduct the exercise sessions for their courses. I thank the Center for Mathematical Economics at the University of Bielefeld for welcoming me during the months I spent there and for making me feel like I was already part of the group. I also want to thank the other PhD students at the University of Milan for sharing both the good and challenging moments of these past years with me. Finally, I am enormously grateful to my family, whose fundamental support and patience made this possible: my father and my mother, my brothers and my sister, and, most importantly, Cristina. Chissà dove finirei senza.

Contents

| | |
|---|-----------|
| Abstract | 3 |
| Notation | 9 |
| Introduction | 11 |
| I.1 Chapter 1: MFGs with general dynamics and costs | 15 |
| I.2 Chapter 2: Linear-quadratic MFGs | 20 |
| I.3 Chapter 3: Stationary MFGs of singular control | 24 |
| 1 Mean field games with general dynamics and costs | 29 |
| 1.1 Standing assumptions | 29 |
| 1.2 Formulation of the N -player game | 30 |
| 1.3 Formulation of the mean field game | 36 |
| 1.4 Approximate N -player coarse correlated equilibria | 42 |
| 1.4.1 Construction of the admissible recommendation profiles to the N -player game | 43 |
| 1.4.2 Proof of Theorem 1.4.1 | 46 |
| 1.5 Existence of a coarse correlated solution of the mean field game | 49 |
| 1.5.1 The auxiliary zero-sum game | 50 |
| 1.5.2 Proof of Theorem 1.5.1 | 53 |
| 1.5.3 Proof of Theorem 1.5.2 | 55 |
| 1.6 An example of coarse correlated solution to the mean field game | 64 |
| 1.6.1 Exhibiting explicit coarse correlated solutions | 65 |
| 1.6.2 Comparison with weak mean field game solutions without com- mon noise of [85] | 68 |
| 1.7 Auxiliary results | 71 |
| 1.7.1 On admissible recommendations | 71 |
| 1.7.2 Propagation of chaos | 73 |
| 1.7.3 Further technical lemmata for the existence of mean field CCEs | 78 |
| 2 Linear-quadratic mean field games | 85 |
| 2.1 Standing assumptions | 85 |
| 2.2 Mean field coarse correlated equilibria for the linear-quadratic MFG | 86 |
| 2.3 Computing mean field coarse correlated equilibria | 88 |
| 2.3.1 Deviating player's optimization problem | 89 |
| 2.3.2 Correlated moment flow | 91 |
| 2.3.3 Optimality condition | 92 |

| | | |
|----------|---|------------|
| 2.4 | Comparison with MFC solution and mean field NE | 95 |
| 2.4.1 | Comparison with MFC solution | 96 |
| 2.4.2 | Comparison with mean field Nash equilibria | 101 |
| 2.5 | Application to an emission abatement game | 103 |
| 2.5.1 | Translation and interpretation of findings in the abatement game | 104 |
| 2.5.2 | A tractable class of mean field CCEs | 107 |
| 2.5.3 | Comparison with mean field NE | 109 |
| 2.6 | A mean field game that does not admit mean field NEs | 114 |
| 2.7 | Auxiliary results: Some standard proofs | 117 |
| 3 | Stationary mean field games of singular control | 125 |
| 3.1 | The N -player game | 125 |
| 3.2 | The ergodic mean field game | 129 |
| 3.3 | Assumptions and preliminary results | 132 |
| 3.4 | Competitive case: Mean field equilibria and approximation | 135 |
| 3.4.1 | The deviating player problem | 136 |
| 3.4.2 | Regular recommendation | 137 |
| 3.4.3 | Singular recommendation | 141 |
| 3.5 | Cooperative case: Mean field solution and approximation | 145 |
| 3.5.1 | Mean field control solution | 145 |
| 3.5.2 | Approximation | 148 |
| 3.6 | Numerical illustrations | 151 |
| 3.7 | Auxiliary results | 154 |
| A | Relaxed controls | 163 |
| B | Weak and strong existence for controlled equations | 165 |

Notation

General notation

For $d \in \mathbb{N}$ with $d \geq 1$ and $x, y \in \mathbb{R}^d$, we denote by $\langle x, y \rangle$ the scalar product in \mathbb{R}^d , as well as by $|\cdot|$ the Euclidean norm in \mathbb{R}^d .

For $d, N \in \mathbb{N}$, $d, N \geq 1$, and $v = (v^1, \dots, v^N) \in \mathbb{R}^{Nd}$, for each $i = 1, \dots, N$, set $v^{-i} = (v^1, \dots, v^{i-1}, v^{i+1}, \dots, v^N) \in \mathbb{R}^{(N-1)d}$ and, for $a \in \mathbb{R}^d$, set

$$(v^{-i}, a) = (v^1, \dots, v^{i-1}, a, v^{i+1}, \dots, v^N) \in \mathbb{R}^{Nd}.$$

For $d, k \in \mathbb{N}$, $d, k \geq 1$, we denote by $\mathbb{R}^{d \times k}$ the set of $d \times k$ matrices with real-valued entries. For any matrix $A \in \mathbb{R}^{d \times k}$, we denote by A^\top its transpose. We denote by $|A|$ the maximum norm of A , i.e. $|A| = \max_{i,j} |a_{i,j}|$. We denote by \mathcal{S}^d the set of $d \times d$ symmetric matrices and by I_d the identity matrix in \mathcal{S}^d .

Unless otherwise stated, C indicates a generic positive constant, which may change from line to line.

Functional and measure spaces

For a metric space (E, d_E) , we denote by \mathcal{B}_E the Borel σ -algebra generated by the topology of E . When the context allows, we will drop the dependence upon E , and just denote it by \mathcal{B} . We denote by $\mathcal{C}_b(E)$ the set of continuous bounded function $f : E \rightarrow \mathbb{R}$.

Given a measure space (M, \mathcal{M}, μ) and $B \subseteq \mathbb{R}$, we denote by $L^2(M, \mathcal{M}, \mu; B)$ the set of measurable functions $f : M \rightarrow B$ so that $\int_M |f(m)|^2 \mu(dm) < \infty$, and we define its L^2 -norm by $\|f\|_{L^2} = (\int_M |f(m)|^2 \mu(dm))^{1/2}$. As usual, we identify functions f_1 and f_2 which are equal μ -a.s..

We will denote by $\mathcal{P}(E)$ the set of probability measures on (E, \mathcal{B}_E) . For $p \geq 1$, we denote by $\mathcal{P}^p(E)$ the set of probability measures $\mu \in \mathcal{P}(E)$ so that, for some point $x_0 \in E$, and thus for any, it holds $\int_E d_E^p(x, x_0) \mu(dx) < \infty$. Let $\mathcal{W}_{p,E}(\mu_1, \mu_2)$ denote the p -Wasserstein distance on $\mathcal{P}^p(E)$, defined as

$$\begin{aligned} \mathcal{W}_{p,E}^p(\mu_1, \mu_2) \\ = \inf \left\{ \int_{E \times E} d_E^p(x, y) \pi(dx, dy) : \pi \in \mathcal{P}(E \times E), \pi \text{ has marginals } \mu_1, \mu_2 \right\}. \end{aligned}$$

Any time we will be given two metric spaces (E, d_E) and $(E', d_{E'})$, we will regard $E \times E'$ as a metric space itself, with the distance $d((e, f), (e', f')) = d_E(e, f) + d_{E'}(e', f')$. The p -Wasserstein distance on $\mathcal{P}^p(E \times E')$ will always be meant with respect to such distance on $E \times E'$. When $E = \mathbb{R}^d$, for any probability measure $\mu \in \mathcal{P}^1(\mathbb{R}^d)$, we set $\bar{\mu} = \int_{\mathbb{R}^d} y \mu(dy)$ to denote its first order moment.

For $T > 0$ fixed, we denote by \mathcal{C}^d the set of continuous functions from $[0, T]$ in \mathbb{R}^d , $d \in \mathbb{N}$, i.e. $\mathcal{C}^d = \mathcal{C}([0, T]; \mathbb{R}^d)$. We endow \mathcal{C}^d with the norm $\|x\|_{\mathcal{C}^d} = \sup_{s \in [0, T]} |x_s|$. Occasionally, we will use the semi-norm $\|x\|_{t, \mathcal{C}^d} = \sup_{s \in [0, t]} |x_s|$, for $x \in \mathcal{C}^d$. We will denote as $\mathbb{W}^d \in \mathcal{P}(\mathcal{C}^d)$ the law of a standard d -dimensional Brownian motion, and by $\mathcal{C}(\mathcal{P}^2)$ the set of continuous functions from $[0, T]$ in $\mathcal{P}^2(\mathbb{R}^d)$, i.e. $\mathcal{C}(\mathcal{P}^2) = \mathcal{C}([0, T]; \mathcal{P}^2(\mathbb{R}^d))$, where $\mathcal{P}^2(\mathbb{R}^d)$ is endowed with the 2-Wasserstein distance. We endow $\mathcal{C}(\mathcal{P}^2)$ with the supremum distance $\sup_{t \in [0, T]} \mathcal{W}_{2, \mathcal{P}^2(\mathbb{R}^d)}(\mu_t^1, \mu_t^2)$, for any $\mu^1 = (\mu_t^1)_{t \in [0, T]}$ and $\mu^2 = (\mu_t^2)_{t \in [0, T]}$ in $\mathcal{C}(\mathcal{P}^2)$.

Probability

When given a filtered probability space $(\Omega, \mathcal{F}, \mathbb{G}, \mathbb{P})$, we regard as the \mathbb{P} -augmentation of the filtration $\mathbb{G} = (\mathcal{G}_t)_t$ the filtration $\mathbb{F} = (\mathcal{F}_t)_t$, where $\mathcal{F}_t = \bigcap_{\varepsilon > 0} \sigma(\mathcal{G}_{t+\varepsilon}, \mathcal{N})$ and \mathcal{N} stands for the \mathbb{P} -null sets of Ω . Such a filtration satisfies the usual assumptions.

Given an arbitrary filtration \mathbb{G} , we will use the notation $\mathbb{H}^2(\mathbb{G})$ for the set of all \mathbb{G} -progressively measurable \mathbb{R}^k -valued processes $\alpha = (\alpha_t)_{t \in [0, T]}$ such that $\mathbb{E}[\int_0^T |\alpha_t|^2 dt] < \infty$.

Introduction

Mean field games (MFGs) have been an active theme of research for almost two decades, started in the mid 2000's from the seminal works of Lasry and Lions [89] and of Huang, Malhamé and Caines [76]. Roughly speaking, MFGs arise as the limit formulation of symmetric N -player games with mean field interactions between the players, and they act as a treatable approximation of such games, when the numbers of players N is large enough. Indeed, despite the symmetry and the mean-field interaction properties of these systems, stochastic differential games with many players are hardly tractable, except in few specific cases (see, e.g. [43]). Thus, the urge to consider a tractable approximation arises. Thanks to the mean field interaction and propagation of chaos type results (see, e.g., the seminal work of Sznitman [104]), one expects that, as the number of players grows to infinity, the empirical distribution of players' states converges to the law of some representative player. Then, when studying a MFG, one studies the behaviour of such representative player of infinitely many, as she interacts with the limit distribution of the population.

The study of MFGs has been tackled in various ways: analytically (see, e.g., [18, 34, 35, 89]), by weak formulation (e.g. [28, 44]), by relaxed controls (e.g [84]), by linear programming approach (see [22, 58, 59]), or by the Pontryagin maximum principle (e.g. [36, 69]). For a detailed presentation of the probabilistic approach to MFGs, we refer to the two-volume book by Carmona and Delarue [38, 39].

To find solutions to MFGs, most of the literature relies on the following scheme: firstly, the representative player reacts optimally to a fixed flow of measures; then, the flow of measures arises as aggregation of all such identical players' best responses at equilibrium. This two-step procedure is the translation, in the limit of infinitely many players, of the concept of Nash equilibrium (NE) in the underlying N -player game. The connection between Nash equilibria and solutions to the MFG is made rigorous in two ways: on the one hand, a solution of the MFG allows to construct approximate Nash equilibria for the corresponding N -player games, if N is sufficiently large, see, e.g., [28, 36, 44, 76]. On the other hand, approximate Nash equilibria can be shown to converge to solutions of the corresponding MFG. The choice of admissible strategies, while always important, is crucial for results of this kind: see [65, 85] for results in open loop strategies and [33, 87, 88] for convergence in closed loop strategies.

Despite their popularity, Nash equilibria present some flaws, well known in game theory literature. First, they raise numerical complexity issues, see for instance [68]. Second, agents are proved to actually behave according to a Nash equilibrium only under strong rationality assumptions. Finally, they can be highly inefficient compared to Pareto optimum. As an alternative to Nash equilibria, correlated equilibria (CEs) (see Aumann's (see [6, 7])) and coarse correlated equilibria (CCEs) (see Moulin's and

Vial's [93]) have been introduced in game theory literature. They can be understood as generalizations of the notion of Nash equilibrium by the introduction of a correlation device, which allows agents to adopt correlated strategies without any cooperation.

CEs and CCEs can be described as follows: a centralized instance, which we refer to as the moderator, provides each players with a recommendation. The recommendation can be seen as a lottery over the set of strategy profiles, which selects a strategy for each player according to a publicly known distribution. In other words, the joint distribution of the recommendations to the players is common knowledge and may correlate players' strategies. Each player must then decide whether to follow the suggested strategy or deviate from it. Deviations can be of two types, leading to either correlated equilibria or coarse correlated equilibria.

In CEs, each player is informed of her (and only her) sampled strategy. After observing it and knowing the distribution of the lottery over strategy profiles, she evaluates whether it is optimal for her to play the suggested strategy, supposing the the other players do so. A moderator's lottery is a CE if no player has any incentive to unilaterally deviate from moderator's recommendation after having received it.

In CCEs, each player must decide whether to commit to the moderator recommendation (whatever it will be) before the lottery is actually run. If a player commits, then she is communicated in private her (and only her) selected strategy, and must follow it. Instead, if a player deviates, she will do so without any information on the outcome of the lottery. Therefore, in CCEs each player evaluates whether it is optimal for her to commit to the suggested strategy, supposing the the other players do so, based only on the distribution of the lottery over strategy profiles and before observing her sampled strategy. A moderator's lottery is a CCE if every player prefers to commit rather than unilaterally deviate, assuming that all others do commit.

The difference between CEs and CCEs resides in the different nature of the deviations from the moderator's recommendation. In CEs, a deviating player is partially informed about moderator's lottery through her received recommendation and can use this information in her decision making. Thus, player's deviations depend on the moderator's recommendation itself. On the contrary, in CCEs each player decides whether to commit to moderator's lottery before being informed of it. As a result, deviations in CCEs are independent of moderator's recommendation itself; a deviating player has no information about the outcome of moderator's lottery and must rely only on the publicly known lottery's distribution to decide whether to commit. Consequently, if a lottery is a CE it is also a CCE, as CEs' possible deviations include those in CCEs, while the contrary is generally false.

When the distribution used by the mediator is a product distribution, CCEs reduce to usual Nash equilibria in mixed strategies, as in this case the mediator's recommendations do not carry any additional information over what is common knowledge. Additionally, when dealing with Nash equilibria, each player autonomously makes her decisions, and the equilibrium property arises by taking into account each player's individual response to the others, giving rise to a fixed point problem. On the contrary, in CEs and CCEs, players react to a moderator's suggestion.

Among the nice features of CEs and CCEs, they are shown to be able to lead to higher payoffs than Nash equilibria in standard game theory (see [57, 92, 91] for examples in static linear quadratic games), even in situations where correlated equilibria

cannot, for instance in potential games [96]. Moreover, they are computationally “easier” as shown by [68]. Finally, they naturally arise from a learning procedure of the players, such as the so called regret-based dynamics (see, e.g., Hart and Mas-Colell [73], Roughgarden [102, Section 17.4]). It must be noticed that the first implicit appearance of CCEs is given in Hannan’s [72], in which learning dynamics were taken into account, and a very natural convergence result to CCEs (Hannan’s set therein) has been proved.

Recently, correlation between players’ strategy choices has been considered in the context of mean field games. In the works [19, 20, 29] by Bonnesini, Campi and Fischer, existence and convergence results are established for correlated equilibria in mean field games with discrete time and finite state and action spaces. Remarkably, the proposed notion of correlated MFG solution features a naturally stochastic flow of measures, as the aggregation of individual behaviours preserves the stochasticity of the correlation device. A second group of papers by Laurière et al. [90] and Muller et al. [94, 95] considers both CEs and CCEs in a similar setting. Interestingly, they propose a different definition of CE for the mean field game from that in [19, 20, 29], and in [94] the authors discuss on how the two definitions of CE can be seen as equivalent. In addition, [94, 95] contain an extensive discussion of learning algorithms for approximating Nash equilibria, CEs and CCEs in the mean field limit.

Notions of equilibria other than Nash have already been considered in literature, both for games with finitely many and infinitely many players. We cite the principal-agent problems and Stackelberg equilibria (see, e.g., the book [48] or the papers [8, 16, 56, 61] and the references therein) and the mean field control (MFC) problem, or stochastic optimization problem of McKean-Vlasov (MKV) type. As for Stackelberg equilibria and principal-agent problems, they feature a leader or a principal who decides first her strategy to anticipate the best response of the followers or agents. Correlated and coarse correlated equilibria are essentially different from the aforementioned problems, for one main reason: Differently from these problems, in CEs and CCEs, the mediator is not in principle an optimizer. The mediator may not even be an actual person or agency, although we have opted for the interpretation of a mediator recommending strategies: it may be just the result of the learning procedure of the players (see again [73]), or the result of a pre-play communication protocol among the players (see, among others, [14]). As for MFC control problems, their solutions can be regarded as the infinitely many players counterparts of Pareto optima strategy profiles in the N -player game, often reached via the solution of a so-called central planner’s optimization problem. Sometimes referred to as cooperative equilibria, the notions of MFC solutions and Pareto efficient strategies can be related via convergence of approximation arguments (see, e.g., [36, 86]), analogously to Nash equilibria and MFG solutions. The MFC problem is essentially different both from MFGs (see [40] for an illuminating discussion of their differences), and from CCEs, as in CCEs the moderator does not have any power to force her decision on the players, but they can always choose to deviate from moderator’s suggested strategy. On the contrary, when considering the central planner’s optimization problem, players do not decide by themselves which strategy to play, but the central planner unilaterally allocates the agents’ resources to maximise a welfare utility functional.

The aim of this thesis is to introduce coarse correlated equilibria in continuous-time MFGs and to study their properties. In particular, the problems addressed by this thesis are:

- (A) Give a proper definition of coarse correlated equilibria both in continuous-time stochastic N -player games and in MFGs;
- (B) Ensure existence of coarse correlated equilibria in general MFGs;
- (C) Provide a methodology for computing coarse correlated equilibria in specific classes of MFGs;
- (D) Analyse the payoffs yielded by coarse correlated equilibria in the MFG, in comparison with MFG solutions and solutions of the MFC problem;
- (E) Justify the definition of coarse correlated equilibria in the MFG, by analyzing their relationship with coarse correlated equilibria in the underlying N -player game.

In Chapter 1, we consider MFGs driven by additive Wiener noise, with general drift and cost functions. The interaction term is given through a flow of probability measures, which appears both in the drift and in the cost functions. We introduce coarse correlated equilibria in both continuous time stochastic differential games and mean field games. The notion of coarse correlated solution to the MFG is justified by proving an approximation result. An existence result is also presented, whose proof relies on a minimax theorem.

In Chapter 2, we consider linear-quadratic MFGs. The interaction term is given by a flow of first order moments, which appears in the payoff functional only. A methodology to compute CCEs in such class of MFGs is provided and, through the study of a simple yet important example with applications in environmental economics, we show that there exist infinitely many CCEs for the MFG which both yield higher payoffs than the classical MFG solutions and are more efficient with respect to the environmental goals, highlighting the benefits provided by CCEs over MFGs solutions. Finally, through the study of a specific MFG, we show that infinitely many CCEs may exist even when the MFG does not admit any MFG solution.

In Chapter 3, we consider a simple class of stationary MFGs of singular control. The reward, of ergodic type, is given by the long-time average of an expected utility functional. The interaction term is given through the stationary mean of the distribution of the representative player, which appears only in the reward. We provide constructive existence results, as well as approximation results and comparison with both MFG solutions and MFC solutions. Moreover, we show that CCEs may exist even when MFG solutions do not.

Referring to previous works [19, 20, 29], the reason for considering CCEs instead of CE is both theoretical and practical. First, they are more general than CE, thus than Nash equilibria, both in mixed and pure strategies. Secondly, the fact that if a player deviates, she is not informed of the outcome of the moderator's lottery makes the treatment of CCEs easier than the one of CE, due to the fact that in CCEs deviations are independent of the outcome of moderator's lottery. On the contrary, in

CEs, every player is informed of the outcome of moderator’s lottery, and then decides whether to play accordingly or not. Thus, in CEs, deviations would depend on that outcome, giving rise to delicate measurability issues, which for the moment we do not know how to handle properly. Moreover, while in [20, 29] restricted closed loop strategies were considered, in this thesis we consider for simplicity stochastic open loop strategies, since we deal with the more challenging problem of continuous time, actions and states.

The results in this thesis are original. In particular:

- The results in Chapter 1 are obtained in collaboration with L. Campi and M. Fischer and can be found in:

L. Campi, F. Cannerozzi, M. Fischer, *Coarse correlated equilibria for continuous time mean field games in open loop strategies*, Electronic Journal of Probability, 29:Paper No. 1, 2024.

- The results in Chapter 2 are obtained in collaboration with L. Campi and F. Cartellier and can be found in:

L. Campi, F. Cannerozzi, and F. Cartellier. *Coarse Correlated Equilibria in Linear Quadratic Mean Field Games and Application to an Emission Abatement Game*, Applied Mathematics and Optimization, 91(1):Paper No. 8, 2025.

- The results in Chapter 3 are obtained in collaboration with G. Ferrari and can be found in:

F. Cannerozzi and G. Ferrari. *Cooperation, correlation and competition in ergodic N -player games and mean-field games of singular controls: A case study*. Preprint: arXiv:2404.15079, 2024.

This thesis consists of three chapters and two appendices. In the following sections, we give a more detailed introduction to each chapter, and we briefly present the results obtained. Regarding the appendices, in Appendix A we recall some results on relaxed controls, while in Appendix B a result regarding weak and strong uniqueness of controlled equations is stated and proven.

I.1 Chapter 1: MFGs with general dynamics and costs

In this chapter, we consider continuous-time MFGs with general dependence on the flow of measures, and we deal both with the existence of CCEs in the MFG and with the relation between CCEs in the N -player game and in the MFG. We now describe the setting, without giving the details or the precise assumptions we work with.

Let $d \in \mathbb{N}$, $d \geq 1$. In the N -player game, on a suitable probability space $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$, large enough to support a sequence $(W^i)_{i \geq 1}$ of independent d -dimensional Brownian motions and a sequence $(\xi^i)_{i \geq 1}$ of independent and identically random variables with values in \mathbb{R}^d , we consider the state dynamics

$$\begin{cases} dX_t^i = b(t, X_t^i, \mu_t^N, \alpha_t^i)dt + dW_t^i, & 0 \leq t \leq T, & i = 1, \dots, N \\ X_0^i = \xi^i, \end{cases} \quad (\text{I.1.1})$$

where μ_t^N is the empirical measure of the state processes of all players at time t :

$$\mu_t^N = \frac{1}{N} \sum_{j=1}^N \delta_{X_t^j}. \quad (\text{I.1.2})$$

Each player is associated with the following cost functional

$$\mathfrak{J}_i^N(\alpha^1, \dots, \alpha^N) = \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^N, \alpha_t^i) dt + g(X_T^i, \mu_T^N) \right]. \quad (\text{I.1.3})$$

Here, $T > 0$ is the fixed time horizon, A denotes the set of actions and $(b, f) : [0, T] \times \mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \times A \rightarrow \mathbb{R}^d \times \mathbb{R}$ are Borel measurable functions satisfying suitable assumptions, and so is $g : \mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \rightarrow \mathbb{R}$. The strategies for the N -player game are open-loop, i.e., the correlation device would recommend the players to use strategies adapted to the filtration generated by noises and initial data. More precisely, if we set $\mathbb{F}^{1,N}$ to denote the filtration generated by the Brownian motions and the initial data, strategies are given by the processes $\alpha = (\alpha_t)_{t \in [0, T]}$ with values in A which are $\mathbb{F}^{1,N}$ -progressively measurable and square integrable. We denote the set of open loop strategy profiles by \mathbb{A}_N^N , and we endow it with a suitable σ -algebra \mathcal{A}^N .

The correlation device, or recommendation to the N players, is modeled as a random variable taking values in the set of open loop strategy profiles; we require it to be independent of the random shocks and the initial states which determine players' states' evolution. More precisely, we have the following definition:

Definition 1.2.1. We call *recommendation profile* to the N players a pair $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ so that the following holds:

1. $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is a complete probability space; Ω^0 is a Polish space and \mathcal{F}^{0-} is its corresponding Borel σ -algebra.
2. $\Lambda = (\Lambda^1, \dots, \Lambda^N)$ is a random vector with values in \mathbb{A}_N^N :

$$\begin{aligned} \Lambda : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) &\longrightarrow (\mathbb{A}_N^N, \mathcal{A}^N) \\ \omega_0 &\longmapsto \Lambda(\omega_0) = (\alpha^1, \dots, \alpha^N) : [0, T] \times \Omega^1 \rightarrow A^N. \end{aligned} \quad (\text{I.1.4})$$

We then build a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ large enough to support both the moderator's recommendation profile Λ and the random shocks and initial data $(\xi^j, W^j)_{j=1}^N$, in such a way that they are independent. On this new probability space, we consider the strategy profile associated to the recommendation Λ , that is, the strategies that will appear in the committing players' dynamics and payoffs. Roughly speaking, each committing player i would use the strategy λ^i given by

$$\lambda_t^i(\omega) = \Lambda^i(\omega_0)_t(\omega_1).$$

The strategies $(\lambda^i)_{i=1}^N$ will depend both on the realisation of moderator's lottery $\Lambda^i(\omega_0)$ and the initial data and idiosyncratic noise that determine the evolution of players' states. We deal very carefully with the measurability properties the recommendation has to fulfill so that the players' states are well-defined and the recommended

strategies are implementable by the players, which leads to the definition of admissible recommendation profile in Definition 1.2.2. We provide instances of admissible recommendations in Example 1.3.1.

If a player follows the moderator's recommendation $\Lambda^i(\omega_0)$, then she will play according to the strategy λ^i associated to the outcome $\Lambda^i(\omega_0)$, while if a player does not commit to moderator's recommendation, she will use an open loop strategy β in \mathbb{A}_N . We notice that there is an asymmetry between the information available to the mediator and the deviating player, since the information available to the deviating player is just given by the smaller Brownian filtration $\mathbb{F}^{1,N}$. In particular, she will use a strategy which is independent of \mathcal{F}^{0-} , thus of the admissible recommendation profile Λ . This models the fact that if a player deviates, she will do so without any knowledge about the outcome $\Lambda^i(\omega_0)$ of the moderator's lottery: she will not exploit any of the additional information the mediator would give away when communicating the recommended strategies to the players. Conversely, the strategy profile associated to the recommendation $\lambda = (\lambda_t)_{t \in [0,T]}$ is by definition dependent of the recommendation Λ , and it carries at least some of the information the the moderator uses to randomize players' strategies. Notice that, although a deviation β is independent of the moderators recommendation Λ , it is not independent of the strategy profile λ associated to the recommendation, as both are dependent on the random shocks and initial data $(\xi^j, W^j)_{j=1}^N$ that determine players' dynamics.

If all players follow the recommendation, then the dynamics are given by (I.1.1) with λ^j instead of α^j , for any $j = 1, \dots, N$, and the cost functional of player $i = 1, \dots, N$ is given by $\mathfrak{J}_i^N(\Lambda)$ according to (I.1.3). If instead player i does not play according to the recommendation Λ^i and plays a different strategy $\beta \in \mathbb{A}_N$, while the other players stick to the recommendation profile Λ^{-i} , then the dynamics are given by (I.1.1) with β instead of α^i and λ^j instead of α^j , for each $j \neq i$, and the cost functional of player i is given by $\mathfrak{J}_i^N(\Lambda^{-i}, \beta)$ according to (I.1.3). The equilibrium property is as usual: An admissible recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is a coarse correlated equilibrium for the N -player game if no player as any incentive to unilaterally deviate from the moderator's recommendation, i.e.

$$\mathfrak{J}_i^N(\Lambda) \leq \mathfrak{J}_i^N(\Lambda^{-i}, \beta)$$

for all open loop strategies β in \mathbb{A}_N and all players $i = 1, \dots, N$. We notice that our definition of CCE extends the one of Nash equilibria in open loop strategies, as every NE in open strategies is a CCE according to our formulation: It is enough to choose $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ as the trivial probability space and Λ as constant and equal to the NE. In other words, we make the moderator run a deterministic lottery over the set of strategy profile, which selects the NE with probability 1.

In the mean field limit, the notion of coarse correlated solution we present corresponds to a pair given by a recommendation with values in the set of open loop strategies for the representative player and a random flow of measures fulfilling the following two properties:

- *Optimality*: the representative player has no incentive to deviate from the recommended strategy *before* the extraction has happened.

- *Consistency*: the flow of measures at any time t equals the marginal law of the representative player's state conditioned on the σ -algebra generated by the whole flow of measures up to terminal time.

More specifically, we consider a suitable probability space $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$, large enough to support a Brownian motion W and an independent random variable ξ . We set \mathbb{F}^* to denote the filtration generated by the Brownian motion and the initial datum. The strategies for the representative player are again open-loop, and the set of open-loop strategies for the MFG is denoted by \mathbb{A} , endowed with suitable σ -algebra. Analogously to the N -player game, the moderator would recommend the representative player to use open-loop strategies. To take into account the impact of moderator's lottery on the flow of measures, we introduce the notion of correlated measure flow:

Definition 1.3.3. A *correlated measure flow* is a triple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ where:

1. $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is a recommendation to the representative player.
2. $\mu : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) \rightarrow (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)})$ is a random continuous flow of measures in $\mathcal{P}^2(\mathbb{R}^d)$.

Notice that we request the flow of measures to be measurable with respect to \mathcal{F}^{0-} , since it is expected to be stochastic as a result of the mediator's randomization only.

We consider again a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ large enough to support both the moderator's recommendation Λ , the flow of measures μ and the pair (ξ, W) , in such a way that the pairs (Λ, μ) and (ξ, W) are independent. On this new probability space, we consider the strategy profile associated to the recommendation Λ , that is, the strategies that will appear in the committing players' dynamics and payoffs, which will be given by

$$\lambda_t(\omega) = \Lambda(\omega_0)_t(\omega_*),$$

provided that the measurability requirements are satisfied. If the representative player decides to play according to the admissible recommendation Λ , the dynamics is given by the following SDE:

$$\begin{cases} dX_t = b(t, X_t, \mu_t, \lambda_t)dt + dW_t, & 0 \leq t \leq T, \\ X_0 = \xi, \end{cases} \quad (\text{I.1.5})$$

and the cost functional is given by

$$\mathfrak{J}(\Lambda, \mu) = \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \lambda_t)dt + g(X_T, \mu_T) \right]. \quad (\text{I.1.6})$$

If instead the representative player decides to ignore the mediator's recommendation and to use a possibly different strategy β in \mathbb{A} , the dynamics is given by (I.1.5), with β replaced by λ . As for the equilibrium property, we have the following: A correlated measure flow $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ is a *coarse correlated solution* of the mean field game if the following properties hold:

- *Optimality*: for every open loop deviation β in \mathbb{A} , it holds

$$\mathfrak{J}(\Lambda, \mu) \leq \mathfrak{J}(\beta, \mu). \quad (\text{I.1.7})$$

- *Consistency*: for every time $t \in [0, T]$, μ_t is a version of the conditional law of X_t given μ , that is,

$$\mu_t(\cdot) = \mathbb{P}(X_t \in \cdot \mid \mu) \quad \mathbb{P}\text{-a.s.} \quad \forall t \in [0, T]. \quad (\text{I.1.8})$$

The consistency condition (I.1.8) should be read in the following way: the mediator imagines what the flow of measures will be, up to the terminal horizon T , before the game starts, and gives a recommendation to each player according to his idea. If all players commit to the mediator's lottery for generating recommendations, then the flow of measures should arise from aggregation of the individual behaviors, consistently with what imagined by the mediator. Regarding the strategy of the deviating player, if the player deviates, she chooses her strategy on her own, without using any of the information carried by Λ or μ : the only information she has about Λ or μ comes from the knowledge of their joint law, which is assumed to be known by the representative player, in analogy to the N -player game.

As for the N -player game, the usual notion of MFG solution (see Definition 1.3.5) is included in our definition of coarse correlated solution to the MFG. Indeed, starting from a MFG solution (α^*, μ^*) , with $\alpha^* \in \mathbb{A}$ an open loop strategy and μ^* a deterministic flow of measures, we choose $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ as the trivial probability space, take μ deterministic and equal to μ^* and make the moderator recommend with probability one to play according to α^* .

In Section 1.6, through the study of a simple example, our notion of coarse correlated solution to the MFG is compared to the more usual notion of MFG solution, (as defined, e.g., in [36]) and the notion of weak MFG solution of [85], as they both may feature a random flow of measures which is not due to a common noise that acts equally on the players. By the study of this specific MFG, we show that there exist infinitely many mean field coarse correlated solutions to the MFG which are not weak solutions to the MFG without common noise. Moreover, we show that weak MFG solutions which satisfy an additional measurability constraint are indeed coarse correlated solution to the MFG as well.

The main contributions of this chapter are as follows:

- i) We justify our notion of coarse correlated solution for the MFG by showing that any coarse correlated solution for the MFG induces a sequence of approximate CCEs in the N -player game, with vanishing error as N goes to infinity. This is the content of Section 1.4.
- ii) Under an additional convexity assumption, in Section 1.5 we prove the existence of a coarse correlated solution for the mean field game.

Both results will be established using a genuinely probabilistic approach. As for the approximation result, we use the limit flow of measures to act as a correlation device between players' strategies in the N -player game, in the same spirit of [20, 29]. More specifically, starting from a coarse correlated solution to the MFG $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda^*, \mu^*)$, we build a probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ large enough to carry any sequence $(\Lambda^i)_{i \geq 1}$ of admissible recommendations such that:

1. for every i , Λ^i is supported on the set of open-loop strategies progressively measurable with respect to player i 's private noise;
2. for every i , Λ^i has the same distribution as Λ^* ;
3. for every $N \geq 2$, $(\Lambda^1, \dots, \Lambda^N)$ is exchangeable.

Then, we define the probability space $(\Omega^{0,N}, \mathcal{F}^{0-,N}, \mathbb{P}^{0,N})$ as $(\bar{\Omega}, \mathcal{F}^{0-,N}, \bar{\mathbb{P}})$, where $\mathcal{F}^{0-,N}$ a suitable sub- σ -field of $\bar{\mathcal{F}}^{0-}$. The proof of the inverse convergence relies on propagation of chaos arguments, which are in part reminiscent of [36, 44], but specific to our situation. In order to exploit the optimality property (I.1.7) of the coarse correlated solution, great care is taken in comparing the cost functional associated to open-loop strategies in the N -player game with the payoff of the coarse correlated solution.

To prove the existence result, we associate a zero-sum game to the search of a coarse correlated solution for the MFG, inspired by the works of Hart and Schmeidler [74] (for static games), Nowak [97, 98] (for continuous time dynamic games) and Bonesini [19, Appendix 1.B] (for MFGs with discrete time and finite states and actions, in the setting of [29]), which requires us to apply a minimax theorem. Loosely speaking, the game should be of the following type: player A, the maximizer, chooses a correlated measure flow (Λ, μ) , while player B chooses a deviating strategy $\beta \in \mathbb{A}$. The payoff functional is the following:

$$F[(\Lambda, \mu), \beta] = \mathfrak{J}(\beta, \mu) - \mathfrak{J}(\Lambda, \mu). \quad (\text{I.1.9})$$

Player A aims at maximizing F , while player B chooses her strategy in order to minimize F . In order to get an equilibrium, one should restrict to correlated measure flows (Λ, μ) so that the consistency condition (I.1.8) is satisfied. If we could show that the game has a positive value and player A has an optimal strategy (Λ^*, μ^*) , then we would have established that such a strategy would satisfy the optimality property (1.3.12) as well, and therefore (Λ^*, μ^*) would be a mean field CCE. In order to get a convenient structure for the sets of strategies and good continuity and convexity properties of the payoff functionals, we define a more general zero-sum game, in which we embed our auxiliary problem. Particular care is needed in dealing with the term depending both on β and μ , since it must reflect independent strategy choices of the opponents. Using Fan's minimax theorem, we will show that the auxiliary game has positive value and admits an optimal strategy for the maximizing player and, finally we are able to prove the existence result, by using such an optimal strategy to induce a coarse correlated solution of the mean field game. To do so, compactness arguments are exploited, adapting some of the techniques used in Lacker's works [84, 87].

I.2 Chapter 2: Linear-quadratic MFGs

In Chapter 1, CCEs have been introduced in both continuous time stochastic differential games and MFGs. The notion of coarse correlated solution to the MFG is justified by proving an approximation result, and an existence result is also proved, by means

of a minimax theorem. Although its generality, this result is not constructive, and the question of how to construct coarse correlated solutions to MFG is left open.

This chapter's goal is to develop a methodology for computing coarse correlated equilibria in MFGs, and to effectively compare them to MFG solutions and MFC solutions, as they respectively are the mean field counterparts of Nash equilibria and Pareto optima strategy profiles. For this reason, we do not consider the N -player game, but we limit our analysis to the mean field game. Since we search for explicit solutions, we restrict our analysis to linear-quadratic stochastic MFGs, working in a setting closely related to [69]. More precisely, we consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, large enough to support a Brownian motion W and an independent random variable ξ , both with values in \mathbb{R}^d , $d \geq 1$. The representative player's state variable evolves with linear dynamics given by

$$dX_t = (A_t X_t + B_t \alpha_t) dt + \sigma_t dW_t, \quad X_0 = \xi, \quad (\text{I.2.1})$$

and it is associated with a linear-quadratic payoff functional, to be maximized,

$$\begin{aligned} \mathfrak{J}(\alpha, \bar{\mu}) = \mathbb{E} \left[\int_0^T \left(\langle L_t, \bar{\mu}_t \rangle - \frac{1}{2} \langle \bar{Q}_t \bar{\mu}_t, \bar{\mu}_t \rangle - \left(\frac{1}{2} \langle Q_t X_t, X_t \rangle + \langle \tilde{Q}_t X_t, \bar{\mu}_t \rangle + \frac{1}{2} \langle R_t \alpha_t, \alpha_t \rangle \right. \right. \right. \\ \left. \left. + \langle q_t, X_t \rangle + \langle S_t X_t, \alpha_t \rangle + \langle r_t, \alpha_t \rangle \right) dt - \frac{1}{2} \langle \bar{H} \bar{\mu}_T, \bar{\mu}_T \rangle - \left(\frac{1}{2} \langle H X_T, X_T \rangle + \langle \tilde{H} X_T, \bar{\mu}_T \rangle \right) \right]. \end{aligned} \quad (\text{I.2.2})$$

Here, $T > 0$ is a fixed finite time horizon and $\bar{\mu} = (\bar{\mu}_t)_{t \in [0, T]}$ is a random variable taking values in \mathcal{C}^d . All the coefficients appearing in the dynamics and in the payoff functional are measurable deterministic matrix valued functions, satisfying suitable assumptions. The strategy α is a square integrable process taking values in \mathbb{R}^k , $k \geq 1$, with measurability requirements to be specified.

Throughout this whole chapter, we will refer to coarse correlated equilibria in MFGs as mean field CCEs and to MFG solutions as mean field Nash equilibria. Although inspired from the notion of coarse correlated solution to the MFG of Chapter 1, the definition of mean field CCE is different from the one in previous chapter and specific to this framework. Here, we rely on the notion of correlated moment flow, as given in the following definition:

Definition 2.2.1. A *correlated moment flow* is a pair $(\lambda, \bar{\mu})$ satisfying the following properties:

- i) $\lambda = (\lambda_t)_{t \in [0, T]}$ is a square integrable \mathbb{F} -progressively measurable process with values in \mathbb{R}^k .
- ii) $\bar{\mu} = (\bar{\mu}_t)_{t \in [0, T]}$ is an \mathcal{F}_0 -measurable random variable with values in \mathcal{C}^d .
- iii) $\bar{\mu}$ is independent of both ξ and W .

We refer to λ as the *recommended strategy* and to $\bar{\mu}$ as the *random flow of moments*.

We can interpret a correlated moment flow $(\lambda, \bar{\mu})$ as follows: moderator's lottery is run before the game starts and independently of the idiosyncratic shocks that

determine the random evolution of representative player's state. The way they are correlated is chosen by the moderator at the beginning of the game as part of the equilibrium. In mean field CCEs, the flow of moments can be stochastic, and the strategy should be correlated to the flow of moments only as a consequence of moderator's lottery. We stress that, while the recommended strategy λ is correlated both to ξ and W and to $\bar{\mu}$, $\bar{\mu}$ is independent of the initial datum and the noise.

If the representative player decides to trust the mediator and therefore accepts to follow her recommendation λ before knowing it, the dynamics is given by equation (I.2.1) with λ instead of α , and the player gets the reward $\mathfrak{J}(\lambda, \bar{\mu})$. If instead she decides to deviate, then she will use an open loop strategy, i.e. a process β with values in \mathbb{R}^k , square integrable and progressively measurable with respect to the Brownian filtration $\mathbb{F}^{\xi, W}$. As in Chapter 1, we denote the set of open loop strategies for the representative player by \mathbb{A} , so that $\mathbb{A} = \mathbb{H}^2(\mathbb{F}^{\xi, W})$. The deviating players dynamics is given by equation (I.2.1) with β instead of α , and her reward is $\mathfrak{J}(\beta, \bar{\mu})$. Observe that when she deviates, her strategy β is measurable only with respect to the initial datum and the idiosyncratic noise, since she has no information on the outcome of the moderator's lottery. The deviating player can only use her knowledge of the law of the correlated moment flow $(\lambda, \bar{\mu})$, which is assumed to be publicly known.

In this context, the equilibrium property is as follows: A correlated moment flow $(\lambda, \bar{\mu})$ is a mean field CCE if the following holds:

- *Optimality*: for every open loop strategy β , it holds

$$\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\beta, \bar{\mu}).$$

- *Consistency*: for every time $t \in [0, T]$, $\bar{\mu}_t$ is a version of the conditional expectation of X_t given $\bar{\mu}$, that is,

$$\bar{\mu}_t = \mathbb{E}[X_t | \bar{\mu}] \quad \mathbb{P}\text{-a.s.} \quad \forall t \in [0, T].$$

The first and main contribution of the paper is to provide a methodology to compute coarse correlated equilibria, which is detailed in Section 2.3. Since the set of CCEs is typically very wide and it is difficult to characterize in a continuous time setting, we focus on a tractable class of correlated moment flows for which we are able to characterize a sufficient condition for being a mean field CCE. To do so, we adopt the following procedure:

- We fix a correlated moment flow $(\lambda, \bar{\mu})$. We suppose that the representative player does not commit to the moderator's lottery and we compute her best deviating strategy $\hat{\beta}$, i.e.

$$\hat{\beta} = \arg \max_{\beta \in \mathbb{A}} \mathfrak{J}(\beta, \bar{\mu}).$$

$\hat{\beta}$ will depend upon the law of $(\lambda, \bar{\mu})$ itself, but not on its actual realization.

- We define a parameterised class of correlated moment flows $(\lambda, \bar{\mu})$ of similar shape as the best deviating strategy $\hat{\beta}$ so that the consistency condition is fulfilled. The correlation is due to a suitable random parameter δ .

- Finally, for $(\lambda, \bar{\mu})$ in such a class, with corresponding parameter δ , we express the optimality condition

$$\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\hat{\beta}, \bar{\mu})$$

as an inequality involving the law of $\bar{\mu}$ and δ only.

As a result, we reduce the search for a mean field CCE to finding a law for $\bar{\mu}$ and δ that verifies an optimality inequality. The choice of focusing on a class of correlated moment flows with shape similar to the best deviation allows for explicit analytical comparison between the two payoffs in the optimality condition.

Then, in Section 2.4 we analyze the relationship between mean field CCEs, mean field NEs and MFC solutions, under additional assumptions ensuring existence and uniqueness of mean field NEs and MFC solutions. In more detail, we prove the following results:

- We show that no mean field CCE can outperform the payoff of the MFC solution. Moreover, if the MFC solution is not a mean field NE, we establish that the MFC payoff is unattainable by a mean field CCE.
- Analogously to what noticed in Chapter 1, any mean field NE is also a mean field CCE, with deterministic flow of moments and $\mathbb{F}^{\xi, W}$ -progressively measurable recommendation. We show that the reverse implication is true as well: any mean field CCE with deterministic flow of moments is in fact a mean field NE.
- Finally, specifically to the parameterized class of correlated moment flows mentioned above, we give a condition so that a mean field CCE $(\lambda, \bar{\mu})$ yields a higher payoff than the mean field NE.

We apply our results to an emission abatement game between countries, inspired by environmental economics literature on international environmental agreements [12, 57]. This is the content of Section 2.5. We show that it is possible to build simple mean field CCEs that both yield much higher payoffs than the mean field NE and guarantee higher average abatement levels. In this one-dimensional MFG, the parameterised class of correlated moment flows $(\lambda, \bar{\mu})$ is so that the random parameter coincides with the first derivative of the flow of moments $\bar{\mu}$ itself. After assuming that such derivative is constant in time, we are able to express the optimality condition, the outperformance condition and the trade-off between higher payoffs and higher abated quantities in terms of the mean and the variance of the (constant) growth rate of $\bar{\mu}$ itself.

This application exemplifies the role of mean field CCEs as a middle ground between mean field NEs and MFC solutions. Moreover, it shows an additional interest of CCEs, which is to help a regulator not only to lead the population to a more optimal payoff than the free-riding NE, but also or otherwise to lead it to match other and potentially payoff-conflicting targets, such as the abatement level of players in this application.

Finally, in Section 2.6 we consider a simple MFG class, studied in [84, Section 7], which does not admit any mean field NE. As in the abatement game, we apply our results to show that mean field CCEs may exist even when mean field NE do not.

The occurrence of this phenomenon is due to the nature of our methodology, which does not involve the usual two steps procedure used to compute mean field NEs: first, optimize with a fixed flow of moments and, second, perform a fixed point argument to determine the flow. Since our results are still valid even when the MFG fixed-point condition fails to hold, existence of mean field CCEs may hold even when mean field NEs fail to exist.

I.3 Chapter 3: Stationary MFGs of singular control

In this chapter, we investigate a simple class of one-dimensional ergodic stochastic games of singular control in both competitive and cooperative settings, considering scenarios with a finitely many players as well as in the mean field limit. The goal is to characterize and compare Nash equilibria, coarse correlated equilibria and Pareto optima in the mean field limit and to analyse their relations with the corresponding equilibria in the underlying N -player game.

In the N -player game, presented in Section 3.1, we introduce the concepts of coarse correlated equilibrium, Nash equilibrium, and Pareto efficiency for this game. The notion of Pareto efficiency is associated with the problem of a central planner who seeks to maximize the average of the rewards of all N agents. As constructing equilibria for N -player games in continuous-time and space is a challenging problem, we study the corresponding MFG. In the mean field limit, the interaction term is given by the long-time average of the distribution of the population, which is represented by a positive scalar parameter m .

The game under study shows strategic complementarities (since the marginal profit is increasing in the interaction variable; cf. [106]) and finds natural applications in dynamic oligopolies, such as in Cournot oligopoly with complementary products or in advertising games (see, e.g., [107, 108, 109]). In this regard, the state variable of each agent can be the output or the goodwill stock, which is increased by irreversible investment or advertising, respectively. The resulting payoff is then derived through an isoelastic inverse demand function, depending on the aggregate level of production or goodwill in the entire market. In particular, the ergodic structure of the reward functional we consider is relevant in the context of investment into public goods, in which it might be important to take care of the payoffs received by successive generations. The class of games with strategic complementarities (also known as submodular/supermodular games) has garnered significant attention in Economics. Particular importance has the potential emergence of multiple equilibria in these games. As Xavier Vives asserts in [109], “Complementarities are intimately linked to multiple equilibria and have a deep connection with strategic situations, and the concept of strategic complementarity is at the center stage of game-theoretic analyses”. Among the myriad contributions, we refer to the deterministic games considered in [106, 107, 108], as well as to the dynamic stochastic formulations and mean field formulations presented in [1, 2, 3, 42, 51, 52, 53, 54].

There is a rapidly increasing number of contributions on MFGs of singular controls, which has focused on abstract results regarding the existence and uniqueness

of equilibria (see [47, 50, 54, 66, 67]), as well as on explicit characterizations of the Nash solution (see [27, 32, 31, 55, 71]). Particularly related to this setting are [31] and [47]. In [31], in the context of a mean field game with singular controls, stationary discounted and ergodic Nash equilibria are explicitly constructed and related via the vanishing discount factor method. In the very recent [47], existence of optimal controls for stationary singular single-agent control problems as well as the existence of mean field game equilibria for stationary singular MFGs are derived through a relaxed approach. In our setting, the stationary one-dimensional setting of the mean field game and control problem allows for explicit characterizations of the equilibria (see also [13, 31, 32, 45] and references therein in the context of singular/impulse control games).

We now describe in more detail the stationary MFG and present the notions of mean field CCEs in this context. As detailed Section 3.2, we work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supporting a standard Brownian motion W and an independent non-negative random variable ξ . We consider singular controls, i.e. processes ν belonging to the set

$$\mathcal{S} := \{ \nu : \Omega \times \mathbb{R}_+ \rightarrow \mathbb{R}_+, \mathbb{F}\text{-adapted and such that } t \mapsto \nu_t \text{ is a.s.} \\ \text{non-decreasing, right-continuous, } \nu_{0-} = 0 \text{ and } \mathbb{E}[\nu_T] < \infty \forall T > 0 \},$$

to be subject to further integrability and measurability restrictions. Given a strategy $\nu \in \mathcal{S}$, we consider the following geometric dynamics:

$$dX_t^\nu = -\delta X_t^\nu dt + \sigma X_t^\nu dW_t + d\nu_t, \quad X_{0-} = \xi. \quad (\text{I.3.1})$$

For any \mathcal{F}_{0-} -measurable non-negative random variable \bar{m} , possibly degenerate, we consider the following reward, to be maximized:

$$\mathfrak{J}(\nu, \bar{m}) = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^\nu)^\alpha \bar{m}^\beta dt - q\nu_T \right]. \quad (\text{I.3.2})$$

If there is no moderation, the representative player will use open loop strategies, i.e. processes $\nu \in \mathcal{S}$ which are adapted to the filtration $\mathbb{F}^{\xi, W}$ generated by ξ and W . As in previous chapters, we denote by \mathbb{A} the set of open loop strategies for the MFG.

As for mean field CCEs, we rely on the following notion, which plays the same role of correlated measure flows in Chapter 1 and correlated moment flows in Chapter 2:

Definition 3.2.2. A *correlated stationary strategy* is a triple $(Z, \lambda, \bar{\mu}_\infty)$ satisfying the following properties:

- (i) Z is a correlation device, i.e a random variable $Z : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ so that Z is \mathcal{F}_{0-} -measurable and independent of ξ, W .
- (ii) $\lambda = (\lambda_t)_{t \geq 0-}$ belongs to \mathcal{S} and it is progressively measurable with respect to the \mathbb{P} -augmentation of the filtration $(\sigma(Z) \vee \mathcal{F}_t^{\xi, W})_{t \geq 0-}$.
- (iii) $\bar{\mu}_\infty$ is a $\sigma(Z)$ -measurable non-negative random variable.

We interpret correlated strategy profiles as in previous chapters: The mediator chooses a correlation device Z to generate a recommendation λ for the representative player. The extraction happens before the game starts, as Z is \mathcal{F}_0 -measurable. Moderator's lottery introduces extra randomness in the game, independently of the initial datum and the noise that determine the random evolution of the representative player's state, as Z is independent of ξ and W . Finally, since $\bar{\mu}_\infty$ is expected to be stochastic only as the result of the mediator's randomization, we request it to be measurable with respect to the correlation device Z .

As for dynamics and payoff, if the representative player decides to trust the mediator and so to follow her recommendation λ , then the dynamics is given by (I.3.1) with λ instead of ν and the payoff is given by $\mathfrak{J}(\lambda, \bar{\mu}_\infty)$, with \mathfrak{J} defined by (I.3.2). If instead the representative player chooses to deviate, she uses a strategy $\nu \in \mathbb{A}$, her dynamics is given by (I.3.1), and she gets the reward $\mathfrak{J}(\nu, \bar{\mu}_\infty)$. Observe that, when the representative player deviates, her strategy ν is adapted to the Brownian filtration and therefore independent of $\bar{\mu}_\infty$, since she has no information on the outcome of moderator's lottery: the deviating player can only use her knowledge of the law of the correlated stationary strategy $(Z, \lambda, \bar{\mu}_\infty)$, which is assumed to be publicly known.

The equilibrium property in the stationary framework is as follows: A correlated stationary strategy $(Z, \lambda, \bar{\mu}_\infty)$ is a coarse correlated equilibrium for the ergodic MFG if the following holds:

- *Optimality*: for every open loop strategy $\nu \in \mathbb{A}$, it holds

$$\mathfrak{J}(\lambda, \bar{\mu}_\infty) \geq \mathfrak{J}(\nu, \bar{\mu}_\infty).$$

- *Consistency*: The process X^λ admits a stationary distribution and it holds

$$\bar{\mu}_\infty = \int_{\mathbb{R}_+} xp_\infty(dx, \bar{\mu}_\infty),$$

where p_∞ is the stochastic kernel so that $\theta_\infty(dx, dm) = p_\infty(dx, m)\rho(dm)$ with $\rho = \mathbb{P} \circ \bar{\mu}_\infty^{-1}$ and $\theta_\infty = \lim_{t \rightarrow \infty} \mathbb{P} \circ (X_t^\lambda, \bar{\mu}_\infty)^{-1}$ in the weak sense.

The consistency condition should be read in the usual way: the mediator imagines what the stationary mean $\bar{\mu}_\infty$ will be, before the game starts, and gives a recommendation to each player according to her idea. If all players commit to the mediator's lottery for generating recommendations, then the long-time average should be consistent with what imagined by the mediator. As in previous chapters, the definition of mean field CCE for the ergodic MFG is consistent with the one of mean NE (see Definition 3.2.4): any mean field NE is a mean field CCE with deterministic stationary mean and $\mathbb{F}^{\xi, W}$ -adapted correlated strategy. We note that considering CCEs in this context is particularly important, since the potential presence of multiple Nash equilibria leads to the question of how players can coordinate towards one of them.

The first contribution of the chapter is to provide sufficient conditions for the existence of coarse correlated equilibria for the competitive mean field game of singular controls. This is the content of Section 3.4. To this extent, we apply the methodology developed in Chapter 2 for linear-quadratic MFGs to the stationary MFG of singular

control. We consider two classes of correlated stationary strategies: while in both classes the correlating device Z is the random mean $\bar{\mu}_\infty$ itself, in the first class the recommendation λ^r is a $\sigma(\bar{\mu}_\infty)$ -measurable regular (absolutely continuous) control, while in the second one, the recommendation λ^s is a policy of reflection type at a random barrier $a(\bar{\mu}_\infty)$. The sufficient conditions for correlated stationary strategies in each of the two classes involve only the distribution of the random stationary mean $\bar{\mu}_\infty$, and they are highly non-linear. Surprisingly, the sufficient conditions of the two classes differ only by a constant. With respect to Chapter 2, we move a step forward, and show that every CCE in these classes induces a sequence of approximate CCEs in the underlying N -player game with vanishing error. This result is in the same spirit as the approximation result in Chapter 1: we use the random stationary mean to act as a correlation device between players' strategies in the N -player game, and we explicitly build the sequence of recommendations to the N players.

Secondly, we completely characterize the Nash equilibria in the ergodic mean field game and prove that their existence and uniqueness depend on the strength of the strategic complementarity, measured by the parameter $\beta \in (0, 1)$. In particular, if $0 < \beta < 1 - \alpha$ or $1 - \alpha < \beta < 1$, then a unique Nash mean field equilibrium exists, where the state process is reflected upwards à la Skorohod at an explicitly given barrier. On the other hand, if $\beta = 1 - \alpha$, either infinitely many equilibria exist, each of reflecting type, or none exist. The existence of multiple equilibria is related to the fact that the mean field game under study faces strategic complementarities (see, again, [52, 53]).

Our last contributions regards the solution of the stationary MFC problem, which is tackled in Section 3.5. We construct the mean field control solution using a Lagrange-multiplier approach, which transforms the original McKean-Vlasov control problem into a two-stage optimization problem: We first restrict to strategies so that the corresponding stationary mean is equal to some prescribed level $m \geq 0$, we compute the optimal strategy within this smaller constrained set and finally we optimize over all possible values of the stationary mean. A somehow similar approach has been used in [45], for a MFC problem of impulse control. Then, we show that the MFC solution can approximate the solution to a central planner problem aiming to achieve Pareto efficiency in the game with N players. To prove the approximation result, we follow the approach of [37, Section 6], although we use different techniques due to the nature of our dynamics and payoff. It is noteworthy that the probabilistic representation of the Lagrange multiplier as the derivative of the mean field control problem's value function with respect to the mean field parameter is the key ingredient for suitably applying the Law of Large Numbers and completing the proof of the approximation result.

Finally, we compare the three notions of mean field equilibria numerically. This is the content of Section 3.6. Since the optimality conditions for the mean field CCEs are non-linear in moments of the correlation device, and present intricate dependence on the model's parameters, we restrict to a specific choice of parameters. We assume that the correlation device has Gamma distribution, and we provide infinitely many instances of Gamma distributions under which coarse correlated equilibria may exist. The analysis is carried out as the strength of the strategic complementarity, measured by the parameter $\beta \in (0, 1)$, varies. In the case where there exist both a unique mean

field NE and a unique MFC solution, we show that a coarse correlated equilibrium with a recommendation in the form of a regular control can outperform the Nash equilibrium, whose equilibrium policy is instead of a reflecting type (and thus singularly continuous), while it can not outperform the MFC solution reward. Finally, we show that infinitely many mean field CCEs may exist even when NE do not, highlighting a feature already explored in Chapter 2.

Despite the specific setting in which the game is formulated (geometric Brownian dynamics and profit function of power type), the analysis reveals a rich structure of the solution while also requiring technical results and arguments. Among these, we highlight as well as the probabilistic representation of the Lagrange multiplier employed in the analysis of the mean field central planner control problem of Section 3.5, as well as the derivation of novel first-order conditions for optimality in ergodic singular stochastic control problems in Section 3.3, which are of independent interest. In this sense, the results in this chapter bring contributions also to the literature on one-dimensional singular stochastic control problems with ergodic performance criteria. Among others, we refer to [5, 15, 46, 78, 81, 110] where explicit solutions to ergodic bounded-variation stochastic control problems have been derived.

Chapter 1

Mean field games with general dynamics and costs

In the framework of continuous time symmetric stochastic differential games in open loop strategies, we introduce a generalization of mean field game solution, called coarse correlated solution. This can be seen as the analogue of a coarse correlated equilibrium in the N -player game. We justify our definition by showing that a coarse correlated solution for the mean field game induces a sequence of approximate coarse correlated equilibria with vanishing error for the underlying N -player games. Existence of coarse correlated solutions for the mean field game is proved by a minimax theorem. An example with explicit solutions is discussed as well.

1.1 Standing assumptions

We here state the standing assumptions on the state dynamics and on the costs of the players in both the N -player game and the limit game, which will be enforced throughout the whole chapter.

We are given a finite time horizon $T > 0$, a control actions space A , an initial state distribution $\eta \in \mathcal{P}(\mathbb{R}^d)$, and the following functions:

$$\begin{aligned}(b, f) &: [0, T] \times \mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \times A \rightarrow \mathbb{R}^d \times \mathbb{R}, \\ g &: \mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \rightarrow \mathbb{R},\end{aligned}$$

which will be referred to, respectively, as the drift function, the running cost and the terminal cost. The following Assumptions **A** will be in force throughout the whole manuscript.

Assumptions **A**.

(A.1) $A \subseteq \mathbb{R}^l$, for some $l \geq 1$, is a compact set.

(A.2) $\eta \in \mathcal{P}^{\bar{p}}(\mathbb{R}^d)$, for some $\bar{p} > 4$.

(A.3) The functions b , f and g are jointly measurable in (t, x, m, a) .

(A.4) $b(t, x, m, a)$ is Lipschitz in $a \in A$, $m \in \mathcal{P}^2(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$, uniformly in t :

$$|b(t, x, m, a) - b(t, x', m', a')| \leq L (|a - a'| + |x - x'| + \mathcal{W}_{2, \mathbb{R}^d}(m, m'))$$

for every $t \in [0, T]$, (x, m, a) and (x', m', a') in $\mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \times A$.

(A.5) The functions $[0, T] \ni t \mapsto (b, f)(t, 0, \delta_0, a_0)$ are bounded, for some $a_0 \in A$ and $\delta_0 \in \mathcal{P}^2(\mathbb{R}^d)$.

(A.6) f and g are locally Lipschitz in (x, m, a) for every fixed $t \in [0, T]$ with at most quadratic growth, i.e., there exists a positive constant $L > 0$ so that

$$\begin{aligned} & |(f, g)(t, x, m, a) - (f, g)(t, x', m', a')| \\ & \leq L \left(1 + |x| + |x'| + \left(\int_{\mathbb{R}^d} |y|^2 m(dy) \right)^{\frac{1}{2}} + \left(\int_{\mathbb{R}^d} |y|^2 m'(dy) \right)^{\frac{1}{2}} + |a| + |a'| \right) \\ & \quad \cdot (|x - x'| + \mathcal{W}_{2, \mathbb{R}^d}(m, m') + |a - a'|), \end{aligned}$$

for every $t \in [0, T]$, (x, m, a) and (x', m', a') in $\mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \times A$.

Assumptions (A.3-6) are fairly standard in MFG literature, see e.g. [36]. In particular, they yield existence and uniqueness of strong solutions to the equations governing the dynamics, as well as finiteness and sufficient regularity of the cost functionals. Assumptions (A.1-2) enable us to apply compactness arguments in the proof of our main results, and they are particularly relevant for the existence in result in Section 1.5.

1.2 Formulation of the N -player game

Consider the following canonical space

$$\Omega^1 = \prod_1^\infty (\mathbb{R}^d \times \mathcal{C}^d), \quad \mathcal{F}^1 = \bigotimes_1^\infty (\mathcal{B}_{\mathbb{R}^d} \otimes \mathcal{B}_{\mathcal{C}^d}), \quad \mathbb{P}^1 = \bigotimes_1^\infty (\eta \otimes \mathbb{W}^d). \quad (1.2.1)$$

We define a sequence of random variables $(\xi^i)_{i \geq 1}$ and of Brownian motions $(W^i)_{i \geq 1}$, by taking the projections:

$$\xi^i(\omega_1) = \xi^i((x^j, w^j)_{j \geq 1}) = x^i, \quad W_t^i(\omega_1) = W_t^i((x^j, w^j)_{j \geq 1}) = w_t^i, \quad t \in [0, T]. \quad (1.2.2)$$

By definition of \mathbb{P}^1 , $(\xi^i)_{i \geq 1}$ and $(W^i)_{i \geq 1}$ are mutually independent, $(\xi^i)_{i \geq 1}$ are independent and identically distributed with law $\eta \in \mathcal{P}^{\bar{\mathcal{P}}}(\mathbb{R}^d)$ and $(W^i)_{i \geq 1}$ are independent d -dimensional standard Brownian motions.

Let $N \in \mathbb{N}$, $N \geq 2$, be the number of players. We define the filtration $\mathbb{F}^{1, N}$ as the \mathbb{P}^1 -augmentation of the filtration generated by the first N random variables $(\xi^i)_{i=1}^N$ and Brownian motions $(W^i)_{i=1}^N$. Therefore, for the N -player game, we work on the space

$$(\Omega^1, \mathcal{F}^1, \mathbb{F}^{1, N}, \mathbb{P}^1). \quad (1.2.3)$$

We stress that, for every $N \geq 2$, we keep the probability space $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$ fixed while the filtration $\mathbb{F}^{1,N}$ varies.

Consider the set \mathbb{A}_N of $\mathbb{F}^{1,N}$ -progressively measurable processes taking values in A :

$$\mathbb{A}_N = \left\{ \alpha : [0, T] \times \Omega^1 \rightarrow A \mid \alpha \text{ is } \mathbb{F}^{1,N}\text{-progressively measurable} \right\}. \quad (1.2.4)$$

Provided that we identify processes which are equal $Leb_{[0,T]} \otimes \mathbb{P}^1$ -a.e., we can regard \mathbb{A}_N as

$$\mathbb{A}_N = L^2([0, T] \times \Omega^1, \mathcal{R}^{1,N}, Leb_{[0,T]} \otimes \mathbb{P}^1; A),$$

where $\mathcal{R}^{1,N}$ stands for the progressive σ -algebra on $[0, T] \times \Omega^1$, using the filtration $\mathbb{F}^{1,N}$. We call any element $\alpha \in \mathbb{A}_N$ an open loop strategy for the N -player game. We regard a vector $(\alpha^1, \dots, \alpha^N) \in \mathbb{A}_N^N = \times_1^N \mathbb{A}_N$ as an open loop strategy profile for the N players, which will be occasionally denoted by α . We endow such a space \mathbb{A}_N with the norm

$$\|\alpha\|_{L^2} = \mathbb{E}^{\mathbb{P}^1} \left[\int_0^T |\alpha_t|^2 dt \right]^{\frac{1}{2}} \quad (1.2.5)$$

and consider the Borel σ -algebra $\mathcal{B}_{\mathbb{A}_N}$ associated to that. We observe that, since $([0, T] \times \Omega^1, \mathcal{B}_{[0,T] \times \Omega^1})$ is Polish and A is closed, \mathbb{A}_N is a separable Banach space. In the following, we will make no distinction between an $\mathbb{F}^{1,N}$ -progressively measurable process α and any other process α' which is equal to it $Leb_{[0,T]} \otimes \mathbb{P}^1$ -almost everywhere.

Definition 1.2.1 (Recommendation profile). We call *recommendation profile* to the N players a pair $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ so that the following holds:

1. $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is a complete probability space; Ω^0 is a Polish space and \mathcal{F}^{0-} is its corresponding Borel σ -algebra.
2. $\Lambda = (\Lambda^1, \dots, \Lambda^N)$ is a random vector with values in \mathbb{A}_N^N :

$$\begin{aligned} \Lambda : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) &\longrightarrow (\mathbb{A}_N^N, \mathcal{B}_{\mathbb{A}_N^N}) \\ \omega_0 &\longmapsto \Lambda(\omega_0) = (\alpha^1, \dots, \alpha^N) : [0, T] \times \Omega^1 \rightarrow A^N. \end{aligned} \quad (1.2.6)$$

We interpret the recommendation profile as follows: A correlation device or a mediator runs a lottery over open loop strategy profiles according to some publicly known distribution \mathbb{P}^0 and communicates privately to each player a strategy according to the selected profile. The extraction of the strategy profile happens before the game starts and it is independent of the idiosyncratic shocks that determine the random evolution of players' states.

For a given recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$, we build a probability space large enough to support both the moderator's recommendation profile Λ and the random shocks and initial data $(\xi^j, W^j)_{j=1}^N$, in such a way that they are independent. On this new probability space, we then consider the strategy profile associated to the recommendation Λ , that is, the strategies that the committing players will use. These strategies will depend both on the realisation of moderator's lottery $\Lambda^i(\omega_0)$ and the

initial data and idiosyncratic noise that determine the evolution of players' states. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be defined by

$$(\Omega, \mathcal{F}, \mathbb{P}) = (\Omega^0 \times \Omega^1, \mathcal{F}^{0-} \otimes \mathcal{F}^1, \mathbb{P}^0 \otimes \mathbb{P}^1). \quad (1.2.7)$$

We complete the σ -algebra \mathcal{F} with the \mathbb{P} -null sets and endow the product probability space with the \mathbb{P} -augmentation of the filtration

$$\mathbb{F} = \mathcal{F}^{0-} \otimes \mathbb{F}^{1,N} = (\mathcal{F}^{0-} \otimes \mathcal{F}_t^{1,N})_{t \in [0, T]}.$$

On the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ we would like to consider the strategy profile associated to a recommendation profile Λ , by setting

$$\lambda_t^i(\omega) = \lambda_t^i(\omega_0, \omega_1) = \Lambda^i(\omega_0)_t(\omega_1), \quad i = 1, \dots, N. \quad (1.2.8)$$

A priori, the process λ constructed may not be progressively measurable, for instance when a recommendation profile Λ takes uncountably many values. The essential reason is that we cannot deduce the measurability of a set in the product σ -algebra from the measurability of its sections, as shown, e.g., in [103, p. 5]. For this reason, we have the following admissibility definition:

Definition 1.2.2 (Admissible recommendation profile). A recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is *admissible* if there exists a process $\lambda = (\lambda_t^1, \dots, \lambda_t^N)_{t \in [0, T]}$ with values in A^N , defined on the product space $(\Omega, \mathcal{F}, \mathbb{P})$ and \mathbb{F} -progressively measurable, so that, for every $i = 1, \dots, N$, it holds

$$\|(\lambda_t^i(\omega_0, \cdot))_{t \in [0, T]} - \Lambda^i(\omega_0)\|_{L^2([0, T] \times \Omega^1)} = 0, \quad \mathbb{P}^0\text{-a.s.}, \quad i = 1, \dots, N. \quad (1.2.9)$$

Any such process λ will be called *strategy profile associated to the admissible recommendation profile* $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$.

Observe that the strategy profile λ is the result of both moderator's recommendation and the noises, and it will appear in the dynamics and the cost of the committing players. We remark that, by Proposition 1.7.2, given any admissible recommendation to the N players $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$, the strategy profile λ associated to it is unique $Leb_{[0, T]} \otimes \mathbb{P}$ -almost everywhere. We give some examples of admissible recommendations in Example 1.3.1 in the following Section 1.3.

Remark 1.2.1. As usual, we can extend random variables defined on Ω^1 to random variables defined on Ω . Indeed, suppose $X : (\Omega^1, \mathcal{F}^1) \rightarrow (E, \mathcal{E})$ is a random variable with values in some measurable space (E, \mathcal{E}) . We can then regard X as defined on the space (Ω, \mathcal{F}) via the identification $\tilde{X}(\omega_0, \omega_1) = X(\omega_1)$, and analogously for Ω^0 . In this sense, via the identification $(\tilde{\xi}^i, \tilde{W}^i)(\omega_0, \omega_1) = (\xi^i, W^i)(\omega_1)$ for every $i = 1, \dots, N$, we can regard the Brownian motions and initial data as defined on Ω ; we observe that $(W^i)_{i=1}^N$ are independent standard Brownian motions with respect to the filtration \mathbb{F} as well. Moreover, we can identify each process $\alpha \in \mathbb{A}_N$, which is defined on Ω^1 , with a process $\tilde{\alpha}$ defined on Ω via the identification $\tilde{\alpha}(\omega_0, \omega_1) = \alpha(\omega_1)$. Such a process is progressively measurable with respect to the filtration \mathbb{F} and independent of \mathcal{F}^{0-} .

Observe that, by construction, we have $\mathcal{F}^{0-} \subseteq \mathcal{F}_t$ for every $t \in [0, T]$, and \mathcal{F}^{0-} is independent of the filtration of noises \mathbb{F}^1 . This models the fact that the extraction of the strategy profile happens before the game starts and is independent of the idiosyncratic shocks that determine the random evolution of players' states. We stress that, by definition, the realization $\Lambda^i(\omega_0)$ is an $\mathbb{F}^{1,N}$ -progressively measurable process in \mathbb{A}_N , for any scenario $\omega_0 \in \Omega^0$ and $i = 1, \dots, N$. Observe that, even though Λ and $(\xi^j, W^j)_{j=1}^N$ are independent, the strategy profile associated to the recommendation $\lambda = (\lambda_t)_{t \in [0, T]}$ is in general not independent of either of them, since it is the result of both the recommendation profile and the random shocks and initial data.

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ be an admissible recommendation profile. On the probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ defined in (1.2.7), we assign players state dynamics and define the cost functionals. If all players follow the recommendation Λ , players' state dynamics are given by the following system of stochastic differential equations:

$$\begin{cases} dX_t^j = b(t, X_t^j, \mu_t^N, \lambda_t^j)dt + dW_t^j, & 0 \leq t \leq T, \\ X_0^j = \xi^j, \end{cases} \quad (1.2.10)$$

for every $j \in \{1, \dots, N\}$, where μ_t^N is the empirical measure of the state processes of all players at time t :

$$\mu_t^N = \frac{1}{N} \sum_{j=1}^N \delta_{X_t^j}. \quad (1.2.11)$$

Suppose player i deviates, while the other players follow the recommendations they receive from the mediator. The deviating player will pick instead an open loop strategy $\beta \in \mathbb{A}_N$. In other words, at every time t and for every scenario ω , player i plays the action $\tilde{\beta}_t(\omega) = \beta_t(\omega_1)$ instead of playing the recommended action $\lambda_t^i(\omega) = \Lambda^i(\omega_0)_t(\omega_1)$. Then, players' state dynamics are given by the following system of stochastic differential equations:

$$\begin{cases} dX_t^j = b(t, X_t^j, \mu_t^N, \lambda_t^j)dt + dW_t^j, & 0 \leq t \leq T, & X_0^j = \xi^j, & j \neq i \\ dX_t^i = b(t, X_t^i, \mu_t^N, \beta_t)dt + dW_t^i, & 0 \leq t \leq T, & X_0^i = \xi^i, \end{cases} \quad (1.2.12)$$

where μ_t^N is defined as in (1.2.11). Assumptions **A** ensure that there always exists an \mathbb{F} -adapted continuous solution to both equations (1.2.10) and (1.2.12) so that

$$\mathbb{E}[\sup_{t \in [0, T]} \max_{1 \leq j \leq N} |X_t^j|^2] < \infty.$$

Moreover, pathwise uniqueness holds so that, by Theorem B.1, uniqueness in law holds as well.

Remark 1.2.2. We notice that there is an asymmetry between the information available to the mediator and the deviating player. If a player deviates, she will use an open loop strategy $\beta \in \mathbb{A}_N$, therefore the information available to the deviating player is just given by the smaller Brownian filtration $\mathbb{F}^{1,N}$. In particular, she will use a strategy which is independent of \mathcal{F}^{0-} , thus of the admissible recommendation profile Λ . This models the fact that if a player deviates, she does not have access to the information

contained in mediator's recommendation: she will not exploit any of the additional information the mediator would give away when communicating the recommended strategies to the players. Conversely, player i 's strategy $\lambda^i = (\lambda_t^i)_{t \in [0, T]}$ associated to the recommendation Λ^i is by definition dependent of the recommendation Λ^i to player i , and it carries at least some of the information the the moderator uses to randomize players' strategies. Notice that, even if a deviation β is independent of the moderators recommendation Λ , it is not independent of the strategy profile λ associated to the recommendation, as both are dependent on the random shocks and initial data $(\xi^j, W^j)_{j=1}^N$ that determine players' dynamics.

As for the cost functional, let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ be an admissible recommendation profile. If all players follow the recommendation, then the cost functional of player $i = 1, \dots, N$ is given by

$$\mathfrak{J}_i^N(\Lambda) = \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^N, \lambda_t^i) dt + g(X_T^i, \mu_T^N) \right],$$

with dynamics given by (1.2.10). If instead player i does not play according to the recommendation Λ^i and plays a different strategy $\beta \in \mathbb{A}_N$, while the other players stick to the recommendation profile Λ^{-i} , we define the cost functional of player i as

$$\mathfrak{J}_i^N(\Lambda^{-i}, \beta) = \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^N, \beta_t) dt + g(X_T^i, \mu_T^N) \right],$$

where the dynamics are given by (1.2.12). We stress that the expectation in the cost functional is taken with respect to the product probability measure $\mathbb{P} = \mathbb{P}^0 \otimes \mathbb{P}^1$, although we omit this dependence for conciseness. Finally, we give the notion of ε -coarse correlated equilibrium:

Definition 1.2.3 (ε -coarse correlated equilibrium). Let $\varepsilon \geq 0$. An admissible recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is an ε -coarse correlated equilibrium for the N -player game (ε -CCE) if

$$\mathfrak{J}_i^N(\Lambda) \leq \mathfrak{J}_i^N(\Lambda^{-i}, \beta) + \varepsilon \tag{1.2.13}$$

for all open loop strategies $\beta \in \mathbb{A}_N$ and all players $i = 1, \dots, N$. We call an admissible recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ a *coarse correlated equilibrium* for the N -player game if it is an ε -coarse correlated equilibrium with $\varepsilon = 0$.

Notice that even if player i deviates, she can actually compute her cost functional $\mathfrak{J}_i^N(\Lambda^{-i}, \beta)$ by only knowing the joint law of the admissible recommendation profile Λ , for any open loop strategy β . Indeed, it is possible to express the equilibrium property (1.2.13) in terms of the law of the recommendation profile Λ on the right-hand side, and in terms of the marginal law of Λ^{-i} on the left-hand side. We refer to the following Remark 1.2.5 for more precise comments.

We compare the definition of coarse correlated equilibria with the more usual notion of Nash equilibria, that we recall:

Definition 1.2.4 (ε -Nash equilibrium). Let $\varepsilon \geq 0$. An open loop strategy profile $\alpha^* = (\alpha^{*,1}, \dots, \alpha^{*,N})$ in \mathbb{A}_N^N is an ε -Nash equilibrium in open loop strategies if

$$\mathfrak{J}_i^N(\alpha^*) \leq \mathfrak{J}_i^N(\alpha^{*,-i}, \beta) + \varepsilon$$

for all open loop strategies $\beta \in \mathbb{A}_N$ and all players $i = 1, \dots, N$. We call an open loop strategy profile $\alpha^* = (\alpha^{*,1}, \dots, \alpha^{*,N})$ in \mathbb{A}_N^N a *Nash equilibrium in open loop strategies* if it is an ε -Nash equilibrium in open loop strategies with $\varepsilon = 0$.

The usual notion of Nash equilibrium in open loop strategies is consistent with the definition of coarse correlated equilibrium: Suppose we are given an ε -Nash equilibrium $\alpha^* = (\alpha^{*,1}, \dots, \alpha^{*,N})$ in open loop strategies. We choose $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ as the trivial probability space and Λ as constant and equal to $(\alpha^{*,1}, \dots, \alpha^{*,N})$. It is then straightforward to see that the triple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is an ε -CCE according to Definition 1.2.3.

Observe that a Nash equilibrium in open loop strategies $(\hat{\alpha}^1, \dots, \hat{\alpha}^N)$ is progressively measurable with respect to the filtration $\mathbb{F}^{1,N}$, while a strategy profile λ associated to an admissible recommendation Λ contains the information carried by Λ itself, which is the information the mediator uses to randomize players' strategies. Moreover, when dealing with Nash equilibria, the deviating player has access to the same information as the other players, since they all use $\mathbb{F}^{1,N}$ -progressively measurable strategies. On the contrary, CCEs present a certain asymmetry between the information available to the committing and the deviating players, as pointed out in Remark 1.2.2.

Remark 1.2.3 (Role of the probability space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$). According to Definition 1.2.2, the probability space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is part of the definition of admissible recommendation. The natural interpretation is that the mediator chooses the auxiliary space he uses to correlate players' strategies. Moreover, according to equations (1.2.10) and (1.2.12), it determines the probability space on which state processes are defined. In order to keep the notation as simple as possible, by abuse of notation, we mostly refer only to Λ as the admissible recommendation instead of the pair $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$.

Remark 1.2.4 (Relationship with correlated equilibria of [20, 29]). It is worth to briefly compare our notion of coarse correlated equilibria with the notion of correlated equilibria of [29] and [20]. In these works, the authors deal with discrete time models with finite sets of individual states and control actions, and consider restricted closed-loop strategies. Most importantly, in their framework, correlated equilibria are considered, and not coarse correlated equilibria: there, each player observes the outcome of moderator's lottery and then decides whether to play it or not. Thus, in that context, if player i deviates, she would use a strategy β which is not a priori independent of the outcome of moderator's lottery, but could depend on the realization $\Lambda^i(\omega_0)$, and she could use that information to choose her deviation. In result, player i 's deviating strategy would be progressively measurable to the filtration generated by both its private recommendation Λ^i and the noises and initial data $(\xi^i, W^i)_{i=1}^N$. This would lead to the additional difficulty of considering a specific information structure for each player, as the private recommendations Λ^i are possibly different one from the

other. On the contrary, in our model, the deviating player chooses a strategy β which is independent of the admissible recommendation profile Λ , as the deviating player has no access to the recommended strategy.

Remark 1.2.5 (Expressing ε -CCE in terms of the law of the admissible recommendation profile). It is possible to restate the equilibrium condition (1.2.13) in terms of the law of Λ , by taking advantage of regular conditional probabilities.

Let Λ be an ε -CCE. Denote by $\gamma_N \in \mathcal{P}(\mathbb{A}_N^N)$ the law of Λ and by $\gamma_N^{-i} \in \mathcal{P}(\mathbb{A}_N^{N-1})$ the law of Λ^{-i} . Since both Ω^0 and Ω^1 are Polish, so are \mathbb{A}_N^N and (Ω, \mathcal{F}) , so that there exists a version of the regular conditional probability of \mathbb{P} given $\Lambda = (\alpha^1, \dots, \alpha^N)$, which we denote by \mathbb{P}^α . By conditioning on $\sigma(\Lambda)$, we get that the process X solution to (1.2.10) satisfies the equation

$$\begin{cases} dX_t^j = b(t, X_t^j, \mu_t^N, \alpha_t^j)dt + dW_t^j, & 0 \leq t \leq T, \\ X_0^j = \xi^j, \end{cases}$$

on the probability space $(\Omega, \mathcal{F}, \mathbb{P}^\alpha)$, for γ_N -a.e. $\alpha \in \mathbb{A}_N^N$, under \mathbb{P}^α . To see this, observe that $(W^j)_{j=1}^N$ are independent \mathbb{F} -Brownian motions under \mathbb{P}^α as well, since by construction $(W^j)_{j=1}^N$ and Λ are independent under \mathbb{P} , and it holds $\mathbb{P}^\alpha(\Lambda = \alpha) = 1$, which implies

$$\mathbb{P}^\alpha(b(t, x, m, \lambda_t^j) = b(t, x, m, \alpha_t^j), \forall x \in \mathbb{R}, \forall j = 1, \dots, N) = 1, \text{ Leb}_{[0, T]}\text{-a.e. } t \in [0, T],$$

for γ_N -a.e. $\alpha \in \mathbb{A}_N^N$. This is enough to conclude that X satisfies equation (1.2.10) under \mathbb{P}^α and that the equality

$$\mathfrak{J}_i^N(\Lambda) = \int_{\mathbb{A}_N^N} \mathfrak{J}_i^N(\alpha) \gamma_N(d\alpha),$$

holds. Here, $\mathfrak{J}_i^N(\alpha)$ stands for the cost functional associated to an admissible recommendation Λ \mathbb{P}^0 -a.s. equal to $\alpha \in \mathbb{A}_N^N$. For any deviation $\beta \in \mathbb{A}_N$, the same reasoning applies to equation (1.2.12) and to the cost $\mathfrak{J}_i^N(\Lambda^{-i}, \beta)$. Then, the equilibrium condition (1.2.13) reads as

$$\int_{\mathbb{A}_N^N} \mathfrak{J}_i^N(\alpha) \gamma_N(d\alpha) \leq \int_{\mathbb{A}_N^{N-1}} \mathfrak{J}_i^N(\alpha^{-i}, \beta) \gamma_N^{-i}(d\alpha^{-i}) + \varepsilon \quad \forall \beta \in \mathbb{A}_N, \quad i = 1, \dots, N. \quad (1.2.14)$$

We stress that the right-hand side depends only on the marginal law of Λ^{-i} , for every $\beta \in \mathbb{A}_N$, for every $i = 1, \dots, N$, so that this shows makes rigorous the expression of the equilibrium property (1.2.13) in terms of the joint law of the recommendation profile. Moreover, although we opted for a formulation of CCEs in terms of random variables, condition (1.2.14) has already been used in the literature (see, e.g., [57, 91, 92] in the static case and [97, 98] in the differential one) to formulate the CCE property.

1.3 Formulation of the mean field game

Consider the following canonical space

$$\Omega^* = \mathbb{R}^d \times \mathcal{C}^d, \quad \mathcal{F}^* = \mathcal{B}_{\mathbb{R}^d} \otimes \mathcal{B}_{\mathcal{C}^d}, \quad \mathbb{P}^* = \eta \otimes \mathbb{W}^d. \quad (1.3.1)$$

Define ξ and $W = (W_t)_{t \in [0, T]}$ as

$$\xi(\omega_*) = \xi(x, w) = x, \quad W_t(\omega_*) = W_t(x, w) = w_t. \quad (1.3.2)$$

By definition of \mathbb{P}^* , ξ and W are independent, ξ is an \mathbb{R}^d -valued random variable with law η and W is a standard Brownian motion. Define the filtration \mathbb{F}^* as the \mathbb{P}^* -augmentation of the filtration generated by ξ and W .

Consider the set \mathbb{A} of \mathbb{F}^* -progressively measurable processes taking values in A :

$$\mathbb{A} = \left\{ \alpha : [0, T] \times \Omega^* \rightarrow A \mid \alpha \text{ is } \mathbb{F}^*\text{-progressively measurable} \right\}. \quad (1.3.3)$$

Provided that we identify processes which are equal $Leb_{[0, T]} \otimes \mathbb{P}^*$ -a.e., we can regard \mathbb{A} as

$$\mathbb{A} = L^2([0, T] \times \Omega^*, \mathcal{R}^*, Leb_{[0, T]} \otimes \mathbb{P}^*; A),$$

where \mathcal{R}^* stands for the progressive σ -algebra on $[0, T] \times \Omega^*$, using the filtration \mathbb{F}^* . We call any element $\alpha \in \mathbb{A}$ an open loop strategy for the mean field game. We endow such a space \mathbb{A} with the norm

$$\|\alpha\|_{L^2} = \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T |\alpha_t|^2 dt \right]^{\frac{1}{2}} \quad (1.3.4)$$

and consider the Borel σ -algebra $\mathcal{B}_{\mathbb{A}}$ associated to that. We observe that, since $([0, T] \times \Omega^*, \mathcal{B}_{[0, T] \times \Omega^*})$ is Polish and A is closed, \mathbb{A} is a separable Banach space. Finally, we will make no distinction between an \mathbb{F}^* -progressively measurable process α and any other process α' which is equal to it $Leb_{[0, T]} \otimes \mathbb{P}^*$ -almost everywhere.

As in the N -player game, we define the recommendation to the representative player, and then the admissibility requirement.

Definition 1.3.1 (Recommendation for the mean field game). We call *recommendation* a pair $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ where:

1. $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is a complete probability space; Ω^0 is a Polish space and \mathcal{F}^{0-} is its corresponding Borel σ -algebra.
2. Λ is a random variable with values in \mathbb{A} :

$$\begin{aligned} \Lambda : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) &\longrightarrow (\mathbb{A}, \mathcal{B}_{\mathbb{A}}) \\ \omega_0 &\longmapsto \Lambda(\omega_0) = \alpha : [0, T] \times \Omega^* \rightarrow A. \end{aligned} \quad (1.3.5)$$

For a given recommendation $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$, we build a probability space large enough to support both the moderator's recommendation Λ and the random shocks and initial datum (ξ^*, W^*) , in such a way that they are independent. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be defined by

$$(\Omega, \mathcal{F}, \mathbb{P}) = (\Omega^0 \times \Omega^*, \mathcal{F}^{0-} \otimes \mathcal{F}^*, \mathbb{P}^0 \otimes \mathbb{P}^*). \quad (1.3.6)$$

We complete the σ -algebra \mathcal{F} with the \mathbb{P} -null sets and endow the product probability space with the \mathbb{P} -augmentation of the filtration

$$\mathbb{F} = \mathcal{F}^{0-} \otimes \mathbb{F}^* = (\mathcal{F}^{0-} \otimes \mathcal{F}_t^*)_{t \in [0, T]}.$$

On the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ we would like to consider the strategy profile associated to a recommendation profile Λ , by setting

$$\lambda_t(\omega) = \lambda_t(\omega_0, \omega_*) = \Lambda(\omega_0)_t(\omega_*), \quad (1.3.7)$$

As in Section 1.2, in order for such process λ to be progressively measurable, we have the following admissibility definition:

Definition 1.3.2 (Admissible recommendation for the mean field game). A recommendation for the mean field game $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is admissible if there exists an A -valued process $\lambda = (\lambda_t)_{t \in [0, T]}$, defined on $(\Omega, \mathcal{F}, \mathbb{P})$ and \mathbb{F} -progressively measurable, so that it holds

$$\|(\lambda_t(\omega_0, \cdot))_{t \in [0, T]} - \Lambda(\omega_0)\|_{L^2([0, T] \times \Omega^*)} = 0, \quad \mathbb{P}^0\text{-a.s.} \quad (1.3.8)$$

Any such process λ will be called *strategy associated to the admissible recommendation* $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$.

We remark that, by Proposition 1.7.2, given any admissible recommendation $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$, the strategy λ associated to it is unique $Leb_{[0, T]} \otimes \mathbb{P}$ -almost everywhere.

Definition 1.3.3 (Correlated measure flow). A *correlated measure flow* is a triple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ where:

1. $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is an admissible recommendation.
2. $\mu : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) \rightarrow (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)})$ is a random continuous flow of measures in $\mathcal{P}^2(\mathbb{R}^d)$.

The same considerations as in Remark 1.2.1 about the extension of random variables on the product space $(\Omega, \mathcal{F}, \mathbb{P})$ hold for correlated measure flows as well.

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ be a correlated measure flow. On the product probability space $(\Omega, \mathcal{F}, \mathbb{P})$ defined in (1.3.6), we assign state dynamics. If the representative player decides to play according to the admissible recommendation Λ , the dynamics is given by the following SDE:

$$\begin{cases} dX_t = b(t, X_t, \mu_t, \lambda_t)dt + dW_t, & 0 \leq t \leq T, \\ X_0 = \xi. \end{cases} \quad (1.3.9)$$

If instead the representative player decides to ignore the mediator's recommendation and to use a possibly different strategy $\beta \in \mathbb{A}$, the dynamics is given by the following SDE:

$$\begin{cases} dX_t = b(t, X_t, \mu_t, \beta_t)dt + dW_t, & 0 \leq t \leq T, \\ X_0 = \xi. \end{cases} \quad (1.3.10)$$

By Assumptions **A**, on any space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ there exists a solution to equation (1.3.9) and pathwise uniqueness holds. By Theorem B.1, uniqueness in law holds. Analogous considerations apply to equation (1.3.10).

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ be a correlated measure flow. The cost functionals for the representative player and the deviating player, whose state dynamics follow (1.3.9) and (1.3.10), respectively, are given by:

$$\begin{aligned}\mathfrak{J}(\Lambda, \mu) &= \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \lambda_t) dt + g(X_T, \mu_T) \right], \\ \mathfrak{J}(\beta, \mu) &= \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \beta_t) dt + g(X_T, \mu_T) \right].\end{aligned}\tag{1.3.11}$$

As in the N -player game, the expectation in the cost functional is taken with respect to the product probability measure $\mathbb{P} = \mathbb{P}^0 \otimes \mathbb{P}^*$; in particular, it depends on the mediator's randomization \mathbb{P}^0 also when the representative player deviates. Finally, we give the definition of coarse correlated solution of the mean field game:

Definition 1.3.4 (Coarse correlated solution). A correlated measure flow $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ is a *coarse correlated solution* of the mean field game if the following properties hold:

- (i) Optimality: for every deviation $\beta \in \mathbb{A}$, it holds

$$\mathfrak{J}(\Lambda, \mu) \leq \mathfrak{J}(\beta, \mu).\tag{1.3.12}$$

- (ii) Consistency: for every time $t \in [0, T]$, μ_t is a version of the conditional law of X_t given μ , that is,

$$\mu_t(\cdot) = \mathbb{P}(X_t \in \cdot \mid \mu) \quad \mathbb{P}\text{-a.s.} \quad \forall t \in [0, T].\tag{1.3.13}$$

We will refer to coarse correlated solutions of the mean field game as coarse correlated mean field solutions and mean field coarse correlated equilibria (CCE) as well.

Remark 1.3.1 (Role of $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$). Analogously as in the N -player game, although the probability space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is part of the definitions of admissible recommendation and correlated measure flow, when it is clear from the context we refer to Λ and (Λ, μ) , instead of the pair $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ and the triple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$, as admissible recommendation and correlated measure flow, respectively.

As in [20, 29], the consistency condition (1.3.13) should be read in the following way: the mediator imagines what the flow of measures will be, up to the terminal horizon T , before the game starts, and gives a recommendation to each player according to his idea. Since the flow of measures is expected to be stochastic as a result of the mediator's randomization only, we request it to be measurable with respect to \mathcal{F}^{0-} , and, since the randomization is performed before the game starts, we have $\mathcal{F}^{0-} \subseteq \mathcal{F}_t$ for any $t \geq 0$. If all players commit to the mediator's lottery for generating recommendations, then the flow of measures should arise from aggregation of the individual behaviors, consistently with what imagined by the mediator. Since the generation of the recommendation is performed on the basis of the whole flow of measures, we formulate consistency condition (1.3.13) with respect to conditioning on the whole flow. Regarding the strategy of the deviating player, as in the N -player game,

if the player deviates, she chooses her strategy on her own, without using any of the information carried by Λ or μ : the only information she has about Λ or μ comes from the knowledge of their joint law, which is assumed to be known by the representative player, in analogy to the N -player game.

Analogously to Nash equilibria in the N -player game, coarse correlated solutions extend the more common notion of mean field game solutions, that we recall:

Definition 1.3.5 (MFG solution). A pair $(\alpha^*, \mu^*) \in \mathbb{A} \times \mathcal{C}(\mathcal{P}^2)$ is said to be a *MFG solution* if the following properties are satisfied:

- (i) Optimality: α^* minimizes $\mathfrak{J}(\cdot, \mu^*)$ over \mathbb{A} , i.e.

$$\mathfrak{J}(\alpha^*, \mu^*) = \min_{\beta \in \mathbb{A}} \mathfrak{J}(\beta, \mu^*). \quad (1.3.14)$$

- (ii) Consistency: let X^* be the state process controlled by α^* . For every time $t \in [0, T]$, μ_t^* is the law of the state X_t^* , that is,

$$\mu_t^*(\cdot) = \mathbb{P}(X_t^* \in \cdot) \quad \forall t \in [0, T]. \quad (1.3.15)$$

We will refer to MFG solutions as mean field Nash equilibria (NE) as well.

Any MFG solution is a mean field CCE as well: indeed, it is enough to choose the trivial probability space for $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$, to take μ deterministic equal to μ^* and to make the moderator recommend with probability one to play according to α^* , i.e. $\Lambda = \alpha^* \mathbb{P}^0$ -a.s.

Remark 1.3.2 (Relation with MFGs with common noise). Given the consistency condition (1.3.13), it is worth comparing coarse correlated solutions to the MFG and solutions to MFGs with common noise (see, e.g., [41, 85, 88] and in [39]). In the latter, the flow of measures is stochastic due to a common noise that equally impacts the state dynamics of all players in the underlying N -player game. As a consequence, the flow of measures is expected to be adapted to the filtration generated by the common noise (the so called *strong solutions*); if this is not the case, compatibility conditions between the noises and the flow of measures itself are needed in order to guarantee that the flow μ picks into the future in a minimal way (the so called *weak solutions*). In the case of a coarse correlated solution to the MFG, on the other hand, the flow of measures is expected to be stochastic as a result of the mediator's randomization only, which is generated before the beginning of the game. More formally, this implies that the flow of measure is \mathcal{F}^{0-} -measurable with $\mathcal{F}^{0-} \subseteq \mathcal{F}_t$ for any $t \geq 0$. Recommendations to the representative player are given according to the mediator's idea of the whole flow, which leads to the consistency condition with conditioning with respect to the whole flow up to terminal time. In this sense, the mediator sees into the future, and consequently no compatibility condition is needed.

One might be tempted to regard the randomness driving the mediator's lottery for selecting recommendations as a common noise that affects the state dynamics only through the control. There are at least two major differences though: First, such a common noise will have no impact on the controls of a deviating player. To put

it differently, only the pre-committing players' dynamics are directly affected by the mediator's lottery over strategy profiles. Second, such a common noise would not be exogenous; instead, it is built into the correlation device used by the mediator, as represented by the auxiliary probability space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$, and as such is part of the solution.

Remark 1.3.3 (Expressing mean field CCEs in terms of the law of the correlated measure flow). Analogously to Remark 1.2.5, we show that it is possible to express the equilibrium property (1.3.12) in terms of the law of the correlated measure flow.

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ be a mean field CCE. Denote by $\gamma \in \mathcal{P}(\mathbb{A} \times \mathcal{C}(\mathcal{P}^2))$ the joint law of (Λ, μ) and by $\rho \in \mathcal{P}(\mathcal{C}(\mathcal{P}^2))$ the marginal law of μ . Since both $(\Omega^0, \mathcal{F}^{0-})$ and $(\Omega^1, \mathcal{F}^1)$ are Polish spaces, so are \mathbb{A} and (Ω, \mathcal{F}) , so that there exists a version of the regular conditional probability of \mathbb{P} given $(\Lambda, \mu) = (\alpha, m)$, which we denote by $\mathbb{P}^{\alpha, m}$. As in Remark 1.2.5, by conditioning on $\sigma(\Lambda, \mu)$, we get that X is a solution to the SDE

$$\begin{cases} dX_t = b(t, X_t, m_t, \alpha_t)dt + dW_t, & 0 \leq t \leq T, \\ X_0 = 0, \end{cases} \quad (1.3.16)$$

on the probability space $(\Omega, \mathcal{F}, \mathbb{P}^{\alpha, m})$, for γ -a.e. $(\alpha, m) \in \mathbb{A} \times \mathcal{C}(\mathcal{P}^2)$ and that the equality

$$\mathfrak{J}(\Lambda, \mu) = \int_{\mathbb{A} \times \mathcal{P}(\mathcal{C}^d)} \mathfrak{J}(\alpha, m) \gamma(d\alpha, dm),$$

holds. Here, $\mathfrak{J}(\alpha, m)$ stands for the cost functional associated to a correlated measure flow (Λ, μ) \mathbb{P}^0 -a.s. equal to $(\alpha, m) \in \mathbb{A} \times \mathcal{C}(\mathcal{P}^2)$. Analogous reasoning holds for $\mathfrak{J}(\beta, \mu)$, except that the conditioning on $\sigma(\Lambda, \mu)$ does not modify the dependence upon the deviating strategy $\beta \in \mathbb{A}$. Therefore, the optimality condition (1.3.12) reads as

$$\int_{\mathbb{A} \times \mathcal{C}(\mathcal{P}^2)} \mathfrak{J}(\alpha, m) \gamma(d\alpha, dm) \leq \int_{\mathcal{C}(\mathcal{P}^2)} \mathfrak{J}(\beta, m) \rho(dm), \quad \forall \beta \in \mathbb{A}.$$

Comparing this inequality with (1.2.14), we observe that the dependence upon the non-deviating players in the left hand-side of (1.2.14) is replaced by the law of the flow of measures μ , which captures the behavior of the population.

Example 1.3.1 (Admissible recommendations). Fix a complete probability space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$. We provide some simple examples of random variables Λ defined on $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ with values in $(\mathbb{A}, \mathcal{B}_{\mathbb{A}})$ which are admissible recommendations in the sense of Definition 1.3.2.

1. Suppose that Λ takes only finitely many values, say $\alpha^1, \dots, \alpha^k \in \mathbb{A}$, $k \geq 1$, i.e. $\mathbb{P}^0(\Lambda = \alpha^i) = p_i$, with $p_i \geq 0$ for every $i = 1, \dots, k$, $\sum_{i=1}^k p_i = 1$. We can easily define the associated strategy $(\lambda_t)_{t \in [0, T]}$ as

$$\lambda_t(\omega_0, \omega_*) = \sum_{i=1}^k \mathbb{1}_{\{\Lambda = \alpha^i\}}(\omega_0) \alpha_t^i(\omega_*).$$

We explicit the dependence upon the scenario $\omega = (\omega_0, \omega_*)$:

$$\Lambda(\omega_0)_t(\omega_*) = \sum_{i=1}^k \mathbb{1}_{\{\alpha^i\}}(\Lambda(\omega_0)) \alpha_t^i(\omega_*).$$

By the same line of reasoning of Remark 1.2.1, we have that this process is \mathbb{F} -progressively measurable, since the processes α^i , $i = 1, \dots, k$, are \mathbb{F} -progressively measurable and the \mathcal{F}^{0-} -measurable real-valued random variables $\mathbb{1}_{\{\alpha^1\}}(\Lambda(\omega_0))$ can be regarded as defined on the product space $\Omega^0 \times \Omega^*$ and $\mathcal{F}^{0-} \otimes \{\emptyset, \Omega^*\}$ -measurable, therefore \mathbb{F} -progressively measurable. Finally, condition (1.3.8) is satisfied by λ itself.

2. Suppose Λ takes at most countably many values. We can define λ as

$$\lambda_t(\omega_0, \omega_*) = \sum_{i=1}^{\infty} \mathbb{1}_{\{\Lambda=\alpha^i\}}(\omega_0) \alpha_t^i(\omega_*).$$

Set $\lambda_t^n(\omega_0, \omega_*) = \sum_{i=1}^n \mathbb{1}_{\{\Lambda=\alpha^i\}}(\omega_0) \alpha_t^i(\omega_*)$ and observe that, by the same argument of the previous point, λ_t^n is an \mathbb{F} -progressively measurable process for each $n \geq 1$. Furthermore, for each $(t, \omega_0, \omega_*) \in [0, T] \times \Omega^0 \times \Omega^*$, the sequence $\lambda_t^n(\omega_0, \omega_*)$ is eventually constant, being $(\{\Lambda = \alpha^i\})_{i \geq 1}$ a partition of Ω^0 . Therefore, the sequence λ^n converges pointwise to $\lambda = (\lambda_t)_{t \in [0, T]}$. Being λ the pointwise limit of λ^n , we deduce that λ is a progressively measurable process with values in A which satisfies (1.3.8), so that Λ is admissible.

3. Let $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ be a complete probability space, with Ω^0 Polish and \mathcal{F}^{0-} the corresponding Borel σ -algebra, and let $(\lambda_t)_{t \in [0, T]}$ be an A -valued process defined on the $\mathbb{P}^0 \otimes \mathbb{P}^*$ -completion of the product space $(\Omega^0 \times \Omega^*, \mathcal{F}^{0-} \otimes \mathcal{F}, \mathbb{P}^0 \otimes \mathbb{P}^*)$ with values in A . Assume that it is progressively measurable with respect to the $\mathbb{P}^0 \otimes \mathbb{P}^*$ -augmentation of the filtration $\mathbb{F} = (\mathcal{F}^{0-} \otimes \mathcal{F}_t^*)_{t \in [0, T]}$. We can define a function $\Lambda : \Omega^0 \rightarrow \mathbb{A}$ by setting

$$\Lambda(\omega_0) = \begin{cases} (\lambda_t(\omega_0, \cdot))_{t \in [0, T]} : [0, T] \times \Omega^* \rightarrow A & \omega_0 \in \Omega^0 \setminus N, \\ (t, \omega_*) \rightarrow \lambda_t(\omega_0, \omega_*), & \omega_0 \in \Omega^0 \setminus N, \\ a_0 & \omega_0 \in N. \end{cases} \quad (1.3.17)$$

where $N \subset \Omega^0$ is a \mathbb{P}^0 -null set and a_0 is an arbitrary point in A . By Lemma 1.7.1 in Section 1.7.1, the pair $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is an admissible recommendation, with strategy associated to the recommendation Λ given by the process λ itself.

1.4 Approximate N -player coarse correlated equilibria

The next result shows how to construct a sequence of approximate N -player coarse correlated equilibria with approximation error tending to zero as $N \rightarrow \infty$, provided we have a coarse correlated solution to the mean field game.

Theorem 1.4.1. *Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda^*, \mu^*)$ be a coarse correlated solution of the mean field game. For each $N \geq 2$, there exist:*

- (i) *an admissible recommendation to the N players $((\Omega^{0, N}, \mathcal{F}^{0-, N}, \mathbb{P}^{0, N}), \Lambda^N)$;*
- (ii) *a real valued $\varepsilon_N \geq 0$, with $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$,*

so that $((\Omega^{0,N}, \mathcal{F}^{0-,N}, \mathbb{P}^{0,N}), \Lambda^N)$ is an ε_N -coarse correlated equilibrium for the N -player game.

The proof of Theorem 1.4.1 has two main steps: First, starting from a coarse correlated solution to the MFG, we build a probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ large enough to carry any sequence $(\Lambda^i)_{i \geq 1}$ of admissible recommendations such that

1. for every i , Λ^i is supported on the set of open-loop strategies progressively measurable with respect to player i 's private noise;
2. for every i , Λ^i has the same distribution as Λ^* ;
3. for every $N \geq 2$, $(\Lambda^1, \dots, \Lambda^N)$ is exchangeable.

Then, we define the probability space $(\Omega^{0,N}, \mathcal{F}^{0-,N}, \mathbb{P}^{0,N})$ as $(\bar{\Omega}, \mathcal{F}^{0-,N}, \bar{\mathbb{P}})$, where $\mathcal{F}^{0-,N}$ is the $\bar{\mathbb{P}}$ -completion of $\sigma(\Lambda^1, \dots, \Lambda^N)$. The construction of the space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ strictly depends on the coarse correlated solution we need to approximate. This is accomplished in Section 1.4.1. Then, in Section 1.4.2, we prove Theorem 1.4.1. In order to exploit the optimality property (1.3.12) of the coarse correlated solution, great care is taken in comparing the cost functional associated to open-loop strategies in the N -player game with the payoff of the coarse correlated solution. We conclude by a propagation of chaos result, which is specific to our situation but quite standard. For completeness, the statement and the proof of such result is deferred to Section 1.7.2.

1.4.1 Construction of the admissible recommendation profiles to the N -player game

With respect to the probability space $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$ defined in (1.2.3), let us denote by $\mathbb{F}^{(i)}$ the \mathbb{P}^1 -augmentation of the filtration generated by (ξ^i, W^i) . Let us introduce the following set of strategies:

$$\mathbb{A}_{(i)} = \{ \alpha \in \mathbb{A}_N \mid \alpha \text{ is } \mathbb{F}^{(i)} \text{ progressively measurable} \}. \quad (1.4.1)$$

We stress that, by construction, for each $N \geq 2$, open loop strategies for the N -player game are defined on the same probability space $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$ and we have the inclusions $\mathbb{A}_N \subseteq \mathbb{A}_{N+1}$ and $\mathbb{A}_{(i)} \subseteq \mathbb{A}_N$ for every $i \leq N$. Let us denote by $\rho \in \mathcal{P}(\mathcal{C}(\mathcal{P}^2))$ the distribution of μ^* . Let

$$K : \mathcal{F}^{0-} \times \mathcal{C}(\mathcal{P}^2) \rightarrow [0, 1]$$

be the regular conditional probability of \mathbb{P}^0 given μ^* , which exists and is unique since both $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ and $(\mathcal{C}(\mathcal{P}^2), \mathcal{B})$ are Polish spaces. Here and in the following, \mathcal{B} stands for the Borel σ -algebra on $\mathcal{C}(\mathcal{P}^2)$. Let γ denote the joint law of (Λ^*, μ^*) under \mathbb{P}^0 , let κ be a version of the regular conditional probability of γ given μ^* , that is, the stochastic kernel $\kappa : \mathcal{B}_{\mathbb{A}} \times \mathcal{C}(\mathcal{P}^2) \rightarrow [0, 1]$ so that it holds

$$\mathbb{P}^0((\Lambda^*, \mu^*) \in C \times B) = \int_B \kappa(C, m) \rho(dm) \quad \forall C \in \mathcal{B}_{\mathbb{A}}, \forall B \in \mathcal{B}. \quad (1.4.2)$$

Define the probability space $(\bar{\Omega}, \bar{\mathcal{F}})$ in the following way:

$$\bar{\Omega} = \left(\prod_1^\infty \Omega^0 \right) \times \mathcal{C}(\mathcal{P}^2), \quad \bar{\mathcal{F}} = \left(\bigotimes_1^\infty \mathcal{F}^{0-} \right) \otimes \mathcal{B}, \quad (1.4.3)$$

and define $\bar{\mathbb{P}}$ so that, for every cylinder R with basis $A_1 \times \cdots \times A_N \times B$, with $A_i \in \mathcal{F}^{0-}$ for every $i = 1, \dots, N$, $N \geq 2$, $B \in \mathcal{B}$, it holds

$$\bar{\mathbb{P}}(R) = \int_B \prod_{i=1}^N K(A_i, m) \rho(dm). \quad (1.4.4)$$

We complete the space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ with the $\bar{\mathbb{P}}$ -null sets. Let $\bar{\omega} = ((\omega_0^i)_{i \geq 1}, m)$ denote a scenario in $\bar{\Omega}$. Let $\mu : (\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}}) \rightarrow (\mathcal{C}(\mathcal{P}^2), \mathcal{B})$ be the projection on $\mathcal{C}(\mathcal{P}^2)$, that is

$$\mu(\bar{\omega}) = m. \quad (1.4.5)$$

Lemma 1.4.2. *There exists a sequence of recommendations $(\Lambda^i)_{i \geq 1}$ from $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ to $\times_{N=1}^{\infty} \mathbb{A}_N$ so that, for each $i \geq 1$, the following holds:*

- (a) Λ^i is an admissible recommendation, and it takes values in $\mathbb{A}_{(i)}$.
- (b) The joint law of $(\Lambda^1, \dots, \Lambda^N)$ under $\bar{\mathbb{P}}$ is supported on $\times_{i=1}^N \mathbb{A}_{(i)} \subseteq \mathbb{A}_N^N$ and it is given by

$$\gamma_N(d\alpha^1, \dots, d\alpha^N) = \int_{\mathcal{C}(\mathcal{P}^2)} \bigotimes_{i=1}^N \kappa(d\alpha^i, m) \rho(dm). \quad (1.4.6)$$

As a consequence, for every $i \geq 1$, (Λ^i, μ) has the same distribution as (Λ^*, μ^*) and $(\Lambda^i)_{i \geq 1}$ are conditionally independent given μ .

Proof. Recall from (1.2.3) and (1.3.1) the definitions of the spaces $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$ and $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$. Observe that, up to completion, it holds

$$(\Omega^1, \mathcal{F}^1, \mathbb{P}^1) = \bigotimes_1^{\infty} (\Omega^*, \mathcal{F}^*, \mathbb{P}^*) \quad (1.4.7)$$

so that a scenario $\omega_1 \in \Omega^1$ can be written as $\omega_1 = (\omega_*^j)_{j \geq 1}$. Moreover, by definition of $(\xi^j, W^j)_{j \geq 1}$ in (1.2.2), for every $i \geq 1$ it holds

$$(\xi^i, W^i)(\omega_1) = (x^i, w^i) = (\xi^*, W^*)(\omega_*^i),$$

so that $(\xi^i, W^i)_{i \geq 1}$ can be seen as a sequence of independent copies of (ξ^*, W^*) . Define the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ as in (1.2.7). Let $\lambda^i = (\lambda_t^i)_{t \in [0, T]}$, $i \geq 1$, be independent copies of $\lambda^* = (\lambda_t^*)_{t \in [0, T]}$, the strategy associated to the admissible recommendation Λ^* according to (1.3.7), so that

$$\lambda_t^i(\bar{\omega}, \omega_1) = \lambda_t^i((\omega_0^j)_{j \geq 1}, m, (\omega_*^j)_{j \geq 1}) = \lambda_t^*(\omega_0^i, \omega_*^i).$$

For every i , λ^i is \mathbb{F} -progressively measurable: indeed, since by definition the measures \mathbb{P} and $\mathbb{P}^0 \otimes \mathbb{P}^*$ coincide on the cylinders A_i of the form

$$A_i = \{(\bar{\omega}, \omega_1) = ((\omega_0^j)_{j \geq 1}, m, (\omega_*^j)_{j \geq 1}) \in \bar{\Omega} \times \Omega^1 \mid (\omega_0^i, \omega_*^i) \in G\}, \quad (1.4.8)$$

for any $G \in \mathcal{F}^{0-} \otimes \mathcal{F}^1$, every $\mathbb{P}^0 \otimes \mathbb{P}^*$ -null set N can be identified with a \mathbb{P} -null cylinder A_i of the form (1.4.8) with basis N . Therefore, for every $t \in [0, T]$, the \mathbb{P} -augmentation of the filtration $\mathcal{F}^{0-} \otimes \mathcal{F}_t^i$ contains all the cylinders with basis

$$A_i = \{(\bar{\omega}, \omega_1) = ((\omega_0^j)_{j \geq 1}, m, (\omega_*^j)_{j \geq 1}) \in \bar{\Omega} \times \Omega^1 \mid (\omega_0^i, \omega_*^i) \in G\},$$

for any G in the $\mathbb{P}^0 \otimes \mathbb{P}^*$ -augmentation of $\mathcal{F}^{0-} \otimes \mathcal{F}_t^*$. This is enough to conclude that λ^i is progressively measurable with respect to the \mathbb{P} -augmentation of $\mathcal{F}^{0-} \otimes \mathcal{F}_t^i$, and so with respect to the filtration \mathbb{F} as well. We define Λ^i as in (1.3.17), that is

$$\Lambda^i(\bar{\omega}) = \begin{cases} (\lambda_t^i(\bar{\omega}, \cdot))_{t \in [0, T]} : [0, T] \times \Omega^1 \rightarrow A & \bar{\omega} \in \bar{\Omega} \setminus \mathcal{N}, \\ (t, \omega_1) \rightarrow \lambda_t^i(\bar{\omega}, \omega_1), & \\ a_0 & \bar{\omega} \in \mathcal{N}, \end{cases} \quad (1.4.9)$$

where $\mathcal{N} \subseteq \bar{\Omega}$ is a $\bar{\mathbb{P}}$ -null set and a_0 is an arbitrary point in A . By Lemma 1.7.1, Λ^i is an admissible recommendation from $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ to $(\mathbb{A}_N, \mathcal{B}_{\mathbb{A}_N})$, for every $N \geq i$. Since the associated strategies coincide pointwise, it holds $\Lambda^i(\bar{\omega}) = \Lambda^*(\omega_0^i)$ $\bar{\mathbb{P}}$ -a.s., as ensured by Proposition 1.7.2. In particular, this implies that Λ^i only takes values in $\mathbb{A}_{(i)}$, since for every fixed $\bar{\omega}$ the control process $(\lambda_t^i(\bar{\omega}, \cdot))_{t \in [0, T]}$ is $\mathbb{F}^{(i)}$ -progressively measurable. This proves point (a).

As for point (b), for every $N \geq 2$, $(\Lambda^1, \dots, \Lambda^N)$ takes values in $\times_{j=1}^N \mathbb{A}_{(j)}$ by construction. Hence, we may restrict the attention to Borel sets $C_j \subseteq \mathbb{A}_{(j)}$, for every $j = 1, \dots, N$. Let $B \in \mathcal{B}$. Since $\Lambda^j(\bar{\omega}) = \Lambda^*(\omega_0^j)$ for every $j = 1, \dots, N$ $\bar{\mathbb{P}}$ -a.s., by definition of $\bar{\mathbb{P}}$, we have

$$\begin{aligned} \bar{\mathbb{P}}(\Lambda^1 \in C_1, \dots, \Lambda^N \in C_N, \mu \in B) &= \bar{\mathbb{P}}\left(\times_{j=1}^N \{\omega_0^j : \Lambda^*(\omega_0^j) \in C_j\} \times \times_{j=N+1}^{\infty} \Omega^0 \times \mathcal{C}(\mathcal{P}^2)\right) \\ &= \int_B \prod_{j=1}^N K(\{\omega_0^j : \Lambda^*(\omega_0^j) \in C_j\}, m) \rho(dm) \\ &= \int_B \prod_{j=1}^N K(\{\omega_0 : \Lambda^*(\omega_0) \in C_j\}, m) \rho(dm) = \int_B \prod_{j=1}^N \kappa(C_j, m) \rho(dm). \end{aligned}$$

This shows also that (Λ^i, μ) are identically distributed as (Λ^*, μ^*) and that $(\Lambda^i)_{i \geq 1}$ are conditionally i.i.d. given μ . \square

Set $\mathcal{F}^{0-, N} = \sigma(\Lambda^1, \dots, \Lambda^N)$ and $(\Omega^{0, N}, \mathcal{F}^{0-, N}, \mathbb{P}^{0, N}) = (\bar{\Omega}, \mathcal{F}^{0-, N}, \bar{\mathbb{P}})$, for $N \geq 2$. Then,

$$(\Lambda^1, \dots, \Lambda^N) : (\bar{\Omega}, \mathcal{F}^{0-, N}, \bar{\mathbb{P}}) \rightarrow (\mathbb{A}_N^N, \mathcal{B}_{\mathbb{A}_N^N})$$

is the candidate ε_N -coarse correlated equilibrium to the N -player game, with ε_N to be determined.

Remark 1.4.1. The construction of the probability spaces $(\Omega^{0, N}, \mathcal{F}^{0-, N}, \mathbb{P}^{0, N})$ is rather involved, but has the advantage of making the admissible recommendation to the N -players $(\Lambda^1, \dots, \Lambda^N)$ easy to define. Besides this technical reason, we notice that, both in the N -player game and in the mean field game, the mediator may choose the space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ he uses to randomize players' strategies, as already pointed out in Sections 1.2 and 1.3. Then, it is natural to use the same space $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ on which the coarse correlated solution to the mean field game is defined to randomize players' strategies in the N -player game as well.

1.4.2 Proof of Theorem 1.4.1

By symmetry, let us consider only possible deviations of player $i = 1$. For every $N \geq 2$, let ε_N be given by

$$\varepsilon_N := \sup_{\beta \in \mathbb{A}_N} (\mathfrak{J}_1^N(\Lambda^N) - \mathfrak{J}_1^N(\Lambda^{N,-1}, \beta)) = \mathfrak{J}_1^N(\Lambda^N) - \inf_{\beta \in \mathbb{A}_N} \mathfrak{J}_N^1(\Lambda^{N,-1}, \beta). \quad (1.4.10)$$

By definition of ε_N , Λ^N is an ε_N -coarse correlated equilibrium for every $N \geq 2$. In order to conclude the proof, we only need to prove that $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$. For each $N \geq 2$, choose $\beta^N \in \mathbb{A}_N$ so that

$$\mathfrak{J}_1^N(\Lambda^{N,-1}, \beta^N) \leq \inf_{\beta \in \mathbb{A}_N} \mathfrak{J}_1^N(\Lambda^{N,-1}, \beta) - \frac{1}{N}.$$

Let $Z = (Z_t)_{t \in [0, T]}$ be the solution of

$$dZ_t = b(t, Z_t, \mu_t, \beta_t^N) dt + dW_t^1, \quad Z_0 = \xi^1, \quad (1.4.11)$$

and define the corresponding cost as

$$\mathfrak{J}(\beta^N, \mu) = \mathbb{E} \left[\int_0^T f(s, Z_s, \mu_s, \beta_s^N) ds + g(Z_T, \mu_T) \right].$$

Let X be the solution of

$$dX_t = b(t, X_t, \mu_t, \lambda_t^1) dt + dW_t^1, \quad X_0 = \xi^1, \quad (1.4.12)$$

with associated cost

$$\mathfrak{J}(\Lambda^1, \mu) = \mathbb{E} \left[\int_0^T f(s, X_s, \mu_s, \lambda_s^1) ds + g(X_T, \mu_T) \right].$$

Observe that, by construction, $(\xi^1, W^1, \mu, \lambda^1)$ under \mathbb{P} is distributed as $(\xi, W, \mu^*, \lambda^*)$ under \mathbb{P}^* . Therefore, by Theorem B.1 the joint distribution of (X, λ^1, μ) under \mathbb{P} is the same as (X^*, λ^*, μ^*) under \mathbb{P}^* , where X^* denotes the state process resulting from the mean field CCE (Λ^*, μ^*) . Moreover, note that by construction λ^1 is $\mathbb{F}^{\mu, \xi^1, W^1}$ -progressively measurable, where

$$\mathcal{F}_t^{\mu, \xi^1, W^1} = \sigma(\mu) \vee \sigma(\xi^1) \vee \sigma(W_s^1 : s \leq t), \quad t \in [0, T].$$

By the Lipschitz continuity of b , X may be taken to be $\mathbb{F}^{\mu, \xi^1, W^1}$ -adapted as well.

To prove the theorem, it is enough to show the following:

$$\mathfrak{J}(\Lambda^1, \mu) \leq \mathfrak{J}(\beta^N, \mu), \quad (1.4.13a)$$

$$\lim_{N \rightarrow \infty} \mathfrak{J}_N^1(\Lambda^N) = \mathfrak{J}(\Lambda^1, \mu), \quad (1.4.13b)$$

$$\lim_{N \rightarrow \infty} |\mathfrak{J}_N^1(\Lambda^{N,-1}, \beta^N) - \mathfrak{J}(\beta^N, \mu)| = 0. \quad (1.4.13c)$$

These imply the conclusion, by noticing that

$$\varepsilon_N \leq \mathfrak{J}_N^1(\Lambda^N) - \mathfrak{J}(\Lambda^1, \mu) + \mathfrak{J}(\Lambda^1, \mu) - \mathfrak{J}(\beta^N, \mu) + \mathfrak{J}(\beta^N, \mu) - \mathfrak{J}_N^1(\Lambda^{N,-1}, \beta^N) - \frac{1}{N}.$$

We start by proving (1.4.13a). We observe that we cannot just deduce it from the optimality property (1.3.12) of (Λ^*, μ^*) : since β^N may belong to $\mathbb{A}_N \setminus \mathbb{A}_{(1)}$, it may not be identifiable with an open loop strategy for the MFG, for which inequality (1.4.13a) would hold. Instead, we prove it by using the regular conditional probability of \mathbb{P} given $(\xi^i, W^i)_{i=2}^N$. Denote by (\mathbf{x}, \mathbf{w}) a point $(x^i, w^i)_{i=2}^N \in (\mathbb{R}^d \times \mathcal{C}^d)^{N-1}$, and let $P_\eta = \bigotimes_1^{N-1} (\eta \otimes \mathbb{W}^d)$ denote the joint law of $(\xi^i, W^i)_{i=2}^N$ under \mathbb{P} . Let $\mathbb{P}^{\mathbf{x}, \mathbf{w}}$ be a version of the regular conditional probability of \mathbb{P} given $(\xi^i, W^i)_{i=2}^N = (x^i, w^i)_{i=2}^N$. We rewrite (1.4.13a) as

$$\begin{aligned}
& \mathfrak{J}(\Lambda^1, \mu) - \mathfrak{J}(\beta^N, \mu) = \\
& = \mathbb{E} \left[\left(\int_0^T f(s, X_s, \mu_s, \lambda_s^1) ds + g(X_T, \mu_T) \right) - \left(\int_0^T f(s, Z_s, \mu_s, \beta_s^N) ds + g(Z_T, \mu_T) \right) \right] \\
& = \mathbb{E} \left[\mathbb{E} \left[\left(\int_0^T f(s, X_s, \mu_s, \lambda_s^1) ds + g(X_T, \mu_T) \right) \right. \right. \\
& \quad \left. \left. - \left(\int_0^T f(s, Z_s, \mu_s, \beta_s^N) ds + g(Z_T, \mu_T) \right) \middle| (\xi^i, W^i)_{i=2}^N \right] \right] \\
& = \int_{(\mathbb{R}^d \times \mathcal{C}^d)^{N-1}} \left(\mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, X_s, \mu_s, \lambda_s^1) ds + g(X_T, \mu_T) \right] \right. \\
& \quad \left. - \mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, Z_s, \mu_s, \beta_s^N) ds + g(Z_T, \mu_T) \right] \right) P_\eta(d\mathbf{x}, d\mathbf{w}).
\end{aligned} \tag{1.4.14}$$

We analyse separately the two terms in the last equality. Let us start with the term depending upon λ^1 . Since μ , λ^1 , W^1 and ξ^1 are independent of $(\xi^i, W^i)_{i=2}^N$ under \mathbb{P} and X is $\mathbb{F}^{\mu, \xi^1, W^1}$ -adapted, X is independent of $(\xi^i, W^i)_{i=2}^N$ as well. We deduce that, under $\mathbb{P}^{\mathbf{x}, \mathbf{w}}$, W^1 is an \mathbb{F} -Brownian motion, X solves equation (1.4.12) and $\mathbb{P}^{\mathbf{x}, \mathbf{w}} \circ (X, \lambda^1, \mu)^{-1} = \mathbb{P} \circ (X, \lambda^1, \mu)^{-1} = \mathbb{P}^* \circ (X^*, \lambda^*, \mu^*)^{-1}$, for P_η -a.e. $(\mathbf{x}, \mathbf{w}) \in (\mathbb{R}^d \times \mathcal{C}^d)^{N-1}$. In particular, this implies that

$$\begin{aligned}
& \mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, X_s, \mu_s, \lambda_s^1) ds + g(X_T, \mu_T) \right] \\
& = \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T f(s, X_s^*, \mu_s^*, \lambda_s^*) ds + g(X_T^*, \mu_T^*) \right] = \mathfrak{J}(\Lambda^*, \mu^*)
\end{aligned} \tag{1.4.15}$$

for P_η -a.e. $(\mathbf{x}, \mathbf{w}) \in (\mathbb{R}^d \times \mathcal{C}^d)^{N-1}$.

As for the term depending upon β^N , we note that, since $\beta^N \in \mathbb{A}_N$, there exists a progressively measurable functional $\hat{\beta} : [0, T] \times (\mathbb{R}^d \times \mathcal{C}^d)^N \rightarrow A$ so that

$$\beta_t^N = \hat{\beta}(t, \xi^1, W^1, \dots, \xi^N, W^N).$$

Under $\mathbb{P}^{\mathbf{x}, \mathbf{w}}$, it holds

$$\beta_t^N = \hat{\beta}(t, \xi^1, W^1, x^2, w^2, \dots, x^N, w^N) \quad \forall t \in [0, T] \quad \mathbb{P}^{\mathbf{x}, \mathbf{w}}\text{-a.s.}, \tag{1.4.16}$$

since $(\xi^i, W^i)_{i=1}^2$ are almost surely constant under $\mathbb{P}^{\mathbf{x}, \mathbf{w}}$. Since the joint law of μ , W^1 and ξ^1 is the same under both \mathbb{P} and $\mathbb{P}^{\mathbf{x}, \mathbf{w}}$, (1.4.16) implies that Z satisfies

$$dZ_t = b(t, Z_t, \mu_t, \hat{\beta}_t^N(\mathbf{x}, \mathbf{w})) dt + dW_t^1, \quad Z_0 = \xi^1, \tag{1.4.17}$$

under $\mathbb{P}^{\mathbf{x}, \mathbf{w}}$, with $\hat{\beta}^N(\mathbf{x}, \mathbf{w})$ $\mathbb{F}^{(1)}$ -progressively measurable. For every $(\mathbf{x}, \mathbf{w}) \in (\mathbb{R}^d \times \mathcal{C}^d)^{N-1}$, define the strategy

$$\tilde{\beta}(\mathbf{x}, \mathbf{w}) = \left(\hat{\beta}(t, \xi, W, x^2, w^2, \dots, x^N, w^N) \right)_{t \in [0, T]}. \quad (1.4.18)$$

Then $\tilde{\beta}(\mathbf{x}, \mathbf{w})$ belongs to \mathbb{A} for every $(\mathbf{x}, \mathbf{w}) \in (\mathbb{R}^d \times \mathcal{C}^d)^{N-1}$, and it depends measurably upon (\mathbf{x}, \mathbf{w}) . For every (\mathbf{x}, \mathbf{w}) , let \tilde{Z} be the solution of

$$d\tilde{Z}_t = b(t, \tilde{Z}_t, \mu_t^*, \tilde{\beta}_t(\mathbf{x}, \mathbf{w}))dt + dW_t, \quad \tilde{Z}_0 = \xi.$$

Since $\mathbb{P}^{\mathbf{x}, \mathbf{w}} \circ (Z, \beta, \mu)^{-1} = \mathbb{P}^{\mathbf{x}, \mathbf{w}} \circ (Z, \hat{\beta}^N(\mathbf{x}, \mathbf{w}), \mu)^{-1} = \mathbb{P}^* \circ (\tilde{Z}, \tilde{\beta}(\mathbf{x}, \mathbf{w}), \mu^*)^{-1}$, it follows that

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, Z_s, \mu_s, \beta_s^N) ds + g(Z_T, \mu_T) \right] \\ &= \mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, Z_s, \mu_s, \hat{\beta}_s^N(\mathbf{x}, \mathbf{w})) ds + g(Z_T, \mu_T) \right] \\ &= \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T f(s, \tilde{Z}_s, \mu_s^*, \tilde{\beta}_s(\mathbf{x}, \mathbf{w})) ds + g(\tilde{Z}_T, \mu_T^*) \right] = \mathfrak{J}(\tilde{\beta}(\mathbf{x}, \mathbf{w}), \mu^*). \end{aligned} \quad (1.4.19)$$

We note that the left-hand side of (1.4.15) depends measurably upon (\mathbf{x}, \mathbf{w}) due to a monotone class argument. Being (Λ^*, μ^*) a mean field CCE by assumption, (1.4.15) and (1.4.19) imply that

$$\begin{aligned} \mathfrak{J}(\Lambda^1, \mu) - \mathfrak{J}(\beta^N, \mu) &= \int \left(\mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, X_s, \mu_s, \lambda_s^1) ds + g(X_T, \mu_T) \right] \right. \\ &\quad \left. - \mathbb{E}^{\mathbb{P}^{\mathbf{x}, \mathbf{w}}} \left[\int_0^T f(s, Z_s, \mu_s, \beta_s^N) ds + g(Z_T, \mu_T) \right] \right) P_\eta(d\mathbf{x}, d\mathbf{w}) \\ &= \int \left(\mathfrak{J}(\Lambda^*, \mu^*) - \mathfrak{J}(\tilde{\beta}(\mathbf{x}, \mathbf{w}), \mu^*) \right) P_\eta(d\mathbf{x}, d\mathbf{w}) \leq 0, \end{aligned}$$

which yields (1.4.13a).

As for (1.4.13b) and (1.4.13c), they must be handled by continuity arguments on the cost functions and propagation of chaos as stated in Lemma 1.7.3. We give the details only for (1.4.13b), since (1.4.13c) is analogous. We have:

$$\begin{aligned} |\mathfrak{J}_N^1(\Lambda^N) - \mathfrak{J}(\Lambda^1, \mu)| &\leq \mathbb{E} \left[\int_0^T \left| f(t, X_t^{1, N}, \mu_t^N, \lambda_t^1) - f(t, X_t, \mu_t, \lambda_t^1) \right| dt \right. \\ &\quad \left. + \left| g(X_T^N, \mu_T^N) - g(X_T, \mu_T) \right| \right] = \mathbb{E} [\Delta f + \Delta g]. \end{aligned}$$

For Δf , Assumptions **A** ensure that f is locally Lipschitz with at most quadratic

growth. Therefore, by straightforward estimates, we have:

$$\begin{aligned}
\mathbb{E}[\Delta f] &\leq C \mathbb{E} \left[1 + \|X^{1,N}\|_{\mathcal{C}^d}^2 + \|X\|_{\mathcal{C}^d}^2 + \int_0^T \int_{\mathbb{R}^d} |y|^2 \mu_t(dy) dt + \int_0^T \int_{\mathbb{R}^d} |y|^2 \mu_t^N(dy) dt \right. \\
&\quad \left. + 2 \int_0^T |\lambda_t^1| dt \right]^{\frac{1}{2}} \cdot \mathbb{E} \left[\left(\|X^{1,N} - X\|_{\mathcal{C}^d}^2 + \int_0^T \mathcal{W}_{2,\mathbb{R}^d}^2(\mu_t^N, \mu_t) dt \right) \right]^{\frac{1}{2}} \\
&\leq C \left(1 + \max_{k=1,\dots,N} \mathbb{E} \left[\|X^{k,N}\|_{\mathcal{C}^d}^2 \right]^{\frac{1}{2}} + \mathbb{E} \left[\|X\|_{\mathcal{C}^d}^2 \right]^{\frac{1}{2}} \right) \\
&\quad \cdot \left(\mathbb{E} \left[\|X^{1,N} - X\|_{\mathcal{C}^d}^2 \right]^{\frac{1}{2}} + \sup_{t \in [0,T]} \mathbb{E} \left[\mathcal{W}_{2,\mathbb{R}^d}^2(\mu_t^N, \mu_t) \right]^{\frac{1}{2}} \right).
\end{aligned}$$

By Lemma 1.7.3, the right-hand side tends to 0 as N goes to infinity. The convergence of $\mathbb{E}[\Delta g]$ is shown analogously.

1.5 Existence of a coarse correlated solution of the mean field game

The main result of this section regards the existence of a coarse correlated solution of the MFG, which requires the following additional assumption:

Assumption B. For every $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d)$, the set

$$K(t, x, m) = \{(b(t, x, m, a), z) : a \in A, f(t, x, m, a) \leq z\} \subseteq \mathbb{R}^d \times \mathbb{R} \quad (1.5.1)$$

is closed and convex.

Assumption **B** is standard when dealing with relaxed controls (see, e.g., [60] and [75] for relaxed controls in control theory, and the series of works [41, 84, 85, 87], or [28], in mean field games literature). Similarly to these works, we will use relaxed controls to apply compactness arguments, and we will use Assumption **B** to come back to strict controls. We refer to Appendix A for a brief recap on relaxed controls.

The main result of this section is the following:

Theorem 1.5.1 (Existence of a coarse correlated solution of the MFG). *In addition to Assumptions **A**, suppose that Assumption **B** holds. Then there exists a coarse correlated solution of the mean field game.*

The road map to prove Theorem 1.5.1 is as follows: Taking inspiration from [74, 97, 98] and [19, Appendix 1.B], we associate a zero-sum game to the search of a mean field CCE. Loosely speaking, the game should be of the following type: player A, the maximizer, chooses a correlated measure flow (Λ, μ) , while player B chooses a deviating strategy $\beta \in \mathbb{A}$. The payoff functional is the following:

$$F[(\Lambda, \mu), \beta] = \mathfrak{J}(\beta, \mu) - \mathfrak{J}(\Lambda, \mu). \quad (1.5.2)$$

Player A aims at maximizing F , while player B chooses her strategy in order to minimize F . In order to get an equilibrium, one should restrict to correlated measure

flows (Λ, μ) so that the consistency condition (1.3.13) is satisfied. If we could show that the game has a positive value and player A has an optimal strategy (Λ^*, μ^*) , then we would have established that such a strategy would satisfy the optimality property (1.3.12) as well, and therefore (Λ^*, μ^*) would be a mean field CCE.

In order to get a convenient structure for the sets of strategies and good continuity and convexity properties of the payoff functionals, in Section 1.5.1 we define a more general zero-sum game, in which we embed our auxiliary problem. As shown in Proposition 1.5.3, the auxiliary zero-sum game extends the payoff functional in equation (1.5.2). Particular care is needed in dealing with the term depending both on β and μ , since it must reflect independent strategy choices of the opponents. Using Fan's minimax theorem, we will show that the auxiliary game has positive value and admits an optimal strategy for the maximizing player. This is the content of Theorem 1.5.2, whose proof is deferred to Section 1.5.3. In Section 1.5.2 we prove Theorem 1.5.1, by using such an optimal strategy to induce a coarse correlated solution of the mean field game. Many technical lemmata are needed to prove this existence result. For reader's convenience, some statements and proofs are deferred to Section 1.7.3.

Remark 1.5.1. Under similar assumptions to Assumptions **A**, the existence of weak or also strong MFG solutions has been established in the literature; see, for instance, [84, 85, 88]. As already noticed, strong MFG solutions are coarse correlated solutions to the MFG as well, and, at least in some cases, this is also true for weak MFG solutions (see upcoming Section 1.6.2). Nevertheless, we think that the existence result of Theorem 1.5.1 is of independent interest, for two main reasons. Firstly, the proof shows directly the existence of a coarse correlated solutions without relying on existence results for stronger notions of equilibria, which are usually based on fixed point theorems. Such a direct proof is instead based on a minimax theorem, which, to the best of our knowledge, is used for the first time in continuous time MFG literature. Secondly, since we reduce the search of a coarse correlated solution to the MFG to finding an equilibrium for an auxiliary linear zero-sum game on spaces of measures, our formulation paves the way to the use of linear programming methods for computing mean field CCEs. We further discuss this point in the following Remark 1.5.2.

1.5.1 The auxiliary zero-sum game

We now formally define the auxiliary zero-sum game.

Definition 1.5.1 (Strategies for player A). A *strategy for player A* is a probability measure $\Gamma \in \mathcal{P}(\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2))$ so that there exists a tuple $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ with the following properties:

- (i) $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a filtered probability space satisfying the usual assumptions; Ω is Polish and \mathcal{F} is its corresponding Borel σ -algebra.
- (ii) W is an \mathbb{F} -Brownian motion and ξ is an \mathcal{F}_0 -measurable independent \mathbb{R}^d -valued random variable with law η .
- (iii) μ is an \mathcal{F}_0 -measurable random variable with values in $\mathcal{C}(\mathcal{P}^2)$; it is independent of both ξ and W .

(iv) \mathbf{r} is an \mathbb{F} -progressively measurable relaxed control $\mathbf{r} = (\mathbf{r}_t)_{t \in [0, T]}$ with values in A .

(v) Let X be the solution of

$$dX_t = \int_A b(t, X_t, \mu_t, a) \mathbf{r}_t(da) dt + dW_t, \quad t \in [0, T], \quad X_0 = \xi. \quad (1.5.3)$$

Then $\mu_t(\cdot) = \mathbb{P}(X_t \in \cdot \mid \mu)$ \mathbb{P} -a.s for every $t \in [0, T]$.

(vi) Γ is the joint law under \mathbb{P} of X , μ and \mathbf{r} : $\Gamma = \mathbb{P} \circ (X, \mathbf{r}, \mu)^{-1}$.

We denote by \mathcal{K} the set of strategies for player A.

We observe that, by Assumptions **A**, there exists a unique solution to equation (1.5.3) for every tuple \mathfrak{U} satisfying properties (i-iv).

Definition 1.5.2 (Strategies for player B). A stochastic kernel Σ from $\mathcal{C}(\mathcal{P}^2)$ to $\mathcal{C}^d \times \mathcal{V}$ is a *strategy for player B* if there exists a tuple $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mathbf{r})$ so that

- (i) $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a filtered probability space satisfying the usual assumptions; Ω is Polish and \mathcal{F} is its corresponding Borel σ -algebra.
- (ii) W is an \mathbb{F} -Brownian motion and ξ is an \mathcal{F}_0 -measurable independent \mathbb{R}^d -valued random variable with law η .
- (iii) \mathbf{r} is an \mathbb{F} -progressively measurable relaxed control $\mathbf{r} = (\mathbf{r}_t)_{t \in [0, T]}$ with values in A .
- (iv) For every $m \in \mathcal{C}(\mathcal{P}^2)$, $\Sigma(\cdot, m) \in \mathcal{P}(\mathcal{C}^d \times \mathcal{V})$ is the joint law under \mathbb{P} of (X^m, \mathbf{r}) , where X^m is the solution to

$$dX_t^m = \int_A b(t, X_t^m, m_t, a) \mathbf{r}_t(da) dt + dW_t, \quad X_0 = \xi, \quad (1.5.4)$$

that is:

$$\Sigma(B, m) = \mathbb{P}((X^m, \mathbf{r}) \in B) \quad \forall m \in \mathcal{C}(\mathcal{P}^2), B \in \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{V}}. \quad (1.5.5)$$

We denote by \mathcal{Q} the set of strategies for player B.

By Lemma 1.7.4, the set of strategies \mathcal{Q} for player B is well defined in the sense that the map Σ is truly a stochastic kernel.

We now define the payoff functional \mathfrak{p} for the zero-sum game. Let us introduce the function $\mathfrak{F} : \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2) \rightarrow \mathbb{R}$ defined by

$$\mathfrak{F}(y, q, m) = \int_0^T \int_A f(t, y_t, m_t, a) q_t(da) dt + g(y_T, m_T). \quad (1.5.6)$$

Definition 1.5.3 (Auxiliary zero-sum game). The *auxiliary zero-sum game* is a zero-sum game where:

- The set of strategies for player A, the maximizer, is the set \mathcal{K} introduced in Definition 1.5.1.
- The set of strategies for player B, the minimizer, is the set \mathcal{Q} introduced in Definition 1.5.2.
- The payoff functional is the function $\mathbf{p} : \mathcal{K} \times \mathcal{Q} \rightarrow \mathbb{R}$ defined as

$$\begin{aligned} \mathbf{p}(\Gamma, \Sigma) = & \int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \rho(dm) \\ & - \int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \mathfrak{F}(y, q, m) \Gamma(dy, dq, dm), \end{aligned} \quad (1.5.7)$$

where ρ denotes the marginal of Γ on $\mathcal{C}(\mathcal{P}^2)$.

We denote the lower and upper values of the game as, respectively, v^A and v^B :

$$v^A = \sup_{\Gamma \in \mathcal{K}} \inf_{\Sigma \in \mathcal{Q}} \mathbf{p}(\Gamma, \Sigma), \quad v^B = \inf_{\Sigma \in \mathcal{Q}} \sup_{\Gamma \in \mathcal{K}} \mathbf{p}(\Gamma, \Sigma).$$

If the lower and upper values of the game are equal, we set $v = v^A = v^B$ and call v the value of the game. We say that a strategy $\Gamma^* \in \mathcal{K}$ is optimal for player A if

$$\inf_{\Sigma \in \mathcal{Q}} \mathbf{p}(\Gamma^*, \Sigma) = \max_{\Gamma \in \mathcal{K}} \inf_{\Sigma \in \mathcal{Q}} \mathbf{p}(\Gamma, \Sigma).$$

The next result ensures existence of an optimal strategy for the maximizing player:

Theorem 1.5.2 (Existence of the value of the game and of an optimal strategy for the maximizing player). *Consider the game described in Definition 1.5.3. The following holds:*

- (i) *The game has a value, i.e. $v^A = v^B$.*
- (ii) *There exists a strategy $\Gamma^* \in \mathcal{K}$ which is optimal for player A.*
- (iii) *The value v of the game is non negative: $v \geq 0$.*

The proof of this theorem is deferred to Section 1.5.3. The last result of this section shows that, for every correlated measure flow (Λ, μ) so that consistency condition (1.3.13) is satisfied and every deviation $\beta \in \mathbb{A}$, there exists a pair of strategies $(\Gamma, \Sigma) \in \mathcal{K} \times \mathcal{Q}$ so that the following equality holds:

$$\mathbf{p}(\Gamma, \Sigma) = \mathfrak{J}(\Lambda, \mu) - \mathfrak{J}(\beta, \mu) = F[(\Lambda, \mu), \beta]. \quad (1.5.8)$$

Proposition 1.5.3. *Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ be a correlated measure flow. Denote by ρ the law of μ . Let $\lambda = (\lambda_t)_{t \in [0, T]}$ be the strategy associated to the admissible recommendation Λ and let $\beta \in \mathbb{A}$.*

- (i) *Let X be the solution to (1.3.9). Suppose that consistency condition (1.3.13) is satisfied. For every $t \in [0, T]$, set $\mathbf{r}_t(da)dt = \delta_{\lambda_t}(da)dt$. Then, the probability measure $\Gamma = \mathbb{P} \circ (X, \mathbf{r}, \mu)^{-1}$ belongs to the set \mathcal{K} .*

(ii) For every $t \in [0, T]$, set $\mathbf{b}_t(da)dt = \delta_{\beta_t}(da)dt$. Denote by Y the solution to (1.3.10). Then, there exists $\Sigma \in \mathcal{Q}$ so that

$$\mathbb{P}((Y, \mathbf{b}, \mu) \in B \times S) = \int_S \Sigma(B, m) \rho(dm), \quad \forall B \in \mathcal{B}_{\mathcal{C}^d \times \mathcal{V}}, S \in \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}. \quad (1.5.9)$$

(iii) The pair of strategies (Γ, Σ) satisfies equation (1.5.8).

Proof. In the following, we work on the probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ defined in (1.3.6). Recall that, as pointed out in Remark 1.2.1, we can think of W , ξ and μ as independent random variables, each of them defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$. Observe that the $\mathcal{P}(A)$ -valued process $\mathbf{r} = (\delta_{\lambda_t})_{t \in [0, T]}$ is \mathbb{F} -progressively measurable since Λ is admissible by assumption. Let X be the solution to equation (1.3.9). Since X obviously satisfies (1.5.3) for such a process \mathbf{r} and the condition $\mu_t(\cdot) = \mathbb{P}(X_t \in \cdot \mid \mu)$ holds by assumption, $\Gamma = \mathbb{P} \circ (X, \mu, \mathbf{r})^{-1}$ belongs to \mathcal{K} .

As for point (ii), recall from Remark 1.2.1 that we can regard β as defined on the product space $(\Omega, \mathcal{F}, \mathbb{P})$, and that β and μ are mutually independent by construction. Therefore, the $\mathcal{P}(A)$ -valued process $\mathbf{b} = (\delta_{\beta_t}(da))_{t \in [0, T]}$ is independent of μ . Let Y be the solution of equation (1.3.10). By Lemma 1.7.5 in Section 1.7.3, equation (1.5.9) holds.

Finally, since X and Y are defined on the same filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, we can write the integrals in \mathbf{p} as expectations:

$$\begin{aligned} \int \mathfrak{F}(y, q, m) \Gamma(dy, dq, dm) &= \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \lambda_t) dt + g(X_T, \mu_T) \right] = \mathfrak{J}(\Lambda, \mu), \\ \int \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \rho(dm) &= \mathbb{E} \left[\int_0^T f(t, Y_t, \mu_t, \beta_t) dt + g(Y_T, \mu_T) \right] = \mathfrak{J}(\beta, \mu). \end{aligned}$$

This proves (1.5.8). \square

1.5.2 Proof of Theorem 1.5.1

Let $\Gamma^* \in \mathcal{K}$ be an optimal strategy for player A, which exists by Theorem 1.5.2. By applying a mimicking argument, stated and proven in Lemma 1.7.6 in Section 1.7.3, there exists a strategy $\hat{\Gamma}^* \in \mathcal{K}$ so that the following holds:

- The marginal distributions of Γ^* and $\hat{\Gamma}^*$ on $\mathcal{C}(\mathcal{P}^2)$ are the same: $\Gamma^*(\mathcal{C}^d \times \mathcal{V} \times \cdot) = \hat{\Gamma}^*(\mathcal{C}^d \times \mathcal{V} \times \cdot)$.
- Let (X, \mathbf{r}, μ) be such that $\hat{\Gamma}^* = \mathbb{P} \circ (X, \mathbf{r}, \mu)^{-1}$. Then \mathbf{r} is of the form $\mathbf{r}_t = \hat{q}_t(X_t, \mu)$, where $\hat{q} : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \rightarrow \mathcal{P}(A)$ is a measurable function.
- For every $\Sigma \in \mathcal{Q}$, it holds

$$\mathbf{p}(\Gamma^*, \Sigma) = \mathbf{p}(\hat{\Gamma}^*, \Sigma).$$

In particular, it holds

$$\inf_{\Sigma \in \mathcal{Q}} \mathbf{p}(\hat{\Gamma}^*, \Sigma) = \inf_{\Sigma \in \mathcal{Q}} \mathbf{p}(\Gamma^*, \Sigma) = \max_{\Gamma \in \mathcal{K}} \inf_{\Sigma \in \mathcal{Q}} \mathbf{p}(\Gamma, \Sigma) \geq 0. \quad (1.5.10)$$

Let $\mathfrak{U} = ((\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}}, \hat{\mathbb{P}}), \hat{\xi}, \hat{W}, \hat{\mu}, \hat{\mathfrak{t}})$ be as in Definition 1.5.1, so that $\hat{\Gamma}^* = \mathbb{P} \circ (\hat{X}, \hat{\mathfrak{t}}, \hat{\mu})^{-1}$. Recall that, by Lemma 1.7.6, $\hat{\mathfrak{t}}_t(da)dt = \hat{q}_t(\hat{X}_t, \hat{\mu})(da)dt \text{ Leb}_{[0,T]} \otimes \hat{\mathbb{P}}\text{-a.s.}$ By Assumption **B**, the set $K(t, x, m_t)$ defined by (1.5.1) is convex for every $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$. Therefore, by a well known measurable selection argument (see, e.g., [75, Lemma A.9]) there exists a measurable function $\hat{\alpha} : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \rightarrow A$ so that

$$\begin{aligned} \int_A b(t, x, m_t, a) \hat{q}_t(x, m)(da) &= b(t, x, m_t, \hat{\alpha}(t, x, m)), \\ f(t, x, m_t, \hat{\alpha}(t, x, m)) &\leq \int_A f(t, x, m_t, a) \hat{q}_t(x, m)(da). \end{aligned} \quad (1.5.11)$$

It follows that \hat{X} is a solution to equation

$$d\hat{X}_t = b(t, \hat{X}_t, \hat{\mu}_t, \hat{\alpha}(t, \hat{X}_t, \hat{\mu}))dt + d\hat{W}_t, \quad \hat{X}_0 = \hat{\xi} \quad (1.5.12)$$

as well, and the consistency condition (1.3.13) is still satisfied. By Lemma 1.7.7, we deduce that the solution \hat{X} to equation (1.5.12) can be taken adapted to the $\hat{\mathbb{P}}$ -augmentation of the filtration $(\sigma(\hat{\mu}) \vee \sigma(\hat{\xi}) \vee \sigma(\hat{W}_s : s \leq t))_{t \in [0, T]}$, and therefore there exists a progressively measurable function $\Phi : \mathcal{C}(\mathcal{P}^2) \times \mathbb{R}^d \times \mathcal{C}^d \rightarrow \mathcal{C}^d$ so that

$$\hat{X} = \Phi(\hat{\mu}, \hat{\xi}, \hat{W}) \quad \hat{\mathbb{P}}\text{-a.s.} \quad (1.5.13)$$

Set

$$\begin{aligned} \hat{\lambda} : [0, T] \times \mathcal{C}(\mathcal{P}^2) \times \mathbb{R}^d \times \mathcal{C}^d &\rightarrow A \\ (t, m, x, w) &\mapsto \hat{\lambda}_t(m, x, w) = \hat{\alpha}_t(\Phi_t(m, x, w), m_t); \\ \lambda = (\lambda_t)_{t \in [0, T]} &= (\hat{\lambda}_t(\hat{\mu}, \hat{\xi}, \hat{W}))_{t \in [0, T]}. \end{aligned} \quad (1.5.14)$$

Then, the progressively measurable processes $(\hat{\alpha}_t(\hat{X}_t, \hat{\mu}))_{t \in [0, T]}$ and $(\hat{\lambda}_t(\hat{\mu}, \hat{\xi}, \hat{W}))_{t \in [0, T]}$ are equal $\text{Leb}_{[0, T]} \otimes \hat{\mathbb{P}}\text{-a.s.}$, which implies that \hat{X} solves

$$d\hat{X}_t = b(t, \hat{X}_t, \hat{\mu}_t, \hat{\lambda}_t(\hat{\mu}, \hat{\xi}, \hat{W}))dt + d\hat{W}_t, \quad \hat{X}_0 = \hat{\xi}$$

as well, and the consistency condition is still satisfied. Set

$$(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) = (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho).$$

By Lemma 1.7.1, there exists a \mathbb{P}^0 -null set $N \subset \Omega^0$ so that the pair (Λ^*, μ^*) defined by

$$\begin{aligned} \Lambda^* : (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho) &\rightarrow (\mathbb{A}, \mathcal{B}_{\mathbb{A}}) \\ m &\mapsto \Lambda^*(m) = \begin{cases} (\hat{\lambda}_t(m, \cdot, \cdot))_{t \in [0, T]}, & m \in \Omega^0 \setminus N, \\ a_0 & m \in N, \end{cases} \\ \mu^* = \text{Id} : (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho) &\rightarrow (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho) \end{aligned} \quad (1.5.15)$$

is a correlated measure flow, where a_0 is an arbitrary point in A . Let X^* be the solution of (1.3.10) on the product probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ defined in (1.3.6). Note that the strategy associated to the admissible recommendation Λ^* strategy λ^*

is equal to $\hat{\lambda}_t(\mu^*, \xi, W) \text{Leb}_{[0,T]} \otimes \mathbb{P}$ -almost surely. Since uniqueness in law holds by Theorem B.1, it follows that

$$\mathbb{P} \circ (X^*, (\delta_{\lambda_t^*}(da))_{t \in [0,T]}, \mu^*)^{-1} = \hat{\mathbb{P}} \circ (\hat{X}, (\delta_{\hat{\alpha}(t, \hat{X}_t, \hat{\mu})}(da))_{t \in [0,T]}, \hat{\mu})^{-1}, \quad (1.5.16)$$

which implies that the consistency condition (1.3.13) is satisfied.

Finally, we verify that the correlated measure flow just defined satisfies the optimality condition (1.3.12). For any $\beta \in \mathbb{A}$, let $\Sigma \in \mathcal{Q}$ be as in point (ii) of Proposition 1.5.3. Then, by (1.5.16), (1.5.11) and (1.5.10), for every $\Sigma \in \mathcal{Q}$ it holds

$$\begin{aligned} \mathfrak{J}(\Lambda^*, \mu^*) &= \mathbb{E}^{\hat{\mathbb{P}}} \left[\int_0^T f(t, \hat{X}_t, \hat{\mu}_t, \hat{\alpha}(t, \hat{X}_t, \hat{\mu})) dt + g(\hat{X}_T, \hat{\mu}_T) \right] \\ &\leq \mathbb{E}^{\hat{\mathbb{P}}} \left[\int_0^T \int_A f(t, \hat{X}_t, \hat{\mu}_t, a) \hat{q}_t(\hat{X}_t, \hat{\mu}) dt + g(\hat{X}_T, \hat{\mu}_T) \right] = \int \mathfrak{F}(y, q, m) \hat{\Gamma}^*(dy, dq, dm) \\ &\leq \int \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \rho(dm) = \mathfrak{J}(\beta, \mu^*), \end{aligned}$$

which proves that (Λ^*, μ^*) satisfies the optimality condition and therefore is a mean field CCE.

Remark 1.5.2. As described in this section, the first step to find a coarse correlated solution is to find an optimal strategy for the maximizing player in the auxiliary zero-sum game in Definition 1.5.3, which exists by Theorem 1.5.2. Since this optimization problem is given by a linear payoff functional defined over a convex set of probability measures, linear programming methods could be employed to find approximate solutions. The study of computational methods for finding coarse correlated solutions to the MFG is beyond the scope of this work. We refer to [79, 99] for a linear programming approach to the computation of (true) correlated equilibria in N -player games.

1.5.3 Proof of Theorem 1.5.2

The main instrument is the following Minimax Theorem, due to K. Fan:

Theorem 1.5.4 ([62], Theorem 2). *Let X be a compact Hausdorff space and Y an arbitrary set (not topologized). Let $f : X \times Y \rightarrow \mathbb{R}$ be a real-valued function such that, for every $y \in Y$, $x \mapsto f(x, y)$ is lower semi-continuous on X . If $f(\cdot, y)$ is concave on X for every $y \in Y$ and $f(x, \cdot)$ convex on Y for every $x \in X$, then*

$$\max_{x \in X} \inf_{y \in Y} f(x, y) = \inf_{y \in Y} \max_{x \in X} f(x, y). \quad (1.5.17)$$

The following results aims at verifying that the auxiliary zero-sum game in Definition 1.5.3 satisfies the assumptions of Theorem 1.5.4. We start with some useful moment estimates for the solution to (1.5.3):

Lemma 1.5.5 (Estimates). *Let $\Gamma \in \mathcal{K}$, let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ be the tuple associated to Γ , as in Definition 1.5.1, and let X be the solution to (1.5.3). Then, for every $2 \leq p \leq \bar{p}$, there exists a constant $C = C(p, T, \eta, b, A)$ so that*

$$\mathbb{E} [\|X\|_{\mathcal{C}^d}^p] \leq C. \quad (1.5.18)$$

The proof is omitted as it is just a straightforward application of Gronwall's lemma. We recall the following fact, which will be used extensively and whose proof can be found in [105, Theorem 7.12]: given a metric space (E, d_E) , a sequence $(\mu_n)_n \subset (\mathcal{P}^p(E), \mathcal{W}_{p,E})$ is relatively compact if and only if it is tight and satisfies

$$\lim_{r \rightarrow \infty} \sup_n \int_{\{x: d_E^p(x, x_0) \geq r\}} d_E^p(x, x_0) \mu_n(dx) = 0. \quad (1.5.19)$$

Lemma 1.5.6. \mathcal{K} is pre-compact in $(\mathcal{P}^2(\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)), \mathcal{W}_{2, \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)})$.

Proof. Let $(\Gamma^n)_{n \geq 1}$ be a sequence in \mathcal{K} , let us show that it is pre-compact, which is equivalent to show that $(\Gamma^n)_{n \geq 1}$ is tight and condition (1.5.19) is satisfied. Moreover, by [84, Lemma A.2], relative compactness of the sequence $(\Gamma^n)_{n \geq 1}$ is equivalent to the relative compactness of each sequence of marginals on \mathcal{C}^d , $\mathcal{C}(\mathcal{P}^2)$ and \mathcal{V} .

Since A is compact by Assumptions **A**, the space \mathcal{V} is compact as well. Then, we automatically get both tightness of the sequence of the marginals on \mathcal{V} of $(\Gamma^n)_{n \geq 1}$ and property (1.5.19).

In the following, for every $n \geq 1$, let $\mathfrak{U}^n = ((\Omega^n, \mathcal{F}^n, \mathbb{F}^n, \mathbb{P}^n), \xi^n, W^n, \mu^n, \mathfrak{r}^n)$ and X^n be as in Definition 1.5.1, so that $\Gamma^n = \mathbb{P}^n \circ (X^n, \mu^n, \mathfrak{r}^n)^{-1}$. Let Γ_1^n be the law of X^n under \mathbb{P}^n . We prove the tightness by means of Kolmogorov-Čentsov criterion, as stated, e.g., in [80, Corollary 16.9]. Let $2 < p \leq \bar{p}$, $0 \leq s < t \leq T$. We have:

$$\begin{aligned} \mathbb{E}^n [|X_t^n - X_s^n|^p] &\leq C \mathbb{E}^n \left[\int_s^t \int_A |b(u, X_u^n, \mu_u^n, a)|^p \mathfrak{r}_u^n(da) du + |W_t - W_s|^p \right] \\ &\leq C \left(|t - s|^{p-1} \int_s^t \mathbb{E}^n \left[\int_A |b(u, X_u^n, \mu_u^n, a)|^p \mathfrak{r}_u^n(da) \right] du + |t - s|^{\frac{p}{2}} \right), \end{aligned}$$

for some positive constant C which is updated from line to line. For every $u \in [0, T]$, we have

$$\begin{aligned} &\mathbb{E}^n \left[\int_A |b(u, X_u^n, \mu_u^n, a)|^p \mathfrak{r}_u^n(da) \right] \\ &\leq C \mathbb{E}^n \left[|X_u^n|^p + \left(\int_{\mathbb{R}^d} |y|^2 \mu_u^n(dy) \right)^{\frac{p}{2}} + \int_A |a - a_0|^p \mathfrak{r}_u^n(da) + |b(u, 0, \delta_0, a_0)|^p \right] \\ &\leq C \left(1 + \mathbb{E}^n \left[|X_u^n|^p + \int_{\mathbb{R}^d} |y|^p \mu_u^n(dy) \right] \right) = C \left(1 + \mathbb{E}^n [|X_u^n|^p] + \mathbb{E}^n [|X_u^n|^p | \mu^n] \right) \\ &= C \left(1 + 2\mathbb{E}^n [|X_u^n|^p] \right) \leq C \left(1 + \mathbb{E}^n \left[\sup_{u \in [0, T]} |X_u^n|^p \right] \right) \leq C, \end{aligned} \quad (1.5.20)$$

where the last inequality follows from Lemma 1.5.5, with C independent of $n \geq 1$. Such a uniform bound implies that

$$\begin{aligned} \mathbb{E}^n [|X_t^n - X_s^n|^p] &\leq C \left(|t - s|^{p-1} \int_s^t \mathbb{E}^n \left[\int_A |b(u, X_u^n, \mu_u^n, a)|^p \mathfrak{r}_u^n(da) \right] du + |t - s|^{\frac{p}{2}} \right) \\ &\leq C \left(|t - s|^{p-1} |t - s| C(p, T, \eta, b, A) + |t - s|^{\frac{p}{2}} \right) \leq C \left(|t - s|^p + |t - s|^{\frac{p}{2}} \right) \\ &\leq C |t - s|^{\frac{p}{2}}. \end{aligned}$$

Set $\beta = p/2 - 1$, so that we get

$$\mathbb{E}^n [|X_t^n - X_s^n|^p] \leq C |t - s|^{1+\beta}, \quad (1.5.21)$$

with $p, \beta > 0$. Since $\mathbb{P}^n \circ (X_0^n)^{-1} = \eta \in \mathcal{P}^p(\mathbb{R}^d)$ for every $n \geq 1$, we have the tightness of the initial laws as well. This concludes of the proof of the tightness of $(\Gamma_1^n)_{n \geq 1}$. As for condition (1.5.19), we have:

$$\begin{aligned} \limsup_{r \rightarrow \infty} \sup_n \int_{\{y: \|y\|_{\mathcal{C}^d}^2 > r\}} \|y\|_{\mathcal{C}^d}^2 \Gamma_1^n(dy) &= \limsup_{r \rightarrow \infty} \sup_n \mathbb{E}^n \left[\|X^n\|_{\mathcal{C}^d}^2 \mathbb{1}_{\{\|X^n\|_{\mathcal{C}^d}^2 > r\}} \right] \\ &\leq \limsup_{r \rightarrow \infty} \sup_n \left(\mathbb{E}^n [\|X^n\|_{\mathcal{C}^d}^4] \right)^{\frac{1}{2}} \mathbb{P}^n (\|X^n\|_{\mathcal{C}^d}^2 > r)^{\frac{1}{2}} \leq C \limsup_{r \rightarrow \infty} \sup_n \mathbb{P}^n (\|X^n\|_{\mathcal{C}^d}^2 > r)^{\frac{1}{2}} \end{aligned}$$

for some positive constant C independent of n . By Markov's inequality and estimate (1.5.18) again, we get

$$\limsup_{r \rightarrow \infty} \sup_n \int_{\{y: \|y\|_{\mathcal{C}^d}^2 > r\}} \|y\|_{\mathcal{C}^d}^2 \Gamma_1^n(dy) \leq C \limsup_{r \rightarrow \infty} \sup_n \mathbb{E}^n [\|X^n\|_{\mathcal{C}^d}^2]^{\frac{1}{2}} r^{-\frac{1}{2}} = 0.$$

Finally, we turn to the sequence $(\rho^n)_{n \geq 1}$, where $\rho^n = \mathbb{P}^n \circ (\mu^n)^{-1}$. Let $\mathbb{P}^{n,m}(\cdot) = \mathbb{P}^n(\cdot \mid \mu = m)$ be the regular conditional distribution of \mathbb{P}^n given $\mu^n = m$. Then, $\mu_t^n = m_t$ $\mathbb{P}^{n,m}$ -a.s. and $\mathbb{P}^{n,m} \circ (X_t^n)^{-1} = m_t$ ρ -a.e. for every $t \in [0, T]$, which implies that, for every $s, t \in [0, T]$, we have

$$\mathbb{E}^{n,m} \left[\mathcal{W}_{2, \mathbb{R}^d}^p(\mu_t^n, \mu_s^n) \right] \leq \mathbb{E}^{n,m} [|X_t^n - X_s^n|^p]$$

for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$. Integrating with respect to ρ yields

$$\mathbb{E}^n \left[\mathcal{W}_{2, \mathbb{R}^d}^p(\mu_t^n, \mu_s^n) \right] \leq \mathbb{E}^n [|X_t^n - X_s^n|^p] \leq C |t - s|^{1+\beta}$$

where the last inequality follows from (1.5.21) with $\beta = p/2 - 1$. Since $\mathbb{P}^n \circ (\mu_0^n)^{-1} = \delta_\eta$, it is enough to apply again Kolmogorov-Čentsov criterion and deduce the tightness of $(\rho^n)_{n \geq 1}$. Finally, we verify condition (1.5.19). To this extent, we note that, for every $n \geq 1$, there exists a continuous modification of the process $(\mathbb{E}^n[|X_t^n|^2 \mid \mu^n])_{t \in [0, T]}$, so that it holds

$$\sup_{t \in [0, T]} \int_{\mathbb{R}^d} |y|^2 \mu_t^n(dy) = \sup_{t \in [0, T]} \mathbb{E}^n [|X_t^n|^2 \mid \mu_t^n] \quad \mathbb{P}^n\text{-a.s.}$$

Indeed, estimate (1.5.21) on the moments of X^n implies that $(\mathbb{E}^n[|X_t^n|^2 \mid \mu^n])_t$ satisfies

$$\begin{aligned} \mathbb{E}^n \left[\left| \mathbb{E}^n [|X_t^n|^2 \mid \mu^n] - \mathbb{E}^n [|X_s^n|^2 \mid \mu^n] \right|^p \right] &\leq \mathbb{E}^n \left[\left| |X_t^n|^2 - |X_s^n|^2 \right|^p \right] \\ &= \mathbb{E}^n [|X_t^n - X_s^n|^p |X_t^n + X_s^n|^p] \leq \mathbb{E}^n [|X_t^n - X_s^n|^{2p}]^{\frac{1}{2}} \mathbb{E}^n [|X_t^n + X_s^n|^{2p}]^{\frac{1}{2}} \leq C |t - s|^{\frac{p}{2}}, \end{aligned}$$

where we have used Cauchy-Schwartz inequality, (1.5.18) and (1.5.21) to bound, respectively, $\mathbb{E}^n[|X_t^n - X_s^n|^{2p}]^{1/2}$ and $\mathbb{E}^n[|X_t^n + X_s^n|^{2p}]^{1/2}$. Therefore, by choosing $2 < p <$

$\bar{p}/2$ and $\beta = p/2 - 1$ as above, we deduce from [80, Theorem 3.3], that there exists a continuous modification of $(\mathbb{E}^n[|X_t^n|^2 \mid \mu^n])_{t \in [0, T]}$. Then, observe that

$$\int_{\mathbb{R}^d} |y|^2 \mu_t^n(dy) = \mathbb{E}^n[|X_t^n|^2 \mid \mu^n] \quad \forall t \in [0, T] \cap \mathbb{Q}, \mathbb{P}^n\text{-a.s.}$$

Since both processes are almost surely continuous, we can take the supremum over every $t \in [0, T]$ to conclude that

$$\sup_{t \in [0, T]} \int_{\mathbb{R}^d} |y|^2 \mu_t^n(dy) = \sup_{t \in [0, T]} \mathbb{E}^n[|X_t^n|^2 \mid \mu^n] \leq \mathbb{E}^n \left[\sup_{t \in [0, T]} |X_t^n|^2 \mid \mu^n \right] \quad \mathbb{P}^n\text{-a.s.} \quad (1.5.22)$$

We are now ready to show that (1.5.19) holds for $(\rho^n)_{n \geq 1}$: by applying (1.5.22) in the first inequality, Cauchy-Schwartz and Markov inequalities, we have

$$\begin{aligned} & \lim_{r \rightarrow \infty} \sup_n \int_{\{m: \sup_{t \in [0, T]} \int_{\mathbb{R}^d} |y|^2 m_t(dy) > r\}} \sup_{t \in [0, T]} \int_{\mathbb{R}^d} |y|^2 m_t(dy) \rho^n(dm) \\ & \leq \lim_{r \rightarrow \infty} \sup_n \mathbb{E}^n \left[\mathbb{E}^n[\|X\|_{\mathcal{C}^d}^2 \mid \mu^n] \mathbf{1}_{\{\mathbb{E}^n[\|X\|_{\mathcal{C}^d}^2 \mid \mu^n] > r\}} \right] \\ & \leq \lim_{r \rightarrow \infty} \frac{1}{r^{\frac{1}{2}}} \sup_n \mathbb{E}^n[\|X^n\|_{\mathcal{C}^d}^4]^{\frac{1}{2}} \mathbb{E}^n[\|X^n\|_{\mathcal{C}^d}^2]^{\frac{1}{2}} \leq \lim_{r \rightarrow \infty} Cr^{-\frac{1}{2}} = 0, \end{aligned}$$

since the suprema over $n \geq 1$ are finite by Lemma 1.5.5. \square

Lemma 1.5.7. \mathcal{K} is closed in $(\mathcal{P}^2(\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)), \mathcal{W}_{2, \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)})$.

Proof. It is enough to prove that, for every sequence $(\Gamma^n)_{n \geq 1} \subseteq \mathcal{K}$ converging to Γ as $n \rightarrow \infty$ in $\mathcal{W}_{2, \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)}$, we have $\Gamma \in \mathcal{K}$. We work on the following canonical space: let $(\bar{\Omega}, \bar{\mathcal{G}})$ be given by

$$(\bar{\Omega}, \bar{\mathcal{G}}) = (\mathcal{C}^d \times \mathcal{C}(\mathcal{P}^2) \times \mathcal{V}, \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{V}} \otimes \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}).$$

We equip such a space with the filtration $\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}$ given by

$$\mathcal{G}_t = \mathcal{B}_{t, \mathcal{C}^d} \otimes \mathcal{F}_t^{\mathcal{V}} \otimes \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)},$$

where $\mathcal{B}_{t, \mathcal{C}^d} = \sigma(\mathcal{C}^d \ni x \mapsto x_s : s \leq t)$. Let x , m and q denote the projection from $\bar{\Omega}$ in \mathcal{C}^d , $\mathcal{C}(\mathcal{P}^2)$ and \mathcal{V} , respectively. Define the process $w = (w_t)_{t \in [0, T]}$ as

$$w_t = w_t(x, q, m) = x_t - x_0 - \int_0^t \int_A b(s, x_s, m_s, a) q_s(da) ds. \quad (1.5.23)$$

Observe that w is a continuous process on $(\bar{\Omega}, \bar{\mathcal{F}})$ and, by [84, Corollary A.5], for every $t \in [0, T]$ w_t is a continuous with at most linear growth function of (x, q, m) .

For every $n \geq 1$, let $\mathfrak{U}^n = ((\Omega^n, \mathcal{F}^n, \mathbb{F}^n, \mathbb{P}^n), \xi^n, W^n, \mu^n, \mathfrak{r}^n)$ and X^n be as in Definition 1.5.1, so that $\Gamma^n = \mathbb{P}^n \circ (X^n, \mu^n, \mathfrak{r}^n)^{-1}$. Since $\Gamma^n \circ (x_0, w, m, q, x)^{-1} = \mathbb{P}^n \circ (\xi^n, W^n, \mu^n, \mathfrak{r}^n, X^n)^{-1}$, we have that the tuple $\mathfrak{U}^n = ((\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{F}}^{\Gamma^n}, \Gamma^n), x_0, w, m, q)$ satisfies the requirements of Definition 1.5.1, where $\bar{\mathbb{F}}^{\Gamma^n}$ denotes the Γ^n -augmentation

of the filtration \mathbb{G} . We show that the tuple $\mathfrak{U} = ((\overline{\Omega}, \overline{\mathcal{F}}, \overline{\mathbb{F}}^\Gamma, \Gamma), x_0, w, m, q)$ satisfies the requirements of Definition 1.5.1, which implies $\Gamma \in \mathcal{K}$.

We start by the independence property of w, m and q under Γ . Let $(t^i)_{i=1}^k \subseteq [0, T]$, $\varphi^i \in \mathcal{C}_b(\mathbb{R}^d)$ for $i = 1, \dots, k$, $\psi \in \mathcal{C}_b(\mathcal{C}(\mathcal{P}^2))$, $\phi \in \mathcal{C}_b(\mathbb{R}^d)$ be bounded continuous functions. Since W^n, μ^n and ξ^n are independent under \mathbb{P}^n and $\Gamma^n \rightarrow \Gamma$ weakly, we have

$$\begin{aligned} & \mathbb{E}^{\Gamma^n} \left[\prod_{i=1}^k \varphi^i(w_{t^i}(x, q, m)) \psi(m) \varphi(x_0) \right] \\ &= \mathbb{E}^{\mathbb{P}^n} \left[\prod_{i=1}^k \varphi^i(W_{t^i}^n) \psi(\mu^n) \varphi(\xi^n) \right] = \mathbb{E}^{\mathbb{P}^n} \left[\prod_{i=1}^k \varphi^i(W_{t^i}^n) \right] \mathbb{E}^{\mathbb{P}^n} [\psi(\mu^n)] \mathbb{E}^{\mathbb{P}^n} [\varphi(\xi^n)] \\ &= \mathbb{E}^{\Gamma^n} \left[\prod_{i=1}^k \varphi^i(w_{t^i}(x, q, m)) \right] \mathbb{E}^{\Gamma^n} [\psi(m)] \mathbb{E}^{\Gamma^n} [\varphi(x_0)] \end{aligned}$$

where the first equality holds since $w_t(X^n, \mathbf{r}^n, \mu^n) = W_t^n$ for every $t \in [0, T]$ \mathbb{P}^n -a.s. Then, since $\phi^i \circ w_{t^i}$ is a continuous function of (x, q, m) for every i , weak convergence implies that

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{E}^{\Gamma^n} \left[\prod_{i=1}^k \varphi^i(w_{t^i}(x, q, m)) \psi(m) \varphi(x_0) \right] = \mathbb{E}^\Gamma \left[\prod_{i=1}^k \varphi^i(w_{t^i}(x, q, m)) \psi(m) \varphi(x_0) \right], \\ & \lim_{n \rightarrow \infty} \mathbb{E}^{\Gamma^n} \left[\prod_{i=1}^k \varphi^i(w_{t^i}^n(x, q, m)) \right] \mathbb{E}^{\Gamma^n} [\psi(m)] \mathbb{E}^{\Gamma^n} [\varphi(\xi^n)] \\ &= \mathbb{E}^\Gamma \left[\prod_{i=1}^k \varphi^i(w_{t^i}(x, q, m)) \right] \mathbb{E}^\Gamma [\psi(m)] \mathbb{E}^\Gamma [\varphi(x_0)]. \end{aligned} \tag{1.5.24}$$

This is enough to ensure the mutual independence under Γ of $(w_{t^i})_{i=1, \dots, k}$, x_0 and m for every $(t^i)_{i=1}^k \subset [0, T]$, which yields the independence of w, x_0 and m . Moreover, by taking ψ and ϕ identically equal to 1, equation (1.5.24) implies that w is natural Brownian motion under Γ , since the finite dimensional distributions of w coincide with the ones of a Brownian motion. Let us verify the independence of increments properties. Let $s > t$, $\varphi \in \mathcal{C}_b(\mathcal{C}^d)$ $\mathcal{B}_{t, \mathcal{C}^d}$ -measurable, $\chi \in \mathcal{C}_b(\mathcal{V})$ $\mathcal{F}_t^\mathcal{V}$ -measurable, $\psi \in \mathcal{C}_b(\mathcal{C}(\mathcal{P}^2))$ and $\phi \in \mathcal{C}_b(\mathbb{R}^d)$. Then, we have:

$$\begin{aligned} & \mathbb{E}^\Gamma [\phi(w_s - w_t) \varphi(x) \chi(q) \psi(m)] = \mathbb{E}^\Gamma [\phi(w_s(x, q, m) - w_t(x, q, m)) \varphi(x) \chi(q) \psi(m)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\Gamma^n} [\phi(w_s(x, q, m) - w_t(x, q, m)) \varphi(x) \chi(q) \psi(m)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}^n} [\phi(w_s(X^n, \mathbf{r}^n, \mu^n) - w_t(X^n, \mathbf{r}^n, \mu^n)) \varphi(X^n) \chi(\mathbf{r}^n) \psi(\mu^n)] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}^n} [\phi(W_s^n - W_t^n) \varphi(X^n) \chi(\mathbf{r}^n) \psi(\mu^n)] = 0, \end{aligned}$$

where the last equality holds since W^n is a \mathbb{F}^n -Brownian motion under \mathbb{P}^n , μ^n is \mathcal{F}_0^n -measurable and X^n and \mathbf{r}^n are both \mathbb{F}^n -adapted. By working with an approximating sequence, this holds also for bounded measurable φ, χ, ψ and ϕ , which is enough to conclude the independence of increments. Finally, since w is \mathbb{G} -Brownian motion, it remains so under the Γ -augmentation of \mathbb{G} .

Since $\Gamma^n \circ x_0^{-1} \equiv \eta$, we have that $\Gamma \circ x_0 = \eta$ as well. Moreover, since w is a $\overline{\mathbb{F}}^\Gamma$ -Brownian motion and x is $\overline{\mathbb{F}}^\Gamma$ -adapted by definition of the filtration, equation (1.5.23) implies that x is a solution to (1.5.3).

As for the consistency condition, observe that, for every $t \in [0, T]$, $\varphi \in \mathcal{C}_b(\mathbb{R}^d)$, $\psi \in \mathcal{C}_b(\mathcal{C}(\mathcal{P}^2))$, we have

$$\begin{aligned} \mathbb{E}^{\Gamma^n} \left[\int_{\mathbb{R}^d} \varphi(y) m_t(dy) \psi(m) \right] &= \mathbb{E}^{\mathbb{P}^n} \left[\int_{\mathbb{R}^d} \varphi(y) \mu_t^n(dy) \psi(\mu^n) \right] \\ &= \mathbb{E}^{\mathbb{P}^n} \left[\mathbb{E}^{\mathbb{P}^n} [\varphi(X_t^n) \psi(\mu^n) \mid \mu^n] \right] \\ &= \mathbb{E}^{\mathbb{P}^n} [\varphi(X_t^n) \psi(\mu^n)] = \mathbb{E}^{\Gamma^n} [\varphi(x_t) \psi(m)], \end{aligned}$$

since μ_t^n is a version of the conditional law under \mathbb{P}^n of X_t^n given μ^n . Therefore, by weak convergence we have both

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{E}^{\Gamma^n} [\varphi(x_t) \psi(m)] &= \mathbb{E}^\Gamma [\varphi(x_t) \psi(m)], \\ \lim_{n \rightarrow \infty} \mathbb{E}^{\Gamma^n} \left[\int_{\mathbb{R}^d} \varphi(y) m_t(dy) \psi(m) \right] &= \mathbb{E}^\Gamma \left[\int_{\mathbb{R}^d} \varphi(y) m_t(dy) \psi(m) \right] \end{aligned}$$

where the second limit holds since the function $m \mapsto \int_{\mathbb{R}^d} \varphi(y) m_t(dy) \in \mathcal{C}_b(\mathcal{C}(\mathcal{P}^2))$, which implies

$$\mathbb{E}^\Gamma \left[\int_{\mathbb{R}^d} \varphi(y) m_t(dy) \psi(m) \right] = \mathbb{E}^\Gamma [\varphi(x_t) \psi(m)].$$

This is enough to conclude that $m_t = \Gamma(x_t \in \cdot \mid m)$ Γ -a.s for every $t \in [0, T]$, since the random element (x_t, m) takes values in a Polish space. \square

Lemma 1.5.8 (Convexity). *\mathcal{K} and \mathcal{Q} are convex.*

Proof. We start by proving that \mathcal{K} is convex. Let Γ^i , $i = 1, 2$, be in \mathcal{K} , and let $\alpha \in (0, 1)$. Let $\mathfrak{U}^i = ((\Omega^i, \mathcal{F}^i, \mathbb{F}^i, \mathbb{P}^i), \xi^i, W^i, \mu^i, \mathfrak{r}^i)$ be as in Definition 1.5.1, so that $\Gamma^i = \mathbb{P}^i \circ (X^i, \mathfrak{r}^i, \mu^i)^{-1}$. Set $\Xi^i = (\xi^i, W^i, \mu^i, \mathfrak{r}^i, X^i)$. Without loss of generality, we can suppose that the tuples are defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ which supports also a Bernoulli random variable $U \sim B(\alpha)$, so that U and $(\Xi^i)_{i=1,2}$ are mutually independent. If needed, we can enlarge the filtration so that U is \mathcal{F}_0 -measurable. Let us consider the following random variables:

$$\begin{aligned} \xi^\alpha &= U\xi^1 + (1-U)\xi^2, & \mu^\alpha &= U\mu^1 + (1-U)\mu^2, \\ W^\alpha &= UW^1 + (1-U)W^2, & \mathfrak{r}^\alpha &= U\mathfrak{r}^1 + (1-U)\mathfrak{r}^2, \\ X^\alpha &= UX^1 + (1-U)X^2. \end{aligned} \tag{1.5.25}$$

Set $\Xi^\alpha = (\xi^\alpha, W^\alpha, \mu^\alpha, \mathfrak{r}^\alpha, X^\alpha)$ and $\Gamma^\alpha = \mathbb{P} \circ (X^\alpha, \mathfrak{r}^\alpha, \mu^\alpha)^{-1}$. Observe that the law of Ξ^1 under \mathbb{P} is the same as the law of Ξ^α conditionally to $U = 1$, as the two tuples coincide on the set $\{U = 1\}$, and analogously for $U = 0$. Therefore, for every Borel set B , we have

$$\begin{aligned} \mathbb{P}(\Xi^\alpha \in B) &= \mathbb{P}(\Xi^\alpha \in B \mid U = 1) \mathbb{P}(U = 1) + \mathbb{P}(\Xi^\alpha \in B \mid U = 0) \mathbb{P}(U = 0) \\ &= \mathbb{P}(\Xi^1 \in B) \mathbb{P}(U = 1) + \mathbb{P}(\Xi^2 \in B) \mathbb{P}(U = 0) \\ &= \alpha \mathbb{P}(\Xi^1 \in B) + (1 - \alpha) \mathbb{P}(\Xi^2 \in B). \end{aligned} \tag{1.5.26}$$

In particular, (1.5.26) implies that $\Gamma^\alpha = \alpha\Gamma^1 + (1 - \alpha)\Gamma^2$. Let us show that the tuple Ξ^α satisfies the requirements of Definition 1.5.1. By (1.5.26), ξ^α has law η and W^α is a natural Brownian motion. To see that it is an \mathbb{F} -Brownian motion, let $s < t$, $G \in \mathcal{F}_s$, $B \in \mathcal{B}_{\mathbb{R}^d}$: then

$$\begin{aligned} \mathbb{E}[\mathbf{1}_B(W_t^\alpha - W_s^\alpha)\mathbf{1}_G] &= \mathbb{E}[\mathbf{1}_B(W_t^\alpha - W_s^\alpha)\mathbf{1}_G\mathbf{1}_{\{U=1\}}] + \mathbb{E}[\mathbf{1}_B(W_t^\alpha - W_s^\alpha)\mathbf{1}_G\mathbf{1}_{\{U=0\}}] \\ &= \mathbb{E}[\mathbf{1}_B(U(W_t^1 - W_s^1) + (1-U)(W_t^2 - W_s^2) \in B)\mathbf{1}_G\mathbf{1}_{\{U=0\}}] \\ &\quad + \mathbb{E}[\mathbf{1}_B(U(W_t^1 - W_s^1) + (1-U)(W_t^2 - W_s^2) \in B)\mathbf{1}_G\mathbf{1}_{\{U=1\}}] \\ &= \mathbb{E}[\mathbf{1}_B(W_t^1 - W_s^1)\mathbf{1}_{G \cap \{U=1\}}] + \mathbb{E}[\mathbf{1}_B(W_t^1 - W_s^1)\mathbf{1}_{G \cap \{U=0\}}] = 0, \end{aligned}$$

since U is \mathcal{F}_0 -measurable by assumption. As for the mutual independence of ξ^α , μ^α and W^α , we have that the joint law factorizes in the product of the marginals: by using (1.5.26), since $(\xi^i, W^i)_{i=1,2}$ share the same joint law, one gets

$$\begin{aligned} &\mathbb{P}(\mu^\alpha \in A, W^\alpha \in B, \xi^\alpha \in C) \\ &= \alpha\mathbb{P}(\mu^1 \in A, W^1 \in B, \xi^1 \in C) + (1 - \alpha)\mathbb{P}(\mu^2 \in A, W^2 \in B, \xi^2 \in C) \\ &= \alpha\mathbb{P}(\mu^1 \in A)\mathbb{P}(W^1 \in B)\mathbb{P}(\xi^1 \in C) + (1 - \alpha)\mathbb{P}(\mu^2 \in A)\mathbb{P}(W^2 \in B)\mathbb{P}(\xi^2 \in C) \\ &= (\alpha\mathbb{P}(\mu^1 \in A) + (1 - \alpha)\mathbb{P}(\mu^2 \in A))\mathbb{W}^d(B)\eta(C) = \mathbb{P}(\mu^\alpha \in A)\mathbb{P}(W^\alpha \in B)\mathbb{P}(\xi^\alpha \in C). \end{aligned}$$

With similar arguments, one can show that for every $t \in [0, T]$, $g : \mathbb{R}^d \rightarrow \mathbb{R}$, $f : \mathcal{C}(\mathcal{P}^2) \rightarrow \mathbb{R}$ bounded and measurable, it holds

$$\mathbb{E}[g(X_t^\alpha)f(\mu^\alpha)] = \mathbb{E}\left[\int_{\mathcal{C}^d} g(y)\mu_t^\alpha(dy)f(\mu^\alpha)\right],$$

which implies that μ_t^α is a version of the conditional distribution of X_t^α given μ^α . Finally, consider the set

$$\Omega^1 = \left\{ X_t^1 = \xi^1 + \int_0^t \int_A b(s, X_s^1, \mu_s^1, a)\mathbf{r}_s^1(da)ds + W_t^1 \quad \forall t \in [0, T] \right\} \cap \{U = 1\},$$

and define analogously Ω^2 . We have that $\Omega^1 \cap \Omega^2 = \emptyset$ and $\mathbb{P}(\Omega^1) = \alpha$, since X^1 satisfies the equation above \mathbb{P} -a.s., and analogously $\mathbb{P}(\Omega^2) = 1 - \alpha$, so that $\mathbb{P}(\Omega^1 \cup \Omega^2) = 1$. On such a set, X^α satisfies the equation

$$X_t^\alpha = \xi^\alpha + \int_0^t \int_A b(s, X_s^\alpha, \mu_s^\alpha, a)\mathbf{r}_s^\alpha(da)ds + W_t^\alpha, \quad t \in [0, T].$$

Since X^α is \mathbb{F} -adapted, X^α is a solution to equation (1.5.3), which concludes this part of the proof.

Let us turn to the convexity of the set \mathcal{Q} . Let Σ^i , $i = 1, 2$, be in \mathcal{Q} , and $\alpha \in (0, 1)$. Let $\mathfrak{U}^i = ((\Omega^i, \mathcal{F}^i, \mathbb{F}^i, \mathbb{P}^i), \xi^i, W^i, \mathbf{r}^i)$ be as in Definition 1.5.2 so that $\Sigma^i(\cdot, m) = \mathbb{P}^i((X^{m,i}, \mathbf{r}^i) \in \cdot)$, where $X^{m,i}$ is the solution to equation (1.5.4) on $(\Omega^i, \mathcal{F}^i, \mathbb{F}^i, \mathbb{P}^i)$ when b is evaluated at $m \in \mathcal{C}(\mathcal{P}^2)$. Let $\Theta^i = \mathbb{P}^i \circ (\xi^i, W^i, \mathbf{r}^i)^{-1}$, and consider the maps \mathcal{I}_{Θ^i} defined by

$$\begin{aligned} \mathcal{I}_{\Theta^i} : \mathcal{C}(\mathcal{P}^2) &\longrightarrow \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}^d \times \mathcal{V}) \\ m &\longmapsto \mathcal{I}_{\Theta^i}(m) = \mathbb{P}^i \circ (\xi^i, W^i, X^{m,i}, \mathbf{r}^i)^{-1}. \end{aligned} \tag{1.5.27}$$

Similarly as for the set \mathcal{K} , suppose that the tuples are defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supporting also a Bernoulli random variable $U \sim B(\alpha)$, so that U and $(\xi^i, W^i, \mathbf{r}^i)_{i=1,2}$ are mutually independent. If needed, we can enlarge the filtration so that U is \mathcal{F}_0 -measurable. Let ξ^α, W^α and \mathbf{r}^α be as in (1.5.25), and, for every $m \in \mathcal{C}(\mathcal{P}^2)$, define

$$X^{\alpha,m} = UX^{1,m} + (1-U)X^{2,m}.$$

Let $\Theta^\alpha = \mathbb{P}^\alpha \circ (\xi^\alpha, W^\alpha, \mathbf{r}^\alpha)^{-1}$ and consider the map $\mathcal{I}_{\Theta^\alpha}$, defined analogously to above. By point (ii) of Lemma 1.7.4, it induces a stochastic kernel $\Sigma^\alpha \in \mathcal{Q}$. By working in the same way as in the case of \mathcal{K} , we can show that $\mathcal{I}_{\Theta^\alpha}(m) = \alpha\mathcal{I}_{\Theta^1}(m) + (1-\alpha)\mathcal{I}_{\Theta^2}(m)$ for each $m \in \mathcal{C}(\mathcal{P}^2)$, which implies that $\Sigma^\alpha = \alpha\Sigma^1 + (1-\alpha)\Sigma^2 \in \mathcal{Q}$. \square

Proposition 1.5.9. *The map $\mathcal{K} \times \mathcal{Q} \ni (\Gamma, \Sigma) \mapsto \mathbf{p}(\Gamma, \Sigma)$ is bilinear. Moreover, $\mathcal{K} \ni \Gamma \mapsto \mathbf{p}(\Gamma, \Sigma)$ is continuous for every $\Sigma \in \mathcal{Q}$.*

Proof. Bilinearity is clear, hence we focus on the continuity of $\mathbf{p}(\cdot, \Sigma)$ for fixed Σ . Take $(\Gamma^n)_{n \geq 1}$, Γ in \mathcal{K} and suppose $\Gamma^n \rightarrow \Gamma$ in the 2-Wasserstein distance. We treat separately the term depending just upon $\Gamma \in \mathcal{K}$ and the term depending also upon $\Sigma \in \mathcal{Q}$ in (1.5.7).

By [105, Theorem 7.12], $\Gamma^n \rightarrow \Gamma$ in 2-Wasserstein metrics if and only if

$$\int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \psi(y, q, m) \Gamma^n(dy, dq, dm) \rightarrow \int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \psi(y, q, m) \Gamma(dy, dq, dm), \quad (1.5.28)$$

for every ψ continuous with at most quadratic growth; hence, we just need to show that the functional \mathfrak{F} defined in (1.5.6) is continuous with at most quadratic growth. By Assumptions **A** and [84, Corollary A.5], we have that $\mathfrak{F}(y, q, m)$ is continuous. It is straightforward to verify that \mathfrak{F} has at most quadratic growth, in the sense that

$$\mathfrak{F}(y, q, m) \leq C \left(1 + \|y\|_{\mathcal{C}^d}^2 + \sup_{t \in [0, T]} \int_{\mathbb{R}^d} |y|^2 m_t(dy) + \int_0^T \int_A |a - a_0|^2 q_s(da) ds \right).$$

Therefore, we get continuity of the term depending only upon Γ .

Denote by ρ^n and ρ the marginal of Γ^n and Γ on $\mathcal{C}(\mathcal{P}^2)$. We can manipulate the term depending both upon Γ and Σ as

$$\begin{aligned} & \int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \rho(dm) \\ &= \int_{\mathcal{C}(\mathcal{P}^2)} \left(\int_{\mathcal{C}^d \times \mathcal{V}} \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \right) \rho(dm) = \int_{\mathcal{C}(\mathcal{P}^2)} g(m) \rho(dm) \end{aligned}$$

where we set

$$g(m) = \int_{\mathcal{C}^d \times \mathcal{V}} \mathfrak{F}(y, q, m) \Sigma(dy, dq, m).$$

We must show that $g : \mathcal{C}(\mathcal{P}^2) \rightarrow \mathbb{R}$ is continuous with at most quadratic growth with respect to the 2-Wasserstein distance. As for the growth condition, estimate (1.7.20) in Lemma 1.7.4 proves that g has at most quadratic growth in $m \in \mathcal{C}(\mathcal{P}^2)$. As for the continuity, let $(m^n)_{n \geq 1}, m \in \mathcal{C}(\mathcal{P}^2)$ so that $m^n \rightarrow m$ in $\mathcal{W}_{2, \mathcal{C}^d}$. Note

that $\Sigma(dy, dq, m^n) \rightarrow \Sigma(dy, dq, m)$ in $\mathcal{W}_{2, \mathcal{C}^d \times \mathcal{V}}$, as implied by Lemma 1.7.4. Define $\phi^n(y, q) = \mathfrak{F}(y, q, m^n)$. Since the cost functions are locally Lipschitz, we have that ϕ^n converges to ϕ uniformly on bounded sets of $\mathcal{C}^d \times \mathcal{V}$. This is enough to conclude that

$$g(m^n) = \int_{\mathcal{C}^d \times \mathcal{V}} \phi^n(y, q) \Sigma(dy, dq, m^n) \rightarrow \int_{\mathcal{C}^d \times \mathcal{V}} \phi(y, q) \Sigma(dy, dq, m) = g(m)$$

as $m^n \rightarrow m$. \square

We can now prove points (i) and (ii) of Theorem 1.5.2: take $X = \mathcal{K}$, $Y = \mathcal{Q}$ and $f(x, y) = \mathfrak{p}(\Gamma, \Sigma)$ in the statement of Theorem 1.5.4. By Lemmata 1.5.6 and 1.5.7, \mathcal{K} is compact with the topology of convergence in 2-Wasserstein distance and both sets \mathcal{K} and \mathcal{Q} are convex by Lemma 1.5.8. By Proposition 1.5.9, the payoff \mathfrak{p} is both concave and continuous in Γ for every fixed $\Sigma \in \mathcal{Q}$ and convex in Σ for every fixed Γ . Therefore, Theorem 1.5.4 yields the existence of both the value v of the auxiliary zero-sum game and an optimal strategy for player A. The next proposition proves point (iii), concluding the proof of Theorem 1.5.2.

Proposition 1.5.10 (Positivity of the value of the auxiliary zero-sum game). *Let v be the value of the zero-sum game defined in Definition 1.5.3 has a value v . Then $v \geq 0$.*

Proof. We show that, for every $\Sigma \in \mathcal{Q}$ there exists a strategy $\Gamma_\Sigma \in \mathcal{K}$ so that $\mathfrak{p}(\Gamma_\Sigma, \Sigma) = 0$. Fix $\Sigma \in \mathcal{Q}$, let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mathfrak{b})$ be a tuple as in Definition 1.5.2 so that $\Sigma(\cdot, m) = \mathbb{P}((X^m, \mathfrak{b}) \in \cdot)$, for every $m \in \mathcal{C}(\mathcal{P}^2)$. On this probability space, consider the following stochastic differential equation of McKean-Vlasov type:

$$\begin{cases} dY_t = \int_A b(t, Y_t, p_t, a) \mathfrak{b}_t(da) dt + dW_t, & t \in [0, T], & Y_0 = \xi; \\ \mathcal{L}(Y_t) = p_t, & t \in [0, T], & p = (p_t)_{t \in [0, T]} \in \mathcal{C}(\mathcal{P}^2). \end{cases} \quad (1.5.29)$$

Under Assumptions **A**, there exists a unique pair (Y, p) satisfying (1.5.29), where $Y = (Y_t)_{t \in [0, T]}$ is an \mathbb{F} -adapted continuous process so that $\mathbb{E}[\sup_{t \in [0, T]} |Y_t|^2] < \infty$, as ensured by, e.g., [38, Theorem 4.21], which implies that p actually belongs to $\mathcal{C}(\mathcal{P}^2)$. Define the deterministic flow of measures μ by setting $\mu = p$. Define Γ_Σ as $\mathbb{P} \circ (Y, \mu, \mathfrak{b})^{-1}$. Since $\mu = p$ is deterministic and (Y, p) is a solution to (1.5.29), μ is \mathcal{F}_0 -measurable and independent of ξ and W , and consistency condition holds trivially. This implies that Γ_Σ belongs to \mathcal{K} . By writing the integrals in \mathfrak{p} as expectations, we have:

$$\begin{aligned} & \mathfrak{p}(\Gamma_\Sigma, \Sigma) \\ &= \int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \rho_\Sigma(dm) - \int_{\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}(\mathcal{P}^2)} \mathfrak{F}(y, q, m) \Gamma_\Sigma(dy, dq, dm) \\ &= \mathbb{E} \left[\int_0^T \int_A f(t, Y_t, p_t, a) \mathfrak{b}_t(da) dt + g(Y_T, p_T) \right] \\ & \quad - \mathbb{E} \left[\int_0^T \int_A f(t, Y_t, p_t, a) \mathfrak{b}_t(da) dt + g(Y_T, p_T) \right] = 0, \end{aligned}$$

where $\rho_\Sigma(\cdot) = \delta_p(\cdot)$ denotes the marginal law of Γ_Σ on $\mathcal{C}(\mathcal{P}^2)$. Since such a construction holds for every $\Sigma \in \mathcal{Q}$, we have

$$\sup_{\Gamma \in \mathcal{K}} \mathbf{p}(\Gamma, \Sigma) \geq \mathbf{p}(\Gamma_\Sigma, \Sigma) = 0 \quad \forall \Sigma \in \mathcal{Q}.$$

Taking the infimum with respect to $\Sigma \in \mathcal{Q}$, we have

$$\inf_{\Sigma \in \mathcal{Q}} \sup_{\Gamma \in \mathcal{K}} \mathbf{p}(\Gamma, \Sigma) \geq \mathbf{p}(\Gamma_\Sigma, \Sigma) \geq 0,$$

which shows that v^B is non-negative. Since $v^A = v^B = v$, this proves that the value of the auxiliary zero-sum game is non-negative. \square

1.6 An example of coarse correlated solution to the mean field game

Taking inspiration from the work of Bardi and Fischer [11] and Lacker's papers [85, 87], we provide a simple example of a mean field game possessing mean field CCEs with non-deterministic flow of measures μ . Consistently with [11, 85, 87] but differently from the rest of the chapter, we consider a payoff, to be maximized, instead of a cost.

The MFG is as follows: We consider $d = 1$, $A = [a, b]$, with $a < 0 < b$, and $\eta = \delta_0$. For $m \in \mathcal{P}(\mathbb{R})$, denote by \bar{m} its mean $\int_{\mathbb{R}} ym(dy)$. Consider the following coefficients and profit functions:

$$b(t, x, m, a) = a, \quad f(t, x, m, a) = 0, \quad g(x, m) = cx\bar{m},$$

with $c > 0$ a positive constant. Observe that they satisfy the requirements of Assumptions **A**. We want to find a coarse correlated solution for the mean field game whose payoff functional, to be maximized, is given by

$$\mathfrak{J}(\Lambda, \mu) = \mathbb{E}[cX_T\bar{\mu}_T], \tag{1.6.1}$$

under the constraint

$$X_t = \int_0^t \lambda_s ds + W_t, \quad 0 \leq t \leq T, \tag{1.6.2}$$

where λ is the strategy associated to an admissible recommendation Λ in the sense of (1.3.7).

The rest of the section is organised as follows: In Section 1.6.1, we show that there exist infinitely many coarse correlated solutions to the MFG with non-deterministic flow of measures. Such solutions are neither classical MFG solutions nor a randomization (or a mixture, in the language of [85, 87]) of the classical MFG solutions. In Section 1.6.2, we compare coarse correlated solutions to the MFG with Lacker's weak solutions to MFG without common noise, as they both may feature a random flow of measures which is not due to a common noise that acts equally on the players. By the study of this specific MFG, we show that there exist infinitely many mean field CCEs which are not weak solutions to the MFG without common noise. Moreover, we show that weak MFG solutions which satisfy an additional measurability constraint are indeed coarse correlated solution to the MFG as well.

1.6.1 Exhibiting explicit coarse correlated solutions

Set $\Omega^0 = \{1, 2\}^2$, $\mathcal{F}^{0-} = 2^{\Omega^0}$ the power set and, given some probability measure $\mathbb{P}^0 \in \mathcal{P}(\Omega^0, \mathcal{F}^{0-})$, we set $\mathbb{P}^0((i, j)) = p_{i,j}$, so that $p_{i,j} \geq 0$ for all i, j and $\sum_{i,j=1}^2 p_{i,j} = 1$. Consider the following open loop strategies and flows of measures:

$$\begin{aligned}\alpha_t^+(\omega_*) &\equiv b, & \mu^+ &= (\mathbb{P}^* \circ (tb + W_t)^{-1})_{t \in [0, T]}; \\ \alpha_t^-(\omega_*) &\equiv a, & \mu^- &= (\mathbb{P}^* \circ (ta + W_t)^{-1})_{t \in [0, T]}.\end{aligned}\tag{1.6.3}$$

It was shown in [11] that the pairs (α^+, μ^+) and (α^-, μ^-) are two non-equivalent open-loop solutions of the mean field game, with initial distribution $\eta = \delta_0$, where by “non-equivalent” we mean that the flows of measures μ^+ and μ^- do not coincide. We point out that this result holds for more general initial distributions $\eta \in \mathcal{P}(\mathbb{R})$, see [11, Definition 3.1 and Theorem 3.1]. Choose $a_1, a_2 \in [0, 1]$, and set:

$$\begin{aligned}\mu^1 &= (a_1 \mu_t^+ + (1 - a_1) \mu_t^-)_{t \in [0, T]}, \\ \mu^2 &= (a_2 \mu_t^+ + (1 - a_2) \mu_t^-)_{t \in [0, T]}.\end{aligned}\tag{1.6.4}$$

Define (Λ, μ) in the following way:

$$(\Lambda, \mu)((i, j)) = \begin{cases} (\alpha^+, \mu^1) & (i, j) = (1, 1), \\ (\alpha^+, \mu^2) & (i, j) = (1, 2), \\ (\alpha^-, \mu^1) & (i, j) = (2, 1), \\ (\alpha^-, \mu^2) & (i, j) = (2, 2). \end{cases}\tag{1.6.5}$$

We claim that, as long as $a < 0 < b$, for every $T, c > 0$ there exists a probability measure $(p_{i,j})_{i,j=1,2}$ and a suitable choice of the parameters $(a_i)_{i=1,2}$ so that the tuple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ is a coarse correlated solution of the mean field game according to Definition 1.3.4.

First of all, as shown in Example 1.3.1 in Section 1.3, since Λ takes only two values, it is admissible. Therefore, the tuple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ is a correlated measure flow.

Let us begin with the consistency condition. We first observe that, when the state equation is controlled by α^+ (respectively, α^-), the law of the state process at time t , X_t , is exactly μ_t^+ (respectively, μ_t^-), for every time $t \in [0, T]$.

Suppose that $p_{1,1} + p_{2,1}$ and $p_{1,2} + p_{2,2}$ are both strictly positive. Then, observe that

$$\mathbb{P}(X_t \in \cdot \mid \mu)(\omega_0) = \mathbb{P}(X_t \in \cdot \mid \mu)(\omega_0) = \begin{cases} \mathbb{P}(X_t \in \cdot \mid \mu = \mu^1) & \text{if } \mu(\omega_0) = \mu^1, \\ \mathbb{P}(X_t \in \cdot \mid \mu = \mu^2) & \text{if } \mu(\omega_0) = \mu^2. \end{cases}$$

We can compute explicitly such a conditional probability. Fix $A \in \mathcal{B}_{\mathbb{R}^d}$:

$$\begin{cases} \mathbb{P}(X_t \in A \mid \mu = \mu^1) \\ \mathbb{P}(X_t \in A \mid \mu = \mu^2) \end{cases} = \begin{cases} \frac{p_{1,1}}{p_{1,1}+p_{2,1}} \mathbb{P}(X_t^+ \in A) + \frac{p_{2,1}}{p_{1,1}+p_{2,1}} \mathbb{P}(X_t^- \in A) & \text{if } \mu(\omega_0) = \mu^1, \\ \frac{p_{1,2}}{p_{1,2}+p_{2,2}} \mathbb{P}(X_t^+ \in A) + \frac{p_{2,2}}{p_{1,2}+p_{2,2}} \mathbb{P}(X_t^- \in A) & \text{if } \mu(\omega_0) = \mu^2. \end{cases}$$

In order to satisfy the consistency condition, it must hold

$$\begin{cases} \frac{p_{1,1}}{p_{1,1}+p_{2,1}} \mathbb{P}(X_t^+ \in A) + \frac{p_{2,1}}{p_{1,1}+p_{2,1}} \mathbb{P}(X_t^- \in A) = \mu_t^1(A) \\ \frac{p_{1,2}}{p_{1,2}+p_{2,2}} \mathbb{P}(X_t^+ \in A) + \frac{p_{2,2}}{p_{1,2}+p_{2,2}} \mathbb{P}(X_t^- \in A) = \mu_t^2(A) \end{cases}$$

for every $A \in \mathcal{B}_{\mathbb{R}^d}$. By definition of μ^1 and μ^2 ,

$$\begin{cases} \frac{p_{1,1}}{p_{1,1}+p_{2,1}}\mu_t^+(A) + \frac{p_{2,1}}{p_{1,1}+p_{2,1}}\mu_t^-(A) = a_1\mu_t^+(A) + (1-a_1)\mu_t^-(A), \\ \frac{p_{1,2}}{p_{1,2}+p_{2,2}}\mu_t^+(A) + \frac{p_{2,2}}{p_{1,2}+p_{2,2}}\mu_t^-(A) = a_2\mu_t^+(A) + (1-a_2)\mu_t^-(A) \end{cases}$$

which holds if and only if

$$\begin{cases} \frac{p_{1,1}}{p_{1,1}+p_{2,1}} = a_1, \\ \frac{p_{1,2}}{p_{1,2}+p_{2,2}} = a_2. \end{cases} \quad (1.6.6)$$

We can regard (1.6.6) as the consistency condition.

We now turn our attention to the optimality condition. Consider $\gamma = \mathbb{P} \circ (\Lambda, \mu)^{-1} = \mathbb{P}^0 \circ (\Lambda, \mu)^{-1} \in \mathcal{P}(\mathbb{A} \times \mathcal{C}(\mathcal{P}^2))$ and $\rho = \mathbb{P} \circ \mu^{-1} = \mathbb{P}^0 \circ \mu^{-1} \in \mathcal{P}(\mathcal{C}(\mathcal{P}^2))$. As in Remark 1.3.3, we can rewrite the optimality condition using disintegration of measures as

$$\int_{\mathbb{A} \times \mathcal{P}^2(\mathcal{C}^d)} \mathfrak{J}(\alpha, m) \gamma(d\alpha, dm) \geq \int_{\mathcal{P}^2(\mathcal{C}^d)} \mathfrak{J}(\beta, m) \rho(dm) \quad \forall \beta \in \mathbb{A},$$

where

$$\mathfrak{J}(\theta, m) = \mathbb{E}[cX_T \bar{m}_T], \quad X_t = \int_0^t \theta_s ds + W_t, \quad 0 \leq t \leq T,$$

for m in $\mu(\Omega^0) := \{\mu^1, \mu^2\}$ and $\theta = \alpha$ in $\Lambda(\Omega^0) := \{\alpha^+, \alpha^-\} \subseteq \mathbb{A}$ on the left-hand side of the inequality above and $\theta = \beta$ in \mathbb{A} on the right-hand side. We rewrite explicitly the inequality as

$$\begin{aligned} \mathfrak{J}(\Lambda, \mu) - \mathfrak{J}(\beta, \mu) &= p_{1,1} (\mathfrak{J}(\alpha^+, \mu^1) - \mathfrak{J}(\beta, \mu^1)) + p_{1,2} (\mathfrak{J}(\alpha^+, \mu^2) - \mathfrak{J}(\beta, \mu^2)) \\ &\quad + p_{2,1} (\mathfrak{J}(\alpha^-, \mu^1) - \mathfrak{J}(\beta, \mu^1)) + p_{2,2} (\mathfrak{J}(\alpha^-, \mu^2) - \mathfrak{J}(\beta, \mu^2)) \geq 0. \end{aligned} \quad (1.6.7)$$

Therefore, using (1.6.1), we have

$$\begin{aligned} \mathfrak{J}(\Lambda, \mu) - \mathfrak{J}(\beta, \mu) &= p_{1,1} (cT^2 b(a_1 b + (1-a_1)a) - cM(\beta)T(a_1 b + (1-a_1)a)) \\ &\quad + p_{1,2} (cT^2 b(a_2 b + (1-a_2)a) - cM(\beta)T(a_2 b + (1-a_2)a)) \\ &\quad + p_{2,1} (cT^2 a(a_1 b + (1-a_1)a) - cM(\beta)T(a_1 b + (1-a_1)a)) \\ &\quad + p_{2,2} (cT^2 a(a_2 b + (1-a_2)a) - cM(\beta)T(a_2 b + (1-a_2)a)), \end{aligned}$$

where $M(\beta) := \mathbb{E}[\int_0^T \beta_t dt] = \mathbb{E}[X_T^\beta]$. We can set $m(\beta) := 1/T M(\beta) = 1/T \mathbb{E}[\int_0^T \beta_t dt]$. Observe that $m(\beta) \in [a, b]$, being the mean of an A -valued process, and $m(\mathbb{A}) = [a, b]$, since for every $c \in [a, b]$ the constant process $\beta \equiv c$ belongs to \mathbb{A} . We divide by cT^2 to obtain the following condition:

$$\begin{aligned} &p_{1,1} (b(a_1 b + (1-a_1)a) - m(\beta)(a_1 b + (1-a_1)a)) \\ &\quad + p_{1,2} (b(a_2 b + (1-a_2)a) - m(\beta)(a_2 b + (1-a_2)a)) \\ &\quad + p_{2,1} (a(a_1 b + (1-a_1)a) - m(\beta)(a_1 b + (1-a_1)a)) \\ &\quad + p_{2,2} (a(a_2 b + (1-a_2)a) - m(\beta)(a_2 b + (1-a_2)a)) \\ &\geq 0. \end{aligned} \quad (1.6.8)$$

The condition above can be seen as a positivity condition for a real affine function of $g(m)$, $m \in [a, b]$, i.e.

$$\begin{aligned} \inf_{m \in [a, b]} g(m) &= \inf_{m \in [a, b]} h((p_{i,j})_{i,j=1,2}, (a_i)_{i=1,2}; a, b)m + k((p_{i,j})_{i,j=1,2}, (a_i)_{i=1,2}; a, b) \geq 0 \\ \begin{cases} h((p_{i,j})_{i,j=1,2}, (a_i)_{i=1,2}; a, b) &= -\{p_{1,1}(a_1b + (1-a_1)a) + p_{1,2}(a_2b + (1-a_2)a) \\ &\quad + p_{2,1}(a_1b + (1-a_1)a) + p_{2,2}(a_2b + (1-a_2)a)\}, \\ k((p_{i,j})_{i,j=1,2}, (a_i)_{i=1,2}; a, b) &= p_{1,1}b(a_1b + (1-a_1)a) + p_{1,2}b(a_2b + (1-a_2)a) \\ &\quad + p_{2,1}a(a_1b + (1-a_1)a) + p_{2,2}a(a_2b + (1-a_2)a). \end{cases} \end{aligned} \quad (1.6.9)$$

We now impose the consistency condition (1.6.6) to get:

$$\begin{aligned} h((p_{i,j})_{i,j=1,2}; a, b) &= -b \left(\frac{p_{1,1}^2 + p_{2,1}p_{1,1}}{p_{1,1} + p_{2,1}} + \frac{p_{1,2}^2 + p_{1,2}p_{2,2}}{p_{1,2} + p_{2,2}} \right) \\ &\quad - a \left(\frac{p_{2,1}^2 + p_{2,1}p_{1,1}}{p_{1,1} + p_{2,1}} + \frac{p_{2,2}^2 + p_{1,2}p_{2,2}}{p_{1,2} + p_{2,2}} \right), \\ k((p_{i,j})_{i,j=1,2}; a, b) &= b^2 \left(\frac{p_{1,1}^2}{p_{1,1} + p_{2,1}} + \frac{p_{1,2}^2}{p_{1,2} + p_{2,2}} \right) + a^2 \left(\frac{p_{2,1}^2}{p_{1,1} + p_{2,1}} + \frac{p_{2,2}^2}{p_{1,2} + p_{2,2}} \right) \\ &\quad + 2ab \left(\frac{p_{1,1}p_{2,1}}{p_{1,1} + p_{2,1}} + \frac{p_{1,2}p_{2,2}}{p_{1,2} + p_{2,2}} \right). \end{aligned} \quad (1.6.10)$$

Observe that imposing the consistency condition (1.6.6) reduces the number of parameters but makes the problem nonlinear in the probabilities $(p_{i,j})_{i,j=1,2}$.

Looking at (1.6.9) and (1.6.10), we observe that it implies that every randomization of the open loop MFG solutions (α^+, μ^+) and (α^-, μ^-) is a coarse correlated solution to the MFG as well, and it covers the case treated in [11], for any choices of $a < 0 < b$. To see this, consider a probability measures $\mathbb{P}^0 = (p_{i,j})_{i,j=1,2}$ so that $p_{1,2} = p_{2,1} = 0$ and, therefore, $p_{2,2} = 1 - p$ and $p_{1,1} = p \in [0, 1]$. For any p , such probability measure is a randomization of the MFG solutions (α^+, μ^+) and (α^-, μ^-) . Equations (1.6.10) take the simpler form

$$\begin{aligned} h((p, 0, 0, 1-p); a, b) &= -bp - a(1-p), \\ k((p, 0, 0, 1-p); a, b) &= b^2p + a^2(1-p). \end{aligned} \quad (1.6.11)$$

If $h((p, 0, 0, 1-p); a, b) \geq 0$, then condition (1.6.9) becomes

$$\inf_{m \in [a, b]} -bm + b^2 = a^2(1-p) - ab(1-p) \geq 0,$$

which is satisfied if $a < 0$ and $b > 0$, for any $p \in [0, 1]$. If instead $h((p, 0, 0, 1-p); a, b) < 0$, we have

$$\inf_{m \in [a, b]} -bm + b^2 = b^2(1-p) - ab(1-p) \geq 0,$$

which is again satisfied if $a < 0$ and $b > 0$, for every $p \in [0, 1]$. Observe that, when $p = 1$, $\mu \equiv \mu^1 = \mu^+$, while, when $p = 0$, $\mu \equiv \mu^2 = \mu^-$. This shows that the deterministic correlated measure flows $(\Lambda, \mu) \equiv (\alpha^+, \mu^+)$ and $(\Lambda, \mu) \equiv (\alpha^-, \mu^-)$ are indeed mean field CCE in the sense of Definition 1.3.4.




Figure 1.1: Existence of mean field CCEs as the probability measure $(p_{i,j})_{i,j=1,2}$ varies. Yellow points correspond to the values of $(p_{i,j})$ associated to coarse correlated solutions. The complement set of CCEs in each pseudo-simplex is indicated in blue. We fix $p_{2,1} \in [0, 1)$ and let the other parameters vary, so that every point $(p_{1,1}, p_{1,2}, p_{2,2})$ in each pseudo 2-dimensional simplex is such that $p_{1,1} + p_{1,2} + p_{2,2} = 1 - p_{2,1}$, $p_{i,j} \geq 0$. We consider $[a, b] = [-1, 1]$ and $p_{2,1}$ equal to 0.0, 0.3, and 0.7, respectively.

Turning to more interesting cases, consider $[a, b] = [-1, 1]$. The choice of a symmetric interval is not necessary, but it has been made to ease the comparison with previous results in the literature (see the next subsection). Figure 1.1 shows the existence of coarse correlated mean field equilibria as the probability measure $(p_{i,j})_{i,j=1,2}$ varies. Yellow spots in Figure 1.1 refer to those probability measures on $(\Omega^0, \mathcal{F}^{0-})$ so that (Λ, μ) is indeed a mean field CCE. In particular, it shows the existence of infinitely many mean field CCEs for the system. Observe that there exist infinitely many coarse correlated solutions of the mean field game so that Λ is not a deterministic function of μ , i.e., they are not a randomization of the solutions (α^+, μ^+) and (α^-, μ^-) : they correspond to those probability measures $(p_{i,j})_{i=1}^2$ so that $p_{1,2} \cdot p_{2,1} > 0$. Referring to Figure 1.1, they correspond to yellow points in the interior of any of the three pseudo-simplexes.

1.6.2 Comparison with weak mean field game solutions without common noise of [85]

Consider $A = [-1, 1]$, $T = 2$. With this choice of control actions and time horizon, the example we proposed matches the setting of Lacker’s “illuminating example” of [85, Section 3.3]. We show that there exists a coarse correlated solution of the MFG which is not a weak MFG solution without common noise as defined in Definition 3.1 therein. In particular, the most important feature is the fact that the recommendation Λ can not be expressed as a deterministic function of the flow of measures. On the other hand, we show that, under an additional assumption, weak MFG solutions induce coarse correlated solutions to the MFG.

To be consistent with the notation and the setup of Lacker’s paper, we use the notion of relaxed controls, which are used extensively in Section 1.5 (see, in particular, Section A for definitions, notation and some important properties). Let $(p_{i,j})_{i,j=1,2}$ be so that $p_{1,1} + p_{2,1}$ and $p_{1,2} + p_{2,2}$ are strictly positive. We introduce the relaxed

controls $\delta^+ = (\delta_t^+)_{t \in [0, T]}$ and $\delta^- = (\delta_t^-)_{t \in [0, T]}$, by setting

$$\begin{aligned}\delta_t^+(\omega_*; da) &= \delta_{\alpha_t^+(\omega_*)}(da) \equiv \delta_1(da), \quad \forall t \in [0, T], \omega_* \in \Omega^*, \\ \delta_t^-(\omega_*; da) &= \delta_{\alpha_t^-(\omega_*)}(da) \equiv \delta_{-1}(da), \quad \forall t \in [0, T], \omega_* \in \Omega^*.\end{aligned}$$

Consider the correlated measure flow (Λ, μ) defined by (1.6.5) and observe that the strategy $\lambda = (\lambda_t)_{t \in [0, T]}$ associated to the admissible recommendation Λ can be rewritten as a relaxed control as

$$\mathbf{r}_t(\omega; da) = \mathbf{r}_t(\omega_0, \omega_*; da) = \mathbb{1}_{\{\Lambda = \alpha^+\}}(\omega_0) \delta_t^+(da) + \mathbb{1}_{\{\Lambda = \alpha^-\}}(\omega_0) \delta_t^-(da). \quad (1.6.12)$$

We point out that \mathbf{r} does not depend on ω_* since δ^+ and δ^- do not depend on ω_* . Starting from (Λ, μ) , we define a random variable $\tilde{\mu}$ with values in $\mathcal{P}(\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d)$ by setting

$$\tilde{\mu}(\cdot) = \mathbb{P}((W, \mathbf{r}, X) \in \cdot \mid \mu). \quad (1.6.13)$$

We observe that $\sigma(\mu) = \sigma(\tilde{\mu})$: we have $\sigma(\tilde{\mu}) \subseteq \sigma(\mu)$ since, by definition of regular conditional probability, $\tilde{\mu}$ must be $\sigma(\mu)$ measurable; to get the opposite inclusion, for every $t \in [0, T]$, let $\tilde{\mu}_t^x$ be the push forward of $\tilde{\mu}$ through the map $\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d \ni (\omega, q, x) \mapsto x_t \in \mathbb{R}^d$. Then, by exploiting the consistency condition (1.3.13), we have

$$\tilde{\mu}_t^x(A) = \tilde{\mu}(\{x \in \mathcal{C}^d : x_t \in A\}) = \mathbb{P}(X_t \in A \mid \mu) = \mu_t(A),$$

for every $A \in \mathcal{B}_{\mathbb{R}^d}$, i.e. $\tilde{\mu}_t^x = \mu_t$ \mathbb{P} -a.s. for every $t \in [0, T]$. Let $(\mathcal{B}_{t, \mathcal{C}^d})_{t \in [0, T]}$ be the natural filtration of the identity process on \mathcal{C}^d , i.e. $\mathcal{B}_{t, \mathcal{C}^d} = \sigma(\mathcal{C}^d \ni x \mapsto x_s \in \mathbb{R}^d : 0 \leq s \leq t)$, and let $(\mathcal{F}_t^{\tilde{\mu}})_{t \in [0, T]}$ be the natural filtration of $\tilde{\mu}$, that is

$$\mathcal{F}_t^{\tilde{\mu}} = \sigma(\tilde{\mu}(C) : C \in \mathcal{B}_{t, \mathcal{C}^d} \otimes \mathcal{F}_t^{\mathcal{V}} \otimes \mathcal{B}_{t, \mathcal{C}^d}).$$

We observe that, for every $t \in (0, T]$, we have $\mathcal{F}_t^{\tilde{\mu}} = \sigma(\mu)$. To see this, observe that

$$\sigma(\mu) \supseteq \mathcal{F}_t^{\tilde{\mu}} \supseteq \sigma(\tilde{\mu}_s^x : s \leq t) = \sigma(\mu_s : s \leq t) = \sigma(\mu),$$

where the last equality holds for every $t > 0$, as can be verified by explicit calculations. Finally, for $t = 0$, we have $\mathcal{F}_0^{\tilde{\mu}} = \{\emptyset, \Omega^0\}$. Having established the relations between such σ -algebras, it is straightforward to verify that the tuple $((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), W, \tilde{\mu}, \mathbf{r}, X)$ satisfies properties (1-4) and (6) of [85, Definition 3.1]. Now, pick a probability measure \mathbb{P}^0 so that $p_{1,2} \cdot p_{2,1} > 0$ and $\bar{\mu}_T^1 > 0$, $\bar{\mu}_T^2 < 0$. Figure 1.1 shows that such a choice is possible (actually, there exist infinitely many measures \mathbb{P}^0 with the desired property). For such a choice of \mathbb{P}^0 , the relaxed control \mathbf{r} does not satisfy the optimality condition (5) of [85, Definition 3.1], since, as shown in [85, Section 3.3], every optimal control \mathbf{r}^* must be of the form $\mathbf{r}_t^*(da)(\omega) = \delta_{\alpha_t^*(\omega)}(da)$ for $Leb_{[0, T]}$ -a.e. t , with

$$\alpha_t^* = \text{sign} \left(\mathbb{E} \left[\tilde{\mu}_T^x \mid \mathcal{F}_t^{\tilde{\mu}} \right] \right).$$

Here, $\text{sign}(0) = 0$. Since $\mathcal{F}_0^{\tilde{\mu}}$ is trivial and $\mathcal{F}_t^{\tilde{\mu}} = \sigma(\mu)$ for $t > 0$, the optimal control α_t^* must be equal to

$$\alpha_t^*(\omega) = \alpha_t^*(\omega_0) = \begin{cases} -1 & \text{if } \bar{\mu}_T(\omega_0) < 0, \\ 1 & \text{if } \bar{\mu}_T(\omega_0) > 0, \end{cases} \quad 0 < t \leq T, \quad (1.6.14)$$

and equal to an arbitrary value at $t = 0$. In particular, observe that such a control is a deterministic function of the measure $\tilde{\mu}$. For every \mathbb{P}^0 so that $p_{1,2} + p_{2,1} > 0$, this is not the case of the correlated measure flow (Λ, μ) defined in (1.6.5), since Λ is not a deterministic function of μ .

The essential reason for the lack of optimality, in the sense of Lacker, of the relaxed control \mathbf{r} defined by (1.6.12) resides in the differences between allowed deviations: on the one hand, for weak mean field games solutions in the sense of [85], all adapted compatible controls $\mathbf{b} = (\mathbf{b}_t)_{t \in [0, T]}$ are allowed, where ‘‘compatible’’ means that $\sigma(\mathbf{b}_s : s \leq t)$ is conditionally independent of $\mathcal{F}_T^{\xi, \tilde{\mu}, W}$ given $\mathcal{F}_t^{\xi, \tilde{\mu}, W}$ for every t , which leads to a very rich class of controls. On the other hand, for coarse correlated solution of the MFG, only \mathbb{F}^* -progressively measurable strategies are allowed as deviations. Therefore, many more solutions exist.

Lastly, we show that under an additional assumption on the measurability of the random measure, it is indeed possible to define a mean field CCE starting from weak MFG solution without common noise. This property is not a-priori granted, due to the difference between the respective consistency conditions. Let $\tilde{\mu}$ be a weak MFG solution without common noise. Let μ_t be the push forward of $\tilde{\mu}$ through the map $\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d \ni (w, q, x) \mapsto x_t \in \mathbb{R}^d$. Define a random flow of measures by setting $\mu = (\mu_t)_{t \in [0, T]}$. Assume that the flow of measures μ carries the same information as the random measure $\tilde{\mu}$, i.e.

$$\sigma(\mu_s : 0 \leq s \leq t) = \mathcal{F}_t^{\tilde{\mu}}, \quad \forall t \in [0, T]. \quad (1.6.15)$$

If a weak MFG solution $\tilde{\mu}$ satisfies condition (1.6.15), then $\tilde{\mu}$ does induce a mean field CCE. Indeed, set $\rho = \mathbb{P} \circ \mu^{-1}$. By (1.6.15), we have $\mu_t = \mathbb{P}(X_t \in \cdot \mid \tilde{\mu}) = \mathbb{P}(X_t \in \cdot \mid \mu)$, i.e. the consistency condition (1.3.13) is satisfied. Moreover, the assumption on equality of the filtrations ensures that there exists a progressively measurable function $\varphi : [0, T] \times \mathcal{C}(\mathcal{P}^2) \rightarrow A$ so that

$$\alpha_t^* = \text{sign} \left(\mathbb{E} \left[\tilde{\mu}_T \mid \mathcal{F}_t^{\tilde{\mu}} \right] \right) = \varphi(t, \mu).$$

Then, we define $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ and $(\hat{\Lambda}, \hat{\mu})$ as

$$\begin{aligned} (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) &= (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho), \\ \hat{\mu} &= \text{Id} : (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho) \rightarrow (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho), \\ \hat{\Lambda} &: (\mathcal{C}(\mathcal{P}^2), \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}, \rho) \rightarrow (\mathbb{A}, \mathcal{B}_{\mathbb{A}}) \\ m &\mapsto \hat{\Lambda}(m) = (\varphi(t, m))_{t \in [0, T]}. \end{aligned} \quad (1.6.16)$$

By Lemma 1.7.1, the tuple $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \hat{\Lambda}, \hat{\mu})$ is a correlated measure flow. Let \hat{X} be the solution of (1.6.2) on the product probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ defined in (1.3.6). Since uniqueness in law holds by Theorem B.1, it follows that $(\hat{X}, \hat{\mu})$ has the same joint law as (X, μ) , which implies that the consistency condition (1.3.13) is satisfied. Since $\hat{\lambda}_t = \varphi(t, \hat{\mu})$, $(\hat{\Lambda}, \hat{\mu})$ satisfies optimality condition (1.3.12) as well and therefore it is a mean field CCE.

We observe that the additional assumption on the filtrations (1.6.15) is satisfied both by the weak MFG solution exhibited in [85, Proposition 3.7] and in our case, as

shown above. We point out that this CCE has been already considered: suppose that the flow of measures as law $\rho = p\delta_{\mu^+} + (1-p)\delta_{\mu^-}$, $p \in (0,1)$, for μ^+ and μ^- given by (1.6.4). Then, the correlated measure flow $(\hat{\Lambda}, \hat{\mu})$ corresponds to the probability measures \mathbb{P}^0 so that $p_{1,2} = p_{2,1} = 0$, $p_{1,1} = p$ and $p_{2,2} = 1-p$, which, as shown, always satisfy the condition (1.6.9). Roughly speaking, it corresponds to the case when $\hat{\Lambda} = \phi(\hat{\mu})$ \mathbb{P}^0 -a.s., for some deterministic measurable ϕ .

1.7 Auxiliary results

1.7.1 On admissible recommendations

Lemma 1.7.1. *Let $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ be a complete probability spaces and $(\Omega^*, \mathcal{F}^*, \mathbb{F}^*, \mathbb{P}^*)$ be a filtered probability space satisfying the usual assumptions. Fix a bounded A -valued process $(\lambda_t)_{t \in [0, T]}$ defined on the completion of the product space $(\Omega^0 \times \Omega^*, \mathcal{F}^{0-} \otimes \mathcal{F}^*, \mathbb{P}^0 \otimes \mathbb{P}^*)$. Assume that it is progressively measurable with respect to the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$, where \mathbb{F} is the $\mathbb{P}^0 \otimes \mathbb{P}^*$ -augmentation of the filtration $(\mathcal{F}^{0-} \otimes \mathcal{F}_t^*)_{t \in [0, T]}$. Define a function $\Lambda : \Omega^0 \rightarrow A$ by setting*

$$\Lambda(\omega_0) = \begin{cases} (\lambda_t(\omega_0, \cdot))_{t \in [0, T]} : [0, T] \times \Omega^* \rightarrow A & \omega_0 \in \Omega^0 \setminus N, \\ (t, \omega_*) \rightarrow \lambda_t(\omega_0, \omega_*), & \\ a_0 & \omega_0 \in N. \end{cases} \quad (1.7.1)$$

where $N \subset \Omega^0$ is a \mathbb{P}^0 -null set and a_0 is some point in A . The function Λ defined in (1.7.1) is an admissible recommendation.

Proof. Take any bounded $(\mathcal{F})_{t \in [0, T]}$ -progressively measurable process $(\lambda_t)_{t \in [0, T]}$ defined on the product space $(\Omega^0 \times \Omega^*, \mathcal{F}^{0-} \otimes \mathcal{F}^*, \mathbb{P}^0 \otimes \mathbb{P}^*)$ taking values in \mathbb{R} , and not necessarily in A .

Observe that it is always possible to define a function Λ from $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ to $(L^2([0, T] \times \Omega^*; \mathcal{R}^*, Leb_{[0, T]} \otimes \mathbb{P}^*), \mathcal{B}_{L^2})$ as in (1.7.1), where \mathcal{R}^* denotes the progressive σ -algebra associated to the filtration \mathbb{F}^* . Indeed, by construction of the filtration \mathbb{F} , since λ is \mathbb{F} -progressively measurable, there exists a set $N \subset \Omega^0$, $\mathbb{P}^0(N) = 0$, so that the section $(\lambda_t(\omega_0, \cdot))_{t \in [0, T]}$ is \mathcal{R}^* -measurable for every $\omega_0 \in \Omega^0 \setminus N$. Set $\Lambda(\omega_0) = (\lambda_t(\omega_0, \cdot))_{t \in [0, T]}$ for $\omega_0 \in \Omega^0 \setminus N$ and $\Lambda(\omega_0) \equiv a_0$ for $\omega_0 \in N$, where a_0 is any point in \mathbb{R} , which is exactly (1.7.1).

Let \mathcal{H} be the set of bounded progressively measurable processes λ so that the function Λ defined according to (1.7.1) is a $\mathcal{F}^{0-} \setminus \mathcal{B}_{L^2}$ measurable random variable. We show that \mathcal{H} is a monotone class which contains the set \mathcal{E} of progressively measurable processes $\lambda : [0, T] \times \Omega^0 \times \Omega^* \rightarrow \mathbb{R}$ of the form

$$\lambda_t = \sum_{i=1}^n \zeta^i \mathbf{1}_{[t_i, t_{i+1})}(t), \quad (1.7.2)$$

where $n \geq 1$, $t_i \in [0, T]$, $t_i < t_{i+1}$ for every $i = 1, \dots, n$, ζ^i are bounded \mathcal{F}_{t_i} -measurable random variables. Having established such properties, we apply monotone class theorem (as stated, e.g., in [101, Theorem II.3.1]) to conclude that \mathcal{H} contains the set of \mathbb{F} -progressively measurable bounded processes defined on the product space $\Omega^0 \times \Omega^*$.

To see that \mathcal{H} is a monotone class, observe that \mathcal{H} is clearly a vector space and contains all processes λ so that $\lambda_t \equiv c$ for every $t \in [0, T]$, for all $c \in \mathbb{R}$. Let $(\lambda^n)_{n \geq 1} \subseteq \mathcal{H}$, with $\lambda^n \uparrow \lambda$ as n goes to infinity, λ^n positive and bounded by the same constant $C \geq 0$ for every n . By monotone convergence, λ is bounded and $\mathcal{F}^{0-} \otimes \mathcal{R}^*$ -measurable as well, so that we can define Λ as in (1.7.1), as previously discussed. Let Λ^n be the L^2 -valued random variables defined starting from λ^n according to (1.7.1), which are $\mathcal{F}^{0-} \setminus \mathcal{B}_{L^2}$ measurable since λ^n belongs to \mathcal{H} for every $n \geq 1$, by assumption. Without loss of generality, we can suppose that the \mathbb{P}^0 -null set N appearing in the definition of Λ^n and Λ is the same for every $n \geq 1$. Notice that, for every $\omega_0 \in \Omega^0 \setminus N$, the sections $(\lambda_t^n(\omega_0, \cdot)_{t \in [0, T]}) \uparrow (\lambda_t(\omega_0, \cdot)_{t \in [0, T]})$ for every $(t, \omega_*) \in [0, T] \times \Omega^*$. Therefore, by monotone convergence, it holds

$$\begin{aligned} \|\Lambda^n(\omega_0) - \Lambda(\omega_0)\|_{L^2}^2 &= \|(\lambda_t^n(\omega_0, \cdot)_{t \in [0, T]}) - (\lambda_t(\omega_0, \cdot)_{t \in [0, T]})\|_{L^2}^2 \\ &= \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T |\lambda^n(t, \omega_0, \omega_*) - \lambda(t, \omega_0, \omega_*)|^2 dt \right] \rightarrow 0 \end{aligned} \quad (1.7.3)$$

for every $\omega_0 \in \Omega^0 \setminus N$, i.e. $\Lambda = \lim_{n \rightarrow \infty} \Lambda^n$ \mathbb{P}^0 -a.s., which implies that Λ is $\mathcal{F}^{0-} \setminus \mathcal{B}_{L^2}$ measurable, since the probability space is complete and $L^2([0, T] \times \Omega^*, \mathcal{R}^*, Leb_{[0, T]} \otimes \mathbb{P}^*)$ is a complete norm space. Finally, Λ is admissible, since the process λ obviously satisfies (1.3.8), choosing the same \mathbb{P}^0 -null set N used in the definition Λ .

To see that $\mathcal{E} \subseteq \mathcal{H}$, suppose first that λ is of the form

$$\lambda_t = \sum_{i=1}^n \mathbb{1}_{A_i}(\omega_0) \mathbb{1}_{B_i}(\omega_*) \mathbb{1}_{[t_i, t_{i+1})}(t),$$

where $n \geq 1$, $t_i \in [0, T]$, $t_i < t_{i+1}$ for every $i = 1, \dots, N$, $A_i \in \mathcal{F}^{0-}$ and $B_i \in \mathcal{F}_{t_i}^*$. We can regard each variable $\mathbb{1}_{B_i}(\omega_*) \mathbb{1}_{[t_i, t_{i+1})}(t)$ as a bounded progressively measurable process α^i . Therefore, (1.7.1) takes the following form:

$$\Lambda(\omega_0) = \begin{cases} \alpha^i & \omega_0 \in A_i, \quad i = 1, \dots, N, \\ 0 & \omega_0 \in (\cup_{i=1}^n A_i)^c \end{cases}$$

which shows that Λ is $\mathcal{F}^{0-} \setminus \mathcal{B}_{L^2}$ -measurable. By Dynkin Lemma, conclusion holds true for progressively measurable simple processes of the form

$$\lambda_t = \sum_{i=1}^n \mathbb{1}_{C_i}(\omega_0, \omega_*) \mathbb{1}_{[t_i, t_{i+1})}(t), \quad (1.7.4)$$

where $n \geq 1$, $t_i \in [0, T]$, $t_i < t_{i+1}$ for every $i \in \{1, \dots, n\}$, $C_i \in \mathcal{F}^{0-} \otimes \mathcal{F}_{t_i}^*$. Finally, let λ be of the form (1.7.2). Thanks to the boundedness assumption on ζ^i , we can find a sequence of simple processes $(\lambda^n)_{n \geq 1}$ of the form (1.7.4) so that $|\lambda^n| \leq |\lambda|$ and $\lambda_t^n(\omega_0, \omega_*) \rightarrow \lambda_t(\omega_0, \omega_*)$ pointwise for every (t, ω_0, ω_*) . Let Λ^n and Λ be defined according to (1.7.1) starting by the processes λ^n . Due to point 2.b) above, conclusion holds true for each Λ^n . Using dominated convergence, we can prove that (1.7.3) holds for \mathbb{P}^0 -a.e. $\omega_0 \in \Omega^0$, so that Λ is the a.s. pointwise limit of Λ^n , which implies that Λ is $\mathcal{F}^{0-} \setminus \mathcal{B}_{L^2}$ measurable. \square

Proposition 1.7.2. *Let $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ be a complete probability space.*

i) Let $\Lambda : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) \rightarrow (\mathbb{A}, \mathcal{B}_{\mathbb{A}})$ be an admissible recommendation. Let λ^1 and λ^2 be two \mathbb{F} -progressively measurable processes with values in A so that (1.3.8) holds. Then $\lambda^1 = \lambda^2$ $Leb_{[0,T]} \otimes \mathbb{P}$ -almost surely.

ii) Let $\Lambda, \Gamma : (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) \rightarrow (\mathbb{A}, \mathcal{B}_{\mathbb{A}})$ be admissible recommendations; let λ, γ be the strategies associated to Λ, Γ , according to (1.3.8). Suppose that $\lambda = \gamma$ $Leb_{[0,T]} \otimes \mathbb{P}$ -almost surely. Then, $\Lambda = \Gamma$ \mathbb{P}^0 -a.s.

Proof. As for point i), let $N^i, i = 1, 2$, be two \mathbb{P}^0 -null sets so that, for every $\omega_0 \in \Omega_0 \setminus N^i$ the sections $(\lambda_t^i(\omega_0, \cdot))_{t \in [0,T]}$ are \mathbb{F}^* -progressively measurable processes and equation (1.3.8) holds true. Without loss of generality, we can assume that $N^1 = N^2 = N$. Since for every $\omega_0 \in \Omega^0 \setminus N$ it holds $\|\Lambda - (\lambda_t^i(\omega_0, \cdot))_{t \in [0,T]}\|_{L^2} = 0, i = 1, 2$, we deduce that

$$\|(\lambda_t^1(\omega_0, \cdot))_{t \in [0,T]} - (\lambda_t^2(\omega_0, \cdot))_{t \in [0,T]}\|_{L^2}^2 = 0$$

for every $\omega_0 \in \Omega^0 \setminus N$. Therefore, by taking the integral with respect to \mathbb{P}^0 , we obtain

$$\begin{aligned} 0 &= \mathbb{E}^{\mathbb{P}^0} \left[\left\| (\lambda_t^1(\omega_0, \cdot))_{t \in [0,T]} - (\lambda_t^2(\omega_0, \cdot))_{t \in [0,T]} \right\|_{L^2}^2 \right] \\ &= \mathbb{E}^{\mathbb{P}^0} \left[\mathbb{E}^{\mathbb{P}^*} \left[\int_0^T |\lambda_s^1(\omega_0, \omega_*) - \lambda_s^2(\omega_0, \omega_*)|^2 ds \right] \right] = \mathbb{E} \left[\int_0^T |\lambda_s^1(\omega_0, \omega_*) - \lambda_s^2(\omega_0, \omega_*)|^2 ds \right] \end{aligned}$$

by Fubini's theorem. This is enough to conclude that $\lambda^1 = \lambda^2$ $Leb_{[0,T]} \otimes \mathbb{P}$ -a.s.

As for point ii), by the same line of reasoning, if $\lambda = \gamma$ $Leb_{[0,T]} \otimes \mathbb{P}$ -a.s., the sections $(\lambda_t(\omega_0, \cdot))_{t \in [0,T]}$ and $(\gamma_t(\omega_0, \cdot))_{t \in [0,T]}$ are $Leb_{[0,T]} \otimes \mathbb{P}^*$ -almost everywhere equal, which implies that

$$\begin{aligned} \|\Lambda(\omega_0) - \Gamma(\omega_0)\|_{L^2}^2 &= \left\| (\lambda_t(\omega_0, \cdot))_{t \in [0,T]} - (\gamma_t(\omega_0, \cdot))_{t \in [0,T]} \right\|_{L^2}^2 \\ &= \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T |\lambda_t(\omega_0, \omega_*) - \gamma_t(\omega_0, \omega_*)|^2 dt \right] = 0 \end{aligned}$$

\mathbb{P}^0 -a.s., so that $\Lambda = \Gamma$ \mathbb{P}^0 -a.s. □

1.7.2 Propagation of chaos

Here, we prove the propagation of chaos type result which is needed in the proof of Theorem 1.4.1. The probability spaces and the random variables we use here are defined in Section 1.4.1.

We work on the product probability space

$$(\Omega, \mathcal{F}, \mathbb{P}) = (\overline{\Omega}, \overline{\mathcal{F}}, \overline{\mathbb{P}}) \otimes (\Omega^1, \mathcal{F}^1, \mathbb{P}^1),$$

with $(\overline{\Omega}, \overline{\mathcal{F}}, \overline{\mathbb{P}})$ defined by (1.4.3) and (1.4.4) and $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$ by (1.2.1) or, equivalently, by (1.4.7). Consider the random measure flow μ defined by (1.4.5) and the recommendations $(\Lambda^i)_{i \geq 1}$ defined by (1.4.9), which we recall are conditionally i.i.d. given μ under $\overline{\mathbb{P}}$. We endow such a probability space with the filtration \mathbb{F} given by the \mathbb{P} -augmentation of the filtration generated by $\overline{\mathcal{F}}$, the initial data $(\xi^i)_{i \geq 1}$ and the

Brownian motions $(W^i)_{i \geq 1}$. We observe that for every $N \geq 2$, each $\beta \in \mathbb{A}_N$ is also \mathbb{F} -progressively measurable and, for every $i \geq 1$, each strategy λ^i associated to the admissible recommendation Λ^i is \mathbb{F} -progressively measurable as well.

Fix $N \geq 2$, $\beta \in \mathbb{A}_N$ and $1 \leq i \leq N$. Let $X = X[\Lambda^{N,-i}, \beta] = (X^j[\Lambda^{N,-i}, \beta])_{j=1}^N$ be the solution of

$$\begin{cases} dX_t^j = b(t, X_t^j, \mu_t^N, \lambda_t^j)dt + dW_t^j, & X_0^j = \xi^j, \quad j \neq i, \\ dX_t^i = b(t, X_t^i, \mu_t^N, \beta_t)dt + dW_t^i, & X_0^i = \xi^i. \end{cases}$$

The process $X[\Lambda^{N,-i}, \beta]$ is the state process of the N -players when every player $j \neq i$ follows the recommendation Λ^i and player i deviates by picking the strategy β , where μ_t^N denotes the empirical measure of the N -players' states at time t defined in (1.2.11). Let us introduce also the empirical measure of the processes $X = X[\Lambda^{N,-i}, \beta]$:

$$\mu^N[\Lambda^{N,-i}, \beta] = \frac{1}{N} \sum_{j=1}^N \delta_{X^j[\Lambda^{N,-i}, \beta]} \in \mathcal{C}(\mathcal{P}^2). \quad (1.7.5)$$

Let us denote $X = X[\Lambda] = X[\Lambda^{N,-i}, \Lambda^i]$ the state process of the N players when every player $i = 1, \dots, N$ follows the recommendation Λ^i . Then, let us consider the following auxiliary processes: let $(Z^j[\Lambda^{N,-i}, \beta])_{j=1}^N$ be the solution of

$$\begin{cases} dZ_t^j = b(t, Z_t^j, \mu_t, \lambda_t^j)dt + dW_t^j, & Z_0^j = \xi^j, \quad j \neq i, \\ dZ_t^i = b(t, Z_t^i, \mu_t, \beta_t)dt + dW_t^i, & Z_0^i = \xi^i \end{cases}$$

and $\nu^N[\Lambda^{-i}, \beta]$ be the empirical measure of the processes $Z[\Lambda^{-i}, \beta]$:

$$\nu^N[\Lambda^{-i}, \beta] = \frac{1}{N} \sum_{j=1}^N \delta_{Z^j[\Lambda^{N,-i}, \beta]} \in \mathcal{C}(\mathcal{P}^2).$$

Lemma 1.7.3. *Let β be either an open-loop strategy in \mathbb{A}_K for some $K \geq 2$, or be equal to λ^i , the strategy associated to the admissible recommendation Λ^i to player i . It holds:*

$$\sup_{t \in [0, T]} \mathbb{E} [\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_t^N[\Lambda^{-i}, \beta], \mu_t)] \xrightarrow{N \rightarrow \infty} 0, \quad (1.7.6)$$

$$\max_{1 \leq j \leq N} \mathbb{E} \left[\|X^{j, N}[\Lambda^{-i}, \beta] - Z^j[\Lambda^{-i}, \beta]\|_{\mathcal{C}^d}^2 \right] \xrightarrow{N \rightarrow \infty} 0, \quad (1.7.7)$$

$$\sup_{N \geq 2} \max_{1 \leq j \leq N} \mathbb{E} \left[\|X^{j, N}[\Lambda^{-i}, \beta]\|_{\mathcal{C}^d}^2 + \|Z^j[\Lambda^{-i}, \beta]\|_{\mathcal{C}^d}^2 \right] < \infty. \quad (1.7.8)$$

Proof. Because of the symmetry properties of the systems of SDEs, we can suppose $i = 1$. Throughout the proof, to make notation as simple as possible, we omit the dependence upon $[\Lambda^{-1}, \beta]$. For the same reason, define, for each $j \geq 1$, the following process γ^j :

$$\gamma_t^j = \begin{cases} \beta_t & j = 1, \\ \lambda_t^j & j \geq 2. \end{cases}$$

Obviously, in the case that β is λ^1 , we have $\gamma^j \equiv \lambda^j$ for every j . Moreover, let us introduce the following auxiliary processes: let $(Y^j)_{j \geq 1}$ be the solution of

$$dY_t^j = b(t, Y_t^j, \mu_t, \lambda_t^j)dt + dW_t^j, \quad Y_0^j = \xi^j.$$

Let η^N be the empirical measure of the processes Y^j :

$$\eta^N = \frac{1}{N} \sum_{j=1}^N \delta_{Y^j} \in \mathcal{C}(\mathcal{P}^2).$$

Denote by X^* the state process resulting from the coarse correlated solution of the MFG, i.e.

$$dX_t^* = b(t, X_t^*, \mu_t^*, \lambda_t^*) dt + dW_t^*, \quad X_0^* = \xi^*.$$

Since, for every $j \geq 1$, $(\xi^j, W^j, \mu, \lambda^j)$ are distributed as $(\xi^*, W^*, \mu^*, \lambda^*)$, by Theorem B.1 the processes $(Y^j)_{j \geq 1}$ are identically distributed copies of X^* ; moreover, the joint distribution of (Y^j, μ) under \mathbb{P} is the same of (X^*, μ^*) under \mathbb{P}^* , which, by marginalizing at every time $t \in [0, T]$, implies that Y^j satisfies the consistency condition (1.3.13) as well.

For every fixed $t \in [0, T]$, by the triangular inequality, it holds

$$\mathbb{E} [\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_t^N, \mu_t)] \leq C \mathbb{E} [\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_t^N, \nu_t^N) + \mathcal{W}_{2, \mathbb{R}^d}^2(\nu_t^N, \eta_t^N) + \mathcal{W}_{2, \mathbb{R}^d}^2(\eta_t^N, \mu_t)]. \quad (1.7.9)$$

We start from the third term in (1.7.9): let \mathbb{P}^m be a version of the regular conditional probability of \mathbb{P} given $\mu = m$, and denote by $\mathbb{E}^m[\cdot]$ the expectation with respect to the measure \mathbb{P}^m . By construction, the strategies $(\lambda^j)_{j \geq 1}$ associated to the admissible recommendations $(\Lambda^j)_{j \geq 1}$ are i.i.d. under \mathbb{P}^m , for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$. Since μ is independent of $(W^j)_{j \geq 1}$ and $(\xi^j)_{j \geq 1}$ under \mathbb{P} , the processes $(Y^j)_{j \geq 1}$ are independent under \mathbb{P}^m . Moreover, since $\mu_t(\cdot) = \mathbb{P}(Y_t^j \in \cdot \mid \mu)$ \mathbb{P} -a.s. for every $t \in [0, T]$ and $\mu_t = m_t$ \mathbb{P}^m -a.s. for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$, we have

$$m_t = \mathbb{P}^m \circ (Y_t^j)^{-1}, \quad \rho\text{-a.e.}, \quad \forall t \in [0, T], \quad (1.7.10)$$

for every $j \geq 1$. We can conclude that the processes $(Y_t^j)_{j \geq 1}$ are independent and identically distributed square integrable random variables with law m_t under \mathbb{P}^m , for every t and for ρ -a.e. m . Therefore, as ensured, e.g., by [38, (5.19)], it holds

$$\lim_{N \rightarrow \infty} \mathbb{E}^m [\mathcal{W}_{2, \mathbb{R}^d}^2(\eta_t^N, \mu_t)] = 0, \quad \rho\text{-a.e.}, \quad \forall t \in [0, T].$$

We observe that there exists a function $g : \mathcal{C}(\mathcal{P}^2) \rightarrow \mathbb{R}$ so that $g \in L^1(\rho)$ and $\mathbb{E}^m[\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_t^N, \mu_t)] \leq g(m)$, ρ -a.e., for every t : indeed, since, under \mathbb{P}^m , Y_t^j are i.i.d with law m_t and $\mu_t = m_t$ a.s., we have

$$\begin{aligned} \mathbb{E}^m [\mathcal{W}_{2, \mathbb{R}^d}^2(\eta_t^N, \mu_t)] &\leq 2 \mathbb{E}^m [\mathcal{W}_{2, \mathbb{R}^d}^2(\eta_t^N, \delta_0) + \mathcal{W}_{2, \mathbb{R}^d}^2(\delta_0, \mu_t)] \\ &\leq 2 \left(\frac{1}{N} \sum_{k=1}^N \mathbb{E}^m [|Y_t^k|^2] + \int_{\mathbb{R}^d} |y|^2 m_t(dy) \right) \leq 2 \left(\frac{1}{N} \sum_{k=1}^N \mathbb{E}^m [|Y_t^1|^2] + \mathbb{E}^m [|Y_t^1|^2] \right) \\ &\leq 4 \mathbb{E}^m [|Y_t^1|^2] \leq 4 \mathbb{E}^m [\|Y^1\|_{\mathcal{C}^d}^2]. \end{aligned}$$

The function $g(m) = \mathbb{E}^m [\|Y^1\|_{\mathcal{C}^d}]$ belongs to $L^1(\rho)$, since

$$\int_{\mathcal{C}(\mathbb{P}^2)} g(m) \rho(dm) = \mathbb{E} \left[\mathbb{E} \left[\|Y^1\|_{\mathcal{C}^d}^2 \mid \mu \right] \right] = \mathbb{E} \left[\|Y^1\|_{\mathcal{C}^d}^2 \right] < \infty. \quad (1.7.11)$$

Therefore, by dominated convergence theorem, we have

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2 (\eta_t^N, \mu_t) \right] = 0 \quad (1.7.12)$$

for every $t \in [0, T]$. The convergence in (1.7.12) is actually uniform in time. Indeed, fix $t, s \in [0, T]$: then

$$\begin{aligned} & \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2 (\eta_t^N, \mu_t) - \mathcal{W}_{2, \mathbb{R}^d}^2 (\eta_s^N, \mu_s) \right] \\ &= \mathbb{E} \left[(\mathcal{W}_{2, \mathbb{R}^d} (\eta_t^N, \mu_t) - \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s)) (\mathcal{W}_{2, \mathbb{R}^d} (\eta_t^N, \mu_t) + \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s)) \right] \\ &\leq C \mathbb{E} \left[\|Y^1\|_{\mathcal{C}^d}^2 \right]^{\frac{1}{2}} \mathbb{E} \left[|\mathcal{W}_{2, \mathbb{R}^d} (\eta_t^N, \mu_t) - \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s)|^2 \right]^{\frac{1}{2}}, \end{aligned}$$

where we used Cauchy-Schwartz inequality together with the uniform in time bound given by (1.7.11). By triangulating with η_s^N and μ_s , we get

$$\begin{aligned} & \mathbb{E} \left[|\mathcal{W}_{2, \mathbb{R}^d} (\eta_t^N, \mu_t) - \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s)|^2 \right] \leq \mathbb{E} \left[|\mathcal{W}_{2, \mathbb{R}^d} (\eta_t^N, \eta_s^N) + \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s) \right. \\ & \quad \left. + \mathcal{W}_{2, \mathbb{R}^d} (\mu_s, \mu_t) - \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s)|^2 \right] \leq \mathbb{E} \left[|\mathcal{W}_{2, \mathbb{R}^d} (\eta_t^N, \eta_s^N) - \mathcal{W}_{2, \mathbb{R}^d} (\eta_s^N, \mu_s)|^2 \right] \\ &\leq C \left(\mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2 (\eta_t^N, \eta_s^N) \right] + \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2 (\mu_t, \mu_s) \right] \right) \\ &\leq C \left(\mathbb{E} \left[\frac{1}{N} \sum_{k=1}^N |Z_t^k - Z_s^k|^2 \right] + \mathbb{E} \left[\mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2 (\mu_t, \mu_s) \mid \mu \right] \right] \right) \\ &\leq C \left(\mathbb{E} \left[\frac{1}{N} \sum_{k=1}^N |Z_t^k - Z_s^k|^2 \right] + \mathbb{E} \left[|Y_t^1 - Y_s^1|^2 \right] \right), \end{aligned}$$

where in the last inequality we used (1.7.10) and tower property. By using Lipschitz continuity of b , the triangular inequality and $\mathbb{E}[\|Y^1\|_{\mathcal{C}^d}] < \infty$, it is straightforward to see that $\mathbb{E}[\|Z^k\|_{\mathcal{C}^d}] \leq C$ for every $k \geq 1$, for some positive constant C independent of k . By the same arguments, we have

$$\begin{aligned} \mathbb{E} \left[|Z_t^k - Z_s^k|^2 \right] &\leq C \mathbb{E} \left[\int_s^t |b(u, Z_u^k, \mu_u, \gamma_u^k)|^2 du \right] \\ &\leq C \mathbb{E} \left[\int_s^t \left(1 + |Z_u^k|^2 + \int_{\mathbb{R}^d} |y|^2 \mu_u(dy) + |\gamma_u^k|^2 \right) du \right] \\ &\leq C \mathbb{E} \left[\int_s^t \left(1 + \|Z^k\|_{\mathcal{C}^d}^2 + \|Y^1\|_{\mathcal{C}^d}^2 + |\gamma_u^k|^2 \right) du \right] \leq C |t - s|, \end{aligned}$$

where the constant C depends upon T , b , $\mathbb{E}[\|Y^1\|_{\mathcal{C}^d}^2] < \infty$ and $\text{diam}(A)$, which is a finite quantity since the set A of actions is compact by Assumptions **A**. Analogously holds for $\mathbb{E}[|Y_t^1 - Y_s^1|]$, which implies that

$$|\mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2 (\eta_t^N, \mu_t) - \mathcal{W}_{2, \mathbb{R}^d}^2 (\eta_s^N, \mu_s) \right]| \leq C |t - s|^{\frac{1}{2}}. \quad (1.7.13)$$

This is enough to conclude, by Arzelà-Ascoli theorem, that the convergence in (1.7.12) is uniform in time.

Remind from Section 1.1 that $\|x\|_{t,C^d} = \sup_{s \in [0,t]} |x_s|$, $t \in [0, T]$. To handle the second term in (1.7.9), we use the coupling of ν^N and η^N given by $\frac{1}{N} \sum_{k=1}^N \delta_{(Z^k, Y^k)}$, together with the Lipschitz continuity of b :

$$\begin{aligned} \mathbb{E} \left[\|Z^k - Y^k\|_{t,C^d}^2 \right] &= \mathbb{E} \left[\sup_{0 \leq s \leq t} \left(\int_0^s (b(u, Z_u^k, \mu_u, \gamma_u^k) - b(u, Y_u^k, \mu_u, \lambda_u^k)) du \right)^2 \right] \\ &\leq C \left(\int_0^t \mathbb{E} \left[\sup_{0 \leq u \leq s} |Z_u^k - Y_u^k|^2 \right] ds + \int_0^t \mathbb{E} \left[|\lambda_s^k - \gamma_s^k|^2 \right] ds \right). \end{aligned}$$

By definition of $(\gamma^k)_{k \geq 1}$, we have

$$\int_0^T \mathbb{E} \left[|\lambda_u^k - \gamma_u^k|^2 \right] du = \begin{cases} \int_0^T \mathbb{E} \left[|\lambda_u^1 - \beta_u|^2 \right] du & k = 1, \\ 0 & k \geq 2. \end{cases}$$

Therefore, by Gronwall's lemma, we sum over $k = 1, \dots, N$ to obtain the estimate

$$\begin{aligned} \sup_{t \in [0, T]} \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\eta_t^N, \nu_t^N) \right] &\leq \mathbb{E} \left[\mathcal{W}_{2, C^d}^2(\eta^N, \nu^N) \right] \leq \frac{C}{N} \sum_{k=1}^N \mathbb{E} \left[\|Y^k - Z^k\|_{C^d}^2 \right] \\ &\leq \frac{C}{N} \sum_{j=1}^N \int_0^T \mathbb{E} \left[|\lambda_u^j - \gamma_u^j|^2 \right] du = \frac{C}{N} \int_0^T \mathbb{E} \left[|\lambda_u^1 - \beta_u|^2 \right] du \leq \frac{C}{N} \xrightarrow{N \rightarrow \infty} 0, \end{aligned} \tag{1.7.14}$$

where the constant C depends only upon T , b and $\text{diam}(A)$.

Finally, for the first term of (1.7.9), we use the coupling of μ_t^N and ν_t^N given by $\frac{1}{N} \sum_{k=1}^N \delta_{(X_t^{k,N}, Z_t^k)}$, together with the Lipschitz continuity of b :

$$\begin{aligned} \mathbb{E} \left[\|X^{k,N} - Z^k\|_{t,C^d}^2 \right] &\leq C \int_0^t \left(\mathbb{E} \left[\sup_{0 \leq u \leq s} |X_u^{k,N} - Z_u^k|^2 \right] + \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_s^N, \mu_s) \right] \right) ds \\ &\leq C \int_0^t \left(\mathbb{E} \left[\sup_{0 \leq u \leq s} |X_u^{k,N} - Z_u^k|^2 \right] + \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_s^N, \nu_s^N) \right] + \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\nu_s^N, \mu_s) \right] \right) ds \\ &\leq C \int_0^t \left(\mathbb{E} \left[\sup_{0 \leq u \leq s} |X_u^{k,N} - Z_u^k|^2 \right] + \frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[|X_s^{j,N} - Z_s^j|^2 \right] + \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\nu_s^N, \mu_s) \right] \right) ds \\ &\leq C \left(\int_0^t \max_{k=1, \dots, N} \mathbb{E} \left[\sup_{0 \leq u \leq s} |X_u^{k,N} - Z_u^k|^2 \right] ds + \sup_{s \in [0, t]} \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\nu_s^N, \mu_s) \right] \right). \end{aligned}$$

By taking the maximum over $k = 1, \dots, N$ on the left-hand side and applying Gronwall's lemma, we get

$$\max_{k=1, \dots, N} \mathbb{E} \left[\|X^{k,N} - Z^k\|_{C^d}^2 \right] \leq C \sup_{t \in [0, T]} \mathbb{E} \left[\mathcal{W}_{2, \mathbb{R}^d}^2(\nu_t^N, \mu_t) \right] \rightarrow 0,$$

by (1.7.12), (1.7.13) and (1.7.14), which proves (1.7.7). Coming back to (1.7.6), we have

$$\begin{aligned} \sup_{t \in [0, T]} \mathbb{E} [\mathcal{W}_{2, \mathbb{R}^d}^2(\mu_t^N, \nu_t^N)] &\leq \sup_{t \in [0, T]} \frac{1}{N} \sum_{k=1}^N \mathbb{E} \left[\left| X_t^{k, N} - Z_t^k \right|^2 \right] \\ &\leq \max_{k=1, \dots, N} \mathbb{E} \left[\left\| X^{k, N} - Z^k \right\|_{\mathcal{C}^d}^2 \right] \rightarrow 0. \end{aligned}$$

This estimate together with estimates (1.7.12), (1.7.13) and (1.7.14) implies (1.7.9) and therefore (1.7.6). Finally, (1.7.8) follows from the above calculations. \square

1.7.3 Further technical lemmata for the existence of mean field CCEs

In this section, we state and prove some auxiliary results that were used in Section 1.5 to prove the existence of a mean field CCE. In particular, Lemmata 1.7.4 and 1.7.5 provide the technical instruments we used in Proposition 1.5.3 to show that, for every deviating strategy $\beta \in \mathbb{A}$ and random flow of measures μ , we can represent the joint law of μ , β and deviating player's state process in terms of a strategy for player B in the zero-sum game 1.5.3 and the law of μ . Lemmata 1.7.6 and 1.7.7 were needed in the proof of Theorem 1.5.1 in order to define a mean field CCE starting from an optimal strategy for player A in the zero-sum game defined in Definition 1.5.3.

Consider any tuple $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$, composed of a filtered probability space satisfying usual assumptions, a d -dimensional \mathbb{F} -Brownian motion, an \mathbb{R}^d -valued \mathcal{F}_0 -measurable random variable, an \mathcal{F}_0 -measurable random continuous flow of measures in $\mathcal{P}^2(\mathbb{R}^d)$ and an \mathbb{F} -progressively measurable $\mathcal{P}(A)$ -valued process. Assume that μ , W and ξ are mutually independent. Let us consider the following equations:

$$dX_t = \int_A b(t, X_t, \mu_t, a) \mathfrak{r}_t(da) dt + dW_t, \quad X_0 = \xi, \quad (1.7.15)$$

$$dX_t^m = \int_A b(t, X_t^m, m_t, a) \mathfrak{r}_t(da) dt + dW_t, \quad X_0 = \xi, \quad (1.7.16)$$

where m is a point of $\mathcal{C}(\mathcal{P}^2)$. In order to stress the dependence upon the deterministic flow of measures m , we write X^m for the solution of (1.7.16).

By Assumptions **A**, on any such tuple \mathfrak{U} there exists a solution to equation (1.7.15) and pathwise uniqueness holds. If needed, we can suppose that the filtration \mathbb{F} on $(\Omega, \mathcal{F}, \mathbb{P})$ is the \mathbb{P} -augmentation of the filtration $\mathbb{F}^{\xi, W, \mu, \mathfrak{r}}$, given by

$$\mathcal{F}_t^{\xi, W, \mu, \mathfrak{r}} = \sigma(\xi) \vee \sigma(\mu) \vee \sigma(W_s : s \leq t) \vee \sigma(\mathfrak{r}(C) : C \in \mathcal{B}_{[0, t] \times A}). \quad (1.7.17)$$

By Theorem B.1, uniqueness in law holds. Analogous reasoning holds for equation (1.7.16) as well, for every $m \in \mathcal{C}(\mathcal{P}^2)$.

Lemma 1.7.4. *Let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ be as above, let $\Theta \in \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{V})$ be the joint law of ξ , W and \mathfrak{r} . Let us define the map*

$$\begin{aligned} \mathcal{I}_\Theta : \mathcal{C}(\mathcal{P}^2) &\longrightarrow \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}^d \times \mathcal{V}) \\ m &\longmapsto \mathcal{I}_\Theta(m) = \mathbb{P} \circ (\xi, W, X^m, \mathfrak{r})^{-1}, \end{aligned} \quad (1.7.18)$$

where X^m is the solution to equation (1.7.16).

(i) The map \mathcal{I}_Θ is continuous, in the sense that

$$\sup_{t \in [0, T]} \mathcal{W}_{2, \mathbb{R}^d}(m_t^n, m_t) \rightarrow 0 \text{ as } n \rightarrow \infty \Rightarrow \mathcal{I}_\Theta(m^n) \xrightarrow{n \rightarrow \infty} \mathcal{I}_\Theta(m) \text{ in } \mathcal{W}_{2, \mathbb{R}^d \times \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d}.$$

(ii) The map \mathcal{I}_Θ induces a stochastic kernel Σ from $\mathcal{C}(\mathcal{P}^2)$ to $\mathcal{C}^d \times \mathcal{V}$, by setting

$$\Sigma(B, m) = \mathbb{P}((X^m, \mathbf{r}) \in B) = \mathcal{I}_\Theta(m)(\mathbb{R}^d \times \mathcal{C}^d \times B) \quad \forall m \in \mathcal{C}(\mathcal{P}^2), B \in \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{V}}.$$

Σ is a strategy for player B , as described in Definition 1.5.2.

Proof. Note that, by Theorem B.1, $\mathcal{I}_\Theta(m)$ is the unique weak solution of (1.7.16) when the joint law of ξ , W and \mathbf{r} is given by Θ and b is evaluated at $m \in \mathcal{C}(\mathcal{P}^2)$. Let $\Sigma \in \mathcal{Q}$, let $(m^n)_{n \geq 1} \subset \mathcal{C}(\mathcal{P}^2)$ so that $m^n \rightarrow m$. For every $n \geq 1$, denote by X and X^n the solution to equation (1.7.16) when b is evaluated at m and m^n , respectively. For every $2 \leq p \leq \bar{p}$, by Lipschitz continuity of b , we have:

$$\mathbb{E} \left[\sup_{0 \leq s \leq T} |X_s^n - X_s|^p \right] \leq C \sup_{t \in [0, T]} \mathcal{W}_{p, \mathbb{R}^d}^2(m_t^n, m_t), \quad (1.7.19)$$

$$\sup_{n \geq 1} \mathbb{E} \left[\sup_{t \in [0, T]} |X_t^n|^p \right] \leq C \left(1 + \sup_{n \geq 1} \sup_{t \in [0, T]} \left(\int_{\mathbb{R}^d} |y|^2 m_t^n(dy) \right)^{\frac{p}{2}} \right). \quad (1.7.20)$$

Therefore, for $p = 2$, we have that $m^n \rightarrow m$ in $\mathcal{C}(\mathcal{P}^2)$ implies that $\|X^n - X\|_{\mathcal{C}^d}^2 \rightarrow 0$ in expectation, which in turns implies that $(\xi, W, \mathbf{r}, X^n) \rightarrow (\xi, W, \mathbf{r}, X)$ in distribution. In order to have convergence in 2-Wasserstein metrics, it is enough to check uniform integrability, according to (1.5.19). Since $(\mathcal{I}_\Theta(m^n))_n$ have the same marginals on $\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{V}$, we just need to check (1.5.19) for the laws of $(X^n)_n$: for every $n \geq 1$, $r > 0$, we have

$$\begin{aligned} \mathbb{E} \left[\|X^n\|_{\mathcal{C}^d}^2 \mathbf{1}_{\{\|X^n\|_{\mathcal{C}^d}^2 > r\}} \right] &\leq (\mathbb{E} [\|X^n\|_{\mathcal{C}^d}^4])^{\frac{1}{2}} \left(\mathbb{E} [\mathbf{1}_{\{\|X^n\|_{\mathcal{C}^d}^2 > r\}}] \right)^{\frac{1}{2}} \\ &\leq (\mathbb{E} [\|X^n\|_{\mathcal{C}^d}^4])^{\frac{1}{2}} \mathbb{E} [\|X^n\|_{\mathcal{C}^d}^2]^{\frac{1}{2}} r^{-\frac{1}{2}} \leq Cr^{-\frac{1}{2}} \end{aligned}$$

by using Cauchy-Schwartz inequality, Markov inequality, (1.7.19) and (1.7.20). By taking the limit as $r \rightarrow \infty$, we get condition (1.5.19) satisfied and so point (i) is proved.

As for point (ii), let $\pi : \mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}^d \times \mathcal{V} \rightarrow \mathcal{C}^d \times \mathcal{V}$ be the projection on the last two components. Note that $(\mathcal{I}_\Theta \circ \pi^{-1})(m) = \mathbb{P} \circ (X^m, \mathbf{r})^{-1}$, which shares the same continuity properties of the map \mathcal{I}_Θ . Therefore, in particular, it is Borel measurable, where $\mathcal{P}(\mathcal{C}^d \times \mathcal{V})$ is endowed with the usual Borel σ -algebra associated with the topology of weak convergence. Then, the thesis follows from the fact that, for a Polish space E , the usual Borel σ -field on $\mathcal{P}(E)$ coincide with the σ -field generated by the maps $\mathcal{P}(E) \ni m \mapsto m(S)$, with $S \in \mathcal{B}_E$, (see, e.g., [17, Corollary 7.29.1]). \square

Lemma 1.7.5. *Let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathbf{r})$ be a tuple so that (ξ, W, \mathbf{r}) and μ are independent. Denote by $\rho \in \mathcal{P}(\mathcal{C}(\mathcal{P}^2))$ the law of μ under \mathbb{P} . Suppose without loss of generality that \mathbb{F} is the \mathbb{P} -augmentation of the filtration $(\mathcal{F}_t^{\xi, W, \mu, \mathbf{r}})_t$ defined by*

(1.7.17). Let X be the unique solution of (1.7.15) on the tuple \mathfrak{U} . Then, the following decomposition of measure holds

$$\mathbb{P}((X, \mathfrak{r}, \mu) \in B \times S) = \int_S \Sigma(B, m) \rho(dm), \quad \forall B \in \mathcal{B}_{\mathcal{C}^d \times \mathcal{V}}, S \in \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}.$$

In particular, $\Sigma(B, m) = \mathbb{P}((X, \mathfrak{r}) \in B \mid \mu = m) = \mathbb{P}((X^m, \mathfrak{r}) \in B)$ for every $B \in \mathcal{B}_{\mathcal{C}^d \times \mathcal{V}}$, ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$.

Proof. Let $\mathbb{P}(\cdot \mid \mu)$ denote the regular conditional probability of \mathbb{P} given μ . Set $\mathbb{P}^m(\cdot) = \mathbb{P}(\cdot \mid \mu = m)$. Since (ξ, W, \mathfrak{r}) and μ are independent by assumption, we have that $\mathbb{P}^m \circ (\xi, W, \mathfrak{r})^{-1} = \mathbb{P} \circ (\xi, W, \mathfrak{r})^{-1}$ for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$. Therefore, it is enough to prove that X is a solution to (1.7.16) on the tuple $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}^m), \xi, W, \mathfrak{r})$ for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$. Then, since uniqueness in law holds for (1.7.16), we deduce that $\mathbb{P}^m \circ (\xi, W, \mathfrak{r}, X)^{-1} = \mathcal{I}_\Theta(m)$ ρ -a.s. Observe that, since the joint law of (ξ, W, \mathfrak{r}) is the same under \mathbb{P} and \mathbb{P}^m for ρ -a.e. m , W is a natural Brownian motion under \mathbb{P}^m as well. By definition of the filtration \mathbb{F} , it can be easily verified that

$$\mathbb{E}^{\mathbb{P}}[\mathbf{1}_A(W_t - W_s)g \mid \mu] = 0 \quad \mathbb{P}\text{-a.s.}$$

for every $0 \leq s < t \leq T$, $A \in \mathcal{B}_{\mathbb{R}^d}$, g bounded and \mathcal{F}_s -measurable. This implies that

$$E^{\mathbb{P}^m}[\mathbf{1}_A(W_t - W_s)g] = \mathbb{E}^{\mathbb{P}}[\mathbf{1}_A(W_t - W_s)g \mid \mu = m] = 0$$

ρ -a.s., for every g bounded and \mathcal{F}_s -measurable. By working with a countable measure determining class of sets, which is possible since the σ -algebra $\mathcal{F}_t^{\xi, W, \mu, \mathfrak{r}}$ is countably generated for every $t \in [0, T]$, the equality holds for every g bounded and \mathcal{F}_s -measurable, for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$, which in turn implies that W remains an \mathbb{F} -Brownian motion under \mathbb{P}^m as well. Under \mathbb{P}^m one has

$$\mathbb{P}^m \left(\int_A b(t, x, \mu_t, a) \mathfrak{r}_t(da) = \int_A b(t, x, m_t, a) \mathfrak{r}_t(da) \quad \forall x \in \mathbb{R}^d \right) = 1$$

for $Leb_{[0, T]}$ -a.e. $t \in [0, T]$. Therefore X solves (1.7.16) for b evaluated at $m \in \mathcal{C}(\mathcal{P}^2)$. The thesis follows from marginalizing as in the proof of point (ii) in 1.7.4. \square

We turn our attention to the mimicking result, needed in the proof of the existence result in Theorem 1.5.1:

Lemma 1.7.6. *Let $\Gamma \in \mathcal{K}$. There exists a measure $\hat{\Gamma} \in \mathcal{K}$ so that the following holds:*

- *The marginal distributions of Γ and $\hat{\Gamma}$ on $\mathcal{C}(\mathcal{P}^2)$ are the same: $\Gamma(\mathcal{C}^d \times \mathcal{V} \times \cdot) = \hat{\Gamma}(\mathcal{C}^d \times \mathcal{V} \times \cdot)$.*
- *Let (X, \mathfrak{r}, μ) be such that $\hat{\Gamma} = \mathbb{P} \circ (X, \mathfrak{r}, \mu)^{-1}$. Then \mathfrak{r} is of the form $\mathfrak{r}_t = \hat{q}_t(X_t, \mu)$, where $\hat{q} : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \rightarrow \mathcal{P}(A)$ is a measurable function.*
- *For every $\Sigma \in \mathcal{Q}$, it holds*

$$\mathfrak{p}(\Gamma, \Sigma) = \mathfrak{p}(\hat{\Gamma}, \Sigma).$$

Proof. In the following, for a metric space (E, d_E) , $\phi : E \rightarrow \mathbb{R}$ continuous and bounded and $m \in \mathcal{P}(E)$, we set $\langle \phi, m \rangle = \int_E \phi(e)m(de)$.

Let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ be as in Definition 1.5.1, so that $\Gamma = \mathbb{P}_\circ(X, \mathfrak{r}, \mu)^{-1}$. As ensured by [87, Lemma C.2], we have that, by choosing $Y_t = (X_t, \mu)$ taking values in $\mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$ as conditioning process, there exists a jointly measurable function $\hat{q} : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \rightarrow \mathcal{P}(A)$ so that, for every $\phi : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \times A \rightarrow \mathbb{R}$ bounded and measurable it holds

$$\int_A \phi(X_t, \mu, a) \hat{q}_t(X_t, \mu)(da) = \mathbb{E} \left[\int_A \phi(X_t, \mu, a) \mathfrak{r}_t(da) | X_t, \mu \right] \quad \mathbb{P}\text{-a.s., a.e. } t \in [0, T], \quad (1.7.21)$$

which we abbreviate as

$$\hat{q}_t(X_t, \mu)(da) = \mathbb{E} [\mathfrak{r}_t(da) | X_t, \mu] \quad \mathbb{P}\text{-a.s., a.e. } t \in [0, T].$$

Next, we manipulate the term of the functional \mathfrak{p} in (1.5.7) which depends only upon Γ :

$$\begin{aligned} \int \mathfrak{F}(y, q, m) \Gamma(dy, dm, dq) &= \mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \mathfrak{r}_t(da) dt + g(X_T, \mu_T) \right] \\ &= \int_0^T \mathbb{E} \left[\mathbb{E} \left[\int_A f(t, X_t, \mu_t, a) \mathfrak{r}_t(da) | X_t, \mu \right] \right] dt + \mathbb{E} [g(X_T, \mu_T)] \\ &= \int_0^T \mathbb{E} \left[\int_A f(t, X_t, \mu_t, a) \hat{q}_t(X_t, \mu)(da) \right] dt + \mathbb{E} [g(X_T, \mu_T)] \quad (1.7.22) \\ &= \int_0^T \mathbb{E} \left[\mathbb{E} \left[\int_A f(t, X_t, \mu_t, a) \hat{q}_t(X_t, \mu)(da) | \mu \right] \right] dt + \mathbb{E} \left[\mathbb{E} [g(X_T, \mu_T) | \mu] \right] \\ &= \int_0^T \mathbb{E} \left[\left\langle \int_A f(t, \cdot, \mu_t, a) \hat{q}_t(\cdot, \mu)(da), \mu_t \right\rangle \right] dt + \mathbb{E} [\langle g(\cdot, \mu_T), \mu_T \rangle]. \end{aligned}$$

Second equality holds by Fubini's theorem and tower property of conditional expectation, third equality holds by definition of the control (1.7.21), third and fourth equalities hold by tower property again, and fifth equality holds since, by consistency condition, $\mu_t(\cdot) = \mathbb{P}(X_t \in \cdot | \mu)$.

We observe that, by choosing $\phi(t, x, m, a) = b(t, x, m_t, a)$ in (1.7.21), we have

$$\int_A b(t, X_t, \mu_t, a) \hat{q}_t(X_t, \mu)(da) = \mathbb{E} \left[\int_A b(t, X_t, \mu_t, a) \mathfrak{r}_t(da) | X_t, \mu \right] \quad \mathbb{P}\text{-a.s., a.e. } t \in [0, T].$$

This is enough to apply [24, Theorem 3.6]: indeed, in its terminology, we can take $\mathcal{E} = \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$, $\Phi : \mathcal{E} \times \mathcal{C}_0^d \rightarrow \mathcal{C}([0, T]; \mathcal{E})$ defined by $\Phi_t(x, m, y) = (x_t + y, m) \in \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$, where $\mathcal{C}_0^d = \{y \in \mathcal{C}^d : x_0 = 0\}$. Set $Z_t = \Phi(X_t - X_0, X_0, \mu) = (X_t, \mu)$, where we note that the second component of Z is constant in time as it is equal to the whole flow $\mu = (\mu_s)_{s \in [0, T]}$. Then, such a result ensures that there exists a probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{F}}, \hat{\mathbb{P}})$, with $\hat{\Omega}$ Polish and $\hat{\mathcal{F}}$ its corresponding Borel σ -algebra, supporting an $\hat{\mathbb{F}}$ -Brownian motion \hat{W} , a continuous \mathcal{E} -valued process \hat{Z} so that there exists an $\hat{\mathbb{F}}$ -adapted process \hat{X} that satisfies

$$\hat{X}_t = \hat{X}_0 + \int_0^t \int_A b(s, \hat{X}_s, \hat{\mu}_s, a) \hat{q}_s(\hat{X}_s, \hat{\mu}_s)(da) ds + \hat{W}_t, \quad \hat{Z} = \Phi(\hat{X}_t - \hat{X}_0, \hat{X}_0, \hat{\mu})$$

so that for every $t \in [0, T]$ it holds $\mathbb{P} \circ Z_t^{-1} = \hat{\mathbb{P}} \circ \hat{Z}_t^{-1}$. This implies both that $\hat{\mu}$ and μ have the same law ρ and that the consistency condition is satisfied, since $\hat{\mathbb{P}} \circ (\hat{X}_t, \hat{\mu})^{-1} = \mathbb{P} \circ (X_t, \mu)^{-1} = m_t(dx)\rho(dm)$. Finally, since Z is $\hat{\mathbb{F}}$ -adapted, we deduce that \hat{X}_0 and $\hat{\mu}$ are \mathcal{F}_0 -measurable and therefore \hat{W} , \hat{X}_0 and $\hat{\mu}$ are mutually independent.

Set $\hat{\Gamma} = \hat{\mathbb{P}} \circ (\hat{X}, \hat{\mu})^{-1}$. Since the last term in the chain of equalities (1.7.22) depends only upon μ and $\hat{\mu}$ share the same law, we can exploit the fact that $\hat{\mu}$ and \hat{X} satisfy the consistency condition as well to get

$$\begin{aligned}
& \int \mathfrak{F}(y, q, m) \Gamma(dy, dm, dq) \\
&= \int_0^T \mathbb{E} \left[\left\langle \int_A f(t, \cdot, \mu_t, a) \hat{q}_t(\cdot, \mu)(da), \mu_t \right\rangle \right] dt + \mathbb{E} [\langle g(\cdot, \mu_T), \mu_T \rangle] \\
&= \int_0^T \mathbb{E}^{\hat{\mathbb{P}}} \left[\left\langle \int_A f(t, \cdot, \hat{\mu}_t, a) \hat{q}_t(t, \cdot, \hat{\mu}_t, a)(da), \hat{\mu}_t \right\rangle \right] dt + \mathbb{E}^{\hat{\mathbb{P}}} [\langle g(\cdot, \hat{\mu}_T), \hat{\mu}_T \rangle] \\
&= \int_0^T \mathbb{E}^{\hat{\mathbb{P}}} \left[\mathbb{E}^{\hat{\mathbb{P}}} \left[\int_A f(t, \hat{X}_t, \hat{\mu}_t, a) \hat{q}_t(\hat{X}_t, \hat{\mu})(da) \right] \right] dt + \mathbb{E}^{\hat{\mathbb{P}}} \left[g(\hat{X}_T, \hat{\mu}_T) \middle| \hat{\mu} \right] \\
&= \int_0^T \mathbb{E}^{\hat{\mathbb{P}}} \left[\int_A f(t, \hat{X}_t, \hat{\mu}_t, a) \hat{q}_t(\hat{X}_t, \hat{\mu})(da) \right] dt + \mathbb{E}^{\hat{\mathbb{P}}} \left[g(\hat{X}_T, \hat{\mu}_T) \right] \\
&= \int \mathfrak{F}(y, q, m) \hat{\Gamma}(dy, dm, dq).
\end{aligned}$$

Analogously, for every $\Sigma \in \mathcal{Q}$, we have

$$\int \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \rho(dm) = \int \mathfrak{F}(y, q, m) \Sigma(dy, dq, m) \hat{\rho}(dm),$$

which proves the desired statement about the payoff functional \mathfrak{p} . \square

Finally, we show that it is always possible to find a strong solution to equation (1.5.3) in the case of a feedback in state control process $\hat{q}_t(x, m)$, as given by Lemma 1.7.6:

Lemma 1.7.7 (Strong solutions for feedback in state controls). *Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space satisfying the usual assumptions, with Ω Polish and \mathcal{F} its Borel σ -algebra, supporting a d -dimensional \mathbb{F} -Brownian motion W , an \mathcal{F}_0 -measurable \mathbb{R}^d -valued random ξ with law η and a \mathcal{F}_0 -measurable random flow of measures μ in $\mathcal{C}(\mathcal{P}^2)$ with law ρ . Assume that ξ , W and μ are mutually independent. Let $\hat{q} : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \rightarrow \mathcal{P}(A)$ be a measurable function, and suppose that there exists a solution of the SDE*

$$dX_t = \int_A b(t, X_t, \mu_t, a) \hat{q}_t(X_t, \mu)(da) dt + dW_t, \tag{1.7.23}$$

so that it holds

$$\mu_t(\cdot) = \mathbb{P}(X_t \in \cdot \mid \mu) \quad \mathbb{P}\text{-a.s.}$$

for every $t \in [0, T]$. Then, X may be taken adapted to the \mathbb{P} -augmentation of the filtration $\mathbb{F}^{\xi, \mu, W} = \sigma(\xi) \vee \sigma(\mu) \vee \mathbb{F}^W$. In particular, there exists a progressively measurable function $\Phi : \mathcal{C}(\mathcal{P}^2) \times \mathbb{R}^d \times \mathcal{C}^d \rightarrow \mathcal{C}^d$ so that $\Phi(\mu, \xi, W) = X$ \mathbb{P} -a.s.

Proof. We set

$$B(t, x, m) = \int_A b(t, x, m_t, a) \hat{q}_t(x, m)(da).$$

B is jointly measurable in (t, x, m) in $[0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$ with at most linear growth in $(x, m) \in \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$ for every $t \in [0, T]$. The following hold:

1. For every $m \in \mathcal{C}(\mathcal{P}^2)$, equation

$$dX_t^m = B(t, X_t^m, m)dt + dW_t, \quad X_0^m = \xi. \quad (1.7.24)$$

admits a unique strong solution. Moreover, let $P^m = \mathbb{P} \circ (X^m)^{-1}$. Then, the map $\mathcal{C}(\mathcal{P}^2) \ni m \mapsto P^m \in \mathcal{P}(\mathcal{C}^d)$ is measurable.

2. There exist a continuous \mathbb{F} -adapted process X solution to

$$dX_t = B(t, X_t, \mu)dt + dW_t, \quad X_0 = \xi. \quad (1.7.25)$$

X is adapted to the \mathbb{P} -augmentation of the filtration $\mathbb{F}^{\xi, \mu, W}$.

3. Pathwise uniqueness holds, in the following sense: suppose there exists a pair of continuous \mathbb{F} -adapted processes (X^1, X^2) which satisfy equation (1.7.25) so that $(X_s^1, X_s^2)_{s \leq t}$ is conditionally independent of $\mathcal{F}_T^{\xi, \mu, W}$ given $\mathcal{F}_t^{\xi, \mu, W}$ for every $t \in [0, T]$. Then, $\mathbb{P}(X_t^1 = X_t^2, 0 \leq t \leq T) = 1$.

4. The joint law of X and μ is given by

$$\mathbb{P} \circ (X, \mu)^{-1} = P^m(dx)\rho(dm).$$

This properties can be proven with the same methods of [87, Appendix A] and [88, Appendix A]. We just point out that the results therein do not hold automatically in our case, since B is not progressively measurable in the measure flow m , in the sense of [87, 88]. Nevertheless, since we require μ to be \mathcal{F}_0 -measurable, the same arguments lead to the results above.

Let X be as in the statement of the lemma. We first show that the joint law of X and μ is given by $P^m(dx)\rho(dm)$. Let $\mathbb{P}^m(\cdot) = \mathbb{P}(\cdot | \mu = m)$ be a version of the regular conditional probability of \mathbb{P} given $\mu = m$. Then, since ξ , W and μ are mutually independent, $\mathbb{P}^m \circ (\xi, W)^{-1} = \mathbb{P} \circ (\xi, W)^{-1}$ for ρ -a.e. m , and, by exploiting the fact the $\mathcal{F}_t^{\xi, \mu, W, X}$ is countably generated for every t , W is an $\mathbb{F}^{\xi, \mu, W, X}$ -Brownian motion under \mathbb{P}^m as well. Therefore, X satisfies equation (1.7.24) on $(\Omega, \mathcal{F}, \mathbb{F}^{\xi, \mu, W, X}, \mathbb{P}^m)$ for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$. By point 1, $\mathbb{P}^m \circ X^{-1} = P^m$ for ρ -a.e. $m \in \mathcal{C}(\mathcal{P}^2)$, which implies that $\mathbb{P} \circ (X, \mu)^{-1} = P^m(dx)\rho(dm)$.

It can be shown by straightforward calculations that $(X_s)_{s \leq t}$ is conditionally independent of $\mathcal{F}_T^{\xi, \mu, W}$ given $\mathcal{F}_t^{\xi, \mu, W}$, for every $t \in [0, T]$. Since pathwise uniqueness holds by point 3, this implies that X is indistinguishable from an $\mathbb{F}^{\xi, \mu, W}$ -adapted solution to equation (1.7.25). \square

Chapter 2

Linear-quadratic mean field games

In this chapter we introduce coarse correlated equilibria in linear-quadratic MFGs. We propose a definition of mean field CCEs specific to the linear-quadratic framework and we develop a methodology to concretely compute mean field coarse correlated equilibria in a linear-quadratic mean field game framework. We compare their performance to mean field control solutions and mean field Nash equilibria. Our approach is implemented in the mean field version of an emission abatement game between greenhouse gas emitters. In particular, we exhibit a simple and tractable class of mean field CCEs which allows to outperform very significantly the mean field NE payoff and abatement levels, bridging the gap between the mean field NE and the social optimum obtained by mean field control. We also consider a simple linear-quadratic MFG, already known in literature, that does not admit any mean field Nash equilibrium, and we show that there exist infinitely many mean field CCEs.

2.1 Standing assumptions

Let $d, k \in \mathbb{N}$, $d, k \geq 1$. Let $T > 0$. The following set of assumptions, regarding coefficients and cost functions, will be in force throughout the whole chapter:

Assumptions LQ. The following matrices or matrix valued functions satisfy the following requirement:

- (1) $A, \sigma \in L^\infty([0, T]; \mathbb{R}^{d \times d})$;
- (2) $B \in L^\infty([0, T]; \mathbb{R}^{d \times k})$;
- (3) $Q, \bar{Q}, \tilde{Q} \in L^\infty([0, T]; \mathcal{S}^d)$, $R \in \mathcal{C}([0, T]; \mathcal{S}^k)$, $H, \bar{H}, \tilde{H} \in \mathcal{S}^d$;
- (4) $H \geq 0$, $Q_t \geq d_1 I_d$ for every $t \in [0, T]$, $d_1 \geq 0$, $R_t \geq d_2 I_k$ for every $t \in [0, T]$, $d_2 > 0$;
- (5) $S \in L^\infty([0, T]; \mathbb{R}^{k \times d})$, $\sup_{t \in [0, T]} |S_t|^2 < d_1 d_2$ if $d_1 > 0$, $S_t = 0$ for every $t \in [0, T]$ otherwise;
- (6) $L, q \in L^\infty([0, T]; \mathbb{R}^d)$, $r \in L^\infty([0, T]; \mathbb{R}^k)$.

Throughout this chapter, we fix a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ satisfying usual assumptions, large enough to support a d dimensional \mathbb{F} -Brownian motion W and an \mathbb{R}^d -valued \mathcal{F}_0 -measurable random variable with law $\eta \in \mathcal{P}^2(\mathbb{R}^d)$ ξ , independent of W . We recall the notation $\bar{\eta}$ to denote the first moment of η . We denote by $\mathbb{F}^{\xi, W} = (\mathcal{F}_t^{\xi, W})_{t \in [0, T]}$ the filtration generated by ξ and W , which we assume without loss of generality to satisfy the usual conditions. We assume the following assumption on the σ -algebra at initial time:

Assumption U. The σ -algebra \mathcal{F}_0 is large enough to support a \mathcal{F}_0 -measurable uniform random variable independent of ξ and W .

2.2 Mean field coarse correlated equilibria for the linear-quadratic MFG

In this section, we introduce the linear-quadratic MFG and propose a suitable definition of coarse correlated equilibria in this context. The notion of mean field CCE proposed here is different from coarse correlated solutions to the MFG in Definition 1.2.3 and it is specific for the linear-quadratic framework. In particular, we rely on the following notion of correlated moment flow.

Definition 2.2.1 (Correlated moment flow). A correlated moment flow is a pair $(\lambda, \bar{\mu})$ satisfying the following properties:

- i) $\lambda = (\lambda_t)_{t \in [0, T]}$ is a process in $\mathbb{H}^2(\mathbb{F})$.
- ii) $\bar{\mu} = (\bar{\mu}_t)_{t \in [0, T]}$ is an \mathcal{F}_0 -measurable \mathcal{C}^d -valued random variable.
- iii) $\bar{\mu}$ is independent of both ξ and W .

We refer to λ as the *recommended strategy* and to $\bar{\mu}$ as the *random flow of moments*.

We will sometimes use the equivalent expressions “correlated strategy” or “suggested strategy” to refer to λ . We can interpret a correlated moment flow $(\lambda, \bar{\mu})$ as follows: moderator’s lottery is run before the game starts and independently of the idiosyncratic shocks that determine the random evolution of representative player’s state. This is made possible by Assumption U, which allows for some independent extra randomness. We stress that, while the recommended strategy λ is correlated both to ξ and W and to $\bar{\mu}$, $\bar{\mu}$ is independent of the initial datum and the noise. We refer to the following Remark 2.2.2 for a discussion of the relationship with the definitions in Chapter 1.

Let us consider a correlated moment flow $(\lambda, \bar{\mu})$. We now assign dynamics and payoff functional. We consider a state variable with linear dynamics given by

$$dX_t = (A_t X_t + B_t \lambda_t) dt + \sigma_t dW_t, \quad X_0 = \xi, \quad (2.2.1)$$

and a linear-quadratic payoff functional

$$\begin{aligned} \mathfrak{J}(\lambda, \bar{\mu}) = \mathbb{E} \left[\int_0^T \left(\langle L_t, \bar{\mu}_t \rangle - \frac{1}{2} \langle \bar{Q}_t \bar{\mu}_t, \bar{\mu}_t \rangle \right) - \left(\frac{1}{2} \langle Q_t X_t, X_t \rangle + \langle \tilde{Q}_t X_t, \bar{\mu}_t \rangle + \frac{1}{2} \langle R_t \lambda_t, \lambda_t \rangle \right. \right. \\ \left. \left. + \langle q_t, X_t \rangle + \langle S_t X_t, \lambda_t \rangle + \langle r_t, \lambda_t \rangle \right) dt - \frac{1}{2} \langle \bar{H} \bar{\mu}_T, \bar{\mu}_T \rangle - \left(\frac{1}{2} \langle H X_T, X_T \rangle + \langle \tilde{H} X_T, \bar{\mu}_T \rangle \right) \right]. \end{aligned} \quad (2.2.2)$$

When needed, we will stress the dependence of the process X on the control λ by using the notation X^λ .

In order to move to the definition of mean field CCE, two cases must be distinguished. If the representative player decides to trust the mediator and therefore accepts to follow her recommendation λ before knowing it, the dynamics is given by equation (2.2.1), and the player gets the reward $\mathfrak{J}(\lambda, \bar{\mu})$. If instead she decides to deviate, she uses an open loop strategy β , i.e. a process $\beta = (\beta_t)_{t \in [0, T]}$ with values in \mathbb{R}^k , square integrable and $\mathbb{F}^{\xi, W}$ -progressively measurable. We denote the set of open loop strategies by \mathbb{A} , so that it holds

$$\mathbb{A} = \mathbb{H}^2(\mathbb{F}^{\xi, W}).$$

Then, the state dynamics of the deviating player is given by equation (2.2.1) with β instead of λ , and her reward is $\mathfrak{J}(\beta, \bar{\mu})$. Observe that when she deviates, her strategy β is measurable only with respect to the initial datum and the idiosyncratic noise, since she has no information on the outcome of the moderator's lottery. The deviating player can only use her knowledge of the law of the correlated moment flow $(\lambda, \bar{\mu})$, which is assumed to be publicly known. As a consequence, when deviating, the state process X of the representative player is independent of the random flow of moments $\bar{\mu}$, which, however, still appears in her payoff.

Definition 2.2.2 (Mean field coarse correlated equilibrium). A correlated moment flow $(\lambda, \bar{\mu})$ is a mean field CCE if the following holds:

- (i) Optimality: for every deviation $\beta \in \mathbb{A}$, it holds

$$\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\beta, \bar{\mu}). \quad (2.2.3)$$

- (ii) Consistency: let $X = (X_t)_{t \in [0, T]}$ be the solution to equation (2.2.1) with the control process λ . For every time $t \in [0, T]$, $\bar{\mu}_t$ is a version of the conditional expectation of X_t given $\bar{\mu}$, that is,

$$\bar{\mu}_t = \mathbb{E}[X_t | \bar{\mu}] \quad \mathbb{P}\text{-a.s.} \quad \forall t \in [0, T]. \quad (2.2.4)$$

The definition of mean field CCE has two fundamental differences with the usual definition of mean field NE for linear-quadratic MFGs. First of all, as already mentioned, the optimality condition features an asymmetry between the suggested strategy, which belongs to $\mathbb{H}^2(\mathbb{F})$, and deviating player's strategies, which belong to the smaller class \mathbb{A} , since the former depends also on the information used by the moderator to run her lottery while the latter does not. As for the consistency condition, we notice that, coherently with $\bar{\mu}$ being stochastic, it is formulated in terms of conditional

expectations, although no common noise is present. It should be interpreted in the following way: if all players commit to the mediator's lottery outcomes before knowing them, then the flow of measures should arise from aggregation of the individual behaviors. In the mean field limit, the influence of the idiosyncratic noise on the flow of moments vanishes, while the influence of moderator's lottery does not. Therefore, $\bar{\mu}$ stays stochastic and its stochasticity should derive from moderator's lottery only.

Remark 2.2.1. The reader might have noticed that $\bar{\mu}$ does not appear in the state dynamics (2.2.1). While computing mean field NEs and MFC solutions in the linear-quadratic case with the flow of moments in the dynamics is standard, computing mean field CCEs can be more delicate when $\bar{\mu}$ appears in the state dynamics. We refer to Section 2.3.1 and to Remark 2.3.2 therein for more explanations.

Remark 2.2.2. In Chapter 1, moderator's lottery was modeled by an auxiliary probability space, which was chosen by the moderator to support the extra randomness for her lottery. As a consequence, the recommended strategy, dynamics and payoff were naturally defined on a suitable product space supporting ξ, W and such extra randomness. Here, thanks to Assumption **U**, the given filtration \mathbb{F} is already big enough to allow for any extra randomization the moderator might want to use. In both formulations moderator's lottery is run independently of ξ and W and deviations are measurable with respect to ξ and W only. Moreover, while Chapter 1 considers a stochastic flow of measures, and the consistency condition is given in terms of conditional probabilities, here it is enough to consider a flow of moments and conditional expectations, due to the linear-quadratic structure of the MFG.

2.3 Computing mean field coarse correlated equilibria

The set of coarse correlated equilibria is typically very wide and it is difficult to characterize in a continuous time setting. We therefore focus on a tractable class of correlated moment flows for which we are able to characterize a sufficient condition for being a mean field CCE. To do so, we adopt the following procedure:

- We fix a correlated moment flow $(\lambda, \bar{\mu})$. We suppose that the representative player does not commit to the moderator's lottery and we compute her best deviating strategy $\hat{\beta}$, i.e.

$$\hat{\beta} = \arg \max_{\beta \in \mathbb{A}} \mathfrak{J}(\beta, \bar{\mu}).$$

This is the content of Proposition 2.3.1. Observe that $\hat{\beta}$ will depend upon the law of $(\lambda, \bar{\mu})$ itself, but not on its actual realization.

- We define a parameterised class of correlated moment flows $(\lambda, \bar{\mu})$ of similar shape as the best deviating strategy $\hat{\beta}$ so that the consistency condition (2.2.4) is fulfilled. The correlation is due to a suitable random parameter δ . This is accomplished in subsection 2.3.2.

- Finally, for $(\lambda, \bar{\mu})$ in such a class, with corresponding parameter δ , we express the optimality condition

$$\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\hat{\beta}, \bar{\mu})$$

as an inequality involving the law of $\bar{\mu}$ and δ only. Such an inequality is established in Theorem 2.3.3.

As a result, we reduce the search for a mean field CCE to finding a law for $\bar{\mu}$ and δ that verifies an optimality inequality. The choice of focusing on a class of correlated moment flows with shape similar to the best deviation allows for explicit analytical comparison between the two payoffs in the optimality condition (2.2.3).

Remark 2.3.1. Interestingly, the outlined procedure does not involve the usual two steps procedure used to compute mean field NEs: first, optimize with a fixed flow of moments and, second, perform a fixed point argument to determine the flow. Indeed, we first impose the consistency condition and then we verify the optimality condition, more in line with a MFC fashion. This sheds light on two important features of mean field CCEs. Firstly, they can be regarded as a middle ground between mean field NEs and MFC solutions. Secondly, mean field CCEs may exist even when mean field NEs do not, since the results of this section are still valid even when the MFG fixed-point condition fails to hold. The comparison between CCEs, NEs and MFC solutions and will be carried out in Section 2.4, and in Section 2.5 through the study of a simple yet important example. In Section 2.6 we will present a simple MFG which does not admit any mean field NE, and we show that it admits infinitely many mean field CCEs.

2.3.1 Deviating player's optimization problem

Suppose that the representative player does not commit to the lottery. Therefore, as anticipated in Section 2.2, she chooses a strategy on her own before the moderator sends his recommendation, hence in particular without any information on the realisation of the correlated moment flow. The only information she has about $(\lambda, \bar{\mu})$ is the joint law of the pair itself, which is assumed to be publicly known. Due to the linear-quadratic structure of the MFG and the fact that any admissible deviation β is independent of $\bar{\mu}$, it turns out that knowing the expectation of $\bar{\mu}_t$ for all $t \in [0, T]$ is enough.

Since the term $\int_0^T (\langle L_t, \bar{\mu}_t \rangle - \frac{1}{2} \langle \bar{Q}_t \bar{\mu}_t, \bar{\mu}_t \rangle) dt - \langle \bar{H} \bar{\mu}_T, \bar{\mu}_T \rangle$ in (2.2.2) can be viewed as an uncontrolled constant for the deviating player's optimization problem, we can focus on the equivalent optimization problem

$$\min_{\beta \in \mathbb{A}} \mathfrak{J}'(\beta, \bar{\mu}),$$

where

$$\begin{aligned} \mathfrak{J}'(\beta, \bar{\mu}) = \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \langle Q_t X_t, X_t \rangle + \langle \tilde{Q}_t X_t, \bar{\mu}_t \rangle + \langle q_t, X_t \rangle + \frac{1}{2} \langle R_t \beta_t, \beta_t \rangle + \langle S_t X_t, \beta_t \rangle \right. \right. \\ \left. \left. + \langle r_t, \beta_t \rangle \right) dt + \frac{1}{2} \langle H X_T, X_T \rangle + \langle \tilde{H} X_T, \bar{\mu}_T \rangle \right] \end{aligned} \tag{2.3.1}$$

under the constraint

$$dX_t = (A_t X_t + B_t \beta_t) dt + \sigma_t dW_t, \quad X_0 = \xi. \quad (2.3.2)$$

Since β is $\mathbb{F}^{\xi, W}$ -progressively measurable, it follows that X is $\mathbb{F}^{\xi, W}$ -adapted, and therefore is independent of the flow of moments $\bar{\mu}$, which implies that deviating player's payoff can be written as:

$$\begin{aligned} \mathfrak{J}'(\beta, \bar{\mu}) &= \int_0^T \mathbb{E} \left[\mathbb{E} \left[\frac{1}{2} \langle Q_t X_t, X_t \rangle + \langle \tilde{Q}_t X_t, \bar{\mu}_t \rangle + \langle q_t, X_t \rangle + \frac{1}{2} \langle R_t \beta_t, \beta_t \rangle + \langle S_t X_t, \beta_t \rangle \right. \right. \\ &\quad \left. \left. + \langle r_t, \beta_t \rangle \mid \mathcal{F}_t^{\xi, W} \right] \right] dt + \mathbb{E} \left[\mathbb{E} \left[\frac{1}{2} \langle H X_T, X_T \rangle + \langle \tilde{H} X_T, \bar{\mu}_T \rangle \mid \mathcal{F}_T^{\xi, W} \right] \right] \\ &= \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \langle Q_t X_t, X_t \rangle + \langle \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] \rangle + q_t, X_t \right) + \frac{1}{2} \langle R_t \beta_t, \beta_t \rangle + \langle S_t X_t, \beta_t \rangle + \langle r_t, \beta_t \rangle \right) dt \\ &\quad \left. + \frac{1}{2} \langle H X_T, X_T \rangle + \langle \tilde{H} \mathbb{E}[\bar{\mu}_T], X_T \rangle \right]. \end{aligned} \quad (2.3.3)$$

This is now a standard linear-quadratic control problem, which can be solved by the stochastic maximum principle.

Proposition 2.3.1 (Optimal strategy for the deviating player). *Let ϕ , ψ and θ be the solutions of the following ODEs:*

$$\begin{cases} \dot{\phi}_t + \phi_t A_t + A_t^\top \phi_t + Q_t - (\phi_t B_t + S_t^\top) R_t^{-1} (B_t^\top \phi_t + S_t) = 0, & \phi_T = H, \\ \dot{\psi}_t + A_t^\top \psi_t + \tilde{Q}_t - (\phi_t B_t + S_t^\top) R_t^{-1} B_t^\top \psi_t = 0, & \psi_T = \tilde{H}, \\ \dot{\theta}_t + \psi_t \frac{d\mathbb{E}[\bar{\mu}_t]}{dt} + A_t^\top \theta_t + q_t - (\phi_t B_t + S_t^\top) R_t^{-1} (B_t^\top \theta_t + r_t) = 0, & \theta_T = 0. \end{cases} \quad (2.3.4)$$

There exists a unique optimal strategy for the deviating player, which is given by

$$\hat{\beta}_t = -R_t^{-1} ((B_t^\top \phi_t + S_t) X_t^{\hat{\beta}} + B_t^\top \psi_t \mathbb{E}[\bar{\mu}_t] + B_t^\top \theta_t + r_t), \quad (2.3.5)$$

where $X^{\hat{\beta}}$ is the solution of equation (2.2.1) with the control $\hat{\beta}$.

We postpone the proof to Section 2.7. We observe only that the optimal control is actually feedback in the state X_t and in the expectation $\mathbb{E}[\bar{\mu}_t]$. Moreover, while the functions ϕ and ψ do not depend upon $\bar{\mu}$ or its expectation, the flow of expectations $\mathbb{E}[\bar{\mu}_t]$ appears in the equation for θ , through its time derivative $\frac{d\mathbb{E}[\bar{\mu}_t]}{dt}$.

Remark 2.3.2. This first step towards calculating mean field CCEs requires a filtering procedure, since the deviating player does not observe the actual realisation of $\bar{\mu}$. If the dynamics of the deviating player were dependent on $\bar{\mu}$, this step would require a much more involved analysis. Indeed, the state process X and $\bar{\mu}$ would not be independent, even if β is $\mathbb{F}^{\xi, W}$ -progressively measurable, which would lead to considering the projections on $\mathbb{F}^{\xi, W}$ of the processes X^i , $X^i X^j$ and $X^j \bar{\mu}^i$, $1 \leq i, j \leq d$. Alternatively, we could take advantage of the linear dynamics of X : Suppose that the drift of X features a linear term $C_t \bar{\mu}_t$, with $C \in L^\infty([0, T]; \mathbb{R}^{d \times d})$. Then, for any

$\beta \in \mathbb{H}^2(\mathbb{F})$, we could use the representation of X as $X_t = \tilde{X}_t + M_t^\mu$, where \tilde{X} and M^μ satisfy the following equations:

$$\begin{aligned} d\tilde{X}_t &= (A_t\tilde{X}_t + B_t\beta_t)dt + \sigma_t dW_t, & \tilde{X}_0 &= \xi, \\ dM_t^\mu &= (A_tM_t^\mu + C_t\bar{\mu}_t)dt, & M_0^\mu &= 0. \end{aligned}$$

If $\beta \in \mathbb{H}^2(\mathbb{F}^1)$, \tilde{X} is independent of $\bar{\mu}$, thus of M^μ , so that we could apply the same filtering procedure as before to the payoff functional to compute the best deviating strategy. Nevertheless, this would lead to much more complicated expressions, both for the class of correlated flows \mathcal{G} and for the optimality inequality (2.3.12), as one would need to take into considerations not only the moments of $\bar{\mu}$ but also the moments of M^μ , as well as mixed terms. This is why we have opted for a flow-free state dynamics, and postponed the analysis of the more general case to future research.

2.3.2 Correlated moment flow

We now consider a class of correlated moment flows $(\lambda, \bar{\mu})$ with a similar structure as the deviating player's best strategy $\hat{\beta}$ in (2.3.5). Our goal is to easily compare the payoff functionals $\mathfrak{J}'(\lambda, \bar{\mu})$ and $\mathfrak{J}'(\hat{\beta}, \bar{\mu})$. Hence we use the same functions ϕ and ψ , whereas we replace $\mathbb{E}[\bar{\mu}_t]$ with $\bar{\mu}_t$ itself and the term $R_t^{-1}(B_t^\top \theta_t + r_t)$ with a free parameter $\delta = (\delta_t)_{t \in [0, T]}$. Given any such δ , we define $\bar{\mu}$ so that the consistency condition (2.2.4) is satisfied, so that we will be left with taking care of the optimality condition only.

More precisely, let \mathcal{G} be the set of all correlated moment flows $(\lambda, \bar{\mu})$ defined as

$$\begin{aligned} \lambda_t &= -R_t^{-1}((B_t^\top \phi_t + S_t)X_t + B_t^\top \psi_t \bar{\mu}_t + \delta_t), \\ \dot{\bar{\mu}}_t &= (A_t - B_t R_t^{-1}(B_t^\top \phi_t + S_t + B_t^\top \psi_t))\bar{\mu}_t - B_t R_t^{-1} \delta_t, & \bar{\mu}_0 &= \bar{\eta}, \end{aligned} \tag{2.3.6}$$

where $\delta = (\delta_t)_{t \in [0, T]}$ is any process in $\mathbb{H}^2(\mathcal{F}_0)$ independent of ξ and W , and ϕ and ψ are as in (2.3.4). The parameter δ represents the extra source of randomness in the correlated moment flow with respect to ξ and W .

Lemma 2.3.2. *Any correlated moment flow $(\lambda, \bar{\mu}) \in \mathcal{G}$ satisfies the consistency condition (2.2.4).*

Proof. Let $(\lambda, \bar{\mu}) \in \mathcal{G}$ corresponding to some $\delta \in \mathbb{H}^2(\mathcal{F}_0)$ independent of ξ and W . To ease the notation, set

$$\Phi_t = R_t^{-1}(B_t^\top \phi_t + S_t), \quad \Psi_t = R_t^{-1}B_t^\top \psi_t, \quad \Theta_t = R_t^{-1}(B_t^\top \theta_t + r_t). \tag{2.3.7}$$

Notice that $\bar{\mu}$ satisfies the measurability requests of Definition 2.2.1. The dynamics of the representative player state is given by

$$\begin{aligned} dX_t &= ((A_t - B_t \Phi_t)X_t - B_t(\Psi_t \bar{\mu}_t + R_t^{-1} \delta_t)) dt + \sigma_t dW_t, \\ X_0 &= \xi, \end{aligned} \tag{2.3.8}$$

which implies that the process $(\bar{\mu}_t - X_t)_{t \in [0, T]}$ satisfies the stochastic differential equation

$$d(\bar{\mu}_t - X_t) = (A_t - B_t \Phi_t)(\bar{\mu}_t - X_t)dt - \sigma_t dW_t, \tag{2.3.9}$$

Since $\delta \in \mathbb{H}^2(\mathcal{F}_0)$, equation (2.3.8) admits a unique continuous adapted solution X satisfying $\mathbb{E}[\sup_{t \in [0, T]} |X_t|^2] < \infty$. Since ξ , W and δ are independent by assumption, by taking the conditional expectation with respect to $\bar{\mu}$ in (2.3.9), we get

$$d\mathbb{E}[\bar{\mu}_t - X_t | \bar{\mu}] = (A_t - B_t \Phi_t) \mathbb{E}[\bar{\mu}_t - X_t | \bar{\mu}] dt, \quad \mathbb{E}[\bar{\mu}_0 - X_0 | \bar{\mu}] = \bar{\mu}_0 - \mathbb{E}[\xi] = 0, \quad \mathbb{P}\text{-a.s.},$$

which implies $\mathbb{E}[\bar{\mu}_t - X_t | \bar{\mu}] = 0$ \mathbb{P} -a.s. for every t , i.e. (2.2.4). \square

Remark 2.3.3. Although the structure of the class \mathcal{G} is simple and quite specific, we will see later in Sections 2.5 and 2.6 that it is rich enough to contain a large set of mean field CCEs with some desirable properties, such as significantly outperforming the mean field NE.

2.3.3 Optimality condition

Let $(\lambda, \bar{\mu}) \in \mathcal{G}$. Since consistency has already been verified in Lemma 2.3.2, the goal is now to restate the optimality condition (2.2.3) in terms of quantities dependent upon the law of $\bar{\mu}$ and δ only.

Theorem 2.3.3. *Let $(\lambda, \bar{\mu}) \in \mathcal{G}$ corresponding to some $\delta \in \mathbb{H}^2(\mathcal{F}_0)$ independent of ξ and W . Let Φ , Ψ and Θ be given by in (2.3.7). Set*

$$M_t = Q_t + \Phi_t^\top R_t \Phi_t - 2\Phi_t^\top S_t, \quad N_t = \tilde{Q}_t + \Psi_t^\top R_t \Phi_t - \Psi_t^\top S_t, \quad G_t = \Psi_t^\top R_t \Psi_t. \quad (2.3.10)$$

Let $f(\bar{\mu}) = (f_t(\bar{\mu}))_{t \in [0, T]}$ be given by

$$\begin{cases} \dot{f}_t(\bar{\mu}) = (A_t - B_t \Phi_t) f_t(\bar{\mu}) + B_t (\Psi_t (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t]) + R_t^{-1} \delta_t - \Theta_t), & 0 \leq t \leq T, \\ f_0(\bar{\mu}) = 0. \end{cases} \quad (2.3.11)$$

Then, $(\lambda, \bar{\mu})$ is a mean field CCE if and only if the following condition is satisfied:

$$\begin{aligned} & \int_0^T \left(\mathbb{E}[\langle N_t (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t]), \bar{\mu}_t - \mathbb{E}[\bar{\mu}_t] \rangle] + \frac{1}{2} (\mathbb{E}[\langle G_t \bar{\mu}_t, \bar{\mu}_t \rangle] - \langle G_t \mathbb{E}[\bar{\mu}_t], \mathbb{E}[\bar{\mu}_t] \rangle) \right. \\ & \quad \left. + \frac{1}{2} \mathbb{E}[\langle R_t^{-1} \delta_t, \delta_t \rangle] - \frac{1}{2} \langle R_t \Theta_t, \Theta_t \rangle \right) dt + \mathbb{E}[\langle \tilde{H} (\bar{\mu}_T - \mathbb{E}[\bar{\mu}_T]), \bar{\mu}_T - \mathbb{E}[\bar{\mu}_T] \rangle] \\ & \leq \int_0^T \left(\frac{1}{2} (\mathbb{E}[\langle M_t (\bar{\mu}_t + f_t(\bar{\mu})), \bar{\mu}_t + f_t(\bar{\mu}) \rangle] - \mathbb{E}[\langle M_t \bar{\mu}_t, \bar{\mu}_t \rangle]) + \mathbb{E}[\langle N_t f_t(\bar{\mu}), \mathbb{E}[\bar{\mu}_t] \rangle] \right. \\ & \quad \left. + \langle q_t - \Phi_t^\top r_t, \mathbb{E}[f_t(\bar{\mu})] \rangle + \langle B_t^\top (\phi_t + \psi_t) \mathbb{E}[\bar{\mu}_t], \Theta_t \rangle - \mathbb{E}[\langle B_t^\top (\phi_t + \psi_t) \bar{\mu}_t, R_t^{-1} \delta_t \rangle] \right. \\ & \quad \left. + \mathbb{E}[\langle B_t^\top \phi_t f_t(\bar{\mu}), \Theta_t \rangle] - \mathbb{E}[\langle r_t, \Theta_t - R_t^{-1} \delta_t \rangle] \right) dt \\ & \quad + \frac{1}{2} (\mathbb{E}[\langle H (\bar{\mu}_T + f_T(\bar{\mu})), \bar{\mu}_T + f_T(\bar{\mu}) \rangle] - \mathbb{E}[\langle H \bar{\mu}_T, \bar{\mu}_T \rangle]) + \langle \tilde{H} \mathbb{E}[f_T(\bar{\mu})], \mathbb{E}[\bar{\mu}_T] \rangle. \end{aligned} \quad (2.3.12)$$

Proof. Since $(\lambda, \bar{\mu}) \in \mathcal{G}$ satisfies the consistency condition (2.2.4) by Lemma 2.3.2, we focus on the optimality condition (2.2.3). Since the term depending on $\bar{\mu}$ only $\int_0^T (\mathbb{E}[\langle L_t, \bar{\mu}_t \rangle] - \frac{1}{2} \mathbb{E}[\langle \tilde{Q}_t \bar{\mu}_t, \bar{\mu}_t \rangle]) dt - \mathbb{E}[\langle \tilde{H} \bar{\mu}_T, \bar{\mu}_T \rangle]$ in (2.2.2) is the same for the payoffs of

the representative player and the deviating player, the optimality condition is satisfied if and only if

$$\mathfrak{J}'(\lambda, \bar{\mu}) \leq \min_{\hat{\beta} \in \mathbb{A}} \mathfrak{J}'(\hat{\beta}, \bar{\mu}) = \mathfrak{J}'(\hat{\beta}, \bar{\mu}), \quad (2.3.13)$$

with $\hat{\beta}$ given by (2.3.5) and $\mathfrak{J}'(\lambda, \bar{\mu})$ and $\mathfrak{J}'(\hat{\beta}, \bar{\mu})$ defined by (2.3.3). Denote by $X^{\hat{\beta}} = (X_t^{\hat{\beta}})_{t \in [0, T]}$ the state of the deviating player when she uses the strategy $\hat{\beta}$ defined in (2.3.5), i.e.

$$dX_t^{\hat{\beta}} = ((A_t - B_t \Phi_t) X_t^{\hat{\beta}} - B_t (\Psi_t \mathbb{E}[\bar{\mu}_t] + \Theta_t)) dt + \sigma_t dW_t, \quad \hat{X}_0 = \xi,$$

and by $X = (X_t)_{t \in [0, T]}$ the state of the representative player corresponding to the correlated moment flow (2.3.6), i.e.

$$dX_t = ((A_t - B_t \Phi_t) X_t - B_t (\Psi_t \bar{\mu}_t + R_t^{-1} \delta_t)) dt + \sigma_t dW_t, \quad X_0 = \xi.$$

We rewrite the cost functionals $\mathfrak{J}'(\lambda, \bar{\mu})$ and $\mathfrak{J}'(\hat{\beta}, \bar{\mu})$ by taking advantage of the explicit form of λ and $\hat{\beta}$ and functions (2.3.10):

$$\begin{aligned} \mathfrak{J}'(\lambda, \bar{\mu}) = & \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \langle M_t X_t, X_t \rangle + \langle N_t X_t, \bar{\mu}_t \rangle + \frac{1}{2} \langle G_t \bar{\mu}_t, \bar{\mu}_t \rangle + \langle q_t - \Phi_t^\top r_t, X_t \rangle \right. \right. \\ & + \langle (R_t \Phi_t - S_t) X_t, R_t^{-1} \delta_t \rangle + \langle R_t \Psi_t \bar{\mu}_t, R_t^{-1} \delta_t \rangle - \langle \Psi_t^\top r_t, \bar{\mu}_t \rangle \\ & \left. \left. + \frac{1}{2} \langle R_t^{-1} \delta_t, \delta_t \rangle - \langle r_t, R_t^{-1} \delta_t \rangle \right) dt + \frac{1}{2} \langle H X_T, X_T \rangle + \langle \tilde{H} X_T, \bar{\mu}_T \rangle \right], \end{aligned}$$

and

$$\begin{aligned} \mathfrak{J}'(\hat{\beta}, \bar{\mu}) = & \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \langle M_t X_t^{\hat{\beta}}, X_t^{\hat{\beta}} \rangle + \langle N_t X_t^{\hat{\beta}}, \mathbb{E}[\bar{\mu}_t] \rangle + \frac{1}{2} \langle G_t \mathbb{E}[\bar{\mu}_t], \mathbb{E}[\bar{\mu}_t] \rangle \right. \right. \\ & + \langle q_t - \Phi_t^\top r_t, X_t^{\hat{\beta}} \rangle + \langle (R_t \Phi_t - S_t) \hat{X}_t, \Theta_t \rangle + \langle R_t \Psi_t \mathbb{E}[\bar{\mu}_t], \Theta_t \rangle - \langle \Psi_t^\top r_t, \mathbb{E}[\bar{\mu}_t] \rangle \\ & \left. \left. + \frac{1}{2} \langle R_t \Theta_t, \Theta_t \rangle - \langle r_t, \Theta_t \rangle \right) dt + \frac{1}{2} \langle H X_T^{\hat{\beta}}, X_T^{\hat{\beta}} \rangle + \langle \tilde{H} X_T^{\hat{\beta}}, \mathbb{E}[\bar{\mu}_T] \rangle \right]. \end{aligned}$$

By Itô's formula, we get

$$\begin{cases} d(X_t^{\hat{\beta}} - X_t) = (A_t - B_t \Phi_t)(X_t^{\hat{\beta}} - X_t) dt + B_t (\Psi_t (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t]) + R_t^{-1} \delta_t - \Theta_t) dt, \\ X_0^{\hat{\beta}} - X_0 = 0, \end{cases}$$

so that it holds

$$X_t^{\hat{\beta}} = X_t + f_t(\bar{\mu}), \quad 0 \leq t \leq T, \quad \mathbb{P}\text{-a.s.} \quad (2.3.14)$$

In particular, we note that $f(\bar{\mu})$ is $\sigma(\bar{\mu})$ -measurable. Then, we have

$$\begin{aligned} \mathbb{E}[\langle M_t X_t^{\hat{\beta}}, X_t^{\hat{\beta}} \rangle] &= \mathbb{E}[\langle M_t (X_t + f_t(\bar{\mu})), X_t + f_t(\bar{\mu}) \rangle] \\ &= \mathbb{E}[\langle M_t X_t, X_t \rangle] + \mathbb{E}[\langle M_t (\bar{\mu}_t + f_t(\bar{\mu})), \bar{\mu}_t + f_t(\bar{\mu}) \rangle] - \mathbb{E}[\langle M_t \bar{\mu}_t, \bar{\mu}_t \rangle], \end{aligned}$$

where we have used the fact that X satisfies the consistency condition (2.2.4). Therefore, we have

$$\begin{aligned}
\mathfrak{J}'(\hat{\beta}, \bar{\mu}) &= \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \langle M_t X_t, X_t \rangle + \frac{1}{2} \langle M_t(\bar{\mu}_t + f_t(\bar{\mu})), \bar{\mu}_t + f_t(\bar{\mu}) \rangle - \frac{1}{2} \langle M_t \bar{\mu}_t, \bar{\mu}_t \rangle \right. \right. \\
&\quad + \langle N_t X_t, \mathbb{E}[\bar{\mu}_t] \rangle + \langle N_t f_t(\bar{\mu}), \mathbb{E}[\bar{\mu}_t] \rangle + \frac{1}{2} \langle G_t \mathbb{E}[\bar{\mu}_t], \mathbb{E}[\bar{\mu}_t] \rangle + \langle q_t - \Phi_t^\top r_t, X_t \rangle \\
&\quad + \langle q_t - \Phi_t^\top r_t, f_t(\bar{\mu}) \rangle + \langle (R_t \Phi_t - S_t) X_t, \Theta_t \rangle + \langle (R_t \Phi_t - S_t) f_t(\bar{\mu}), \Theta_t \rangle \\
&\quad + \langle R_t \Psi_t \mathbb{E}[\bar{\mu}_t], \Theta_t \rangle - \langle \Psi_t^\top r_t, \mathbb{E}[\bar{\mu}_t] \rangle + \frac{1}{2} \langle R_t \Theta_t, \Theta_t \rangle - \langle r_t, \Theta_t \rangle \Big) dt \\
&\quad + \frac{1}{2} \langle H X_T, X_T \rangle + \frac{1}{2} \langle H(\bar{\mu}_T + f_T(\bar{\mu})), \bar{\mu}_T + f_T(\bar{\mu}) \rangle - \frac{1}{2} \langle H \bar{\mu}_T, \bar{\mu}_T \rangle \\
&\quad \left. + \langle \tilde{H} X_T, \mathbb{E}[\bar{\mu}_T] \rangle + \langle \tilde{H} f_T(\bar{\mu}), \mathbb{E}[\bar{\mu}_T] \rangle \right].
\end{aligned}$$

Since the correlated moment flow $(\lambda, \bar{\mu})$ satisfies the consistency condition (2.2.4), and noticing that

$$\begin{aligned}
\mathbb{E}[\langle N_t \bar{\mu}_t, \mathbb{E}[\bar{\mu}_t] - \bar{\mu}_t \rangle] &= -\mathbb{E}[\langle N_t(\mathbb{E}[\bar{\mu}_t] - \bar{\mu}_t), \mathbb{E}[\bar{\mu}_t] - \bar{\mu}_t \rangle], \\
\mathbb{E}[\langle (R_t \Phi_t - S_t) \bar{\mu}_t, \Theta_t \rangle] &= \langle (R_t \Phi_t - S_t) \mathbb{E}[\bar{\mu}_t], \Theta_t \rangle,
\end{aligned}$$

we obtain

$$\begin{aligned}
\mathfrak{J}'(\hat{\beta}, \bar{\mu}) - \mathfrak{J}'(\lambda, \bar{\mu}) &= \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \langle M_t(\bar{\mu}_t + f_t(\bar{\mu})), \bar{\mu}_t + f_t(\bar{\mu}) \rangle - \frac{1}{2} \langle M_t \bar{\mu}_t, \bar{\mu}_t \rangle \right. \right. \\
&\quad - \langle N_t(\mathbb{E}[\bar{\mu}_t] - \bar{\mu}_t), \mathbb{E}[\bar{\mu}_t] - \bar{\mu}_t \rangle + \langle N_t f_t(\bar{\mu}), \mathbb{E}[\bar{\mu}_t] \rangle + \frac{1}{2} \langle G_t \mathbb{E}[\bar{\mu}_t], \mathbb{E}[\bar{\mu}_t] \rangle - \frac{1}{2} \langle G_t \bar{\mu}_t, \bar{\mu}_t \rangle \\
&\quad + \langle q_t - \Phi_t^\top r_t, f_t(\bar{\mu}) \rangle + \langle (R_t(\Phi_t + \Psi_t) - S_t) \mathbb{E}[\bar{\mu}_t], \Theta_t \rangle - \langle (R_t(\Phi_t + \Psi_t) - S_t) \bar{\mu}_t, R_t^{-1} \delta_t \rangle \\
&\quad + \langle (R_t \Phi_t - S_t) f_t(\bar{\mu}), \Theta_t \rangle + \frac{1}{2} \langle R_t \Theta_t, \Theta_t \rangle - \frac{1}{2} \langle R_t^{-1} \delta_t, \delta_t \rangle - \langle r_t, \Theta_t - R_t^{-1} \delta_t \rangle \Big) dt \\
&\quad + \frac{1}{2} \langle H(\bar{\mu}_T + f_T(\bar{\mu})), \bar{\mu}_T + f_T(\bar{\mu}) \rangle - \frac{1}{2} \langle H \bar{\mu}_T, \bar{\mu}_T \rangle \\
&\quad \left. - \langle \tilde{H}(\mathbb{E}[\bar{\mu}_T] - \bar{\mu}_T), \mathbb{E}[\bar{\mu}_T] - \bar{\mu}_T \rangle + \langle \tilde{H} f_T(\bar{\mu}), \mathbb{E}[\bar{\mu}_T] \rangle \right].
\end{aligned}$$

Therefore, the correlated moment flow $(\lambda, \bar{\mu})$ defined by (2.3.6) is a mean field CCE if and only if the RHS above is non-negative. By rearranging the terms and using the equalities

$$R_t \Phi_t - S_t = B_t^\top \phi_t, \quad R_t \Psi_t = B_t^\top \psi_t, \quad (2.3.15)$$

we get condition (2.3.12). \square

The condition for a correlated moment flow of class \mathcal{G} to be a mean field CCE is now reduced to an optimality condition which only depends on the population average state and the correlating device of the mediator, i.e. on the joint law of $(\delta, \bar{\mu})$. Even though the inequality looks quite long, it can become very tractable and easy to interpret when one specifies some class of dynamics for $\bar{\mu}$ as done in Sections 2.5 and 2.6.

2.4 Comparison with MFC solution and mean field NE

In this section, we analyze the relationship between mean field CCEs, mean field NEs and MFC solutions. In more detail, we prove the following results:

- Under the additional Assumptions **LQ-MFC**, we compute the MFC solution $\hat{\alpha}$ and we show that it is unique and optimal in the broader class of controls $\mathbb{H}^2(\mathbb{F})$. This is accomplished in Proposition 2.4.1 and Lemma 2.4.2. Then, we show that no mean field CCE can outperform the payoff of the MFC solution. Moreover, if the MFC solution is not a mean field NE, we establish that the MFC payoff is unattainable by a mean field CCE. This is accomplished in Theorem 2.4.3.
- As for mean field NE, under the additional Assumptions **LQ-NE**, we show in Proposition 2.4.4 that there exists a unique mean field NE in our setting. Then, we show that for a correlated moment flow $(\lambda, \bar{\mu})$ to be a mean field CCE different from the mean field NE, it is necessary that the flow of moments $\bar{\mu}$ is stochastic, which is the content of Theorem 2.4.5.
- Finally, Theorem 2.4.6 gives a condition so that a mean field CCE $(\lambda, \bar{\mu}) \in \mathcal{G}$ yields a higher payoff than the mean field NE.

We remark that the results in the first two points above are fully general, in the sense that they do not restrict to correlated moment flows in the class \mathcal{G} defined by (2.3.6), while the condition on a mean field CCE $(\lambda, \bar{\mu})$ to outperform the payoff of the mean field NE is provided only for correlated moment flows in \mathcal{G} .

We recall here for reader's convenience the definitions of both mean field NE and MFC solution.

Definition 2.4.1. We say that a pair $(\alpha^*, \bar{m}^*) \in \mathbb{A} \times \mathcal{C}^d$ is a mean field Nash equilibrium if the following properties hold:

- (i) Optimality: α^* maximizes $\mathfrak{J}(\cdot, \bar{m}^*)$ over \mathbb{A} , i.e.,

$$\mathfrak{J}(\alpha^*, \bar{m}^*) = \max_{\beta \in \mathbb{A}} \mathfrak{J}(\beta, \bar{m}^*). \quad (2.4.1)$$

- (ii) Consistency: let $X^* = (X_t^*)_{t \in [0, T]}$ be the solution to equation (2.2.1) with the control process α^* . For every time $t \in [0, T]$, \bar{m}_t^* equals the expectation of X_t^* , i.e.,

$$\bar{m}_t^* = \mathbb{E}[X_t^*], \quad \forall t \in [0, T]. \quad (2.4.2)$$

Definition 2.4.2. For any $\beta \in \mathbb{A}$, let X^β be the solution of equation (2.2.1) with β instead of λ . Denote by $\mathbb{E}[X^\beta] = (\mathbb{E}[X_t^\beta])_{t \in [0, T]}$ the corresponding flow of first order moments. We say that a strategy $\hat{\alpha}$ is a MFC solution, if

$$\mathfrak{J}(\hat{\alpha}, \mathbb{E}[X^{\hat{\alpha}}]) = \max_{\beta \in \mathbb{A}} \mathfrak{J}(\beta, \mathbb{E}[X^\beta]). \quad (2.4.3)$$

2.4.1 Comparison with MFC solution

In this subsection we compare the expected payoffs of mean field CCEs and the MFC solution. To this extent, we assume the following additional assumptions:

Assumptions LQ-MFC. Referring to the matrix valued functions introduced in Assumptions **LQ**, we assume:

- (1) $Q_t + 2\tilde{Q}_t + \bar{Q}_t \geq \hat{d}_1 I_d$, for every $t \in [0, T]$, $\hat{d}_1 \geq 0$;
- (2) $\sup_{t \in [0, T]} |S_t|^2 < \hat{d}_1 d_2$ if $\hat{d}_1 > 0$, $S_t = 0$ for every $t \in [0, T]$ otherwise;
- (3) $H + 2\tilde{H} + \bar{H} \geq 0$.

Under Assumptions **LQ** and **LQ-MFC**, there exists a unique MFC solution. Since computations are very standard, we postpone them to Section 2.7.

Proposition 2.4.1. *Let $\hat{\phi}$ and $\hat{\theta}$ be the solutions of the following equations:*

$$\begin{cases} \dot{\hat{\phi}}_t + \hat{\phi}_t A_t + A_t^\top \hat{\phi}_t + (Q_t + 2\tilde{Q}_t + \bar{Q}_t) - (\hat{\phi}_t B_t + S_t^\top) R_t^{-1} (B_t^\top \hat{\phi}_t + S_t) = 0, \\ \hat{\phi}_T = H + 2\tilde{H} + \bar{H}, \\ \dot{\hat{\theta}}_t + A_t^\top \hat{\theta}_t + q_t - L_t - (\hat{\phi}_t B_t + S_t^\top) R_t^{-1} (B_t^\top \hat{\theta}_t + r_t) = 0, \\ \hat{\theta}_T = 0. \end{cases} \quad (2.4.4)$$

Define \bar{A} and \bar{B} as

$$\bar{A}_t = A_t - B_t R_t^{-1} B_t^\top \hat{\phi}_t - B_t R_t^{-1} S_t, \quad \bar{B}_t = B_t R_t^{-1} (B_t^\top \hat{\theta}_t + r_t). \quad (2.4.5)$$

Let $\bar{\psi}$ and $\bar{\theta}$ be the solutions of the following equations:

$$\begin{cases} \dot{\bar{\psi}}_t + \bar{A}_t^\top \bar{\psi}_t + A_t^\top \bar{\psi}_t + (\bar{Q}_t + 2\tilde{Q}_t) - (\bar{\phi}_t B_t + S_t^\top) R_t^{-1} B_t^\top \bar{\psi}_t = 0, & \bar{\psi}_T = \bar{H} + 2\tilde{H}, \\ \dot{\bar{\theta}}_t - \bar{\psi}_t \bar{B}_t + A_t^\top \bar{\theta}_t + q_t - (\bar{\phi}_t B_t + S_t^\top) R_t^{-1} (B_t^\top \bar{\theta}_t + r_t) = 0, & \bar{\theta}_T = 0. \end{cases} \quad (2.4.6)$$

Let ϕ be the solution of the matrix Riccati equation in (2.3.4). There exists a unique MFC solution $\hat{\alpha}$, which is given by

$$\hat{\alpha}_t = -R_t^{-1} ((B_t^\top \phi_t + S_t) \hat{X}_t + B_t^\top \bar{\psi}_t \hat{x}_t + (B_t^\top \bar{\theta}_t + r_t)), \quad (2.4.7a)$$

$$\dot{\hat{x}}_t = (A_t - B_t R_t^{-1} B_t^\top \phi_t - B_t R_t^{-1} S_t) \hat{x}_t - B_t R_t^{-1} (B_t^\top \bar{\theta}_t + r_t), \quad \hat{x}_0 = \bar{\eta}, \quad (2.4.7b)$$

where \hat{X} is the solution of

$$\begin{cases} d\hat{X}_t = ((A_t - R_t^{-1} (B_t^\top \phi_t + S_t)) \hat{X}_t - R_t^{-1} B_t^\top \bar{\psi}_t \hat{x}_t - R_t^{-1} (B_t^\top \bar{\theta}_t + r_t)) dt + \sigma_t dW_t, \\ \hat{X}_0 = \xi. \end{cases} \quad (2.4.8)$$

In particular, it holds $\hat{x}_t = \mathbb{E}[\hat{X}_t]$ for every $t \in [0, T]$.

Assumptions **LQ-MFC** guarantee that the payoff functional is strictly concave jointly in the state and the measure arguments, and that the Riccati equation in (2.4.4) is uniquely solvable.

Showing that no mean field CCE can outperform the payoff of the MFC solution requires first to show that the MFC solution is actually optimal over the larger control set $\mathbb{H}^2(\mathbb{F})$, as it is done in the following preliminary lemma:

Lemma 2.4.2. *Let $\hat{\alpha}$ be the solution of the MFC problem. Then, for any β in $\mathbb{H}^2(\mathbb{F})$, $\beta \neq \hat{\alpha}$, it holds*

$$\mathfrak{J}(\hat{\alpha}, \hat{x}) > \mathfrak{J}(\beta, \mathbb{E}[X^\beta]). \quad (2.4.9)$$

Proof. To ease the notation, we set

$$\mathfrak{J}^{MFC}(\alpha) = \mathfrak{J}(\alpha, \mathbb{E}[X^\alpha]), \quad (2.4.10)$$

where the process X has dynamics given by (2.2.1), for any $\alpha \in \mathbb{H}^2(\mathbb{F})$. We observe that

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{2} \langle \bar{Q}_t \mathbb{E}[X_t], \mathbb{E}[X_t] \rangle + \langle \tilde{Q}_t X_t, \mathbb{E}[X_t] \rangle + \frac{1}{2} \langle Q_t X_t, X_t \rangle + \frac{1}{2} \langle R_t \alpha_t, \alpha_t \rangle + \langle S_t X_t, \alpha_t \rangle \right] \\ &= \mathbb{E} \left[\frac{1}{2} \langle (\bar{Q}_t + 2\tilde{Q}_t) \mathbb{E}[X_t], \mathbb{E}[X_t] \rangle + \frac{1}{2} \langle Q_t X_t, X_t \rangle + \frac{1}{2} \langle R_t \alpha_t, \alpha_t \rangle + \langle S_t X_t, \alpha_t \rangle \right] \\ &= \mathbb{E} \left[\frac{1}{2} \langle (\bar{Q}_t + 2\tilde{Q}_t + Q_t) \mathbb{E}[X_t], \mathbb{E}[X_t] \rangle + \frac{1}{2} \langle Q_t (X_t - \mathbb{E}[X_t]), (X_t - \mathbb{E}[X_t]) \rangle \right. \\ & \quad \left. + \frac{1}{2} \langle R_t \alpha_t, \alpha_t \rangle + \langle S_t (X_t - \mathbb{E}[X_t]), \alpha_t \rangle + \langle S_t \mathbb{E}[X_t], \alpha_t \rangle \right]. \end{aligned} \quad (2.4.11)$$

By Assumptions **LQ-MFC**, this equality implies that the running payoff in the functional \mathfrak{J}^{MFC} is strictly concave jointly in $\mathbb{E}[X_t]$, $X_t - \mathbb{E}[X_t]$ and α_t , for every $t \in [0, T]$. Since \mathfrak{J}^{MFC} is also upper semi-continuous, this implies that the maximum exists and that it is unique over the broader class $\mathbb{H}^2(\mathbb{F})$.

We are left to show that the maximum point is indeed $\hat{\alpha}$. For the sake of clarity, we set

$$\begin{aligned} f(t, x, m, a) &= \langle L_t, m \rangle - \frac{1}{2} \langle \bar{Q}_t m, m \rangle - \frac{1}{2} \langle Q_t x, x \rangle - \langle \tilde{Q}_t x, m \rangle - \langle q_t, x \rangle \\ & \quad - \frac{1}{2} \langle R_t a, a \rangle - \langle S_t x, a \rangle - \langle r_t, a \rangle, \\ g(x, m) &= - \left(\frac{1}{2} \langle \bar{H} m, m \rangle + \frac{1}{2} \langle H x, x \rangle + \langle \tilde{H} x, m \rangle \right). \end{aligned} \quad (2.4.12)$$

Let β in $\mathbb{H}^2(\mathbb{F})$. We define the following process:

$$\tilde{\beta}_t = \mathbb{E}[\beta_t | \mathcal{F}_t^{\xi, W}], \quad t \in [0, T]. \quad (2.4.13)$$

Since $\mathbb{F}^{\xi, W}$ satisfies the usual assumptions, $\tilde{\beta}$ can be taken $\mathbb{F}^{\xi, W}$ -progressively measurable (see, e.g., [23, Section 2]), so that we have $\tilde{\beta} \in \mathbb{A}$. Let $\tilde{X} = (\tilde{X}_t)_{t \in [0, T]}$ be the solution of

$$d\tilde{X}_t = (A_t \tilde{X}_t + B_t \tilde{\beta}_t) dt + \sigma_t dW_t, \quad \tilde{X}_0 = \xi.$$

Then, using the explicit expression for the solution \tilde{X} of the SDE above, it can be shown by direct computation that

$$\tilde{X}_t = \mathbb{E}[X_t^\beta | \mathcal{F}_t^{\xi, W}] \quad \mathbb{P}\text{-a.s.}, \quad t \in [0, T].$$

Due to the concave linear-quadratic structure of f , we have the following:

$$\begin{aligned} \mathfrak{J}^{MFC}(\beta) &= \mathbb{E} \left[\int_0^T f(t, X_t^\beta, \mathbb{E}[X_t^\beta], \beta_t) dt + g(X_T^\beta, \mathbb{E}[X_T^\beta]) \right] \\ &= \int_0^T \mathbb{E} \left[\mathbb{E} \left[f(t, X_t^\beta, \mathbb{E}[X_t^\beta], \beta_t) \mid \mathcal{F}_t^{\xi, W} \right] \right] dt + \mathbb{E} \left[\mathbb{E} \left[g(X_T^\beta, \mathbb{E}[X_T^\beta]) \mid \mathcal{F}_T^{\xi, W} \right] \right] \\ &\leq \mathbb{E} \left[\int_0^T f(t, \mathbb{E}[X_t^\beta | \mathcal{F}_t^{\xi, W}], \mathbb{E}[X_t^\beta], \mathbb{E}[\beta_t | \mathcal{F}_t^{\xi, W}]) dt + g(\mathbb{E}[X_T^\beta | \mathcal{F}_T^{\xi, W}], \mathbb{E}[X_T^\beta]) \right] \\ &= \mathbb{E} \left[\int_0^T f(t, \tilde{X}_t, \mathbb{E}[\tilde{X}_t], \tilde{\beta}_t) dt + g(\tilde{X}_T, \mathbb{E}[\tilde{X}_T]) \right] = \mathfrak{J}^{MFC}(\tilde{\beta}), \end{aligned}$$

where we have used the fact that $\mathbb{E}[X_t^\beta] = \mathbb{E}[\tilde{X}_t]$ for every time t . Since $\tilde{\beta}$ belongs to \mathbb{A} , Proposition 2.4.1 implies

$$\mathfrak{J}^{MFC}(\beta) \leq \mathfrak{J}^{MFC}(\tilde{\beta}) \leq \mathfrak{J}^{MFC}(\hat{\alpha}).$$

By strict concavity, we deduce that the inequality is strict for any $\beta \neq \hat{\alpha}$. \square

In the next theorem we prove that the MFC solution provides an upper bound to the payoffs of any mean field CCEs. Moreover, this upper bound can not be attained unless the MFC solution is a mean field NE.

Theorem 2.4.3 (No outperformance over the MFC solution). *Let $(\lambda, \bar{\mu})$ a mean field CCE. Then, the following holds:*

- (i) *If $\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\hat{\alpha}, \hat{x})$, then $(\lambda, \bar{\mu}) = (\hat{\alpha}, \hat{x})$, so $\mathfrak{J}(\lambda, \bar{\mu}) = \mathfrak{J}(\hat{\alpha}, \hat{x})$;*
- (ii) *If the MFC solution is not a mean field NE, then $\mathfrak{J}(\lambda, \bar{\mu}) < \mathfrak{J}(\hat{\alpha}, \hat{x})$. In particular, the MFC solution is not a mean field CCE either.*

Proof of (i). By using the payoff functional \mathfrak{J}^{MFC} defined by (2.4.10), the payoffs' inequality reads as

$$\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\hat{\alpha}, \hat{x}) = \mathfrak{J}^{MFC}(\hat{\alpha}). \quad (2.4.14)$$

We reformulate the MFC problem weakly, by taking advantages of the results of [38, Paragraph 6.6]. We define the set $\mathcal{A} \subseteq \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times L^2([0, T]; \mathbb{R}^k))$ of admissible probability measures in the following way: take any filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ satisfying the usual assumptions, equipped with a d -dimensional \mathbb{F} -Brownian motion W and an \mathcal{F}_0 -measurable random variable ξ independent of W . Let $\alpha \in \mathbb{H}^2(\mathbb{F})$, which we regard as random variable taking values in $L^2([0, T]; \mathbb{R}^k)$. Let $X = X^\alpha$ be the solution of

$$dX_t^\alpha = (A_t X_t^\alpha + B_t \alpha_t) dt + dW_t, \quad X_0^\alpha = \xi. \quad (2.4.15)$$

Then, a probability measure P belongs to \mathcal{A} if $P = \mathbb{P} \circ (\xi, X^\alpha, \alpha)^{-1}$. For any $P \in \mathcal{A}$, set $\bar{x}_t = \int_{\mathbb{R}^d} y(P \circ x_t^{-1})(dy)$, where $\mathcal{C}([0, T]; \mathbb{R}^d) \ni x \mapsto x_t \in \mathbb{R}^d$ is the projection of a

trajectory x at time t . By recalling the definitions of f and g in (2.4.12), define the payoff functional

$$\begin{aligned}\mathcal{J}(P) &= \int_{\mathbb{R}^d \times \mathcal{C}^d \times L^2([0, T]; \mathbb{R}^k)} \left(\int_0^T f(t, x_t, \bar{x}_t, a) dt + g(x_T, \bar{x}_T) \right) P(dz, dx, da) \\ &= \mathbb{E}^{\mathbb{P}} \left[\int_0^T (f(t, X_t, \mathbb{E}^{\mathbb{P}}[X_t], \alpha_t)) dt + g(X_T, \mathbb{E}^{\mathbb{P}}[X_T]) \right].\end{aligned}\tag{2.4.16}$$

By [38, Theorem 6.37], there exists a probability measure P^* in \mathcal{A} so that

$$\mathcal{J}(P^*) \geq \mathcal{J}(P) \quad \forall P \in \mathcal{A}.\tag{2.4.17}$$

Let $(\xi, \hat{X}, \hat{\alpha})$ be the MFC solution given by Proposition 2.4.1 and let \hat{P} be its law. Let $(\lambda, \bar{\mu})$ be a mean field CCE and $(\xi, X^\lambda, \lambda)$ be the corresponding initial state, state process and correlated strategy. We show the following properties:

1. The maximum point P^* is unique and it is equal to \hat{P} .
2. For every m in the support of $\bar{\mu}$, there exists a version of the regular conditional probability of $(\xi, \lambda, X^\lambda)$ given $\bar{\mu} = m$; if we set $P^m = \mathbb{P}((\xi, \lambda, X^\lambda) \in \cdot \mid \bar{\mu} = m)$, then P^m belongs to \mathcal{A} , and it holds

$$\mathfrak{J}(\lambda, \bar{\mu}) = \int_{\mathcal{C}^d} \mathcal{J}(P^m) \rho(dm),$$

where ρ denotes the law of $\bar{\mu}$.

3. We use the above equality to show that $\bar{\mu} = \hat{x}$ \mathbb{P} -a.s. and deduce $\lambda = \hat{\alpha} d\mathbb{P} \otimes dt$ -a.e.

As for point 1, let P^* be the admissible probability measure that maximizes \mathcal{J} . Let $(\Omega^*, \mathcal{F}^*, \mathbb{F}^*, \mathbb{P}^*)$, W^* , ξ^* , α^* and X^* be so that $P^* = \mathbb{P}^* \circ (\xi^*, \alpha^*, X^*)^{-1}$. By applying Proposition 2.4.1 in this probability space, there exists an optimal control $\hat{\beta}$ which maximizes \mathfrak{J} over $\mathbb{H}^2(\mathbb{F}^*)$. Since the flow of moments of $X^{\hat{\beta}}$ is still given by (2.4.7b) and (2.4.8) admits a strong solution, we have $\mathbb{P}^* \circ (\xi^*, X^{\hat{\beta}}, \hat{\beta})^{-1} = \mathbb{P} \circ (\xi, \hat{X}, \hat{\alpha})^{-1} = \hat{P}$. Therefore, we can conclude that

$$\begin{aligned}\mathcal{J}(P^*) &= \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T (f(t, X_t^*, \mathbb{E}^{\mathbb{P}^*}[X_t^*], \alpha_t^*)) dt + g(X_T^*, \mathbb{E}^{\mathbb{P}^*}[X_T^*]) \right] \\ &\geq \mathbb{E}^{\mathbb{P}^*} \left[\int_0^T (f(t, X_t^{\hat{\beta}}, \mathbb{E}^{\mathbb{P}^*}[X_t^{\hat{\beta}}], \hat{\beta}_t)) dt + g(X_T^{\hat{\beta}}, \mathbb{E}^{\mathbb{P}^*}[X_T^{\hat{\beta}}]) \right] = \mathcal{J}(\hat{P}),\end{aligned}$$

with the inequality being strict if $\beta^* \neq \hat{\beta}$. This shows point 1.

As for point 2, we can suppose without loss of generality that $(\Omega, \mathcal{F}, \mathbb{P})$ is a Polish probability space. We note that the state process X^λ is adapted to the filtration generated by ξ , W and λ , which is countably generated. This implies that there exists a version of the regular conditional probability of \mathbb{P} given $\bar{\mu} = m$, that we denote by \mathbb{P}^m . Since ξ and W are independent of $\bar{\mu}$, it is straightforward to see that

W is a Brownian motion under \mathbb{P}^m as well, that the law of ξ under \mathbb{P}^m is η and that X^λ still satisfies equation (2.2.1). Let $P^m = \mathbb{P}^m \circ (\xi, X^\lambda, \lambda)^{-1}$ and observe that P^m belongs to \mathcal{A} for ρ -a.e. m in \mathcal{C}^d . The consistency condition implies that $\mathbb{E}^{\mathbb{P}^m}[X_t^\lambda] = m_t$ for ρ -a.e. m , which in turn implies that

$$\begin{aligned} \mathfrak{J}(\lambda, \bar{\mu}) &= \mathbb{E}^{\mathbb{P}} \left[\int_0^T (f(t, X_t^\lambda, \bar{\mu}_t, \lambda_t)) dt + g(X_T, \bar{\mu}_T) \right] \\ &= \mathbb{E}^{\mathbb{P}} \left[\mathbb{E}^{\mathbb{P}} \left[\int_0^T (f(t, X_t^\lambda, \bar{\mu}_t, \lambda_t)) dt + g(X_T, \bar{\mu}_T) \middle| \bar{\mu} \right] \right] \\ &= \int_{\mathcal{C}^d} \mathbb{E}^{\mathbb{P}^m} \left[\int_0^T (f(t, X_t^\lambda, \mathbb{E}^{\mathbb{P}^m}[X_t^\lambda], \lambda_t)) dt + g(X_T, \mathbb{E}^{\mathbb{P}^m}[X_T^\lambda]) \right] \rho(dm) \\ &= \int_{\mathcal{C}^d} \mathcal{J}(P^m) \rho(dm). \end{aligned}$$

By (2.4.14) and (2.4.17), we have

$$\int_{\mathcal{C}^d} (\mathcal{J}(\hat{P}) - \mathcal{J}(P^m)) \rho(dm) \leq 0, \quad \mathcal{J}(\hat{P}) \geq \mathcal{J}(P^m),$$

which implies $\mathcal{J}(\hat{P}) = \mathcal{J}(P^m)$ for ρ -a.e. m . Since \hat{P} is the unique maximizer of \mathcal{J} by point 1, we get $P^m = \hat{P}$ ρ -a.e. In particular, this implies

$$m_t = \mathbb{E}^{\mathbb{P}^m}[X_t^\lambda] = \int_{\mathbb{R}^d} y(P^m \circ x_t^{-1})(dy) = \int_{\mathbb{R}^d} y(\hat{P} \circ x_t^{-1})(dy) = \mathbb{E}[\hat{X}_t] = \hat{x}_t \text{ for } \rho\text{-a.e. } m.$$

Thus, $\bar{\mu}$ is a.s. equal to \hat{x} , so that the consistency condition (2.2.4) for the mean field CCE $(\lambda, \bar{\mu})$ rewrites as $\hat{x}_t = \mathbb{E}[X_t^\lambda]$. Therefore, we have

$$\mathfrak{J}(\lambda, \bar{\mu}) = \mathbb{E}^{\mathbb{P}} \left[\int_0^T (f(t, X_t^\lambda, \mathbb{E}[X_t^\lambda], \lambda_t)) dt + g(X_T, \mathbb{E}[X_T^\lambda]) \right] = \mathfrak{J}^{MFC}(\lambda).$$

Since, by Lemma 2.4.2, $\hat{\alpha}$ is unique, the previous equality implies that λ is equal to $\hat{\alpha} d\mathbb{P} \otimes dt$ -a.e., which concludes the proof. \square

Proof of (ii). Let us assume that the MFC solution $(\hat{\alpha}, \hat{x})$ is not a mean field NE. By item (i) of Theorem 2.4.3, every mean field CCE yields a lower payoff than the MFC solution; moreover, if there was a mean field CCE yielding the same payoff as the MFC solution, it would be the MFC solution itself. Therefore, we just need to prove that the MFC solution is not a mean field CCE.

The pair $(\hat{\alpha}, \hat{x})$ is a correlated moment flow which satisfies the consistency condition in the definition of mean field CCE. Moreover, since $\hat{\alpha}$ is $\mathbb{F}^{\xi, W}$ -progressively measurable and \hat{x} is deterministic, it satisfies the consistency condition of the definition of the mean field NE as well. Since by assumption the MFC solution is not a mean field NE, it is the optimality condition (2.4.1) in the definition of mean field NE which is not satisfied. Therefore, there exists $\beta \in \mathbb{A}$ so that $\mathfrak{J}(\beta, \hat{x}) > \mathfrak{J}(\hat{\alpha}, \hat{x})$. Since such β is an admissible deviation to the correlated moment flow $(\hat{\alpha}, \hat{x})$, the optimality condition (2.2.3) in the definition of mean field CCE is not satisfied either. This means that the MFC solution is not a mean field CCE. \square

2.4.2 Comparison with mean field Nash equilibria

On top of Assumptions **LQ**, we need the following additional assumptions to guarantee existence and uniqueness of the mean field NE:

Assumptions LQ-NE. Referring to the matrix valued functions introduced in Assumptions **LQ**, we assume:

- (1) $Q_t + \tilde{Q}_t \geq d_1^* I_d$, for every $t \in [0, T]$, $d_1^* \geq 0$;
- (2) $\sup_{t \in [0, T]} |S_t|^2 < d_1^* d_2$ if $d_1^* > 0$, $S_t = 0$ for every $t \in [0, T]$ otherwise;
- (3) $H + \tilde{H} \geq 0$.

We first show that the only mean field CCE with deterministic flow of moments is the mean field NE itself. In particular, this implies that randomization of the flow of moments is needed for mean field CCEs to reach higher payoffs than the mean field NE. Differently from the MFC case, no general outperformance result can be established for mean field CCEs. Instead, one can derive an outperformance condition for correlated moment flows in the class \mathcal{G} , in a similar approach as for the optimality condition in subsection 2.3.3.

As shown by the next proposition, there exists a unique mean field NE. The proof is a standard application of the Pontryagin maximum principle approach together with the fixed point argument of [38, Chapter 4]. We include it in Section 2.7 for the sake of completeness.

Proposition 2.4.4. *Let ϕ^* and θ^* the solutions of the following system:*

$$\begin{cases} \dot{\phi}_t^* + \phi_t^* A_t + A_t^\top \phi_t^* + (Q_t + \tilde{Q}_t) - (\phi_t^* B_t + S_t^\top) R_t^{-1} (B_t^\top \phi_t^* + S_t) = 0, \\ \phi_T^* = H + \tilde{H}, \\ \dot{\theta}_t^* + A_t^\top \theta_t^* + q_t - (\phi_t^* B_t + S_t^\top) R_t^{-1} (B_t^\top \theta_t^* + r_t) = 0, \\ \theta_T^* = 0. \end{cases} \quad (2.4.18)$$

Let ϕ and ψ be as in (2.3.4). There exists a unique mean field NE (α^*, \bar{m}^*) , which is given by

$$\dot{\bar{m}}_t^* = (A_t - B_t R_t^{-1} B_t^\top \phi_t^* - B_t R_t^{-1} S_t) \bar{m}_t^* - B_t R_t^{-1} (B_t^\top \theta_t^* + r_t), \quad \bar{m}_0^* = \bar{\eta}, \quad (2.4.19a)$$

$$\alpha_t^* = -R_t^{-1} ((B_t^\top \phi_t^* + S_t) X_t^* + B_t^\top \psi_t \bar{m}_t^* + B_t^\top \theta_t^* + r_t), \quad (2.4.19b)$$

where $\theta^{\bar{m}^*}$ is the solution of the following equation:

$$\dot{\theta}_t^{\bar{m}^*} - \psi_t \frac{d\bar{m}_t^*}{dt} + A_t^\top \theta_t^{\bar{m}^*} + q_t - (\phi_t B_t + S_t^\top) R_t^{-1} (B_t^\top \theta_t^{\bar{m}^*} + r_t) = 0, \quad \theta_T^{\bar{m}^*} = 0, \quad (2.4.20)$$

and X^* is the solution of equation (2.2.1) with the control α^* .

We observe that, by definition, a mean field NE is a mean field CCE with deterministic flow of measures \bar{m}^* . The converse is true as well, as shown by the following Theorem:

Theorem 2.4.5. *Let $(\lambda, \bar{\mu})$ be a mean field CCE with deterministic $\bar{\mu}$. Then, $(\lambda, \bar{\mu})$ is the mean field NE.*

Proof. We start by observing that, by using the same concavity and projections arguments as in the proof of Lemma 2.4.2, we have

$$\mathfrak{J}(\alpha^*, \bar{m}^*) > \mathfrak{J}(\beta, \bar{m}^*) \quad \forall \beta \in \mathbb{H}^2(\mathbb{F}), \beta \neq \alpha^*.$$

Let $(\lambda, \bar{\mu})$ be a coarse correlated equilibrium with deterministic flow of moments $\bar{\mu}$. Then, the consistency condition (2.2.4) becomes $\bar{\mu}_t = \mathbb{E}[X_t^\lambda]$ for every time t . By optimality, it holds $\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\beta, \bar{\mu})$ for every $\beta \in \mathbb{A}$. By reasoning as in the proof of Lemma 2.4.2, there exists a strategy $\tilde{\lambda}$ $\mathbb{F}^{\xi, W}$ -progressively measurable so that

$$\tilde{\lambda}_t = \mathbb{E}[\lambda_t | \mathcal{F}_t^{\xi, W}], \quad X_t^{\tilde{\lambda}} = \mathbb{E}[X_t^\lambda | \mathcal{F}_t^{\xi, W}], \quad \mathbb{P}\text{-a.s.}, \quad \forall t \in [0, T],$$

where $X^{\tilde{\lambda}}$ is the solution of equation (2.2.1) corresponding to the strategy $\tilde{\lambda}$. Since $\bar{\mu}$ is deterministic, by exploiting the concave linear-quadratic structure of the payoff functional \mathfrak{J} , we have

$$\mathfrak{J}(\beta, \bar{\mu}) \leq \mathfrak{J}(\lambda, \bar{\mu}) \leq \mathfrak{J}(\tilde{\lambda}, \bar{\mu}), \quad \forall \beta \in \mathbb{A}. \quad (2.4.21)$$

Since $\bar{\mu}$ is deterministic by assumption, the consistency condition holds true for the correlated moment flow $(\tilde{\lambda}, \bar{\mu})$ as well, so that (2.4.21) implies that $(\tilde{\lambda}, \bar{\mu})$ is itself a mean field NE. By uniqueness of the mean field NE, we deduce $(\tilde{\lambda}, \bar{\mu}) = (\alpha^*, \bar{m}^*)$, so that in particular $\bar{\mu} = \bar{m}^*$ \mathbb{P} -a.s. Since α^* is the unique maximizer of $\mathfrak{J}(\cdot, \bar{m}^*)$ over $\mathbb{H}^2(\mathbb{F})$, we deduce $\mathfrak{J}(\alpha^*, \bar{m}^*) \geq \mathfrak{J}(\lambda, \bar{m}^*)$. Since $(\lambda, \bar{\mu})$ is a mean field CCE by assumption, it holds $\mathfrak{J}(\alpha^*, \bar{m}^*) \leq \mathfrak{J}(\lambda, \bar{m}^*)$, which, by uniqueness, implies that $\lambda = \alpha^*$ $dt \otimes \mathbb{P}$ -a.e. as well. \square

Finally, consider again correlated moment flows $(\lambda, \bar{\mu})$ in the class \mathcal{G} . By using their specific structure as described in (2.3.6), we are able to provide a condition under which they yield a higher payoff than the mean field NE.

Theorem 2.4.6. *Let $\theta^{\bar{m}^*}$ be the solution of (2.4.20). Set*

$$\Theta_t^{\bar{m}^*} = R_t^{-1}(B_t^\top \theta_t^{\bar{m}^*} + r_t). \quad (2.4.22)$$

Let $(\lambda, \bar{\mu}) \in \mathcal{G}$ corresponding to some $\delta \in \mathbb{H}^2(\mathcal{F}_0)$. Then, $\mathfrak{J}(\lambda, \bar{\mu})$ is higher than the payoff $\mathfrak{J}(\alpha^, \bar{m}^*)$ given by the mean field NE if and only if the following inequality is satisfied:*

$$\begin{aligned} & \int_0^T \left(\frac{1}{2} (\langle \bar{Q} \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle \bar{Q} \bar{\mu}_t, \bar{\mu}_t \rangle]) + \frac{1}{2} (\langle M_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle M_t \bar{\mu}_t, \bar{\mu}_t \rangle]) \right. \\ & + \frac{1}{2} (\langle G_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle G_t \bar{\mu}_t, \bar{\mu}_t \rangle]) + \langle N_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle N_t \bar{\mu}_t, \bar{\mu}_t \rangle] \\ & + \langle L_t - q_t + R_t^{-1}(B_t^\top (\phi_t + \psi_t) + S_t)^\top r_t, \mathbb{E}[\bar{\mu}_t] - \bar{m}_t^* \rangle + \langle B_t^\top \phi_t \bar{m}_t^*, \Theta_t^{\bar{m}^*} \rangle \\ & - \mathbb{E}[\langle B_t^\top \phi_t \bar{\mu}_t, R_t^{-1} \delta_t \rangle] - \mathbb{E}[\langle B_t^\top \psi_t \bar{\mu}_t, R_t^{-1} \delta_t \rangle] \\ & + \langle B_t^\top \psi_t \bar{m}_t^*, \Theta_t^{\bar{m}^*} \rangle + \frac{1}{2} (\langle R_t \Theta_t^{\bar{m}^*}, \Theta_t^{\bar{m}^*} \rangle - \mathbb{E}[\langle R_t^{-1} \delta_t, \delta_t \rangle]) - \langle r_t, \Theta_t^{\bar{m}^*} - R_t^{-1} \mathbb{E}[\delta_t] \rangle \Big) dt \\ & + \frac{1}{2} (\langle \bar{H} \bar{m}_T^*, \bar{m}_T^* \rangle - \mathbb{E}[\langle \bar{H} \bar{\mu}_T, \bar{\mu}_T \rangle]) + \frac{1}{2} (\langle H \bar{m}_T^*, \bar{m}_T^* \rangle - \mathbb{E}[\langle H \bar{\mu}_T, \bar{\mu}_T \rangle]) \\ & + \langle \tilde{H} \bar{m}_T^*, \bar{m}_T^* \rangle - \mathbb{E}[\langle \tilde{H} \bar{\mu}_T, \bar{\mu}_T \rangle] \geq 0. \end{aligned} \quad (2.4.23)$$

The proof is similar to the one of Theorem 2.3.3. For the sake of completeness, we include it in Section 2.7. We observe that, although the inequality (2.4.23) is not easy to interpret, it involves only the law of $\bar{\mu}$ and its associated δ . Moreover, it can be verified separately from the optimality condition (2.3.12), giving some room for mean field CCEs to outperform the mean field NE payoff. This will be accomplished for the abatement game in Section 2.5.

2.5 Application to an emission abatement game

In this section we consider an emission abatement game inspired by environmental economics literature on international environmental agreements, in line with the very popular model of [12]. Previous section's findings allow us to exhibit a simple class of coarse correlated equilibria which (highly) outperforms the mean field NE in this game.

The emission abatement game has the following payoff and dynamics of the representative player state:

$$\mathfrak{J}(\alpha, \bar{\mu}) = \mathbb{E} \left[\int_0^T \left(a\bar{\mu}_t - \frac{b}{2}\bar{\mu}_t^2 - \frac{1}{2}\alpha_t^2 - \frac{\varepsilon}{2}(\bar{\mu}_t - X_t)^2 \right) dt \right], \quad (2.5.1a)$$

$$dX_t = \alpha_t dt + dW_t, \quad X_0 = \xi, \quad (2.5.1b)$$

with a, b non-negative constants and $\varepsilon > 0$. The strategy α_t represents the abatement rate of the player at time t , while X_t models the cumulated abatement over the interval $[0, t]$.

We translate a slightly modified version of the abatement game [12] into a dynamic stochastic mean field game. We follow the N -player formulation of [57] by considering symmetric players, and a normalization of the number of players is implicitly added by replacing the sum of abatement efforts by the flow of moments $\bar{\mu}$. We also add the last term in ε , inspired by further developments of this model in the literature (see [70]), which can be interpreted as a reputational cost. It appears to be necessary when one wants mean field CCEs outperforming the mean field NE at the mean field limit. Indeed, when $\varepsilon = 0$, there exists only a unique mean field CCE, corresponding to the mean field NE. This is straightforward by direct computations and can be also deduced from Proposition 2.5.1 (see upcoming Remark 2.5.2).

Following [12], the other terms of the payoff can be interpreted as follows. The term $a\bar{\mu}_t - \frac{b}{2}\bar{\mu}_t^2$, which depends solely on the mean field component $\bar{\mu}$, is the ‘‘abatement benefit’’. It represents the individual benefit of global climate change mitigation allowed by aggregate abatement efforts, with a decreasing marginal benefit. The quadratic term in the control, i.e. $-\frac{1}{2}\alpha_t^2$, is an ‘‘abatement cost’’ that the representative country privately pays for its abatement effort.

We do not claim that a mean field approximation of the abatement game of [12] is a right way to approach the problem of international environmental agreements economically. We rather use this payoff functional as a toy example that allows us to illustrate very efficiently the interest of mean field CCEs in a context of common good, and to contribute to the findings of [57].

Remark 2.5.1. Going from static to dynamic games also induces some additional assumptions that were not included in reference models [12, 57]. We chose to represent the “abatement benefit” as a running payoff rather than a terminal one, considering that environmental objectives are not only to reach a given level of emissions at a terminal time, but also to abate as much as possible, as early as possible.

2.5.1 Translation and interpretation of findings in the abatement game

In this subsection we apply the theory developed in the previous section to compute mean field CCEs in the abatement game. In the next subsection, we will make a step further and exhibit a simple but interesting subclass of correlated moment flows $(\lambda, \bar{\mu})$ which verify both the optimality inequality (2.3.12) and the NE outperformance inequality (2.4.23).

We use the setting of Section 2.2 with $d = k = 1$. The parameters are given by

$$A_t = 0, B_t = 1, \sigma_t = 1, L_t = a, \bar{Q}_t = b + \varepsilon, Q_t = \varepsilon, \tilde{Q}_t = -\varepsilon, R_t = 1, \quad \forall t \in [0, T], \quad (2.5.2)$$

and remaining parameters equal 0. According to Proposition 2.3.1 with the abatement game parameters as in (2.5.2), for a given correlated moment flow $(\lambda, \bar{\mu})$, the best deviating strategy and the corresponding state process are given by

$$\begin{aligned} \hat{\beta}_t &= \phi_t(\mathbb{E}[\bar{\mu}_t] - X_t^{\hat{\beta}}) - \theta_t, \\ dX_t^{\hat{\beta}} &= \hat{\beta}_t dt + dW_t, \quad X_0^{\hat{\beta}} = \xi, \end{aligned} \quad (2.5.3)$$

with ϕ and θ satisfying equations

$$\begin{cases} \dot{\phi}_t + \varepsilon - \phi_t^2 = 0, & \phi_T = 0, \\ \dot{\theta}_t - \phi_t \left(\theta_t + \frac{d\mathbb{E}[\bar{\mu}_t]}{dt} \right) = 0, & \theta_T = 0. \end{cases} \quad (2.5.4)$$

Note that ψ does not appear as in this case $\psi = -\phi$. We stress that, as only unilateral deviation is allowed, the deviating player can not act on the abatement benefit, and therefore does not consider a and b in her optimal strategy.

The family \mathcal{G} of correlated moment flows defined by (2.3.6) is composed of any correlated moment flow $(\lambda, \bar{\mu})$ so that:

$$\begin{aligned} \lambda_t &= \phi_t(\bar{\mu}_t - X_t) - \delta_t, \\ \dot{\bar{\mu}}_t &= -\delta_t, \quad \bar{\mu}_0 = \bar{\eta}, \end{aligned}$$

for some $\delta \in \mathbb{H}^2(\mathcal{F}_0)$ and where X solves

$$dX_t = \lambda_t dt + dW_t, \quad X_0 = \xi.$$

In particular, we note that it holds $\delta_t = -\frac{d\bar{\mu}_t}{dt}$. Therefore, in this model, the class \mathcal{G} is composed of correlated moment flows $(\lambda, \bar{\mu})$ verifying

$$\lambda_t = \phi_t(\bar{\mu}_t - X_t) + \frac{d\bar{\mu}_t}{dt}, \quad \mathbb{E} \left[\int_0^T \left(\frac{d\bar{\mu}_t}{dt} \right)^2 dt \right] < \infty. \quad (2.5.5)$$

As the correlated strategy depends on $\bar{\mu}$ itself, we remark that the state variable becomes actually mean-reverting. The extra term $\frac{d\bar{\mu}}{dt}$ allows the state of the representative player following the suggested strategy to satisfy the consistency condition by following the suggested variations of $\bar{\mu}$.

As shown by the next proposition, in the abatement game, the optimality condition only depends on the law of $\bar{\mu}$, the reputational cost parameter ε , and the final time horizon T .

Proposition 2.5.1 (Optimality condition for the abatement game). *Let $(\lambda, \bar{\mu})$ be a correlated moment flow in \mathcal{G} . Let $f(\bar{\mu}) = (f_t(\bar{\mu}))_{t \in [0, T]}$ be given by*

$$\begin{cases} \dot{f}_t(\bar{\mu}) = - \left(\phi_t (f_t(\bar{\mu}) + \bar{\mu}_t - \mathbb{E}[\bar{\mu}_t]) + \frac{d\bar{\mu}_t}{dt} + \theta_t \right), & 0 \leq t \leq T, \\ f_0(\bar{\mu}) = 0. \end{cases} \quad (2.5.6)$$

Then, $(\lambda, \bar{\mu})$ is a mean field CCE if and only if the following condition is satisfied:

$$\int_0^T \mathbb{E} \left[\left(\frac{d\bar{\mu}_t}{dt} \right)^2 \right] dt \leq \int_0^T \mathbb{E} \left[(\phi_t f_t(\bar{\mu}) + \theta_t)^2 + \phi_t^2 (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t] + f_t(\bar{\mu}))^2 + (\varepsilon - \phi_t^2) f_t(\bar{\mu})^2 \right] dt. \quad (2.5.7)$$

Proof. Referring to (2.3.7) and (2.3.10), the auxiliary functions for the abatement game are as follows:

$$\begin{aligned} \Phi_t &= \phi_t, & \Psi_t &= -\phi_t, & \Theta_t &= \theta_t, \\ M_t &= \varepsilon + \phi_t^2, & N_t &= -\varepsilon - \phi_t^2, & G_t &= \phi_t^2. \end{aligned} \quad (2.5.8)$$

This implies that $f(\bar{\mu})$ given in (2.3.11) takes the form of equation (2.5.6), recalling that $\delta_t = -d\bar{\mu}_t/dt$ by (2.5.5). After a few computations, we get that the optimality condition (2.3.12) rewrites as

$$\begin{aligned} \int_0^T \mathbb{E} [\delta_t^2] dt &\leq \int_0^T \mathbb{E} \left[(\phi_t f_t(\bar{\mu}) + \theta_t)^2 + \phi_t^2 (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t] + f_t(\bar{\mu}))^2 + (\varepsilon - \phi_t^2) f_t^2(\bar{\mu}) \right. \\ &\quad \left. + 2\varepsilon (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t]) (\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t] + f_t(\bar{\mu})) \right] dt, \end{aligned}$$

using that

$$\mathbb{E}[\bar{\mu}_t^2 - \mathbb{E}[\bar{\mu}_t]^2] = \mathbb{E}[(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t])^2].$$

Since $f_t(\bar{\mu}) = X_t^{\hat{\beta}} - X_t$ \mathbb{P} -a.s., for every time t by (2.3.14) and $f(\bar{\mu})$ is $\sigma(\bar{\mu})$ -measurable by definition, we have

$$f_t(\bar{\mu}) = \mathbb{E}[f_t(\bar{\mu}) | \bar{\mu}] = \mathbb{E}[X_t^{\hat{\beta}} - X_t | \bar{\mu}] = \mathbb{E}[X_t^{\hat{\beta}}] - \bar{\mu}_t$$

where we used the consistency condition (2.2.4) and the fact that $X^{\hat{\beta}}$ and $\bar{\mu}$ are independent. This implies that

$$\mathbb{E}[(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t])(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t] + f_t(\bar{\mu}))] = \mathbb{E}[\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t]](-\mathbb{E}[\bar{\mu}_t] + \mathbb{E}[X_t^{\hat{\beta}}]) = 0.$$

Therefore, $(\lambda, \bar{\mu})$ is a mean field CCE if and only if condition (2.5.7) is satisfied. \square

As for the mean field NE, we have $Q_t + \tilde{Q}_t \equiv 0$, $S_t \equiv 0$ and $H = \tilde{H} = 0$, so that the abatement game satisfies Assumptions **LQ-NE** with $d_1^* = 0$. By Proposition 2.4.4, there exists a unique mean field NE (α^*, \bar{m}^*) , which is given by

$$\alpha_t^* = \phi_t(\bar{m}_t^* - X_t^*), \quad (2.5.9a)$$

$$\bar{m}_t^* = \bar{\eta}, \quad \forall t \in [0, T], \quad (2.5.9b)$$

since we have $\phi^* = \theta^* = 0$ in (2.4.18), which implies $\theta^{\bar{m}^*} = 0$ in (2.4.20) as well.

The mean field NE consists, on average, to null abatement, as \bar{m}_t^* stays constant equal to its initial value. This corresponds to a free-riding equilibrium, where everybody does as little as possible, and prefers to take advantage of the others' efforts. As a result, nobody does anything.

Remark 2.5.2. One can easily see from the optimality condition in equation (2.5.7) that, if $\varepsilon = 0$, the only mean field CCE is the mean field NE. Indeed, in this case $\phi_t = \theta_t = 0$ and $f_t(\bar{\mu}) = d\bar{\mu}_t/dt$, for all $t \in [0, T]$. Hence the right-hand side term in (2.5.7) is null, forcing $d\bar{\mu}_t/dt = 0$, $t \in [0, T]$. As $\lambda_t = d\bar{\mu}_t/dt$, we get $\lambda_t = \alpha_t^* = 0$, $\bar{m}_t^* = \bar{\eta}$, which is the mean field NE when $\varepsilon = 0$. This seems consistent with the findings of [57]. Indeed, in an equivalent N -player static deterministic game without the reputational cost ($\varepsilon = 0$), the authors find that, the more players, the less the payoff-maximising CCE outperforms the payoff of the NE. This probably comes from the fact that, at the mean field limit, there is only one mean field CCE, which is the mean field Nash equilibrium itself.

Coming to the MFC problem, we have $Q_t + 2\tilde{Q}_t + \bar{Q}_t \equiv b$, $S_t \equiv 0$ and $H = \tilde{H} = \bar{H} = 0$, so Assumptions **LQ-MFC** with $\hat{d}_1 = b \geq 0$. By Proposition 2.4.1, there exists a unique MFC solution $(\hat{\alpha}, \hat{x})$ which reads:

$$\hat{\alpha}_t = \phi_t(\hat{x}_t - \hat{X}_t) - \bar{\phi}_t \hat{x}_t - \bar{\theta}_t, \quad (2.5.10a)$$

$$\dot{\hat{x}}_t = -\hat{\phi}_t \hat{x}_t - \hat{\theta}_t, \quad \hat{x}_0 = \bar{\eta}, \quad (2.5.10b)$$

with

$$\begin{cases} \dot{\hat{\phi}}_t = \hat{\phi}_t^2 - b, & \hat{\phi}_T = 0, \\ \dot{\hat{\theta}}_t = \hat{\phi}_t \hat{\theta}_t + a, & \hat{\theta}_T = 0. \end{cases} \quad (2.5.11)$$

The MFC solution adds to the mean-reversion two terms which depend on a and b , i.e. on the coefficients of the abatement benefit. One can note that the MFC solution and the mean field NE coincide if and only if $a = b = 0$. To the contrary, when the ‘‘common good’’ aspect of climate is accounted for in the payoff through the abatement benefit, the central planner can reach higher payoffs by preventing any inefficient free-riding behaviour. This gives some room for mean field CCEs to bridge the gap between the free-riding mean field NE and the central planner optimum.

To find mean field CCEs outperforming the mean field NE, the following condition should be fulfilled.

Proposition 2.5.2 (Outperformance condition over mean field NE). *Let $(\lambda, \bar{\mu})$ be a correlated moment flow in \mathcal{G} . Then, $\mathfrak{J}(\lambda, \bar{\mu}) \geq \mathfrak{J}(\alpha^*, \bar{m}^*)$ if and only if*

$$\int_0^T \mathbb{E} \left[a(\bar{\mu}_t - \bar{m}_t^*) - \frac{b}{2}(\bar{\mu}_t^2 - (\bar{m}_t^*)^2) \right] dt \geq \frac{1}{2} \int_0^T \mathbb{E} \left[\left(\frac{d\bar{\mu}_t}{dt} \right)^2 \right] dt.$$

Proof. By recalling the identities in (2.5.8), inequality (2.4.23) takes the following form:

$$\begin{aligned} \mathfrak{J}(\lambda, \bar{\mu}) - \mathfrak{J}(\alpha^*, \bar{m}^*) &= \int_0^T \left(\frac{1}{2}(b + \varepsilon)((\bar{m}_t^*)^2 - \mathbb{E}[\bar{\mu}_t^2]) + \frac{1}{2}(\varepsilon + \phi_t^2)((\bar{m}_t^*)^2 - \mathbb{E}[\bar{\mu}_t^2]) \right. \\ &\quad + \frac{1}{2}\phi_t^2((\bar{m}_t^*)^2 - \mathbb{E}[\bar{\mu}_t^2]) - (\varepsilon + \phi_t^2)((\bar{m}_t^*)^2 - \mathbb{E}[\bar{\mu}_t^2]) + a(\mathbb{E}[\bar{\mu}_t] - \bar{m}_t^*) \\ &\quad \left. - \phi_t \mathbb{E}[\bar{\mu}_t \delta_t] + \phi_t \mathbb{E}[\delta_t \bar{\mu}_t] - \frac{1}{2} \mathbb{E}[\delta_t^2] \right) \geq 0. \end{aligned}$$

By rearranging the terms and recalling that $\delta_t = -d\bar{\mu}_t/dt$, we get the desired inequality. \square

The equivalence in Proposition 2.5.2 clearly illustrates that, when $a = b = 0$, the best payoff mean field CCE is actually the mean field NE. This was also implied by the fact that, when $a = b = 0$, the MFC solution is a mean field NE as we already mention above.

2.5.2 A tractable class of mean field CCEs

In this subsection we show that, when $a \neq 0$ or $b \neq 0$, the optimality and outperformance conditions are not empty, and neither is their intersection. In this case, the MFC solution is distinct from the mean field NE, which implies, according to Theorem 2.4.3, that the MFC solution is not a mean field CCE, as the required control does not resist any unilateral deviation which tends to a less costly free-riding option. However, by introducing correlation through correlated moment flows, one can manage to drive the population at quite high abatement levels, leading to more desirable social outcomes than the one of the mean field NE.

The optimality condition (2.5.7) is very convenient to use when one focuses on a specific class of dynamics for $\bar{\mu}$. In this subsection, we consider a subclass $\mathcal{G}_l \subseteq \mathcal{G}$ where the flows of moments are linear in time.

More precisely, let \mathcal{G}_l be the set of all correlated moment flows $(\lambda, \bar{\mu}) \in \mathcal{G}$ such that

$$\bar{\mu}_t = \bar{\eta} + tZ, \quad t \in [0, T], \quad (2.5.12)$$

for some $Z \in L^2(\mathcal{F}_0)$ independent of ξ and W . Then, for all correlated moment flows $(\lambda, \bar{\mu}) \in \mathcal{G}_l$ we have

$$\lambda_t = \phi_t(\bar{\mu}_t - X_t) + Z, \quad t \in [0, T].$$

In the rest of the chapter, we will use the notations $z_1 = \mathbb{E}[Z]$, $z_2 = \mathbb{E}[Z^2]$, $\sigma_z^2 = \mathbb{V}[Z]$.

Proposition 2.5.3 (Optimality condition for \mathcal{G}_l). *Let $(\lambda, \bar{\mu}) \in \mathcal{G}_l$. Then $(\lambda, \bar{\mu})$ is a mean field CCE if and only if*

$$z_1^2 c_M + \sigma_z^2 c_V \geq 0 \quad (2.5.13)$$

with

$$c_M = \int_0^T ((\phi_t r_t + g_t)^2 + \varepsilon r_t^2) dt - T, \quad c_V = \frac{\varepsilon}{3} T^3 - T, \quad (2.5.14)$$

and

$$g_t = \int_t^T \phi_s e^{-\int_t^s \phi_u du} ds, \quad r_t = \int_0^t (1 - g_s) e^{-\int_s^t \phi_u du} ds. \quad (2.5.15)$$

Proof. For any given $(\lambda, \bar{\mu}) \in \mathcal{G}_l$ we have $\frac{d\bar{\mu}_t}{dt} \equiv Z$ a.s. and $\mathbb{E}[\bar{\mu}_t] = \bar{\eta} + tz_1$, so that

$$\begin{aligned} \int_0^T \mathbb{E} \left[\left(\frac{d\bar{\mu}_t}{dt} \right)^2 \right] dt &= Tz_2, & \dot{\theta}_t &= \phi_t (\theta_t + z_1), \quad \theta_T = 0, \\ \dot{f}_t(\bar{\mu}) &= -(\phi_t(f_t(\bar{\mu}) + t(Z - z_1)) + Z + \theta_t), & f_0(\bar{\mu}) &= 0. \end{aligned}$$

Let us set

$$p_t = \int_0^t e^{-\int_s^t \phi_u du} ds, \quad v_t = \int_0^t (s\phi_s + 1) e^{-\int_s^t \phi_u du} ds = t.$$

To see that $v_t = t$, we notice that v satisfies $\dot{v}_t + \phi_t v_t = 1 + t\phi_t$, $\phi_0 = 0$, which is solved by $v_t = t$. By using p_t , v_t and the auxiliary functions defined in (2.5.15), $f_t(\bar{\mu})$ and θ_t can be rewritten as follows:

$$\theta_t = -z_1 g_t, \quad f_t(\bar{\mu}) = -Z p_t - (Z - z_1)(v_t - p_t) + z_1(p_t - r_t).$$

We compute the different terms appearing in the integral of the right-hand side of the optimality condition:

$$\begin{aligned} \mathbb{E}[(f_t(\bar{\mu}))^2] &= \sigma_z^2 v_t^2 + z_1^2 r_t^2 = \sigma_z^2 t^2 + z_1^2 r_t^2, \\ \mathbb{E}[(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t] + f_t(\bar{\mu}))^2] &= \sigma_z^2 (v_t - t)^2 + z_1^2 r_t^2 = z_1^2 r_t^2, \\ \mathbb{E}[(\phi_t f_t(\bar{\mu}) + \theta_t)^2] &= \sigma_z^2 \phi_t^2 v_t^2 + z_1^2 (\phi_t r_t + g_t)^2 = \sigma_z^2 \phi_t^2 t^2 + z_1^2 (\phi_t r_t + g_t)^2. \end{aligned}$$

After summing, simplifying and factorising, the optimality condition becomes an inequality on the moments of Z as follows:


$$Tz_2 \leq z_1^2 \int_0^T ((\phi_t r_t + g_t)^2 + \varepsilon r_t^2) dt + \sigma_z^2 \int_0^T \varepsilon t^2 dt.$$

As $z_2 = \sigma_z^2 + z_1^2$, we get (2.5.13) and (2.5.14). \square

Thanks to this simple optimality condition, the set \mathcal{G}_l of mean field CCEs can be easily explored numerically and analytically. In Figure 2.1a we represent the running expected payoffs (time derivative of the payoff) of a mean field CCE, the MFC solution and the mean field NE as curves, and their total payoffs as dots. Figure 2.1b represents the average of the state variables at each time for the same equilibria. As one can see, the mean field CCE in the figure outperforms the mean field NE in terms of both payoff and abatement levels. Moreover, Figure 2.1b shows that this mean field CCE also outperforms the MFC solution in terms of average level of cumulated abatement at the end of the period, i.e. $\mathbb{E}[\bar{\mu}_T] > \bar{x}_T^{MFC}$.

Implications of the optimality condition for \mathcal{G}_l can be further analysed by stating some of its analytical properties.

Proposition 2.5.4. *The coefficients c_M, c_V defined in Proposition 2.5.3 verify the following:*



(a) Running expected utility and payoff of a mean field CCE, the MFC solution and the mean field NE.

(b) Average level of cumulated abatement at time t for this same mean field CCE, the MFC solution and the mean field NE.

Figure 2.1: A mean field CCE in \mathcal{G}_l bridging the gap between the mean field NE and the MFC solution. Parameter values: $T = 5$, $a = 2$, $b = 1$, $\varepsilon = 1$, $\bar{\eta} = 0.1$, $\mathbb{V}[\eta] = 1$, $z_1 = 0.6$, $\sigma_z^2 = 0.06$.

(i) $c_V > 0$ if and only if $\varepsilon T^2 > 3$.

(ii) $c_M < 0$,

Proof. (i) comes from explicit expression of c_V given in (2.5.8). As for (ii), we argue by contradiction. Suppose $c_M \geq 0$. By (2.5.13), this is equivalent to the existence of a mean field CCE in \mathcal{G}_l so that the associated random variable Z satisfies $\sigma_Z^2 = 0$ and $z_1 > 0$. Since $\sigma_Z^2 = 0$, $(\lambda, \bar{\mu})$ is a mean field CCE with deterministic flow of moments $\bar{\mu}_t = \bar{\eta} + tz_1$, for any t in $[0, T]$. By Theorem 2.4.5, this implies that $(\lambda, \bar{\mu}) = (\alpha^*, \bar{m}^*)$. Since $\bar{m}_t^* = \bar{\eta}$ for every time t , this implies that $z_1 = 0$, leading to a contradiction. \square

Proposition 2.5.4 implies that, if the reputational cost coefficient ε and time horizon T are small enough, the only mean field CCE in \mathcal{G}_l is the mean field NE. On the contrary, when T, ε are big enough, for any expectation of Z there exists a variance level so that any correlated moment flow with same expectation and higher variance is a mean field CCE.

2.5.3 Comparison with mean field NE

We have seen above that increasing the variance of Z is a way to build mean field CCEs easily. However, increasing the variance of Z comes at the cost of lowering the odds to outperform the mean field NE, as shown in the next Proposition.

Proposition 2.5.5 (Outperformance over the mean field NE in \mathcal{G}_l). *A correlated moment flow $(\lambda, \bar{\mu}) \in \mathcal{G}_l$ outperforms the mean field NE in terms of payoff if and only if*

$$Tz_1(a - b\bar{\eta}) - (z_1^2 + \sigma_z^2) \left(b\frac{T^2}{3} + 1 \right) \geq 0 \quad (2.5.16)$$

Proof. This result follows directly from Proposition 2.5.2. The inequality is assessed in the specific case of correlated moment flows in \mathcal{G}_l , using the following equalities:

$$\bar{m}_t^* = \bar{\eta}, \quad \mathbb{E}[\bar{\mu}_t] = \bar{\eta} + tz_1, \quad \mathbb{E}[\bar{\mu}_t^2] = \bar{\eta}^2 + 2\bar{\eta}tz_1 + t^2(z_1^2 + \sigma_z^2), \quad \frac{d\bar{\mu}_t}{dt} = Z, \quad t \in [0, T]. \quad (2.5.17)$$

□

The optimality and outperformance conditions for correlated moment flows in \mathcal{G}_l in, respectively, Proposition 2.5.3 and Proposition 2.5.5, are both expressed in terms of the first and second moments of associated variable Z . This allows us to characterize analytically a region of mean field CCEs outperforming the mean field NE in the plane (z_1, σ_z^2) . We first deal with the case $\varepsilon T^2 = 3$: as in this case it holds $c_V = 0$ by Proposition 2.5.3, the inequality (2.5.3) is trivially satisfied by $z_1 = 0$ and any $\sigma_z^2 \geq 0$. Thus, any random variable Z with $z_1 = 0$ would lead to a mean field CCE in the class \mathcal{G}_l . We notice that, by Proposition 2.5.5, these mean field CCEs would not lead to higher payoffs nor to higher average expected abatement, as $\mathbb{E}[\bar{\mu}_t] \equiv 0$. Therefore, in the following, we focus on the case $\varepsilon T^2 > 3$, so that the set of mean field CCEs in \mathcal{G}_l is interesting enough.

Proposition 2.5.6. *Assume $\varepsilon T^2 > 3$. Then, a correlated moment flow in \mathcal{G}_l is a mean field CCE outperforming the mean field NE in terms of payoff if and only if the associated random variable Z verifies*

$$-\frac{c_M}{c_V} z_1^2 \leq \sigma_z^2 \leq z_1 \frac{3T(a - b\bar{\eta})}{bT^2 + 3} - z_1^2. \quad (2.5.18)$$

Moreover, the set of mean field CCEs outperforming the mean field NE is not reduced to the mean field NE if and only if $a - b\bar{\eta} > 0$.

Proof. By combining Propositions 2.5.1, 2.5.4 and 2.5.5, we can see that a correlated moment flow in \mathcal{G}_l with moments z_1, σ_z^2 for Z is a mean field CCE outperforming the mean field NE in terms of expected payoff if and only if equation (2.5.18) is verified. Let us denote by $f, g : \mathbb{R} \rightarrow \mathbb{R}$ respectively the left hand-side and the right-hand side of that equation as function of z_1 . They are both parabola intersecting at the point $(0, 0)$. The second derivative of f , f'' , is strictly positive as c_V is positive according to Proposition 2.5.4, while g'' is strictly negative. Simple arguments therefore imply that the region between the two curves characterized in equation (2.5.18), i.e.,

$$\{(x, y) \in \mathbb{R} \times \mathbb{R}_+ : f(x) \leq y \leq g(x)\},$$

is not equal to the point $(0, 0)$ if and only if $g'(0) > f'(0)$. This is the case if and only if $a - b\bar{\eta} > 0$. If the region between the two curves was reduced to the point $(0, 0)$, the only mean field CCEs outperforming the mean field NE would verify $z_1 = \sigma_z^2 = 0$, which corresponds to the mean field NE. □

Figure 2.2 represents the region of mean field CCEs outperforming the payoff of the mean field NE in the plane (z_1, σ_z^2) for the same parameters as in Figure 2.1. The outperformance condition parabola (“upper parabola”) is represented in red, while the optimality condition parabola (“lower parabola”) is in blue.

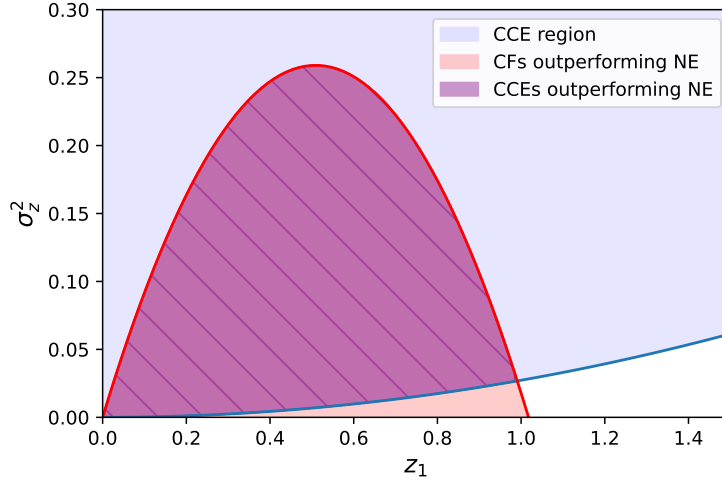


Figure 2.2: Region of mean field CCEs in \mathcal{G}_l which outperform the Nash equilibrium in the plane (z_1, σ_z^2) . “CFs” stands for correlated flows.

Proposition 2.5.6 shows that each of the parameters a , b and ε of the payoff plays specific roles in identifying mean field CCEs that outperform the payoff of the mean field NE. The upper parabola comes from the outperformance condition and only depends on a, b while the lower parabola comes from the optimality condition and only depends on ε . The existence of the abatement benefit leaves space for more correlated moment flows to outperform the free-riding equilibrium payoff, as the upper parabola increases in a and decreases in b . Moreover, Figure 2.3 shows that the ratio $-c_M/c_V$ and hence the lower parabola is decreasing in ε . Therefore, the reputational cost helps correlated moment flows to be CCEs. Indeed, with a higher reputational cost, countries have more interest in staying close to one another, and therefore the correlation device is more enforcing.




Figure 2.3: Representation of $-c_M/c_V$ in function of ε , where ε verifies $\varepsilon T^2 > 3$. The vertical dashed line represents the asymptote $\varepsilon = \frac{3}{T^2}$, corresponding to $c_V = 0$. Parameterisation: $T = 5$.

Figure 2.4 represents the payoffs of the mean field CCEs belonging to \mathcal{G}_l and which

outperform the mean field NE in terms of payoff, i.e., verifying equation (2.5.18). According to this graph, using the simple and tractable class of correlated moment flows \mathcal{G}_l , one is able to explore a large part of the payoffs attainable by mean field CCEs in this game. Indeed, mean field CCEs payoffs get pretty close to the unattainable bound provided by the MFC solution payoff, relatively to the payoff of the mean field NE.




Figure 2.4: Expected payoff of mean field CCEs in \mathcal{G}_l which outperform the mean field NE in terms of payoff (i.e. verifying equation (2.5.18)) in the 3D space with z_1 as the x-axis, σ_z^2 as the y-axis and expected payoff as the z-axis. Parameters: same as Figure 2.1.

The mean field CCEs in \mathcal{G}_l which are optimal in terms of expected payoffs can be identified analytically.

Proposition 2.5.7. *Assume $\varepsilon T^2 > 3$. Then, the expected payoff of mean field CCEs in \mathcal{G}_l is maximised by a correlated moment flow $(\lambda, \bar{\mu})$ so that the associated random variable Z satisfies*

$$z_1 = \frac{T(a - b\bar{\eta})}{2(1 - \frac{c_M}{c_V})(b\frac{T^2}{3} + 1)}, \quad \sigma_z^2 = -z_1^2 \frac{c_M}{c_V}. \quad (2.5.19)$$

Proof. We note that

$$\mathfrak{J}(\lambda, \bar{\mu}) = \mathfrak{J}(\alpha^*, \bar{m}^*) + \frac{T^2}{2} z_1 (a - b\bar{\eta}) - \frac{(\sigma_z^2 + z_1^2)}{2} \left(b\frac{T^3}{3} + T \right).$$

Therefore, $\mathfrak{J}(\lambda, \bar{\mu})$ is strictly decreasing in σ_z^2 . Moreover, since $c_V > 0$ according to Proposition 2.5.4, the optimality condition for correlated moment flows in \mathcal{G}_l of Proposition 2.5.3 implies that $(\lambda, \bar{\mu})$ is a mean field CCE if and only if

$$\sigma_z^2 \geq -z_1^2 \frac{c_M}{c_V}.$$

As $-c_M/c_V > 0$, for any given z_1 the mean field CCE with the highest expected payoff verifies $\sigma_z^2 = -z_1^2 c_M/c_V$. From now on, let us set σ_z^2 to this value. We get

$$\mathfrak{J}(\lambda, \bar{\mu}) = \mathfrak{J}(\alpha^*, \bar{m}^*) + \frac{T^2}{2} z_1 (a - b\bar{\eta}) - \frac{z_1^2 (1 - \frac{c_M}{c_V})}{2} \left(b \frac{T^3}{3} + T \right),$$

which is a polynomial in z_1 whose maximum point is given by z_1 as in (2.5.19). \square

Figure 2.5a shows the payoffs of mean field CCEs in \mathcal{G}_l with payoff-maximizing variance for Z , i.e. verifying $\sigma_z^2 = -z_1^2 c_M/c_V$. These payoffs are expressed as a function of z_1 , on the x -axis, and they are compared to the payoffs of the MFC solution and the mean field NE, with same parameter settings as in the other figures. Figure 2.5b represents the average cumulated abatement over the whole time interval for the same equilibria, in the same fashion.

One can see out of Figure 2.5 that for “little ambitious” mean field CCEs in \mathcal{G}_l , i.e., with relatively small z_1 , there is actually a significant increase in both abatement levels and payoffs with regards to the mean field NE. However, there is a critical value of z_1 , given by the payoff-maximising value of Proposition 2.5.7 and represented by the grey vertical line, after which increasing abatement comes at the cost of decreasing the payoff. This is in partial contrast with the results of [57] where a much stronger trade-off was observed. The difference is due to the presence of the reputational cost.

Characterizing a surface of mean field CCEs allows any moderator to choose the mean field CCE which corresponds to its goal, which might be to maximise expected payoff, or to consider positive externalities of abatement which are not “priced” in \mathfrak{J} , and therefore to favor high abatement over maximising payoffs.




Figure 2.5: (a) Payoff of mean field CCEs of \mathcal{G}_l with optimal variance as a function of z_1 , and (b) their average cumulated abatement $\mathbb{E}[\bar{\mu}_T]$, compared to the same quantities for the MFC solution and mean field NE. Parameter values are the same as in Figure 2.1.

2.6 A mean field game that does not admit mean field NEs

In this section, we consider the linear quadratic MFG described in [84, Section 7]. The author showed therein that, up to a suitable choice of the parameters, this MFG does not admit a mean field NE. The goal of this section is to show that, for the same choice of the parameters, there exist infinitely many mean field CCEs, as anticipated in Remark 2.3.1. We repeat the reasoning of Section 2.5: after applying the results of Section 2.3 to this specific MFG, we consider again the class of linear-in-time flow of moments. We are then able to prove that there exists infinitely many mean field CCEs in this class of correlated moment flows.

The game considered in [84, Section 7] is as follows: let $\mu = (\mu_t)_{t \in [0, T]}$ be a deterministic a flow of measures with values in $\mathcal{P}(\mathbb{R})$. The representative player's dynamics evolve accordingly to under the constraint

$$dX_t = \alpha_t dt + dW_t, \quad X_0 = \xi,$$

and the payoff, to be maximized, is given by

$$\mathfrak{J}(\alpha, \mu) = \mathbb{E} \left[- \int_0^T \alpha_t^2 dt - (X_T - c\bar{\mu}_T)^2 \right]. \quad (2.6.1)$$

The notion of solution considered in [84] is the notion of MFG solution as in Definition 1.3.5 in Chapter 1. Notice that, given the linear-quadratic structure of the MFG, there exists a MFG solution in the sense of Definition 1.3.5 if and only if there exists a mean field NE in the sense of Definition 2.4.1. It is shown in [84, Section 7] that, if the parameters satisfy

$$c = \frac{1+T}{T}, \quad \bar{\eta} \neq 0, \quad (2.6.2)$$

then there exists no MFG solution. Moreover, this MFG does not admit any weak MFG solution without common noise in the sense of [85, Definition 3.1] either. To see this fact, it is enough to start with a weak MFG solution, so that in particular the flow of measures μ is stochastic, and repeat the arguments of [84, Section 7]. Exploiting the independence of ξ , W and μ and by taking the expectation of $\mathbb{E}[\bar{\mu}_t]$ in place of $\bar{\mu}_t$, we get to the same contradiction as for the MFG solution case.

The MFG above fits the setting of Section 2.2 with $d = k = 1$. The parameters are given by

$$B_t = 1, \quad \sigma_t = 1, \quad R_t = 2, \quad H = 2, \quad \tilde{H} = -2c, \quad \bar{H} = 2c^2$$

and the remaining parameters being equal to 0. Notice that Assumption **LQ** are satisfied, while Assumptions **LQ-NE** are violated, since $H + \tilde{H} = -\frac{2}{T} < 0$ for any $T > 0$.

According to Proposition 2.3.1, for a given correlated moment flow $(\lambda, \bar{\mu})$, the best deviating strategy and the corresponding state process are given by

$$\hat{\beta}_t = -\frac{1}{2}(\phi_t X_t^{\hat{\beta}} + \psi_t \mathbb{E}[\bar{\mu}_t] + \theta_t), \quad (2.6.3)$$

where ϕ , ψ and θ satisfy the following ODEs:

$$\begin{cases} \dot{\phi}_t = \frac{1}{2}\phi_t^2, & \phi_T = 2, \\ \dot{\theta}_t = \frac{1}{2}\phi_t\theta_t - \psi_t\frac{d\mathbb{E}[\bar{\mu}_t]}{dt}, & \theta_T = 0. \end{cases} \quad (2.6.4)$$

Note that ψ does not appear as $\psi = -\phi$. Moreover, ϕ is explicitly given by $\phi_t = \frac{2}{1+(T-t)}$. The family \mathcal{G} of correlated moment flows defined by (2.3.6) is composed of any correlated moment flow $(\lambda, \bar{\mu})$ so that:

$$\begin{aligned} \lambda_t &= -\frac{1}{2}(\phi_t X_t + \psi_t \bar{\mu}_t + \delta_t), \\ \dot{\bar{\mu}}_t &= -\frac{1}{2}(\phi_t + \psi_t)\bar{\mu}_t - \frac{1}{2}\delta_t, \quad \mu_0 = \bar{\eta}, \end{aligned} \quad (2.6.5)$$

for some $\delta \in \mathbb{H}^2(\mathcal{F}_0)$, where X solves

$$dX_t = \lambda_t dt + dW_t, \quad X_0 = \xi.$$

By Lemma 2.3.2, any correlated moment flow $(\lambda, \bar{\mu}) \in \mathcal{G}$ satisfies the consistency condition. Analogously to Proposition 2.5.1, we have the following expression for the optimality condition:

Proposition 2.6.1. *Let $(\lambda, \bar{\mu})$ be a correlated moment flow in \mathcal{G} . Then, (λ, μ) is a mean field CCE if and only if the following inequality is satisfied:*

$$\begin{aligned} & \frac{1}{4} \int_0^T (\mathbb{E}[(\delta_t + (\phi_t + \psi_t)\bar{\mu}_t)^2] - (\phi_t + \psi_t)^2 \mathbb{E}[\bar{\mu}_t^2] + (\psi_t^2 + 2\phi_t\psi_t)(\mathbb{E}[(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t])^2])) dt \\ & - 2c(\mathbb{E}[(\bar{\mu}_T - \mathbb{E}[\bar{\mu}_T])^2]) \leq \frac{1}{4} \int_0^T (\phi_t^2 (\mathbb{E}[f_t^2(\bar{\mu})] + 2\mathbb{E}[f_t(\bar{\mu})\bar{\mu}_t]) + 2\phi_t\psi_t\mathbb{E}[\bar{\mu}_t]\mathbb{E}[f_t(\bar{\mu})] \\ & + \theta_t^2 + 2\phi_t\theta_t(\mathbb{E}[f_t(\bar{\mu})] + \mathbb{E}[\bar{\mu}_t]) + 2\psi_t\theta_t\mathbb{E}[\bar{\mu}_t]) dt + \mathbb{E}[f_T^2(\bar{\mu})] + 2\mathbb{E}[f_T(\bar{\mu})\bar{\mu}_T] \\ & - 2c\mathbb{E}[\bar{\mu}_T]\mathbb{E}[f_T(\bar{\mu})], \end{aligned} \quad (2.6.6)$$

where $f(\bar{\mu}) = (f_t(\bar{\mu}))_{t \in [0, T]}$ as the solution of the following equation:

$$\dot{f}_t(\bar{\mu}) = -\frac{1}{2}\phi_t f_t(\bar{\mu}) - \frac{1}{2}\psi_t(\mathbb{E}[\bar{\mu}_t] - \bar{\mu}_t) - \frac{1}{2}(\theta_t - \delta_t), \quad f_0(\bar{\mu}) = 0, \quad (2.6.7)$$

Proof. Referring to (2.3.7) and (2.3.10), the auxiliary functions for the MFG are as follows:

$$\Phi_t = \frac{1}{2}\phi_t, \quad \Psi_t = \frac{1}{2}\psi_t, \quad \Theta_t = \frac{1}{2}\theta_t, \quad M_t = \frac{1}{2}\phi_t^2, \quad N_t = \frac{1}{2}\phi_t\psi_t, \quad G_t = \frac{1}{2}\psi_t^2.$$

Then, the optimality condition (2.3.12) becomes

$$\begin{aligned} & \int_0^T \left(\frac{1}{2}\phi_t\psi_t\mathbb{E}[(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t])^2] + \frac{1}{4}\psi_t^2(\mathbb{E}[\bar{\mu}_t^2] - \mathbb{E}[\bar{\mu}_t]^2) + \frac{1}{4}\mathbb{E}[\delta_t^2] - \frac{1}{4}\theta_t^2 \right) dt \\ & - 2c\mathbb{E}[(\bar{\mu}_T - \mathbb{E}[\bar{\mu}_T])^2] \leq \int_0^T \left(\frac{1}{4}\phi_t^2(\mathbb{E}[(\bar{\mu}_t + f_t(\bar{\mu}))^2] - \mathbb{E}[\bar{\mu}_t^2]) + \frac{1}{2}\phi_t\psi_t\mathbb{E}[f_t(\bar{\mu})\bar{\mu}_t] \right. \\ & \left. + \frac{1}{2}(\phi_t + \psi_t)\theta_t\mathbb{E}[\bar{\mu}_t] - \frac{1}{2}(\phi_t + \psi_t)\mathbb{E}[\bar{\mu}_t\delta_t] + \frac{1}{2}\phi_t\theta_t\mathbb{E}[\bar{\mu}_t] \right) dt \\ & + \frac{1}{2}(2\mathbb{E}[(\bar{\mu}_T + f_T^2(\bar{\mu}))^2] - 2\mathbb{E}[\bar{\mu}_T^2]) - 2c\mathbb{E}[f_T(\bar{\mu})]\mathbb{E}[\bar{\mu}_T]. \end{aligned}$$

By rearranging the terms of the previous inequality, we get to (2.6.6). \square

As in Section 2.5.2, we look for conditions on δ so that equation (2.6.6) becomes more tractable. Let $\tilde{\mathcal{G}}_l$ be the set of all correlated moment flows $(\lambda, \bar{\mu}) \in \mathcal{G}$ such that

$$\delta_t = -2Z - (\phi_t + \psi_t)\bar{\mu}_t, \quad (2.6.8)$$

for some $Z \in L^2(\mathcal{F}_0)$ independent of ξ and W . Then, we have

$$\frac{d\bar{\mu}_t}{dt} = -\frac{1}{2}(\phi_t + \psi_t)\bar{\mu}_t - \frac{1}{2}(-2Z - (\phi_t + \psi_t)\bar{\mu}_t) = Z, \quad \bar{\mu}_0 = \bar{\eta}. \quad (2.6.9)$$

Set $z_1 = \mathbb{E}[Z]$, $\sigma_Z^2 = \mathbb{V}[Z]$, $z_2 = \mathbb{E}[Z^2]$.

Lemma 2.6.2. *Let $(\lambda, \bar{\mu})$ be a correlated moment flow in $\tilde{\mathcal{G}}_l$. Then, condition (2.6.6) is satisfied if and only if*

$$c_V \sigma_Z^2 + c_{M,2} z_1^2 + c_{M,1} z_1 + c_0 \geq 0, \quad (2.6.10)$$

where c_V , $c_{M,2}$, $c_{M,1}$ and c_0 are deterministic constants depending only on T and c .

The proof of this Lemma is deferred to Section 2.7. In order to show the existence of a mean field CCE, it is enough to focus on the sign of the coefficient c_V , as ensured by the following Proposition:

Proposition 2.6.3. *Let c and T be as in (2.6.2). Let $(\lambda, \bar{\mu})$ be a correlated moment flow in $\tilde{\mathcal{G}}_l$. Consider the coefficients c_V , $c_{M,1}$, $c_{M,2}$ and c_0 given by Lemma 2.6.2.*

(i) *For any $z_1 \in \mathbb{R}$, it holds*

$$c_{M,2} z_1^2 + c_{M,1} z_1 + c_0 < 0.$$

In particular, $c_0 < 0$.

(ii) *A mean field CCE exists if and only if $c_V > 0$. In this case, any random variable Z so that $z_1 = 0$ and $\sigma_Z^2 \geq -\frac{c_0}{c_V}$ defines a mean field CCE in the class $\tilde{\mathcal{G}}_l$.*

Proof. To prove point (i), we do the same reasoning as in Proposition 2.5.4: suppose that such a value z_1 exists, and consider Z deterministic equal to z_1 . By Proposition 2.6.1, the correlated moment flow $\tilde{\mathcal{G}}_l$ associated to such Z is a mean field CCE with deterministic flow of moments. In particular, this implies that λ is $\mathbb{F}^{\xi, W}$ -measurable, implying $\lambda \in \mathbb{A}$. Therefore, such a correlated moment flow $(\lambda, \bar{\mu})$ is a mean field NE. Since there does not exist any mean field NE by [84, Section 7], this leads to a contradiction.

To see (ii), notice that, if $c_V = 0$, then any random variable Z so that $z_1 = 0$ and $\sigma_Z^2 > 0$ defines a mean field CCE in the class $\tilde{\mathcal{G}}_l$. If instead $c_V > 0$, then, any $(\lambda, \bar{\mu})$ in $\tilde{\mathcal{G}}_l$ with $z_1 = 0$ and $\sigma_Z^2 \geq -\frac{c_0}{c_V}$ defines a mean field CCE. Suppose now that $Z \in L^2(\mathcal{F}_0)$ with mean z_1 and variance σ_Z^2 yields a mean field CCE. By Lemma 2.6.2 we have

$$0 \leq c_V \sigma_Z^2 + c_{M,2} z_1^2 + c_{M,1} z_1 + c_0 \leq c_V \sigma_Z^2,$$

where last inequality holds true by point (i), which implies $c_V \geq 0$. \square

Lemma 2.6.4. *The coefficient c_V is given by*

$$c_V = T^2(2c - 1) - T. \quad (2.6.11)$$

The proof of the Lemma is postponed to Section 2.7. We go back to the framework of [84, Section 7]. Recall that, by (2.6.2), $c = \frac{1+T}{T} = 1 + \frac{1}{T}$. Then, we have

$$c_V = T^2 \left(2\left(1 + \frac{1}{T}\right) - 1 \right) - T = T^2 + T, \quad (2.6.12)$$

which is always positive. Thus, for any choice of $T > 0$ there exist infinitely many mean field CCEs for this MFG.

2.7 Auxiliary results: Some standard proofs

In this section, we gather the proofs of some results. Although they are very standard, we include them for reader's convenience.

Proof of Proposition 2.3.1. We follow the approach of [111, Chapter 6]. We start by noticing that the equation for ϕ is a matrix Riccati equation, which admits a unique solution $\phi \in \mathcal{C}^1([0, T], \mathcal{S}^d)$ by Chapter 6, Theorem 7.2 therein. This implies the existence and uniqueness for ψ and θ as well as they satisfy linear ODEs.

First, thanks to Assumptions **LQ**(4), the cost functional \mathfrak{J}' is strictly convex and therefore has a unique minimizer. Indeed, by looking at (2.3.3), we have

$$\begin{aligned} & \frac{1}{2} \langle Q_t X_t, X_t \rangle + \frac{1}{2} \langle R_t \beta_t, \beta_t \rangle + \langle S_t X_t, \beta_t \rangle + \langle \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] + q_t, X_t \rangle + \langle r_t, \beta_t \rangle \\ & \geq \frac{1}{2} d_1 |X_t|^2 + \frac{1}{2} d_2 |\beta_t|^2 - \sup_t |S_t| |X_t| |\beta_t| + \langle \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] + q_t, X_t \rangle + \langle r_t, \beta_t \rangle \\ & > \frac{1}{2} d_1 |X_t|^2 + \frac{1}{2} d_2 |\beta_t|^2 - \sqrt{d_1} \sqrt{d_2} |S_t| |X_t| |\beta_t| + \langle \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] + q_t, X_t \rangle + \langle r_t, \beta_t \rangle. \end{aligned}$$

This inequality and the assumption $H \geq 0$ imply that the cost functional is strictly convex and lower semicontinuous, which yields that the minimizer exists and it is unique. Observe that this holds for any $(\mathbb{E}[\bar{\mu}_t])_{t \in [0, T]}$, since it appears only in the linear terms $\langle \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] + q_t, X_t \rangle$ and $\langle r_t, \beta_t \rangle$.

Let $\mathcal{H}(t, x, y, \beta)$ be the the reduced Hamiltonian of the system, defined as

$$\begin{aligned} \mathcal{H}(t, x, y, \beta) &= \langle A_t x + B_t \beta, y \rangle + \frac{1}{2} \langle Q_t x, x \rangle + \langle \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] + q_t, x \rangle \\ & \quad + \langle S_t x, \beta \rangle + \frac{1}{2} \langle R_t \beta, \beta \rangle + \langle r_t, \beta \rangle. \end{aligned}$$

By the stochastic maximum principle (see [111, Chapter 6, Corollary 5.7]), there exists a unique 4-tuple $(X, \hat{\beta}, Y, Z)$ solution to

$$\begin{cases} dX_t = (A_t X_t + B_t \hat{\beta}_t) dt + \sigma_t dW_t, & X_0 = \xi, \\ dY_t = - \left(A_t^\top Y_t + Q_t X_t + \tilde{Q}_t \mathbb{E}[\bar{\mu}_t] + q_t + S_t^\top \hat{\beta}_t \right) dt + Z_t dW_t, & Y_T = H X_T + \tilde{H} \mathbb{E}[\bar{\mu}_T], \\ B_t^\top Y_t + S_t X_t + R_t \hat{\beta}_t + r_t = 0, \end{cases} \quad (2.7.1)$$

and the process $\hat{\beta} = (\hat{\beta}_t)_{t \in [0, T]}$ is the unique optimal control. To explicitly find the solution of the FBSDEs system above, we make the following ansatz on Y :

$$Y_t = \phi_t X_t + \psi_t \mathbb{E}[\bar{\mu}_t] + \theta_t,$$

with ϕ , ψ and θ deterministic functions taking values in $\mathbb{R}^{d \times d}$, $\mathbb{R}^{d \times d}$ and \mathbb{R}^d respectively. Since R is invertible for every time t by assumption **LQ**(4), by comparing the stochastic differential of the ansatz with (2.7.1), we find that $Z_t = \phi_t \sigma$ and that ϕ , ψ and θ must satisfy equations (2.3.4). In particular, the matrix Riccati equation for ϕ admits a unique solution in $\mathcal{C}([0, T]; \mathcal{S}^d)$ by [111, Chapter 6, Theorem 7.2]. This implies the existence and uniqueness for ψ and θ as well. \square

Proof of Proposition 2.4.1. We follow the Pontryagin maximum principle approach for linear-quadratic MFC problems of [38, Section 6.7.1] (see also [69, Section 2]). We recall that, by (2.4.11) and Assumptions **LQ-MFC**, the payoff functional \mathfrak{J}^{MFC} is upper semi-continuous and strictly concave jointly in X_t , $X_t - \mathbb{E}[X_t]$ and α_t so that the MFC solution exists and it is unique. Let \mathcal{H} be the Hamiltonian of the system:

$$\begin{aligned} \mathcal{H}(t, x, y, m, \alpha) = & \langle A_t x + B_t \alpha, y \rangle - \frac{1}{2} \langle \bar{Q}_t m, m \rangle + \langle L_t, m \rangle - \left(\frac{1}{2} \langle Q_t x, x \rangle + \langle \tilde{Q}_t m, x \rangle \right. \\ & \left. + \langle q_t, x \rangle + \langle S_t x, \alpha \rangle + \frac{1}{2} \langle R_t \alpha, \alpha \rangle + \langle r_t, \alpha \rangle \right), \end{aligned} \quad (t, x, y, m, \alpha) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^k.$$

Under Assumptions **LQ** and **LQ-MFC**, the MFC control problem verifies assumptions **Control of MVK Dynamics** in [38, Section 6.4.1, p. 555], which ensure that the following mean field FBSDE system

$$\begin{cases} dX_t = (A_t X_t + B_t \hat{\alpha}_t) dt + \sigma_t dW_t, \\ X_0 = \xi, \\ dY_t = - \left(A_t^\top Y_t - Q_t X_t - S_t^\top \hat{\alpha}_t - (\bar{Q}_t + 2\tilde{Q}_t) \mathbb{E}[X_t] + L_t - q_t \right) dt + Z_t dW_t, \\ Y_T = -(H X_T + (\bar{H} + 2\tilde{H}) \mathbb{E}[X_T]), \\ B_t^\top Y_t - S_t X_t - R_t \hat{\alpha}_t - r_t = 0, \end{cases} \quad (2.7.2)$$

admits a unique solution $(X, \hat{\alpha}, Y, Z)$ and the control $\hat{\alpha} = (\hat{\alpha}_t)_{t \in [0, T]}$ is optimal. To explicitly find the solution, we first set $x_t = \mathbb{E}[X_t]$, $y_t = \mathbb{E}[Y_t]$ and $a_t = \mathbb{E}[\hat{\alpha}_t]$. Then, by taking expectation, we get the following system

$$\begin{cases} \dot{x}_t = A_t x_t + B_t a_t, & x_0 = \mathbb{E}[\xi], \\ \dot{y}_t = - \left(A_t^\top y_t - (Q_t + 2\tilde{Q}_t + \bar{Q}_t) x_t - S_t^\top a_t + L_t - q_t \right), & y_T = -(H + 2\tilde{H} + \bar{H}) x_T, \\ B_t^\top y_t - S_t x_t - R_t a_t - r_t = 0. \end{cases} \quad (2.7.3)$$

By relying on the deterministic LQ control problem associated with the system (2.7.3) and applying the results of [111, Chapter 6, Section 2] (see in particular Corollary 2.10 therein), we get that the solution (x, y, a) to (2.7.3) is unique. To explicitly find it, we make the following ansatz on y :

$$y_t = -\hat{\phi}_t x_t - \hat{\theta}_t,$$

with $\hat{\phi}$ and $\hat{\theta}$ suitable deterministic functions taking values in $\mathbb{R}^{d \times d}$ and \mathbb{R}^d respectively. Since R is invertible for every time t by assumption **LQ**(4), by comparing the differential of the ansatz with (2.7.3), we get to equations (2.4.4). Thanks to Assumptions **LQ-MFC**, there exists a unique solution for the matrix Riccati equation for $\hat{\phi} \in \mathcal{C}([0, T], \mathcal{S}^d)$, by [111, Chapter 6, Theorem 7.2]. We note that the flow of expectations $\hat{x} = (\hat{x}_t)_{t \in [0, T]}$ satisfies the differential equation (2.4.7b).

To prove the existence of a solution to the forward backward system (2.7.2), we can make the ansatz

$$Y_t = -\bar{\phi}_t X_t - \bar{\psi}_t \hat{x}_t - \bar{\theta}_t.$$

with $\bar{\phi}$, $\bar{\psi}$ and $\bar{\theta}$ deterministic functions taking values in $\mathbb{R}^{d \times d}$, $\mathbb{R}^{d \times d}$ and \mathbb{R}^d respectively. By differentiating the ansatz, comparing it with (2.7.2) and using the invertibility of R for any time t , we find that $Z_t = -\bar{\phi}_t \sigma$ and that $\bar{\phi}$ satisfies the same equation as ϕ , so that $\bar{\phi} = \phi$, and equations (2.4.6) must be satisfied by $\bar{\psi}$ and $\bar{\theta}$. \square

Proof of Proposition 2.4.4. We follow the Pontryagin maximum principle approach together with the fixed point argument of [38, Section 3.5]. Let \mathcal{H} be the Hamiltonian of the system:

$$\begin{aligned} \mathcal{H}(t, x, y, m, \alpha) = & \langle A_t x + B_t \alpha, y \rangle - \frac{1}{2} \langle \bar{Q}_t m, m \rangle + \langle L_t, m \rangle - \left(\frac{1}{2} \langle Q_t x, x \rangle + \langle \tilde{Q}_t m, x \rangle \right. \\ & \left. + \langle q_t, x \rangle + \langle S_t x, \alpha \rangle + \frac{1}{2} \langle R_t \alpha, \alpha \rangle + \langle r_t, \alpha \rangle \right), \end{aligned} \quad (t, x, y, m, \alpha) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^k.$$

For a fixed flow of moments $\hat{m} = (\hat{m}_t)_{t \in [0, T]}$, the following FBSDEs system

$$\begin{cases} dX_t = (A_t X_t + B_t \hat{\alpha}_t) dt + \sigma_t dW_t, & X_0 = \xi, \\ dY_t = - \left(A_t^\top Y_t - Q_t X_t - \tilde{Q}_t \hat{m}_t - q_t - S_t^\top \hat{\alpha}_t \right) dt + Z_t dW_t, & Y_T = -(H X_T + \tilde{H} \hat{m}_T), \\ B_t^\top Y_t - S_t X_t - R_t \hat{\alpha}_t - r_t = 0, \end{cases} \quad (2.7.4)$$

admits a unique solution $(X, \hat{\alpha}, Y, Z)$ and the control $\hat{\alpha} = (\hat{\alpha}_t)_{t \in [0, T]}$ is optimal for the LQ control problem with fixed \hat{m} (see again [111, Chapter 6, Corollary 5.7]).

We now impose the fixed-point condition. Set $x_t = \mathbb{E}[X_t]$, $y_t = \mathbb{E}[Y_t]$ and $a_t = \mathbb{E}[\hat{\alpha}_t]$. Then, by [38, Theorem 3.34], there exists a unique mean field NE if and only if the following forward-backward ODE system

$$\begin{cases} \dot{x}_t = A_t x_t + B_t a_t, & x_0 = \mathbb{E}[\xi], \\ \dot{y}_t = - \left(A_t^\top y_t - (Q_t + \tilde{Q}_t) x_t - q_t - S_t^\top a_t \right), & y_T = -(H + \tilde{H}) x_T, \\ B_t^\top y_t - S_t x_t - R_t a_t - r_t = 0, \end{cases} \quad (2.7.5)$$

admits a unique solution. Under Assumptions **LQ** and **LQ-NE**, this is ensured by [111, Chapter 6, Section 2], by relying on the deterministic LQ control problem associated with the system (2.7.5) as in the proof of Proposition 2.4.1. We explicitly find such solution by making the ansatz

$$y_t = -\phi_t^* x_t - \theta_t^*,$$

with ϕ^* and θ^* suitable deterministic functions taking values in $\mathbb{R}^{d \times d}$ and \mathbb{R}^d respectively. Since R is invertible for every time t by assumption **LQ**(4), by comparing the differential of the ansatz with (2.7.5), we get to equations (2.4.18). Assumptions **LQ-NE** guarantee that the Riccati equation in (2.4.18) admits a unique solution $\phi^* \in \mathcal{C}([0, T], \mathcal{S}^d)$ by invoking [111, Chapter 6, Theorem 7.2]. We note that the flow of moments $\bar{m}^* = (\bar{m}_t^*)_{t \in [0, T]}$ satisfies the differential equation (2.4.19a).

The last step it to prove the existence of a solution to the forward backward system (2.7.4). We make the ansatz

$$Y_t = -\tilde{\phi}_t X_t - \tilde{\psi}_t \bar{m}_t^* - \theta_t^{\bar{m}^*},$$

with $\tilde{\phi}$, $\tilde{\psi}$ and $\theta^{\bar{m}^*}$ deterministic functions taking values in $\mathbb{R}^{d \times d}$, $\mathbb{R}^{d \times d}$ and \mathbb{R}^d respectively. By the same reasoning of Proposition 2.3.1, we find that $Z_t = -\tilde{\phi}_t \sigma$, that $\tilde{\phi}$ and $\tilde{\psi}$ satisfy the same equations as ϕ and ψ , so that $\tilde{\phi} = \phi$ and $\tilde{\psi} = \psi$, and that equation (2.4.20) must be satisfied by $\theta^{\bar{m}^*}$. \square

Proof of Theorem 2.4.6. By using $(\phi, \psi, \theta^{\bar{m}^*})$, we write the dynamics of the state process X^* as

$$\begin{aligned} dX_t^* &= ((A_t - B_t R_t^{-1} (B_t^\top \phi_t + S_t)) X_t^* - B_t R_t^{-1} (B_t^\top \psi_t \bar{m}_t^* + B_t^\top \theta_t^{\bar{m}^*} + r_t)) dt + \sigma_t dW_t \\ &= ((A_t - B_t \Phi_t) X_t^* - B_t (\Psi_t \bar{m}_t^* + \Theta_t^{\bar{m}^*})) dt + \sigma_t dW_t, \end{aligned}$$

with Φ and Ψ defined by (2.3.10) and $\Theta^{\bar{m}^*}$ by (2.4.22). We remark that $\theta^{\bar{m}^*}$ and thus $\Theta^{\bar{m}^*}$ depend on \bar{m}^* through its time derivative $d\bar{m}^*/dt$. Let $f^{\bar{m}^*, \bar{\mu}} = (f_t^{\bar{m}^*, \bar{\mu}})_{t \in [0, T]}$ be the solution of

$$\dot{f}_t^{\bar{m}^*, \bar{\mu}} = (A_t - B_t \Phi_t) f_t^{\bar{m}^*, \bar{\mu}} + B_t (\Psi_t (\bar{m}_t^* - \bar{\mu}_t) + \Theta_t^{\bar{m}^*} - \delta_t), \quad f_0^{\bar{m}^*, \bar{\mu}} = 0.$$

By Itô's formula, we have that

$$f_t^{\bar{m}^*, \bar{\mu}} = X_t - X_t^*.$$

Since $f(\bar{\mu})$ is $\sigma(\bar{\mu})$ -measurable and X^* is $\mathbb{F}^{\xi, W}$ -progressively measurable, we have both that X^* and $f^{\bar{m}^*, \bar{\mu}}$ are independent and that

$$f_t^{\bar{m}^*, \bar{\mu}} = \mathbb{E}[f_t^{\bar{m}^*, \bar{\mu}} | \bar{\mu}] = \mathbb{E}[X_t - X_t^* | \bar{\mu}] = \bar{\mu}_t - \bar{m}_t^*, \quad (2.7.6)$$

by consistency condition. Since it holds

$$\begin{aligned} \mathfrak{J}(\lambda, \bar{\mu}) - \mathfrak{J}(\alpha^*, \bar{m}^*) &= \int_0^T \left(\frac{1}{2} \langle \bar{Q}_t \bar{m}_t^*, \bar{m}_t^* \rangle - \frac{1}{2} \mathbb{E}[\langle \bar{Q}_t \bar{\mu}_t, \bar{\mu}_t \rangle] + \langle L_t, \mathbb{E}[\bar{\mu}_t] - \bar{m}_t^* \rangle \right) dt \\ &\quad + \frac{1}{2} \langle \bar{H} \bar{m}_T^*, \bar{m}_T^* \rangle - \frac{1}{2} \mathbb{E}[\langle \bar{H} \bar{\mu}_T, \bar{\mu}_T \rangle] + \mathfrak{J}'(\alpha^*, \bar{m}^*) - \mathfrak{J}'(\lambda, \bar{\mu}), \end{aligned}$$

we focus on the difference $\mathfrak{J}'(\alpha^*, \bar{m}^*) - \mathfrak{J}'(\lambda, \bar{\mu})$. In a very similar way as in the proof

of Theorem 2.3.3 we obtain

$$\begin{aligned}
\mathfrak{J}'(\alpha^*, \bar{m}^*) - \mathfrak{J}'(\lambda, \bar{\mu}) &= \int_0^T \left(\frac{1}{2} (\langle M_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle M_t(\bar{m}_t^* + f_t^{\bar{m}^*, \bar{\mu}}), \bar{m}_t^* + f_t^{\bar{m}^*, \bar{\mu}} \rangle]) \right. \\
&+ \frac{1}{2} (\langle G_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle G_t \bar{\mu}_t, \bar{\mu}_t \rangle]) + \langle N_t \bar{m}_t^*, \bar{m}_t^* - \mathbb{E}[\bar{\mu}_t] \rangle - \mathbb{E}[\langle N_t f_t^{\bar{m}^*, \bar{\mu}}, \bar{\mu}_t \rangle] \\
&- \langle q_t - \Phi_t^\top r_t, \mathbb{E}[f_t^{\bar{m}^*, \bar{\mu}}] \rangle - \langle \Psi_t^\top r_t, \bar{m}_t^* - \mathbb{E}[\bar{\mu}_t] \rangle + \langle (R_t \Phi_t - S_t) \bar{m}_t^*, (\Theta_t^{\bar{m}^*} - R_t^{-1} \mathbb{E}[\delta_t]) \rangle \\
&- \mathbb{E}[\langle (R_t \Phi_t - S_t) f_t^{\bar{m}^*, \bar{\mu}}, R_t^{-1} \delta_t \rangle] - \mathbb{E}[\langle R_t \Psi_t \bar{\mu}_t, R_t^{-1} \delta_t \rangle] + \langle R_t \Psi_t \bar{m}_t^*, \Theta_t^{\bar{m}^*} \rangle \\
&+ \frac{1}{2} (\langle R_t \Theta_t^{\bar{m}^*}, \Theta_t^{\bar{m}^*} \rangle - \mathbb{E}[\langle R_t^{-1} \delta_t, \delta_t \rangle]) - \langle r_t, \Theta_t^{\bar{m}^*} - R_t^{-1} \mathbb{E}[\delta_t] \rangle \Big) dt + \frac{1}{2} \langle H \bar{m}_T^*, \bar{m}_T^* \rangle \\
&- \frac{1}{2} \langle H(\bar{m}_T^* + \hat{f}_T(\bar{\mu})), \bar{m}_T^* + \hat{f}_T(\bar{\mu}) \rangle + \langle \tilde{H} \bar{m}_T^*, \bar{m}_T^* - \mathbb{E}[\bar{\mu}_T] \rangle - \mathbb{E}[\langle \tilde{H} \hat{f}_T(\bar{\mu}), \bar{\mu}_T \rangle].
\end{aligned}$$

Finally, we observe that, by using (2.7.6), we have

$$\begin{aligned}
\langle M_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle M_t(\bar{m}_t^* + f_t^{\bar{m}^*, \bar{\mu}}), \bar{m}_t^* + f_t^{\bar{m}^*, \bar{\mu}} \rangle] &= \langle M_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle M_t \bar{\mu}_t, \bar{\mu}_t \rangle], \\
\langle N_t \bar{m}_t^*, \bar{m}_t^* - \mathbb{E}[\bar{\mu}_t] \rangle - \mathbb{E}[\langle N_t f_t^{\bar{m}^*, \bar{\mu}}, \bar{\mu}_t \rangle] &= \langle N_t \bar{m}_t^*, \bar{m}_t^* \rangle - \mathbb{E}[\langle N_t \bar{\mu}_t, \bar{\mu}_t \rangle], \\
\langle (R_t \Phi_t - S_t) \bar{m}_t^*, (\Theta_t^{\bar{m}^*} - R_t^{-1} \mathbb{E}[\delta_t]) \rangle - \mathbb{E}[\langle (R_t \Phi_t - S_t) f_t^{\bar{m}^*, \bar{\mu}}, R_t^{-1} \delta_t \rangle] \\
&= \langle (R_t \Phi_t - S_t) \bar{m}_t^*, \Theta_t^{\bar{m}^*} \rangle - \mathbb{E}[\langle (R_t \Phi_t - S_t) \bar{\mu}_t, R_t^{-1} \delta_t \rangle], \\
\langle q_t - \Phi_t^\top r_t, \mathbb{E}[f_t^{\bar{m}^*, \bar{\mu}}] \rangle + \langle \Psi_t^\top r_t, \bar{m}_t^* - \mathbb{E}[\bar{\mu}_t] \rangle &= \langle q_t - (\Phi_t + \Psi_t)^\top r_t, \mathbb{E}[\bar{\mu}_t] - \bar{m}_t^* \rangle.
\end{aligned}$$

By using these identities together with (2.3.15), we get to (2.4.23). \square

Proof of Lemma 2.6.2. Referring to (2.6.6), we compute separately the integrals in the LHS and in the RHS, as well as the terminal conditions in the LHS and RHS. We show that each term can be expressed as a quadratic function of σ_Z^2 and z_1 , which implies (2.6.10).

We start by expressing $\bar{\mu}$ and δ in terms of Z :

$$\bar{\mu}_t = \bar{\eta} + tZ, \quad \delta_t = -Z(2 + t(\phi_t + \psi_t)) - \bar{\eta}(\phi_t + \psi_t),$$

so that the integrand in the LHS of (2.6.6) is given by

$$\begin{aligned}
&\mathbb{E}[(\delta_t + (\phi_t + \psi_t) \bar{\mu}_t)^2] - (\phi_t + \psi_t)^2 \mathbb{E}[\bar{\mu}_t^2] + (\psi_t^2 + 2\phi_t \psi_t) (\mathbb{E}[(\bar{\mu}_t - \mathbb{E}[\bar{\mu}_t])^2]) \\
&= \sigma_Z^2 (t^2 \phi_t^2 (c^2 - 2c) + 4 - (1 - c)^2 t^2 \phi_t^2) + z_1^2 (4 - (1 - c)^2 t^2 \phi_t^2) \\
&- z_1 (2\bar{\eta}(1 - c)^2 t \phi_t^2) - \bar{\eta}^2 (1 - c)^2 \phi_t^2.
\end{aligned}$$

Therefore, the LHS of (2.6.6) can be written as

$$\begin{aligned}
&\sigma_Z^2 \left(\frac{1}{4} \int_0^T (t^2 \phi_t^2 (c^2 - 2c) - (1 - c)^2 t^2 \phi_t^2) dt + T - 2cT^2 \right) \\
&+ z_1^2 \left(\frac{1}{4} \int_0^T (4 - (1 - c)^2 t^2 \phi_t^2) dt \right) + z_1 \left(-\frac{1}{4} \int_0^T (2\bar{\eta}(1 - c)^2 t \phi_t^2) dt \right) \\
&- \left(\frac{1}{4} \int_0^T (\bar{\eta}^2 (1 - c)^2 \phi_t^2) dt \right). \tag{2.7.7}
\end{aligned}$$

We turn our attention to the RHS of (2.6.6). Consider the auxiliary functions

$$\begin{aligned}
g_t &= \int_t^T \phi_s e^{-\frac{1}{2} \int_t^s \phi_u du} ds, & v_t &= \int_0^t s \phi_s e^{-\frac{1}{2} \int_s^t \phi_u du} ds, & r_t &= \int_0^t g_s e^{-\frac{1}{2} \int_s^t \phi_u du}, \\
k_t &= \int_0^t \phi_s e^{-\frac{1}{2} \int_s^t \phi_u du} ds, & q_t &= \int_0^t e^{-\frac{1}{2} \int_s^t \phi_u du} ds \\
p_t &= \int_0^t \left(1 + \frac{1-c}{2} s \phi_s\right) e^{-\frac{1}{2} \int_s^t \phi_u du} ds = q_t + \frac{1-c}{2} v_t.
\end{aligned} \tag{2.7.8}$$

We rewrite the equations for θ and $f(\bar{\mu})$ (2.6.4) and (2.6.7) as

$$\begin{aligned}
\dot{\theta}_t - \frac{1}{2} \phi_t \theta_t &= -\psi_t z_1, & \theta_T &= 0, \\
\dot{f}_t(\bar{\mu}) + \frac{1}{2} \phi_t f_t(\bar{\mu}) &= \frac{c}{2} t \phi_t (z_1 - Z) - Z \left(1 + \frac{1-c}{2} t \phi_t\right) - \frac{\bar{\eta}(1-c)}{2} \phi_t + z_1 \frac{c}{2} g_t, & f_0 &= 0.
\end{aligned}$$

so that

$$\begin{aligned}
\theta_t &= z_1 \int_t^T \psi_s e^{-\frac{1}{2} \int_t^s \phi_u du} ds = -z_1 c g_t, \\
f_t(\bar{\mu}) &= \frac{c}{2} (z_1 - Z) \int_0^t s \phi_s e^{-\frac{1}{2} \int_s^t \phi_u du} ds - Z \int_0^t \left(1 + \frac{1-c}{2} s \phi_s\right) e^{-\frac{1}{2} \int_s^t \phi_u du} ds \\
&\quad - \frac{\bar{\eta}(1-c)}{2} \int_0^t \phi_s e^{-\frac{1}{2} \int_s^t \phi_u du} ds + z_1 \frac{c}{2} \int_0^t g_s e^{-\frac{1}{2} \int_s^t \phi_u du} ds.
\end{aligned}$$

By using the expressions for $\bar{\mu}$, δ , $f(\bar{\mu})$ and θ , we are able to express each product, square and expectation appearing in the RHS in terms of z_1 and σ_Z^2 , so that it becomes

$$\begin{aligned}
&\frac{\sigma_Z^2}{4} \int_0^T \phi_t^2 \left(\frac{c^2}{4} v_t^2 + c v_t p_t + p_t^2 - c t v_t - 2 t p_t \right) dt \\
&+ \frac{z_1^2}{4} \int_0^T \left[\phi_t^2 \left(\frac{c^2}{4} r_t^2 - c p_t r_t + p_t^2 + c t r_t - 2 t p_t \right) - 2 c \phi_t^2 t \left(\frac{c}{2} r_t - p_t \right) \right. \\
&+ c^2 g_t^2 - 2 \phi_t c g_t \left(\frac{c}{2} r_t - p_t + t \right) + 2 c^2 \phi_t t g_t \left. \right] dt \\
&+ \frac{z_1}{4} \int_0^T \left[\phi_t^2 \left(\bar{\eta}(1-c) p_t k_t - \frac{\bar{\eta} c (1-c)}{2} k_t r_t + \bar{\eta} c r_t - 2 \bar{\eta} p_t - \bar{\eta}(1-c) t k_t \right) \right. \\
&- 2 c \phi_t^2 \left(\bar{\eta} \left(\frac{c}{2} r_t - p_t \right) - \frac{\bar{\eta}(1-c)}{2} t k_t \right) - 2 \phi_t c g_t \left(\bar{\eta} - \frac{\bar{\eta}(1-c)}{2} k_t \right) + 2 \bar{\eta} c^2 \phi_t g_t \left. \right] dt \\
&+ \frac{1}{4} \int_0^T \left[\phi_t^2 \left(\frac{\bar{\eta}^2 (1-c)^2}{4} k_t^2 - \bar{\eta}^2 (1-c) k_t \right) + 2 c \phi_t^2 \frac{\bar{\eta}^2 (1-c)}{2} k_t \right] dt.
\end{aligned} \tag{2.7.9}$$

Analogous manipulations lead to express the terminal term in the RHS of (2.6.6) as

$$\begin{aligned}
& \sigma_Z^2 \left(\frac{c^2}{4} v_T^2 + c v_T p_T + p_T^2 - c T v_T - 2 T p_T \right) + z_1^2 \left(\bar{\eta}(1-c) p_T k_T - \frac{\bar{\eta} c(1-c)}{2} k_T r_T \right. \\
& \quad \left. + \bar{\eta} c r_T - 2 \bar{\eta} p_T - \bar{\eta}(1-c) T k_T - c^2 T r_T + 2 c T p_T \right) \\
& + z_1 \left(\bar{\eta}(1-c) p_T k_T - \frac{\bar{\eta} c(1-c)}{2} k_T r_T + \bar{\eta} c r_T - 2 \bar{\eta} p_T - \bar{\eta}(1-c) T k_T \right. \\
& \quad \left. + \bar{\eta}(-c^2 r_T + 2 c p_T) + \bar{\eta} c(1-c) T k_T \right) \\
& + \left(\frac{\bar{\eta}^2(1-c)^2}{4} k_T^2 - \bar{\eta}^2(1-c) k_T + \bar{\eta}^2 c(1-c) k_T \right).
\end{aligned} \tag{2.7.10}$$

This is enough to conclude that Lemma 2.6.2 holds true. \square

Proof of Lemma 2.6.4. By (2.7.7), (2.7.9) and (2.7.10), the coefficient multiplying σ_Z^2 is given by

$$\begin{aligned}
c_V &= \frac{1}{4} \int_0^T \phi_t^2 \left(\frac{c^2}{4} v_t^2 + c v_t p_t + p_t^2 - c t v_t - 2 t p_t - t^2(c^2 - 2c) + (1-c)^2 t^2 \right) dt \\
& + 2cT^2 - T + \left(\frac{c^2}{4} v_T^2 + c v_T p_T + p_T^2 - c T v_T - 2 T p_T \right).
\end{aligned}$$

Since $p_t = q_t + \frac{1-c}{2} v_t$ by (2.7.8), we have:

$$\begin{aligned}
& \frac{c^2}{4} v_t^2 + c v_t p_t + p_t^2 - c t v_t - 2 t p_t - t^2(c^2 - 2c) + (1-c)^2 t^2 \\
& = \left(q_t + \frac{1}{2} v_t \right)^2 - t v_t - 2 t q_t + t^2 = \frac{1}{4} h_t^2 - t h_t + t^2 = \left(\frac{1}{2} h_t - t \right)^2,
\end{aligned}$$

where we have set

$$h_t = v_t + 2q_t = \int_0^t (s\phi_s + 2) e^{-\frac{1}{2} \int_s^t \phi_u du} ds.$$

Analogously, we have

$$\frac{c^2}{4} v_T^2 + c v_T p_T + p_T^2 - c T v_T - 2 T p_T = \frac{1}{4} h_T^2 - T h_T.$$

We notice that h is the solution to the following ODE:

$$\dot{h}_t + \frac{1}{2} \phi_t h_t = t \phi_t + 2, \quad h_0 = 0, \tag{2.7.11}$$

which implies $h_t = 2t$. Therefore, we have

$$\begin{aligned}
c_V &= \frac{1}{4} \int_0^T \left(\phi_t^2 \left(\frac{1}{2} h_t - t \right)^2 \right) dt + 2cT^2 - T + \frac{1}{4} h_T^2 - T h_T \\
& = 2cT^2 - T + T^2 - 2T^2 = T^2(2c - 1) - T.
\end{aligned}$$

\square

Chapter 3

Stationary mean field games of singular control

We consider ergodic symmetric N -player and mean field games of singular control in both competitive and cooperative settings. The state process dynamics of a representative player follow geometric Brownian motion, controlled additively through a non-decreasing process. Agents aim to maximize a long-time average reward functional with instantaneous profit of power type. The game shows strategic complementarities, in that the marginal profit function is increasing with respect to the dynamic average of the states of the other players, when $N < \infty$, or with respect to the stationary mean of the players' distribution, in the mean field case. In the mean field formulation, we explicitly construct Nash and coarse correlated equilibria (with singular and regular recommendations), as well as the solution to the mean field control problem associated with central planner optimization. We show that coarse correlated equilibria may exist even when Nash equilibria do not. Additionally, we show that a coarse correlated equilibrium with a regular (absolutely continuous) recommendation can outperform a Nash equilibrium where the equilibrium policy is of reflecting type (thus singularly continuous). Furthermore, we prove that the constructed mean field equilibria and mean field control can approximate the competitive and cooperative equilibria, respectively, in the corresponding N -player game when N is sufficiently large.

3.1 The N -player game

Let $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \geq 0-}, \mathbb{P})$ be a filtered probability space satisfying the usual assumptions. Let $(W^i)_{i \geq 1}$, W be a sequence of independent \mathbb{F} -Brownian motions, and let ξ , $(\xi^i)_{i \geq 1}$ be a sequence of i.i.d. positive random variables with distribution $\eta \in \mathcal{P}(\mathbb{R}_+)$. We assume that they are independent from W and $(W^i)_{i \geq 1}$, and they are \mathcal{F}_0 -measurable.

We consider the following set of strategies, to be subject to further restrictions in the following:

$$\mathcal{S} := \{ \nu : \Omega \times \mathbb{R}_+ \rightarrow \mathbb{R}_+, \mathbb{F}\text{-adapted and such that } t \mapsto \nu_t \text{ is a.s.} \\ \text{non-decreasing, right-continuous, } \nu_{0-} = 0 \text{ and } \mathbb{E}[\nu_T] < \infty \forall T > 0 \}.$$

Let $N \geq 2$. We denote a vector of strategies $(\nu^1, \dots, \nu^N) \in \mathcal{S}^N$ by ν^N . We refer to $\nu^N \in \mathcal{S}^N$ as a strategy profile.

Let δ, σ be in \mathbb{R}_+ . For any strategy profile $\nu \in \mathcal{S}^N$, we consider the following dynamics:

$$dX_t^{\nu^i} = -\delta X_t^{\nu^i} dt + \sigma X_t^{\nu^i} dW_t^i + d\nu_t^i, \quad X_{0^-}^{\nu^i} = \xi^i, \quad (3.1.1)$$

for any $i = 1, \dots, N$. Observe that, for any $\nu^N \in \mathcal{S}^N$, there exists a unique strong solution to (3.1.1) (see, e.g., [100, Theorem 7, Chapter V]). Actually, one has

$$X_t^{\nu^i} = X_t^{i,0} \left(\xi^i + \int_0^t \frac{d\nu_s^i}{X_s^{i,0}} \right),$$

where $X^0 = (X^{1,0}, \dots, X^{N,0})$ denotes the uncontrolled solution of (3.1.1), i.e. the one associated to $\nu^i \equiv 0$, and so that $X_{0^-}^{i,0} \equiv 1$. Moreover, for any $i = 1, \dots, N$, we define the flow of empirical averages of players $j \neq i$ by

$$\bar{\mu}_t^{N, \nu^{-i, N}} := \frac{1}{N-1} \sum_{j \neq i} X_t^{\nu^j}, \quad t \geq 0^-.$$

Let α, β in $(0, 1)$, $q > 0$. Each player is associated with the following reward functional:

$$\mathfrak{J}^N(\nu^i, \nu^{-i, N}) = \liminf_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\nu^i})^\alpha (\bar{\mu}_t^{N, \nu^{-i, N}})^\beta dt - q\nu_T^i \right], \quad (3.1.2)$$

which can possibly be infinite. Occasionally, we use the notation $\pi(x, m) = x^\alpha m^\beta$, for any $(x, m) \in \mathbb{R}_+^2$, and we write $\pi_x(x, m) = \partial_x \pi(x, m)$ and analogously $\pi_m(x, m) = \partial_m \pi(x, m)$.

When dealing with N -player games, we consider open-loop strategies. Roughly speaking, we allow each player to observe the noises of all players, as well as their initial position. To this extent, denote by $\mathbb{F}^N = (\mathcal{F}_t^N)_{t \geq 0^-}$ be the \mathbb{P} -augmentation of the filtration generated by the Brownian motions $(W^i)_{i=1}^N$ and initial data $(\xi^i)_{i=1}^N$.

Definition 3.1.1 (Open-loop strategies for the N -player game). We say that a process $\nu \in \mathcal{S}$ is an open-loop strategy for the N -player game if ν is \mathbb{F}^N -progressively measurable. We denote the set of open-loop strategies for the N -player game by \mathbb{A}_N .

We are interested in different kinds of equilibria in the N -player system. We deal both with the competitive framework and the cooperative case.

Competitive framework

We consider the notion of coarse correlated equilibria in the N -player game, whose definition in the ergodic N -player game is specified below.

As for linear-quadratic MFGs in Chapter 2, we assume a structural condition on the σ -algebra \mathcal{F}_{0^-} , which guarantees that \mathcal{F}_{0^-} is large enough to support the extra randomization that the mediator uses to run his lottery:

Assumption U. The σ -algebra \mathcal{F}_{0-} is large enough to support a \mathcal{F}_{0-} -measurable uniform random variable independent of the initial data ξ , $(\xi^i)_{i \geq 1}$ and the noises W , $(W^i)_{i \geq 1}$.

Next, we introduce correlation between players' strategies.

Definition 3.1.2 (Correlating device). A correlation device is any random variable $Z : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ so that Z is \mathcal{F}_{0-} -measurable and independent of ξ , W , $(\xi^i)_{i \geq 1}$ and $(W^i)_{i \geq 1}$.

The correlation device Z models the moderator's lottery over players' strategy profiles, and it explicitly represents the extra stochasticity introduced by the moderator's to randomize and correlate players' strategies. We further discuss this definition in Remark 3.2.2.

Definition 3.1.3 (Correlated strategy profile). We define a correlated strategy profile as a pair $(Z, \boldsymbol{\lambda})$ so that

- (i) Z is a correlation device;
- (ii) $\boldsymbol{\lambda} = (\lambda^i)_{i=1}^N$ is an admissible recommendation to the N players; that is, for each $i = 1, \dots, N$, $\lambda^i = (\lambda_t^i)_{t \geq 0-}$ belongs to \mathcal{S} and it is progressively measurable with respect to the \mathbb{P} -augmentation of the filtration $(\sigma(Z) \vee \mathcal{F}_t^N)_{t \geq 0-}$.

We now assign dynamics and rewards of each player. To do so, we must distinguish two cases: Suppose that each player i follows the recommendation λ^i . Then, players' state dynamics are given by (3.1.1), and each player gets the reward $\mathfrak{J}^N(\lambda^i, \boldsymbol{\lambda}^{-i})$, with \mathfrak{J}^N given by (3.1.2). Suppose now player i deviates, while other players stick to the correlated strategy profile $\boldsymbol{\lambda}^{-i}$. Then, the deviating player will pick instead an open loop strategy $\nu \in \mathbb{A}_N$. Her dynamics are given by (3.1.1) and her reward is given by $\mathfrak{J}^N(\nu^i, \boldsymbol{\lambda}^{-i})$, with \mathfrak{J}^N given by (3.1.2).

Definition 3.1.4 (ε -coarse correlated equilibrium within the set of strategies \mathcal{B}). Let $\varepsilon \geq 0$, $\mathcal{B} \subseteq \mathbb{A}_N$. A correlated strategy profile $(Z, \boldsymbol{\lambda})$ is an ε -coarse correlated equilibrium (ε -CCE) of the ergodic N -player game within the set of strategies \mathcal{B} , if for any $i = 1, \dots, N$, we have

$$\mathfrak{J}^N(\lambda^i, \boldsymbol{\lambda}^{-i}) \geq \mathfrak{J}^N(\nu, \boldsymbol{\lambda}^{-i}) - \varepsilon, \quad \forall \nu \in \mathcal{B}.$$

If $\varepsilon = 0$, we say that the correlated strategy profile $(Z, \boldsymbol{\lambda})$ is a coarse correlated equilibrium (CCE) of the ergodic N -player game within the set of strategies \mathcal{B} .

We interpret correlation devices and the correlated strategy profiles as in Chapter 1: A correlation device or a mediator runs a lottery over the set of open loop strategy profiles. The extraction of the strategy happens before the game starts and it is independent of the of the initial data and idiosyncratic shocks that determine the random evolution of players' states. These features are captured by the fact that the random variable Z is \mathcal{F}_{0-} -measurable and it is independent of $(\xi^i)_{i \geq 1}$ and $(W^i)_{i \geq 1}$. The correlation device Z introduces extra randomness in the game, but is not exogenous in the sense of a *common noise* (see, e.g., [41]): Indeed, it is picked by the moderator

as part of her recommendation to the players. Finally, notice that the existence of the correlation device Z is guaranteed by Assumption **U**.

We interpret deviations in the usual way: Each player must decide whether to commit to moderator's lottery before the extraction happens, only by relying on the information given by the law of the correlated strategy $(Z, \boldsymbol{\lambda})$, which is assumed to be common knowledge between the players. If a player does not commit, she will pick a strategy without any information on the outcome of the extraction. Notice that, since the deviating player has only knowledge of the law of the correlated strategy profile $(Z, \boldsymbol{\lambda})$, she will use a strategy $\nu \in \mathbb{A}_N$, which is, in particular, independent of the correlation device Z ; consequently, her state process is independent of Z .

The definition of ε -coarse correlated equilibria for the N -player game extends the one of Nash equilibria in open loop strategies, that we recall in this specific framework:

Definition 3.1.5 (ε -Nash equilibrium within the set of strategies \mathcal{B}). Let $\varepsilon \geq 0$, $\mathcal{B} \subseteq \mathbb{A}_N$. A strategy profile $\boldsymbol{\nu}^* = (\nu^{1,*}, \dots, \nu^{N,*}) \in \mathcal{B}^N$ is an ε -Nash equilibrium (ε -NE) of the ergodic N -player game within the set of strategies \mathcal{B} , if for any $i = 1, \dots, N$, we have

$$\mathfrak{J}^N(\nu^{i,*}, \boldsymbol{\nu}^{-i,*}) \geq \mathfrak{J}^N(\nu, \boldsymbol{\nu}^{-i,*}) - \varepsilon, \quad \forall \nu \in \mathcal{B}.$$

If $\varepsilon = 0$, we say that the strategy profile $\boldsymbol{\nu}^*$ is a Nash equilibrium (NE) of the ergodic N -player game within the set of strategies \mathcal{B} .

Every ε -CCE $(Z, \boldsymbol{\lambda})$ with deterministic correlation device Z is an ε -NE: It is enough to notice that, since Z is deterministic, the correlated strategy profile $\boldsymbol{\lambda}$ reduces to an open-loop strategy profile (ν^1, \dots, ν^N) in \mathbb{A}_N^N . Conversely, any ε -NE induces an ε -CCE with deterministic correlation device.

Cooperative framework

In the cooperative case, we look for Pareto efficient strategy profiles, according to the following definition:

Definition 3.1.6 (Pareto efficiency). Let $\mathcal{C} \subseteq \mathbb{A}_N^N$. A strategy profile $\widehat{\boldsymbol{\nu}} \in \mathcal{C}$ is Pareto efficient in the class \mathcal{C} if there does not exist any other $\boldsymbol{\nu} \in \mathcal{C}$ so that

$$\begin{aligned} \mathfrak{J}^N(\boldsymbol{\nu}^j, \boldsymbol{\nu}^{-j}) &\geq \mathfrak{J}^N(\widehat{\boldsymbol{\nu}}^j, \widehat{\boldsymbol{\nu}}^{-j}), \quad \forall j = 1, \dots, N, \\ \mathfrak{J}^N(\boldsymbol{\nu}^i, \boldsymbol{\nu}^{-i}) &> \mathfrak{J}^N(\widehat{\boldsymbol{\nu}}^i, \widehat{\boldsymbol{\nu}}^{-i}), \quad \text{for some } i. \end{aligned}$$

In other words, a strategy profile is Pareto efficient in \mathcal{C} if there does not exist any other strategy profile in \mathcal{C} which makes each player at least as well off and one player strictly better off. To search for Pareto efficient strategy profiles, we associate to the dynamics (3.1.1) and payoff functionals (3.1.2) an N -dimensional control problem. We consider the following functional

$$\bar{\mathfrak{J}}^N(\boldsymbol{\nu}) := \frac{1}{N} \sum_{i=1}^N \mathfrak{J}^N(\nu^i, \boldsymbol{\nu}^{-i}), \quad \boldsymbol{\nu} \in \mathcal{S}_N^N, \quad (3.1.3)$$

which can be regarded as a welfare utility for the N -player system.

Definition 3.1.7. Let $\varepsilon \geq 0$, $\mathcal{C} \subseteq \mathbb{A}_N^N$. A strategy profile $\hat{\nu} = (\hat{\nu}^1, \dots, \hat{\nu}^N) \in \mathcal{C}$ is ε -optimal for the central planner optimization problem within the set of strategy profiles \mathcal{C} if

$$\tilde{\mathfrak{J}}^N(\hat{\nu}) \geq \tilde{\mathfrak{J}}^N(\nu) - \varepsilon, \quad \forall \nu \in \mathcal{C}.$$

If $\varepsilon = 0$, we say that the strategy profile $\hat{\nu}$ is optimal for the central planner within the set of strategy profiles \mathcal{C} .

When dealing with the central planner's optimization problem, the players are referred to as agents, since there is no competition between them: the central planner picks herself a strategy for each player in order to maximize the welfare utility functional $\tilde{\mathfrak{J}}^N$. As a consequence, agents are not allowed to unilaterally deviate from the central planner's strategy profile. It can be easily show that if a strategy profile is an optimum of the central planner maximization problem, it is Pareto efficient as well.

3.2 The ergodic mean field game

In order to determine ε -optimal solutions to the central planner problem and ε -equilibria in the competitive setting, we consider the mean field counterparts of the optimization problem and game considered before. We will then show in Sections 3.4 and 3.5 the relation between mean field solutions to the N -player competitive and cooperative problems respectively.

We work on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ defined in the previous section. Given a strategy $\nu \in \mathcal{S}$, we consider the following dynamics:

$$dX_t^\nu = -\delta X_t^\nu dt + \sigma X_t^\nu dW_t + d\nu_t, \quad X_{0^-} = \xi. \quad (3.2.1)$$

For any \mathcal{F}_{0^-} -measurable non-negative random variable \bar{m} , possibly degenerate, we consider the following payoff functional to be maximized:

$$\tilde{\mathfrak{J}}(\nu, \bar{m}) = \liminf_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^\nu)^\alpha \bar{m}^\beta dt - q\nu_T \right]. \quad (3.2.2)$$

Let $\mathbb{F}^{\xi, W} = (\mathcal{F}_t^{\xi, W})_{t \geq 0^-}$ be the \mathbb{P} -augmentation of the filtration generated by ξ and W . Analogously to Definition 3.1.1, we consider the following strategies:

Definition 3.2.1 (Open-loop strategy for the ergodic MFG). A process $\nu \in \mathcal{S}$ is an open-loop strategy if it is progressively measurable with respect to the filtration $\mathbb{F}^{\xi, W}$. We denote the set of open-loop strategies by \mathbb{A} .

Competitive framework

We consider the mean field analogues of CCEs and NEs in the N -player game. To this extent, we give the following definition:

Definition 3.2.2 (Correlated stationary strategy). A correlated stationary strategy is a triple $(Z, \lambda, \bar{\mu}_\infty)$ satisfying the following properties:

- (i) Z is a correlation device;

- (ii) $\lambda = (\lambda_t)_{t \geq 0^-}$ belongs to \mathcal{S} and it is progressively measurable with respect to the \mathbb{P} -augmentation of the filtration $(\sigma(Z) \vee \mathcal{F}_t^{\xi, W})_{t \geq 0^-}$;
- (iii) $\bar{\mu}_\infty$ is a $\sigma(Z)$ -measurable non-negative random variable.

In the following, we will denote the law of $\bar{\mu}_\infty$ by $\rho \in \mathcal{P}(\mathbb{R}_+)$.

Let $(Z, \lambda, \bar{\mu}_\infty)$ be a correlated stationary strategy. We now assign dynamics and payoff functional. We distinguish the following two cases: If the representative player decides to trust the mediator and so to follow her recommendation λ , then the dynamics is given by (3.2.1) with λ instead of ν and the payoff is given by $\mathfrak{J}(\lambda, \bar{\mu}_\infty)$, with \mathfrak{J} defined by (3.2.2). If instead the representative player chooses to deviate, she uses a strategy $\nu \in \mathbb{A}$, her dynamics is given by (3.2.1), and she gets the reward $\mathfrak{J}(\nu, \bar{\mu}_\infty)$. Observe that, when the representative player deviates, her strategy ν is $\mathbb{F}^{\xi, W}$ -progressively measurable and therefore independent of $\bar{\mu}_\infty$, since she has no information on the outcome of moderator's lottery. As in the N -player game, the deviating player can only use her knowledge of the law of the correlated stationary strategy $(Z, \lambda, \bar{\mu}_\infty)$, which is assumed to be publicly known. Nevertheless, $\bar{\mu}_\infty$ still appears in her payoff.

Definition 3.2.3 (Coarse correlated equilibrium for the ergodic MFG). A correlated stationary triple $(Z, \lambda, \bar{\mu}_\infty)$ is a coarse correlated equilibrium (CCE) for the ergodic MFG if the following holds:

- (i) Optimality: for any deviation $\nu \in \mathbb{A}$, it holds

$$\mathfrak{J}(\lambda, \bar{\mu}_\infty) \geq \mathfrak{J}(\nu, \bar{\mu}_\infty).$$

- (ii) Consistency: the process X^λ admits a stationary distribution and it holds

$$\bar{\mu}_\infty = \int_{\mathbb{R}_+} x p_\infty(dx, \bar{\mu}_\infty), \quad (3.2.3)$$

where p_∞ is the stochastic kernel so that $\theta_\infty(dx, dm) = p_\infty(dx, m)\rho(dm)$ with $\rho = \mathbb{P} \circ \bar{\mu}_\infty^{-1}$ and $\theta_\infty = \lim_{t \rightarrow \infty} \mathbb{P} \circ (X_t^\lambda, \bar{\mu}_\infty)^{-1}$ in the weak sense.

We will refer to CCEs for the ergodic MFG as mean field CCEs as well.

Remark 3.2.1. Property (ii) in Definition 3.2.3 is equivalent to

$$\bar{\mu}_\infty \sim w - \lim_{t \rightarrow \infty} \mathbb{E}[X_t^\lambda | \bar{\mu}_\infty]. \quad (3.2.4)$$

The consistency condition (ii) in Definition 3.2.3 should be read in the usual way: the mediator imagines what the stationary mean $\bar{\mu}_\infty$ will be, before the game starts, and gives a recommendation to each player according to her idea. Since $\bar{\mu}_\infty$ is expected to be stochastic only as the result of the mediator's randomization, we request it to be measurable with respect to the correlation device Z that the moderator uses to generate both the recommendation λ and the random stationary mean $\bar{\mu}_\infty$ itself. If all players commit to the mediator's lottery for generating recommendations, then the long-time average should be consistent with what imagined by the mediator.

The notion of CCE for the ergodic MFG extends the notion of Nash equilibrium for the ergodic MFG, that we borrow from [31]:

Definition 3.2.4 (Nash equilibrium of the ergodic MFG). A pair $(\nu^*, m_\infty^*) \in \mathbb{A} \times \mathbb{R}_+$ is said to be a Nash equilibrium of the ergodic MFG if

- (i) Optimality: ν^* maximizes $\mathfrak{J}(\cdot, \bar{m}_\infty^*)$ over \mathbb{A} , i.e.

$$\mathfrak{J}(\nu^*, m_\infty^*) = \max_{\nu \in \mathbb{A}} \mathfrak{J}(\nu, m_\infty^*).$$

- (ii) Consistency: the optimally controlled process X^{ν^*} admits a limiting distribution $p_\infty^* \in \mathcal{P}(\mathbb{R}_+)$ satisfying

$$m_\infty^* = \int_{\mathbb{R}_+} xp_\infty^*(dx).$$

We will refer to NEs for the ergodic MFG as mean field NEs as well. We stress that, differently from mean field CCEs, when looking for Nash equilibria, m_∞^* is assumed to be deterministic.

As in the N -player game, and actually by exactly the same reasoning, every CCE for the ergodic MFG $(Z, \bar{\mu}_\infty, \lambda)$ with deterministic correlation device Z is an Nash equilibrium for the ergodic MFG as well. Conversely, any Nash equilibrium for the ergodic MFG (ν^*, m_∞^*) induces a mean field CCE with deterministic correlation device.

Remark 3.2.2. As one may notice, the notion of CCE used in this chapter is somehow more restrictive than the one used in the previous chapter, as the correlation device is modeled through a real valued random variable. This is motivated by the desire of a more explicit description of the extra source of stochasticity introduced by the moderator. In particular, by using this explicit representation, it is straightforward to see that a (mean field) CCE is a (mean field) NE if and only if it features a deterministic correlation device Z , simplifying the comparison of the two definitions. Moreover, referring in particular to the MFG, this is consistent with our interpretation of the consistency condition (3.2.3): as the moderator imagines what the random stationary mean $\bar{\mu}_\infty$ will be and gives her recommendation according to it, it is expected that the correlation device is the same as the interaction term itself, which in this framework is given by the stationary mean $\bar{\mu}_\infty$, a real valued random variable. This intuition is also coherent with the findings of Chapter 2, where the correlation was provided through the (time derivative of) the interaction term $(\bar{\mu}_t)_{t \in [0, T]}$. Finally, as it will be shown in Section 3.4, this class of mean field CCEs is rich enough to our scopes.

The study of the existence of CCEs in the ergodic MFG and its relation with the N -player game is the content of Section 3.4: We identify specific classes of correlated stationary strategies for which it is possible to state a sufficient condition to be a CCE. Then we use CCEs in those classes to build a sequence of approximate CCEs the underlying N -player system, with vanishing approximating error. As a byproduct, we get the same approximation result for NEs as well.

Cooperative framework

We address the mean field counterpart of the central planner's maximization problem, which is given by the mean field control (MFC) maximization problem. Roughly

speaking, this problem consists in maximizing the reward (3.2.2) under the additional constraint $\bar{m} = \lim_{t \rightarrow \infty} \mathbb{E}[X_t^\nu] =: \mathbb{E}[X_\infty^\nu]$, for every control choice ν .

In order to properly define the reward, we need to restrict the class of admissible controls.

Definition 3.2.5. We say that a strategy ν is admissible for the mean field control problem if $\nu \in \mathbb{A}$ and the process $(X_t^\nu)_{t \geq 0}$ admits a unique stationary distribution $p_\infty^\nu \in \mathcal{P}(\mathbb{R}_+)$. We denote the set of admissible strategies for the stationary MFC problem by \mathbb{A}_{MFC} .

For any $\nu \in \mathbb{A}_{MFC}$, denote by $\mathbb{E}[X_\infty^\nu]$ the first moment of the corresponding limit measure p_∞^ν . The payoff functional associated to a strategy $\nu \in \mathbb{A}_{MFC}$ is given by

$$\mathfrak{J}(\nu, \mathbb{E}[X_\infty^\nu]) = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^\nu)^\alpha (\mathbb{E}[X_\infty^\nu])^\beta dt - q\nu_T \right].$$

Definition 3.2.6. An admissible control $\hat{\nu} \in \mathbb{A}_{MFC}$ is an optimal control for the mean field control problem if

$$\mathfrak{J}(\hat{\nu}, \mathbb{E}[X_\infty^{\hat{\nu}}]) \geq \mathfrak{J}(\nu, \mathbb{E}[X_\infty^\nu]), \quad \forall \nu \in \mathbb{A}_{MFC}.$$

The study of the central planner's optimization problem and its relation with the N -agent system is the content of Section 3.5: We show in Theorem 3.5.1 that it is possible to completely characterize solutions of the MFC problem. Then, in Theorem 3.5.3, we use the solution of the MFC problem to build a sequence of approximate central planner's optima in the underlying N -agent system, with vanishing approximating error.

3.3 Assumptions and preliminary results

On top of Assumption **U**, we assume the following structural condition on the coefficients of the SDE:

Assumption D. The parameters δ and σ satisfy the following condition:

$$2\delta - \sigma^2 > 0.$$

Moreover, the initial distribution η belongs to $\mathcal{P}^2(\mathbb{R}_+)$.

Remark 3.3.1. Let X^0 be the solution of (3.2.1) when the policy ν is identically equal to 0. Assumption **D** is a dissipativity assumption on the square of X^0 : indeed, by Itô's formula, we have

$$(X_t^0)^2 = \xi^2 + (\sigma^2 - 2\delta) \int_0^t (X_s^0)^2 ds + 2\sigma \int_0^t (X_s^0)^2 dW_s,$$

which is then square-integrable and dissipative. Notice that the same assumption is assumed in [31, Section 6].

We state some important properties of the diffusion X^{ν^a} reflected upwards at some positive barrier a , which will be used extensively through the whole Chapter.

Lemma 3.3.1. *i) For any $a > 0$, let $p_a \in \mathcal{P}(\mathbb{R}_+)$ be given by*

$$p_a(dx) = \frac{2\delta + \sigma^2}{2} a^{\frac{2\delta}{\sigma^2} + 1} x^{-\frac{2\delta}{\sigma^2} - 2} \mathbb{1}_{\{x \geq a\}} dx. \quad (3.3.1)$$

For any $0 \leq k \leq 2$, the measure p_a admits finite k -moment. In particular, it holds

$$\int_{\mathbb{R}_+} x p_a(dx) = \frac{2\delta + \sigma^2}{2\delta} a. \quad (3.3.2)$$

Moreover, the map $\mathbb{R}_+ \times \mathcal{B}_{\mathbb{R}_+} \ni (a, B) \mapsto p_a(B)$ defines a stochastic kernel from \mathbb{R}_+ to $\mathcal{B}_{\mathbb{R}_+}$.

ii) For any $a > 0$, there exists a unique strong solution $(X_t^{\nu^a}, \nu_t^a)_{t \geq 0^-}$ of the Skorohod reflection problem at the barrier a , i.e. a pair of processes so that equation (3.2.1) is satisfied for any t , $\nu_{0^-} = 0$, $t \mapsto \nu_t$ is non-decreasing, and $\int_0^\infty \mathbb{1}_{\{X_s^{\nu^a} > a\}} d\nu_s^a = 0$, \mathbb{P} -a.s.. The process ν^a is given by

$$\nu_t^a = \int_0^t X_s^0 d \left(\sup_{0 \leq u \leq s} \left(\frac{a - X_u^0}{X_u^0} \right)^+ \right), \quad \nu_{0^-}^a = 0.$$

Moreover, the process X^{ν^a} is positively recurrent with stationary measure given by p_a .

iii) There exists a positive constant c so that

$$\sup_{t \geq 0} \mathbb{E}[(X_t^{\nu^a})^2] \leq c(1 + a^2), \quad \sup_{T > 0} \mathbb{E} \left[\left(\frac{1}{T} \nu_T^a \right)^2 \right] \leq c(1 + a^2).$$

The proof is postponed to Section 3.7. For later use, we introduce the real function $C(a, p)$, for $(a, p) \in \mathbb{R}_+^2$, given by

$$C(a, p) = (2\delta + \sigma^2) \left(\frac{p}{2\delta + \sigma^2(1 - \alpha)} a^\alpha - \frac{q}{2} a \right). \quad (3.3.3)$$

Let ν^a is the policy that reflects the process X^{ν^a} solution to (3.2.1) upwards à la Skorohod at the level $a > 0$. By [5, Lemma 2.1], it holds

$$C(a, p) = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T p \cdot (X_t^{\nu^a})^\alpha dt - q \nu_T^a \right] = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T p \cdot (X_t^{\nu^a})^\alpha dt - q \nu_T^a \right]$$

In the following, we will need to solve several ergodic optimization problems of singular controls. The existence of a unique solution to such problems is ensured by following Lemma:

Lemma 3.3.2. *Let $\gamma \in \mathbb{R}$ so that $q\delta - \gamma > 0$ and $p > 0$. Define the following function*

$$g(x, p, \gamma) := x^\alpha p + \gamma x, \quad (3.3.4)$$

and consider the reward functional

$$\tilde{\mathfrak{J}}(\nu, p, \gamma) := \liminf_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^\nu, p, \gamma) dt - q\nu_T \right], \quad (3.3.5)$$

where $X^\nu = (X_t^\nu)_{t \geq 0^-}$ evolves accordingly to (3.2.1). Then, there exists a unique optimal control $\hat{\nu} \in \mathcal{S}$ so that

$$\tilde{\mathfrak{J}}(\hat{\nu}, p, \gamma) = \sup_{\nu \in \mathcal{S}} \tilde{\mathfrak{J}}(\nu, p, \gamma).$$

Moreover, the process $\hat{\nu}$ reflects the state process upwards à la Skorohod at the barrier $\hat{a}(p, \gamma)$ given by

$$\hat{a}(p, \gamma) = \left(\frac{2\alpha\delta}{2\delta + \sigma^2(1-\alpha)} \frac{p}{q\delta - \gamma} \right)^{\frac{1}{1-\alpha}}. \quad (3.3.6)$$

The proof is postponed to Section 3.7.

With respect to a smaller class of strategies, we state a first order optimality condition for a control $\hat{\nu}$, inspired by [63] (see also [10]). Although only the necessary part will be needed, for the sake of completeness, we also show that it is sufficient under additional assumptions.

Lemma 3.3.3 (First order optimality condition). *Let $p > 0$, $\gamma > q\delta$. Let $1 \leq q, q' \leq \infty$ be Young conjugates, and define the set \mathcal{S}_{2q} as the set of controls $\nu \in \mathcal{S}$ so that*

$$\sup_{T > 0} \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^\nu|^{2q} dt \right] < \infty. \quad (3.3.7)$$

Let $\hat{\nu} \in \mathcal{S}_{2q}$ so that

$$\sup_{T > 0} \frac{1}{T} \mathbb{E} \left[\int_0^T |(X_t^{\hat{\nu}})^{\alpha-2}|^{q'} dt \right] < \infty, \quad (3.3.8)$$

if $q' < \infty$, and so that

$$\sup_{T > 0} \inf \{ C \geq 0 : (X_t^{\hat{\nu}})^{\alpha-2} \leq C \text{ for } dt \otimes d\mathbb{P}\text{-a.e. } (t, \omega) \in [0, T] \times \Omega \} < \infty, \quad (3.3.9)$$

if $q' = \infty$.

(a) Suppose that $\hat{\nu}$ is optimal within the set \mathcal{S}_{2q} for the control problem with dynamics (3.2.1) and reward (3.3.5). Then, for every $\nu \in \mathcal{S}_{2q}$, it holds

$$\liminf_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g_x(X_t^{\hat{\nu}}, p, \gamma)(X_t^{\hat{\nu}} - X_t^\nu) dt - q(\hat{\nu}_T - \nu_T) \right] \geq 0. \quad (3.3.10)$$

(b) Suppose that either $\hat{\nu}$ satisfies (3.3.10) and

$$\lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^{\hat{\nu}}, p, \gamma) dt - q\hat{\nu}_T \right] = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^\nu, p, \gamma) dt - q\nu_T \right], \quad (3.3.11)$$

or that $\hat{\nu}$ satisfies

$$\liminf_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g_x(X_t^{\hat{\nu}}, p, \gamma)(X_t^{\hat{\nu}} - X_t^\nu) dt - q(\hat{\nu}_T - \nu_T) \right] \geq 0. \quad (3.3.12)$$

Then $\hat{\nu}$ is optimal within the set \mathcal{S}_{2q} .

The proof is postponed to Section 3.7.

Remark 3.3.2. Let $a > 0$ and let ν^a be the policy that reflects the process $X^{\hat{\nu}}$ upwards à la Skorohod at the barrier a . Then, the control ν^a is so that $\mathbb{P}((X_t^{\nu^a})^{\alpha-2} \leq c \forall t \geq 0) = 1$, for a constant $c > 0$, since the reflected process is so that $X_t^{\nu^a} \geq a$ for any t , and $X^{\hat{\nu}}$ belongs to \mathcal{S}_2 by Lemma 3.3.1. Therefore, (3.3.9) is satisfied, and we take $q = 1$ in (3.3.7). Assumption **D** and Lemma 3.3.1 imply that X^{ν^a} satisfies (3.3.7). By Lemma 3.3.1, (3.3.11) is satisfied as well.

Remark 3.3.3. We can restate Lemma 3.3.3 in terms of linear conditions involving optional projections. Consider the probability measure $\tilde{\mathbb{P}}$ equivalent to \mathbb{P} defined via the Radon-Nykodim derivative

$$\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} \Big|_{\mathcal{F}_t} = e^{\sigma W_t - \frac{\sigma^2}{2}t}, \quad t \geq 0.$$

It can be shown that, for every $\nu \in \mathcal{S}$, the state process X^ν can be represented as

$$X_t^\nu = e^{-\delta t} M_t(\xi + \bar{\nu}_t),$$

where $\bar{\nu} = (\bar{\nu}_t)_{t \geq 0}$ is defined via the identity $\nu_t = \int_0^t e^{-\delta s} M_s d\bar{\nu}_s$. Denote by $\tilde{\mathbb{E}}$ the expectation with respect to $\tilde{\mathbb{P}}$. By taking advantage of Fubini's theorem and optional projections as in [49, Theorem 57, Chapter VI, p. 122], we can restate necessary condition (3.3.10) as

$$\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \tilde{\mathbb{E}} \left[\int_0^T \left(\tilde{\mathbb{E}} \left[\int_s^T e^{-\delta t} g_x(X_t^{\hat{\nu}}, p, \gamma) dt \Big| \mathcal{F}_s \right] - q e^{-\delta s} \right) d(\hat{\nu}_s - \bar{\nu}_s) \right] \geq 0, \quad \forall \nu \in \mathcal{S}_{2q},$$

and analogously holds for (3.3.11) and (3.3.12). While first order conditions for optimality are well known for both finite and infinite time horizon singular control problems (see, e.g., [9, 10, 63, 64] among others), to the best of our knowledge, they had never been derived for singular control problems with ergodic reward functionals.

3.4 Competitive case: Mean field equilibria and approximation

In this section, we focus on coarse correlated equilibria for the MFG, as defined by Definition 3.2.3. We follow the same procedure outlined in Chapter 2: since the set of CCEs is typically very wide and it is difficult to characterize in a continuous time setting, we restrict our analysis to specific classes of correlated stationary strategies for which we are able to state a sufficient condition for being a CCE for the ergodic MFG. With respect to Chapter 2, we move a step forward, and show that every CCE in these classes induces a sequence of approximate CCEs in the underlying N -player game with vanishing error. This result is in the same spirit as Theorem 1.4.1 in Chapter 1, although specific to this context. More specifically, we establish the following:

- We fix a correlated stationary strategy $(Z, \lambda, \bar{\mu}_\infty)$. We suppose that the representative player decides to ignore the moderator's recommendation, and compute her best deviating strategy, i.e.

$$U^* = \arg \max_{\nu \in \mathbb{A}} \mathfrak{J}(\nu, \bar{\mu}_\infty).$$

This is the content of Proposition 3.4.1.

- We define specific classes of correlated flows $(Z, \lambda, \bar{\mu}_\infty)$ so that the consistency condition (ii) is satisfied. This is established in Propositions 3.4.2 and 3.4.5.
- For $(Z, \lambda, \bar{\mu}_\infty)$ in such classes, we express the optimality condition (i) as an inequality involving the law of $\bar{\mu}_\infty$ only, thus deriving a sufficient condition for the existence of CCEs. This is the content of Propositions 3.4.3 and 3.4.6.
- We show that every mean field CCEs in each class induces a sequence of approximate CCEs in the underlying N -player game with vanishing error. This is the content of Theorems 3.4.4 and 3.4.7.

We consider two classes of correlated stationary strategies: while in both classes the correlating device is the random mean $\bar{\mu}_\infty$ itself, in the first class the recommendation λ^r is a $\sigma(\bar{\mu}_\infty)$ -measurable regular control, while in the second one, the recommendation λ^s is a policy of reflection type at a random barrier $a(\bar{\mu}_\infty)$. Surprisingly, the sufficient condition of the two classes differ only by a constant. Moreover, in Theorem 3.4.8, we explicitly characterize existence and uniqueness of Nash equilibria. We find that they belong to class of CCEs with recommendation of reflection type, so that the same approximation result applies.

3.4.1 The deviating player problem

Suppose that the representative player decides to ignore the moderator's recommendation. By definition of CCE for the ergodic MFG, the deviating player must choose her strategy $\nu \in \mathbb{A}$ only by knowing the joint law of the correlated stationary strategy $(Z, \lambda, \bar{\mu}_\infty)$, which is assumed to be publicly known, and not by observing its realizations. Since $\nu \in \mathbb{A}$, it follows that X^ν is $\mathbb{F}^{\xi, W}$ -adapted as well and thus independent of the random variable $\bar{\mu}_\infty$, which implies that deviating player's payoff can be written as:

$$\begin{aligned} \mathfrak{J}(\nu, \bar{\mu}_\infty) &= \lim_{T \uparrow \infty} \frac{1}{T} \left(\int_0^T \mathbb{E} \left[\mathbb{E}[(X_t^\nu)^\alpha \bar{\mu}_\infty^\beta | \mathcal{F}_t^{\xi, W}] \right] - q \mathbb{E} \left[\mathbb{E}[\nu_T | \mathcal{F}_T^{\xi, W}] \right] \right) \\ &= \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^\nu)^\alpha \mathbb{E}[\bar{\mu}_\infty^\beta] dt - q \nu_T \right]. \end{aligned} \quad (3.4.1)$$

Observe that deviating player's payoff functional depends on $\bar{\mu}_\infty$ only through its expectation.

Proposition 3.4.1. *There exists a unique optimal strategy for the deviating player $U^* \in \mathbb{A}$ which reflects the process X^{U^*} upwards á la Skorohod at the level a^* , where a^* is given by*

$$a^* = \left(\frac{2\alpha}{q(2\delta + \sigma^2(1 - \alpha))} \right)^{\frac{1}{1-\alpha}} (\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}}. \quad (3.4.2)$$

Proof. Since the payoff functional of the deviating player is given by (3.4.1), it is enough to apply Lemma 3.3.2 with $\gamma = 0$ and $p = \mathbb{E}[\bar{\mu}_\infty^\beta]$. \square

In the following, we set

$$K := \left(\frac{2\alpha}{q(2\delta + \sigma^2(1 - \alpha))} \right)^{\frac{1}{1-\alpha}}, \quad (3.4.3)$$

so that the optimal policy of the deviating player is the reflection at the level $a^* = K(\mathbb{E}[\bar{\mu}_\infty^\beta])^{1/1-\alpha}$.

3.4.2 Regular recommendation

Let \mathcal{G}^r be the set of correlated stationary strategies $(Z, \lambda^r, \bar{\mu}_\infty)$ so that $\bar{\mu}_\infty \in L^2(\mathcal{F}_{0-})$ independent of ξ and W , $Z = \bar{\mu}_\infty$ and

$$d\lambda_t^r = \delta \bar{\mu}_\infty dt, \quad t \geq 0, \quad (3.4.4)$$

so that, in particular, the recommended strategy λ^r is a $\sigma(\bar{\mu}_\infty)$ -measurable regular control.

Proposition 3.4.2. *Let $(Z, \bar{\mu}_\infty, \lambda^r) \in \mathcal{G}^r$. Define*

$$\mathbf{m}'_{\infty, m}(x) = \frac{2}{\sigma^2} \exp\left(\frac{2\delta m}{\sigma^2}\right) x^{-\frac{2\delta}{\sigma^2}-2} \exp\left(-\frac{2\delta}{\sigma^2} \frac{m}{x}\right), \quad (3.4.5)$$

and let $p_\infty^r(dx, m)$ be the stochastic kernel from \mathbb{R}_+ to $\mathcal{B}_{\mathbb{R}_+}$ defined by

$$p_\infty^r(dx, m) = \frac{x^{-\frac{2\delta}{\sigma^2}-2} \exp\left(-\frac{2\delta}{\sigma^2} \frac{m}{x}\right)}{\int_0^\infty y^{-\frac{2\delta}{\sigma^2}-2} \exp\left(-\frac{2\delta}{\sigma^2} \frac{m}{y}\right) dy} dx. \quad (3.4.6)$$

Then, the triple $(Z, \bar{\mu}_\infty, \lambda^r)$ is a correlated stationary strategy so that the consistency condition (3.2.3) is satisfied. In particular, it holds $\theta_\infty^r(dx, dm) = \lim_{t \rightarrow \infty} \mathbb{P} \circ (X_t^{\lambda^r}, \bar{\mu}_\infty)^{-1} = p_\infty^r(dx, m)\rho(dm)$.

Proof. The measurability requirements on the triple $(Z, \bar{\mu}_\infty, \lambda^r)$ are clearly satisfied. Let X^{λ^r} the state process controlled by λ^r , which satisfies

$$dX_t^{\lambda^r} = \delta(\bar{\mu}_\infty - X_t^{\lambda^r})dt + \sigma X_t^{\lambda^r} dW_t, \quad X_0^{\lambda^r} = \xi. \quad (3.4.7)$$

We show that the joint law of $(X_t^{\lambda^r}, \bar{\mu}_\infty)$ converges weakly to θ_∞^r as $t \rightarrow \infty$. To this extent, it is enough to verify that the regular conditional probability of $X_t^{\lambda^r}$ with respect to $\bar{\mu}_\infty = m$, which we denote by $p_t^r(dx, m)$, converges weakly to $p_\infty^r(dx, m)$ as $t \rightarrow \infty$, for ρ -a.e. $m \in \mathbb{R}_+$. Conditionally to $\bar{\mu}_\infty = m$, $X_t^{\lambda^r}$ satisfies the following equation:

$$dX_t^{\lambda^r, m} = \delta(m - X_t^{\lambda^r, m})dt + \sigma X_t^{\lambda^r, m} dW_t, \quad X_0^{\lambda^r, m} = \xi.$$

By Lemma 3.7.1, we have $\int_0^\infty \mathbf{m}'_{\infty, m}(x) dx < \infty$, which implies that the diffusion $X^{\lambda^r, m}$ is positively recurrent for every $m > 0$. Thus, the measure $p_\infty^r(dx, m)$ is the unique

stationary distribution and $p_t^r(dx, m) \rightarrow p_\infty^r(dx, m)$ in total variation norm (see, e.g., [21, Paragraph 36]). As for equality (3.2.3), define $\varphi = (\varphi_t(m))_{t \geq 0, m > 0}$ as

$$\varphi_t(m) = e^{-\delta t} \mathbb{E}[\xi] + m(1 - e^{\delta t}). \quad (3.4.8)$$

By Itô's formula, it follows that $\varphi_t(\bar{\mu}_\infty) = \mathbb{E}[X_t^{\lambda^r} | \bar{\mu}_\infty]$ for every $t \geq 0$, \mathbb{P} -a.s., which implies that $\bar{\mu}_\infty = \lim_{t \rightarrow \infty} \mathbb{E}[X_t | \bar{\mu}_\infty]$ \mathbb{P} -a.s., and therefore condition (3.2.4) is satisfied, and so (3.2.3). \square

Proposition 3.4.3. *A correlated stationary strategy $(Z, \bar{\mu}_\infty, \lambda^r)$ in the class \mathcal{G}^r is a mean field CCE if and only if the following inequality is satisfied:*

$$c_\beta (\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}} + c_1 \mathbb{E}[\bar{\mu}_\infty] \leq c_{\alpha+\beta} \mathbb{E}[\bar{\mu}_\infty^{\alpha+\beta}], \quad (3.4.9)$$

where c_β , $c_{\alpha+\beta}$ and c_1 are positive constants defined by

$$\begin{aligned} c_\beta &:= \frac{(2\delta + \sigma^2)q}{2} \left(\frac{2\alpha}{q(2\delta + \sigma^2(1-\alpha))} \right)^{\frac{1}{1-\alpha}} \frac{1-\alpha}{\alpha}, \quad c_1 := \delta q, \\ c_{\alpha+\beta} &:= \left(\frac{2\delta}{\sigma^2} \right)^\alpha \frac{\Gamma(\frac{2\delta}{\sigma^2} + 1 - \alpha)}{\Gamma(\frac{2\delta}{\sigma^2} + 1)}. \end{aligned} \quad (3.4.10)$$

Proof. By Proposition 3.4.2, $(Z, \bar{\mu}_\infty, \lambda^r)$ satisfies the consistency condition (3.2.3). Let U^* be the optimal control for the deviating player given by Proposition 3.4.1. Since $\mathfrak{J}(U^*, \bar{\mu}_\infty) = \max_{\nu \in \mathbb{A}} \mathfrak{J}(\nu, \bar{\mu}_\infty)$, we just need to verify that the inequality $\mathfrak{J}(\lambda^r, \bar{\mu}_\infty) \geq \mathfrak{J}(U^*, \bar{\mu}_\infty)$ is equivalent to (3.4.9). Since U^* is a reflection policy at the level a^* given by (3.4.2), formulae (3.3.3) and (3.4.1) yield

$$\begin{aligned} \mathfrak{J}(U^*, \bar{\mu}_\infty) &= C(a^*, \mathbb{E}[\bar{\mu}_\infty^\beta]) \\ &= \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1-\alpha)} \mathbb{E}[\bar{\mu}_\infty^\beta] \left(K(\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}} \right)^\alpha - q \frac{2\delta + \sigma^2}{2} K(\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}} \\ &= (2\delta + \sigma^2) \left(\frac{1}{2\delta + \sigma^2(1-\alpha)} K^\alpha - \frac{q}{2} K \right) (\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}} = c_\beta (\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}}, \end{aligned} \quad (3.4.11)$$

by noticing that, with K given by (3.4.3), it holds

$$(2\delta + \sigma^2) \left(\frac{1}{2\delta + \sigma^2(1-\alpha)} K^\alpha - \frac{q}{2} K \right) = \frac{2\delta + \sigma^2}{2} q K \left(\frac{1-\alpha}{\alpha} \right) = c_\beta.$$

As for the payoff associated to the representative player, conditionally to $\bar{\mu}_\infty = m$ and exploiting ergodicity, it holds

$$\begin{aligned} \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\lambda^r})^\alpha \bar{\mu}_\infty^\beta dt - q \lambda_T^r | \bar{\mu}_\infty = m \right] &= \int_0^\infty m^\beta x^\alpha p_\infty^r(dx, m) - \delta q m \\ &= \left(\frac{2\delta}{\sigma^2} \right)^\alpha \frac{\Gamma(\frac{2\delta}{\sigma^2} - \alpha + 1)}{\Gamma(\frac{2\delta}{\sigma^2} + 1)} m^{\alpha+\beta} - \delta q m, \quad \text{for } \rho\text{-a.e. } m \in \mathbb{R}_+, \end{aligned}$$

where last equality follows from the definition of $p_\infty^r(dx, m)$ and Lemma 3.7.1. Moreover, by (3.4.8), we have the bound

$$\begin{aligned} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\lambda^r})^\alpha \bar{\mu}_\infty^\beta dt - q\lambda_T^r | \bar{\mu}_\infty \right] &\leq \frac{\bar{\mu}_\infty^\beta}{T} \mathbb{E} \left[\int_0^T (1 + X_t^{\lambda^r}) dt | \bar{\mu}_\infty \right] + \delta q \bar{\mu}_\infty \\ &\leq \bar{\mu}_\infty^\beta \left(1 + \sup_{t \geq 0} \mathbb{E} [X_t^{\lambda^r} | \bar{\mu}_\infty] \right) + \delta q \bar{\mu}_\infty \leq C(1 + \bar{\mu}_\infty^2), \quad \forall T > 0, \mathbb{P}\text{-a.s.}, \end{aligned}$$

which is integrable by assumption. Therefore, by dominated convergence theorem, we can exchange limit and expectation to conclude

$$\begin{aligned} \mathfrak{J}(\lambda^r, \bar{\mu}_\infty) &= \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\mathbb{E} \left[\int_0^T (X_t^{\lambda^r})^\alpha \bar{\mu}_\infty^\beta dt - q\lambda_T^r | \bar{\mu}_\infty \right] \right] \\ &= \mathbb{E} \left[\lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\lambda^r})^\alpha \bar{\mu}_\infty^\beta dt - q\lambda_T^r | \bar{\mu}_\infty \right] \right] \\ &= \left(\frac{2\delta}{\sigma^2} \right)^\alpha \frac{\Gamma(\frac{2\delta}{\sigma^2} - \alpha + 1)}{\Gamma(\frac{2\delta}{\sigma^2} + 1)} \int_0^\infty m^{\alpha+\beta} \rho(dm) - \delta q \int_0^\infty m \rho(dm), \end{aligned}$$

By comparing this equation with equation (3.4.11), we get equation (3.4.9). \square

Remark 3.4.1. Observe that, if $\alpha + \beta \leq 1$, for this result to hold true it is enough to require $\bar{\mu}_\infty$ to be in L^1 . Assumption **D** is not needed as well.

Finally, we show how to use a CCE in the class \mathcal{G}^r to build a sequence λ^N of ε_N -CCE in the N -player game with $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$. To this extent, we consider the following set of strategies $\mathcal{B} \subseteq \mathbb{A}_N$:

Definition 3.4.1 (c -admissible strategies). Let $c > 0$. A strategy ν be in \mathbb{A}_N is c -admissible if

$$\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^\nu|^2 dt \right] \leq c. \quad (3.4.12)$$

We denote by $\mathbb{A}_{N,c}$ the set of c -admissible strategies for the N -player game.

Starting from a correlated stationary strategy $(Z, \bar{\mu}_\infty, \lambda^r)$ in the class \mathcal{G}^r , we define the following correlated strategy profiles for the N -player game: take $Z = \bar{\mu}_\infty$ as correlation device, and set, for any $i \geq 1$,

$$d\lambda_t^{i,r} = \delta \bar{\mu}_\infty dt. \quad (3.4.13)$$

Then, when player i plays accordingly to moderator's suggestion, her dynamics are hold by the following equation:

$$dX_t^{i,r} = \delta(\bar{\mu}_\infty - X_t^{i,r})dt + \sigma X_t^{i,r} dW_t^i, \quad X_0^{i,r} = \xi^i. \quad (3.4.14)$$

Observe that, for each $i \geq 1$, the triple $(X^{i,r}, \lambda^{i,r}, \bar{\mu}_\infty)$ has the same law as $(X, \lambda, \bar{\mu}_\infty)$. Moreover, while not independent, the processes $(X^{i,r})_{i \geq 1}$ are conditionally independent given $\bar{\mu}_\infty$.

Theorem 3.4.4 (Approximation of CCEs - regular case). *Let $(Z, \bar{\mu}_\infty, \lambda^r)$ be a CCE in the class \mathcal{G}^r . Let $\lambda^N = (\lambda^{i,r})_{i=1}^N$, with $\lambda^{i,r}$ defined by (3.4.13). Then, for any $c > 0$, the correlated strategy profile $(\bar{\mu}_\infty, \lambda^N)$ defines an ε_N -CCE for the N -player game within the set of strategies $\mathbb{A}_{N,c}$, with $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$.*

Proof. For every $N \geq 2$, set

$$\varepsilon_N := \sup_{\nu \in \mathbb{A}_{N,c}} (\mathfrak{J}_N(\nu, \lambda^{-i,N}) - \mathfrak{J}_N(\lambda^{i,r}, \lambda^{-i,N})).$$

Notice that, by symmetry, ε_N is independent of $i = 1, \dots, N$. Clearly $(\bar{\mu}_\infty, \lambda^N)$ is an ε_N -CCE for the N -player game within the set of strategies $\mathbb{A}_{N,c}$. We show that ε_N vanishes as $N \rightarrow \infty$. Note that, for any ν in $\mathbb{A}_{N,c}$, we have

$$\begin{aligned} \mathfrak{J}_N(\nu, \lambda^{-i,N}) - \mathfrak{J}_N(\lambda^{i,r}, \lambda^{-i,N}) &= (\mathfrak{J}_N(\nu, \lambda^{-i,N}) - \mathfrak{J}(\nu, \bar{\mu}_\infty)) \\ &\quad + (\mathfrak{J}(\nu, \bar{\mu}_\infty) - \mathfrak{J}(\lambda^{i,r}, \bar{\mu}_\infty)) + (\mathfrak{J}(\lambda^{i,r}, \bar{\mu}_\infty) - \mathfrak{J}_N(\lambda^{i,r}, \lambda^{-i,N})) \end{aligned} \quad (3.4.15)$$

where \mathfrak{J} is defined by (3.2.2). We treat separately each of the three terms in the right-hand side of (3.4.15).

As for the first term, by Cauchy-Schwartz inequality and using the inequality $\underline{\lim}_{n \rightarrow \infty} \alpha_n - \underline{\lim}_{n \rightarrow \infty} \beta_n \leq \overline{\lim}_{n \rightarrow \infty} (\alpha_n - \beta_n)$, we have the following estimates:

$$\begin{aligned} &|\mathfrak{J}(\nu, \bar{\mu}_\infty) - \mathfrak{J}_N(\nu, \lambda^{-i,N})| \\ &\leq \overline{\lim}_{T \uparrow \infty} \left(\frac{1}{T} \int_0^T \mathbb{E} [(X_t^{i,\nu})^{2\alpha}] dt \right)^{\frac{1}{2}} \left(\frac{1}{T} \int_0^T \mathbb{E} \left[\left| (\bar{\mu}_t^{N, \lambda^{-i,N}})^\beta - \bar{\mu}_\infty^\beta \right|^2 \right] dt \right)^{\frac{1}{2}} \\ &\leq \left(1 + \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} [|X_t^{i,\nu}|^2] dt \right)^{\frac{1}{2}} \\ &\quad \cdot \left(\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\mathbb{E} \left[\left| (\bar{\mu}_t^{N, \lambda^{-i,N}})^\beta - \bar{\mu}_\infty^\beta \right|^2 \middle| \bar{\mu}_\infty \right] \right] dt \right)^{\frac{1}{2}} \\ &\leq (1+c)^{\frac{1}{2}} \left(\mathbb{E} \left[\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left| (\bar{\mu}_t^{N, \lambda^{-i,N}})^\beta - \bar{\mu}_\infty^\beta \right|^2 \middle| \bar{\mu}_\infty \right] dt \right] \right)^{\frac{1}{2}}, \end{aligned} \quad (3.4.16)$$

where in the last inequality we exchanged limsup and expectation by reverse Fatou's lemma with the integrable upper bound $C(1 + \bar{\mu}_\infty^2)$. Indeed, by recalling that $(X^j)_{j \neq i}$ are i.i.d. as X conditionally to $\bar{\mu}_\infty$, we have the following estimates

$$\mathbb{E} \left[\left| (\bar{\mu}_t^{N, \lambda^{-i,N}})^\beta - \bar{\mu}_\infty^\beta \right|^2 \middle| \bar{\mu}_\infty \right] \leq C \left(1 + \bar{\mu}_\infty^2 + \frac{1}{N-1} \sum_{j \neq i} \mathbb{E} [|X_t^j|^2 \middle| \bar{\mu}_\infty] \right) \leq C(1 + \bar{\mu}_\infty^2),$$

where last inequality holds thanks to Lemma 3.7.2. By Lemma 3.7.3, we then have

$$\begin{aligned} &\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left| (\bar{\mu}_t^{N, \lambda^{-i,N}})^\beta - \bar{\mu}_\infty^\beta \right|^2 \middle| \bar{\mu}_\infty \right] dt \\ &= \int_{\mathbb{R}_+^{N-1}} \left| \left(\frac{1}{N-1} \sum_{j \neq i} x_j \right)^\beta - \bar{\mu}_\infty^\beta \right|^2 \bigotimes_{j \neq i} p_\infty^r(dx_j, \bar{\mu}_\infty), \quad \mathbb{P}\text{-a.s.} \end{aligned}$$

By Lemma 3.7.4, the right hand-side converges to 0 as N goes to ∞ . By applying dominated convergence theorem with $C(1 + \bar{\mu}_\infty^2)$ as the integrable upper bound, the convergence holds in expectation as well.

As for the second term, we claim that $\mathfrak{J}(\nu, \bar{\mu}_\infty) - \mathfrak{J}(\lambda^{i,r}, \bar{\mu}_\infty) \leq 0$ for any $\nu \in \mathbb{A}_{N,c}$. Indeed, observe that, since $\nu \in \mathbb{A}_N$ is independent of $\bar{\mu}_\infty$ by definition of admissible strategies for the N -player game, the proof of Proposition 3.4.1 shows that

$$\sup_{\nu \in \mathbb{A}_N} \mathfrak{J}(\nu, \bar{\mu}_\infty) = \mathfrak{J}(U^{i,*}, \bar{\mu}_\infty),$$

where $U^{i,*}$ is the policy that reflects the process $X^{i,U^{i,*}}$ upward à la Skorohod at the level a^* given by (3.4.2). In particular, $X^{i,U^{i,*}}$ has the same distribution as the process X^{U^*} . Therefore, we have

$$\sup_{\nu \in \mathbb{A}_{N,c}} (\mathfrak{J}(\nu, \bar{\mu}_\infty) - \mathfrak{J}(\lambda^{i,r}, \bar{\mu}_\infty)) \leq \sup_{\nu \in \mathbb{A}_N} \mathfrak{J}(\nu, \bar{\mu}_\infty) - \mathfrak{J}(\lambda, \bar{\mu}_\infty) \leq \mathfrak{J}(U^*, \bar{\mu}_\infty) - \mathfrak{J}(\lambda, \bar{\mu}_\infty) \leq 0$$

where we used the inclusion $\mathbb{A}_{N,c} \subseteq \mathbb{A}_N$, the fact that $(X^{i,r}, \lambda^{i,r}, \bar{\mu}_\infty)$ has the same distribution as $(X, \lambda, \bar{\mu}_\infty)$ and the optimality property (i) of CCE of the ergodic MFG.

As for the third term, taking advantage of the conditional independence and identical distribution of $(X^{\lambda^{i,r}})_{i \geq 1}$ and by analogous estimates as in (3.4.16), we have

$$\begin{aligned} & |\mathfrak{J}(\lambda^{i,r}, \bar{\mu}_\infty) - \mathfrak{J}_N(\lambda^{i,r}, \lambda^{-i,N})| \\ & \leq C(1 + \mathbb{E}[\bar{\mu}_\infty^2])^{\frac{1}{2}} \left(\mathbb{E} \left[\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left| (\bar{\mu}_t^{N, \lambda^{-i,N}})^\beta - \bar{\mu}_\infty^\beta \right|^2 \middle| \bar{\mu}_\infty \right] dt \right] \right)^{\frac{1}{2}} \end{aligned}$$

and we conclude the proof by applying Lemmata 3.7.3 and 3.7.4 as before. \square

Remark 3.4.2 (On the integrability condition). It is worth to notice that the integrability condition (3.4.12) that defines c -admissible strategies can be weakened at the price of more integrability requirements on $\bar{\mu}_\infty$, ξ and the diffusion X^0 . Indeed, let $q = \beta/1 - \alpha$, $k = 1 + \lceil q \rceil$, and suppose that $2\delta - (k-1)\sigma > 0$, $\mathbb{E}[\xi^k] < \infty$ and $\mathbb{E}[\bar{\mu}_\infty^k] < \infty$, which imply that the estimates in point iii) of Lemma 3.3.1 holds up to the k -th moment. Then, up to little modification of the proof, one could consider strategies $\nu \in \mathbb{A}_N$ so that

$$\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T X_s^\nu ds \right] \leq c.$$

3.4.3 Singular recommendation

We now look for a policy λ^s of reflection type at a random barrier $a(\bar{\mu}_\infty)$. Let \mathcal{G}^s be the set of correlated stationary strategies $(Z, \lambda^s, \bar{\mu}_\infty)$ so that $Z = \bar{\mu}_\infty$, $\bar{\mu}_\infty \in L^2(\mathcal{F}_{0-})$ independent of ξ and W , and λ^s is the control that reflects the process X^{λ^s} upwards à la Skorohod at the random level

$$a(\bar{\mu}_\infty) = \frac{2\delta}{2\delta + \sigma^2} \bar{\mu}_\infty. \quad (3.4.17)$$

Proposition 3.4.5. *Let $(Z, \lambda^s, \bar{\mu}_\infty)$ in \mathcal{G}^s . Let $p_\infty^s(dx, m)$ be the stochastic kernel from \mathbb{R}_+ to $\mathcal{B}_{\mathbb{R}_+}$ defined by*

$$p_\infty^s(dx, m) = p_{a(m)}(dx), \quad (3.4.18)$$

where p_a is the family of measures defined by (3.3.1). Then, the triple $(Z, \bar{\mu}_\infty, \lambda^s)$ is a correlated stationary strategy so that the consistency condition (3.2.3) is satisfied. In particular, it holds $\theta_\infty^s(dx, dm) = \lim_{t \rightarrow \infty} \mathbb{P} \circ (X_t^{\lambda^s}, \bar{\mu}_\infty)^{-1} = p_\infty^s(dx, m)\rho(dm)$.

Proof. Since the map $\mathbb{R}_+ \times \mathcal{B}_{\mathbb{R}_+} \ni (a, B) \mapsto p_a(B)$ defines a stochastic kernel from \mathbb{R}_+ to $\mathcal{B}_{\mathbb{R}_+}$, the kernel (3.4.18) is well defined. As in the proof of Proposition 3.4.2, we show that the joint law of $X_t^{\lambda^s}$ and $\bar{\mu}_\infty$ converges weakly to θ_∞^s as $t \rightarrow \infty$. Indeed, conditionally to $\bar{\mu}_\infty = m$, $X_t^{\lambda^s}$ satisfies the equation (3.2.1) with ν replaced by $\lambda^{s,m}$, where $\lambda^{s,m}$ reflects the process $X^{\lambda^{s,m}}$ upwards at the level $a(m)$, for ρ -a.e. $m \in \mathbb{R}_+$. By Lemma 3.3.1, the reflected process $X^{\lambda^{s,m}}$ admits $p_{a(m)}$ given by (3.3.1) as the unique invariant distribution. This implies that the regular conditional probability of $X_t^{\lambda^s}$ with respect to $\bar{\mu}_\infty$, that we denote by $p_t^s(dx, m)$, converges weakly to $p_\infty^s(dx, m)$ as $t \rightarrow \infty$ for ρ -a.e. $m > 0$. Consistency condition (3.2.3) follows from the definition of $a(\bar{\mu}_\infty)$ and (3.3.2). \square

Proposition 3.4.6. *A correlated stationary strategy $(Z, \bar{\mu}_\infty, \lambda^s)$ in \mathcal{G}^s is a mean field CCE for the ergodic MFG if and only if the following inequality is satisfied:*

$$c_\beta(\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}} + c_1\mathbb{E}[\bar{\mu}_\infty] \leq \tilde{c}_{\alpha+\beta}\mathbb{E}[\bar{\mu}_\infty^{\alpha+\beta}], \quad (3.4.19)$$

where c_1 and c_β are given by (3.4.10) and $\tilde{c}_{\alpha+\beta}$ is given by

$$\tilde{c}_{\alpha+\beta} := \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1-\alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha. \quad (3.4.20)$$

Proof. As in the proof of Proposition 3.4.3, it is enough to verify that the inequality $\mathfrak{J}(\lambda^s, \bar{\mu}_\infty) \geq \mathfrak{J}(U^*, \bar{\mu}_\infty)$ is equivalent to (3.4.19). By (3.4.11), the payoff of the deviating player is equal to

$$\mathfrak{J}(U^*, \bar{\mu}_\infty) = c_\beta(\mathbb{E}[\bar{\mu}_\infty^\beta])^{\frac{1}{1-\alpha}}.$$

We turn our attention to $\mathfrak{J}(\lambda^s, \bar{\mu}_\infty)$. We note that, conditionally to $\bar{\mu}_\infty = m$, it holds

$$\lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\lambda^s})^\alpha \bar{\mu}_\infty^\beta dt - q\lambda_T^s \middle| \bar{\mu}_\infty = m \right] = C(a(m), m^\beta), \quad \rho\text{-a.e. } m \in \mathbb{R}_+,$$

where $C(a, p)$ is given by (3.3.3). To see this, it is enough to recall that, by Proposition 3.4.5, for ρ -a.e. m in \mathbb{R}_+ , we have $p_t^s(dx, m) \rightarrow p_\infty^s(dx, m)$ weakly as $t \rightarrow \infty$, with $p_\infty^s(dx, m)$ given by (3.4.18). Since, conditionally to $\bar{\mu}_\infty = m$, the control λ^s is a reflection at the barrier $a(m)$, we apply [5, Lemma 2.1] to get the equality above. By applying Lemma 3.3.1 with $a = a(\bar{\mu}_\infty)$, and exploiting square-integrability of $\bar{\mu}_\infty$, at any time $T > 0$ we can bound the left hand-side with $C(1 + \bar{\mu}_\infty^2)$, for some positive

constant C independent of $\bar{\mu}_\infty$. Therefore, by dominated convergence theorem, we can exchange limit and expectation, to get

$$\begin{aligned}\mathfrak{J}(\lambda^s, \bar{\mu}_\infty) &= \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\mathbb{E} \left[\int_0^T (X_t^{\lambda^s})^\alpha \bar{\mu}_\infty^\beta dt - q\lambda_T^s \middle| \bar{\mu}_\infty \right] \right] = \mathbb{E}[C(a(\bar{\mu}_\infty), \bar{\mu}_\infty^\beta)] \\ &= \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1 - \alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha \mathbb{E}[\bar{\mu}_\infty^{\alpha+\beta}] - q\delta \mathbb{E}[\bar{\mu}_\infty].\end{aligned}$$

By rearranging the terms, we have that $\mathfrak{J}(U^*, \bar{\mu}_\infty) \leq \mathfrak{J}(\lambda, \bar{\mu}_\infty)$ if and only if equation (3.4.19) is satisfied. \square

Finally, consider a correlated stationary strategy $(Z, \bar{\mu}_\infty, \lambda^s)$ in the family \mathcal{G}^s and define a correlated strategy profile for the N -player game starting from it: we take $Z = \bar{\mu}_\infty$ as correlation device, and, for any $i \geq 1$, we consider the policy $\lambda^{i,s} = (\lambda_t^{i,s})_{t \geq 0}$ according to which the state $X^{i,s}$ is reflected upward at the random barrier $a(\bar{\mu}_\infty)$ given by (3.4.17). As in the case of a regular recommendation, for each $i \geq 1$, the triple $(X^{i,s}, \lambda^{i,s}, \bar{\mu}_\infty)$ has the same law as $(X, \lambda^s, \bar{\mu}_\infty)$, and the processes $(X^{i,s})_{i \geq 1}$ are conditionally independent given $\bar{\mu}_\infty$.

Theorem 3.4.7 (Approximation of CCEs - singular case). *Let $(Z, \bar{\mu}_\infty, \lambda^r)$ be a CCE in the class \mathcal{G}^s . Let $\boldsymbol{\lambda}^N = (\lambda^{i,s})_{i=1}^N$, with $\lambda^{i,s}$ the policy according to which the state is reflected upward at the random barrier $a(\bar{\mu}_\infty)$ given by (3.4.17). Then, for any $c > 0$, the correlated strategy profile $(\bar{\mu}_\infty, \boldsymbol{\lambda}^N)$ defines an ε_N -CCE for the N -player game within the set of strategies $\mathbb{A}_{N,c}$, with $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$.*

We omit the proof since it is completely analogous to the proof of Theorem 3.4.4: it is enough to repeat the proof of Theorem 3.4.4, invoking Lemma 3.3.1 instead of Lemma 3.7.2 to ensure integrability.

Nash equilibrium for the ergodic mean field game

Since the MFG considered here does not satisfy the assumptions of [31, Theorem 4.4], we cannot directly deduce existence and uniqueness of NE for the MFG can not be applied. Nevertheless, we have the following result:

Proposition 3.4.8. *If $\alpha + \beta \neq 1$, there exists a unique Nash equilibrium (ν^*, m_∞^*) of the ergodic MFG. Moreover, the process ν^* reflects the state process at a barrier a^* , and the pair (a^*, m_∞^*) is given by*

$$\begin{aligned}a^* &= \left(\frac{2\delta + \sigma^2}{2\delta} \right)^{\frac{\beta}{1-\alpha-\beta}} \left(\frac{2\alpha}{q(2\delta + \sigma^2(1-\alpha))} \right)^{\frac{1}{1-\alpha-\beta}}, \\ m_\infty^* &= \left(\frac{2\delta + \sigma^2}{2\delta} \right)^{\frac{1-\alpha}{1-\alpha-\beta}} \left(\frac{2\alpha}{q(2\delta + \sigma^2(1-\alpha))} \right)^{\frac{1}{1-\alpha-\beta}}.\end{aligned}\tag{3.4.21}$$

If $\alpha + \beta = 1$ and the relation

$$1 + \frac{\sigma^2}{2\delta} = \left(\frac{q\delta}{\alpha} \right)^{\frac{1}{1-\alpha}} \left(1 + (1-\alpha) \frac{\sigma^2}{2\delta} \right)^{\frac{1}{1-\alpha}}\tag{3.4.22}$$

holds, then there exist infinitely many mean field Nash equilibria given by the pair $(\nu^{a(m)}, m)$, with ν a reflection at the level $a(m) = Km^\beta$ and $m > 0$; otherwise, it does not exist any Nash equilibrium for the MFG.

Proof. Fix $m > 0$. By applying Lemma 3.3.2 with $p = m^\beta$ and $\gamma = 0$, the payoff functional $\mathfrak{J}(\nu, m)$ is maximized by the strategy ν^m which reflect the process X^{ν^m} upwards à la Skorohod at the point $a(m)$ given by

$$a(m) = \left(\frac{2\alpha}{q(2\delta + \sigma^2(1 - \alpha))} \right)^{\frac{1}{1-\alpha}} m^{\frac{\beta}{1-\alpha}}.$$

In view of (3.3.2), in order to get the consistency condition ((ii)) satisfied, we impose

$$m_\infty^* = \frac{2\delta}{2\delta + \sigma^2} K (m_\infty^*)^{\frac{\beta}{1-\alpha}}. \quad (3.4.23)$$

If $\alpha + \beta \neq 1$, this is equivalent to

$$(m_\infty^*)^{\frac{1-\alpha-\beta}{1-\alpha}} = \frac{2\delta + \sigma^2}{2\delta} K.$$

If $\alpha + \beta \neq 1$, the function $g(m) := m^{\frac{1-\alpha-\beta}{1-\alpha}}$ is always non negative, strictly monotone and with image equal to \mathbb{R}_+ , which implies that there exists a unique m_∞^* so that the above equality is verified; by direct computation, it can be verified that m_∞^* is given by (3.4.21). Finally, if $\alpha + \beta = 1$, condition (3.4.23) becomes

$$m_\infty^* = \frac{2\delta + \sigma^2}{2\delta} K m_\infty^*.$$

Thus, the MFG admits infinitely many Nash equilibria if $\frac{2\delta + \sigma^2}{2\delta} K = 1$, and none otherwise. By explicit calculations, this is equivalent to

$$\frac{2\delta + \sigma^2}{2\delta} = \left(\frac{q\delta}{\alpha} \right)^{\frac{1}{1-\alpha}} \left(\frac{2\delta + \sigma^2}{2\delta} - \alpha \frac{\sigma^2}{2\delta} \right)^{\frac{1}{1-\alpha}}.$$

By rearranging the terms, we get to (3.4.22). \square

Remark 3.4.3. Assumptions **U** and **D** are not needed to prove Proposition 3.5.1. On the other hand, according to the previous result, restrictions on the parameters are needed in order to have existence or uniqueness of the optimal control. On top of those considerations, we note that, when $\alpha + \beta = 1$ and condition (3.4.22) is not satisfied, we do not have existence of a Nash equilibrium for the ergodic MFG, while the optimality conditions for both classes \mathcal{G}^r and \mathcal{G}^s are still valid. As already noticed in Chapter 2 (see, in particular, Remark 2.3.1 and Section 2.6), the ultimate reason is that the procedure outlined in Section 3.4 does not involve the usual two steps scheme used to compute mean field NEs: first, optimize with a fixed flow of moments and, second, perform a fixed point argument to determine the flow. Actually, we first impose the consistency condition and then we restate the optimality condition.

We notice that the pair (ν^*, m_∞^*) is a correlated stationary strategy in \mathcal{G}^s with deterministic correlation device. In particular, it satisfies the optimality condition (3.4.19). As a consequence of Theorem 3.4.7, we also deduce that every NE for the ergodic MFG induces a sequence of approximate Nash equilibria with vanishing error in the N -player game:

Corollary 3.4.8.1. *Let (ν^*, m_∞^*) be a Nash equilibrium for the MFG, as given by Proposition 3.4.8. For any $i \geq 1$, let $\nu^{i,*}$ be the policy according to which the state is reflected upward at the random barrier a^* given by (3.4.21). Then, for any $c > 0$, the open-loop strategy profile $\boldsymbol{\nu}^{*,N} = (\nu^{i,*})_{i=1}^N$ defines an ε_N -NE for the N -player game within the set of strategies $\mathbb{A}_{N,c}$, with $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$.*

Proof. It is enough to notice that, when starting from the NE (ν^*, m_∞^*) the recommendation $\boldsymbol{\lambda}^N$ defined in Theorem 3.4.7 is actually an open-loop strategy profile for the N -player game. \square

3.5 Cooperative case: Mean field solution and approximation

In this section, we tackle the mean field solution of the central planner's optimization problem defined in Definition 3.2.6.

3.5.1 Mean field control solution

The first result, Theorem 3.5.1, regards existence and uniqueness of optimal solutions of the MFC problem. In order to compute the optimal control, we use a Lagrangian multiplier type approach, to take care of the constraint on the stationary first order moment. We first restrict to strategies so that the corresponding stationary mean is equal to some prescribed level $m \geq 0$, we compute the optimal strategy within this smaller constrained set and finally we optimize over all possible values of the stationary mean. A somehow similar approach has been used in [45], for a MFC problem of impulse control.

Theorem 3.5.1. *If $\alpha + \beta < 1$, there exists a unique optimal control $\hat{\nu}$ for the MFC problem. Moreover, the process upwards à la Skorohod $\hat{\nu}$ reflects the state process at the barrier \hat{a} given by*

$$\hat{a} = \left[\frac{2(\alpha + \beta)}{q(2\delta + \sigma^2(1 - \alpha))} \left(\frac{2\delta + \sigma^2}{2\delta} \right)^\beta \right]^{\frac{1}{1-\alpha-\beta}}, \quad (3.5.1)$$

and the corresponding stationary mean is given by

$$\hat{m}_\infty = \left(\frac{2\delta + \sigma^2}{2\delta} \right)^{\frac{1-\alpha}{1-\alpha-\beta}} \left[\frac{2(\alpha + \beta)}{q(2\delta + \sigma^2(1 - \alpha))} \right]^{\frac{1}{1-\alpha-\beta}}. \quad (3.5.2)$$

If instead $\alpha + \beta > 1$, the problem is ill-posed, in the sense that

$$\sup_{\nu \in \mathbb{A}_{MFC}} \mathfrak{J}(\nu, \mathbb{E}[X_\infty^\nu]) = +\infty.$$

Finally, if $\alpha + \beta = 1$ and

$$\frac{2\delta + \sigma^2}{2\delta + \sigma^2(1 - \alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha < q\delta,$$

the null control $\nu \equiv 0$ is optimal; otherwise, the problem is ill-posed.

Proof. We note that

$$\begin{aligned} & \sup_{\nu \in \mathbb{A}_{MFC}} \mathfrak{J}(\nu, \mathbb{E}[X_\infty^\nu]) \\ &= \sup_{m>0} \sup_{\substack{\nu \in \mathbb{A}_{MFC} \\ \mathbb{E}[X_\infty^\nu]=m}} \mathfrak{J}(\nu, m) = \sup_{m>0} \sup_{\substack{\nu \in \mathbb{A}_{MFC} \\ \mathbb{E}[X_\infty^\nu]=m}} (\mathfrak{J}(\nu, m) + \gamma(m)(\mathbb{E}[X_\infty^\nu] - m)), \end{aligned} \quad (3.5.3)$$

where $\gamma : \mathbb{R}_+ \rightarrow \mathbb{R}$ is any real function of the mean m . Moreover, for $\nu \in \mathbb{A}_{MFC}$ so that $\mathbb{E}[X_\infty^\nu] = m$, we rewrite the right-hand side term of (3.5.3) using ergodicity:

$$\begin{aligned} \mathfrak{J}(\nu, m) + \gamma(m)(\mathbb{E}[X_\infty^\nu] - m) &= \mathfrak{J}(\nu, m) + \lim_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E}[\gamma(m)X_t^\nu - \gamma(m)m] dt \\ &= -\gamma(m)m + \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T ((X_t^\nu)^\alpha m^\beta + \gamma(m)X_t^\nu) dt - q\nu_T \right] \\ &= -\gamma(m)m + \tilde{\mathfrak{J}}(\nu, m^\beta, \gamma(m)), \end{aligned} \quad (3.5.4)$$

where $\tilde{\mathfrak{J}}$ is defined in (3.3.5). We split the problem in the following three steps:

1) For fixed m and γ , show that there exists a unique optimal control of barrier type which maximizes $\tilde{J}(\nu, m^\beta, \gamma)$ over \mathbb{A} . Denote by $\hat{a}(m, \gamma)$ the optimal reflection barrier.

2) Show that for any $m > 0$ there exists a value $\gamma(m)$ so that $m = \mathbb{E}[X_\infty^{\nu^{\hat{a}(m, \gamma(m))}}]$ and deduce that

$$\sup_{\substack{\nu \in \mathbb{A}_{MFC} \\ \mathbb{E}[X_\infty^\nu]=m}} \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T ((X_t^\nu)^\alpha m^\beta + \gamma(m)X_t^\nu) dt - q\nu_T \right] = \tilde{\mathfrak{J}}(\nu^{\hat{a}(m, \gamma(m))}, m^\beta, \gamma(m)).$$

3) Perform the optimization over $m \in \mathbb{R}_+$.

As for Step 1), we restrict to $\gamma < q\delta$, so that, by applying Lemma 3.3.2 with $p = m^\beta$, the unique optimal control is the reflection policy at the level $\hat{a}(m, \gamma) = \hat{a}(m^\beta, \gamma)$ given by (3.3.6).

As for Step 2), we look for $\gamma(m)$ so that $\mathbb{E}[X_\infty^{\nu^{\hat{a}(m, \gamma(m))}}] = m$ and the condition $q\delta - \gamma(m) > 0$ holds. In view of (3.3.2), this is equivalent to imposing

$$\frac{2\delta + \sigma^2}{2\delta} \left(\frac{2\delta + \sigma^2(1 - \alpha)}{2\alpha\delta} \frac{q\delta - \gamma}{m^\beta} \right)^{\frac{1}{\alpha-1}} = m, \quad q\delta - \gamma(m) > 0$$

which are both satisfied by

$$\gamma(m) = q\delta - \left(\frac{2\delta + \sigma^2}{2\delta} \right)^{1-\alpha} \frac{2\delta\alpha}{2\delta + \sigma^2(1 - \alpha)} m^{\alpha+\beta-1}. \quad (3.5.5)$$

Therefore, by choosing $\gamma(m)$ as the Lagrangian multiplier in (3.5.3), we have that

$$\begin{aligned} \sup_{\nu \in \mathbb{A}_{MFC}} \mathfrak{J}(\nu, \mathbb{E}[X_\infty^\nu]) &= \sup_{m>0} \left(-\gamma(m)m + \sup_{\substack{\nu \in \mathbb{A}_{MFC} \\ \mathbb{E}[X_\infty^\nu]=m}} \tilde{\mathfrak{J}}(\nu; m, \gamma(m)) \right) \\ &= \sup_{m>0} \left(-\gamma(m)m + \tilde{\mathfrak{J}}(\nu^{\hat{a}(m, \gamma(m))}; m, \gamma(m)) \right) = \sup_{m>0} \mathfrak{J}(\nu^{\hat{a}(m, \gamma(m))}, m) \\ &= \sup_{m>0} \mathfrak{J}(\nu^{\hat{a}(m, \gamma(m))}, \mathbb{E}[X_\infty^{\nu^{\hat{a}(m, \gamma(m))})]). \end{aligned}$$

We are left with performing the optimization over $m \in \mathbb{R}_+$. By exploiting (3.3.2) and (3.3.3), for every $m > 0$ we have

$$\begin{aligned} \mathfrak{J}(\nu^{\hat{a}(m, \gamma(m))}, \mathbb{E}[X_\infty^{\nu^{\hat{a}(m, \gamma(m))})]) &= C(\hat{a}(m, \gamma(m)), m) \\ &= \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1 - \alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha m^{\alpha+\beta} - q\delta m. \end{aligned}$$

Set

$$f(m) := \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1 - \alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha m^{\alpha+\beta} - q\delta m. \quad (3.5.6)$$

If $\alpha + \beta < 1$, we have

$$f'(m) = (\alpha + \beta) \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1 - \alpha)} m^{\alpha+\beta-1} - q\delta,$$

so that $f''(m) < 0$ for every $m > 0$, i.e. f is strictly concave in \mathbb{R}_+ . This implies that there exists a unique maximizer \hat{m}_∞ . By imposing $f'(m) = 0$, we find the expression of \hat{m}_∞ given by (3.5.2), and by (3.3.2) we find the expression of \hat{a} in (3.5.1).

If $\alpha + \beta > 1$, the function f defined in (3.5.6) is unbounded, and therefore the MFC problem does not admit a maximizer. Finally, suppose $\alpha + \beta = 1$. Then, the function f is just given by

$$f(m) = \left(\frac{2\delta + \sigma^2}{2\delta + \sigma^2(1 - \alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha - q\delta \right) m.$$

Since f is linear in m , we either have $\sup_{m>0} f(m)$ equal to 0 or $+\infty$ depending on the sign of the coefficient. \square

Remark 3.5.1. As for mean field Nash equilibria (see previous Remark 3.4.3), Assumptions **U** and **D** are not needed to prove Proposition 3.5.1. On the other hand, according to the previous result, restrictions on the parameters are needed in order to have existence of the optimal control. Notice that, if an optimal control exists, it is always unique, by strict concavity of the reward functional.

For the sake of completeness and for later use, we derive the relationship between the Lagrangian multiplier $\gamma(m)$ and the constraint parameter m .

Lemma 3.5.2. *Let $\nu = \nu^{a(m)}$ be the strategy which reflect the process $X^{\nu^{a(m)}}$ upwards at the barrier $a(m) = \frac{2\delta}{2\delta + \sigma^2}m$. Let $\gamma(m)$ be given by (3.5.5) and f given by (3.5.6). Then, for any $m > 0$, it holds*

$$f'(m) + \gamma(m) = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \pi_m(X_t^{\nu^{a(m)}}, m) dt \right]. \quad (3.5.7)$$

The proof is postponed to Section 3.7.

Remark 3.5.2. The same calculations of Lemma 3.5.2 show that the following relation holds:

$$\gamma(\widehat{m}_\infty) = \partial_m \mathfrak{J}(\widehat{\nu}, \widehat{m}_\infty) = \partial_m \left(\sup_{\nu \in \mathbb{A}_{MFC}} \mathfrak{J}(\nu, \mathbb{E}[X_\infty^\nu]) \right).$$

In this sense, the Lagrange multiplier $\gamma(\widehat{m}_\infty)$ can be regarded as the derivative of the value function with respect to the measure argument.

3.5.2 Approximation

We show that any solution to the MFC problem as given by Theorem 3.5.1 induces a sequence $(\widehat{\nu}^N)_{N \geq 1}$ of approximate optimal strategy profiles for the central planner, with vanishing error. Our approach is mainly inspired by [37, Section 6], although it uses different techniques due to the nature of our dynamics and payoff.

Let $N \geq 2$. We consider the following set $\mathcal{C} \subseteq \mathbb{A}_N^N$ of strategies for the central planner.

Definition 3.5.1 (Admissible strategies for the central planner optimization problem). A strategy profile $\beta^N = (\beta^{1,N}, \dots, \beta^{N,N})$ is an admissible strategy profile for the central planner optimization problem if $\beta^{i,N} \in \mathbb{A}_N$ for any $i = 1, \dots, N$, if they are exchangeable, in the sense that the triples $(\xi^i, \beta^{i,N}, W^i)_{i=1}^N$ are exchangeable random elements, and ergodic, in the sense that the resulting N -dimensional process $(X_t^{\beta^{1,N}}, \dots, X_t^{\beta^{N,N}})_{t \geq 0}$ is ergodic, where $X^{\beta^{i,N}}$ is given by (3.1.1) for every $i = 1, \dots, N$. Moreover, every strategy β^j is such that

$$\sup_{T > 0} \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^{\beta^j}|^2 dt \right] \leq c, \quad \forall j = 1, \dots, N. \quad (3.5.8)$$

We denote the set of strategy profiles for the central planner by $\mathbb{A}_N^{cp,c}$.

Observe that the inclusion $\mathbb{A}_N^{cp,c} \subseteq \mathbb{A}_N^N$ holds strictly. The exchangeability assumption is well known in the MFC literature, when dealing with approximation results. We refer to [39, Paragraph 6.1.3] for more comments. Observe that, in particular, exchangeability implies

$$\bar{\mathfrak{J}}^N(\beta^N) = \mathfrak{J}^N(\beta^{i,N}, \beta^{-i,N}),$$

for every $i = 1, \dots, N$, for any strategy profile $\beta^{(N)} \in \mathbb{A}_N^{cp,c}$.

Theorem 3.5.3. *Let $\widehat{\nu}^N = (\widehat{\nu}^1, \dots, \widehat{\nu}^N)$ be the strategy profile that reflects each process $X^{\widehat{\nu}^i}$ upwards à la Skorohod at the level \widehat{a} given by (3.5.1). It holds*

$$\lim_{N \rightarrow \infty} \sup_{\beta^{(N)} \in \mathbb{A}_N^{cp,c}} \bar{\mathfrak{J}}^N(\beta^{(N)}) = \lim_{N \rightarrow \infty} \bar{\mathfrak{J}}^N(\nu^{(N)}) = \mathfrak{J}(\widehat{\nu}, \mathbb{E}[X_\infty^{\widehat{\nu}}]). \quad (3.5.9)$$

Proof. Notice that, for every $N \geq 2$, $\widehat{\nu}^N$ belongs to \mathbb{A}_N^{cp} : since the sequence $(\widehat{\nu}^i, X^{\widehat{\nu}^i})_{i \geq 1}$ is i.i.d. as $(\widehat{\nu}, X^{\widehat{\nu}})$, the system is exchangeable. Moreover, by [80, Lemmata 23.17-19], the N -dimensional process $(X^{\widehat{\nu}^i})_{i=1}^N$ is a positively recurrent regular diffusion with ergodic measure $\bigotimes_{i=1}^N \widehat{p}_\infty(dx_i)$, where we set $\widehat{p}_\infty = p_{\widehat{a}}$, with p_a given by (3.3.1). Moreover, up to choosing c large enough, the processes $(X^{\widehat{\nu}^i})_{i=1}^N$ satisfy (3.5.1) by point iii) of Lemma 3.3.1.

We show that

$$\lim_{N \rightarrow \infty} \inf_{\beta^{(N)} \in \mathbb{A}_N^{cp,c}} (\mathfrak{J}^{MFC}(\widehat{\nu}) - \bar{\mathfrak{J}}^N(\beta^{(N)})) \geq 0, \quad \lim_{N \rightarrow \infty} (\mathfrak{J}^{MFC}(\widehat{\nu}) - \bar{\mathfrak{J}}^N(\nu^{(N)})) = 0.$$

Observe that, by Lemma 3.3.1, the inferior limit in the definition of $\mathfrak{J}^{MFC}(\widehat{\nu}^i)$ is actually a limit. Then, by using the inequalities $\overline{\lim}_n z_n - \underline{\lim}_n x_n \geq \overline{\lim}_n (z_n - y_n)$ and $\overline{\lim}_n (z_n + y_n) \geq \overline{\lim}_n z_n + \underline{\lim}_n (y_n)$ and concavity of $\pi(x, m) = x^\alpha m^\beta$ jointly in (x, m) , it holds

$$\begin{aligned} & \mathfrak{J}^{MFC}(\widehat{\nu}^i) - \bar{\mathfrak{J}}^N(\beta^{i,N}, \beta^{-i,N}) \\ & \geq \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \left(\pi(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty) - \pi(X_t^{\beta^{i,N}}, \bar{\mu}_t^{N, \beta^{-i,N}}) \right) dt - q(\widehat{\nu}_T^i - \beta_T^{i,N}) \right] \\ & \geq \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \left(\pi_x(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty) + \gamma(\widehat{m}_\infty) \right) (X_t^{\widehat{\nu}^i} - X_t^{\beta^{i,N}}) dt - q(\widehat{\nu}_T^i - \beta_T^{i,N}) \right] \\ & \quad + \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \left(-\gamma(\widehat{m}_\infty)(X_t^{\widehat{\nu}^i} - X_t^{\beta^{i,N}}) + \pi_m(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty)(\widehat{m}_\infty - \bar{\mu}_t^{N, \beta^{-i,N}}) \right) dt \right] \\ & \geq \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T -\gamma(\widehat{m}_\infty)(X_t^{\widehat{\nu}^i} - X_t^{\beta^{i,N}}) dt \right] \\ & \quad + \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \pi_m(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty)(\widehat{m}_\infty - \bar{\mu}_t^{N, \beta^{-i,N}}) dt \right] \end{aligned} \tag{3.5.10}$$

where we added and subtracted $\gamma(\widehat{m}_\infty)(X_t^{\widehat{\nu}^i} - X_t^{\beta^{i,N}})$ inside the time integral. Last inequality follows from sublinearity of the $\underline{\lim}$ and from Lemma 3.3.3, as the pair $(X^{\widehat{\nu}^i}, \widehat{\nu}^i)$ has the same distribution as $(X^{\widehat{\nu}}, \widehat{\nu})$. Moreover, since $(X^{\widehat{\nu}^i})_{i \geq 1}$ are i.i.d. copies of $X^{\widehat{\nu}}$, by Lemma 3.5.2, it holds

$$\gamma(\widehat{m}_\infty) = \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \pi_m(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty) dt \right] \quad \forall i \geq 1.$$

Since the second $\underline{\lim}$ in the last line of (3.5.10) is actually a limit, by using ergodicity and the identical distribution of $(X^{\widehat{\nu}^i})_{i=1}^N$ and $(X^{\beta^{i,N}})_{i=1}^N$ respectively, we deduce

$$\begin{aligned} & \mathfrak{J}^{MFC}(\widehat{\nu}^i) - \bar{\mathfrak{J}}^N(\beta^{i,N}, \beta^{-i,N}) \\ & \geq \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \left(\pi_m(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty) - \gamma(\widehat{m}_\infty) \right) \left(\widehat{m}_\infty - \bar{\mu}_t^{N, \beta^{-i,N}} \right) dt \right] \\ & = \frac{N}{N-1} \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \left(\pi_m(X_t^{\widehat{\nu}^i}, \widehat{m}_\infty) - \gamma(\widehat{m}_\infty) \right) \left(\frac{1}{N} \sum_{j=1}^N (\widehat{m}_\infty - X_t^{\beta^{j,N}}) \right. \right. \\ & \quad \left. \left. - \frac{1}{N} (\widehat{m}_\infty - X_t^{\beta^{i,N}}) \right) dt \right], \end{aligned} \tag{3.5.11}$$

where we used the identity $\frac{1}{N-1} \sum_{j \neq 1} y_j = \frac{N}{N-1} (\frac{1}{N} \sum_{i=1}^N y_i - \frac{1}{N} y_1)$. By taking the average over N , we have

$$\begin{aligned} \mathfrak{J}^{MFC}(\hat{\nu}) - \bar{\mathfrak{J}}^N(\beta^{(N)}) &\geq \frac{N}{N-1} \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left(\frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right) \right. \\ &\quad \cdot \left. \left(\frac{1}{N} \sum_{j=1}^N (\hat{m}_\infty - X_t^{\beta^{j,N}}) \right) - \frac{1}{N^2} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) (\hat{m}_\infty - X_t^{\beta^{i,N}}) \right] dt \\ &\geq -c \left(\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \int_0^T \left[\left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right| \left| \frac{1}{N} \sum_{j=1}^N (\hat{m}_\infty - X_t^{\beta^{j,N}}) \right| \right] dt \right. \\ &\quad \left. + \frac{1}{N} \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) (\hat{m}_\infty - X_t^{\beta^{i,N}}) \right| \right] dt \right). \end{aligned}$$

We study separately the two $\overline{\lim}$. As for the the second one, by taking advantage the exchangeability of $(X_t^{\hat{\nu}^i})_{i \geq 1}$ and $(X_t^{\beta^{i,N}})_{i=1}^N$, we get

$$\begin{aligned} &\frac{1}{N} \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) (\hat{m}_\infty - X_t^{\beta^{i,N}}) \right| \right] dt \\ &\leq \frac{1}{N} \overline{\lim}_{T \uparrow \infty} \left(\frac{1}{T} \int_0^T \mathbb{E} \left[\left| \pi_m(X_t^{\hat{\nu}}, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right|^2 \right] dt \right)^{\frac{1}{2}} \left(\frac{1}{T} \int_0^T \mathbb{E} \left[\left| \hat{m}_\infty - X_t^{\beta^{1,N}} \right|^2 \right] dt \right)^{\frac{1}{2}} \\ &\leq \frac{c}{N} \overline{\lim}_{T \uparrow \infty} \left(\frac{1}{T} \int_0^T \mathbb{E} \left[\left| \pi_m(X_t^{\hat{\nu}}, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right|^2 \right] dt \right)^{\frac{1}{2}}, \end{aligned}$$

that goes to 0 as $N \rightarrow \infty$, by definition of the set of strategy profiles $\mathbb{A}_N^{cp,c}$. As for the first term, by analogous computations, we get

$$\begin{aligned} &\overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \int_0^T \left[\left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right| \left| \frac{1}{N} \sum_{j=1}^N (\hat{m}_\infty - X_t^{\beta^{j,N}}) \right| \right] dt \\ &\leq \overline{\lim}_{T \uparrow \infty} \left(\frac{1}{T} \mathbb{E} \int_0^T \left[\left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right|^2 \right] dt \right)^{\frac{1}{2}} \\ &\quad \cdot \left(\frac{1}{T} \int_0^T \mathbb{E} \left[\left| \frac{1}{N} \sum_{j=1}^N (\hat{m}_\infty - X_t^{\beta^{j,N}}) \right|^2 \right] dt \right)^{\frac{1}{2}} \\ &\leq c \overline{\lim}_{T \uparrow \infty} \left(\frac{1}{T} \mathbb{E} \int_0^T \left[\left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right|^2 \right] dt \right)^{\frac{1}{2}}, \end{aligned}$$

for a constant c independent of N , by using the bound (3.5.8) and exchangeability of $\beta^{(N)} \in \mathbb{A}_N^{cp,c}$. By the ergodic ratio theorem, it holds

$$\begin{aligned} &\lim_{T \uparrow \infty} \frac{1}{T} \int_0^T \left| \frac{1}{N} \sum_{i=1}^N (\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right|^2 dt \\ &= \int_{\mathbb{R}_+^N} \left| \frac{1}{N} \sum_{i=1}^N (\pi_m(x_i, \hat{m}_\infty) - \gamma(\hat{m}_\infty)) \right|^2 \bigotimes_{i=1}^N \hat{p}_\infty(dx_i), \quad (3.5.12) \end{aligned}$$

\mathbb{P} -a.s.. We show that the right hand-side is uniformly integrable. Take $r = 1/\alpha > 1$. By Jensen inequality and identical distribution of $(X^{\hat{\nu}^i})_{i \geq 1}$, we have

$$\begin{aligned}
& \mathbb{E} \left[\left(\frac{1}{T} \int_0^T \left| \frac{1}{N} \sum_{i=1}^N \left(\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right) \right|^2 dt \right)^r \right] \\
& \leq \frac{C}{T} \mathbb{E} \left[\int_0^T \left| \left(\frac{1}{N} \sum_{i=1}^N \pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right) \right|^{2r} dt \right] \\
& \leq C \left(|\gamma(\hat{m}_\infty)|^{2r} + \frac{1}{N} \sum_{i=1}^N \frac{1}{T} \int_0^T \mathbb{E} \left[|\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty)|^{2r} \right] dt \right) \\
& \leq C \left(1 + \frac{1}{T} \int_0^T \mathbb{E} \left[|\pi_m(X_t^{\hat{\nu}}, \hat{m}_\infty)|^{2r} \right] dt \right) \leq C \left(1 + \frac{1}{T} \int_0^T \mathbb{E} \left[|(X_t^{\hat{\nu}})^\alpha|^{2r} \right] dt \right) \\
& = C \left(1 + \frac{1}{T} \int_0^T \mathbb{E} \left[|(X_t^{\hat{\nu}})|^2 \right] dt \right) \leq C(1 + \hat{a}^2),
\end{aligned}$$

where we used Lemma 3.3.1 in the last estimate. Since last estimate holds for any $T > 0$ and $r > 1$, we deduce that the right hand-side of (3.5.12) is uniformly integrable. By, e.g., [80, Lemma 4.12]) this yields

$$\begin{aligned}
& \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \left| \frac{1}{N} \sum_{i=1}^N \left(\pi_m(X_t^{\hat{\nu}^i}, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right) \right|^2 dt \right] \\
& = \int_{\mathbb{R}_+^N} \left| \frac{1}{N} \sum_{i=1}^N \left(\pi_m(x_i, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right) \right|^2 \bigotimes_{i=1}^N \hat{p}_\infty(dx_i).
\end{aligned}$$

We conclude by invoking the law of large numbers: let $(X_i)_{i \geq 1}$ be a sequence of i.i.d. random variables with law $\hat{p}_\infty(dx)$. Therefore, by Lemmata 3.3.1 and 3.5.2, the sequence $(Y_i)_{i \geq 1}$ defined by $Y_i = \pi_m(X_i, \hat{m}_\infty) - \gamma(\hat{m}_\infty)$ is i.i.d., square integrable and centered, which yields

$$\int_{\mathbb{R}_+^{N-1}} \left| \frac{1}{N} \sum_{i=1}^N \left(\pi_m(x_i, \hat{m}_\infty) - \gamma(\hat{m}_\infty) \right) \right|^2 \bigotimes_{i=1}^N \hat{p}_\infty(dx_i) = \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N Y_i \right|^2 \right],$$

which goes to 0 as $N \rightarrow \infty$. Finally, $\lim_{N \rightarrow \infty} \tilde{\mathfrak{J}}^N(\boldsymbol{\nu}^{(N)}) = \mathfrak{J}^{MFC}(\hat{\nu})$ follows from ergodicity of the processes $(X^{\hat{\nu}^i})_{i=1}^N$ and the law of large numbers, analogously as before. \square

3.6 Numerical illustrations

In this section, we numerically illustrate our previous findings. In particular, we exhibit possible choices of distribution of $\bar{\mu}_\infty$ so that the correlated stationary strategies $(Z, \bar{\mu}_\infty, \lambda^r)$ in \mathcal{G}^r and $(Z, \bar{\mu}_\infty, \lambda^s)$ in \mathcal{G}^s are mean field CCEs, i.e., according to Propositions 3.4.3 and 3.4.6, the inequalities (3.4.9) and (3.4.19) respectively are verified.

We suppose that $\bar{\mu}_\infty$ is distributed as a Gamma with $u > 0$ and scale parameter $v > 0$. Then, for any $k \geq 0$ the k -th moment of $\bar{\mu}_\infty \sim \Gamma(u, v)$ is given by

$$\mathbb{E}[\bar{\mu}_\infty^k] = \frac{1}{\Gamma(u)v^u} \int_0^\infty x^k x^{u-1} e^{-\frac{x}{v}} dx = \frac{\Gamma(u+k)}{\Gamma(u)} v^k.$$

By assuming Gamma distribution on $\bar{\mu}_\infty$, the optimality inequalities (3.4.9) and (3.4.19) for the regular and the singular recommendation become, respectively,

$$c_\beta \left(\frac{\Gamma(\beta+u)}{\Gamma(u)} \right)^{\frac{1}{1-\alpha}} v^{\frac{\beta}{1-\alpha}} \leq c_{\alpha+\beta} \frac{\Gamma(\alpha+\beta+u)}{\Gamma(u)} v^{\alpha+\beta} - c_1 uv, \quad (3.6.1)$$

$$c_\beta \left(\frac{\Gamma(\beta+u)}{\Gamma(u)} \right)^{\frac{1}{1-\alpha}} v^{\frac{\beta}{1-\alpha}} \leq \tilde{c}_{\alpha+\beta} \frac{\Gamma(\alpha+\beta+u)}{\Gamma(u)} v^{\alpha+\beta} - c_1 uv, \quad (3.6.2)$$

Given the non-linear structure of the optimality inequalities and the intricate dependence on the parameters in the constants c_1 , c_β , $c_{\alpha+\beta}$ and $\tilde{c}_{\alpha+\beta}$, we limit ourselves to a specific choice of the parameters. For the sake of illustrations, we fix $\delta = 0.1$ and $\sigma = 0.2$. Notice that this choice satisfies Assumption **D**. We also set $q = 2$.

The case $\alpha + \beta < 1$

In this case, there exist both a unique mean field NE and a unique optimal control for the MFC problem. For the sake of comparison with the payoffs of the NE and the MFC solution, we are also interested in finding values of $(u, v) \in \mathbb{R}_+^2$ so that the reward of the associated CCE is higher than the reward of the NE. Therefore, we pair equations (3.6.1) and (3.6.2) with, respectively,

$$c_{\alpha+\beta} \frac{\Gamma(\alpha+\beta+u)}{\Gamma(u)} v^{\alpha+\beta} - c_1 uv \geq \tilde{c}_{\alpha+\beta} (m_\infty^*)^{\alpha+\beta} - c_1 m_\infty^*, \quad (3.6.3)$$

$$\tilde{c}_{\alpha+\beta} \frac{\Gamma(\alpha+\beta+u)}{\Gamma(u)} v^{\alpha+\beta} - c_1 uv \geq \tilde{c}_{\alpha+\beta} (m_\infty^*)^{\alpha+\beta} - c_1 m_\infty^*. \quad (3.6.4)$$

Figure 3.1 shows that there exist infinitely many mean field CCEs both in \mathcal{G}^r and \mathcal{G}^s which yield an higher reward than the Nash equilibrium (ν^*, m_∞^*) . Here, we set $\alpha = 0.3$ and $\beta = 0.5$.

Figure 3.2 shows the reward associated to those mean field CCEs in \mathcal{G}^r and \mathcal{G}^s that yield an higher reward than the Nash equilibrium (ν^*, m_∞^*) . The improvement on the Nash equilibrium is $\approx 17\%$ of the reward yielded by the mean field control solution $\hat{\nu}$ in the singular case, and $\approx 12\%$ in the regular case. We notice that the payoff associated to the Nash equilibrium for the ergodic MFG is strictly less than the reward given by the solution of the MFC problem. This can be directly deduced from the fact that the stationary mean \hat{m}_∞ associated to the MFC solution $\hat{\nu}$ is the unique maximizer of the function $f(m)$ given by (3.5.6), which can be equivalently expressed as $f(m) = \tilde{c}_{\alpha+\beta} m^{\alpha+\beta} - c_1 m$. Since by Proposition 3.4.8 the value of the ergodic MFG at the Nash equilibrium (ν^*, m_∞^*) can be expressed as $f(m_\infty^*)$, we deduce

$$\mathfrak{J}(\nu^*, m_\infty^*) = f(m_\infty^*) < f(\hat{m}_\infty) = \mathfrak{J}(\hat{\nu}, \mathbb{E}[X_\infty^{\hat{\nu}}]). \quad (3.6.5)$$




Figure 3.1: Values of the parameters $(u, v) \in \mathbb{R}_+^2$ so that the $(Z, \lambda^s, \bar{\mu}_\infty) \in \mathcal{G}^s$ (on the left) and $(Z, \lambda^r, \bar{\mu}_\infty) \in \mathcal{G}^r$ (on the right), $\bar{\mu}_\infty \sim \Gamma(u, v)$ is a mean field CCE outperforming the NE.

The reward of the MFC solution $\mathfrak{J}(\hat{\nu}, \mathbb{E}[X_\infty^{\hat{\nu}}])$ appears to be an unattainable upper bound for the set of mean field CCEs payoffs. While we limit to empirically observe this phenomenon, we point out that it is widely expected and that it is coherent with the findings of Chapter 2 for linear-quadratic mean field games.

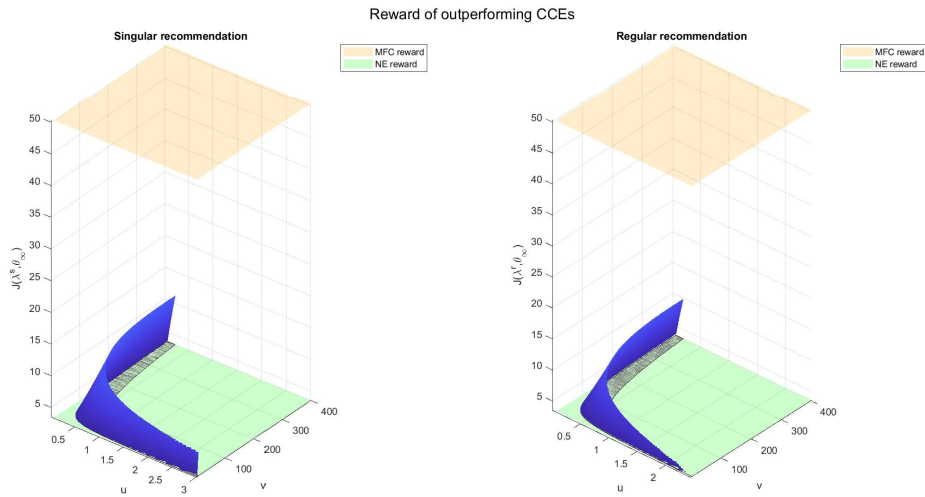


Figure 3.2: Reward associated to mean field CCEs $(Z, \lambda^s, \bar{\mu}_\infty) \in \mathcal{G}^s$ (on the left) and $(Z, \lambda^r, \bar{\mu}_\infty) \in \mathcal{G}^r$ (on the right) which outperform the reward of the mean field NE, when $\bar{\mu}_\infty \sim \Gamma(u, v)$, $(u, v) \in \mathbb{R}_+^2$.

The case $\alpha + \beta = 1$

In this case, by Theorem 3.4.8, either there does not exist any mean field NE or there exist infinitely many, depending on whether the relation (3.4.22) is satisfied. As noticed in Remark 3.4.3, the optimality inequalities (3.4.9) and (3.4.19) are still valid. By imposing $\beta = 1 - \alpha$, and noticing that the parameter v is not anymore relevant in the inequalities, equations (3.6.1) and (3.6.2) can be written in terms of u only:

$$c_\beta \left(\frac{\Gamma(1 - \alpha + u)}{\Gamma(u)} \right)^{\frac{1}{1-\alpha}} \leq (c_{\alpha+\beta}\Gamma(u) - c_1) u, \quad (3.6.6)$$

$$c_\beta \left(\frac{\Gamma(1 - \alpha + u)}{\Gamma(u)} \right)^{\frac{1}{1-\alpha}} \leq (\tilde{c}_{\alpha+\beta}\Gamma(u) - c_1) u. \quad (3.6.7)$$

We observe that there exist maximal values u_r^* and u_s^* , depending on α , so that the inequalities (3.6.6) and (3.6.7) are verified by any $0 < u \leq u_r^*$ and $0 < u \leq u_s^*$, respectively. Figure 3.3 plots such maximal values u_r^* and u_s^* as functions of $\alpha \in (0, 1)$. Therefore, we observe existence of mean field CCEs even in the case in which there does not exist any mean field NE. We notice that, when considering correlated stationary strategies $(Z, \lambda^s, \bar{\mu}_\infty)$, there exist a value $\bar{\alpha}$ so that that the inequality (3.6.6) is verified for any $u > 0$, i.e. $u^* = \infty$. It can be numerically shown that such value $\bar{\alpha}$ is the unique solution of (3.4.22), for fixed δ , σ and q . We can explain this phenomenon as follows: for $\alpha = \bar{\alpha}$, by Theorem 3.4.8, for any $m > 0$ the pair $(\nu^{a(m)}, m)$ is a mean field NE, where $\nu^{a(m)}$ is the policy that reflects the process $X^{\nu^{a(m)}}$ upwards at the level $a(m) = \frac{2\delta}{2\delta + \sigma^2} m$. Therefore, any correlated stationary strategy $(Z, \lambda^s, \bar{\mu}_\infty) \in \mathcal{G}^s$ is just a randomization, or a mixture, of mean field NEs, since λ^s reflects the process X^{λ^s} at the same barrier $a(\bar{\mu}_\infty) = \frac{2\delta}{2\delta + \sigma^2} \bar{\mu}_\infty$. To put in other terms, the pair $(a(\bar{\mu}_\infty), \bar{\mu}_\infty)$ is supported on the set of mean field NEs. This implies that the optimality condition is satisfied by any $\bar{\mu}_\infty$ so that the optimality inequality (3.4.19) holds true, and so by any $(Z, \lambda^s, \bar{\mu}_\infty) \in \mathcal{G}^s$.

The case $\alpha + \beta > 1$

The case $\alpha + \beta > 1$ is completely analogous to the case $\alpha + \beta < 1$. We just observe that, in this case, the MFC problem is ill-posed, in the sense of Theorem 3.5.1 and therefore a-priori we do not have any upper-bound on the set of mean field CCEs payoffs.

3.7 Auxiliary results

Control-theoretic results

Proof of Lemma 3.3.1. Point i) follows from direct calculations. As for point ii), the solution of the Skorohod problem follows from standard arguments (see, e.g. [83, Proposition 3.6.16]). As for ergodicity, note that the derivative of the speed measure of the process X^{ν^a} is given by $\mathbf{m}'(x) = \frac{2}{\sigma^2} x^{-\frac{2\delta}{\sigma^2}} \mathbf{1}_{[a, \infty)}(x)$, which is integrable over $[0, \infty)$ for any $a > 0$. By [21, Paragraph 36], the process X^{ν^a} is ergodic and admits $\frac{1}{\mathbf{m}([a, \infty))} \mathbf{m}'(x) dx = p_a(dx)$ given by (3.3.1) as the unique invariant distribution.




Figure 3.3: Value of u^* as α varies in $[0, 1]$, both for regular and singular recommendations. The blue dashed line is located at the value of $\alpha = \bar{\alpha}$ which satisfies (3.4.22) for fixed δ , σ and q .

As for iii), set $L = 1 + a$ and observe that, since $a < L$ and the control ν^a never acts when the process lies in the region $\{x : x > a\}$, it holds $\text{supp}(d\nu^a) \cap \{X^{\nu^a} \geq L\} = \emptyset$, \mathbb{P} -a.s.. The proof of [31, Lemma 2] implies that

$$\mathbb{E}[|X_t^{\nu^a}|^2] \leq 2L^2 + \mathbb{E}[\xi^2].$$

By definition of L and Assumption **D**, the estimate follows. Finally, By (3.2.1), we have, for any $T > 0$

$$\nu_T^a = X_T^{\nu^a} - \xi + \delta \int_0^T X_s^{\nu^a} ds - \int_0^T \sigma X_s^{\nu^a} dW_s.$$

By taking the expectation, applying Jensen inequality and taking advantage of Itô isometry, we find

$$\begin{aligned} \frac{1}{T^2} \mathbb{E}[|\nu_T^a|^2] &\leq \frac{c}{T^2} \left(1 + \mathbb{E}[|X_T^{\nu^a}|^2] + \mathbb{E} \left[\left(\int_0^T X_s^{\nu^a} ds \right)^2 \right] + \mathbb{E} \left[\left(\int_0^T X_s^{\nu^a} dW_s \right)^2 \right] \right) \\ &\leq \frac{c}{T^2} \left(1 + (T + T^2) \sup_{T \geq 0} \mathbb{E}[(X_T^{\nu^a})^2] \right) \leq c(1 + a^2), \end{aligned}$$

where last inequality follows from previous estimate. \square

Proof of Lemma 3.3.2. We sketch the proof, which is essentially similar to the proof of [55, Theorem 2]. Recall from 3.3.4 the definition of the function $g(x, p, \lambda)$. Let \mathcal{T} be the set of \mathbb{F} -stopping times. Consider the auxiliary optimal stopping problem

$$u(x, p, \gamma) := \inf_{\tau \in \mathcal{T}} \mathbb{E}_x \left[\int_0^\tau e^{-\delta t} g_x(\hat{X}_t, p, \gamma) dt + q e^{-\delta \tau} \right], \quad (3.7.1)$$

where $\hat{X} = (\hat{X}_t)_{t \geq 0}$ is defined by

$$d\hat{X}_t = (-\delta + \sigma^2)\hat{X}_t dt + \sigma\hat{X}_t dW_t,$$

and $\mathbb{E}_x[\cdot]$ denotes $\mathbb{E}[\cdot | \hat{X}_0 = x]$, $x \in \mathbb{R}_+$. Let $\hat{\phi}_0$ is the non-increasing fundamental solution of

$$\frac{1}{2}\sigma^2 x^2 u_{xx}(x) + (-\delta + \sigma^2)xu_x(x) - \delta u(x) = 0,$$

and let $\hat{\mathbf{m}}'$ be the density of the speed measure $\hat{\mathbf{m}}$ of the process \hat{X} . By the same reasoning of [4, Theorem 5], it can be shown that, if there exists a unique $\hat{a}(p, \gamma) > 0$ solution to

$$\int_{\hat{a}(p, \gamma)}^{+\infty} \hat{\phi}_0(y) (\alpha p y^{\alpha-1} - (q\delta - \gamma)) \hat{\mathbf{m}}'(y) dy = 0, \quad (3.7.2)$$

then the value function $u(\cdot, p, \gamma)$ is $C^1(\mathbb{R}_+)$ with $u_{xx}(\cdot, p, \gamma) \in L_{\text{loc}}^\infty(\mathbb{R}_+)$, and that the optimal stopping time is given by $\hat{\tau}(x, p, \gamma) := \inf\{t \geq 0 \mid \hat{X}_t \leq \hat{a}(p, \gamma)\}$.

By explicit calculations, we have $\hat{\phi}_0(y) = y^{-1}$ and $\hat{\mathbf{m}}'(y) = \frac{2}{\sigma^2} y^{-\frac{2\delta}{\sigma^2}}$, so that (3.7.2) becomes

$$\int_{\hat{a}(p, \gamma)}^{+\infty} (\alpha p y^{\alpha-1} - (q\delta - \gamma)) y^{-\frac{2\delta}{\sigma^2}-1} dy = 0. \quad (3.7.3)$$

Since, by assumption, $q\delta - \gamma > 0$, there exists a unique solution $\hat{a}(p, \gamma)$ given by (3.3.6). The proof can be then completed by the same methods of Step B.1 in [55, Appendix B]. \square

Proof of Lemma 3.3.3. We deal with the case $q' < \infty$. The case $q' = \infty$ is completely analogous. We start by proving (a). Let $\hat{\nu}$ be optimal for the control problem with dynamics (3.2.1) and payoff functional $\tilde{J}(\cdot, p, \gamma)$. Recall the definition of the function $g(x, p, \gamma)$ in (3.3.4). For any $\nu \in \mathcal{S}$, $\varepsilon \in (0, \frac{1}{2}]$, set $\nu^\varepsilon = \varepsilon\nu + (1 - \varepsilon)\hat{\nu}$. Set

$$f(\varepsilon, T) = \frac{1}{\varepsilon} \left(\frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^{\nu^\varepsilon}, p, \gamma) dt - q\nu_T^\varepsilon \right] - \frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^{\hat{\nu}}, p, \gamma) dt - q\hat{\nu}_T \right] \right). \quad (3.7.4)$$

For any $\varepsilon \in (0, \frac{1}{2}]$ and $T > 0$, it holds

$$\begin{aligned} f(\varepsilon, T) &= \frac{1}{\varepsilon} \frac{1}{T} \left(\mathbb{E} \left[\int_0^T \left(\int_0^1 g_x(X_t^{\hat{\nu}} + \tau(X_t^{\nu^\varepsilon} - X_t^{\hat{\nu}}), p, \gamma) d\tau \right) (X_t^{\nu^\varepsilon} - X_t^{\hat{\nu}}) dt - q(\nu_T^\varepsilon - \hat{\nu}_T) \right] \right) \\ &= \frac{1}{T} \left(\mathbb{E} \left[\int_0^T \left(\int_0^1 g_x(X_t^{\hat{\nu}} + \tau(X_t^{\nu^\varepsilon} - X_t^{\hat{\nu}}), p, \gamma) d\tau \right) (X_t^{\nu^\varepsilon} - X_t^{\hat{\nu}}) dt - q(\nu_T^\varepsilon - \hat{\nu}_T) \right] \right). \end{aligned}$$

We claim

$$\lim_{\varepsilon \downarrow 0} f(\varepsilon, T) = \frac{1}{T} \mathbb{E} \left[\int_0^T g_x(X_t^{\hat{\nu}}, p, \gamma) (X_t^{\nu^\varepsilon} - X_t^{\hat{\nu}}) dt - q(\nu_T^\varepsilon - \hat{\nu}_T) \right] \quad (3.7.5)$$

uniformly in T . Indeed,

$$\begin{aligned}
& \left| f(\varepsilon, T) - \frac{1}{T} \mathbb{E} \left[\int_0^T g_x(X_t^{\hat{\nu}}, p, \gamma)(X_t^\nu - X_t^{\hat{\nu}}) dt - q(\nu_T - \hat{\nu}_T) \right] \right| \\
& \leq \left| \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\hat{\nu}} - X_t^\nu) \int_0^1 (g_x(X_t^{\nu^\varepsilon} + \tau(X_t^{\hat{\nu}} - X_t^{\nu^\varepsilon}), p, \gamma) - g_x(X_t^{\hat{\nu}}, p, \gamma)) d\tau dt \right] \right| \\
& \leq \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^{\hat{\nu}} - X_t^\nu| \int_0^1 |g_x(X_t^{\nu^\varepsilon} + \tau(X_t^{\hat{\nu}} - X_t^{\nu^\varepsilon}), p, \gamma) - g_x(X_t^{\hat{\nu}}, p, \gamma)| d\tau dt \right].
\end{aligned}$$

By continuity of $g_x(x, p, \gamma)$ in x , the inner integral converges to 0 as $\varepsilon \downarrow 0$ for any $t \geq 0$, \mathbb{P} -a.s.. By linearity of the dynamics (3.2.1), since X_t^ν is positive for any t , it holds

$$\begin{aligned}
X_t^{\nu^\varepsilon} + \tau(X_t^{\hat{\nu}} - X_t^{\nu^\varepsilon}) &= \tau X_t^{\hat{\nu}} + (1 - \tau) X_t^{\nu^\varepsilon} \\
&\geq \frac{1}{2} X_t^{\nu^\varepsilon} = \frac{1}{2} (\varepsilon X_t^\nu + (1 - \varepsilon) X_t^{\hat{\nu}}) \geq \frac{1}{2} X_t^{\hat{\nu}} \geq \frac{1}{4} X_t^{\hat{\nu}}.
\end{aligned}$$

Since $|g_{xx}(y, p, \gamma)| = \alpha(1 - \alpha) p x^{\alpha-2} \leq c|x|^{\alpha-2}$ for any $y \geq x$, $g_x(y, p, \gamma)$ is Lipschitz on $[x, +\infty)$ with Lipschitz constant $c|x|^{\alpha-2}$. Then, it follows

$$\begin{aligned}
|f(\varepsilon, T) - g(T)| &\leq \frac{c}{T} \mathbb{E} \left[\int_0^T |X_t^{\hat{\nu}} - X_t^\nu| |X_t^{\hat{\nu}}|^{\alpha-2} \int_0^1 |X_t^{\nu^\varepsilon} + \tau(X_t^{\hat{\nu}} - X_t^{\nu^\varepsilon}) - X_t^{\hat{\nu}}| d\tau dt \right] \\
&\leq c \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^{\hat{\nu}} - X_t^\nu| \cdot |X_t^{\nu^\varepsilon} - X_t^{\hat{\nu}}| |X_t^{\hat{\nu}}|^{\alpha-2} dt \right] \\
&\leq c \left(\sup_{T>0} \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^{\hat{\nu}}|^{2q} dt \right] + \sup_{T>0} \frac{1}{T} \mathbb{E} \left[\int_0^T |X_t^\nu|^{2q} dt \right] \right) \\
&\leq \sup_{T>0} \left(\frac{1}{T} \mathbb{E} \left[\int_0^T |(X_t^{\hat{\nu}})^{\alpha-2}|^{q'} dt \right] \right)^{\frac{1}{q}} \varepsilon,
\end{aligned} \tag{3.7.6}$$

where we used Hölder's inequality together with conditions (3.3.7) and (3.3.8). On the other hand, by taking the limit with respect to T , it holds

$$\begin{aligned}
\lim_{T \uparrow \infty} f(\varepsilon, T) &\leq \frac{1}{\varepsilon} \left(\lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^{\nu^\varepsilon}, p, \gamma) dt - q\nu_T^\varepsilon \right] \right. \\
&\quad \left. - \lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g(X_t^{\hat{\nu}}, p, \gamma) dt - q\hat{\nu}_T \right] \right) = \frac{1}{\varepsilon} (\tilde{\mathfrak{J}}(\nu^\varepsilon) - \tilde{\mathfrak{J}}(\hat{\nu})) \leq 0,
\end{aligned}$$

by using the inequality $\liminf_n a_n - \liminf_n b_n \geq \liminf_n (a_n - b_n)$ and by optimality, for any $\varepsilon \in (0, \frac{1}{2}]$. Lemma 3.7.5 then implies

$$\lim_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g_x(X_t^{\hat{\nu}}, p, \gamma)(X_t^\nu - X_t^{\hat{\nu}}) dt - q(\nu_T - \hat{\nu}_T) \right] \leq 0.$$

This concludes the proof of point (a).

As for point (b), assume that both (3.3.11) and (a) hold. Then, by using the inequality $\overline{\lim}_n z_n - \underline{\lim}_n x_n \geq \overline{\lim}_n (z_n - y_n)$ (see [82, Equation (4.25)]) and concavity of g jointly in x , it holds

$$\begin{aligned} \tilde{\mathfrak{J}}(\hat{\nu}, p, \gamma) - \tilde{\mathfrak{J}}(\nu, p, \gamma) &\geq \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (g(X_t^\nu, p, \gamma) - g(X_t^{\hat{\nu}}, p, \gamma)) dt - q(\hat{\nu}_T - \nu_T) \right] \\ &\geq \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T g_x(X_t^\nu, p, \gamma)(X_t^{\hat{\nu}} - X_t^\nu) dt - q(\hat{\nu}_T - \nu_T) \right] \geq 0, \end{aligned}$$

which implies that $\hat{\nu}$ is optimal within \mathcal{S}_{2q} . If instead we assume (3.3.12), the claim follows from sub-linearity of the inferior limit. \square

Proof of Lemma 3.5.2. We verify the identity by explicit calculations. First notice that

$$\begin{aligned} \gamma(m) &= q\delta - \left(\frac{2\delta + \sigma^2}{2\delta} \right)^{1-\alpha} \frac{2\delta\alpha}{2\delta + \sigma^2(1-\alpha)} m^{\alpha+\beta-1} \\ &= q\delta - \alpha \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\nu^{a(m)}})^\alpha m^{\beta-1} \right]. \end{aligned} \quad (3.7.7)$$

On the other hand, we have

$$\begin{aligned} f'(m) &= -q\delta + (\alpha + \beta) \frac{2\delta + \sigma^2}{2\delta + \sigma^2(1-\alpha)} \left(\frac{2\delta}{2\delta + \sigma^2} \right)^\alpha m^{\alpha+\beta-1} \\ &= -q\delta + (\alpha + \beta) \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\nu^{a(m)}})^\alpha m^{\beta-1} \right]. \end{aligned}$$

By summing the two terms, we find

$$f'(m) + \gamma(m) = \beta \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (X_t^{\nu^{a(m)}})^\alpha m^{\beta-1} \right] = \underline{\lim}_{T \uparrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \pi_m(X_t^{\nu^{a(m)}}, m) dt \right].$$

\square

Integrability results

Lemma 3.7.1. *Let $m > 0$ and let $\mathbf{m}'_{\infty, m}(x)$ be given by (3.4.5). For any $k \geq 0$, the $\int_0^\infty x^k \mathbf{m}'_{\infty, m}(x) dx$ is finite if and only if $2\delta - (k-1)\sigma^2 > 0$. If so, it holds*

$$\int_0^\infty x^k \mathbf{m}'_{\infty, m}(x) dx = m^{k - \frac{2\delta}{\sigma^2} - 1} \left(\frac{2\delta}{\sigma^2} \right)^{k-1 - \frac{2\delta}{\sigma^2}} \Gamma \left(\frac{2\delta}{\sigma^2} - k + 1 \right). \quad (3.7.8)$$

Proof. Let $k \geq 0$ so that $2\delta - (k-1)\sigma^2 > 0$. By setting $z = x/m$ in the integral in (3.7.8), we have:

$$\int_{\mathbb{R}_+} x^k x^{-\frac{2\delta}{\sigma^2} - 2} \exp \left(-\frac{2\delta}{\sigma^2} \frac{m}{x} \right) dx = m^{k - \frac{2\delta}{\sigma^2} - 1} \int_0^\infty z^{k - \frac{2\delta}{\sigma^2} - 2} \exp \left(-\frac{2\delta}{\sigma^2} \frac{1}{z} \right) dz,$$

and, by making the change of variables $t = 2\delta/\sigma^2 \cdot 1/z$, we have

$$\begin{aligned} \int_0^\infty z^{k-\frac{2\delta}{\sigma^2}-2} \exp\left(-\frac{2\delta}{\sigma^2} \frac{1}{z}\right) dz &= \left(\frac{2\delta}{\sigma^2}\right)^{k-1-\frac{2\delta}{\sigma^2}} \int_0^\infty t^{\frac{2\delta}{\sigma^2}-k} e^{-t} dt \\ &= \left(\frac{2\delta}{\sigma^2}\right)^{k-1-\frac{2\delta}{\sigma^2}} \Gamma\left(\frac{2\delta}{\sigma^2} - k + 1\right), \end{aligned}$$

which yields (3.7.8). \square

Lemma 3.7.2. *For any $m > 0$, let $X^m = (X_t^m)_{t \geq 0}$ be the solution of*

$$dX_t^m = \delta(m - X_t^m)dt + \sigma X_t^m dW_t, \quad X_0^m = \xi.$$

There exists a constant C independent of m so that it holds

$$\sup_{t \geq 0} \mathbb{E}[(X_t^m)^2] \leq C(1 + m^2).$$

Proof. By Itô formula, we have

$$d(X_t^m)^2 = [-(2\delta - \sigma^2)(X_t^m)^2 + 2\delta m(X_t^m)] dt + 2\sigma(X_t^m)^2 dW_t, \quad (X_0^m)^2 = \xi^2,$$

so that, by taking the expectation, we get

$$\mathbb{E}[(X_t^m)^2] = e^{-(2\delta - \sigma^2)t} \left(\mathbb{E}[\xi^2] + 2\delta m \int_0^t \mathbb{E}[X_s^m] e^{-(2\delta - \sigma^2)s} ds \right).$$

By (3.4.8), we have $\mathbb{E}[X_s^m] \leq C(1 + m)$ for any $s \geq 0$ and, by Assumption **D**, $2\delta - \sigma^2 > 0$. Thus, it holds

$$\begin{aligned} \mathbb{E}[(X_t^m)^2] &\leq e^{-(2\delta - \sigma^2)t} \left(\mathbb{E}[\xi^2] + C(1 + m)2\delta m \int_0^t e^{(2\delta - \sigma^2)s} ds \right) \\ &\leq C(1 + m^2) + e^{-(2\delta - \sigma^2)t} (\mathbb{E}[\xi^2] - C_1(1 + m^2)) \leq C(1 + m^2). \end{aligned}$$

This completes the proof. \square

Auxiliary results for the backward convergence problem

Lemma 3.7.3. *Suppose $\mathbb{E}[\bar{\mu}_\infty^2] < \infty$. Let $(X^i)_{i=1}^N$ be i.i.d. as Y conditionally to $\bar{\mu}_\infty$, with either $Y = X^{\lambda^r}$ or $Y = X^{\lambda^s}$, and let $\kappa(dx, m)$ be equal to $p_\infty^r(dx, m)$ given by (3.4.6) or equal to $p_\infty^s(dx, m)$ given by (3.4.18). Let $\bar{\mu}_t^{-i, N} = \frac{1}{N-1} \sum_{j \neq i} X_t^j$, for $t \geq 0$, $1 \leq i \leq N$. Then, it holds*

$$\begin{aligned} \lim_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E} \left[\left| (\bar{\mu}_t^{-i, N})^\beta - \bar{\mu}_\infty^\beta \right|^2 \middle| \bar{\mu}_\infty \right] dt \\ = \int_{\mathbb{R}_+^{N-1}} \left| \left(\frac{1}{N-1} \sum_{j \neq i} x_j \right)^\beta - \bar{\mu}_\infty^\beta \right|^2 \bigotimes_{j \neq i} \kappa(dx_j, \bar{\mu}_\infty), \quad \mathbb{P}\text{-a.s.} \quad (3.7.9) \end{aligned}$$

Proof. Suppose without loss of generality that $(\Omega, \mathcal{F}, \mathbb{P})$ is a Polish space (if it is not, we work on the canonical space). Consider the regular conditional probability of \mathbb{P} given $\bar{\mu}_\infty$. Denote the regular conditional probability of \mathbb{P} given $\bar{\mu}_\infty = m$ by $\mathbb{P}^m(\cdot) = \mathbb{P}(\cdot | \bar{\mu}_\infty = m)$, and by $\mathbb{E}^m[\cdot]$ the expectation with respect to the probability measure \mathbb{P}^m . Conditionally to $\bar{\mu}_\infty = m$, we have that $(X^j)_{j \neq i}$ are i.i.d. as Y ; thus, in particular, the process $\bar{\mu}^{-i, N}$ is a positively recurrent regular diffusion with ergodic measure $\bigotimes_{j \neq i} \kappa(dx_j, m)$ (see, e.g. [80, Lemmata 23.17-19]). By the ergodic ratio theorem, it holds

$$\begin{aligned} \lim_{T \uparrow \infty} \frac{1}{T} \int_0^T \left| (\bar{\mu}_t^{-i, N})^\beta - m^\beta \right|^2 dt \\ = \int_{\mathbb{R}_+^{N-1}} \left| \left(\frac{1}{N-1} \sum_{j \neq i} x_j \right)^\beta - m^\beta \right|^2 \bigotimes_{j \neq i} \kappa(dx_j, m), \quad \mathbb{P}^m\text{-a.s.} \end{aligned} \quad (3.7.10)$$

Therefore, convergence in probability with respect to the probability measure \mathbb{P}^m holds as well. In order to get convergence in L^1 as well, we show that the family of random variables in the left hand-side of (3.7.10) is uniformly integrable. By, e.g., [80, Lemma 4.12], this implies that we can take the expectation with respect to \mathbb{P}^m and exchange the limit and expectation, to get

$$\begin{aligned} \overline{\lim}_{T \uparrow \infty} \frac{1}{T} \int_0^T \mathbb{E}^m \left[\left| (\bar{\mu}_t^{-i, N})^\beta - \bar{\mu}_\infty^\beta \right|^2 \right] dt \\ = \int_{\mathbb{R}_+^{N-1}} \left| \left(\frac{1}{N-1} \sum_{j \neq i} x_j \right)^\beta - m^\beta \right|^2 \bigotimes_{j \neq i} \kappa(dx_j, m) \quad \text{for } \rho\text{-a.e. } m > 0, \end{aligned}$$

which is equivalent to (3.7.9). To verify uniform integrability, take $r = 1/\beta > 1$. By standard estimates, we have

$$\begin{aligned} \mathbb{E}^m \left[\left| \frac{1}{T} \int_0^T \left| (\bar{\mu}_t^{-i, N})^\beta - m^\beta \right|^2 dt \right|^r \right] &\leq \frac{2^{2r-1}}{T} \int_0^T \left(m^{2r\beta} + \mathbb{E}^m \left[(\bar{\mu}_t^{-i, N})^{2r\beta} \right] \right) dt \\ &\leq C \left(m^2 + \sup_{t \geq 0} \mathbb{E}^m[|Y_t|^2] \right) \leq C(1 + m^2), \end{aligned}$$

where last inequality holds true by Lemma 3.7.2 if $Y = X^{\lambda^r}$, and by Lemma 3.3.1 if $Y = X^{\lambda^s}$. This implies that such family is bounded in L^r -norm, thus, since $r > 1$, uniformly integrable. \square

Lemma 3.7.4. *Let $\kappa(dx, m)$ be either equal to $p_\infty^r(dx, m)$ given by (3.4.6) or equal to $p_\infty^s(dx, m)$ given by (3.4.18). Then, for any $m > 0$, it holds*

$$\lim_{N \rightarrow \infty} \int_{\mathbb{R}_+^{N-1}} \left| \left(\frac{1}{N-1} \sum_{j \neq i} x_j \right)^\beta - m^\beta \right|^2 \bigotimes_{j \neq i} \kappa(dx_j, m) = 0. \quad (3.7.11)$$

Proof. Let $(Y_i)_{i \geq 1}$ be a sequence of i.i.d random variables with law $\kappa(dx, m)$, defined on some probability space (X, \mathcal{X}, P) . Up to reindexing, the integral in (3.7.11) can be

expressed in terms of the expectation of a function of $\bar{Y}^n = 1/n \sum_{i=1}^n Y_i$. Since $(Y^i)_{i \geq 1}$ are i.i.d. as $\kappa(dx, m)$ and integrable, we have

$$(\bar{Y}^n)^\beta \rightarrow m^\beta \quad P\text{-a.s.}$$

by the strong law of large number and continuity of the function $x \mapsto x^\beta$. Therefore, convergence in probability holds as well. To conclude, we show that the sequence $(|(\bar{Y}^n)^\beta - m^\beta|^2)_{n \geq 1}$ is uniformly bounded in L^r -norm, for some $r > 1$, which implies that the sequence is uniformly integrable, and thus the convergence in L^2 -norm holds by, e.g., [80, Proposition 4.12]. Let $r = 1/\beta > 1$. By standard estimates, we have

$$\mathbb{E}[(|(\bar{Y}^n)^\beta - m^\beta|^2)^r] \leq 2^{2r-1} (m^2 + \mathbb{E}[(\bar{Y}^n)^2]) \leq C (m^2 + \mathbb{E}[Y_1^2]),$$

where in the last inequality we used the identical distribution of the sequence $(Y_i)_{i \geq 1}$. The expectation of Y_1^2 is finite by Lemma 3.7.1 if $\kappa = p_\infty^r$ and by Lemma 3.3.1 if $\kappa = \tilde{p}$. This concludes the proof. \square

A technical result on the exchange of limits

Lemma 3.7.5. *Let $(a_{h,l})_{h,l \geq 1}$ be a real valued sequence. Suppose that the following holds:*

1. $\lim_{h \rightarrow \infty} a_{h,l} = b_l$ uniformly in m , and
2. $\underline{\lim}_{l \rightarrow \infty} a_{h,l} = c_h$ for every $h \geq 1$.

Then, it holds

$$\underline{\lim}_{l \rightarrow \infty} \lim_{h \rightarrow \infty} a_{h,l} \leq \lim_{h \rightarrow \infty} \underline{\lim}_{l \rightarrow \infty} a_{h,l}. \quad (3.7.12)$$

Proof. For any $h \geq 1$, consider a subsequence $(a_{h,l_k})_{h,k \geq 1}$ so that $c_h = \underline{\lim}_{l \rightarrow \infty} a_{h,l} = \lim_{k \rightarrow \infty} a_{h,l_k}$. Since 1. is satisfied by the subsequence $(a_{h,l_k})_{h,k \geq 1}$ as well, Moore-Osgood theorem implies that there exists $A \in \mathbb{R}$ so that

$$\lim_{h \rightarrow \infty} \lim_{k \rightarrow \infty} a_{h,l_k} = A = \lim_{k \rightarrow \infty} \lim_{h \rightarrow \infty} a_{h,l_k}. \quad (3.7.13)$$

In particular, (3.7.13) implies that $(b_{l_k})_{k \geq 1}$ is a convergent subsequence of $(b_l)_{l \geq 1}$. Therefore, we have

$$\underline{\lim}_{l \rightarrow \infty} \lim_{h \rightarrow \infty} a_{h,l} = \underline{\lim}_{l \rightarrow \infty} b_l \leq \lim_{k \rightarrow \infty} b_{l_k} = \lim_{k \rightarrow \infty} \lim_{h \rightarrow \infty} a_{h,l_k} = \lim_{h \rightarrow \infty} \lim_{k \rightarrow \infty} a_{h,l_k} = \lim_{h \rightarrow \infty} \underline{\lim}_{l \rightarrow \infty} a_{h,l},$$

where last equality holds by definition of $(a_{h,l_k})_{h,l \geq 1}$. This concludes the proof. \square

Appendix A

Relaxed controls

Here, we recall some facts on relaxed controls. Denote by \mathcal{V} the set of positive measures q on $[0, T] \times A$ so that the time marginal is equal to the Lebesgue measure, i.e., $q([s, t] \times A) = t - s$ for every $0 \leq s \leq t \leq T$. We endow \mathcal{V} with the topology of weak convergence of measures, which makes \mathcal{V} a Polish space. It is a well known fact that, when the set A is compact, \mathcal{V} is compact as well. Moreover, for every measure $q \in \mathcal{V}$, there exists a measurable map $[0, T] \ni t \mapsto q_t \in \mathcal{P}(A)$ so that $q(da, dt) = q_t(da)dt$, with $(q_t)_{t \in [0, T]}$ unique up to $Leb_{[0, T]}$ -a.e. equality. We can equip the measurable space $(\mathcal{V}, \mathcal{B}_{\mathcal{V}})$ with the filtration $(\mathcal{F}_t^{\mathcal{V}})_{t \in [0, T]}$ defined by

$$\mathcal{F}_t^{\mathcal{V}} = \sigma(\mathcal{V} \ni q \mapsto q(C) : C \in \mathcal{B}_{[0, t] \times A}).$$

We observe that $\mathcal{F}_t^{\mathcal{V}}$ is countably generated for every $t \in [0, T]$, by reasoning as in the proof of [17, Proposition 7.25]. Finally, one can prove that there exists an $\mathcal{F}_t^{\mathcal{V}}$ -predictable process $\bar{q} : [0, T] \times \mathcal{V} \rightarrow \mathcal{P}(A)$ such that, for each $q \in \mathcal{V}$, $\bar{q}(t, q) = q_t$ for a.e. $t \in [0, T]$ (see, e.g., [84, Lemma 3.2]). By an abuse of notation, we write $q_t(da) = \bar{q}(t, q)(da)$.

Consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$. A relaxed control \mathbf{r} is a \mathcal{V} -valued random variable. We say that \mathbf{r} is \mathbb{F} -adapted if $\mathbf{r}(C)$ is a real valued \mathcal{F}_t -measurable random variable for every $C \in \mathcal{B}_{[0, t] \times A}$. Observe that every A -valued progressively measurable process $\alpha = (\alpha_t)_{t \in [0, T]}$, which is often referred to as strict control, induces a relaxed control by setting

$$\mathbf{r}_t(da)dt = \delta_{\alpha_t}(da)dt.$$

Finally, using the map \bar{q} described above, we can safely identify every \mathbb{F} -adapted relaxed control \mathbf{r} with the unique (up to $Leb_{[0, T]}$ -a.e. equality) \mathbb{F} -progressively measurable process $(\mathbf{r}_t)_{t \in [0, T]}$ with values in $\mathcal{P}(A)$ so that

$$\mathbb{P}(\mathbf{r}(da, dt) = \mathbf{r}_t(da)dt) = 1.$$

In the following, we will use mostly the notation $(\mathbf{r}_t)_{t \in [0, T]}$ for a relaxed control and will make no distinction between a \mathcal{V} -valued random variable and a $\mathcal{P}(A)$ -valued process.

Appendix B

Weak and strong existence for controlled equations

We state and prove a Yamada-Watanabe type result for stochastic differential equations with random coefficients as the ones encountered in Chapter 1. Recall from Appendix A the definition of the space \mathcal{V} and of relaxed controls.

Let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ be a tuple composed by a filtered probability space satisfying usual assumptions, an \mathbb{F} -Brownian motion W , an \mathbb{R}^d -valued \mathcal{F}_0 -measurable random variable ξ , an \mathcal{F}_0 -measurable random flow of measures μ taking values in $\mathcal{C}(\mathcal{P}^2)$ and an \mathbb{F} -adapted \mathcal{V} -valued random variable \mathfrak{r} , in the sense that the random variables $\mathfrak{r}(C)$ are \mathcal{F}_t -measurable for every $C \in \mathcal{B}_{[0,t] \times A}$. Let us consider the following stochastic differential equation:

$$dX_t = G(t, X_t, \mu, \mathfrak{r})dt + dW_t, \quad X_0 = \xi. \quad (\text{B.1})$$

where $G : [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2) \times \mathcal{V} \rightarrow \mathbb{R}^d$ is jointly measurable and progressively measurable in \mathcal{V} ; progressive measurability must be understood in the following sense: for every $q, q' \in \mathcal{V}$, for every $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{C}(\mathcal{P}^2)$, it holds:

$$q(C) = q'(C) \quad \forall C \in \mathcal{B}_{[0,t] \times A} \implies G(t, x, m, q) = G(t, x, m, q').$$

Definition B.1 (Strong solution and pathwise uniqueness). Let $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ be a tuple as above. A strong solution to equation (B.1) on \mathfrak{U} is a continuous \mathbb{F} -adapted process $X = (X_t)_{t \in [0, T]}$ adapted to the \mathbb{P} -augmentation of \mathbb{F} so that

$$X_t = \xi + \int_0^t G(s, X_s, \mu, \mathfrak{r})ds + W_t, \quad 0 \leq t \leq T, \quad (\text{B.2})$$

holds \mathbb{P} -almost surely.

Pathwise uniqueness holds for equation (B.1) if, given two strong solutions X and X' to (B.1) on \mathfrak{U} , they are indistinguishable:

$$\mathbb{P}(X_t = X'_t \quad \forall t \in [0, T]) = 1.$$

Definition B.2 (Weak solution and uniqueness in law). A weak solution to equation (B.1) is a tuple $\mathfrak{U} = ((\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}), \xi, W, \mu, \mathfrak{r})$ as above so that there exists a continuous \mathbb{F} -adapted process $X = (X_t)_{t \in [0, T]}$ satisfying equation (B.1).

Weak uniqueness holds for equation (B.1) if for any two weak solution of (B.1) \mathfrak{U}^i , $i = 1, 2$, so that $\mathbb{P}^1 \circ (\xi^1, W^1, \mu^1, \mathfrak{r}^1)^{-1} = \mathbb{P}^2 \circ (\xi^2, W^2, \mu^2, \mathfrak{r}^2)^{-1}$, it holds

$$\mathbb{P}^1 \circ (X^1, \xi^1, W^1, \mu^1, \mathfrak{r}^1)^{-1} = \mathbb{P}^2 \circ (X^2, \xi^2, W^2, \mu^2, \mathfrak{r}^2)^{-1},$$

where X^i are the continuous \mathbb{F}^i -adapted processes that satisfy equation (B.1) on \mathfrak{U}^i , $i = 1, 2$.

Theorem B.1. *Suppose pathwise uniqueness holds for equation (B.1), in the sense of Definition B.1. Then, uniqueness in law in the sense of Definition B.2 holds as well.*

Proof. Let \mathfrak{U}^1 and \mathfrak{U}^2 be two weak solutions of equation (B.1) in the sense of Definition B.2 above. Since pathwise uniqueness holds for equation (B.1) by assumption, our goal is to bring together the solution on the same filtered probability space. Let us define the following probability measures:

$$\begin{aligned}\hat{\mathbb{Q}}^i &= \mathbb{P}^i \circ (\xi^i, W^i, \mu^i, \mathfrak{r}^i, X^i)^{-1} \in \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}(\mathcal{P}^2) \times \mathcal{V} \times \mathcal{C}^d), \quad i = 1, 2, \\ \mathbb{Q} &= \mathbb{P}^i \circ (\xi^i, W^i, \mu^i, \mathfrak{r}^i)^{-1} \in \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}(\mathcal{P}^2) \times \mathcal{V}), \\ \tilde{\mathbb{Q}} &= \mathbb{P}^i \circ (\xi^i, W^i, \mu^i)^{-1} \in \mathcal{P}(\mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}(\mathcal{P}^2)).\end{aligned}$$

Observe that \mathbb{Q} and $\tilde{\mathbb{Q}}$ are well defined, since $(\xi^i, W^i, \mu^i, \mathfrak{r}^i)$ share the same joint law by assumption. Let us consider the following space:

$$\begin{aligned}\Omega^{can} &= \mathcal{C}^d \times \mathcal{C}^d \times \mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}(\mathcal{P}^2) \times \mathcal{V}; \\ \mathcal{F}^{can} &= \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathbb{R}^d} \otimes \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)} \otimes \mathcal{B}_{\mathcal{V}}; \\ \mathcal{G}_t^{can} &= \mathcal{B}_{t, \mathcal{C}^d} \otimes \mathcal{B}_{t, \mathcal{C}^d} \otimes \mathcal{B}_{\mathbb{R}^d} \otimes \mathcal{B}_{t, \mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)} \otimes \mathcal{F}_t^{\mathcal{V}},\end{aligned}$$

where

$$\mathcal{B}_{t, \mathcal{C}^d} = \sigma(\mathcal{C}^d \ni x \mapsto x_s \in \mathbb{R}^d : s \leq t), \quad \mathcal{F}_t^{\mathcal{V}} = \sigma(\mathcal{V} \ni q \mapsto q(C) \in \mathbb{R} : C \in \mathcal{B}_{[0, t] \times A}).$$

In order to equip the space $(\Omega^{can}, \mathcal{F}^{can}, (\mathcal{G}_t^{can})_{t \in [0, T]})$ with a probability measure, we disintegrate the measures $\hat{\mathbb{Q}}^i$, $i = 1, 2$, in the following way: let $K^i : \mathcal{B}_{\mathcal{C}^d} \times \mathbb{R}^d \times \mathcal{C}^d \times \mathcal{C}(\mathcal{P}^2) \times \mathcal{V} \rightarrow [0, 1]$ be a regular conditional probability of $\hat{\mathbb{Q}}^i$ for $\mathcal{B}_{\mathcal{C}^d}$ given (x, w, m, q) , so that it holds

$$\hat{\mathbb{Q}}^i(A \times B) = \int_B K^i(A, x, m, w, q) \mathbb{Q}(dx, dm, dw, dq),$$

for every $A \in \mathcal{B}_{\mathcal{C}^d}$, $B \in \mathcal{B}_{\mathbb{R}^d} \otimes \mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)} \otimes \mathcal{B}_{\mathcal{V}}$, or more briefly

$$\hat{\mathbb{Q}}^i(dx, dw, dm, dq, dy) = K^i(dy, x, m, q, w) \mathbb{Q}(dx, dw, dm, dq), \quad i = 1, 2.$$

Then, we set

$$\overline{\mathbb{Q}}(dy^1, dy^2, dx, dm, dw, dq) = K^1(dy^1, x, m, q, w) K^2(dy^2, x, m, q, w) \mathbb{Q}(dx, dm, dw, dq).$$

Observe that the joint law under $\overline{\mathbb{Q}}$ of the coordinate projections y^1 , x , m , w and q is exactly $\hat{\mathbb{Q}}^1$, and analogously when considering the coordinate process y^2 instead of

y^1 . Finally, complete the σ -algebra \mathcal{F}^{can} with the $\overline{\mathbb{Q}}$ -null sets $\mathcal{N}^{\overline{\mathbb{Q}}}$ and consider the complete right continuous filtration $(\mathcal{F}_t^{can})_{t \in [0, T]}$ given by

$$\mathcal{F}_t^{can} = \bigcap_{\varepsilon > 0} \sigma(\mathcal{G}_{t+\varepsilon}, \mathcal{N}^{\overline{\mathbb{Q}}}).$$

By Lemma B.2, the coordinate process w is a $(\mathcal{F}_t^{can})_{t \in [0, T]}$ -Brownian motion under $\overline{\mathbb{Q}}$. Furthermore, it holds

$$y_t^i = x + \int_0^t G(s, y_s^i, m, q) ds + w_t, \quad \forall t \in [0, T], \quad \overline{\mathbb{Q}}\text{-a.s.}$$

for $i = 1, 2$. Since pathwise uniqueness in the sense of Definition B.1 holds by assumption, it follows that y^1 and y^2 are indistinguishable under $\overline{\mathbb{Q}}$, which implies $\hat{\mathbb{Q}}^1 = \hat{\mathbb{Q}}^2$. This proves the desired result. \square

Lemma B.2. *In the construction of Theorem B.1, $w = (w_s)_{s \in [0, T]}$ is a Brownian motion under $\overline{\mathbb{Q}}$ with respect to the filtration $(\mathcal{F}_s^{can})_{s \in [0, T]}$.*

Proof. Observe that w is a natural Brownian motion under $\overline{\mathbb{Q}}$. In order to show that it is a Brownian motion with respect to the filtration $(\mathcal{G}_t^{can})_{t \in [0, T]}$, we just need to prove that its increments are independent, and the conclusion follows.

Fix $A_1, A_2 \in \mathcal{B}_{t, \mathcal{C}^d}$, $B \in \mathcal{B}_{\mathbb{R}^d}$, $C \in \mathcal{B}_{t, \mathcal{C}^d}$, $D \in \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}$ and $F \in \mathcal{F}_t^\nu$. By Cauchy-Schwartz inequality, we have, for every $H \in \mathcal{B}_{\mathbb{R}^d}$ and $s > t$:

$$\begin{aligned} & \mathbb{E}^{\overline{\mathbb{Q}}} \left[\mathbb{1}_H(w_s - w_t) \mathbb{1}_{A_1 \times A_2 \times B \times D \times C \times F}(y^1, y^2, x, m, w, q) \right]^2 \\ & \leq \mathbb{E}^{\overline{\mathbb{Q}}} \left[\mathbb{1}_H(w_s - w_t) \mathbb{1}_{A_1 \times B \times D \times C \times F}(y^1, x, m, w, q) \right] \\ & \quad \cdot \mathbb{E}^{\overline{\mathbb{Q}}} \left[\mathbb{1}_H(w_s - w_t) \mathbb{1}_{A_2 \times B \times D \times C \times F}(y^2, x, m, w, q) \right]. \end{aligned}$$

Therefore, it suffices to show that

$$\mathbb{E}^{\overline{\mathbb{Q}}} \left[\mathbb{1}_H(w_s - w_t) \mathbb{1}_{A_1 \times B \times D \times C \times F}(y^1, x, m, w, q) \right] = 0. \quad (\text{B.3})$$

Since the integrand does not depend upon y^2 , we may rewrite such an expectation only with respect to $\hat{\mathbb{Q}}^1$:

$$\begin{aligned} & \mathbb{E}^{\overline{\mathbb{Q}}} \left[\mathbb{1}_H(w_s - w_t) \mathbb{1}_{A_1 \times B \times D \times C \times F}(y^1, x, m, w, q) \right] \\ & = \int \mathbb{1}_H(w_s - w_t) \mathbb{1}_{A_1 \times B \times D \times C \times F}(y^1, x, m, w, q) \hat{\mathbb{Q}}^1(dy^1, dx, dm, dw, dq). \end{aligned}$$

Then, we introduce another disintegration of the measure $\hat{\mathbb{Q}}^1$: let Θ^1 be a regular conditional probability for $\mathcal{B}_{\mathcal{C}^d} \otimes \mathcal{B}_{\mathcal{V}}$ given (x, w, m) :

$$\hat{\mathbb{Q}}^1(A \times B \times C \times D \times F) = \int_{B \times C \times D} \Theta^1(A \times F, x, m, w) \tilde{\mathbb{Q}}(dx, dm, dw), \quad (\text{B.4})$$

for every $A \in \mathcal{B}_{\mathcal{C}^d}$, $B \in \mathcal{B}_{\mathbb{R}^d}$, $C \in \mathcal{B}_{\mathcal{C}^d}$, $D \in \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)}$ and $F \in \mathcal{B}_{\mathcal{V}}$, or more briefly

$$\hat{\mathbb{Q}}^1(dy^1, dq, dx, dw, dm) = \Theta^1(dy^1, dq, x, m, w) \tilde{\mathbb{Q}}(dx, dw, dm).$$

As in [77, Lemma IV.1.1], it can easily be shown that, for every $A \times F \in \mathcal{B}_{s, \mathcal{C}^d} \otimes \mathcal{F}_s^\mathcal{V}$, the map

$$(x, m, w) \mapsto \Theta^1(A \times F, x, m, w)$$

is $\mathcal{B}_{\mathbb{R}^d} \otimes \mathcal{B}_{\mathcal{C}(\mathcal{P}^2)} \otimes \mathcal{B}_{s, \mathcal{C}^d}$ -measurable, for every $s \in [0, T]$. Therefore, we can compute the left-hand side of (B.3):

$$\begin{aligned} & \mathbb{E}^{\bar{\mathbb{Q}}} \left[\mathbf{1}_H(w_s - w_t) \mathbf{1}_{A_1 \times B \times D \times C \times F}(y^1, x, m, w, q) \right] \\ &= \int \mathbf{1}_H(w_s - w_t) \mathbf{1}_{A_1 \times B \times D \times C \times F}(y^1, x, m, w, q) \hat{\mathbb{Q}}^1(dy^1, dx, dm, dw, dq) \\ &= \int \mathbf{1}_H(w_s - w_t) \Theta^1(A_1 \times F, x, m, w) \mathbf{1}_{B \times D \times C}(x, m, w) \tilde{\mathbb{Q}}^1(dx, dm, dw) \\ &= \mathbb{E}^{\mathbb{P}^1} \left[\mathbf{1}_H(W_s^1 - W_t^1) \Theta^1(A_1 \times F, \xi^1, \mu^1, W^1) \mathbf{1}_{B \times D \times C}(\xi^1, \mu^1, W^1) \right] = 0, \end{aligned}$$

since $\Theta^1(A_1 \times F, \xi^1, \mu^1, W^1) \mathbf{1}_{B \times D \times C}(\xi^1, \mu^1, W^1)$ is \mathcal{F}_s^1 -measurable and W^1 is an \mathbb{F}^1 -Brownian motion under \mathbb{P}^1 by assumption. \square

Bibliography

- [1] S. Adlakha and R. Johari. Mean field equilibrium in dynamic games with strategic complementarities. *Oper. Res.*, 61(4):971–989, 2013.
- [2] F. Alvarez, F. Lippi, and P. Souganidis. Price setting with strategic complementarities as a mean field game. *Econometrica*, 91(6):2005–2039, 2023.
- [3] F. E. Alvarez, D. Argente, F. Lippi, E. Méndez, and D. Van Patten. Strategic complementarities in a dynamic model of technology adoption: P2p digital payments. Technical report, National Bureau of Economic Research, 2023.
- [4] L. H. R. Alvarez. Reward functionals, salvage values, and optimal stopping. *Math. Methods Oper. Res.*, 54(2):315–337, 2001.
- [5] L. H. R. Alvarez. A class of solvable stationary singular stochastic control problems. *arXiv preprint arXiv:1803.03464*, 2018.
- [6] R. J. Aumann. Subjectivity and correlation in randomized strategies. *J. Math. Econom.*, 1(1):67–96, 1974.
- [7] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987.
- [8] A. Aurell, R. Carmona, G. Dayanikli, and M. Laurière. Optimal incentives to mitigate epidemics: a Stackelberg mean field game approach. *SIAM J. Control Optim.*, 60(2):S294–S322, 2022.
- [9] P. Bank. Optimal control under a dynamic fuel constraint. *SIAM J. Control Optim.*, 44(4):1529–1541, 2005.
- [10] P. Bank and F. Riedel. Optimal consumption choice with intertemporal substitution. *Ann. Appl. Probab.*, 11(3):750–788, 2001.
- [11] M. Bardi and M. Fischer. On non-uniqueness and uniqueness of solutions in finite-horizon mean field games. *ESAIM Control Optim. Calc. Var.*, 25:Paper No. 44, 33, 2019.
- [12] S. Barrett. Self-enforcing international environmental agreements. *Oxford economic papers*, 46(Supplement_1):878–894, 1994.
- [13] M. Basei, H. Cao, and X. Guo. Nonzero-sum stochastic games and mean-field games with impulse controls. *Math. Oper. Res.*, 47(1):341–366, 2022.
- [14] E. Ben-Porath. Correlation without mediation: expanding the set of equilibrium outcomes by “cheap” pre-play procedures. *J. Econom. Theory*, 80(1):108–122, 1998.

- [15] V. E. Beneš, L. A. Shepp, and H. S. Witsenhausen. Some solvable stochastic control problems. In *Analysis and optimisation of stochastic systems (Proc. Internat. Conf., Univ. Oxford, Oxford, 1978)*, Inst. Math. Appl. Conf. Ser., pages 3–10. Academic Press, London-New York, 1980.
- [16] P. Bergault, P. Cardaliaguet, and C. Rainer. Mean Field Games in a Stackelberg Problem with an Informed Major Player. *SIAM J. Control Optim.*, 62(3):1737–1765, 2024.
- [17] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control*, volume 139 of *Mathematics in Science and Engineering*. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1978. The discrete time case.
- [18] C. Bertucci, J.-M. Lasry, and P.-L. Lions. Some remarks on mean field games. *Comm. Partial Differential Equations*, 44(3):205–227, 2019.
- [19] O. Bonesini. *Four essays in between Probability Theory and Financial Mathematics*. PhD thesis, Università degli Studi di Padova, Padua, Italy, 2023.
- [20] O. Bonesini, L. Campi, and M. Fischer. Correlated equilibria for mean field games with progressive strategies. *Mathematics of Operations Research*, 2024.
- [21] A. N. Borodin and P. Salminen. *Handbook of Brownian motion—facts and formulae*. Probability and its Applications. Birkhäuser Verlag, Basel, second edition, 2002.
- [22] G. Bouveret, R. Dumitrescu, and P. Tankov. Mean-field games of optimal stopping: a relaxed solution approach. *SIAM J. Control Optim.*, 58(4):1795–1821, 2020.
- [23] P. Brémaud and M. Yor. Changes of filtrations and of probability measures. *Z. Wahrsch. Verw. Gebiete*, 45(4):269–295, 1978.
- [24] G. Brunick and S. Shreve. Mimicking an Itô process by a solution of a stochastic differential equation. *Ann. Appl. Probab.*, 23(4):1584–1628, 2013.
- [25] L. Campi, F. Cannerozzi, and F. Cartellier. Coarse Correlated Equilibria in Linear Quadratic Mean Field Games and Application to an Emission Abatement Game. *Appl. Math. Optim.*, 91(1):Paper No. 8, 2025.
- [26] L. Campi, F. Cannerozzi, and M. Fischer. Coarse correlated equilibria for continuous time mean field games in open loop strategies. *Electron. J. Probab.*, 29:Paper No. 1, 2024.
- [27] L. Campi, T. De Angelis, M. Ghio, and G. Livieri. Mean-field games of finite-fuel capacity expansion with singular controls. *Ann. Appl. Probab.*, 32(5):3674–3717, 2022.
- [28] L. Campi and M. Fischer. N -player games and mean-field games with absorption. *Ann. Appl. Probab.*, 28(4):2188–2242, 2018.
- [29] L. Campi and M. Fischer. Correlated equilibria and mean field games: a simple model. *Math. Oper. Res.*, 47(3):2240–2259, 2022.
- [30] F. Cannerozzi and G. Ferrari. Cooperation, correlation and competition in ergodic N -player games and mean-field games of singular controls: A case study. *ArXiv preprint arXiv:2404.15079*, 2024.

- [31] H. Cao, J. Dianetti, and G. Ferrari. Stationary discounted and ergodic mean field games with singular controls. *Math. Oper. Res.*, 48(4):1871–1898, 2023.
- [32] H. Cao and X. Guo. MFGs for partially reversible investment. *Stochastic Process. Appl.*, 150:995–1014, 2022.
- [33] P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions. *The master equation and the convergence problem in mean field games*, volume 201 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2019.
- [34] P. Cardaliaguet and P. J. Graber. Mean field games systems of first order. *ESAIM Control Optim. Calc. Var.*, 21(3):690–722, 2015.
- [35] P. Cardaliaguet, P. J. Graber, A. Porretta, and D. Tonon. Second order mean field games with degenerate diffusion and local coupling. *NoDEA Nonlinear Differential Equations Appl.*, 22(5):1287–1317, 2015.
- [36] R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SIAM J. Control Optim.*, 51(4):2705–2734, 2013.
- [37] R. Carmona and F. Delarue. Forward-backward stochastic differential equations and controlled McKean-Vlasov dynamics. *Ann. Probab.*, 43(5):2647–2700, 2015.
- [38] R. Carmona and F. Delarue. *Probabilistic theory of mean field games with applications. I*, volume 83 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2018. Mean field FBSDEs, control, and games.
- [39] R. Carmona and F. Delarue. *Probabilistic theory of mean field games with applications. II*, volume 84 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2018. Mean field games with common noise and master equations.
- [40] R. Carmona, F. Delarue, and A. Lachapelle. Control of McKean-Vlasov dynamics versus mean field games. *Math. Financ. Econ.*, 7(2):131–166, 2013.
- [41] R. Carmona, F. Delarue, and D. Lacker. Mean field games with common noise. *Ann. Probab.*, 44(6):3740–3803, 2016.
- [42] R. Carmona, F. Delarue, and D. Lacker. Mean field games of timing and models for bank runs. *Appl. Math. Optim.*, 76(1):217–260, 2017.
- [43] R. Carmona, J.-P. Fouque, and L.-H. Sun. Mean field games and systemic risk. *Commun. Math. Sci.*, 13(4):911–933, 2015.
- [44] R. Carmona and D. Lacker. A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.*, 25(3):1189–1231, 2015.
- [45] S. Christensen, B. A. Neumann, and T. Sohr. Competition versus cooperation: a class of solvable mean field impulse control problems. *SIAM J. Control Optim.*, 59(5):3946–3972, 2021.
- [46] A. Cohen, A. Hening, and C. Sun. Optimal ergodic harvesting under ambiguity. *SIAM J. Control Optim.*, 60(2):1039–1063, 2022.
- [47] A. Cohen and C. Sun. Existence of optimal stationary singular controls and mean field game equilibria. *arXiv preprint arXiv:2404.07945*, 2024.

- [48] J. Cvitanić and J. Zhang. *Contract theory in continuous-time models*. Springer Finance. Springer, Heidelberg, 2013.
- [49] C. Dellacherie and P.-A. Meyer. *Probabilities and potential. B*, volume 72 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1982. Theory of martingales, Translated from the French by J. P. Wilson.
- [50] R. Denkert and U. Horst. Extended mean-field games with multi-dimensional singular controls and non-linear jump impact. *arXiv preprint arXiv:2402.09317*, 2024.
- [51] J. Dianetti. Strong solutions to submodular mean field games with common noise and related mckean-vlasov fbsdes. *arXiv preprint arXiv:2212.12413*, 2022.
- [52] J. Dianetti, S. Federico, G. Ferrari, and G. Floccari. Multiple equilibria in mean-field game models of firm competition with strategic complementarities. *Quantitative Finance*, 0(0):1–15, 2025.
- [53] J. Dianetti, G. Ferrari, M. Fischer, and M. Nendel. Submodular mean field games: existence and approximation of solutions. *Ann. Appl. Probab.*, 31(6):2538–2566, 2021.
- [54] J. Dianetti, G. Ferrari, M. Fischer, and M. Nendel. A unifying framework for submodular mean field games. *Math. Oper. Res.*, 48(3):1679–1710, 2023.
- [55] J. Dianetti, G. Ferrari, and I. Tzouanas. Ergodic mean-field games of singular control with regime-switching (extended version). *arXiv preprint arXiv:2307.12012*, 2023.
- [56] M. F. Djete. Stackelberg mean field games: convergence and existence results to the problem of principal with multiple agents in competition. *ArXiv preprint arXiv:2309.00640*, 2023.
- [57] T. Dokka, H. Moulin, I. Ray, and S. SenGupta. Equilibrium design in an n -player quadratic game. *Rev. Econ. Des.*, 27(2):419–438, 2023.
- [58] R. Dumitrescu, M. Leutscher, and P. Tankov. Control and optimal stopping mean field games: a linear programming approach. *Electron. J. Probab.*, 26:Paper No. 157, 49, 2021.
- [59] R. Dumitrescu, M. Leutscher, and P. Tankov. Linear programming fictitious play algorithm for mean field games with optimal stopping and absorption. *ESAIM Math. Model. Numer. Anal.*, 57(2):953–990, 2023.
- [60] N. El Karoui, D. H Nguyen, and M. Jeanblanc-Picqué. Compactification methods in the control of degenerate diffusions: existence of an optimal control. *Stochastics*, 20(3):169–219, 1987.
- [61] R. Elie, T. Mastrolia, and D. Possamaï. A tale of a principal and many, many agents. *Math. Oper. Res.*, 44(2):440–467, 2019.
- [62] K. Fan. Minimax theorems. *Proc. Nat. Acad. Sci. U.S.A.*, 39:42–47, 1953.
- [63] S. Federico, G. Ferrari, F. Riedel, and M. Röckner. On a class of infinite-dimensional singular stochastic control problems. *SIAM J. Control Optim.*, 59(2):1680–1704, 2021.
- [64] G. Ferrari and P. Salminen. Irreversible investment under Lévy uncertainty: an equation for the optimal boundary. *Adv. in Appl. Probab.*, 48(1):298–314, 2016.

- [65] M. Fischer. On the connection between symmetric N -player games and mean field games. *Ann. Appl. Probab.*, 27(2):757–810, 2017.
- [66] G. Fu. Extended mean field games with singular controls. *SIAM J. Control Optim.*, 61(1):283–312, 2023.
- [67] G. Fu and U. Horst. Mean field games with singular controls. *SIAM J. Control Optim.*, 55(6):3833–3868, 2017.
- [68] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [69] P. J. Graber. Linear quadratic mean field type control and mean field games with common noise, with application to production of an exhaustible resource. *Appl. Math. Optim.*, 74(3):459–486, 2016.
- [70] C. Grüning and W. Peters. Can justice and fairness enlarge international environmental agreements? *Games*, 1(2):137–158, 2010.
- [71] X. Guo and R. Xu. Stochastic games for fuel follower problem: N versus mean field game. *SIAM J. Control Optim.*, 57(1):659–692, 2019.
- [72] J. Hannan. Approximation to Bayes risk in repeated play. In *Contributions to the theory of games, vol. 3*, Annals of Mathematics Studies, no. 39, pages 97–139. Princeton University Press, Princeton, N.J., 1957.
- [73] S. Hart and A. Mas-Colell. Regret-based continuous-time dynamics. *Games Econom. Behav.*, 45(2):375–394, 2003. Special issue in honor of Robert W. Rosenthal.
- [74] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Math. Oper. Res.*, 14(1):18–25, 1989.
- [75] U. G. Haussmann and J.-P. Lepeltier. On the existence of optimal controls. *SIAM J. Control Optim.*, 28(4):851–902, 1990.
- [76] M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.*, 6(3):221–251, 2006.
- [77] N. Ikeda and S. Watanabe. *Stochastic differential equations and diffusion processes*, volume 24 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam-New York; Kodansha, Ltd., Tokyo, 1981.
- [78] A. Jack and M. Zervos. A singular control problem with an expected and a pathwise ergodic performance criterion. *J. Appl. Math. Stoch. Anal.*, pages Art. ID 82538, 19, 2006.
- [79] A. X. Jiang and K. Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games Econom. Behav.*, 91:347–359, 2015.
- [80] O. Kallenberg. *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002.
- [81] I. Karatzas. A class of singular stochastic control problems. *Adv. in Appl. Probab.*, 15(2):225–254, 1983.

- [82] I. Karatzas and S. E. Shreve. Connections between optimal stopping and singular stochastic control. I. Monotone follower problems. *SIAM J. Control Optim.*, 22(6):856–877, 1984.
- [83] I. Karatzas and S. E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991.
- [84] D. Lacker. Mean field games via controlled martingale problems: existence of Markovian equilibria. *Stochastic Process. Appl.*, 125(7):2856–2894, 2015.
- [85] D. Lacker. A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields*, 165(3-4):581–648, 2016.
- [86] D. Lacker. Limit theory for controlled McKean-Vlasov dynamics. *SIAM J. Control Optim.*, 55(3):1641–1672, 2017.
- [87] D. Lacker. On the convergence of closed-loop Nash equilibria to the mean field game limit. *Ann. Appl. Probab.*, 30(4):1693–1761, 2020.
- [88] D. Lacker and L. Le Flem. Closed-loop convergence for mean field games with common noise. *Ann. Appl. Probab.*, 33(4):2681–2733, 2023.
- [89] J.-M. Lasry and P.-L. Lions. Mean field games. *Jpn. J. Math.*, 2(1):229–260, 2007.
- [90] M. Laurière, S. Perrin, J. Pérolat, S. Girgin, P. Muller, R. Élie, M. Geist, and O. Pietquin. Learning in mean field games: A survey. *ArXiv preprint arXiv:2205.12944*, 2024.
- [91] H. Moulin, I. Ray, and S. S. Gupta. Coarse Correlated Equilibria in an Abatement Game. Cardiff Economics Working Papers E2014/24, Cardiff, 2014.
- [92] H. Moulin, I. Ray, and S. Sen Gupta. Improving Nash by coarse correlation. *J. Econom. Theory*, 150:852–865, 2014.
- [93] H. Moulin and J.-P. Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *Internat. J. Game Theory*, 7(3-4):201–221, 1978.
- [94] P. Muller, R. Elie, M. Rowland, M. Lauriere, J. Perolat, S. Perrin, M. Geist, G. Piliouras, O. Pietquin, and K. Tuyls. Learning Correlated Equilibria in Mean-Field Games, 2022.
- [95] P. Muller, M. Rowland, R. Elie, G. Piliouras, J. Perolat, M. Lauriere, R. Marinier, O. Pietquin, and K. Tuyls. Learning equilibria in mean-field games: Introducing mean-field psro, 2021.
- [96] A. Neyman. Correlated equilibrium and potential games. *Internat. J. Game Theory*, 26(2):223–227, 1997.
- [97] A. Nowak. Correlated relaxed equilibria in nonzero-sum linear differential games. *Journal of Mathematical Analysis and Applications*, 163(1):104–112, 1992.
- [98] A. Nowak. Correlated equilibria in nonzero-sum differential games. *Journal of Mathematical Analysis and Applications*, 174(2):539–549, 1993.

- [99] C. H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):Art. 14, 29, 2008.
- [100] P. E. Protter. *Stochastic integration and differential equations*, volume 21 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, second edition, 2005. Corrected third printing.
- [101] L. C. G. Rogers and D. Williams. *Diffusions, Markov processes, and martingales. Vol. 1*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2000. Foundations, Reprint of the second (1994) edition.
- [102] T. Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, Cambridge, 2016.
- [103] J. M. Stoyanov. *Counterexamples in probability*. Wiley Series in Probability and Statistics. John Wiley & Sons, Ltd., Chichester, second edition, 1997.
- [104] A.-S. Sznitman. Topics in propagation of chaos. In *École d'Été de Probabilités de Saint-Flour XIX—1989*, volume 1464 of *Lecture Notes in Math.*, pages 165–251. Springer, Berlin, 1991.
- [105] C. Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003.
- [106] X. Vives. Nash equilibrium with strategic complementarities. *J. Math. Econom.*, 19(3):305–321, 1990.
- [107] X. Vives. Complementarities and games: New developments. *Journal of Economic Literature*, 43(2):437–479, 2005.
- [108] X. Vives. Games with strategic complementarities: New applications to industrial organization. *International Journal of Industrial Organization*, 23(7-8):625–637, 2005.
- [109] X. Vives. Strategic complementarities in oligopoly. In *Handbook of Game Theory and Industrial Organization, Volume I*, pages 9–39. Edward Elgar Publishing, 2018.
- [110] A. P. N. Weerasinghe. Stationary stochastic control for Itô processes. *Adv. in Appl. Probab.*, 34(1):128–140, 2002.
- [111] J. Yong and X. Y. Zhou. *Stochastic controls*, volume 43 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1999. Hamiltonian systems and HJB equations.