Topics in Stochastic Games and Networks

Notes from ORF 569, First Draft

Please do not share!

RENE A. CARMONA
Paul M. Wythes '55 Professor of Engineering and Finance
Bendheim Center for Finance
ORFE & PACM
Princeton University
Princeton, N.J. 08544, USA

Contents

Preface / Disclaimer

This is the first draft of the notes for the course ORF 569. I prepared them from my handwritten notes and the transcripts of the lectures typed by Nicolas Tan.

This set contains most of what I covered in the course. If you happen to read these notes, please send me all the comments and suggestions you are willing to share. They will help me upgrade the quality of this document, potentially making it useful to others, especially if I happen to teach the course one more time.

René Carmona
Princeton
May 2, 2019.

Part I

STATIC GAMES

1

A Crash Course on Static Games

The purpose of this chapter is to offer a crash course on the theory of static games with a finite number of players, acting simultaneously and competing to achieve individual goals. In this introductory chapter, players have complete information and no element of the game is due to chance. Mathematically, this means that no stochastic element is part of the definition of the game models.

1.1 INTRODUCTION TO STATIC GAMES WITH FINITELY MANY PLAYERS

1.1.1 Notations and Definitions

We study games between N players. We denote by $[N] = \{1, \dots, N\}$ the set of the N different agents/players. Because of the ludic connotation of the word player, like most economists, we will often prefer the terminology agent. For each agent i, the set of actions that they are allowed to take is denoted by A^i. For most of the examples considered in these lectures, A^i will be a subset of $\mathbb{R}^{k_i}$ for some $k_i > 0$. We call A^i the set of admissible actions available to player i, or in short, the action space of agent i. The choices of elements of A^i are sometimes called pure strategies to distinguish them from the so-called mixed strategies which we introduce later on. The elements of the set $A = A^1 \times A^2 \times \cdots \times A^N$ of admissible actions for all of the N agents will be called **strategy profiles**. For any $\boldsymbol{\alpha} \in A$, we denote by $\boldsymbol{\alpha} = (\alpha^1, \alpha^2, \dots, \alpha^N)$ with $\alpha^i \in A^i$ such a strategy profile.

We also associate to each agent $i \in [N]$ a cost function:

$$J^i : A \ni \boldsymbol{\alpha} \mapsto J^i(\boldsymbol{\alpha}) = J^i(\alpha^i, \boldsymbol{\alpha}^{-i}) \in \mathbb{R}$$

where we use the notation $\boldsymbol{\alpha}^{-i} = (\alpha^1, \dots, \alpha^{i-1}, \alpha^{i+1}, \dots, \alpha^N)$ which is commonly used in game theory to denote the actions of all the agents but agent i, and single out the role of the action α^i of agent i in their own cost.

Definition 1.1 *For each $i \in [N]$, the* ***best responses of agent*** *i to the actions $\boldsymbol{\alpha}^{-i} = (\alpha^1, \dots, \alpha^{i-1}, \alpha^{i+1}, \dots, \alpha^N)$ of the other agents, is the subset of A^i, denoted by $br^i(\boldsymbol{\alpha}^{-i})$, comprising the minimizers of the cost function $A^i \ni \alpha \mapsto J^i(\alpha, \boldsymbol{\alpha}^{-i})$. In other words:*

$$br^i(\boldsymbol{\alpha}^{-i}) = \arg\min_{\alpha \in A^i} J^i(\alpha, \boldsymbol{\alpha}^{-i}). \tag{1.1}$$

More generally, we define the best response function (for all the agents simultaneously) as the function:

$$br : A \ni \boldsymbol{\alpha} = (\alpha^1, \dots, \alpha^N) \mapsto br(\boldsymbol{\alpha}) = br^1(\boldsymbol{\alpha}^{-1}) \times br^2(\boldsymbol{\alpha}^{-2}) \times \cdots \times br^N(\boldsymbol{\alpha}^{-N}) \subset A. \tag{1.2}$$

Clearly, this is a set value function (often called *correspondence*. A strategy profile $\boldsymbol{\alpha}^*$ is said to be a fixed point of the map br if

$$\boldsymbol{\alpha}^* \in br(\boldsymbol{\alpha}^*). \tag{1.3}$$

We can now define the notion of Nash equilibrium as follows:

Definition 1.2 *A strategy profile $\boldsymbol{\alpha}^* \in A$ is said to be a **Nash equilibrium** if it is a fixed point of the best response function br. Equivalently, if for each $i \in [N]$,*

$$J^i(\boldsymbol{\alpha}^*) \leqslant J^i(\alpha^i, \boldsymbol{\alpha}^{-i,*}), \qquad \text{for all } \alpha^i \in A^{(i)}. \tag{1.4}$$

The collection of all Nash equilibria will be denoted by $\mathcal{N}$:

$$\mathcal{N} = \{\boldsymbol{\alpha} \in A;\ \boldsymbol{\alpha} \in br(\boldsymbol{\alpha})\}. \tag{1.5}$$

In plain words $\boldsymbol{\alpha}^* \in \mathcal{N}$ if and only if, for each $i \in [N]$, if all the other agents take actions $\boldsymbol{\alpha}^{*-i}$, agent i cannot be better off by choosing an action departing from α^{*i}.

1.1.2 Social Cost and Price of Anarchy

Definition 1.3 *We define the **social cost** $J(\boldsymbol{\alpha})$ of a strategy profile $\boldsymbol{\alpha} \in A$ as the sum of the individual costs $J^i(\boldsymbol{\alpha})$, i.e.*

$$J(\boldsymbol{\alpha}) = \sum_{i=1}^N J^i(\boldsymbol{\alpha}) = \sum_{i=1}^N J^i(\alpha^i, \boldsymbol{\alpha}^{-i}). \tag{1.6}$$

*A strategy profile $\boldsymbol{\alpha}$ that minimizes the social cost among all admissible strategy profiles $\boldsymbol{\alpha} \in A$ of the N players is called a **social optimal** strategy profile, and it is denote by*

$$\boldsymbol{\alpha}^{*,SC} \in \arg\min_{\boldsymbol{\alpha} \in A} J(\boldsymbol{\alpha}) \tag{1.7}$$

We shall often divide the social cost $J(\boldsymbol{\alpha})$ by N when we vary N and especially when we try to study the limit $N \to \infty$. Obviously, this does not affect the set of minimizers of the social cost function. We shall also use the terminology *central planner* for someone attempting to search for a social optimal strategy profile. This type of global optimization is to be compared with the search for Nash equilibria where individual players are optimizing (selfishly) their own individual costs. To compare quantitatively the effects of these two forms of optimization, we introduce a couple of measures whose names try to capture their goals. We first give the commonly accepted definition of the *Price of Anarchy* as a measure of the inefficiency of selfish optimization in the above set-up of static strategic game. This measure tries to quantify the relative effect of individual optimizers trying to minimize their own individual costs without considering the effect of their choices on the overall social cost.

Definition 1.4 *The **Price of Anarchy** (PoA) in a static strategic game is defined as the ratio between the worst social cost of a Nash equilibrium and the optimal social cost among all admissible strategy profiles, namely*

$$PoA = \frac{\sup_{\boldsymbol{\alpha}\in\mathcal{N}} J(\boldsymbol{\alpha})}{\inf_{\boldsymbol{\alpha}\in A} J(\boldsymbol{\alpha})} = \frac{\sup\{J(\boldsymbol{\alpha}^*);\ \boldsymbol{\alpha}^*\in\mathcal{N}\}}{J(\boldsymbol{\alpha}^{*,SC})}. \tag{1.8}$$

A similar notion, though less popular, was introduced under the name of **Price of Stability** (PoS). It is defined as the ratio between the best social cost of a Nash equilibrium and the optimal social cost among all admissible strategy profiles:

$$PoS = \frac{\inf_{\boldsymbol{\alpha}\in\mathcal{N}} J(\boldsymbol{\alpha})}{\inf_{\boldsymbol{\alpha}\in A} J(\boldsymbol{\alpha})} = \frac{\inf\{J(\boldsymbol{\alpha}^*) : \boldsymbol{\alpha}^*\in\mathcal{N}\}}{J(\boldsymbol{\alpha}^{*,SC})}. \tag{1.9}$$

When the aggregate cost function J is non-negative (which is obviously the case when all the individual cost functions J^i are non-negative), the denominator is never greater than the numerator and as a result PoA is greater than or equal to 1. When it is equal to 1, Nash equilibria are not any worse than the result of centralized optimization. However, when $PoA > 1$, how greater than 1 it is tells us are worse Nash equilibria are when it comes to the overall social cost.

We shall compute explicitly some of these quantities, for example PoA, in some network games in which agents interact under the restriction of a protocol that leads to collective solutions producing unexpected (and very unnatural) results. This will shed some light on the very nature of Nash equilibra. See Subsection 2.3.4.

To close our discussion of optimality, we state the definition of the notion of Pareto optimality for the sake of completeness.

Definition 1.5 *A strategy profile $\boldsymbol{\alpha}^* \in A$ is said to be Pareto optimal if there exists no $\boldsymbol{\alpha} \in A$ such that*

- $\forall i \in [N], J^i(\boldsymbol{\alpha}) \leqslant J^i(\boldsymbol{\alpha}^*)$
- $\exists i_0 \in [N], J^{i_0}(\boldsymbol{\alpha}) < J^i(\boldsymbol{\alpha}^*)$

1.1.3 Search for a Nash Equilibrium

We review quickly the most popular methods used in the search for Nash equilibria.

1. Fixed Point theorem.
   The very definition of a Nash equilibrium is based on the notion of a fixed point for a specific map, the so-called best response function. So it is not surprising that most of the mathematical proofs of existence are based on the use of *fixed point theorems*. We shall give multiple examples in what follows, and in each case, we shall give precise references to the specific theorems we appeal to. For the time being we list some of them for the sake of definiteness.
   - The Banach fixed point theorem is also known as the contraction mapping theorem. It can be used when the mapping is a strict contraction and it guarantees not only existence, but also uniqueness of the fixed point. It is often used to show existence of ordinary and stochastic differential equations. We shall used it in several cases in what follows.

- Brower's fixed point theorem is one of the earliest ones. It guarantees existence of a fixed point for continuous maps on a compact set. It was extended by Schauder to more general settings. We shall use Schauder's fixed point theorem in several existence arguments in the sequel. For the sake of later reference, we state the form of Schauder's fixed point theorem which we shall use in the sequel.

  Theorem 1.6 *Let E be a Banach space, and let C be a non-empty closed convex set in E. Let $F : C \mapsto C$ be a continuous map such that $F(C) \subset K$ where K is a compact subset of C. Then F has a fixed point in K.*

  The proof can be found in [9, Exercise 6.26 p.179].
- Schauder's fixed point theorem was generalized by Kakutani to set valued functions. The appendix at the end of this chapter states the important facts on set-valued functions which we will use in these lectures, including Kakutani's fixed point theorem [20]. This theorem was used by John Nash in his seminal work [28]. We give the details in the next section.
- other fixed point theorems have been designed and used to take advantage of the special structure of some game models.
  - for example an order structure amenable to lattice theory is present in the model as in the case of unimodular games for which Tarski's fixed point theorem [38] is often used;
  - or in the case of some potential games for which direct iterative procedures can converge to equilibria.

For the sake of completeness, we also mention that in the case of differential games, variational inequality techniques have been used successfully when the cost functions exhibit strong convexity or monotonicity properties. See for example [36] for an introduction.

1.1.4 Mixed Strategies and Nash Theorem

This subsection is devoted to the statement and the proof of John Nash's famous existence theorem for equilibria in mixed strategies. But before we define such equilibria, we prove a first application of Kakutani's fixed point theorem to the existence of Nash equilibria in pure strategies.

Theorem 1.7 *Consider $\mathcal{G} = ([N], A = A^1 \times \cdots \times A^N, J = (J^i)_{i\in[N]})$ and let us assume that:*

- *for all $i \in [N]$, A^i is a non-empty convex compact subset of $\mathbb{R}^k$;*
- *for all $i \in [N]$, $J^i : A \to \mathbb{R}$ is continuous;*
- *for all $i \in [N]$, for all $\boldsymbol{\alpha}^{-i} \in A^{-i}$, $A^i \ni \alpha^i \mapsto J^i(\alpha^i, \boldsymbol{\alpha}^{-i})$ is quasi-convex, i.e., for all $a \in \mathbb{R}$, $\{\alpha \in A^i, J^i(\alpha, \boldsymbol{\alpha}^{-i}) \leqslant a\}$ is convex.*

Then $\mathcal{G}$ has at least one Nash equilibrium in pure strategies.

Proof: The proof relies on Kakutani fixed point theorem (see Theorem 1.29). Recall that the best response *br* is defined as the set valued function[1]

[1] A set value functions is often called a correspondence and it is denoted with a special arrow $\twoheadrightarrow$.

$$br(\boldsymbol{\alpha}) = br^1(\boldsymbol{\alpha}^{-1}) \times \cdots \times br^N(\boldsymbol{\alpha}^{-N}),$$

from A into the set 2^A of all the subsets of A, where

$$br^i(\boldsymbol{\alpha}^{-i}) = \arg\inf_{\alpha \in A^i} J^i(\alpha, \boldsymbol{\alpha}^{-i}),$$

and that our goal is to find a fixed point for br, i.e., $\boldsymbol{\alpha} \in A$ such that $\boldsymbol{\alpha} \in br(\boldsymbol{\alpha})$. To apply Kakutani fixed point theorem, we need to check that for all $\boldsymbol{\alpha} \in A$, $br(\boldsymbol{\alpha})$ is a non-empty, compact, and convex set, and that moreover, the graph of br is closed.

- For any $\boldsymbol{\alpha} \in A$ and every $i \in [N]$, since A^i is compact and the mapping $\alpha \mapsto J^i(\alpha, \boldsymbol{\alpha}^{-i})$ is continuous, the image (range) $J^i(A^i, \boldsymbol{\alpha}^{-i})$ is compact. So for every $i \in [N]$, there exists $\alpha^{i,*} \in A^i$ such that $\alpha^{i,*} \in br(\boldsymbol{\alpha}^{-i})$, which implies that $br(\boldsymbol{\alpha})$ is non-empty.
  Using the quasi-convexity assumption, for $i \in [N]$, if we choose $a = J^i(\alpha^{i,*}, \boldsymbol{\alpha}^{-i})$, the set of best response to $\boldsymbol{\alpha}^{-i}$ for player i, namely $br^i(\boldsymbol{\alpha}^{-i}) \subset A^i$ is closed (by continuity of $J^i(\cdot, \boldsymbol{\alpha}^{-i})$) hence compact, and because of the quasi-convexity assumption, convex as well.
- To prove that the graph of br is closed we start with a sequence $\{(\boldsymbol{\alpha}_n, \boldsymbol{\beta}_n)\}_{n\in\mathbb{N}}$ of pairs of strategy profiles in $A \times A$ satisfying $\boldsymbol{\beta}_n \in br(\boldsymbol{\alpha}_n)$ for every $n \in \mathbb{N}$, we suppose that $\boldsymbol{\alpha}_n \to \boldsymbol{\alpha} \in A$ and $\boldsymbol{\beta}_n \to \boldsymbol{\beta}$, and we prove that $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is in the graph of br, namely that $\boldsymbol{\beta} \in br(\boldsymbol{\alpha})$. For every $n \in \mathbb{N}$ and $i \in [N]$, we know that for every $\alpha'^i \in A^i$,
  $$J^i(\beta^i_n, \boldsymbol{\alpha}^{-i}_n) \leqslant J^i(\alpha'^i, \boldsymbol{\alpha}^{-i}_n)$$
  Taking the limit as $n \to \infty$ and using the continuity of J^i, we obtain that for every $\alpha'^i \in A^i$,
  $$J^i(\beta^i, \boldsymbol{\alpha}^{-i}) \leqslant J^i(\alpha'^i, \boldsymbol{\alpha}^{-i}).$$
  Thus, $\beta^i \in br^i(\boldsymbol{\alpha}^{-i})$ for every $i \in [N]$, or equivalently, $\boldsymbol{\beta} \in br(\boldsymbol{\alpha})$, showing that the graph of br is closed.

Kakutani's fixed point theorem (see Theorem 1.29) implies the desired conclusion. □

Definition 1.8 *Given a game $\mathcal{G} = ([N], A = A^1 \times \cdots \times A^N, J = (J^i)_{i\in[N]})$, we define the extended game $\tilde{\mathcal{G}} = ([N], \tilde{A} = \tilde{A}^1 \times \cdots \times \tilde{A}^N, \tilde{J} = (\tilde{J}^i)_{i\in[N]})$ as follows: for all $i \in [N]$, $\tilde{A}^i = \mathcal{P}(A^i)$ and*

$$\tilde{J}^i(\pi^1, \ldots, \pi^N) = \int_{A^1} \cdots \int_{A^N} J^i(\alpha^1, \ldots, \alpha^N)\pi^1(d\alpha^1) \ldots \pi^N(d\alpha^N).$$

The elements π^i of $\mathcal{P}(A^i)$ are called mixed strategies, and the elements $\boldsymbol{\pi} = (\pi^1, \ldots, \pi^N)$ of $\tilde{A}$ are called mixed strategy profiles. A Nash equilibrium for for the game $\tilde{\mathcal{G}}(N, \tilde{A}, \tilde{J})$ is called a Nash equilibrium in mixed strategies for $\mathcal{G}$.

Nash's original theorem on the subject [28] was to prove existence of such Nashe equilibrium in mixed strategies, without any convexity assumptions. We now prove this famous result in a slightly more general context than the finite action space case.

Theorem 1.9 (Nash Existence Theorem) *Let us assume that the game $\mathcal{G} = ([N], A = A^1 \times \cdots \times A^N, J = (J^i)_{i\in[N]})$ is such that:*

- *for all $i \in [N]$, A^i is a non-empty, compact subset of a metric space;*
- *for all $i \in [N]$, $J^i : A \to \mathbb{R}$ continuous.*

Then $\mathcal{G}$ has at least one Nash equilibrium in mixed strategies.

Proof: The proof consists in applying Theorem 1.7 above to the extended game. So, at least in an indirect way, it relies on Kakutani's fixed point theorem. Since A^i is assumed to be a compact metric space, $\mathcal{P}(A^i)$ is also a compact metric space when endowed with the topology of weak convergence. A sequence of probability measures $(\mu_n)_{n\geqslant 1}$ is said to converge weakly to the probability measure μ if and only if:

$$\forall f : A^i \to \mathbb{R} \text{ continuous}, \int f d\mu_n \to \int f d\mu.$$

Notice that $\tilde{A}^i = \mathcal{P}(A^i)$ is obviously convex and non-empty. The next step is to prove the continuity of $\tilde{J}^i : \tilde{A} \to \mathbb{R}$. This is an easy consequence of the continuity of J^i. Indeed, if we assume that $\boldsymbol{\pi}_n = (\pi_n^1, \dots, \pi_n^N)$ converges weakly to $\boldsymbol{\pi} = (\pi^1, \dots, \pi^N)$, each π_n^i converges weakly to π^i and:

$$\begin{aligned}\tilde{J}^i(\pi_n^1, \dots, \pi_n^N) &= \int_{A^1} \cdots \int_{A^N} J^i(\alpha^1, \dots, \alpha^N)\pi_n^1(d\alpha^1)\dots\pi_n^N(d\alpha^N) \\ &= \int_{A^1} \pi_n^1(d\alpha^1) \int_{A^2} \pi_n^2(d\alpha^2) \dots \int_{A^N} J^i(\alpha^1, \dots, \alpha^N)\pi_n^N(d\alpha^N).\end{aligned}$$

Treating $\alpha^1, ..., \alpha^{N-1}$ as fixed, one sees that the right most integral converges to $\int_{A^N} J^i(\alpha^1, \dots, \alpha^N)\pi^N(d\alpha^N)$, then, treating $\alpha^1, ..., \alpha^{N-2}, \alpha^N$ as fixed, one sees that the integral with respect to $\pi_n^{N-1}(d\alpha^{N-1})$ converges toward the corresponding integral with respect to $\pi^{N-1}(d\alpha^{N-1})$, and repeating the argument N times we conclude that $\tilde{J}^i(\pi_n)$ converges toward

$$\int_{A^1} \pi^1(d\alpha^1) \int_{A^2} \pi^2(d\alpha^2) \dots \int_{A^N} J^i(\alpha^1, \dots, \alpha^N)\pi^N(d\alpha^N)$$

which is equal to $\tilde{J}^i(\pi)$, proving the continuity of $\tilde{J}^i$.

FInally, we pick $i \in [N]$ and $\boldsymbol{\pi}^{-i} = (\pi_j)_{j\neq i} \in \tilde{A}^{-i}$ where $\tilde{A}^{-i} = \prod_{j\neq i} \mathcal{P}(A^j)$. and we prove that the mapping

$$\mathcal{P}(A^i) \ni \pi^i \to \tilde{J}^i(\pi^i, \boldsymbol{\pi}^{-i}) \in \mathbb{R}$$

is quasi-convex, namely that for each $a \in \mathbb{R}$, the set $\{\pi \in \mathcal{P}(A^i) : \tilde{J}^i(\pi, \boldsymbol{\pi}^{-i}) \leqslant a\}$ is convex. But if $\theta, \nu \in \{\pi \in \mathcal{P}(A^i) : \tilde{J}^i(\pi, \boldsymbol{\pi}^{-i}) \leqslant a\}$ and $\lambda \in [0, 1]$, using the definition of $\tilde{J}^i$, we easily check that

$$\tilde{J}^i(\lambda\theta + (1-\lambda)\nu, \boldsymbol{\pi}^{-i}) = \lambda\tilde{J}^i(\theta, \boldsymbol{\pi}^{-i}) + (1-\lambda)\tilde{J}^i(\nu, \boldsymbol{\pi}^{-i}) \leqslant a$$

which completes the check of the assumptions of Theorem 1.7, hence, completing the proof of the present Theorem 1.9. □

Remark 1.10 • *The weak convergence in the space of probability $\mathcal{P}(A^i)$ can be metrizised by the Prokhorov distance or the Wasserstein distance.*

- *If A^i is simply a metric space and $\tilde{A}^i \subset \mathcal{P}(A^i)$, it is possible to check the compactness of $\tilde{A}^i$ using Prokhorov's theorem. For example, if A^i is a separable complete metric space, Prokhorov's theorem says that $\tilde{A}^i$ is tight (namely relatively compact) if and only if the closure of $\tilde{A}^i$ in $(\mathcal{P}(A^i))$ for the topology of the weak convergence is compact.*

1.1.5 **Games with Mean Field Interactions**

We shall say that a game model has mean field interactions (or in short that it is a mean field game model) if all the admissible strategy sets A^i are equal to a single set A^0 and if the individual cost functions J^i are of the following form:

$$J^i(\boldsymbol{\alpha}) = J^i(\alpha^i, \boldsymbol{\alpha}^{-i}) = J\big(\alpha^i, \frac{1}{N-1}\sum_{j\neq i}\delta_{\alpha^j}\big) \tag{1.10}$$

for some function $J : A^0 \times \mathcal{P}(A^0) \ni (\alpha, \mu) \mapsto J(\alpha, \mu) \in \mathbb{R}$. Here and in the following we use the notation $\mathcal{P}(A^0)$ to denote the set of probability measures on A^0 which is implicitly assumed to be a subset of a Euclidean space $\mathbb{R}^k$ and equipped with its Borel σ-field. Also, we used the notation δ_x for the point measure with unit mass concentrated at the point x. So the measure appearing as the second argument of the function J in formula (1.10) is the empirical measure of the actions taken by the players different from player i. For the sake of later reference, we state the following definition of an empirical measure.

Definition 1.11 *If $\boldsymbol{x} = (x^1, \cdots, x^k) \in E^k$ where $(E, \mathcal{E})$ is a measure space, we denote by:*

$$\overline{\mu}^{\boldsymbol{x}} = \frac{1}{k}\sum_{i=1}^{k}\delta_{x^i} \tag{1.11}$$

the empirical measure of $\boldsymbol{x}$. Recall that the point measure δ_x is defined for $x \in E$ by $\delta_x(A) = 1$ if $x \in A$ and 0 otherwise for any $A \in \mathcal{E}$.

The following example was proposed by Lasry and Lions as a simple example of mean field games. It is usually called "*where should I put my towel on the beach?*" and it is presented in full detail in [11]. Imagine that A is a compact set representing a beach, and $\alpha^0 \in A$ is the location of a refreshment stand. Given two positive numbers a and b each individual beach goer i will choose a location α^i on the beach to minimize:

$$J^i(\boldsymbol{\alpha}) = ad(\alpha^i, \alpha^0) - b\sum_{j\neq i} d(\alpha^i, \alpha^j)$$

where d is a distance on A. So each individual wants to minimize the distance to the refreshment stand while at the same time avoiding to be too close to the other beach goers. This is clearly a game with mean field interaction in the sense of the above definition if we use the function:

$$J(\alpha, \mu) = ad(\alpha, \alpha^0) - b\int_A d(\alpha, \alpha')\, \mu(d\alpha').$$

The standard approach to mean field games is to work in the limit $N \to \infty$ of large games for which one expects that the empirical measures $\overline{\mu}^{\boldsymbol{\alpha}}$ converge toward a measure μ, in which case, the search for Nash equilibria for the N-player game is replaced by an approximation obtained by having a generic player respond optimally to the distribution μ of actions of the other players. Mathematically, the search for these approximate Nash equilibria can be summarized as:

- compute the **Best Response Function**: $\mu \mapsto \hat{\boldsymbol{\alpha}}^\mu = \arg\inf_{\alpha\in A} J(\alpha, \mu)$;
- find a **Fixed Point** $\hat\mu$ satisfying supp $\hat\mu \subset \arg\inf_{\alpha\in A} J(\alpha, \hat\mu)$,

where supp $\hat\mu$ denotes the *support* of the measure μ. In words, this second condition means that the measure $\hat\mu$ needs to be concentrated on the best responses. We shall come back to these types of games later on in the lectures.

1.2 Potential Games

As before, we denote by $A = A^1 \times \dots A^N$ the set of strategy profiles, and for any $i \in \{1, \dots, N\}$, the couple $(\alpha, \boldsymbol{\alpha}^{-i}) \in A$ with $\alpha \in A^i$ and $\boldsymbol{\alpha}^{-i} \in A^1 \times \dots A^{i-1} \times A^{i+1} \times \dots \times A^N$ stands for the strategy profile $(\alpha^1, \dots, \alpha^{i-1}, \alpha, \alpha^{i+1}, \dots, \alpha^N)$. Also, if φ is any function $\varphi : A \ni \boldsymbol{\alpha} \mapsto \varphi(\boldsymbol{\alpha}) \in \mathbb{R}$, we shall use the same notation φ if for a given player $i \in [N]$, we re-order the arguments and evaluate the function at the strategy profile $\boldsymbol{\alpha} = (\alpha^i, \boldsymbol{\alpha}^{-i})$, using the notation $\varphi(\alpha^i, \boldsymbol{\alpha}^{-i})$.

This section is devoted to the presentation of the important class of potential games. It was introduced by Monderer and Shapley in 1996.

Definition 1.12 *A game $\mathcal{G} = \big([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1\dots N}\big)$ is said to be a* ***Potential Game*** *if there exists a map $\varphi : A \ni \boldsymbol{\alpha} \mapsto \varphi(\boldsymbol{\alpha}) \in \mathbb{R}$, called a potential function, such that for every player $i \in [N]$, and any $\boldsymbol{\alpha}^{-i} \in A^{-i}$, we have:*

$$J^i(\beta, \boldsymbol{\alpha}^{-i}) - J^i(\alpha, \boldsymbol{\alpha}^{-i}) = \varphi(\beta, \boldsymbol{\alpha}^{-i}) - \varphi(\alpha, \boldsymbol{\alpha}^{-i}), \qquad \forall\, \alpha, \beta \in A^i \tag{1.12}$$

In words, this definition means that the change in a single player's cost function due to a change in their own strategy only, is given by exactly the same change of the potential function φ.

Remark 1.13 *If the sets of admissible strategies $(A^i)_{i=1,\dots,N}$ are intervals in $\mathbb{R}$ and each cost function J^i is continuous and differentiable, then the fact that the game is a potential game means that for every $i \in [N]$ we have:*

$$\frac{\partial}{\partial \alpha^i} J^i(\alpha, \boldsymbol{\alpha}^{-i}) = \frac{\partial}{\partial \alpha^i} \varphi(\alpha, \boldsymbol{\alpha}^{-i}), \qquad \alpha \in A^i, \boldsymbol{\alpha}^{-i} \in A^{-i}.$$

In the set-up of the above remark, standard calculus can be used to prove the following characterization.

Theorem 1.14 *Let $\mathcal{G} = \big([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1\dots N}\big)$ be a game in which the strategy sets A^i are intervals of real numbers and let us assume that the cost functions are twice continuously differentiable. Then $\mathcal{G}$ is a potential game if and only if:*

$$\frac{\partial^2 J^i}{\partial \alpha^i \partial \alpha^j} = \frac{\partial^2 J^j}{\partial \alpha^i \partial \alpha^j}, \quad \text{for every } i, j \in \{1, \dots, N\}. \tag{1.13}$$

Moreover, if the cost functionals satisfy (1.13) *and $\boldsymbol{\alpha}^0$ is an arbitrary strategy profile in A, then a potential for $\mathcal{G}$ is given by the function:*

$$\varphi(\boldsymbol{\alpha}) = \sum_{i=1}^{N} \int_0^1 (\beta^i)'(t) \frac{\partial J^i}{\partial \alpha^i}\big(\beta(t)\big) dt \tag{1.14}$$

where $\beta : [0,1] \to A$ is any piecewise continuously differentiable path in A that connects $\boldsymbol{\alpha}^0$ to $\boldsymbol{\alpha}$ (i.e. $\beta(0) = \boldsymbol{\alpha}^0$ and $\beta(1) = \boldsymbol{\alpha}$).

The interested reader can find a proof in [26, Theorem 4.5].

Important Consequence of the Definition

If φ is a potential function for a game $\mathcal{G} = \big([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1\dots N}\big)$ and if

$$\boldsymbol{\alpha}^* \in \arg\inf_{\boldsymbol{\alpha}\in A} \varphi(\boldsymbol{\alpha}),$$

then for any $i \in [N]$ and any $\boldsymbol{\alpha} \in A^{-i}$,

$$J^i(\alpha, \boldsymbol{\alpha}^{*-i}) - J^i(\alpha^{*,i}, \boldsymbol{\alpha}^{*-i}) = \varphi(\alpha, \boldsymbol{\alpha}^{*-i}) - \varphi(\alpha^{*,i}, \boldsymbol{\alpha}^{*-i}) \geqslant 0, \tag{1.15}$$

which shows that $\boldsymbol{\alpha}^*$ is a Nash equilibrium for the game. Conversely, if $\boldsymbol{\alpha}^*$ is a Nash equilibrium, it is clear that for each $i \in [N]$, α^{*i} is a minimum for the function $A^i \ni \alpha \mapsto \varphi(\alpha, \alpha^{*-i})$, which in many cases will be enough to conclude that $\boldsymbol{\alpha}^*$ is a minimum of φ.

Consequences

- If A^i is finite for every $i \in [N]$, so is the number of admissible strategies $\boldsymbol{\alpha} \in A$ and there is at least one Nash equilibrium when the game is potential.
- If A is compact and the potential function φ is continuous, then there exists a Nash equilibrium.
- Moreover, if A is convex and φ is strictly convex, then the Nash equilibrium is unique.

Definition 1.15 *A sequence $(\boldsymbol{\alpha}_n)_{n\geqslant 0}$ of strategy profiles is called a* ***path*** *if for every $n \geqslant 0$, $\boldsymbol{\alpha}_{n+1}$ is obtained from $\boldsymbol{\alpha}_n$ by allowing one player, say $i_n \in [N]$, to change they strategy.*

For example, if $\alpha_n^{i_n}$ changes into $\beta \in A^{i_n}$, then the strategy profile $\boldsymbol{\alpha}_n$ changes to a new strategy profile $\boldsymbol{\alpha}_{n+1}$:

$$\boldsymbol{\alpha}_n = (\alpha_n^1, \dots, \alpha_n^N) \longrightarrow \boldsymbol{\alpha}_{n+1} := (\alpha_n^1, \dots, \alpha_n^{i_n-1}, \beta, \alpha_n^{i_n+1}, \dots, \alpha_n^N).$$

Definition 1.16 *A path is called an* ***improvement path*** *if*

$$J^{i_n}(\boldsymbol{\alpha}_{n+1}) < J^{i_n}(\boldsymbol{\alpha}_n).$$

The following result is quite useful when the set of feasible strategies are finite.

Proposition 1.17 *The end-point of any finite improvement path is a Nash equilibrium.*

Proof: Let $(\boldsymbol{\alpha}_n)_{0\leqslant n\leqslant k}$ be a finite improvement path. If $\boldsymbol{\alpha}_k$ is not a Nash equilibrium, then there exists a player $i(k)$ who can deviate from $\alpha_k^{i(k)}$ and lower they cost by using a different strategy, say $\beta^{i(k)}$. Consequently, we could add the strategy profile $(\beta^{i(k)}, \boldsymbol{\alpha}_k^{-i(k)})$ to the path and extend it while still improving the cost. This is a contradiction. □

Remark 1.18 *Obviously,* myopic best response *dynamics, or even* better response dynamics*, whether they are deterministic or random, create improvement paths.So for finite potential games for which the set A of strategy profiles is finite, every improvement path is finite and every sequence of better or best responses converge to a Nash equilibrium. While attractive in principle, this result does not always lead to a practical algorithm to compute Nash equilibria because the cardinality of the set A may be very large and the search for improvements may be slow.*

Algorithmic Computation of Equilibria

As long as we are concentrating on game models for which the set A of admissible strategy profiles is finite, it is easy to design an algorithm to compute Nash equilibria in pure strategies for potential games.

Algorithm 1 Computation of a Pure Nash Equilibrium for a Potential Game

---

Choose an initial value $\boldsymbol{\alpha}_0 \in A$ (arbitrary)
Set $\ell = 0$
while $\boldsymbol{\alpha}_\ell$ is not a Nash Equilibrium **do**
  Find a player i_ℓ and an admissible action $\alpha \in A^{i_\ell}$ such that $J^{i_\ell}(\alpha, \boldsymbol{\alpha}_{i_\ell}^{-i_\ell}) < J^{i_\ell}(\boldsymbol{\alpha}_{i_\ell})$
  Set $\boldsymbol{\alpha}_{\ell+1} = (\alpha, \boldsymbol{\alpha}_{i_\ell}^{-i_\ell})$
  Set $\ell = \ell + 1$
end while
Return $\boldsymbol{\alpha}_\ell$

1.2.1 Approximate Nash Equilibria

Prompted by the previous remarks, the following notion of approximate Nash equilibrium is often used to alleviate the computational difficulties.

Definition 1.19 *If $\epsilon > 0$, a strategy profile $\boldsymbol{\alpha}^* \in A$ is said to be an ϵ-approximate Nash equilibrium if:*

$$\forall i \in [N], \quad J^i(\boldsymbol{\alpha}^*) \leqslant J^i(\alpha, \boldsymbol{\alpha}^{*-i}) + \epsilon, \qquad \textit{for all } \alpha \in A^i.$$

Clearly, if $\epsilon > 0$ and if the cost functions J^i are bounded (or equivalently if the potential function is bounded):

- every ϵ-improvement path (for which each step reduces the cost by an amount at least ϵ) is finite;
- every better response or best response ϵ-improvement path converges to an ϵ-Nash equilibrium in a finite number of steps.

The next result shows that a potential function, if it exists, is uniquely determined up to a constant.

Proposition 1.20 *The potential function is determined up to a constant. In other words, if φ_1 and φ_2 are potential functions for the same potential game $\mathcal{G} = ([N], (A^i)_{i=1,\dots,N}, (J^i)_{i=1,\dots,N})$, then there exists a constant $c \in \mathbb{R}$ such that*

$$\varphi_1 = \varphi_2 + c.$$

Proof: Pick an arbitrary $\boldsymbol{\alpha}^0 = (\alpha^{0,1}, \dots, \alpha^{0,N}) \in A$, then define a function $H : A \ni \boldsymbol{\alpha} = (\alpha^1, \dots, \alpha^N) \mapsto H(\boldsymbol{\alpha}) \in \mathbb{R}$ by

$$H(\boldsymbol{\alpha}) = \sum_{i=1}^{N} J^i(\boldsymbol{\alpha}_i) - J^i(\boldsymbol{\alpha}_{i+1}) \tag{1.16}$$

where $\boldsymbol{\alpha}_1 = \boldsymbol{\alpha}$, and $\boldsymbol{\alpha}_{i+1} = (\alpha^{0,1}, \ldots, \alpha^{0,i}, \alpha^{i+1}, \ldots, \alpha^N)$ for $i = 1, \ldots N-1$, and $\boldsymbol{\alpha}_{N+1} = \boldsymbol{\alpha}^0$. If φ_1 and φ_2 are potential functions, we can have

$$H(\boldsymbol{\alpha}) = \varphi_1(\boldsymbol{\alpha}) - \varphi_1(\boldsymbol{\alpha}^0), \quad \text{and} \quad H(\boldsymbol{\alpha}) = \varphi_2(\boldsymbol{\alpha}) - \varphi_2(\boldsymbol{\alpha}^0).$$

Let $c = \varphi_1(\boldsymbol{\alpha}^0) - \varphi_2(\boldsymbol{\alpha}^0)$, the conclusion then follows. □

1.2.2 **Example 1: Cournot Competition**

Let us assume that N firms, labelled by the integer numbers in $\{1, \ldots, N\}$ produce the same good which they try to sell on the market. The quantity produced by firm i is denoted by a positive real number $q^i \in [0, \infty)$. The choice of such a quantity is the control of a firm on the economy outcome. The production profile of these firms is then denoted by $\boldsymbol{q} = (q^1, \ldots, q^N) \in [0, \infty)^N$. The cost for producing q units of good (q can be real numbers) is the same for all the firms. It is given by a function $c : [0, \infty) \ni q \mapsto c(q) \in \mathbb{R}$. We denote by Q the total production in the economy, i.e.

$$Q = \sum_{i=1}^{N} q^i.$$

The price p of the product is determined using a time honored procedure in economics: satisfying a *clearing constraint* by matching supply and demand on the market. This is one by positing an inverse demand function $f : [0, \infty) \to \mathbb{R}$, in which case the clearing condition reads, $p(\boldsymbol{q}) = f(Q)$. For the sake of simplicity, we choose an affine function:

$$f(Q) = a - bQ,$$

for some constants $a, b > 0$. Each firm $i \in [N]$ tries to minimize its costs as given by cost functions J^i representing the negative of the net revenues generated from its production and its sales. More precisely, the cost function $J^i : [0, \infty)^N \to \mathbb{R}$ is given by:

$$J^i(\boldsymbol{q}) = -\left[p(\boldsymbol{q}) \cdot q^i - c(q^i)\right] = c(q^i) - aq^i + bq^i \sum_{j=1}^{N} q^j.$$

The Cournot competition game is a potential game. In order to prove this claim, we compute the partial derivative $\partial^2 J^i / \partial q^i \partial q^j$ for every $i, j \in \{1, \ldots, N\}$. From the above definition of the cost functions, it is easy to see that:

$$\frac{\partial^2 J^i(\boldsymbol{q})}{\partial q^i \partial q^j} = \frac{\partial}{\partial q_i}(bq_i) = b = \frac{\partial^2 J^j(\boldsymbol{q})}{\partial q^i \partial q^j}, \qquad \forall\, i, j \in \{1, \ldots N\},\ \forall\, \boldsymbol{q} \in [0, \infty)^N.$$

which proves that the Cournot competition game is a potential game. Moreover, we can also use Rosenthal's result to construct a potential function:

$$\varphi(\boldsymbol{q}) = \sum_{i=1}^{N} \int_0^1 (\beta^i(t))' \frac{\partial^2 J^i}{\partial q^i} (\beta^i(t), \boldsymbol{\beta}^{-i}(t)) dt$$

from a linear path $\boldsymbol{\beta} : [0,1] \ni t \mapsto \boldsymbol{\beta}(t) = t\boldsymbol{q} \in [0,\infty)^N$. A straight computation gives:

$$(\beta(t)^i)' = q^i, \qquad \text{and} \qquad \frac{J^i(\boldsymbol{\beta}(t))}{\partial q^i} = \frac{\partial c(tq^i)}{\partial q^i} - a + 2btq^i + bt\sum_{j\neq i}^{N} q^j$$

so that

$$\boxed{\varphi(\boldsymbol{q}) = \sum_{i=1}^{N} c(q^i) - a\sum_{i=1}^{N} q^i + b\sum_{i=1}^{N} \cdot (q^i)^2 + \frac{b}{2}\sum_{i=1}^{N} q^i \sum_{j\neq i}^{N} q^j.}$$

1.2.3 Example 2: Congestion Games

Here, we list the major components of a static congestion game. We shall revisit this class of game models in next chapter under the name of network congestion games when a specific graph structure is underpinning the game model. For the time being, we assume that a congestion game model comprises:

- N players enumerated by $\{1,\ldots,N\}$.
- A set of resources denoted by E.
- For each player $i \in [N]$, a set A^i of feasible strategies such that $A^i \subseteq 2^E$ the collection of subsets of E. Namely, for player i, a strategy $\alpha^i \in A^i$ is a set of resources $\alpha^i \subseteq E$. As usual, we denote by $A = A^1 \times \ldots A^N$ the collection of feasible strategy profiles for the N players.
- For each resource $e \in E$, a load function $k_e : A \to \{0,\ldots,N\}$ defined as $k_e(\boldsymbol{\alpha}) = |\{i; e \in \alpha^i\}|$ giving, for every strategy profile $\boldsymbol{\alpha} \in A$ the number of players using the resource e.
- For each resource $e \in E$, a cost function $c_e : \{0,\ldots,N\} \to \mathbb{R}$.
- For each strategy profile $\boldsymbol{\alpha} \in A$, player i experiences a cost $J^i(\boldsymbol{\alpha})$ defined as the sum of the costs associated to the resources they are using, the cost for each such resource being a function of the number of agents using this particular resource. Mathematically, this means:
$$J^i(\boldsymbol{\alpha}) = \sum_{e\in\alpha^i} c_e(k_e(\boldsymbol{\alpha})). \tag{1.17}$$

We denote such a congestion game of N players by the t-uple

$$(E, (A^i)_{i=1,\cdots,N}, (c_e)_{e\in E}, (k_e)_{e\in E}, (J^i)_{i=1,\ldots,N}).$$

We shall study the important class of network congestion games in the next chapter. They are also known under the name of *undirected Shapley design games*. The following existence result is the most important theoretical result on this class of games. It is due to R.W. Rosenthal in 1973 [33].

Proposition 1.21 *Every congestion game with a finite set of resources has a pure Nash equilibrium.*

Proof: We first prove that the function:

$$\varphi(\boldsymbol{\alpha}) = \sum_{e \in E} \sum_{k=1}^{k_e(\boldsymbol{\alpha})} c_e(k) \tag{1.18}$$

is a potential function for the game. For every $\boldsymbol{\alpha} \in A$, $i \in \{1, \ldots, N\}$, and $\tilde{\alpha} \in A^i$,

$$\begin{aligned}
\varphi(\tilde{\alpha}, \boldsymbol{\alpha}^{-i}) &= \sum_{e \in E} \sum_{k=1}^{k_e(\tilde{\alpha}, \boldsymbol{\alpha}^{-i})} c_e(k) \\
&= \sum_{e \in E} \left(\sum_{k=1}^{k_e(\boldsymbol{\alpha})} c_e(k) + \mathbf{1}_{\{e \in \tilde{\alpha} \setminus \alpha^i\}} c_e(k_e(\tilde{\alpha}, \boldsymbol{\alpha}^{-i})) - \mathbf{1}_{\{e \in \alpha^i \setminus \tilde{\alpha}^i\}} c_e(k_e(\boldsymbol{\alpha})) \right) \\
&= \sum_{e \in E} \sum_{k=1}^{k_e(\boldsymbol{\alpha})} c_e(k) + \sum_{e \in \tilde{\alpha} \setminus \alpha^i} c_e(k_e(\tilde{\alpha}, \boldsymbol{\alpha}^{-i})) - \sum_{e \in \alpha^i \setminus \tilde{\alpha}^i} c_e(k_e(\boldsymbol{\alpha})) \\
&= \sum_{e \in E} \sum_{k=1}^{k_e(\boldsymbol{\alpha})} c_e(k) + \sum_{e \in \tilde{\alpha}} c_e(k_e(\tilde{\alpha}, \boldsymbol{\alpha}^{-i})) - \sum_{e \in \alpha^i} c_e(k_e(\boldsymbol{\alpha})) \\
&= \varphi(\boldsymbol{\alpha}) + J^i(\tilde{\alpha}, \boldsymbol{\alpha}^{-i}) - J^i(\boldsymbol{\alpha})
\end{aligned} \tag{1.19}$$

which proves that φ is indeed a potential function for the game. To conclude the proof, we construct a finite improvement path. Recall the result of Proposition 1.17 and Algorithm 1. Let $\boldsymbol{\alpha}_0, \boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \cdots$ of strategy profiles such that for each $\ell = 0, 1, 2, \cdots$ there is a player i_ℓ such that $\boldsymbol{\alpha}_{\ell+1} = (\alpha, \boldsymbol{\alpha}_{i_\ell}^{-i_\ell})$ for some $\alpha \in A^{i_\ell}$ such that $J^{i_\ell}(\alpha, \boldsymbol{\alpha}_{i_\ell}^{-i_\ell}) < J^{i_\ell}(\boldsymbol{\alpha}_{i_\ell})$. Then, the sequence $\varphi(\boldsymbol{\alpha}_{i_\ell}$ is strictly decreasing, and since the set A is finite, this sequence must terminate after a finite number of steps. □

1.3 STRATEGIC GAME ZOOLOGY 101

The remaining of this chapter is devoted to a mathematical classification of the various (finite) strategic games, at least of those models of interest to us in these lectures.

Definition 1.22 *An N-player strategic game $\mathcal{G} = ([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1,\ldots,N})$ is said to be:*
• *a **coordination game** if:*

$$\forall i, j \in \{1, \ldots, N\}, \forall \boldsymbol{\alpha} \in A, \quad J^i(\boldsymbol{\alpha}) = J^j(\boldsymbol{\alpha}) : \tag{1.20}$$

• *a **dummy game** if:*

$$\forall i \in \{1, \ldots, N\}, \forall \boldsymbol{\alpha} \in A, \forall \tilde{\alpha} \in A^i, \quad J^i(\boldsymbol{\alpha}) = J^i(\tilde{\alpha}, \boldsymbol{\alpha}^{-1}). \tag{1.21}$$

Our first theoretical result identifies potential games as *aggregates* of games of the above types.

Proposition 1.23 *An N-player strategic game $\mathcal{G} = ([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1,\dots,N})$ is a potential game if and only if there exist N-tuples of functions $(\varphi_i^c)_{i=1,\dots,N}$ and $(\varphi_i^d)_{i=1,\dots,N}$ such that: for every $i \in \{1,\dots,N\}$:*

$$J^i = \varphi_i^c + \varphi_i^d$$

and

- *the game $((A^i)_{i=1,\cdots,N}, (\varphi_i^c)_{i=1,\dots,N})$ is a coordination game;*
- *the game $((A^i)_{i=1,\cdots,N}, (\varphi_i^d)_{i=1,\dots,N})$ is a dummy game.*

Proof: ($\Leftarrow$): Let φ be a potential function for the original game, and define for all $i \in \{1,\dots,N\}$ and $\boldsymbol{\alpha} \in A$ the functions φ_i^c and φ_i^d by:

$$\begin{cases} \varphi_i^c(\boldsymbol{\alpha}) = \varphi(\boldsymbol{\alpha}) \\ \varphi_i^d(\boldsymbol{\alpha}) = J^i(\boldsymbol{\alpha}) - \varphi(\boldsymbol{\alpha}). \end{cases}$$

Clearly $J^i = \varphi_i^c + \varphi_i^d$ by definition. Moreover, $\varphi_i^c(\boldsymbol{\alpha}) = \varphi_j^c(\boldsymbol{\alpha})$ for all $i \neq j$, proving that the game $((A^i)_{i=1,\cdots,N}, (\varphi_i^c)_{i=1,\dots,N})$ is a coordination game. Finally:

$$\varphi_i^d(\boldsymbol{\alpha}) - \varphi_i^d(\alpha, \boldsymbol{\alpha}^{-i}) = \varphi(\boldsymbol{\alpha}) + \varphi(\alpha, \boldsymbol{\alpha}^{-i}) + J^i(\boldsymbol{\alpha}) - J^i(\alpha, \boldsymbol{\alpha}^{-i}) = 0$$

since φ is a potential function for the original game, proving that $((A^i)_{i=1,\cdots,N}, (\varphi_i^d)_{i=1,\dots,N})$ is a dummy game.

($\Rightarrow$): We define the function φ by $\varphi = \varphi_i^c$ for any $i \in [N]$. Such a function φ is a potential function for the original game because, if $i \in [N]$ and $\boldsymbol{\alpha} \in A$, for any $\alpha \in A^i$ we have:

$$\begin{aligned} J^i(\boldsymbol{\alpha}) - J^i(\alpha, \boldsymbol{\alpha}^{-i}) &= \varphi_i^c(\boldsymbol{\alpha}) + \varphi_i^d(\boldsymbol{\alpha}) - \varphi_i^c(\alpha, \boldsymbol{\alpha}^{-i}) - \varphi_i^d(\alpha, \boldsymbol{\alpha}^{-i}) \\ &= \varphi_i^c(\boldsymbol{\alpha}) - \varphi_i^c(\alpha, \boldsymbol{\alpha}^{-i}) \\ &= \varphi(\boldsymbol{\alpha}) - \varphi(\alpha, \boldsymbol{\alpha}^{-i}) \end{aligned}$$

which proves the desired result. □

We now introduce the notion of isomorphism between games which will allow us to classify game models.

Definition 1.24 *Two N-player games, $\mathcal{G} = ([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1,\dots,N})$ and $\mathcal{G}' = ([N], (A'^i)_{i=1,\cdots,N}, (J'^i)_{i=1,\dots,N})$ are said to be isomorphic if there exists N bijective mappings $\{\phi_i\}_{i=1\dots,N}$:*

$$\phi_i : A^i \ni \alpha^i \mapsto \phi_i(\alpha^i) \in A'^i$$

satisfying for every $i \in \{1,\dots,N\}$ and every $\boldsymbol{\alpha} = (\alpha^1,\dots,\alpha^n) \in A$,

$$J^i(\alpha^1,\dots,\alpha^N) = J'^i(\phi_1(\alpha^i),\dots,\phi_N(\alpha^N)).$$

With this definition in hand we can prove the first classification result.

Proposition 1.25 *1. Every coordination game is isomorphic to a congestion game.*
2. Every dummy game is isomorphic to a congestion game.
3. Every potential game is isomorphic to a congestion game.

Proof:

1. Let $\mathcal{G} = ([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1,\dots,N})$ be a coordination game and let us denote by $J(\boldsymbol{\alpha}) = J^i(\boldsymbol{\alpha})$ the cost function common to all players $i \in \{1,\dots,N\}$. We now construct a congestion game $(E, (A'^i)_{i=1,\cdots,N}, (c_e)_{e\in E}, (k_e)_{e\in E}, (J'^i)_{i=1,\dots,N})$ isomorphic to the game we started from by the following steps:
   - each strategy profile $\boldsymbol{\alpha} \in A$ is associated to a different resource $e(\boldsymbol{\alpha})$. The collection of resources is denoted by $E = \{e(\boldsymbol{\alpha}); \boldsymbol{\alpha} \in A\}$. In other words, the set E is indexed by the set of strategy profiles;
   - for each resource $e = e(\boldsymbol{\alpha}) \in E$, we define its cost function $c_e : \{0,\dots,N\} \to \mathbb{R}$ by:
   $$c_e(k) = c_{e(\boldsymbol{\alpha})}(k) = \mathbf{1}_{\{k=N\}} J(\boldsymbol{\alpha}), \qquad k \in \{0,\dots,N\};$$
   - for each player i, the feasible strategy set A'^i is defined by:
   $$A'^i = \{\phi_i(\alpha^i) : \alpha^i \in A^i\}$$
   where the mapping $\phi_i : A^i \to 2^E$ takes the form:
   $$\phi_i(\alpha^i) = \bigcup_{\boldsymbol{\alpha}^{-i}\in A^{-i}} \left\{e(\alpha^i, \boldsymbol{\alpha}^{-i})\right\}, \qquad \alpha^i \in A^i.$$
   - for each player i, the cost function $J'^i : A' \to \mathbb{R}$ is defined by
   $$J'^i(\boldsymbol{\alpha}') = \sum_{e\in\alpha'^i} c_e(k_e(\boldsymbol{\alpha}')),$$
   where $k_e(\boldsymbol{\alpha}') = |\{i : e \in \alpha'^i\}|$ is the load function associated to resource $e \in E$ evaluated at strategy profile $\boldsymbol{\alpha}' \in A'$.

   By construction, $(E, (A'^i)_{i=1,\cdots,N}, (c_e)_{e\in E}, (k_e)_{e\in E}, (J'^i)_{i=1,\dots,N})$ is a congestion game. We already used the notation $A' = A'^1 \times \dots A'^N$ for the collection of strategy profiles for the N players. To show that it is isomorphic to the game we started from, we notice first that by construction, the mapping $\{\phi_i\}_{i=1,\dots,N}$ are bijective from A^i onto A'^i for every player i. Moreover, for a given $i \in \{1,\dots,N\}$, for every resource $e = e(\boldsymbol{\beta}) \in E$ with some $\boldsymbol{\beta} \in A$, and for every strategy $\alpha'^i = \phi_i(\alpha^i)) \in A'^i$ with some $\alpha^i \in A^i$, we have
   $$e(\boldsymbol{\beta}) = e \in \alpha'^i \iff \beta^i = \alpha^i.$$
   Thus, if a resource $e = e(\boldsymbol{\beta}) \in E$ is fully loaded with a strategy profile $\boldsymbol{\alpha}' = (\phi_1(\alpha^1),\dots,\phi_N(\alpha^N) \in A'$ for some $\boldsymbol{\alpha} = (\alpha^1,\dots,\alpha^N) \in A$, we must have:
   $$\begin{aligned} k_e(\boldsymbol{\alpha}') = N &\iff e = e(\boldsymbol{\beta}) \in \tilde{\alpha}^i, \qquad i = 1,\dots,N \\ &\iff \beta^i = \alpha^i, \qquad i = 1,\dots,N \\ &\iff \boldsymbol{\beta} = \boldsymbol{\alpha} \end{aligned}$$
   Finally, for every $\boldsymbol{\alpha} \in A$, the cost function for player i evaluated at the strategy profile $\boldsymbol{\alpha}' = (\phi_1(\alpha^1),\dots,\phi_N(\alpha^N)) \in A'$ satisfies
   $$J'^i(\boldsymbol{\alpha}') = \sum_{e\in\alpha'^i} c_e(k_e(\boldsymbol{\alpha}')) = \sum_{e(\boldsymbol{\beta})\in\phi_i(\alpha^i)} c_e(k_e(\tilde{\boldsymbol{\alpha}})) = \sum_{\{e(\boldsymbol{\beta})\in E:\beta^i=\alpha^i\}} \mathbf{1}_{k_{e(\boldsymbol{\beta})}(\tilde{\boldsymbol{\alpha}})=N} \cdot J(\boldsymbol{\beta})$$
   so that
   $$J'^i(\phi_1(\alpha^1),\dots,\phi_N(\alpha^N)) = \sum_{\{e(\boldsymbol{\beta})\in E:\beta^i=\alpha^i\}} \mathbf{1}_{\boldsymbol{\beta}=\boldsymbol{\alpha}} \cdot J(\boldsymbol{\beta}) = J(\boldsymbol{\alpha}) = J^i(\boldsymbol{\alpha}),$$
   which completes the proof of the desired isomorphism.

2. Let $\mathcal{G} = ([N], (A^i)_{i=1,\cdots,N}, (J^i)_{i=1,\ldots,N})$ be a dummy game. For each $i \in [N]$ and each $\boldsymbol{\alpha}^{-i} \in A^{-i}$ we introduce the resource $e(\boldsymbol{\alpha}^{-i})$ and the corresponding cost function $c_{e(\boldsymbol{\alpha}^{-i})}$ defined by:
$$c_{e(\boldsymbol{\alpha}^{-i})}(j) = \begin{cases} J^i(\alpha, \boldsymbol{\alpha}^{-i}) & \text{if } j = 1 \\ 0 & \text{otherwise} \end{cases}$$
for an arbitrary $\alpha \in A^i$ whose choice does not really matter since we are dealing with a dummy game. Let E be the set of resources introduced this way and for $i \in [N]$ and each $\alpha^i \in A^i$ let us set:
$$\phi_i(\alpha^i) = \bigcup_{\boldsymbol{\alpha}^{-i} \in A^{-i}} \{e(\boldsymbol{\alpha}^{-i})\} \cup \bigcup_{j \neq i} \bigcup_{\boldsymbol{\beta}^{-j} \in A^{-j}, \alpha^i \neq \beta^i} \{e(\beta^{-j}\}$$
and as before:
$$A'^i = \{\phi_i(\alpha^i) : \alpha^i \in A^i\}.$$
We now check that the game $(E, (A'^i)_{i=1,\cdots,N}, (c_e)_{e \in E}, (k_e)_{e \in E}, (J'^i)_{i=1,\ldots,N})$ with $A' = A'^1 \times \ldots A'^N$ and
$$J'^i(\boldsymbol{\alpha}') = \sum_{e \in \alpha'^i} c_e(k_e(\boldsymbol{\alpha}')),$$
which is a congestion game by construction, is in fact isomorphic to the game we started from. Let $i \in [N]$, $\alpha^i \in A^i$, $\boldsymbol{\alpha}^{-i} \in A^{-i}$ and let us consider the strategy profile $\boldsymbol{\alpha} = (\alpha^i, \boldsymbol{\alpha}^{-i}) \in A$. We claim that for any resource $e \in E$:
   a) $\{j;\ e \in \phi_j(\alpha^j)\} = \{i\}$ if $e = e(\boldsymbol{\alpha}^{-i})$ for some i;
   b) $|\{j;\ e \in \phi_j(\alpha^j)\}| \geqslant 2$ otherwise.

   Notice that, by construction of $\phi_i(\alpha^i)$, we have $e(\boldsymbol{\alpha}^{-i}) \in \phi_i(\alpha^i)$. So if $j \neq i$:
$$e(\boldsymbol{\alpha}^{-i}) \notin \bigcup_{\boldsymbol{\beta}^{-j} \in A^{-j}, \alpha^i \neq \beta^i} \{e(\beta^{-j}\}$$
and hence $e(\boldsymbol{\alpha}^{-i}) \notin \phi_j(\alpha^j)$. On the other hand:
$$\begin{aligned} e \in E \backslash \Big(\bigcup_i \{e(\boldsymbol{\alpha}^{-i})\} &\Rightarrow e = e(\boldsymbol{\beta}^{-j}) \text{ for some } j \in [N] \text{ and } \boldsymbol{\beta}^{-j} \in A^{-j} \\ &\Rightarrow e \in \phi_j(\alpha^j). \end{aligned}$$
Now since $\boldsymbol{\alpha}^{-j} \neq \boldsymbol{\beta}^{-j}$, there exists $k \in [N] \backslash \{j\}$ with $\alpha^k \neq \beta^k$, implying that $e(\boldsymbol{\beta}^{-j}) \in \phi_k(\alpha^k)$. Consequently, if $\tilde{\boldsymbol{\alpha}} = \big(\phi_1(\alpha^1, \cdots, \phi_N(\alpha^N)\big)$ we must have:
$$\begin{aligned} J'^i(\tilde{\boldsymbol{\alpha}}) &= \sum_{e \in \phi_i(\alpha^i)} c_e\big(k_e(\tilde{\boldsymbol{\alpha}})\big) \\ &\sum_{\boldsymbol{\beta}^{-i} \in A^{-i}} c_{e(\boldsymbol{\beta}^{-i})}\big(k_{e(\boldsymbol{\beta}^{-i})}(\tilde{\boldsymbol{\alpha}})\big) + \sum_{j \neq i} \sum_{\boldsymbol{\beta}^{-j} \in A^{-j}, \beta^i \neq \beta^j} c_{e(\boldsymbol{\beta}^{-i})}\big(k_{e(\boldsymbol{\beta}^{-i})}(\boldsymbol{\alpha})\big) \\ &J^i(\alpha^i, \boldsymbol{\alpha}^{-i}) = J^i(\boldsymbol{\alpha}) \end{aligned}$$
completing the proof of the desired isomorphism.
3. We now start from a potential game. We know from Proposition 1.23 that its cost functions can be written in the form $J^i = \varphi_i^d + \varphi_i^c$ in such a way that
$$([N], (A^i)_{i=1,\cdots,N}, (\varphi_i^c)_{i=1,\ldots,N}) \quad ([N], (A^i)_{i=1,\cdots,N}, (\varphi_i^d)_{i=1,\ldots,N})$$
are a coordination and a dummy game respectively. From the first two steps of the proof, we know that $([N], (A^i)_{i=1,\cdots,N}, (\varphi_i^c)_{i=1,\ldots,N})$ and $([N], (A^i)_{i=1,\cdots,N}, (\varphi_i^d)_{i=1,\ldots,N})$ are isomorphic to congestion games

$$([N], (A'^i)_{i=1,\cdots,N}, (J'^i)_{i=1,\ldots,N}) \quad \text{and} \quad ([N], (A''^i)_{i=1,\cdots,N}, (J"^i)_{i=1,\ldots,N})$$

with resource spaces E' and $E"$, and isomorphisms $\phi'_i : A^i \mapsto A'^i$ and $\phi"_i : A^i \mapsto A"^i$ respectively. Without any loss of generality we can assume that $E' \cap E" = \emptyset$ (otherwise, we can rename some of the resources to get in such a situation). We set $E = E \cup E"$ and to each $e \in E$ we associate the cost associate to it as an element of E' or $E"$. Next for each $i \in [N]$ we define the function $\phi_i : A^i \mapsto 2^E$ by $\phi_i(\alpha^i) = \phi'_i(\alpha^i) \cup \phi"_i(\alpha^i)$ and $A'^i = \{\phi_i(\alpha^i);\ \alpha^i \in A^i\}$, and if $\boldsymbol{\alpha}' = (\alpha'^1 \cup \alpha"^1, \ldots, \alpha'^N \cup \alpha"^N)$ we set:

$$J'(\boldsymbol{\alpha}') = J'^i(\alpha'^1, \ldots, \alpha'^N) + J"^i(\alpha"^1, \ldots, \alpha"^N).$$

Accordingly:

$$\begin{aligned} J'^i\big(\phi_1(\alpha^1), \cdots, \phi_N(\alpha^N))\big) &= J'^i\big(\phi'_1(\alpha^1), \ldots, \phi'_N(\alpha^N)\big) + J"^i\big(\phi"_1(\alpha^1), \ldots, \phi"_N(\alpha^N)\big) \\ &= \phi^d_i(\alpha^1, \ldots, \alpha^N) + \phi^c_i(\alpha^1, \ldots, \alpha^N) \\ &= J^i(\alpha^1, \ldots, \alpha^N) \end{aligned}$$

which completes the proof of the desired isomorphism and the proposition. □

1.4 Appendix: Correspondences and Selection Theorems

This appendix collects without proof a couple of abstract mathematical results which are useful in proving existence and regularity of extrema when dealing with optimization problems.

Definition 1.26 *A correspondence φ between two spaces S and X, denoted by $\varphi : S \twoheadrightarrow X$, is a map $\varphi : S \to 2^X$, where 2^X is the collection of all the subsets of X.*

We said that S is the ***domain*** *of φ, and X is the* ***range space****. The image of a set $A \subset S$ under φ is the set*

$$\varphi(A) = \bigcup_{s \in A} \varphi(s).$$

The ***range*** *of a correspondence φ is the image of S. We can identify φ with its* ***graph*** *$gr(\varphi)$, the subset of $S \times X$ given by*

$$gr(\varphi) = \{(s, x) \in S \times X : x \in \varphi(s)\}.$$

Definition 1.27 *A function ψ is a selection of φ if $\psi : S \to X$ is such that for every $s \in S$, $\psi(s) \in \varphi(s)$.*

Remark 1.28 • *One of the many differences between functions and correspondences is the definition of inverse image. For a correspondence φ, two generalizations of the inverse image of a set $A \subset X$ are the "upper inverse" (also called the* ***strong inverse****), which is defined as $\varphi^u(A) = \{x : \varphi(x) \subset A\}$, and the "lower inverse" (or the* ***weak inverse****) defined by $\varphi^\ell(A) = \{x : \varphi(x) \cap A \neq \emptyset\}$. These two definitions of inverse image for a correspondence give raise to two notions of continuity. A correspondence is said to be "upper hemicontinuous" if the upper inverse image of any open set is open. Similar definition for it to be "lower hemicoutinuous".*

- *Every correspondence $\varphi : S \twoheadrightarrow X$ has an inverse correspondence $\varphi^{-1} : X \twoheadrightarrow S$ defined by*
$$\varphi^{-1}(x) = \{s \in S : x \in \varphi(s)\} = \varphi^{\ell}(\{x\}), \qquad \forall x \in X.$$
The set $\varphi^{-1}(x)$ is also called the lower section of φ at x.
- *A correspondence $\varphi : S \twoheadrightarrow X$ between two topological vector spaces S and X is said to be **closed** if its graph $gr(\varphi)$ is a closed subset of $S \times X$. A correspondence $\varphi : S \twoheadrightarrow X$ is said to be **closed-valued** if $\varphi(s)$ is a closed set for every $s \in S$.*
- *The notion of "upper hemicontinuous" comes in the **Closed Graph theorem** which identifie the property of being closed and being upper hemicontinuous for a closed-valued correspondence with compact Hausdorff range space. Moreover, the compactness is preserved under upper hemicontinuous property, [1, Lemma 17.8].*

Theorem 1.29 (Kakutani-Fan-Glicksberg fixed point theorem) *Let K be a non-empty compact convex subset of a topological Hausdorff vector space. Let $\varphi : K \twoheadrightarrow K$ be a correspondence such that $\varphi(k)$ is non-empty and convex for all $k \in K$, and φ is closed ($gr(\varphi)$ is closed in $K \times K$). Then the set of fixed points $\{k : k \in \varphi(k)\}$ is a non-empty compact set.*

We shall use several times a minimization theorem of a topological nature. It is known under the name of Berge's maximum theorem [1, Theorem 17.31], stated here for a minimum.

Theorem 1.30 (Berge Minimum Selection Theorem) *Let X and Y be topological spaces, let $\varphi : X \twoheadrightarrow Y$ be a correspondence between X and Y with non-empty compact values, and let $f : gr(\varphi) \to \mathbb{R}$ be a continuous function. We define*
$$m(x) = \inf_{y \in \varphi(x)} f(x, y),$$
and
$$\mu(x) = \arg\inf_{y \in \varphi(x)} f(x, y) = \left\{ z \in \varphi(x) : \ f(x, z) = \inf_{y \in \varphi(x)} f(x, y) \right\}.$$
Then,

1. *m is continuous;*
2. *μ has non-empty compact values.*

We conclude this appendix with the statement of a selection theorem which we shall use in the sequel.

Definition 1.31 *A function $f : S \times X \to \mathbb{R}$ is a **Carathéodory function** if*

- *for all $s \in S$, $x \mapsto f(s, x)$ is continuous,*
- *and for all $x \in X$, $s \mapsto f(s, x)$ is measurable.*

*A correspondence $\varphi : S \twoheadrightarrow X$ is **weakly measurable** if $(s, x) \mapsto \delta(s, x) = d(x, \varphi(s))$ is Carathéodory.*

Theorem 1.32 (Minimum Selection Theorem) *Let (S, d) be a measurable space, X a separable metric space, and $\varphi : S \twoheadrightarrow X$ a weakly measurable correspondence with non-empty compact values. Let $f : S \times X \to \mathbb{R}$ be a Carathéodory function. We define the value function $m : S \to \mathbb{R}$ by*

$$m(s) = \inf_{x \in \varphi(s)} f(s, x), \tag{1.22}$$

and the correspondence $\mu : S \twoheadrightarrow X$ of minimizers by

$$\mu(s) = \left\{ x \in \varphi(s); f(s, x) = \inf_{y \in \varphi(s)} f(s, y) \right\}. \tag{1.23}$$

Then

1. *$s \mapsto m(s)$ is $\mathcal{S}$-measurable;*
2. *$s \mapsto \mu(s)$ has non-empty compact values;*
3. *$s \mapsto \mu(s)$ has a measurable selection.*

2

Introduction to Graphs & Network Games

2.1 Minimalist Introduction to Finite Graphs

2.1.1 Basic Notations

An **undirected graph** G consists of a non-empty finite set V of elements called **vertices** (or **nodes**), and a finite set E of unordered pairs of vertices called **edges**. To include the case of directed graphs under the same notation set, we shall often write $E \subset V \times V$ even if technically speaking,$V \times V$ is the set of ordered pairs. An undirected graph with a vertex set V and an edge set E will be denoted $G = (V, E)$. We will often use the labels $1, 2, \dots, N$ to represent the nodes or vertices so that $V = \{1, 2, \dots, N\}$, sometimes denoted $V = [N]$, and accordingly, an edge will be an unordered pair $\{i, j\}$ for $i, j \in V$. We say that node i is "connected" to node j if $\{i, j\} \in E$. Note that this notation can be a source of confusion as the edges are unordered. In other words an edge is a "set" with two elements and it should be denoted $\{i, j\}$ instead of (i, j) since the order does not matter. Also, notice that we shall work only with graphs without loops in the sense that (i, i) will never be an edge.

We say that two vertices i and j of a graph G are **adjacent** if there is an edge joining them, namely $\{i, j\} \in E$, and the vertices i and j are then **incident** with such an edge. The **degree** of a vertex i of G, denoted $d(i)$, is the number of edges incident with i.

A useful tool in the analysis of a finite graph $G = (V, E)$ is its **adjacency matrix** $A^{(G)} = [a^{(G)}_{ij}]_{i,j \in V}$. If a graph G has N vertices, its adjacency matrix $A^{(G)}$ is defined as the $N \times N$ matrix whose $ij - th$ entry indicates if there is an edge between vertex i and vertex j. To be specific:

$$a^{(G)}_{ij} = \begin{cases} 1 & \text{if } \{i, j\} \in E \\ 0 & \text{o.w.} \end{cases} \tag{2.1}$$

Notice that, as long as we are dealing with undirected graphs, the adjacency matrix $A^{(G)}$ is symmetric.

Remark 2.1 *In some models, the entries of the adjacency matrix are real numbers. In these cases, the entry $a^{(G)}_{ij}$ is intended to give the actual strength of the connection between vertices i and j. Unless specified otherwise, we shall assume that the entries $a^{(G)}_{ij}$ of the adjacendy matrix are either* 0 *or* 1.

As usual in linear algebra, the **spectrum** of a matrix A, denoted by $\Sigma(A)$, is the collection of the eigenvalues of A. As we already noticed, by definition of an undirected graph, the adjacency matrix $A^{(G)}$ is symmetric, so its eigenvalues are all real numbers. Thus, the spectrum of the adjacency matrix $A^{(G)}$ associated to the graph G with N vertices is a subset of $\mathbb{R}$. It has exactly N elements if we count the eigenvalues repeating them according to their multiplicity. We denote by $\lambda_{max}(A^{(G)}) = \max\{\Sigma(A^{(G)})\}$ and $\lambda_{min}(A^{(G)}) = \min\{\Sigma(A^{(G)})\}$ the maximum and minimum eigenvalues of $A^{(G)}$.

2.1.2 Measures of Centrality

One fundamental problem in network analysis is to identify important nodes in complex networks. Different measures of importance or *centrality* have been proposed to achieve this objective. We review a few of the centrality measures which we shall use in these lectures.

Typically, measures of centrality try to quantify the importance of a node based on the importance of the nodes it is connected to, including the neighbor nodes, the two-hop neighbors, the three-hop neighbors, ... Notice that, since $a_{ij}^{(G)} = 0$ or 1, we have:

$$\left[\left(A^{(G)}\right)^2\right]_{ij} = \sum_{k=1}^{N} a_{ik}^{(G)} a_{kj}^{(G)}$$

which is equal to the number of paths from vertex i to vertex i with one hop. Similarly

$$\left[\left(A^{(G)}\right)^3\right]_{ij} = \sum_{k,\ell=1}^{N} a_{ik}^{(G)} a_{k\ell}^{(G)} a_{\ell j}^{(G)}$$

which is equal to the number of paths from vertex i to vertex i with two hops. More generally, $\left[(A^{(G)})^\ell\right]_{i,j} \neq 0$ if and only if there exists a path joining vertex i to vertex j with exactly $\ell - 1$ hops. We should keep these simple facts in mind to understand the rationale behind the definitions of the following measures of centrality.

Katz centrality [22]:

Definition 2.2 *Let $G = (V, E)$ be an undirected graph with N nodes and adjacency matrix $A^{(G)}$, and let a be an arbitrary positive parameter strictly smaller than $1/\lambda_{max}(A^{(G)})$. Then the* ***Katz centrality*** *of G is defined as:*

$$k_a(i) = \sum_{k=1}^{\infty} \sum_{j=1}^{N} \left(a^k \left[\left(A^{(G)}\right)^k\right]_{ij} \right), \quad \textit{for all } i \in V. \tag{2.2}$$

k_a can be viewed as a function on the set V of nodes, or equivalently as an N-vector $k_a = [k_a(i)]_{i=1,\cdots,N}$. We shall denote by $\mathbf{1}_N$ the N-vector with entries all equal to 1. Since $\left[\left(A^{(G)}\right)^k\right]_{ij}$ is equal to the i-th entry of the vector $\left[\left(A^{(G)}\right)^k\right]\mathbf{1}_N$, we can rewrite the definition of the Katz centrality as:

$$k_a = \sum_{k=1}^{\infty} a^k \big(A^{(G)}\big)^k \mathbf{1}_N = aA^{(G)}\mathbf{1}_N + a^2\big(A^{(G)}\big)^2\mathbf{1}_N + a^3\big(A^{(G)}\big)^3\mathbf{1}_N + \cdots .$$

Since $0 < a < 1/\lambda_{max}(A^{(G)})$, the matrix $I_N - aA^{(G)}$ is invertible and its value is given by the series expansion:

$$[I_N - aA^{(G)}]^{-1} = I_N + aA^{(G)} + a^2[A^{(G)}]^2 + a^3[A^{(G)}]^3 + \cdots$$

which proves that the Katz centrality measure can be rewritten as:

$$k_a = \big([I_N - aA^{(G)}]^{-1} - I_N\big)\mathbf{1}_N - \mathbf{1}_N$$

and if we use the fact that

$$[I_N - aA^{(G)}]k_a = \mathbf{1}_N - [I_N - aA^{(G)}]\mathbf{1}_N = aA^{(G)}\mathbf{1}_N$$

we get:

$$\boxed{k_a = a\big[I_N - aA^{(G)}\big]^{-1} A^{(G)}\mathbf{1}_N.}$$

We shall often skip the subscript $_N$ from the notations I_N and $\mathbf{1}_N$ when there will be no ambiguity concerning the dimension.

Bonacich centrality and **Degree centrality** [6]:

Definition 2.3 *Let a, b be arbitrary parameters such that $b < 1/\lambda_{max}(A^{(G)})$. The **Bonacich centrality** measure is defined implicitly as the solution of the equation:*

$$\mathbf{b}_{a,b}(i) = \sum_{j=1}^{N} (a + b\,\mathbf{b}_{a,b}(j)) \cdot a_{ij}^{(G)}. \tag{2.3}$$

In matrix notation, if we use the notation $\mathbf{b}_{a,b}$ for the vector $\mathbf{b}_{a,b} = [\mathbf{b}_{a,b}(i)]_{i=1,\cdots,N}$, the equation used in the definition states:

$$\mathbf{b}_{a,b} = aA^{(G)}\mathbf{1} + bA^{(G)}\mathbf{b}_{a,b} \tag{2.4}$$

which can be rewritten as:

$$[I - bA^{(G)}]\mathbf{b}_{a,b} = aA^{(G)}\mathbf{1} \tag{2.5}$$

which implies that, since the assumption implies that $I - bA^{(G)}$ is invertible:

$$\boxed{\mathbf{b}_{a,b} = a\left(I - bA^{(G)}\right)^{-1} A^{(G)}\mathbf{1} \quad (2.6)}$$

The parameter b in the Bonacich centrality measure is a form of radius of influence allowing us to scale up or down the range of nodes effectually influencing the score. The role of the parameter a is not much more than a normalization factor. In fact, it is often chosen so that:

$$\|\mathbf{b}_{a,b}\|_2^2 = \sum_{i=1}^{N} |\mathbf{b}_{a,b}|^2 = 1.$$

Notice that:

- If $a = 1, b = 0$, we then have $\mathbf{b}_{1,0} = A^{(G)}\mathbf{1}$, so that $\mathbf{b}_{1,0}(i) = d(i)$ is the degree of vertex i (namely the number of nodes immediately connected to i). For this reason, $\mathbf{b}_{1,0}$ is often called the **degree centrality**;;
- if $b > 0$, $\mathbf{b}_{a,b}(i)$ measures the degree of connection to well connected nodes;
- if $b < 0$, $\mathbf{b}_{a,b}(i)$ measures the degree of connection to weakly connected (less central) nodes.

Eigenvector centrality [7] [8] This idea behind the eigenvector centrality is to give a high score to a node if it is connected to other nodes with high scores.

Definition 2.4 *If $\lambda > 0$ is fixed, we define the* ***eigenvector centrality*** *measure of a node i as:*

$$v(i) = \frac{1}{\lambda} \sum_{j=1}^{N} a_{ij}^{(G)} v(j) \tag{2.7}$$

In other words, the importance of node i is proportional to the sum of the importances of its neighbors. In matrix notation, the vector v of centralities satisfies:

$$\lambda v = A^{(G)} v \tag{2.8}$$

which shows that λ is an eigenvalue and v is an associated eigenvector for the adjacency matrix $A^{(G)}$. A natural requirement we may want to impose on a measure is to be non-negative. For this reason, we choose $\lambda = \lambda_{max}(A^{(G)})$ the maximal eigenvalue of $A^{(G)}$, because in this case, one can find an eigenvector with non-negative entries. This is essentially the statement of the classical Perron-Frobenius theorem in matrix theory.

Other measures of centrality. There are many other measures of centrality, and they have been put to good use in specific applications. Beyond degree, Katz, Bonacich and eigenvector measures, *betweenness* and *closedness* measures of centrality have also been used with some success. Moreover, there are also well known extensions to directed graphs such as the measure of *prestige influence* and the famous *Google page rank* measure [29] that ranks the importance of web pages in online searches.

2.1.3 Density and Sparsity

Roughly speaking, a graph is said to be dense when the number of its edges is of the same order as the maximal number of edges it could have given the number of vertices. In contrast, a graph is said to be sparse if it only has a small number of edges. These definitions are vague and need some work to become rigorous mathematical definitions.

Definition 2.5 *The density $d(G)$ of a graph $G = (V, E)$ is defined as the ratio*

$$d(G) = \frac{|E|}{\binom{|V|}{2}} = \frac{2|E|}{|V|(|V|-1)}. \tag{2.9}$$

If we wanted to define the notions of density and sparsity for a fixed graph, we could say that a graph G is dense if $d(G) > 1/2$ and that it is sparse if $d(G) < 1/2$. However, we would like to define and use these notions for large graphs as asymptotic properties of the density. Essentially dense graphs are graphs for which the average degree grows like the number of vertices. We shall give a more precise definition later on.

Definition 2.6 *A sequence $(G_n)_{n\geqslant 1}$ of graphs with $\lim_{n\to\infty}|V(G_n)| = \infty$ is said to be sparse if*

$$|E(G_n)| \in O\big(|V(G_n)|\big) \qquad \text{as } n \to \infty.$$

The sequence is said to be dense if

$$|E(G_n)| \in O\big(|V(G_n)|^2\big) \qquad \text{as } n \to \infty.$$

For example, if G_n has n nodes and the degree of each node is a fixed constant, say k, then $|E(G_n)| \sim k|V(G_n)|$ and the sequence is sparse. On the other hand, if the degree of each node is a fixed fraction γ of n, then $|E(G_n)| \sim \gamma|V(G_n)|^2$ and the sequence is dense.

2.2 Random Graphs

We now review simple examples of random graph models which will play an important role in the sequel, providing examples and counter-examples to some of the abstract limiting theory touted for the analysis of large networks.

2.2.1 Erdös-Renyi Graphs

At this early stage of our discussion of games *over graphs*, we only present two types of simple random graph models which are referred to as Erdös-Rényi because these authors introduced and studied them in a series of papers in the sixties. See for example [13].

- Given two integers N and M, consider the set Ω of graphs with N vertices and M edges. The random graph model $G(N, M)$ is obtained by sampling the elements of Ω with the same probability equal to

$$p = \frac{1}{|\Omega|} = \binom{\binom{N}{2}}{M}^{-1}.$$

- Even though this model of random graphs is usually referred to as the Erdös-Rényi graph model, it was independently and simultaneously proposed by Gillbert. See [15]. Given a set of N (fixed, non-random) vertices, assume that edges are present or absent, independently of each other, with the same probability $p \in [0, 1]$. This kind of random graph is denoted by $G(N, p)$. Its set of edges is denoted by $E(G(N, p))$. In the $G(N, p)$ model, the number of edges is random. Its expectation is:

$$\mathbb{E}[E(G(N, p))] = p\binom{N}{2}.$$

Notice also that for a given integer M, the random graph $G(N,p)$ has exactly M edges with probability:

$$\mathbb{P}\big[|E(G(N,p)))| = M\big] = p^M(1-p)^{N(N-1)/2-M}.$$

2.2.2 Stochastic Block Model

We now described the stochastic block model introduced in 1983 by Holland, Laskey and Leinhardt in [18]. The building blocks of the model are:

- an integer N representing the number of vertices $[N] := \{1,\cdots,N\}$;
- an integer $n \in \{1,\cdots,N\}$ representing the number of *communities* (of vertices);
- $\boldsymbol{p} = (p_1,\ldots,p_n)$ a probability vector on $[n] := \{1,\ldots,n\}$ (the prior on the n communities). So the p_k are non-negative real numbers summing up to one, i.e. $\sum p_k = 1$. They represent a prior probability on the communities.
- a N-tuple $(c_i)_{i=1,\cdots,N}$ of $[n]$-valued independent identically distributed (i.i.d. for short) random variables with values in $[n]$ and common probability distribution $\boldsymbol{p}$. In particular for $i \in \{1,\cdots,N\}$, $\mathbb{P}[c_i = k] = p_k$ for all $k \in \{1,\cdots,n\}$. For each sample of the i.i.d. random variables, $c_1,\cdots,c_N$, we define the partition $\mathcal{C} = \{C_1,\cdots,C_n\}$ of the set $[N] = \{1,\cdots,N\}$ of vertices into n disjoint subsets given by $C_k = \{i \in [N];\ c_i = k\}$. The elements of this (random) partition are called *communities*. Notice that Np_k is the expected number of elements in the k-th community, i.e. $\mathbb{E}[|C_k|] = Np_k$.
- a symmetric $n \times n$ matrix $\tilde{\boldsymbol{p}} = [p_{k,\ell}]_{k,\ell=1,\cdots,n}$ of real numbers in the interval $[0,1]$;

With all these elements in hand, the stochastic block model is defined in the following way.

Definition 2.7 *For positive integers N and $n \leqslant N$, a probability vector $\boldsymbol{p}$ of dimension n, and a symmetric $n \times n$ matrix $\tilde{\boldsymbol{p}}$ with entries in $[0,1]$, samples of the stochastic block model, $SBM(N,n,\boldsymbol{p},\tilde{\boldsymbol{p}})$ in notation, are generated in the following way. After each of the N vertices is assigned a community label in $[n] := \{1,\ldots,n\}$ independently of each other and with probability given by the* ***community prior p****, conditional on this prior assignment, all pairs $\{i,j\}$ of vertices are considered independently of each others, and a vertex $i \in C_k$ is connected to a vertex $j \in C_\ell$ with probability $\tilde{p}_{k,\ell}$.*

In other words, a random sample G is drawn under $SBM(N,n,\boldsymbol{p},\tilde{\boldsymbol{p}})$ if $\mathbf{c}$ is a random variable in $[n]^N$ with i.i.d. components $(c_i)_{i=1,\cdots,N}$ with common distribution $\boldsymbol{p}$, and $G = (\{1,\ldots,N\}, E(G))$ is an N-vertex simple graph where vertices i and j are connected with probability $\tilde{p}_{c_i,c_j}$, independently of other pairs of vertices. Even more explicitly:

$$\mathbb{P}[c_i = k] = p_k, \qquad i \in [N] \text{ and } k \in [n],$$

and

$$\mathbb{P}[e_{ij} \in E(G)|\mathbf{c}] = \tilde{p}_{c_i,c_j}, \qquad i,j \in [N].$$

As already explained earlier, the (random) partition by community sets $\mathcal{C} = \{C_1,\ldots,C_n\}$ can be defined for each $k \in [n]$ by $C_k = C_k(\mathbf{c}) := \{i \in [N];\ c_i = k\}$.

Examples.

- If $\tilde{p}_{k,\ell} = p$ for some fixed number $p \in [0,1]$, then the partition in communities is irrelevant and the stochastic block model gives the Erdös-Rényi model $G(N,p)$.
- Let us assume that p and q are given numbers in $[0,1]$, let $n \leqslant N$, and let us define the matrix $\tilde{\boldsymbol{p}}$ by:

$$\tilde{p}_{k,l} = \begin{cases} p & \text{if} \quad k = l \\ q & \text{if} \quad k \neq l \end{cases}$$

  In this case:
  - two vertices in the same community share an edge with probability p;
  - two vertices in different communities share an edge with probability q;

  We recover the so-called **planted partition model** which is said to be *assortative* if $p > q$, and *disassortative* if $p < q$.

The stochastic block model has been the object of very many theoretical and statistical analyses. In its study, a major challenge is to design algorithms which given a sample graph, can detect if the sample comes from a Erdös-Rényi model or a generic stochastic block model. Also, active research is concerned with the determination of the community structure from graph samples.

For the sake of completeness we give the formulas providing the probabilities of the sample graphs and their communities. For example, if $\boldsymbol{x} \in [n]^N$ it is clear that we have:

$$\mathbb{P}[\mathbf{c} = \boldsymbol{x}] = \prod_{i=1}^{N} p_{x_i} = \prod_{k=1}^{n} p_k^{|C_k(\boldsymbol{x})|} \tag{2.10}$$

Moreover, if We define a vector $\mathbf{y} \in \{0,1\}^{\binom{N}{2}}$ such that $y_{ij} = \mathbf{1}_{e_{ij} \in E(G)}$ for any $i, j \in \{1, \ldots, N\}$, then:

$$\mathbb{P}[E(G) = \boldsymbol{y} | \mathbf{c} = \boldsymbol{x}] = \prod_{1 \leqslant i \leqslant j \leqslant N} \tilde{p}_{x_i,x_j}^{y_{ij}} (1 - \tilde{p}_{x_i,x_j})^{1-y_{ij}} \tag{2.11}$$

$$= \prod_{1 \leqslant k \leqslant \ell \leqslant n} \tilde{p}_{k,\ell}^{N_{k,\ell}(\boldsymbol{x},\boldsymbol{y})} (1 - \tilde{p}_{k,\ell})^{N_{k,\ell}^c(\boldsymbol{x},\boldsymbol{y})} \tag{2.12}$$

where:

$$N_{k,\ell}(\boldsymbol{x},\boldsymbol{y}) = \sum_{i<j, x_i=k, x_j=\ell} 1_{\{y_{ij}=1\}} \tag{2.13}$$

$$N_{k,\ell}^c(\boldsymbol{x},\boldsymbol{y}) = \sum_{i<j, x_i=k, x_j=\ell} 1_{\{y_{ij}=0\}} = |C_k(\boldsymbol{x})| \cdot |C_\ell(\boldsymbol{x})| - N_{k,\ell}(\boldsymbol{x},\boldsymbol{y}), \quad k \neq \ell \tag{2.14}$$

$$N_{k,k}^c(\boldsymbol{x},\boldsymbol{y}) = \sum_{i<j, x_i=x_j=k} 1_{\{y_{ij}=0\}} = \binom{|C_k(\boldsymbol{x})|}{2} - N_{k,k}(\boldsymbol{x},\boldsymbol{y}) \tag{2.15}$$

2.3 Network Games

There are several natural ways to bring together graphs and games into a single model. A first natural possibility is to identify the N players/agents to the nodes $V = [N] =$

$\{1, \cdots, N\}$ of a graph, the presence of an edge between two nodes signaling a direct interaction between the corresponding players. A class of network games of a different nature occurs when the graph underpinning the game has a physical significance, like for example in the routing games, the players being independent individuals who cannot be identified to the nodes of the physical graph, and being in much larger numbers than the vertices in most applications. We shall study instances of both classes of network games.

2.3.1 **Graph Interaction Games**

Let us consider a graph $G = (V, E)$ with a vertex set $V = [N]$ representing N different agents, and an edge set $E \subset V \times V$, the adjacency matrix being denoted by $A^{(G)}$.

Next, we consider here an N player strategic game and we identify the players with the nodes of the graph G. For each player i, the set of actions that they are allowed to choose is denoted by A^i, assumed to be a subset of $\mathbb{R}^{k_i}$ for some $k_i > 0$. Next, we introduce the product $A = A^1 \times A^2 \times \cdots \times A^N$ and we call its elements $\boldsymbol{\alpha} = (\alpha^1, \alpha^2, \ldots, \alpha^N)$ with $\alpha^i \in A^{(i)}$ **strategy profiles**.

Looking at the nature of the cost functions J^i of the different players $i \in V$, one can see the first idiosyncrasy of games over a graph. Indeed, while the cost function of player $i \in V$ is still of the form

$$J^i : A \ni \boldsymbol{\alpha} \mapsto J^i(\boldsymbol{\alpha}) = J^i(\alpha^i, \boldsymbol{\alpha}^{-i}) \in \mathbb{R}$$

the dependence on the second argument is usually on $\boldsymbol{\alpha}^{G,-i} = (\alpha_j)_{j \neq i, a_{ij}^{(G)} \neq 0}$, namely the actions of the immediate neighbors j of player i in the graph.

As for the notions of best response and Nash equilibrium, they are identical to those introduced in the previous chapter.

2.3.2 **A First Example of Nash Equilibrium**

For each player $i \in [N]$, let us assume that $A^i = \mathbb{R}$ and consider the cost function:

$$J^i(\alpha^i, \boldsymbol{\alpha}^{-i}) = f_i\big(\alpha^i - \alpha^{i,0} + \delta \sum_{j \neq i} a_{ij}^{(G)} \alpha^j\big) + g_i(\alpha^{-i}) \tag{2.16}$$

where:

- $f_i : \mathbb{R} \mapsto \mathbb{R}$ is even (i.e. $f(-x) = f(x)$),
  - f_i is decreasing on $(-\infty, 0]$,
  - f_i is increasing on $[0, \infty)$;
- δ gives the strength of the interaction between players (increasing with δ);
- $\alpha^{i,0}$ is an (autarkic) action preferred by player i in the absence of any interaction between players (i.e. when $\delta = 0$).
- g_i can be a general function.

By plain differentiation we see that for any $i \in \{1, \ldots, N\}$, if $\boldsymbol{\alpha}^{-i}$ is given, the best response of player i is:

$$br^i(\alpha^{-i}) = \alpha^{i,0} - \delta \cdot \sum_{j \neq i} a_{ij}^{(G)} \cdot \alpha^j \tag{2.17}$$

since the minimum of f is attained at 0. In words, player i compares their autarkic action to the average of the actions of her neighbors, their best response being the difference between the two. A Nash equilibrium being a fixed point of the best response function, $\boldsymbol{\alpha}^* \in A$ is a Nash equilibrium if it satisfies:

$$\alpha^{*i} = \alpha^{i,0} - \delta \cdot \sum_{j \neq i} a_{ij}^{(G)} \cdot \alpha^{*j}, \qquad i = 1, \cdots, N. \tag{2.18}$$

This is a system of N equations with N unknowns. We study the case of unconstrained actions, i.e. when $A^i = \mathbb{R}$ for all $i \in [N]$, later on. However, A^i can be smaller. For example, one has $|A^i| = 2$ in games with binary actions, typically when the action of a player is *to adopt or not to adopt*, like in games modeling new technology adoption, or vaccination, etc. Other types of constraints can be accommodated. For example if the actions α^i can only be non-negative, then $A^i = [0, \infty)$ and the fixed point equation for a Nash equilibrium becomes:

$$\alpha^{*i} = \max\Big(0, \alpha^{i,0} - \delta \cdot \sum_{j \neq i} a_{ij}^{(G)} \cdot \alpha^{*j}\Big), \qquad i = 1, \cdots, N,$$

or when α^i is limited to be in a bounded interval, say $A_i = [0, L]$, in which case one has to solve a system of equations:

$$\alpha^{*i} = \min\Big[L, \max\Big(0, \alpha^{i,0} - \delta \cdot \sum_{j \neq i} a_{ij}^{(G)} \cdot \alpha^{*j}\Big)\Big], \qquad i = 1, \cdots, N.$$

Unconstrained Action Games

We close this subsection with the analysis of the unconstrained case for which $A^i = \mathbb{R}$ for all $i = 1, \cdots, N$.

Using the fact that $a_{ii}^{(G)} = 0$ because we do not want our graph to have simple loops, we see that the system of N equations (2.18) characterizing Nash equilibriums can be rewritten in matrix form as:

$$[\mathbf{I} + \delta A^{(G)}]\boldsymbol{\alpha}^* = \boldsymbol{\alpha}^0$$

from which we deduce:

$$\boldsymbol{\alpha}^* = [\mathbf{I} + \delta A^{(G)}]^{-1}\boldsymbol{\alpha}^0$$

provided the matrix $\mathbf{I} + \delta A^{(G)}$ is invertible, which will be the case if $|\delta|$ is small enough to guarantee the convergence of the Taylor series expansion:

$$[\mathbf{I} + \delta A^{(G)}]^{-1} = \mathbf{I} - \delta A^{(G)} + \delta^2 (A^{(G)})^2 + \cdots + (-1)n\delta^n (A^{(G)})^n + \cdots .$$

Remember that the adjacency matrix $A^{(G)}$ captures the geometry of the network structure of the graph underpinning the interactions, and δ captures the cost impact.

Further, we assume that the game is symmetric by assuming that $\alpha^{i,0} = \alpha^0 \in \mathbb{R}$ for all $i \in \{1, \ldots, N\}$. Since $\boldsymbol{\alpha}^0 = \alpha^0 \mathbf{1}$, we have:

$$
\begin{aligned}
\boldsymbol{\alpha}^* &= \alpha^0[\mathbf{I} + \delta A^{(G)}]^{-1}\mathbf{1} \\
&= \alpha^0 \left(\mathbf{I} - [\mathbf{I} + \delta A^{(G)}]^{-1}[\mathbf{I} + \delta A^{(G)}] + [\mathbf{I} + \delta A^{(G)}]^{-1}\right)\mathbf{1} \\
&= \alpha^0 \left(\mathbf{I} - \delta[\mathbf{I} + \delta A^{(G)}]^{-1}A^{(G)}\right)\mathbf{1} \\
&= \alpha^0(\mathbf{1} - \mathbf{b}_{\delta,-\delta})
\end{aligned} \tag{2.19}
$$

where $\mathbf{b}_{a,b}$ is the Bonacich centrality measure.

So in the symmetric case, the unique Nash equilibrium is entirely determined by the graph structure as given by the Bonacich centrality measure since:

$$
\boxed{\boldsymbol{\alpha}^* = \alpha^0(\mathbf{1} - \mathbf{b}_{\delta,-\delta})}
$$

and that it is independent of the specific forms of the cost functions f_i and g_i.

Remark 2.8 *In equation* (2.17)*,*

- *if $\delta < 0$, the best response br^i of player i is increasing in the actions of their neighbors. The game is called a game with **strategic complements**, as the actions chosen by the different players mutually reinforce one another.*
- *if $\delta > 0$, br^i is decreasing in the actions of their neighbors. The game is called a game with **strategic substitutes**, as the actions chosen by the players mutually offset one another.*

Remark 2.9 *The simple relationship which we identified between equilibrium actions and Bonacich centrality does not hold in general for heterogenous agents when the symmetry assumption is not satisfied. Still, if all the agents change their autarkic actions α^{i0} by the same amount, say $\boldsymbol{\alpha}^0 \mapsto \boldsymbol{\alpha}^0 + \alpha\mathbf{1}$, then the change in the Nash equilibrium in still given by the Bonacich centrality since:*

$$
\boldsymbol{\alpha}^*(\boldsymbol{\alpha}^0 + \alpha\mathbf{1}) - \boldsymbol{\alpha}^*(\boldsymbol{\alpha}^0) = \alpha(\mathbf{1} - \boldsymbol{b}_{\delta,-\delta}).
$$

2.3.3 Network Congestion Games

We revisit the class of congestion games introduced in Subsection 1.2.3 and add a graph structure to the model. Recall that congestion games with a finite set of resources are potential games and have pure Nash equilibria, see Proposition 1.21 and that coordination, dummy and potential games are isomorphic to congestion games, recall Proposition 1.25.

In this subsection, we consider an N player game, the set of players being $\{1, \ldots, N\}$. We also assume the presence of a graph $G = (V, E)$, we view each edge $e \in E$ as a resource, and we define a mapping $c : E \ni e \mapsto c(e) \in \mathbb{R}$ representing the cost for using edge $e \in E$. We also assume that to each player $i \in \{1, \ldots, N\}$ is associated an ordered pair of two vertices $(s_i, t_i) \in V \times V$ and that the set of feasible strategies for player i is the collection of paths of the graph G from vertex s_i to t_i:

$$
A^i = \{\alpha \subset E;\ \alpha \text{ is a path from } s_i \text{ to } t_i\}.
$$

For each strategy profile $\boldsymbol{\alpha} = (\alpha^1, \cdots, \alpha^N) \in A = A_1 \times \cdots \times A_N$ we let

$$e(\boldsymbol{\alpha}) = \bigcup_{i=1}^{N} \alpha^i \tag{2.20}$$

be the set of all the paths used in $\boldsymbol{\alpha}$ by all the players. For every $e \in E$, we define a load function $k_e : A \ni \boldsymbol{\alpha} \mapsto k_e(\boldsymbol{\alpha}) \in \{0, \ldots, N\}$ by:

$$k_e(\boldsymbol{\alpha}) = |\left\{i \in \{1, \cdots, N\};\ e \in \alpha^i\right\}|.$$

It gives the number of players using the edge e in $\boldsymbol{\alpha}$. To each player $i \in \{1, \ldots, N\}$, we associate the cost function $J^i : A \ni \boldsymbol{\alpha} \mapsto J^i(\boldsymbol{\alpha}) \in \mathbb{R}$ given by:

$$J^i(\boldsymbol{\alpha}) = \sum_{e \in \alpha^i} \frac{c(e)}{k_e(\boldsymbol{\alpha})}.$$

The network with these cost functions $\{J^i\}_{i=1\ldots,N}$ is called a **cost sharing network** since the cost of an edge is divided by the number of players using the edge. For each edge $e \in E$, we define the cost function $c_e : \{0, \ldots, N\} \mapsto \mathbb{R}$ defined as

$$c_e(k) = \frac{c(e)}{k}, \qquad \forall k = 1, \ldots, N, \qquad \text{and } c_e(0) = 0,$$

then we can see that the game is indeed a congestion game in the sense introduced in Subsection 1.2.3. We call it a **network congestion game**.

In this spirit, we define the social cost for N players employing a strategy profile $\boldsymbol{\alpha} \in A$ as the sum of individual costs of each player incurred by taking collectively the strategy profile $\boldsymbol{\alpha}$, namely:

$$J(\boldsymbol{\alpha}) = \sum_{i=1}^{N} J^i(\boldsymbol{\alpha}).$$

and using the function $e : A \ni \boldsymbol{\alpha} \to e(\boldsymbol{\alpha}) \in 2^E$ defined by (2.20) we see that:

$$J(\boldsymbol{\alpha}) = \sum_{i=1}^{N} J^i(\boldsymbol{\alpha}) = \sum_{i=1} \sum_{e \in \alpha^i} \frac{c(e)}{k_e(\boldsymbol{\alpha})} = \sum_{e \in e(\boldsymbol{\alpha})} k_e(\boldsymbol{\alpha}) \frac{c(e)}{k_e(\boldsymbol{\alpha})} = \sum_{e \in e(\boldsymbol{\alpha})} c(e)$$

and the proof of Proposition 1.21 shows that this function J is a potential function for the game $(A^1 \times \ldots \times A^N, \{J^i\}_{i=1,\ldots,N})$ defined above.

2.3.4 **The Braess Paradox**

This paradox is usually presented with lighter notation. Still, we chose to introduce it as an instance of the models of the above section, hence the heavier than necessary notation. We consider the graph $G = (V = \{s, t, a, b\}, E = \{e_{sa}, e_{at}, e_{sb}, e_{bt}\})$. One should think of the node s as the start of a trip and t as the terminal location (target) of the trip. To go from s to t, cars have to go through city a or city b. Ignoring the dashed vertical arrow, the diagram below gives a representation of the possible itineraries.

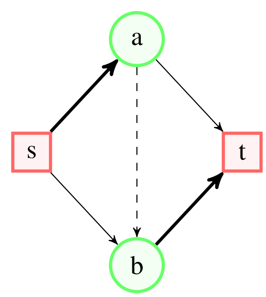

The edges of the graph give the individual legs of the possible trips. We assume that there is an even number N (we shall use $N = 100$ for the purpose of illustration) of car traveling simultaneously from s to t. Each of them has an identical feasible strategy set $A_0 = \{\{e_{sa}, e_{at}\}, \{e_{sb}, e_{bt}\}\}$. For edges $e \in E$, we define the cost functions:

$$c_{e_{sa}}(k) = c_{e_{bt}}(k) = k, \quad c_{e_{at}}(k) = c_{e_{sb}}(k) = c, \qquad \forall$$

for $k \in \{0, \ldots, 100\}$, k being the number of cars traveling through the edge.

In words, the cost incurred for traveling through edges e_{sa} and e_{bt} is proportional to the number of cars on the edge. They are drawn as thick arrows in the above plot. On the other end, the cost incurred for traveling on edges e_{sb} and e_{at} is constant. We should think ot the cost as the travel time to drive from point s to point t. The cost function of traveler $i \in [N]$ is defined as:

$$J^i(\boldsymbol{\alpha}) = \sum_{e \in \alpha^i} c_e(k_e(\boldsymbol{\alpha})) = \begin{cases} c_{e_{sa}}(k_{e_{sa}}(\boldsymbol{\alpha})) + c_{e_{at}}(k_{e_{at}}(\boldsymbol{\alpha})) & if\ \alpha^i = \{e_{sa}, e_{at}\} \\ c_{e_{sb}}(k_{e_{sb}}(\boldsymbol{\alpha})) + c_{e_{bt}}(k_{e_{bt}}(\boldsymbol{\alpha})) & if\ \alpha^i = \{e_{sb}, e_{bt}\}. \end{cases} \tag{2.21}$$

Thus,

$$\begin{aligned} J^i(\boldsymbol{\alpha}) &= c + \text{ the number of cars going through the same route as car } i \\ &= c + \sum_{j=1}^{N} \mathbf{1}_{\alpha^j = \alpha^i} \end{aligned} \tag{2.22}$$

Proposition 2.10 *$\boldsymbol{\alpha}^* \in A$ is a Nash equilibrium if and only if 50% of cars go through path $\{e_{sa}, e_{at}\}$ and the other 50% on path $\{e_{sb}, e_{bt}\}$. More precisely*

$$\boldsymbol{\alpha}^* \in A \text{ is a Nash equilibrium} \quad \Leftrightarrow \quad \left|\{i : \alpha^{*,i} = \{e_{sa}, e_{at}\}\}\right| = \frac{N}{2}$$

Proof: Let $\boldsymbol{\alpha}^* \in A$ such that $\left|\{i : \alpha^{*,i} = \{e_{sa}, e_{at}\}\}\right| = N/2$. Then for any rider $i \in \{1, \ldots, N\}$, assume w.l.o.g. that she goes through path $\alpha^{*,i} = \{e_{sa}, e_{at}\}$. The cost is:

$$J^i(\boldsymbol{\alpha}^*) = c + \sum_{j=1}^{N} \mathbf{1}_{\alpha^{*,j} = \alpha^{*,i}} = c + \frac{N}{2}.$$

If she changes to path $\tilde{\alpha} = \{s_{sb}, e_{bt}\}$, then her cost becomes:

$$J^i(\tilde{\alpha}, \boldsymbol{\alpha}^{*,-i}) = c + \left(\sum_{j \neq i} \mathbf{1}_{\alpha^{*,j} = \tilde{\alpha}} + 1\right) = c + \frac{N}{2} + 1.$$

This show that $\boldsymbol{\alpha}^*$ is a Nash equilibrium.

Inversely, for any Nash equilibrium $\boldsymbol{\alpha}^* \in A$, we must have for every $i \in \{1, \ldots, N\}$ and every $\tilde{\alpha} \in A_0$,

$$J^i(\boldsymbol{\alpha}^*) \leqslant J^i(\tilde{\alpha}, \boldsymbol{\alpha}^{*,-i}),$$

namely

$$\sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\alpha^{*,i}} \leqslant \sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\tilde{\alpha}}.$$

Since there are only two strategies in A_0, we denote by $\tilde{\alpha}_-$ the alternative strategy in the set A_0 compared to $\tilde{\alpha}$. We must have for every $i \in \{1, \ldots, N\}$

$$2\sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\alpha^{*,i}} \leqslant \sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\tilde{\alpha}} + \sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\tilde{\alpha}_-} = N. \tag{2.23}$$

Thus, there must exists another car $k \in \{1, \ldots, N\}$ such that $\alpha^{*,k} \neq \alpha^{*,i}$, so that we also have

$$\sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}\neq\alpha^{*,i}} = \sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\alpha^{*,k}} \leqslant \frac{N}{2}.$$

On the other hand

$$\sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}=\alpha^{*,i}} + \sum_{j=1}^N \mathbf{1}_{\alpha^{*,j}\neq\alpha^{*,i}} = N$$

so that the inequality (2.23) holds for an arbitrary $i \in \{1, \ldots, N\}$, namely

$$\left|\left\{j : \alpha^{*,j} = \alpha^{*,i}\right\}\right| = \frac{N}{2}.$$

Again, because there is only two strategy in A_0, we have

$$\left|\left\{i : \alpha^{*,i} = \{e_{sa}, e_{at}\}\right\}\right| = \frac{N}{2}.$$

which completes the proof. □

Remark 2.11 *1. The above Nash equilibrium is unique up to a permutation of the travelers within each group.*

2. Notice that if k_1 drivers choose the route $\{e_{sa}, e_{at}\}$ and k_2 drivers choose the route $\{e_{sb}, e_{bt}\}$, then the aggregate travel time (the social cost in the terminology of Definition 1.3) is

$$k_1(k_1 + c) + k_2(k_2 + c) = 2k_1^2 - 2Nk_1 + N^2 + cN$$

which is minimum for $k_1 = N/2$, in which case the social cost is $cN + N^2/2$ and the average travel time per driver is $c + N/2$ which is the same cost as the individual cost in the Nash equilibrium identified above. In other words, according to Definition 1.4, the Price of Anarchy is one, i.e. $PoA = 1$.

Now we get to the actual paradox by adding a new route (edge) between vertices a and b, denoted by e_{ab}, and we assign absolutely no cost to the use of this edge, i.e. $c_{e_{ab}}(k) = 0$ for every $k \in \mathbb{N}$. In terms of travel time, one could imagine that Elon Musk's Boring Company finally delivered by constructing a tunnel from a to b in which cars could instantly *beamed* from a to b.

Back from the future, in what follow we assume $c > N$. For the sake of definiteness we choose $c = N + 1$. In any case, the new feasible strategy set, still denoted by A_0, is now:

$$A_0 = \{\{e_{sa}, e_{at}\}, \{e_{sb}, e_{bt}\}, \{e_{sa}, e_{ab}, e_{bt}\}\}.$$

In this new set up, if we choose a driver $i \in [N]$ and define the integers k_1, k_2 and k_3 by:

$$k_1 = |\{j \neq i;\ \alpha^j = \{e_{sa}, e_{at}\}\}, \qquad k_2 = |\{j \neq i;\ \alpha^j = \{e_{sb}, e_{bt}\}\}$$

$$k_3 = |\{j \neq i;\ \alpha^j = \{e_{sa}, e_{ab}, e_{bt}\}\}$$

so that $k_1 + k_2 + k_3 = N - 1$, the travel time (cost) of driver i is

- $1 + k_1 + k_3 + c = N + c - k_2$ if $\alpha^i = \{sa, at\}$;
- $c + 1 + k_2 + k_3 = N + c - k_1$ if $\alpha^i = \{sa, at\}$;
- $1 + k_1 + k_3 + 1 + k_2 + k_3 = N + 1 + k_3$ if $\alpha^i = \{sa, ab, bt\}$.

From this, we easily deduce that the unique (uniqueness taken in the same sense as before) Nash equilibrium $\alpha^* \in A$ is when all the cars travel on $\{e_{sa}, e_{ab}, e_{bt}\}$, in which case the individual travel time is $2N$ which is strictly greater than $c + N/2$ as long as $c < 3N/2$. A first facet of the paradox is that even though the addition of an extra leg with no travel time could be viewed as an improvement, the drivers in a Nash equilibrium experience a significantly worse outcome: *everybody is worse off* in the present situation. This example is an illustration that in some cases, Nash equilibria lead to unnatural and undesirable outcomes.

The social cost is minimized in the same way as before, the drivers splitting themselves between the two itineraries as if the extra link had not been added, so the minimal per driver travel time (i.e. social cost) is $c + N/2$. As a result, the Price of Anarchy is given by:

$$\text{PoA} = \frac{2N}{c + N/2} = \frac{1}{4} + \frac{c}{2N}$$

which could be arbitrary close to $4/3$ if $c \in (N, 3N/2)$ is close to N. As argued before, PoA says how worse Nash equilibria can be when compared to a social optimum.

3

Games with Incomplete Information and Auctions

The purpose of this chapter is to discuss games with incomplete information and study the most frequently used models for auctions. It is important to emphasize that incomplete information *is different from* imperfect information. *A perfect information game assumes that each player knows everything perfectly (the spaces of feasible actions, the cost functions, etc.). This is the situation we considered so far. In an* imperfect information *setting, the players still have information about the components of the game, but this information is not complete. For instance, the information a player receives or acquires about the other players' actions or cost functions may be noisy. However, in the* incomplete information *setting, a single player does not see what the other players are doing.*
In this chapter, we concentrate on problems with incomplete information. In particular, we study in some detail the standard examples of auctions. We follow the approach initiated by J. Harsanyi [16] whose fundamental contribution was rewarded with the Nobel Prize in Economics which he shared with J. Nash and R. Selten in 1994.

3.1 BAYESIAN GAMES

We shall distinguish between information that is known *a priori* and information that is acquired *a posteriori*. These are two elements of the model which could be stochastic. There may be more. When compared with the game models considered so-far, Bayesian games bring to the table a set of important new elements. First, nature can affect the costs incurred by the players, and consequently their actions. The states of nature (or the world) will be captured by the elements of a set Ω. Players do react to the state of the world differently according to their types. So each player can be of a certain type in a set of admissible types. The distribution of the states of the world and the types of the players are weighted by prior probabilities. At the beginning of the game, a state of the world and a profile of N types, one for each player, are drawn according to the prior probability. Before making a decision and taking action, each player is informed of their own type, but not of the types of the other players and the state of the world. Players choose actions *simultaneously*, conditioned on the knowledge of their respective types. Then, they pay their individual costs.

Definition 3.1 *An N - player Bayesian game, denoted by $\mathcal{G}(\Omega, \Theta, A, J)$, consists of the following elements:*

- *a set Ω comprising the* states of the world*;*
- *for each player $i \in [N]$, a set Θ^i of* types *which we shall denote θ^i. We use the notation $\Theta = \Theta^1 \times \cdots \times \Theta^N$ for the product of the spaces of types of the individual players;*

- *a* prior probability *on* $\Omega \times \Theta$, *denoted by* ρ;
- *for each player* $i \in [N]$, *a set* A^i *of feasible actions. As usual, we denote by* $A = A^1 \times \cdots \times A^N$ *the set of admissible strategy profiles;*
- *for each player* $i \in [N]$, *a cost function* $J^i : \Omega \times \Theta \times A \to \mathbb{R}$. *The collection of all cost functions* $(J^i)_{i=1,\ldots,N}$ *is denoted by* J.

We shall denote by $\mathcal{P}(\Omega \times \Theta)$ the set of probability measures on the product $\Omega \times \Theta$. There is no ambiguity in this statement as long as both Ω and Θ are finite. If this is not the case, we need to specify σ-fields of subsets of Ω and Θ for this set of probability measures to be well defined. We shall disregard this issue in this chapter to avoid measure theoretic complications. Even though many if not all the results of this chapter are true in greater generality, we may want to think of the sets Ω and Θ as finite.

Definition 3.2 *For every player* $i \in [N]$, *a pure strategy* $\underline{\alpha}^i$, *is a function from* Θ^i *into* A^i. *The set of pure strategy for player* i *is denoted by* $\mathbb{A}^i$. *A pure strategy profile, say* $\underline{\boldsymbol{\alpha}}$, *is a collection of* N *pure strategies, one for each player, namely*

$$\underline{\boldsymbol{\alpha}} = (\underline{\alpha}^1, \ldots, \underline{\alpha}^N) \in \mathbb{A}^1 \times \ldots \times \mathbb{A}^N.$$

The set of pure strategy profiles is denote by $\mathbb{A} = \mathbb{A}^1 \times \ldots \times \mathbb{A}^N$. *The Bayesian cost incurred by player* i *is the real valued function* $\hat{J}^i : \mathbb{A} \to \mathbb{R}$ *defined by:*

$$\hat{J}^i(\underline{\boldsymbol{\alpha}}) = \int_\Omega \int_{\Theta^{-i}} J^i(\omega, \boldsymbol{\theta}, \underline{\boldsymbol{\alpha}}(\boldsymbol{\theta}))\rho(d\omega, d\boldsymbol{\theta}^{-i}|\theta_i)$$

where $\boldsymbol{\theta} = (\theta^1, \ldots, \theta^N)$ *and* $\underline{\boldsymbol{\alpha}}(\boldsymbol{\theta}) = (\underline{\alpha}^1(\theta^1), \ldots, \underline{\alpha}^N(\theta^N))$.

Remark 3.3 *In the above definition of the Bayesian cost incurred by player* i, *the integration is with respect to the conditional prior* $\rho(dw, d\boldsymbol{\theta}^{-i}|\theta^i)$ *to emphasize the fact that each player is informed of their own type* θ^i *before making a decision.*

Definition 3.4 *A pure strategy profile* $\underline{\boldsymbol{\alpha}}^* \in \mathbb{A}$ *is said to be a Bayesian Nash equilibrium (BNE for short) of the Bayesian game* $\mathcal{G}(\Omega, \Theta, A, J)$ *if for every* $i \in \{1, \ldots, N\}$ *and every* $\underline{\alpha}_i \in \mathbb{A}_i$, *we have*

$$\hat{J}^i(\boldsymbol{\alpha}^*) \leqslant \hat{J}^i(\underline{\alpha}^i, \underline{\boldsymbol{\alpha}}^{*-i})$$

where $(\underline{\alpha}^i, \underline{\boldsymbol{\alpha}}^{*-i}) = (\underline{\alpha}^{1*}, \ldots, \underline{\alpha}^{i-1*}, \underline{\alpha}^i, \underline{\alpha}^{i+1*}, \ldots, \underline{\alpha}^{N*})$.

Remark 3.5 *1. Notice that the definition of a Bayesian Nash equilibrium uses the Bayesian costs* $\hat{J}^i$, *namely the costs integrated with respect to the conditional prior probability. In other words, this is a Nash equilibrium for the* expanded game *for which the feasible actions are pure strategies in* $\mathbb{A}^i$ *instead of being mere actions in* A^i.

2. Since a pure strategy is a map α^i *from* Θ^i *into* A^i, *one could define the notion of* Bayesian Nash equilibrium in mixed strategies *by considering mixed strategies as maps from* Θ^i *into the space* $\mathcal{P}(A^i)$ *of probability measures on* A^i. *We shall use mixed strategies extensively in Chapter 5 and Chapter 6*

3. As defined, Bayesian games are not really games of incomplete information because the sets A, Θ, *the cost functions* J *and the prior are* common knowledge *to all the players. Only the information about the realizations of the types are not shared. So these games could be viewed as games of imperfect information.*

3.2 AUCTIONS

Let us now turn to auctions in which N agents who bid to buy an object. We consider here a single-stage sealed-bid auctions in which players propose their bids individually and simultaneously (sealed-bids). Depending on the bids, only one player, called the winner, will get the object and this winner will have to face a certain cost to purchase the object.

Using the set-up and the definitions introduced above, we model this single-stage sealed-bid auction as an N - player Bayesian game $\mathcal{G}(\Theta, A, J, \rho)$. More precisely,

- there is no component Ω describing the state of the world;
- the type of player i is their valuation of the object, denoted by v_i. We assume that there is a maximum value $\bar{v}$ so that $\theta^i \in [0, \bar{v}]$ for each i;
- the action of player i is their bid for the object, denoted by α_i. The set of feasible actions for player i is all positive bid, i.e. $A^i = [0, \infty)$.
- We assume that, under the prior probability distribution on the types, all types (valuations in the present situation) are i.i.d. with distribution ρ_0 on $[0, \infty)$.

Remark 3.6 *Notice that the prior probability which we denoted by ρ in the previous section on our introduction to Bayesian games is here the product probability measure $\rho_0 \times \cdots \times \rho_0$ of N copies of the probability measure ρ_0 mentioned in the last bullet point above. This common distribution ρ_0 on $[0, \infty)$ reflects a form of symmetry among all the actors involved in the auction. The auction models considered in this chapter can be viewed as games with incomplete information because individual players do not know the valuations (types) of the other players.*

We now describe the auction mechanism, the choice of the winner, and the cost they are required to face to acquire the object. While there are many different types of auctions, those of interest to us in these lectures have the following in common.

- As we already mentioned, nature is not present in auction models and the first step is the revelation of the types: each player *receives* the information about their valuation, and nothing else.
- In the second step, players make their moves by simultaneously placing sealed bids for the object.
- In the third step, the winner is chosen among the players who placed the highest bid, and the winner pays for the object.

So we are only considering auctions for which the object goes to the highest bidder. In case of a tie (several players proposed the same highest bid), the winner is chosen at random among the highest bidders, with equal probabilities. As for the price the winner has to pay, we shall consider two specific auction mechanisms which we shall study in detail in this chapter.

- In a **the first price auction** the winner pays the highest bid, which is their own bid since they won the auction.
- In a **second price auction** also called a **Vickrey auction**, the winner pays the second highest bid.

In any case, the winner needs to compare their payment for the object to their a priori valuation to decide whether or not the operation results in a net gain or net loss. We formalize this point below in the mathematical definitions of these two types of auction.

There are many other types of auction beyond the above sealed-bid auction models. In contrast, they are open format. Among the most popular, the **English auction** and the **Dutch auction** stand out. The English auction is a multi-stage ascending auction. The auctioneer starts from a very low price, and at each round, buyers accept the proposed price and remain in the game, and go into the next round, or exit the game if the price is regarded as already too high. The auctioneer increases the proposed price progressively at each stage, and the last player remaining in the game wins. In contrast, the Dutch auction is a multi-stage descending auction. It starts from a very high price and the proposed price is lowered progressively until one of the players accepts the price. The first to do that wins the auction and has to pay this specific price.

We shall not investigate these two types of auctions because in some rigorous sense (which will not be elaborated in these lectures)

- Descending (Dutch) auctions are equivalent to First Price auctions
- Ascending (English) auctions are equivalent to Second Price auctions

Before we define mathematically the first and second price auctions as Bayesian games, and proceed to the identification of large classes of equilibria, we introduce a class of models for which the analysis can be pushed further.

3.2.1 Symmetric Auction Models

Definition 3.7 *An N-player sealed-bid auction is said to be symmetric if:*

- *all players share a common type space $\Theta_0 = [0, \overline{v}]$, so that now, $\Theta = (\Theta_0)^N$;*
- *all players share a common feasible action space $A_0 = [0, \infty)$, and the set of all admissible strategy profile is $A := (A_0)^N$;*
- *all players use the same pure strategy, namely for any $\underline{\boldsymbol{\alpha}} = (\underline{\alpha}^1, \ldots, \underline{\alpha}^N) \in \mathbb{A}$, there exists a pure strategy $\underline{\alpha} : \Theta_0 \to A_0$ such that $\underline{\alpha}^i = \underline{\alpha}$ for all $i \in \{1, \ldots, N\}$. The set of all admissible common pure strategies is denoted by $\mathbb{A}_0$.*

It is natural to assume that pure strategies in $\mathbb{A}_0$ are non-decreasing. Indeed, it is reasonable to expect that the higher the valuation, the higher the bid should be. Also, and mostly for the sake of convenience, we shall restrict ourselves to continuous and increasing pure strategies. $\underline{\alpha} \in \mathbb{A}_0$ is continuous if it is continuous on the open interval $(0, \bar{v})$, and left and right continuous at 0 and $\overline{v}$ repectiveley, i.e.$\lim_{\theta \searrow 0} \underline{\alpha}(\theta) = \underline{\alpha}(0)$ and $\lim_{\theta \nwarrow \bar{v}} \underline{\alpha}(\theta) = \underline{\alpha}(\bar{v})$.

We denote by $F : [0, \overline{v}] \to [0, 1]$ the cumulative distribution function (c.d.f. for short) of the individual prior probability $\rho_0 \in \mathcal{P}([0, \overline{v}])$. Let us assume that player i wins the auction. There would be no loss of generality assuming that i is a specific player, say 1, because of the symmetry assumption. If all admissible pure strategies are continuous and increasing, then for any $\underline{\alpha} \in \mathbb{A}_0$, if we define $\underline{\boldsymbol{\alpha}}$ by $\underline{\boldsymbol{\alpha}}(\boldsymbol{\theta}) = (\underline{\alpha}(\theta^1), \cdots, \underline{\alpha}(\theta_N)) \in A$ for $\boldsymbol{\theta} \in (\Theta_0)^N$, we have:

$$\underline{\alpha}(\theta_i) > \max_{j \neq i} \underline{\alpha}(\theta^j) \iff \theta_i > \max_{j \neq i} \theta^j.$$

Let us denote by Z the highest type of all other players ($j \neq i$)

$$Z = \max_{j \neq i} \theta^j \ ,$$

and by $G : \Theta_0 \to [0,1]$ its c.d.f. which, because of the symmetry assumption does not depend upon the particular value of i. Since the types are i.i.d. for the prior probability, we have:

$$G(z) = \mathbb{P}[Z \leqslant z] = \mathbb{P}[\theta^j \leqslant z,\ j \neq i] = F(z)^{N-1}.$$

Accordingly, the conditional expectation of Z given the event $\{Z < \theta\}$ for some $\theta \in [0, \overline{v}]$ is given by:

$$\mathbb{E}[Z|Z < \theta] = \frac{1}{G(\theta)} \int_0^\theta z\ G(dz).$$

When studying auctions, our first concern will be to determine pure strategy profiles forming a Bayesian Nash equilibrium. In doing so, we shall be concerned with the point of view of the bidders. But we shall also take the point of view of the auctioneer and compute the payment he or she can expect by running the auction. To this effect, we introduce the following definition and notation.

Definition 3.8 *Given a type profile $\boldsymbol{\theta} \in \Theta$, we denote by $p_i(\boldsymbol{\theta})$ the payment from player i to the auctioneer. Accordingly, we denote by $\mathbb{E}[p_i(\boldsymbol{\theta})|\theta_i]$. the expected payment of player i conditioned on the knowledge of their valuation θ_i.*

3.2.2 Equilibrium Analysis of Vickrey Auctions

We first give a precise mathematical definition of this type of auction.

Definition 3.9 *An N-player second price auction is a 4-tuple (Θ, A, J, ρ_0) where:*

- *$\Theta = \Theta^1 \times \cdots \times \Theta^N$ with $\Theta^i = [0, \overline{v}]$ for all $i \in [N]$;*
- *$\rho_0 \in \mathcal{P}([0, \overline{v}])$ is the common individual valuation prior probability;*
- *$A = A^1 \times \cdots \times A^N$ with $A^i = [0, \infty)$ for all $i \in [N]$;*
- *$J = (J^i)_{i=1,\cdots,N}$ where for each $i \in [N]$, $J^i : \Theta^i \times A \ni (\theta_i, \boldsymbol{\alpha}) \mapsto J^i(\theta_i, \boldsymbol{\alpha}) \in \mathbb{R}$ is defined by:*

$$J^i(\theta_i, \boldsymbol{\alpha}) = \begin{cases} 0 & if\ \alpha^i < \alpha^{(1)} \\ \frac{\alpha^{(2)} - \theta_i}{|\{j|\alpha^j = \max_k \alpha^k\}|} & if\ \alpha^i = \alpha^{(1)} \end{cases}$$

where $\alpha^{(1)} = \max_k \alpha^k$ and $\alpha^{(2)} = \max\{\alpha^j :\ \alpha_j < \max_k \alpha^k\}$.

In the above definition of the cost J^i, the first line corresponds to a situation where player i does not win the auction since thier bid α^i is smaller than the highest bid, so their cost is 0. The second line covers two different cases. If the bid α^i of player i is strictly greater than all the other bids, player i wins the auction and pays exactly $\alpha^{(2)}$ to acquire the object. To obtain their real cost, we subtract their valuation θ_i to this payment. The second case covered by the second line is when player i is not the only one to have the highest bid. For the sake of definiteness, we chose a specific convention to *break the ties*. We assume that the winner is chosen via an independent draw, uniformly at random among the $|\{j|\alpha_j =$

$\max_k \alpha^k\}|$ players sharing the highest bid, and that the cost faced by player i is the expected cost over this random draw given that as before, in case of winning, we subtract their valuation to determine her actual cost.

Since we are in the framework of Bayesian games, a pure strategy for player $i \in [N]$ is a function $\underline{\alpha}^i$ from $\Theta^i = [0, \overline{v}]$ intol $A^i = [0, \infty)$. As before, we denote by $\mathbb{A}^i$ the set of pure strategies for player i and by $\mathbb{A} = \mathbb{A}^1 \times \cdots \times \mathbb{A}^N$ the set of pure strategy profiles. Accordingly, the Bayesian costs are defined by the expectation (or multiple integral):

$$\hat{J}^i(\underline{\boldsymbol{\alpha}}) = \int \cdots \int J^i\big(\theta^i, \underline{\boldsymbol{\alpha}}(\boldsymbol{\theta})\big)\rho_0(d\theta^1)\cdots \rho_0(\check{d\theta^i})\cdots \rho_0(d\theta_N) \tag{3.1}$$

where we used the notation $\underline{\boldsymbol{\alpha}}(\boldsymbol{\theta})$ for the strategy profile $(\underline{\alpha}^1(\theta^1), \cdots, \underline{\alpha}^N(\theta^N))$, and where the check ˇ over $\rho_0(d\theta^i)$ indicates that this term $\rho_0(d\theta^i)$ is not present, so the multiple integral is only over the θ^j for $j \neq i$. So except for θ_i which is known to player i, the types θ^j for $j \neq i$ are regarded as independent random variables with common distribution given by the prior probability ρ_0 and the Bayesian cost integrate this randomness by including the above multiple integrals in the cost to player i.

Remember that players try to maximize their individual rewards as quantified by the difference between their valuation of the object and the actual price they have to pay for it, or equivalently minimize the corresponding costs.

In a strategic game, an easy way to identify a Nash equilibrium is to find, when they exist, dominant strategies for all the players. We define these strategies as follows in the setting of the expanded Bayesian games.

Definition 3.10 *For a single player $i \in [N]$, a pure strategy $\underline{\alpha}^{*i} \in \mathbb{A}^i$ is said to be a dominant strategy if for every other pure strategy $\underline{\alpha}^i \in \mathbb{A}^i$, and every $\underline{\boldsymbol{\alpha}}^{-i} \in \mathbb{A}^{-i}$, we have:*

$$\hat{J}^i(\underline{\alpha}^{*i}, \underline{\boldsymbol{\alpha}}^{-i}) \leqslant \hat{J}^i(\underline{\alpha}^i, \underline{\boldsymbol{\alpha}}^{-i}).$$

In words, no matter what other players are doing, the Bayesian cost incurred by player i using the pure strategy $\underline{\alpha}^{*i}$ is no worse than the Bayesian cost when they pick another pure strategy $\underline{\alpha}^i$.

Clearly, the above definition implies the following simple identification of Bayesian Nash equilibria.

Proposition 3.11 *If $\underline{\boldsymbol{\alpha}}^* = (\underline{\alpha}^{*1}, \cdots, \underline{\alpha}^{*N})$ is a pure strategy profile such that for each $i \in [N]$, $\underline{\alpha}^{*i} \in \mathbb{A}^i$ is a dominant strategy for player i, then $\underline{\boldsymbol{\alpha}}^*$ is a Bayesian Nash equilibrium for the game.*

Proof: Indeed, for each $i \in [N]$ and $\underline{\alpha}^i \in \mathbb{A}_i$, we have:

$$\hat{J}^i(\underline{\alpha}^{*i}, \underline{\boldsymbol{\alpha}}^{*-i}) \leqslant \hat{J}^i(\underline{\alpha}^i, \underline{\boldsymbol{\alpha}}^{*-i})$$

by definition of a dominant strategy for player i. This proves the desired claim. □

Theorem 3.12 *In a second price auction $\mathcal{G}(\Theta, A, J, \rho_0)$, for any player $i \in [N]$, the pure strategy $\underline{\alpha}^*_i \in \mathbb{A}^i$ defined by:*

$$\underline{\alpha}^{*i}(\theta^i) = \theta^i, \qquad \forall \theta^i \in \Theta^i$$

is a dominant strategy for player i.

In words, a dominant strategy for player i is to be truthful and bid their own valuation. Note that according to Proposition 3.11, the auction is in a Bayesian Nash equilibrium if all the players use this dominant strategy.

Proof: Let us choose a player $i \in [N]$, $\underline{\alpha}^i \in \mathbb{A}^i$ a pure strategy for this player, $\underline{\boldsymbol{\alpha}}^{-i} \in \mathbb{A}^{-i}$ pure strategies for the other players, $\boldsymbol{\theta} = (\theta^1, \cdots, \theta^N)$ a set of types for all the players, and let us define:

$$\hat{\beta} := \max_{j \neq i} \underline{\hat{\alpha}}^j(\theta^j).$$

We show that it is not profitable for player i to bid $\beta < \theta^i$. Indeed,

- if $\theta^i > \beta > \hat{\beta}$, player i winns the auction with net cost $\hat{\beta} - \theta^i$, the same thing they would get by bidding θ^i.
- if $\theta^i > \beta = \hat{\beta}$, player i may or may not get the object since $k \geqslant 2$ players have the highest bid. The expected (over the randomization needed to choose the winner) cost is $(\hat{\beta} - \theta^i)/k$, and player i is worse off than by bidding θ^i;
- if $\hat{\beta} \geqslant \theta_i > \beta$, neither θ^i nor β are strong enough bids to win the auction and the net cost to player i is 0 for both bids;
- if $\theta^i > \hat{\beta} > \beta$, bidding β is once more worse than bidding θ^i.

So irrespective of what the other players are doing, bidding $\beta < \theta^i$ is worse than bidding θ^i. Similarly, one shows that bids greater than θ^i are not any better. In conclusion, we showed that player i bidding their own valuation θ^i is a dominant strategy. □

The previous result, together with Proposition 3.11 imply the following result.

Corollary 3.13 *The pure strategy profile* $\underline{\boldsymbol{\alpha}}^* = (\underline{\alpha}^{*1}, \ldots, \underline{\alpha}^{*N})$ *satisfying*

$$\underline{\alpha}^{*i}(\theta^i) = \theta^i, \qquad \forall \theta^i \in \Theta^i$$

for every $i \in [N]$ *is a Bayesian Nash equilibrium for the Vickrey auction.*

In other words, *being truthful and bidding our own private valuation* gives a Bayesian Nash equilibrium. In the next subsection, we show that it is essentially the only such Bayesian Nash equilibrium if one restrict ourself to a reasonable class of pure strategies.

Symmetric Equilibria

In this subsection, we restrict ourselves to pure common strategies in $\mathbb{A}_0$ which are non-decreasing and continuous. We shall also assume that the prior of $\rho_0 \in \mathcal{P}([0, \overline{v}])$ has full support in the sense that there is no open interval $(a, b) \subseteq [0, \overline{v}]$ such that $\rho((a, b)) = 0$.

Under these condition the truthful equilibrium identified earlier is the only Bayesian Nash equilibria in this family.

Proposition 3.14 *In a symmetric Vickrey auction, if the prior* ρ_0 *is continuous and has full support, the pure strategy profile* $\underline{\boldsymbol{\alpha}}^* \in \mathbb{A}$ *defined by*

$$\underline{\alpha}^{*i}(\theta) = \underline{\alpha}^*(\theta) = \theta, \quad \forall \theta \in [0, \infty),\ i \in [N],$$

is the unique symmetric continuous increasing Bayesian Nash equilibrium.

Proof: Let $\hat{\beta} : [0, \bar{v}] \to [0, \infty)$ be the common pure strategy mapping of a symmetric continuous increasing Bayesian Nash equilibrium. We show that $\hat{\beta}(\theta) = \underline{\alpha}^*(\theta) = \theta$ for all $\theta \in [0, \infty)$.

While the assumption of continuity of ρ_0 is not needed for the result to be true, we added it to make our life easier. Indeed, under this continuity assumption, since the valuations (types) are drawn independently according to the prior distribution ρ_0, no two valuations will be equal and we do not have to worry about possible ties between the valuations, hence between the bids since we limit ourselves to symmetric auctions given by increasing pure strategy functions.

$\diamond$ Step 1: we show that for every $\theta \in [0, \bar{v}]$, $\hat{\beta}(\theta) \leqslant \theta$. Otherwise, there must exist a $\theta \in [0, \bar{v}]$ such that $\hat{\beta}(\theta) > \theta$, and by continuity of $\hat{\beta}$, there exists $t \in (0, \theta)$ such that $\hat{\beta}(t) > \theta$. Choose a player $i \in [N]$. By symmetry, it does not matter which one.Then the full support assumption for ρ_0 and the independence of the individual prior probabilities imply that:

$$\mathbb{P}[\forall j \neq i, \theta^j \in (t, \theta)] > 0,$$

so if the type θ^i of player i were to be θ, player i would win the auction with positive probability because $\hat{\beta}(\theta^i) = \hat{\beta}(\theta) > \hat{\beta}(\theta^j)$ for all $j \neq i$ since $\hat{\beta}$ is increasing, and her net cost would be smaller than if she had bid θ since $\hat{\beta}(\theta) > \theta$. This would contradict the fact that we proved earlier that bidding θ is a dominant strategy.

$\diamond$ Step 2: we show that for every $\theta \in [0, \bar{v}]$, $\hat{\beta}(\theta) \geqslant \theta$. Otherwise, there exists a $\theta \in (0, \bar{v})$ such that $\hat{\beta}(\theta) < \theta$. If that is the case, again by continuity of $\hat{\beta}$, there exists $t \in (\theta, \bar{v})$ such that $\hat{\beta}(t) < \theta$. As before, we fix a player $i \in [N]$ and the full support assumption on ρ_0 implies that:

$$\mathbb{P}[\forall j \neq i, \theta^j \in (0, t)] > 0,$$

so if the type θ^i of player i were to be θ, player i would not win the auction with positive probability because $\hat{\beta}(\theta^i) = \hat{\beta}(\theta) < \hat{\beta}(\theta^j)$ for all $j \neq i$ since $\hat{\beta}$ is increasing. Now if all the other players $j \neq i$ keep the same bids $\hat{\beta}(\theta^j)$, by bidding θ instead of $\hat{\beta}(\theta)$, player i would be better off. Indeed, since for $j \neq i$, $\theta^j < t$, we have $\hat{\beta}(\theta^j) < \hat{\beta}(t) < \theta$ since $\hat{\beta}$ is increasing, so player i would win the auction by bidding θ and have a smaller cost than by bidding $\hat{\beta}(\theta)$ contradicting the fact that the symmetric pure strategy profile given by the individual strategy function $\hat{\beta}$ is a Bayesian Nash equilibrium. □

We conclude this section with the computation of the expected payment to the auctioneer.

Proposition 3.15 *In a sealed-bid symmetric second price auction Bayesian Nash equilibrium, the expected payment to the auctioneer is given by:*

$$\mathbb{E}[p_i(\boldsymbol{\theta})|\theta^i = \theta] = G(\theta)\mathbb{E}[Z|Z < \theta] \tag{3.2}$$

which does not depend upon i because of the symmetry assumption.

Proof: The expected payment to the auctioneer by the winner, say i, equals the probability that i wins the auction times the conditional expectation of the second highest bid. If the auction is symmetric and with increasing and continuous admissible pure strategies, then

$$\begin{aligned}
\mathbb{E}[p_i(\boldsymbol{\theta})|\theta^i = \theta] &= \mathbb{P}\big[\underline{\alpha}^*(\theta) > \max_{j\neq i} \underline{\alpha}^*(\theta^j)\big]\mathbb{E}\big[\max_{j\neq i} \underline{\alpha}^*(\theta^j)|\max_{j\neq i} \underline{\alpha}^*(\theta^j) < \underline{\alpha}^*(\theta)\big] \\
&= \mathbb{P}[\theta > Z]\ \mathbb{E}[Z|Z < \theta] \\
&= G(\theta)\mathbb{E}[Z|Z < \theta]
\end{aligned}$$

where the second equality is justified by the truthful bidding in second price auction. □

3.2.3 Equilibrium Analysis of First Price Auctions

As in the case of Vickrey auctions, we first give a precise mathematical definition of the model.

Definition 3.16 *An N-player first price auction is a 4-tuple (Θ, A, J, ρ_0) where Θ, A and ρ_0 are as above in the case of second price auctions, namely:*

- $\Theta = \Theta_1 \times \cdots \times \Theta_N$ *with* $\Theta_i = [0, \overline{v}]$ *for all* $i = 1, \cdots, N$*;*
- $\rho_0 \in \mathcal{P}([0, \overline{v}])$ *is the common individual valuation prior probability;*
- $A = A^1 \times \cdots \times A^N$ *with* $A^i = [0, \infty)$ *for all* $i = 1, \cdots, N$*;*

and the individual costs are given by:

$$J^i(\theta^i, \boldsymbol{\alpha}) = \begin{cases} 0 & if\, \alpha^i < \alpha^{(1)} \\ \frac{\alpha^i - \theta_i}{|\{j|\alpha^j = \max_k \alpha^k\}|} & if\, \alpha^i = \alpha^{(1)} \end{cases}$$

The sets of pure strategies $\mathbb{A}^i$ and the set $\mathbb{A}$ of pure strategy profiles are defined in the same way, and the same goes for the Bayesian costs $\hat{J}^i$.

In order to simplify the analysis, we shall restrict ourselves from now on to symmetric auctions with continuous and increasing individual common pure strategy functions. We shall also assume that the common individual prior probability ρ_0 is continuous. This rules out ties and for that reason, the cost function for player i takes the form:

$$J^i(\boldsymbol{\theta}, \boldsymbol{\alpha}) = (\alpha^i - \theta_i)\mathbf{1}_{\alpha^i \geqslant \max_{j \neq i} \alpha^j}.$$

for all $\boldsymbol{\theta} \in (\Theta_0)^N$ and $\boldsymbol{\alpha} \in A$.

Proposition 3.17 *In a symmetric sealed-bid first price auction with a continuous common individual prior probability ρ_0, the pure strategy profile $\underline{\alpha}^* \in \mathbb{A}_0$ defined by:*

$$\underline{\alpha}^*(\theta) = \mathbb{E}[Z | Z < \theta] \tag{3.3}$$

is the unique symmetric continuously differentiable increasing Bayesian Nash equilibrium.

Proof: In our search for Bayesian Nash equilibria, we consider only strategy profiles $\underline{\boldsymbol{\alpha}} = (\underline{\alpha}^1, \cdots, \underline{\alpha}^N)$ where $\underline{\alpha}^i = \underline{\alpha}$ for all $i = 1, \cdots, N$ with $\underline{\alpha} : [0, \overline{v}] \mapsto [0, \infty)$ increasing and continuously differentiable. For notational convenience, we restrict ourselves to cases where player 1 wins the auction. As explained earlier, this does not create any loss of generality in the proof because of the symmetry assumption. Let $\beta : [0, \overline{v}] \mapsto [0, \infty)$ be such an admissible strategy profile. We assume that players $2, 3, \cdots, N$ use strategy β and we investigate when β happens to be the best response of player 1 to these choices of the other players. If she wins with bid $b = \beta(\theta^1)$, then

$$b > \max_{j \neq 1} \beta(\theta^j) = \beta(\max_{j \neq 1} \theta^j)$$

where the second equality comes from the monotonicity of the pure strategy $\underline{\alpha}^*$. So she wins if and only if:

$$max_{j \neq 1} \theta^j < \beta^{-1}(b)$$

since being increasing, β is invertible and its inverse is also monotone increasing and differentiable. β will be the best response of player 1 if it minimizes her expected costs given her valuation θ^1 and the fact that the other players use strategy β. These expected costs are given by:

$$
\begin{aligned}
\hat{J}^1(\underline{\boldsymbol{\beta}}) &= \mathbb{E}[J^1(\boldsymbol{\theta}, \underline{\boldsymbol{\beta}})|\theta^1] \\
&= \mathbb{E}[(\beta(\theta^1) - \theta^1)\mathbf{1}_{\beta(\theta^1) \geqslant \max_{j \neq 1} \beta(\theta^j)}) \\
&= (b - \theta^1) \cdot \mathbb{P}[\beta(\theta^1) \geqslant \max_{j \neq 1} \beta(\theta^j|\theta^1] \\
&= (b - \theta^1) G((\beta^{-1}(b)).
\end{aligned}
$$

β is the best response when this function of b is minimized. The first order condition is obtained by taking the derivative w.r.t. b, the derivative equals to 0 when $\theta^* = \theta^1$ (FOC), namely

$$
G(\beta^{-1}(b)) + (b - \theta^1) \frac{G'(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))} = 0,
$$

and setting $\theta = \theta^1$, $\theta = \beta^{-1}(b)$ or equivalently $\beta(\theta) = b$, the first order condition rewrites:

$$
\begin{aligned}
0 &= G(\theta)\beta'(\theta) + (\beta(\theta) - \theta)G'(\theta) \\
&= G(\theta)\beta'(\theta) + \beta(\theta)G'(\theta) - \theta G'(\theta) \\
&= [\beta G]'(\theta) - \theta G'(\theta)
\end{aligned}
$$

and integrating between 0 and θ we get: If $\underline{\alpha}^*(0) = 0$, then we can have for every $\theta \in \Theta_0$,

$$
-G(\theta)\beta(\theta) + \int_0^\theta zG(z)dz = 0.
$$

and finally:

$$
\beta(\theta) = \frac{1}{G(\theta)} \int_0^\theta zG(z)dz = \mathbb{E}[Z|Z < \theta].
$$

Note that we used $\beta(0) = 0$ which we can assume without any loss of generality because of the symmetry assumption. In conclusion, we showed that any symmetric Bayesian Nash equilibrium in continuously differentiable increasing pure strategy was necessarily given by the conditional expectation (3.3). Conversely, one sees by running the argument backward that this conditional expectation actually gives such an equilibrium. □

Remark 3.18 *In the equilibrium determined above, the expected payment to the auctioneer becomes*

$$
\mathbb{E}[p_1(\boldsymbol{\theta})|\theta^1 = \theta] = \underline{\alpha}^*(\theta)G(\theta)
$$

where the first term corresponds to the amount actually paid by the winner to the auctioneer, and the second term corresponded to the probability that a player wins the auction. This is exactly the same as the expected payment (3.2) *found in the case of the Vickrey auction. This is a particular case of the revenue equivalence theorem which we state without proof next.*

3.2.4 **Revenue Equivalence**

We conclude this chapter with an important result in the mathematical theory of auctions. It is at the origin of many of the developments of the theory of mechanism design. We state it without proof.

Theorem 3.19 *For any sealed-bid auction where the object goes to the highest bidder, if the types (values) are i.i.d. (in other words if the prior of the Bayesian game model is a product of copies of the same distribution), and if the players are risk neutral (i.e. maximize their own cost functions), then any Bayesian Nash equilibrium in pure symmetric and increasing strategies gives the same expected payment to the auctioneer.*

Remark 3.20 *Even though we did not say it explicitly in the statement of the Revenue Equivalence theorem, we need to need to restrict ourselves to strategies satisfying* $\underline{\alpha}(0) = 0$*, namely which bid nothing when the value is* 0*.*

4

Static Games with a Continuum of Players

The purpose of this chapter is to present the elements of the mathematical theory of games with a continuum of players. These games are often called non-atomic games. In a first section, we review the typical framework used in the existing literature, and we emphasize the abstract mathematical nature of this framework. Next, we discuss graphon games. In anticipation of the detailed presentation of the main results, we devote an entire section to the basic facts from graphon theory. We then generalize the models of Bayesian games to the continuum of players set-up and we introduce several specific models of games of co-ordination and information acquisition which rely on the use of noisy signals. The chapter ends with a discussion of the measure theoretical difficulties created by the manipulations of a continuum of random variables, especially when we wish them to be statistically inde-pendent, and we present the theory of Fubini extensions and the so-called exact law of large numbers.

4.1 Non-atomic Games

This section is devoted to the definition of non-atomic games and to the analysis of a spe-cific subclass of models called *anonymous games*.

4.1.1 Generalities

The basic components of a non-atomic game comprise the following elements.

- A set of players: $(I, \mathcal{I}, \lambda)$ is a probability space where I represents the set of players, $\mathcal{I}$ is a σ-algebra of subsets of I viewed as *coalitions*, and λ is a probability measure that quantifies the weights attached to the coalitions;
- A real separable Banach space E (for example $\mathcal{C}_0([0,1])$, $\mathbb{R}^k$, etc). It represents the set of all possible actions / moves of the individual players;
- A function $A : I \to 2^E$ such that for λ - almost every $i \in I$, $A(i)$ is the set of feasible actions for player i.
- A complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where Ω represents the set of possible states of nature, the σ-field $\mathcal{F}$ comprises the events of interest, and $\mathbb{P}$ is a probability distribution.

Here and in the following, we denote by 2^E the set of subsets of E. We define the set of admissible strategy profiles by

$$L_A = \{\boldsymbol{\alpha} \in L^1(I, \mathcal{I}, \lambda; E);\ \ \alpha(i) \in A(i) \quad \text{for } \ \lambda - a.e.\, i \in I\},$$

as the space of (equivalence classes of) λ-integrable functions from I into E which for λ-a.e. $i \in I$, take values in $A(i)$. When the Banach space E is infinite dimensional, we use the notion of *Bochner integral* when we integrate functions with values in the Banach space E. In particular, if $\boldsymbol{\alpha} \in L_A$ is an admissible control, the quantity:

$$\int_I \alpha(i) d\lambda(i)$$

which can be interpreted as an aggregate action, defines an element of E. The reader unwilling to deal with the generality of infinite dimensional Banach spaces can think of $E = \mathbb{R}^k$ and deal usual integrals only.

Definition 4.1 *A preference function is a function $\Pi : I \times \Omega \times L_A \ni (i, \omega, \boldsymbol{\alpha}) \mapsto \Pi(i, \omega, \boldsymbol{\alpha}) \in 2^E$ such that:*

$$\Pi(i, \omega, \boldsymbol{\alpha}) \subset A(i), \quad \textit{for } \lambda - a.e.\, i \in I \textit{ and } \mathbb{P} - a.e.\, \omega \in \Omega.$$

For each $i \in I$, the set $\Pi(i, \omega, \boldsymbol{\alpha})$ represents the set of all actions that are feasible for player i, and preferred to the actions $\alpha(i) \in A(i)$, given that the state of the world is $\omega \in \Omega$ and the other players $j \neq i$ play with actions $(\alpha(j))_{j \neq i}$.

We assume that players first observe the state of nature $\omega \in \Omega$, and then decide to take actions according to a decision rule:

$$\boldsymbol{\alpha} : \Omega \ni \omega \mapsto \boldsymbol{\alpha}(\omega) \in L_A.$$

The introduction of the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ renders the strategy profiles $\boldsymbol{\alpha}$, and hence the actions $\alpha(i)$ of the individual players random as they depend upon the *state of nature* $\omega \in \Omega$. In part because strategy profiles were only unambiguously defined for λ-almost every player, we omit the effect of zero measure subsets of players in the definition of equilibria.

Definition 4.2 *A decision rule $\boldsymbol{\alpha}^*$ is a (random) Cournot-Nash equilibrium (CNE for short) if for $\mathbb{P}$-a.e. $\omega \in \Omega$ and λ−a.e. $i \in I$*

$$\Pi(i, \omega, \boldsymbol{\alpha}^*(\omega)) = \varnothing.$$

Example 4.3 *Like in all the game models considered so far, let us assume that each player i has a cost function $J^i : \Omega \times A(i) \times L_A \ni (\omega, \alpha^i, \boldsymbol{\alpha}) \mapsto J^i(\omega, \alpha^i, \boldsymbol{\alpha}) \in \mathbb{R}$. If we adapt the classical definition of a Nash equilibrium to this context, it is natural to say that $\boldsymbol{\alpha}^*$ is a CNE if for $\mathbb{P}$-a.e. $\omega \in \Omega$ and λ - a.e. $i \in I$, we have for all $\alpha^i \in A(i)$,*

$$J^i(\omega, \boldsymbol{\alpha}^*(\omega)(i), \boldsymbol{\alpha}^*(\omega)) \leqslant J^i(\omega, \alpha^i, \boldsymbol{\alpha}^*(\omega)).$$

This is consistent with above definition if for any decision rule $\boldsymbol{\alpha} : \Omega \to L_A$, we define for λ-a.e. $i \in I$ and $\mathbb{P}$-a.e. $\omega \in \Omega$,

$$\Pi(i, \omega, \boldsymbol{\alpha}) = \left\{ \alpha^i \in A(i);\;\; J^i(\omega, \alpha^i, \boldsymbol{\alpha}) < J^i(\omega, \boldsymbol{\alpha}(\omega)(i), \boldsymbol{\alpha}(\omega)) \right\}.$$

Remark 4.4 *The terminology* non-atomic *is fully justified when hen the measure λ is continuous, i.e. when $\lambda(\{i\}) = 0$ for all $i \in I$. Recall that in measure theory an atom is defined as a set $A \in \mathcal{I}$ such that $\lambda(A) > 0$, and for any $B \subseteq A$ and $B \in \mathcal{I}$, we have $\lambda(B) = 0$ or $\lambda(B) = \lambda(A)$. Notice that then, I cannot be countable so we are dealing with a game model with a* continuum *of players.*

We can try to recast the set-up of games with finitely many players, say N, into the current framework by setting $I = \{1, \ldots, N\}$ and defining the probability measure λ by $\lambda(A) = \frac{1}{N}|A|$ for any $A \subset I$ in which case:

$$\int_I f(x) d\lambda(x) = \frac{1}{N} \sum_{i=1}^{N} f(i). \tag{4.1}$$

However, such a measure λ is definitely not continuous since it has a discrete finite support with N atoms.

For more on (static) non-atomic games the reader is referred to [35, 2, 3, 17, 25, 4].

4.1.2 **Anonymous Games**

We now introduce an important subclass of non-atomic games. It is especially important because it predates and includes what we came to call Mean Field Games. For each admissible profile $\boldsymbol{\alpha} \in L_A$, we denote by $\lambda_{\boldsymbol{\alpha}}$ the push-forward of the measure λ by $\boldsymbol{\alpha}$. This is the measure on E defined by:

$$\lambda_{\boldsymbol{\alpha}}(B) = \lambda(\{i \in I;\ \alpha(i) \in B\}), \qquad B \in \mathcal{B}(E)$$

where $\mathcal{B}(E)$ denotes the Borel σ-field of E.

Remark 4.5 *Technically speaking, the push forward is defined for an everywhere-defined $(\mathcal{I}, \mathcal{B}(E))$-measurable function $\boldsymbol{\alpha}$ from I into E, the condition $\alpha(i) \in A(i)$ for λ-almost every $i \in I$ being only relevant to the game definition. However, it is clear that if $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ are two such functions which are equal λ-almost everywhere in I, then the measures which are pushed forward by $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ are the same. In other words, the measure $\lambda_{\boldsymbol{\alpha}}$ only depends upon the equivalent class of the strategy profile $\boldsymbol{\alpha}$.*

The situation is simple in the case of finite player games. Indeed, in this case, the support of the measure λ is the full set $I = \{1, \cdots, N\}$ and λ is the normalized counting measure, so that $\lambda_{\boldsymbol{\alpha}}(B)$ is the proportion of players whose action is in B, that is $\lambda_{\boldsymbol{\alpha}}$ is the empirical measure of the action profile $\boldsymbol{\alpha}$ since (4.1) implies:

$$\int_A \varphi(a) \lambda_{\boldsymbol{\alpha}}(da) = \int_I \varphi(\alpha(i)) \lambda(di) = \frac{1}{N} \sum_{i=1}^{N} \varphi(\alpha(i)).$$

Definition 4.6 *We say that the game is anonymous if for every player $i \in I$, there exists a function $\mathcal{J}^i : A^i \times \mathcal{P}(E) \mapsto \mathbb{R}$ such that for any given strategy profile $\boldsymbol{\alpha}$ the cost to player i is given by $J^i(\boldsymbol{\alpha}) = \mathcal{J}^i(\alpha(i), \lambda_{\boldsymbol{\alpha}})$.*

In the case of a finitely many player anonymous game, the cost to a player is a function of its own action and the empirical distribution of the actions of all the players. This is what we called a game with mean field interactions.

In the case of anonymous games, the cost function $\mathcal{J}^i$, when evaluated at (α^i, μ) with $\alpha^i \in A(i)$ and $\mu \in \mathcal{P}(E)$, should be interpreted as the cost to player i when their action is α^i and they face the distribution μ of actions chosen by all the players. When λ is non-atomic (continuous), the choice of action by player i does not affect this distribution, so we can as well say that it is the distribution of the actions chosen by all the other players. The following comments can be helpful in understanding the notion of anonymous games.

- Even though we did not include it in the above definition, it is possible to include the state of nature $\omega \in \Omega$ in the above definition which accordingly, extends in an obvious manner to this more general class of game models.
- As stated above, the cost function $\mathcal{J}^i$ needs to be defined on the set $\mathcal{P}(E)$ of ALL the probability measures on E. However, for the purpose of the analysis of the game, this is demanding too much. Indeed, it is enough for the function $\mathcal{J}^i$ to be defined on the set of probability measures which happen to be the push-forward of an admissible strategy profile in L_A.
- While the cost function $\mathcal{J}^i$ in a general non-atomic game usually depends upon the state of nature ω, the action $\alpha^i \in A(i)$ of player i and the entire strategy profile $\boldsymbol{\alpha}$, in an anonymous game, the dependence upon the strategy profile comes in through the pushed forward measure $\lambda_{\boldsymbol{\alpha}}$ only. Since in the case of finite player games the latter is the empirical measure of the actions, we recover what we called game models with *mean field* interactions.

As in the case of the general non-atomic games discussed earlier, we omit the effect of zero measure subsets of players in the definition of equilibria.

Definition 4.7 *An admissible strategy profile $\boldsymbol{\alpha}^* \in L_A$ is said to be a Nash equilibrium if*

$$\alpha^*(i) \in \; arg \inf_{\alpha^i \in A(i)} \mathcal{J}^i(\alpha^i, \lambda_{\boldsymbol{\alpha}^*}) \tag{4.2}$$

for λ-almost every $i \in I$. Accordingly, the distribution $\lambda_{\boldsymbol{\alpha}^*}$ is called a Nash (equilibrium) distribution.

4.1.3 Existence of Nash Equilibria

The proof of general existence results for equilibria is beyond the scope of this first set of lectures. We shall restrict ourselves to the graphon games discussed next.

Remark on Uniqueness. Inspired by the theory of stochastic differential mean field games developed by Lasry and Lions, one can introduce a notion of monotonicity for the cost functions, and under this monotonicity condition, it is possible to prove uniqueness of the Nash equilibrium.

4.2 INTRODUCTION TO GRAPHONS

The mathematical theory of graphons is concerned with the study of kernels on general measure spaces when we view these kernels as limits of random finite graphs [19]. In its most functional analytic form, it focuses on the spectral theory of the operators associated to these kernels, especially when the kernels are symmetric (i.e. the graphs in question are undirected), and the graph limits can be characterized and captured by the results of these analyses. Here, we restrict ourselves to kernels on the special measure space given by the unit interval $[0, 1]$, its Borel σ-field $\mathcal{B}_{[0,1]}$ and the Lebesgue measure on $[0, 1]$. The most important reference on the subject is still Lovász's book [24].

First, we recall some standard results from functional analysis which we shall use freely in our discussion of the abstract theory of graphons.

4.2.1 Functional Analysis Background

A Banach space is a complete normed linear vector space. Here, we shall only consider separable Banach spaces. For such spaces, the Borel σ-field is the smallest σ-field containing the balls. A Hilbert space is a Banach space whose norm is derived from an inner product. We recall some standard properties of operators on Banach spaces, and we refer to [34] or [31] for details, proofs and complements.

- A linear operator $A : B_1 \to B_2$ from a Banach space B_1 into a Banach space B_2 is said to be **bounded** if there exists an $M > 0$ such that for all $f \in B_1$,

  $$\|Af\|_{B_2} \leqslant M\|f\|_{B_1}, \qquad f \in B_1.$$

  So a linear operator is bounded if the image of the unit ball is bounded, namely, it is contained in a ball. The smallest of the constants M satisfying the above inequality is called the norm of the operator. In other words:

  $$\|A\| = \sup\{\|Af\| : f \in H_1, \|f\| \leqslant 1\}.$$

- We denote by $L(B_1, B_2)$ the collection of all bounded linear operators from B_1 into B_2. Clearly, it is a vector space. It is in fact a Banach space for the above norm. Notice that for linear operators, being bounded or being continuous is the same thing.
- If B is a Banach space, a linear operator $T \in L(B, B)$ is said to be *invertible* if there exists $S \in L(B, B)$ such that $ST = I = TS$ where I denotes the identity operator of B, that is $If = f$ for all $f \in B$. In this case, we write $S = T^{-1}$. For linear operators, we use the product notation ST for the composition of mappings, i.e. $STf = S[Tf]$ for $f \in B$.
- T is invertible if and only if $ker(T) = \{f \in B : Tf = 0\} = \{0\}$ and $rg(T) = \{g \in B : \exists\, f \in B, \quad Tf = g\} = B$.
- If B_1 and B_2 are Banach spaces, a linear operator $A : B_1 \to B_2$ is said to be *compact* if the closure in B_2 of the image by A of the unit ball of B_1 is compact in B_2.
- The **spectrum** $\Sigma(A)$ of an operator $A \in L(B, B)$ is the set of all complex numbers λ such that $A - \lambda I$ is not invertible. Thus, $\lambda \in \Sigma(A)$ if and only if at least one of the following two statements is true:

1. The range of $A - \lambda I$ is not all of B (A is not surjective).
2. $A - \lambda I$ is not one-to-one (A is not injective).

If $T - \lambda I$ is not one-to-one, then λ is said to be an **eigenvalue** of A and $ker(T - \lambda I)$ is called the corresponding eigenspace associated to λ.

- If H is a Hilbert space, an operator $A \in L(H, H)$ is said to be *symmetric* if $< Af, g >=< f, Ag >$ for all f and g in H. This is the natural generalization of the notion of symmetric matrix. Not surprisingly, the spectrum of a symmetric operator is a subset of the real line.
- If H is a Hilbert space, an operator $A \in L(H, H)$ is said to be a *Hilbert-Schmidt* operator if there exists an orthonormal basis $\{f_i\}_{i \geqslant 1}$ of H and a set of square summable complex numbers $\{\lambda_i\}_{i \geqslant 1}$ such that $Af_i = \lambda_i f_i$ for all $i \geqslant 1$.In other words, there exists an orthonormal basis of eigenvectors of the operator A, the corresponding eigenvalues being square summable. Hilbert-Schmidt operators are compact.

In what follows, we shall use extensively the Lebesgue spaces on the unit interval, namely the Banach spaces $L^p([0,1])$ and $L^p([0,1]; \mathbb{R}^k)$ for $1 \leqslant p \leqslant \infty$. Except for $p = \infty$, these Banach spaces are separable. For $p < \infty$, $L^p([0,1]; \mathbb{R}^k)$ is the space of (equivalence classes of) measurable functions f on $[0,1]$ with values in $\mathbb{R}^k$ satisfying:

$$\int_{[0,1]} |f(x)|^p \, dx < \infty,$$

the norm $\|f\|_{L^p}$ of such a function f being defined as the power $1/p$ of the above integral. $L^p([0,1])$ is a short notation for the case $k = 1$. $L^\infty([0,1]; \mathbb{R}^k)$ is the space of (equivalence classes of) essentially bounded measurable functions f on $[0,1]$ with values in $\mathbb{R}^k$ the norm being the essential supremum of $|f(x)|$ over $x \in [0,1]$. We shall use the same symbol $|\cdot|$ to denote the modulus of a complex number, the absolute value of a real number or the Euclidean norm in $\mathbb{R}^k$.

The case $p = 2$ is very special because $H = L^2([0,1]; \mathbb{R}^k)$ is a Hilbert space for the inner product $\langle f, g \rangle = \int_{[0,1]} f(x)^\top g(x) dx$.

4.2.2 **Graphon Theory**

Recall that $L^2([0,1])$ is the Hilbert space of square integrable real-valued functions defined on $[0,1]$, and $L^2([0,1]; \mathbb{R}^k)$ is the space of square integrable $\mathbb{R}^k$ - valued functions defined on $[0,1]$. We shall denote them both by L^2 when no confusion is possible. Square integrability is understood with respect to the Lebesgue measure on $[0,1]$ and as usual, we use the term function instead of equivalence class of almost everywhere equal functions for the elements of these Hilbert spaces. The norms $\|\cdot\|_{L^2}$ of these Hilbert spaces were defined earlier.

Recall that for any bounded linear operator A, we use the symbol $\Sigma(A)$ to denote its spectrum. As in previous chapters, the symbol $\mathbf{1}_N \in \mathbb{R}^N$ denotes the vector of all ones, and $\mathbf{1}_{[0,1]}(\cdot)$ the constant function equal to one on $[0,1]$. Here, $\mathbf{I}$ will denote the identify operator, most often of L^2.

Part of the theory reviewed below applies to a general probability space $(I, \mathcal{I}, \lambda_I)$, but for the sake of definiteness, we shall use $I = [0,1]$, $\mathcal{I} = \mathcal{B}(I)$ its Borel σ-field, and λ_I the Lebesgue measure.

Definition 4.8 *A* ***kernel*** *is an integrable measurable function $k : [0,1] \times [0,1] \to \mathbb{R}$. We shall say that k is a* ***graphon*** *if it is symmetric (i.e. $k(x,y) = k(y,x)$ for all x and y in I) and with values in $[0,1]$. The collection of graphons is denoted by $\tilde{\mathcal{W}}$.*

In the context of graphon games which we shall study in the next section, the variable $x \in I$ generalizes the notion of player, and $k(x,y)$ will quantify the strength of the interaction between player x and player y.

Definition 4.9 *To each graphon $w \in \tilde{\mathcal{W}}$, we associate the graphon operator $A^w : L^2 \to L^2$ defined by:*

$$[A^w f](x) = \int_{[0,1]} w(x,y) f(y) dy, \qquad f \in L^2, \quad x \in [0,1]. \tag{4.3}$$

Notice that $[A^w f](x)$ is defined for every $x \in [0,1]$ if w is defined everywhere, and that its value does not change if f is replaced by a function which is equal to f almost everywhere. A^w is a bounded operator on L^2 and by definition, A^w is what is called a kernel operator. In fact, sine w takes values in $[0,1]$, A^w as defined by (4.3) can be viewed as a bounded operator from L^p to L^q of norm at most 1 for every $p, q \in [0,1]$. Since the kernel w is square integrable in the sense that:

$$\int_{[0,1]} \int_{[0,1]} |w(x,y)|^2 \, dxdy \ < \infty,$$

A^w is a Hilbert-Schmidt operator on the Hilbert space L^2. As such, its spectrum $\Sigma(A^w)$ is a countable set of square summable eigenvalues λ_i , the sum of their squares giving the square of the Hilbert-Schmidt norm of the operator A^w:

$$\|A^w\|_{HS}^2 = \sum_{i \geqslant 1} |\lambda_i|^2 < \infty.$$

The function w being symmetric and real valued, the operator A^w is symmetric and as a result, all of its eigenvalues are real. We shall use all these properties of the graphon operator, and in particular, the fact that it is a compact operator.

Centrality Measure

We already hinted at an analogy between graphs and graphons. In this spirit one can define the *Bonacich centrality* of a graphon as:

$$b_\lambda = [I - \lambda A^w]^{-1} \mathbf{1}_{[0,1]} \tag{4.4}$$

in other words, the resolvent operator evaluated at the constant function equal to 1. This definition makes sense for $\lambda \in [0, 1/\rho(A^w))$ where $\rho(A^w)$ denotes the spectral radius of the graphon operator A^w.

4.2.3 Finite Graphs Naturally Associated to Graphons

Graphons should be thought of as finite graph limits. Incidentally, graphons can provide natural procedures to generate sequences of deterministic finite graphs and samples from random graph models as we are about to demonstrate. Also, we shall emphasize the fact that finite graphs should be viewed as *step function graphons*.

Definition 4.10 *We say that a graphon $w : [0,1]^2 \to \mathbb{R}$ is a step function if there is a partition $\mathcal{D} = \{D_1, \ldots, D_n\}$ of $[0,1]$ into measurable sets such that w is constant on every product set $D_k \times D_\ell$. In other words, if there exists a symmetric real matrix $A = (a_{k\ell})_{1 \leqslant k,\ell \leqslant N}$ such that:*

$$w(x, y) = a_{k\ell}, \qquad \forall (x, y) \in D_k \times D_\ell. \tag{4.5}$$

The sets $D_k \times D_\ell$ are called the steps of w.

Example 4.11 *For each integer N we denote by $\mathcal{D}^N = \{D_1, \ldots, D_N\}$ the regular partition of $[0,1]$ with $D_i = [\frac{i-1}{N}, \frac{i}{N})$ for $i = 1, \ldots, N-1$, $D_N = [\frac{N-1}{N}, 1]$.*

Now if $G = (V, E)$ is a graph with $V = [N]$ and $A^{(G)} = (a_{ij}^{(G)})_{1 \leqslant i,j \leqslant N}$ is its adjacency matrix, we use the regular partition $\mathcal{D}^N$ of $[0,1]$ and define the graphon $w^{(G)}$ by $w^{(G)}(x, y) = a_{ij}^{(G)}$ if $(x, y) \in D_i \times D_j$. Note that if all the entries of the matrix $[a_{ij}^{(G)}]_{ij=1,\cdots,N}$ are non-negative and not greater than 1, then $w^{(G)} \in \tilde{\mathcal{W}}_0$. Also, note that these entries do not need to be 0 or 1.

We now explain how to construct random graph models from a graphon.

Definition 4.12 *(Sampling procedure) Given a graphon $w \in \tilde{\mathcal{W}}$ and an integer N, for each realization $\{u_i\}_{i=1,\ldots,N}$ of i.i.d. random variables uniformly distributed on $[0,1]$, we define the* weight matrix *A_w as follows:*

$$[A_w]_{ij} = \mathbf{1}_{i \neq j} w(u_i, u_j). \tag{4.6}$$

Then, starting from A_w, we construct the random graph with set of vertices $[N] = \{1, \cdots, N\}$ by connecting different vertices i and j with probability $[A_w]_{ij} = w(u_i, u_j)$. In other words, given the realizations $\{u_i\}_{i=1,\ldots,N}$, we have:

$$\mathbb{P}[a_{ij}^{(G)} = 1] = [A_w]_{ij} = w(u_i, u_j).$$

Notice that $[A_w]_{ij} \in [0,1]$ while $a_{ij}^{(G)} \in \{0, 1\}$.

- By definition of the above sampling procedure, for any $p \in [0,1]$, the constant graphon $w(x, y) = p$ corresponds to the Erdos-Renyi random graph model $G(N, p)$ with edge probability p.
- If we consider a prior probability vector $p \in [0,1]^n$ with $\sum_k p_k = 1$, any step function graphon $w^{[n]}$ gives raise to a stochastic block model $SBM(N, p, W^{[n]})$. A sample (X, G) with $x_i \in [0,1]$ and $G = ([N], E(G))$ can be characterized by the weight adjacency matrix A_w computed from $w^{[n]}$.

We now go back to the part where we explained how to see a graphon as a random graph model. Given $[N] = \{1, \dots, N\}$, generate $x_1, \dots, x_N$ i.i.d. $U(0, 1)$ and let

$$a_{i,j}^{(w)} = \begin{cases} w(x_i, x_j), & \text{if } i \neq j, \\ 0, & \text{if } i = j. \end{cases}$$

We then define the random variables $a^{i,j} \in \{0, 1\}$ such that $\mathbb{P}[a^{i,j} = 1] = a_{i,j}^{(w)}$.
Let us now examine the sampling procedure if the graphon is a step graphon. So let us assume that there exists a partition $\mathcal{D} = \{D_1, \dots, D_n\}$ of $[0, 1]$, and a symmetric matrix $p = [p_{k,\ell}]_{k,\ell=1,\dots,n}$ of non-negative numbers such that $w(x, y) = p_{k,\ell}$ if $(x, y) \in D_k, \times D_\ell$. Set $\alpha_k = |D_k|$ for the Lebesgue measue of the st D_k. This is the probability that a random variable U_k uniforrnly distributed over $[0, 1]$ belongs D_k. Then the above sampling proceddure for the step graphon w gives rise to the stochastic block model. Indeed, the probability that there is an edge between the vertices i and j is

$$\begin{aligned} \mathbb{P}[a^{i,j} = 1] &= \text{probability there is an edge between } i \text{ and } j \\ &= w(u_i, u_j) \\ &= p_{k\ell} \;\text{ if }\; u_k \in C_k \text{ and } \underline{\ell} \in C_\ell, \\ &= p_{k\ell} \;\text{ with probability } \alpha_k \alpha_\ell \end{aligned}$$

$\alpha_k \alpha_\ell$ being the probability that $\{U_k \in D_k \;\text{ and } U_\ell \in D_\ell\}$ if the uniform random variables $(U_k)_{1 \leqslant k \leqslant n}$ are independent.

4.2.4 **Cut Norms**

For each measurable function w on $I \times I$, we denote its $L^p(I \times I)$ - norm by $\|w_p\|$ for $p \in [1, \infty)$, i.e.

$$\|w\|_p = \Big[\int_{I \times I} |w(x, y)|^p \; dxdy\Big]^{1/p}.$$

Obviously, this norm is of interest when it is finite. The ∞-norm $\| \cdot \|_\infty$ is defined as the essential supremum as usual. We shall use these norms to manipulate the functional analytic properties of a graphon kernel and its associated operator. However, for their role as limits of finite graphs, graphons are better studied using the so-called cut norm.

Definition 4.13 *The cut norm of a graphon $w \in \tilde{\mathcal{W}}$ is denoted by $\|w\|_\square$ and is defined as follows:*

$$\|w\|_\square = \sup_{D_1, D_2} \left| \int_{D_1 \times D_2} w(x, y) dxdy \right| \tag{4.7}$$

where the supremum is taken over the measurable subsets D_1 and D_2 of $I = [0, 1]$.

Remark 4.14 *By a standard monotone class argument it is easy to show thatL*

$$\|w\|_\square = \sup_{0 \leqslant f \leqslant 1, 0 \leqslant g \leqslant 1} \int \int f(x) g(y) w(x, y) dxdy.$$

Despite its name, the cut-norm is not a norm in the usual sense. It is used to define the notion of cut-distance between graphons.

Definition 4.15 *The cut metric between two graphons v and w in $\tilde{\mathcal{W}}$ is defined as:*

$$d_\square(w,v) = \inf_{\pi \in \Pi_{[0,1]}} \|w^\pi - v\|_\square,$$

where $w^\pi(x,y) = w(\pi(x),\pi(y))$ and $\Pi_{[0,1]}$ is the class of bi-measurable, measure preserving bijections of $[0,1]$ onto itself.

Again, $d_\square$ is not a real distance on $\tilde{\mathcal{W}}$. Indeed $d_\square(w,v) = 0$ if $v = w^\pi$ for some $\pi \in \Pi_{[0,1]}$. So we define the equivalence relation $\sim$ on the space $\tilde{\mathcal{W}}$ of graphons by

$$w \sim w \text{ if there exists } \pi \in \Pi_{[0,1]} \text{ such that } v = w^\pi,$$

and on the quotient space, $\mathcal{W} = \tilde{\mathcal{W}}/\sim$, once appropriately defined, $d_\square$ becomes a real distance. This quotient space is in fact complete and compact for the distance $d_\square$. For the sake of completeness we state such a result even though its proof is beyond the scope of these lectures. We refer to [19] and [12] for details.

Theorem 4.16 *The graphon (quotient) space $\mathcal{W}$ is a compact metric space (hence complete) for the metric $d_\square$.*

Properties such as approximation, convergence, etc of graphons are naturally stated and interpreted in terms of the cut norm distance. On the other hand the classical theory of operators offers numerous tools to control the properties of the graphon operators. The following estimate will allow us to use the best of both worlds.

Proposition 4.17 *A graphon $w \in \tilde{\mathcal{W}}$ and its associated graphon operator A^w satisfy, for any $p,q \in [1,\infty]$ and $q' = (1-1/q)^{-1}$,*

$$\|w\|_\square \leqslant \|A^w\|_{L^p,L^q} \leqslant \sqrt{2}(4\|w\|_\square)^{(1-\frac{1}{p})\wedge\frac{1}{q}}$$

Recall that the above operator norms are defined as:

$$\begin{aligned}\|A^w\|_{L^p,L^q} = \sup_{\|f\|_{L^p}\leqslant 1} \|A^w f\|_{L^q} &= \sup_{\|f\|_{L^p}\leqslant 1} \sup_{\|g\|_{L^{q'}}\leqslant 1} \int A^w f(x)\, g(x)dx \\ &= \sup_{\|f\|_{L^p}\leqslant 1} \sup_{\|g\|_{L^{q'}}\leqslant 1} \int\int w(x,y) f(y) g(x) dy dx.\end{aligned} \tag{4.8}$$

The special case is when $p = q = 2$ is particularly simple. It reads:

$$\|w\|_\square \leqslant \|A^w\|_{L^2,L^2} \leqslant \sqrt{8\|w\|_\square}.$$

This proposition helps shifting the burden of proof of smoothness or convergence, from properties of functions to properties of operators. It syas that if one needs to prove convergence of a sequence $(w_n)_{n\geqslant 1}$ toward a graphon w, it can be approached by proving that $d_\square(w,w_n)$ tends to 0, or equivalently that $\|A^w - A^{w_n}\|_{L^p,L^q} = \|A^{w-w_n}\|_{L^p,L^q}$ tends to 0.

Proof: Step 1. We first prove the result in the particular case $p = \infty, q = 1$, namely that:

$$\|w\|_\square \leqslant \|A^w\|_{L^\infty, L^1} \leqslant 2\|w\|_\square.$$

Notice that:

$$\begin{aligned}\|A^w\|_{L^\infty, L^1} &= \sup_{\|f\|_{L^\infty} \leqslant 1} \sup_{\|g\|_{L^\infty} \leqslant 1} \int A^w f(x)g(x)dx \\ &= \sup_{-1 \leqslant f \leqslant 1} \sup_{-1 \leqslant g \leqslant 1} \int\int w(x,y)f(y)g(x)dxdy,\end{aligned}$$

so if we take $f = \mathbf{1}_{D_1}$ and $g = \mathbf{1}_{D_2}$ for D_1 and D_2 Borel subsets of $[0,1]$, then we have:

$$\|A^w\|_{L^\infty, L^1} \geqslant \|w\|_\square.$$

Next, decompose the functions in the unit balls into the difference of two non-negative functions. We introduce functions f_1 and f_2 valued in $[0,1]$ such that $f = f_1 - f_2$, as well as g_1 and g_2 such that $g = g_1 - g_2$. We get:

$$\|A^w\|_{L^\infty, L^1} = \sup_{0 \leqslant f_1, f_2 \leqslant 1} \sup_{0 \leqslant g_1, g_2 \leqslant 1} \int\int w(x,y)[f_1(y) - f_2(y)] \cdot [g_1(x) - g_2(x)]dxdy. \quad (4.9)$$

Since

$$\begin{aligned}&\int\int w(x,y)[f_1(y) - f_2(y)] \cdot [g_1(x) - g_2(x)]dxdy \\ &\qquad = \langle A^w(f_1 - f_2), (g_1 - g_2)\rangle \\ &\qquad = \langle A^w f_1, g_1\rangle + \langle A^w f_2, g_2\rangle - \langle A^w f_2, g_1\rangle - \langle A^w f_1, g_2\rangle \\ &\qquad \leqslant \langle A^w f_1, g_1\rangle + \langle A^w f_2, g_2\rangle,\end{aligned}$$

we must have:

$$\|A^w\|_{L^\infty, L^1} \leqslant \sup_{0 \leqslant f_1, f_2 \leqslant 1} \sup_{0 \leqslant g_1, g_2 \leqslant 1} \langle A^w f_1, g_1\rangle + \langle A^w f_2, g_2\rangle = 2\|w\|_\square.$$

which concludes the proof of the first step.

Step 2. We prove that:

$$\|A^w\|_{L^\infty, L^1} \leqslant \|A^w\|_{L^p, L^q}.$$

To prove this, we notice that:

$$\begin{aligned}\|A^w\|_{L^p, L^q} &= \sup_{\|f\|_{L^p} \leqslant 1} \|A^w f\|_{L^q} \\ &\geqslant \sup_{\|f\|_{L^\infty} \leqslant 1} \|A^w f\|_{L^q} \\ &\geqslant \sup_{\|f\|_{L^\infty} \leqslant 1} \|A^w f\|_{L^1} = \|A^w\|_{L^\infty, L^1}.\end{aligned}$$

Step 3. To conclude the proof of the lemma, we use the Riesz-Thorin interpolation theorem. See for example Theorem 4.32 in [9] Now, we use this theorem in the following way. Let $p_1 = \infty$ and $q_1 = 1$. Set $r = (1 - \frac{1}{p}) \wedge \frac{1}{q}$, so that $1 - r = \frac{1}{p} \vee (1 - \frac{1}{q})$. Define $\bar{p} = p(1-r)$ and $\bar{q}$ such that $1 - \frac{1}{q} = (1-r)(1 - \frac{1}{\bar{q}})$. Then $\bar{p}, \bar{q}$ are in $[1, \infty]$,

$$(\frac{1}{p}, \frac{1}{q}) = r(\frac{1}{p_1}, \frac{1}{q_1}) + (1-r)(\frac{1}{\bar{p}}, \frac{1}{\bar{q}}).$$

Riesz-Thorin's theorem tells us that

$$\|A^w\|_{L^p, L^q} \leqslant \|A^w\|^r_{L^{p_1}, L^{q_1}} \|A^w\|^{(1-r)}_{L^{\bar{p}}, L^{\bar{q}}}$$

and the right hand side is bounded by $2\|w\|^r_\square$. Indeed

$$\|A^w\|_{L^{\bar{p}}, L^{\bar{q}}} \leqslant \|A^w\|_{L^1, L^\infty} = \|w\|_\infty \leqslant 1.$$

This concludes the proof of 4.8. □

Finally, we give a convergence result of importance in practical applications.

Proposition 4.18 *Let $\{w^n\}_{n \geqslant 1}$ be a sequence of graphons such that*

$$\lim_{n \to \infty} d_\square(w^n, w) = 0$$

for some graphon w. Then there exists a sequence $\{\pi_n\}_{n \geqslant 1}$ in the space $\Pi_{[0,1]}$ of permutations of $[0,1]$ such that:

$$\lim_{n \to \infty} \|A^{(w^n)^{\pi_n}} - A^w\| = 0.$$

Proof: For each $\epsilon > 0$ there exists an integer N_ϵ such that if $n \geqslant N_\epsilon$, the $d_\square(w^n, w) < \epsilon$. So for each $n \geqslant N_\epsilon$, there exists $\pi_n \in \Pi_{[0,1]}$ such that $\|(w^n)^{\pi_n} - w\| < \epsilon$. We conclude with the equivalence of the cut norm and the cut distance proven in the above proposition. □

4.3 GRAPHON GAMES

In this section, we study the class of non-atomic games associated to graphons as defined in the previous section. We present the main results of [30], though with different proofs.

In the chapters devoted to finite games, a player was represented by an index in $\{1, \ldots, N\}$. In the beginning of this chapter, we discussed non-atomic games for which players were indexed by the elements of a measured space $(I, \mathcal{I}, \lambda)$. Graphon games use the same general set-up as non-atomic games for the particular case where $I = [0,1]$, $\mathcal{I} = \mathcal{B}_I$ the Borel σ-field of I, $\lambda = \lambda_I$ being the Lebesgue measure dx. Consistent with our discussion of non-atomic games, strategy profiles are now functions $\boldsymbol{\alpha} : I \to \mathbb{R}^k$ with $\boldsymbol{\alpha}(x) \in A(x) \subset \mathbb{R}^{k_x}$ for all $x \in I$, and we assume that $k_x \leqslant k$, in other words that all the A^x are contained in the same Euclidean space $\mathbb{R}^k$. Remember that in the case of games with finitely many players, a strategy profile was a vector $\boldsymbol{\alpha} = (\alpha^1, \ldots, \alpha^N)$, where $\alpha^i \in A(i) \subset \mathbb{R}^{k_i}$. In several network games, player i was interacting with player j with strength given by the entry $a_{i,j}$ of a matrix, possibly the adjacency matrix of the underlying graph. In a graphon game, player x feels the interaction with the ensemble of the other players in the game through an aggregate

$$z(x|\boldsymbol{\alpha}) = \int w(x,y)\alpha(y)dy, \tag{4.10}$$

given by a weighted average of the actions $\alpha(y)$ of all the players as weighted by the values $w(x,y)$ of a graphon $w \in \tilde{\mathcal{W}}$. In order to be admissible, strategy profiles will have to be at least Lebesgue measurable and square integrable so that the above integral makes sense. Since $\alpha(y) \in A(y) \subset \mathbb{R}^{k_y}$, in order for the definition and the manipulations of the above integral to be sound, we assume that all the sets $A(y)$ are contained in the same $\mathbb{R}^k$, which amount to assuming that all the k_y's are the same, say k. As a result, the above integral is understood as an $\mathbb{R}^k$ valued integral defining the aggregate $z(x|\boldsymbol{\alpha})$ as an element of $\mathbb{R}^k$.

Admissible Strategy Profiles

So according to the above discussion, we would like to say that a function $\boldsymbol{\alpha} : I \mapsto \mathbb{R}^k$ is an admissible strategy profile if it is measurable, sqquare integrable and $\boldsymbol{\alpha}(x) \in A(x)$ for every $x \in I$. Notice that we assume that $\boldsymbol{\alpha}$ is a function defined everywhere, not merely an equivalence class of Lebesgue measurable functions. Still, since the integral defining the aggregate $z(x|\boldsymbol{\alpha})$ is with respect to the Lebesgue measure, changing $\boldsymbol{\alpha}$ at one point, or more generally on a set of Lebesgue measure 0, does not change the value of z. So $z(x|\boldsymbol{\alpha}) = z(x|\boldsymbol{\alpha}')$ if $\boldsymbol{\alpha}(y) = \boldsymbol{\alpha}'(y)$ for almost every $y \in I$, and despite our best efforts, since $z(x|\boldsymbol{\alpha})$ only depends upon the equivalence class of the function $\boldsymbol{\alpha}$, like in the case of non-atomic games, the set of admissible strategy profiles could be defined as the set L_A defined in our introductory presentation of non-atomic games, namely the subset of elements $\boldsymbol{\alpha}$ of $L^1([0,1], dx)$ for which $\boldsymbol{\alpha}(x) \in A(x)$ for almost every $x \in [0,1]$. However, because our analysis will strongly rely on the properties of the graphon operator on L^2 - Hilbert spaces, we shall define the set of admissible strategy profiles as:

$$L^2_A = \{\boldsymbol{\alpha} \in L^2([0,1], dx; \mathbb{R}^k);\ \alpha(x) \in A(x) \quad \text{for } \ a.e.\, x \in [0,1]\},$$

In any case, if player x decides to change the value $\alpha(x)$ of their action while all the other players choose not to, the value of the aggregate $z(x|\boldsymbol{\alpha})$ is not affected.

Costs

In the graphon game model, the costs to the individual players are determined by a single real valued function $J : \mathbb{R}^k \times \mathbb{R}^k \mapsto \mathbb{R}$. The interpretation of $J(\alpha, z)$ is the following: this value will be the cost to player $x \in I$ if the player chooses action α (which implicitly requires that $\alpha \in A(x)$) while the actions of all the (other) players amount to the aggregate z through formula (4.10).

The graphon game defined by the quantities introduced above will be denoted by

$$\mathcal{G} = (\mathbf{A} = (A(x))_{x\in[0,1]}, w, J).$$

Remark 4.19 *While very similar in spirit, a graphon game is not an anonymous game as we defined it earlier, the differences stemming from the definitions of the costs. Indeed, even if one assumes that all the cost functions J^i in the anonymous game model are the same, say J, the cost to a given player has two important characteristics.*

1. *The significance (or insignificance) of the action of a player on the costs depends upon the measure λ. In particular, it is nil if this measure is continuous or atomless.*
2. *It depends upon the actions of the other players in a very specific way entirely determined by the measure λ since it is a function of the push-forward of this measure by the players' actions.*

On the other hand, the parallels to these properties in the case of graphon games highlight the differences between the models.

1. *Since the role of the measure λ is played by the Lebesgue measure, the action of a single player does not impact the aggregate from which the costs of the other players are computed.*
2. *For each player x, while the aggregate influencing their cost is* ***linear*** *in the actions of the other players, it is not symmetric since it is the weighted average of the actions of the other players, say y, weighted by the values $w(x, y)$ of the graphon: in other words, unless w is a constant graphon, different players should feel different aggregates since they are potentiallhy different weighted averages for different players!*

Example 4.20 *The* ***min-max graphon*** *is defined by the formula*

$$w(x, y) = x \wedge y(1 - x \vee y).$$

It can be used to model the interactions between players in games for which agent states can be represented by points on a line and agents with central locations are more affected by their close neighbors. This can be illustrated by the fact that if $x = 1/2$, we have:

$$w(1/2, y) = \begin{cases} \frac{1}{2}y, & \text{if } y < \frac{1}{2}, \\ \frac{1}{2}(1-y), & \text{if } y \geqslant \frac{1}{2}, \end{cases}$$

confirming, at least for in this case, that the interaction is stronger with nearby players. A useful property of this graphon is proven in the following lemma.

Lemma 4.21 *For the min-max graphon we have:*

$$|w(x, y) - w(x', y')| \leqslant 2(|x - x'| + |y - y'|.$$

Proof: Since $2x \wedge y = x + y - |x - y|$, we have:

$$\begin{aligned} 2|x \wedge y - x' \wedge y'| &= \big|x + y - |x - y| - x' - y' + |x' - y'|\big| \\ &\leqslant |x - x'| + |y - y'| + \big|-|x - y| + |x' - y'|\big| \\ &\leqslant |x - x'| + |y - y'| + |-x + y + x' - y'| \\ &\leqslant 2(|x - x'| + |y - y'|, \end{aligned}$$

where we used the inequality $|a - b| \geqslant \big||a| - |b|\big|$. So we proved that the map $(x, y) \mapsto x \wedge y$ is Lipschitz with Lipschitz constant at most 1. We porve similarly that the function $(x, y) \mapsto 1 - x \vee y$ has the same property. Finally:

$$\begin{aligned} |w(x, y) - w(x', y')| &= \big|(x \wedge y)(1 - x \vee y) - (x' \wedge y')(1 - x' \vee y')\big| \\ &\leqslant \big|(x \wedge y)(1 - x \vee y) - (x' \wedge y')(1 - x \vee y)\big| \\ &\qquad + \big|(x' \wedge y')(1 - x \vee y) - (x' \wedge y')(1 - x' \vee y')\big| \\ &= \big|(x \wedge y - x' \wedge y')\big|(1 - x \vee y) + (x' \wedge y')\big|(1 - x \vee y) - (1 - x' \vee y')\big| \\ &= \big(|x - x'| + |y - y'|\big) + \big(|x - x'| + |y - y'|\big) \end{aligned}$$

which proves the desired result. □

Best Responses and Nash Equilibria

In a graphon game, the best response function is the function $\boldsymbol{\alpha}_{BR} : L_A \ni \boldsymbol{\alpha} \to \boldsymbol{\alpha}_{BR}(\boldsymbol{\alpha}) \in (\mathbb{R}^k)^{[0,1]}$ defined by:

$$\boldsymbol{\alpha}_{BR}(\boldsymbol{\alpha})(x) = \arg\inf_{\alpha \in A^x} J(\alpha, z(x|\boldsymbol{\alpha})), \qquad x \in [0,1]. \tag{4.11}$$

So the best response $\boldsymbol{\alpha}_{BR}(\boldsymbol{\alpha})$ to any admissible strategy profile $\boldsymbol{\alpha} \in L_A$ is a set value function on $[0,1]$ whose values are subsets of $\mathbb{R}^k$, including possibly the empty set in case the minimum of the function J in the first variable is not attained. The definition of a Nash equilibrium naturally follows.

Definition 4.22 *An admissible strategy profile $\boldsymbol{\alpha}^*$ is a Nash equilibrium for the graphon game if for almost every $x \in [0,1]$ (with respect to the Lebesgue measure of $[0,1]$), $\boldsymbol{\alpha}^*(x) \in \boldsymbol{\alpha}_{BR}(\boldsymbol{\alpha}^*)(x)$.*

Given the above remarks about the values of the aggregate, this means that for almost every $x \in [0,1]$, for every $\alpha \in A(x)$,

$$J(\boldsymbol{\alpha}^*(x), z(x|\boldsymbol{\alpha}^*)) \leqslant J(\alpha, z(x|\boldsymbol{\alpha}^*)).$$

Next, we tackle the question of existence and uniqueness of such an equilibrium. But before doing so, we emphasize some of the properties of the graphon operator which may not have been clear. Recall that given a graphon $w : [0,1] \times [0,1] \to [0,1]$, we defined the graphon operator A^w by $[A^w f](x) = \int f(y)w(x,y)dy$ for $f : [0,1] \to \mathbb{R}$ measurable and square integrable and we hinted at the fact that the definition applied as well to $\mathbb{R}^k$-valued functions f. This extension is important because since admissible strategy profiles take values in $\mathbb{R}^k$, we want to have the graphon operator act on such functions and consider $A^w : L^2([0,1];\mathbb{R}^k) \to L^2([0,1];\mathbb{R}^k)$ defined by

$$(A^w \boldsymbol{\alpha})(x) = \int \boldsymbol{\alpha}(y)w(x,y)dy = \begin{pmatrix} \int \alpha_1(y)w(x,y)dy \\ \dots \\ \int \alpha_k(y)w(x,y)dy \end{pmatrix}.$$

As we mentioned earlier, using the same notation for the operator acting on $L^2([0,1];\mathbb{R})$ or $L^2([0,1];\mathbb{R}^k)$ should not be an issue. We state its main properties as a lemma for future references.

Lemma 4.23 *The extension A^w of the graphon operator to $L^2([0,1];\mathbb{R}^k)$ is a symmetric Hilbert-Schmidt (hence compact) operator whose eigenvalues are, except for their multiplicities (which are at most multiplied by k), the same as those of the original graphon operator A^w when defined on $L^2([0,1];\mathbb{R})$.*

Proof: The proof of these claim is obvious. We merely illustrate the changes in multiplicity of the eigenvalues. If $f \in L^2([0,1];\mathbb{R})$ and $\lambda \in \mathbb{R}$ are such that $A^w f = \lambda f$, then:

$$A^w(f, 0, \cdots, 0) = (A^w f, 0, \cdots, 0) = (\lambda f, 0, \cdots, 0) = \lambda(f, 0, \cdots, 0)$$

and similarly:

$$A^w(0, f, 0, \cdots, 0) = (0, A^w f, 0, \cdots, 0) = (0, \lambda f, 0, \cdots, 0) = \lambda(0, f, 0, \cdots, 0)$$

and so on. So λ is an eigenvalue of the extended operator A^w with multiplicity k if it was an eigenvalue with multiplicity 1 of the original graphon operator A^w. $\square$

Assumptions and Preliminary Estimate

We first state the assumptions we shall use in the various steps of the existence analysis. They are taken from [30].

A1 The function $J : \mathbb{R}^k \times \mathbb{R}^k \ni (\alpha, z) \to J(\alpha, z) \in \mathbb{R}$ is continuously differentiable and strongly convex in α with a constant $\ell_c > 0$ uniform in $z \in \mathbb{R}^k$. This means that for every $\alpha', \alpha \in \mathbb{R}^k$ and $z \in \mathbb{R}^k$, we have for every $\epsilon \in [0,1]$

$$J(\epsilon\alpha + (1-\epsilon)\alpha', z) \leqslant \epsilon J(\alpha, z) + (1-\epsilon)J(\alpha', z) - \frac{\ell_c}{2}\epsilon(1-\epsilon)\|\alpha' - \alpha\|^2.$$

A2 $\nabla_\alpha J(\alpha, z)$ is Lipschitz in z with a Lipschitz constant ℓ_J uniform in $\alpha \in \mathbb{R}^k$. This means that for every $z', z \in \mathbb{R}^k$ and $\alpha \in \mathbb{R}^k$, we have

$$\|\nabla_\alpha J(\alpha, z') - \nabla_\alpha J(\alpha, z)\| \leqslant \ell_J \|z' - z\|.$$

A3 For every $x \in [0,1]$, the set of feasible strategy $A(x)$ is a closed convex subset of $\mathbb{R}^k$.

The following operator will play a crucial role in the existence proof. If $\tilde{z} : [0,1] \mapsto \mathbb{R}^k$ is a $\mathbb{R}^k$-valued function defined everywhere on $[0,1]$, for each $x \in [0,1]$ we set:

$$[B\tilde{z}](x) = \arg\inf_{\alpha \in A^x} J(\alpha, \tilde{z}(x)). \tag{4.12}$$

Notice that under assumption A1, for $x \in [0,1]$ fixed, the function $\mathbb{R}^k \ni \alpha \mapsto J(\alpha, \tilde{z}(x))$ is strongly convex, and if A3 holds, this function has a unique minimum on the closed convex set $A(X)$. This implies that the set $[B\tilde{z}](x)$ is a singleton, which defines a function $[0,1] \ni x \mapsto [B\tilde{z}](x) \in A(X) \subset \mathbb{R}^k$. Next, we remark that if $\tilde{z}$ and $\tilde{z}'$ are two $\mathbb{R}^k$-valued functions on $[0,1]$ which are equal almost everywhere, then $[B\tilde{z}](x) = [B\tilde{z}'](x)$ for almost every $x \in [0,1]$.

Lemma 4.24 *Under assumptions A1, A2 and A3, for every $\mathbb{R}^k$-valued functions f and g which are square integrable on $[0,1]$ we have:*

$$\|Bf - Bg\|_{L^2} \leqslant \frac{\ell_J}{\ell_c}\|f - g\|_{L^2}. \tag{4.13}$$

It is important to keep in mind the fact that $f \mapsto Bf$ is not linear.

Proof: The strong convexity assumption A1 is equivalent to the fact that for every α, α_0 and z in $\mathbb{R}^k$ we have:

$$[\nabla_\alpha J(\alpha, z) - \nabla_\alpha J(\alpha_0, z)] \cdots (\alpha - \alpha_0) \geqslant \ell_c|\alpha - \alpha_0|^2. \tag{4.14}$$

which implies

$$|\nabla_\alpha J(\alpha, z) - \nabla_\alpha J(\alpha_0, z)| \geqslant \ell_c|\alpha - \alpha_0|. \tag{4.15}$$

Let us assume that $x \in [0,1]$ is fixed momentarily. Using (4.15) with $z = g(x)$, $\alpha = [Bf](x)$ and $\alpha_0 = [Bg](x)$ we get:

$$|[Bf](x) - [Bg](x)| \leqslant \frac{1}{\ell_c}\big|\nabla_\alpha J([Bf](x), g(x)) - \nabla_\alpha J([Bg](x), g(x))\big|. \tag{4.16}$$

Since $[Bf](x) = \arg\inf_{\alpha \in A^x} J(\alpha, f(x))$, by convexity we have:

$$\nabla_\alpha J([Bf](x), f(x)) \cdot (\alpha - [Bf](x)) \geqslant 0$$

for all α, so using $\alpha = [Bg](x)$ in this inequality we get:

$$\nabla_\alpha J([Bf](x), f(x)) \cdot ([Bg](x) - [Bf](x)) \geqslant 0. \tag{4.17}$$

Similarly, we get:

$$\nabla_\alpha J([Bg](x), g(x)) \cdot ([Bf](x) - [Bg](x)) \geqslant 0. \tag{4.18}$$

Adding (4.17) and (4.18) we get:

$$\begin{aligned}
[\nabla_\alpha J([Bf](x), f(x)) &- \nabla_\alpha J([Bf](x), g(x)) \cdot ([Bg](x) - [Bf](x) \\
&\geqslant [\nabla_\alpha J([Bg](x), g(x)) - \nabla_\alpha J([Bf](x), g(x)) \cdot ([Bg](x) - [Bf](x)) \\
&\geqslant \ell_c |[Bg](x) - [Bf](x)|^2
\end{aligned} \tag{4.19}$$

where we used the strong convexity (4.14) to derive (4.19). The scalar product of two elements of $\mathbb{R}^k$ being smaller than the product of their norms, we see that the left hand side of (4.19) is bounded from above by:

$$|\nabla_\alpha J([Bf](x), f(x)) - \nabla_\alpha J([Bf](x), g(x))| \, |[Bg](x) - [Bf](x)|,$$

which, together with (4.19) gives:

$$\begin{aligned}
|[Bg](x) - [Bf](x)| &\leqslant \frac{1}{\ell_c} |[\nabla_\alpha J([Bf](x), f(x)) - \nabla_\alpha J([Bf](x), g(x))| \\
&\leqslant \frac{\ell_J}{\ell_c} |g(x) - f(x)|
\end{aligned} \tag{4.20}$$

where we used assumption A2 to derive (4.20). Since this inequality between non-negative real numbers is true for all $x \in [0, 1]$, we can square both sides, integrate both sides between 0 and 1 and take square roots of both sides, proving the desired estimate (4.13). □

The estimate (4.13) of Lemma 4.24 only depends upon the definition (4.12) of the operator B and the nature of assumptions A1, A2, and A3.

4.3.0.1 ***Existence of Equilibria***

Using the definition of the graphon operator A^w, we easily see that (recall the definition (4.11) of the best response function): if $\boldsymbol{\alpha}$ is an admissible stragegy profile in L^2_A,

$$\boldsymbol{\alpha} \text{ is a Nash equilibrium for the graphon game} \Longleftrightarrow \boldsymbol{\alpha} \text{ is a fixed point of } BA^w.$$

Indeed, if $\boldsymbol{\alpha} \in L^2_A$, $[A^w\boldsymbol{\alpha}](x) = z(x|\boldsymbol{\alpha})$ for almost every $x \in [0, 1]$ and as a result $[BA^w\boldsymbol{\alpha}](x) = \arg\inf_{\alpha \in A^x} J(\alpha, z(x|\boldsymbol{\alpha}))$ for almost every $x \in [0, 1]$, and demanding that

$$\boldsymbol{\alpha}(x) = \arg\inf_{\alpha \in A^x} J(\alpha, z(x|\boldsymbol{\alpha}))$$

for almost every $x \in [0, 1]$ is the very definition of a Nash equilibrium for the graphon game.

We now use Schauder's fixed point theorem in the form recalled in Theorem 1.6 to prove existence of Nash equilibria for graphon games. Our proof will require the following extra assumption saying that all the sets A^x of feasible actions are contained in the same bounded subset of $\mathbb{R}^k$. This will provide the compactness needed for the application of Schauder's theorem.

A4 The number α_{max} defined by:

$$\alpha_{max} = \sup_{x \in [0,1]} \sup_{\alpha \in A^x} |\alpha|$$

is finite.

Notice that under assumption A4, for every $f \in L^2([0,1], dx; \mathbb{R}^k)$ and almost every $x \in [0,1]$, $|[Bf](x)| \leqslant \alpha_{max}$.

Proposition 4.25 *If the graphon game $\mathcal{G}(\mathbf{A}, J, W)$ satisfies assumptions A1-A4, then it admits at least one Nash equilibrium.*

Proof: To be consistent with the notations of the statement of Theorem 1.6, we introduce the set $C = \{f \in L^2([0,1], dx; \mathbb{R}^k);\ \|f\| \leqslant \alpha_{max}\}$. C so defined is a closed convex subset of the Hilbert space $E = L^2([0,1], dx; \mathbb{R}^k)$. Also, for every $f \in L^2([0,1], dx; \mathbb{R}^k)$, $Bf \in C$, and consequently $BA^w f \in C$.

Since the (vector valued extension of the) graphon operator A^w is a Hilbert-Schmidt operator on $L^2([0,1], dx; \mathbb{R}^k)$, recall Lemma 4.23, it is a fortiori a compact operator and since C is bounded, the closure $\tilde{K}$ of $A^w C$ is a compact subset of $L^2([0,1], dx; \mathbb{R}^k)$. Being Lipschitz, B is also continuous and $K = B\tilde{K}$ is also compact. Consequently, the map $F = BA^w$ satisfies the assumption of Theorem 1.6, proving the existence of a fixed point. □

4.3.0.2 ***Uniqueness***

Proposition 4.26 *If the graphon game $\mathcal{G}(\mathbf{A}, J, w)$ satisfies assumptions A1-A3, and:*

$$\frac{\ell_J}{\ell_c} \lambda_{max}(A^w) < 1,$$

where $\lambda_{max}(w)$ is the largest eigenvalue of A^w, then there exists one and only one Nash equilibrium $\boldsymbol{\alpha}^* \in L^2_A$.

Proof: If f and g are elements of $L^2([0,1], dx; \mathbb{R}^k)$, using the Lipschitz property (4.13) and the fact that the norm of the operator A^w is its largest eigenvalue, we get:

$$\|BA^w f - BA^w g\|_{L^2} \leqslant \frac{\ell_J}{\ell_c} \|A^w f - A^w g\|_{L^2} \leqslant \frac{\ell_J}{\ell_c} \lambda_{max} < 1$$

which proves that the map BA^w is a strict contraction on $L^2([0,1], dx; \mathbb{R}^k)$, implying existence and uniqueness of a fixed point. □

4.3.1 **Example of the LQ Graphon Games**

Like in classical control theory and the theory of differential games linear quadratic models provide one of the rare classes of explicitly solvable models. Let us consider a graphon game $\mathcal{G}(\mathbf{A}, J, w)$ such that $A^x = [0, \infty)$ for every $x \in [0,1]$ so that assumption A3 is satisfied, and the cost function J is given by:

$$J(\alpha, z) = \frac{1}{2}\alpha^2 - \alpha(az + b)$$

for constants $a \in \mathbb{R}$ and $b > 0$. Clearly, the function J satisfies assumption A1 with constant $\ell_c = 1$ and since:

$$\frac{\partial J(\alpha, z)}{\partial \alpha} = \alpha - (az + b).$$

it also satisfies assumption A2 with constant $\ell_J = a$. So if the graphon w is such that:

$$|a|\lambda_{\max}(A^w) < 1, \tag{4.21}$$

then Proposition 4.26 implies that the graphon game has a unique Nash equilibrium. We proceed to compute it explicitly. The best response of player $x \in [0,1]$ to a dtrategy profile $\boldsymbol{\alpha}$ is given by:

$$\underline{\alpha}_{BR}(\boldsymbol{\alpha})(x) = \max\{0, az(x|\boldsymbol{\alpha}) + b\} \tag{4.22}$$

where $z(x|\boldsymbol{\alpha}) = \int_{[0,1]} \boldsymbol{\alpha}(y)w(x,y)dy$.

We distinguish two cases.

- If $a > 0$, the best response of each agent $x \in [0,1]$ is an increasing function of the aggregate $z(x|\boldsymbol{\alpha})$, so the game is a ***game of strategic complements***. Indeed, for every $x \in [0,1]$, $a > 0$ and $b > 0$, then $z(x|\boldsymbol{\alpha}) \geqslant 0$ so that $\partial_\alpha J(0, z(x|\boldsymbol{\alpha})) < 0$. Hence

  $$J(\alpha, z(x|\boldsymbol{\alpha})) = J(0, z(x|\boldsymbol{\alpha})) + \partial_\alpha J(0, z(x|\boldsymbol{\alpha}))\alpha + o(\alpha) < J(0, z(x|\boldsymbol{\alpha}))$$

  for $\alpha > 0$ small enough. So 0 cannot be the best response since the cost could be lowered by a small perturbation. As a consequence $\boldsymbol{\alpha}_{BR}(\boldsymbol{\alpha}) > 0$, and the Nash equilibrium $\boldsymbol{\alpha}^*$ is internal in the sense that

  $$\boldsymbol{\alpha}^*(x) > 0, \qquad \text{for all } x \in [0,1].$$

  From the form (4.22) of the best response function, now that we know that the Nash equilibrium is internal, we get

  $$\boldsymbol{\alpha}^*(x) = az(x|\boldsymbol{\alpha}^*) + b = a[A^w\boldsymbol{\alpha}^*](x) + b$$

  for almost every $x \in [0,1]$, which implies that:

  $$[I - aA^w]\boldsymbol{\alpha}^* = b\mathbf{1}_{[0,1]}$$

  and if the operator $[I - aA^w]$ is invertible, which is the case since we assume that $|a|\lambda_{max}(A^w)$, we have:

  $$\boxed{\boldsymbol{\alpha}^* = b[I - aA^w]^{-1}\mathbf{1}_{[0,1]} = bb_a}$$

  where b_a is the Bonacich measure of centrality of the graphon. Recall (4.4). The importance of this result is that, as it was already the case for finitely many player games, the values of the Nash equilibrium are entirely determined by the geometric information contained in the underlying graph.

- If $a < 0$, the best response of each agent $x \in [0,1]$ is an decreasing function of the aggregate $z(x|\boldsymbol{\alpha})$, so the game is a ***game of strategic substitutes***. We cannot guarantee that equilibria are internal, but under the uniqueness condition (4.21), we see that the unique Nash equilibrium is internal if and only if equation:

$$[I + |a|A^w]\boldsymbol{\alpha}^* = b\mathbf{1}_{[0,1]}$$

has a solution $\boldsymbol{\alpha}^*$ such that $\boldsymbol{\alpha}^*(x) > 0$ for almost every $x \in [0,1]$. Condition (4.21) implies that the operator $I + |a|A^w$ is invertible, so the candidate to be the unique Nash equilibrium is given by:

$$\boldsymbol{\alpha} = b[I + |a|A^w]^{-1}\mathbf{1}_{[0,1]},$$

but we cannot be sure that it is internal, and hence the fixed point of the best response function, since the sign of the sum of the Taylor series expansion:

$$\boldsymbol{\alpha} = b\Big(I - |a|A^w\mathbf{1}_{[0,1]} + -|a|^2(A^w)^2\mathbf{1}_{[0,1]} - |a|^3(A^w)^3\mathbf{1}_{[0,1]} + \cdots\Big)$$

cannot be determined for every $x \in [0,1]$ without more information.

4.4 GAMES OF COORDINATION AND INFORMATION ACQUISITION

4.4.1 Coordination and Information Acquisition: Beauty Contest

We present a game model studied in [27]. We chose it to illustrate the use of families of independent random variables to model signals informing players in a game. This particular example is an instance of a game model studied by Morris and Shin under the name of *beauty contest*. As all the games we considered so far, the game is a one-shot game in which a continuum of players take actions simultaneously. When information is costly, agents must balance the cost of information against its benefits. The latter depend on the likely actions of the other agents and the information they acquired.

We model an industry in which the supplier's product demand depends upon:

- an uncertainty state of the market place (e.g. size of customer base);
- the supplier own price;
- average price of the supplier's competitors.

To improve decision making a supplier relies on a survey of the market conditions. The information is transmitted through signals which can be private or public. So the information can be exogenous or endogenous.

This is the description of the specific example from [27].

- Players are represented by real numbers $\ell \in [0,1]$. They act simultaneously.
- Player ℓ chooses an information acquisition policy $z_\ell \in \mathbb{R}^n_+$. The integer n is the number of sources of information. For each source $i \in \{1,\dots,n\}$, $z_{\ell i}$ represents the attention paid by player ℓ to the source i.

- After the choice of z_ℓ, the player observes a vector $x_\ell = (x_{\ell 1}, \dots, x_{\ell n}) \in \mathbb{R}^n$ of signals which inform the player about an unobservable variable θ. The precision of the signals depending upon the choice of $z_{\ell i}$.
- On the basis of this information, player ℓ then takes action $\alpha_\ell \in A = \mathbb{R}$.
- Player ℓ uses pure strategies (z_ℓ, a_ℓ) with $z_\ell \in \mathbb{R}^n_+$ and $a_\ell : \mathbb{R}^n \mapsto \mathbb{R}$. Their actions are of the form $\alpha_\ell = a_\ell(x_\ell)$, so the action α_ℓ taken by payer ℓ is a *feedback* function of the signal x_ℓ acquired by the player.
- The cost to player $\ell \in [0, 1]$ is defined as:

$$u_\ell = C(z_\ell) + (1-\gamma)(\alpha_\ell - \theta)^2 + \gamma(\alpha_\ell - \overline{\alpha})^2$$

  where $C : \mathbb{R}_+ \to \mathbb{R}$ is differentiable and increasing, $\gamma \in [-1, 1]$ and $\overline{\alpha}$ represents the average action, namely

$$\overline{\alpha} = \int_0^1 \alpha_\ell d\ell.$$

  $\gamma \in (-1, 1)$ is a parameter which balances the cost of aligning the action of player ℓ with the others (coordination) and matching the state variable (the fundamental motive of the game). The term $(\alpha_\ell - \overline{\alpha})^2$ represents the coordination with the other players: if $\gamma > 0$, each player will want to choose an action as similar as possible to the average action of their peers . On the other end, if $\gamma < 0$, each player tries to behave differently than others.

Remark 4.27 *Note that from a pure mathematical point of view, in order for the average action $\overline{\alpha}$ to be well defined, we need to assume that the function $[0, 1] \ni \ell \mapsto \alpha_\ell$ is measurable and integrable.*

According to our terminology, the present game model should be called a *mean field game* because the actions of the other players enter the cost to a given player through their mean as given by their plain aggregate. Indeed, as we argued earlier, the above integral should be understood in the context of games with a continuum of players, as the equivalent of the empirical mean in the case of games with finitely many players. We shall come back to this interpretation in the next section.

Signals: the Sources of Information

Players start the game with no knowledge of the value of θ. Obviously we could put a prior distribution on the set of possible values of θ. We shall not do that here for the sake of simplicity. The signal observed by player ℓ from source i is assumed to be of the form:

$$x_{\ell i} = \theta + \eta_i + \epsilon_{\ell i}$$

where $\eta_i \sim \mathcal{N}(0, \kappa_i^2)$ and $\epsilon_{\ell i} \sim \mathcal{N}(0, \frac{\xi_i^2}{z_{\ell i}})$, all these Gaussian random variables being **independent**, and κ_i, ξ_i are constant.

- In other words, we assume that each signal source has its own sender noise which we denote by $\tilde{x}_i = \theta + \eta_i$ and which are assumed to be Gaussian and have precision $1/\kappa_i^2$. The noise term η_i stands the intrinsic noise of the information source i.

- If player ℓ chooses to pay attention to source i, they do so imperfectly, adding an idiosyncratic noise term $\epsilon_{\ell i}$, so that $x_{\ell i} = \tilde{x}_i + \epsilon_{\ell i}$, which represents the part of the noise brought by the player. We also notice that the variance of $\epsilon_{\ell i}$, namely ξ_i^2/z_{li} depends upon the constant ξ_i and is inversely proportional to the attention $z_{\ell i}$ that player ℓ pays to the source of information i. In other words, the more attention player ℓ gives to signal source i, the less the noise they incur.

We highlight the independence assumption because we intend to revisit it within next section.

Assumption 1. *We suppose that η_i and $\epsilon_{\ell i}$ for $i \in \{1 \ldots, N\}$ and $\ell \in [0,1]$ are all independent.*

Dependencies between the Signals

Let us look at the covariance of the signal observations of the same information source from different players. If $\ell \neq \ell'$, we have:

$$\begin{aligned}
\mathrm{cov}(x_{\ell i}, x_{\ell' i}) &= \mathbb{E}\left[(x_{\ell i} - \mathbb{E}[x_{\ell i}])(x_{\ell' i} - \mathbb{E}[x_{\ell' i}])\right] \\
&= \mathbb{E}[(\eta_i + \epsilon_{\ell i})(\eta_i + \epsilon_{\ell' i})] \\
&= \mathbb{E}[\eta_i^2] + \mathbb{E}[\eta_i \epsilon_{\ell' i}] + \mathbb{E}[\epsilon_{\ell i}\eta_i] + \mathbb{E}[\epsilon_{\ell i}\epsilon_{\ell' i}] \\
&= \kappa_i^2.
\end{aligned}$$

and if we denote by $\rho_{\ell\ell' i}$ the correlation between $x_{\ell i}$ and $x_{\ell' i}$, since $x_{\ell i} \sim \mathcal{N}(\theta, \sigma_{\ell i}^2)$ with $\sigma_{\ell i}^2 = \kappa_i^2 + \xi_i^2/z_{\ell i}$, the fact that $\mathrm{cov}(x_{\ell i}, x_{\ell' i}) = \rho_{\ell\ell' i}\sigma_{\ell i}\sigma_{\ell' i}$, we conclude that:

$$\rho_{\ell\ell' i} = \frac{\kappa_i^2}{\sqrt{\left(\kappa_i^2 + \frac{\xi_i^2}{z_{\ell i}}\right)\left(\kappa_i^2 + \frac{\xi_i^2}{z_{\ell' i}}\right)}}.$$

Remark 4.28

- *If a player pays more attention to information, then the precisions of the corresponding observed signals increase ($z_{\ell i}$ increases), which will also lead to an increase of the correlation $\rho_{\ell\ell' i}$ between this player and the others.*
- *If $\xi_i = 0$ or $z_{\ell i} \to \infty$, then $\rho_{\ell\ell' i} \to 1$. In this case, we say that the signals are* public, *i.e. the observations $x_{\ell i}$ are the same for all the players.*
- *When $\kappa_i = 0$, then $\rho_{\ell\ell' i} = 0$, and since the signals are jointly Gaussian, they are independent and in this situation, we say that the signals are* private.
- *In this model, in general we have $0 < \rho_{\ell\ell' i} < 1$ and the signals are* endogenous *since they depend upon the decisions of the players.*

Equilibrium Analysis

First we recall the definition of the strategies in the model.

Definition 4.29 *For player $\ell \in [0,1]$, a strategy is a couple (z_ℓ, a_ℓ) where $z_\ell \in \mathbb{R}^n_+$ and $a_\ell : \mathbb{R}^n \to \mathbb{R}$. A strategy (z_ℓ, a_ℓ) is of feedback form if the action taken by player ℓ satisfies*

$$\alpha_\ell = a_\ell(x_\ell).$$

where x_ℓ is the signal observed by player ℓ.

Searching for Nash equilibria in the full generality of the model is highly difficult and technical, so we shall limit our search to a subclass of equilibria for which we can take advantage of special features of the models.

Definition 4.30 *The game is said to be symmetric if all the players use a common strategy (z, a).*

For any given $\ell \in [0,1]$, suppose that player ℓ choses a strategy (z_ℓ, a_ℓ) while all the other players $\ell' \neq \ell$, use strategy (z, a), then the cost to player ℓ is given by:

$$J^\ell((z_\ell, a_\ell), (z_{-\ell}, a_{-\ell})) = \mathbb{E}[u_\ell] = C(z_\ell) + (1-\gamma)\mathbb{E}[(\alpha_\ell - \theta)^2] + \gamma\mathbb{E}[(\alpha_\ell - \overline{\alpha})^2] \quad (4.23)$$

where

$$\overline{\alpha} = \int_0^1 a(x_{\ell'})d\ell'. \quad (4.24)$$

since all the players but possibly one use the same strategy (z, a).

Remark 4.31 *The definition of $\overline{\alpha}$ seems to say that it should be a random variable because x_ℓ are random variables for all $\ell \in [0,1]$. In order to understand clearly the issue which needs to be addressed here, we momentarily come back to the case of a game with N players with observations $x_1, \dots, x_N$ and $\ell \in [N]$. We can still assume that $x_{\ell i} = \theta + \eta_i + \epsilon_{\ell i}$. Then in this case:*

$$\overline{\alpha} = \frac{1}{N}\sum_{e=1}^{N} \alpha_\ell = \frac{1}{N}\sum_{e=1}^{N} a(x_\ell) \sim \mathbb{E}[a(x)]$$

if we use the standard law of large numbers to replace the last average by the expectation of a random variable having the same distribution as those random variables appearing in the averages. Because of the law of large numbers, $\overline{\alpha}$ is deterministic in the limit of large games.

If we come back to the definition (4.24) *of $\overline{\alpha}$ in the continuum of players case, for $\overline{\alpha}$ to be deterministic, it should be equal to its expectation and if the random variables x_ℓ are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and if we can use Fubini's theorem to interchange integrals, we should have:*

$$\begin{aligned}\overline{\alpha} = \mathbb{E}[\overline{\alpha}] = \int_\Omega \Big(\int_0^1 a(x_\ell)d\ell\Big)\mathbb{P}(d\omega) &= \int_0^1 \Big(\int_\Omega a(x_\ell)\mathbb{P}(d\omega)\Big)d\ell \\ &= \int_0^1 \mathbb{E}[a(x_\ell)]d\ell = \mathbb{E}[a(x)]\end{aligned}$$

if x is a random variable with distribution equal to the common distribution of all the x_ℓ. So it seems that if we have a ***continuum of independent*** *random variables with the same* *distribution, the law of large numbers seems to give the same result as the application of Fubini's theorem. However, before one can apply any form of Fubini's theorem, one would need to make sure that the function* $(\ell, \omega) \mapsto a\big(x_\ell(\omega)\big)$ *is jointly measurable on* $[0,1] \times \Omega$. *This is not a given and we address this issue in the next section.*

While restricting ourselves to symmetric models already made the analysis more manageable, the identification of Nash equilibria is still out of reach. For this reason, we further limit the search to a subclass of strategies. From now on, we shall search for symmetric Nash equilibria among linear strategies. To be specific we limit ourselves to feedback functions of the form:

$$a_\ell(x) = \sum_{i=1}^n a_{\ell i} x_i \tag{4.25}$$

with $\sum_{i=1}^n a_{\ell i} = 1$. Then since we look for symmetric equilibria we compute the average $\overline{\alpha}$ with a common strategy a for all players (except possibly one). Consequently we compute $\overline{\alpha}$ with a strategy $a : \mathbb{R}^n \ni x \mapsto \sum_{i=1}^n a_i x_{li} \in \mathbb{R}$ and $x_{\ell i} = \theta + \eta_i + \epsilon_{\ell i}$, so that:

$$\begin{aligned}
\overline{\alpha} = \int_0^1 a(x_{\ell'}) d\ell' &= \int_0^1 \left(\sum_{i=1}^n a_i x_{\ell' i} \right) d\ell' \\
&= \int_0^1 \left(\sum_{i=1}^n a_i (\theta + \eta_i + \epsilon_{\ell' i}) \right) d\ell' \\
&= \theta + \sum_{i=1}^n a_i \eta_i + \int_0^1 \left(\sum_{i=1}^n a_i \epsilon_{\ell' i} \right) d\ell' \\
&= \theta + \sum_{i=1}^n a_i \eta_i
\end{aligned}$$

where the last equality is due to the fact that

$$\int_0^1 \epsilon_{\ell' i} d\ell' = 0$$

which cannot be justified yet, still, recall the content of the above Remark 4.31, but which follows from the *exact law of large numbers* which we shall prove in the next section. Consequently:

$$\begin{aligned}
\mathbb{E}\Big[\big(\sum_{i=1}^n a_{\ell i} x_{\ell i} - \overline{\alpha} \big)^2 \Big] &= \mathbb{E}\Big[\big(\sum_{i=1}^n a_{\ell i} \epsilon_{\ell i} + \sum_{i=1}^n (a_{\ell i} - a_i) \eta_i \big)^2 \Big] \\
&= \sum_{i=1}^n \frac{a_{\ell i}^2 \xi_i^2}{z_{\ell i}} + \sum_{i=1}^n (a_{\ell i} - a_i)^2 \kappa_i^2
\end{aligned}$$

So returning to formula (4.23) and using the fact that the function a is linear, recall (4.25), we get:

$$
\begin{aligned}
&J^\ell((z_\ell, a_\ell), (z_{-\ell}, a_{-\ell})) \\
&\quad = C(z_\ell) + (1-\gamma)\mathbb{E}\Big[\big(\sum_{i=1}^n a_{\ell i} x_{\ell i} - \theta\big)^2\Big] + \gamma\mathbb{E}\Big[\big(\sum_{i=1}^n a_{\ell i} x_{\ell i} - \overline{\alpha}\big)^2\Big] \\
&\quad = C(z_\ell) + (1-\gamma)\mathbb{E}\Big[\big(\sum_{i=1}^n a_{\ell i}(x_{\ell i} - \theta\big)^2\Big] + \gamma \sum_{i=1}^n \frac{a_{\ell i}^2 \xi_i^2}{z_{\ell i}} + \gamma \sum_{i=1}^n (a_{\ell i} - a_i)^2 \kappa_i^2 \\
&\quad = C(z_\ell) + \sum_{i=1}^n a_{\ell i}^2 [(1-\gamma)\kappa_i^2 + \frac{\xi_i^2}{z_{\ell i}}] + \gamma \sum_{i=1}^n (a_{\ell i} - a_i)^2 \kappa_i^2
\end{aligned}
\tag{4.26}
$$

One of the main reasons for the presentation of this game model was to motivate the study of Fubini extensions and the exact law of large numbers presented in the next section. This goal being essentially attained, we do not pursue the detailed analysis of the model. For the sake of completeness, we state without proofs the main results of the paper [27] giving an almost complete characterization of the linear symmetric Nash equilibria of the game.

If we rewrite the right hand side of (4.26) as $C(z_\ell) + J^{(1)}(z_\ell, a_\ell) + J^{(2)}(z_\ell, a_\ell)$, using the fact that the function C is convex, one can show that in a symmetric equilibrium (for which the last part $J^{(2)}(z_\ell, a_\ell)$ disappears from the optimization), each players tries to minimize $C(z_\ell) + J^{(1)}(z_\ell, a_\ell)$, leading to the useful characterization of these equilibria.

Lemma 4.32 *A strategy (z, a) is a symmetric Nash equilibrium if and only if it solves the minimization problem*

$$
(z, a) \in \arg \inf_{(z,\tilde{a})} \sum_{i=1}^n \tilde{a}_i^2 [(1-\gamma)\kappa_i^2 + \frac{\xi_i^2}{\tilde{z}_i}] + C(\tilde{z})
$$

under the constraint $\sum_{i=1}^n \tilde{a}_i = 1$.

Finally, the result of the above lemma can be used to prove:

Proposition 4.33 *The exists a unique symmetric linear Nash equilibrium (z^*, a^*), and in this equilibrium, the influence a_i^* of the i-th signal and the attention z_i^* paid to it satisfy:*

$$
a_i^* = \frac{\psi_i}{\sum_{i=1}^n \psi_i} \quad \textit{and} \quad z_i^* = \frac{a_i^* \xi_i}{\sqrt{C'(z^*)}} \quad \textit{with} \quad \psi_i = \frac{1}{(1-\gamma)\kappa_i^2 + \xi_i^2 / z_i^*}
$$

where ψ_i is set to 0 whenever $z_i^* = a_i^* = 0$.

4.5 Fubini Extensions and the Exact Law of Large Numbers

The models of non-atomic games presented in the previous section raise a few delicate mathematical questions. Indeed, by considering games for which the players are labelled by elements i of an uncountable set I, this set I being equipped with a σ-field $\mathcal{I}$ and a probability measure λ which is assumed to be continuous (i.e. non-atomic), one may wonder if objects like $(\epsilon_i)_{i \in I}$ exist if one requires the ϵ_i's to be independent identically distributed mean-zero random variables, and if one expects to be measurable functions of the variable i (for example if we want to integrate over i to consider continuous averages) and possible use a form of law of large numbers to capture the properties of the empirical distribution of these random variables.

This type of *uncountable sequence* (please pardon me for the oxymoron) of random variables is what is usually called a *white noise*. While very popular in modeling circles, it is much less so among mathematical purists, and we are about to see the reasons why.

The goal of this section is to present the mathematical tools from real analysis and measure theory which were developed in part to address these questions.

Over fifty years ago, economists suggested that the appropriate model for perfectly competitive markets is a model with a continuum of traders represented by elements of a measurable space. In such a set-up, traders are treated as *price takers* as it is assumed that their individual influence on the prices is negligible. In such a set-up, summation or aggregation is generalized by the notion of integral. In games with a continuum of players, the latter are labelled by the elements $i \in I$ of an arbitrary set I (often assumed to be uncountable, and most often chosen to be the unit interval $[0, 1]$) equipped with a σ-field $\mathcal{I}$, and a probability measure λ. So if the state of each player $i \in I$ is given by a random variable X^i on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, in analogy with the countable case leading to formula (4.28), the quantity:

$$F_\omega(x) = \lambda\big(\{i \in I;\ X^i(\omega) \leqslant x\}\big) \tag{4.27}$$

appears as a natural generalization of the *proportion of $X^i(\omega)$'s which are not greater than x*, in other words, of the cumulative distribution function of the empirical distribution. And if the X^i's were to be independent with the same distribution, it would be reasonable to expect that a generalization of the Law of Large Numbers to this setting could hold. However, as we explain in the next subsection, measurability issues get in the way, and such a generalization is, when it does exist, far from trivial.

4.5.1 Kolmogorov Extension Theorem

Definition 4.34 *Let E be a Polish space (i.e. a space which is homeomorphic to a separable complete metric space) and let $\mathcal{E} = \mathcal{B}(E)$ be its Borel $\sigma-$field. If I is an arbitrary index set, and if for each finite tuple $(i_1, \dots, i_k) \in I^k$ we have a probability measure $\mu_{(i_1,\dots,i_k)}$ on $(E_{i_1} \times \cdots \times E_{i_k}, \mathcal{E}_{i_1} \otimes \ldots \otimes \mathcal{E}_{i_k})$ where $E_{i_j} = E$ and $\mathcal{E}_{i_j} = \mathcal{E}$ for $j = 1, \dots, k$, we say that the family of probability measures $\{\mu_{(i_1,\dots,i_k)}\}$ is consistent if:*

1. *for each $\{i_1, \dots, i_k\} \subset I$ and each permutation π of $\{i_1, \dots, i_k\}$, we have*

$$\mu_{(i_1,\dots,i_k)} = \mu_{(\pi(i_1),\dots,\pi(i_k))} \circ \varphi_\pi^{-1}$$

where $\varphi_\pi(x_{\pi(i_1)}, \dots, x_{\pi(i_k)}) = (x_{i_1}, \dots, x_{i_k})$.

2. *for each $\{i_1, \dots, i_k, i_{k+1}\} \subset I$, for any $A \in \mathcal{E}_{i_1} \otimes \cdots \otimes \mathcal{E}_{i_k}$,*

$$\mu_{(i_1,\dots,i_k)}(A) = \mu_{(i_1,\dots,i_k,i_{k+1})}(A \times E).$$

To be more specific one may want to think of E as one of the spaces we will use the result of Kolmogorov's theorem stated without proof below. These spaces could be, and will be, $\mathbb{R}$, $\mathbb{R}^d$, a separable Hilbert space like L^2 or a more general separable Banach space like L^p with $1 \leqslant p < \infty$ cor the space $C([0,T];B)$ of continuous (hence bounded) functions from the interval $[0,T]$ into a real separable Banach space B. We state without proof the fundamental existence theorem of Kolmogorov. Most introductory textbooks of measure theory based probability provide a proof. For the sake of definiteness we refer the interested reader to [5].

Theorem 4.35 (Kolmogorov existence theorem) *If $\Omega = E^I$ is the product of copies of a Polish space E and can be interpreted as the set of all the functions from I into E, if $\mathcal{F} = \mathcal{E}^I$ with $\mathcal{E}^I = \bigotimes_{i\in I}\mathcal{E}_i$ is the product $\sigma-$field of copies of the Borel σ-field of E, and if $\boldsymbol{\mu} = \{\mu_{(i_1,\dots i_k)}\}$ is a consistent family of probability measures, then there exists a probability measure $\mathbb{P}_{\boldsymbol{\mu}}$ on $(\Omega,\mathcal{F})$ such that if for every $i \in I$ and every $\omega \in \Omega$ we set*

$$X^i(\omega) = \omega(i)$$

for the E-valued i-th coordinate map, then for every $\{i_1,\dots,i_k\} \subseteq I$ we have

$$\mathcal{L}_{\mathbb{P}_{\boldsymbol{\mu}}}(X^{i_1},\dots,X^{i_k}) = \mu_{(i_1,\dots,i_k)}.$$

As usual, we denote by $\mathcal{L}_{\mathbb{P}}(Z)$ the law (distribution) of the random variable or random vector Z under the probability $\mathbb{P}$.

Remark 4.36 *So given a set of consistent finite dimensional probability distributions over an arbitrary index set I, ONE CAN ALWAYS construct a probability space $(\Omega,\mathcal{F},\mathbb{P})$ and a stochastic process $(X^i)_{i\in I}$ on this probability space such that the finite dimensional marginals of the process are the probability measures we started from.*

Kolmogorov's theorem laid the foundations of the theory of stochastic processes on a firm mathematical ground, and gave mathematicians the option to make completely rigorous a good number of intuitive arguments about the properties of these processes. Indeed, the latter are often introduced through specific properties of their finite dimensional marginal distributions, and Kolmogorov's existence theorem guarantees the existence of processes defined in this way. Basic examples include:

- The process of Brownian motion.
- The Poisson process.
- White noise processes introduced as independent and identically distributed families $(X^i)_{i\in I}$ of random variables with given common distribution $\mu \in \mathcal{P}(E)$.

4.5.2 **Structure of Product Sigma Fields**

Kolmogorov's theorem has the merit of resolving the existence problem for stochastic processes. Still it is not without shortcomings, and as we are about to see, it is extremely unsatisfactory when it comes to manipulations of the process so constructed. To illustrate this important point, we first discuss the fine structure of the product σ-field $\mathcal{E}^I = \bigotimes_{i\in I}\mathcal{E}_i$ with $\mathcal{E}_i = \mathcal{E}$ for all $i \in I$.

Recall that the product σ-field is the smallest σ-field containing the measurable cylinders defined as the subsets of $\Omega = E^I$ of the form:

$$\big\{\omega \in \Omega;\ \omega(i_j) \in A_{i_j},\ j = 1, \dots, k\big\} \in \mathcal{E}^I.$$

for a finite subset $\{i_1, \dots, i_k\}$ of I and $A_{i_1}, \dots, A_{i_k}$ in $\mathcal{E}$. As such, the product σ-field appears as the σ-field generated by the marginal coordinate projections.

For convenience, we recall the definition of a σ-field generated by a family of sets or functions.

Definition 4.37 *For an arbitrary space Ω, an arbitrary index set I, and a measurable space $(E, \mathcal{E})$, let us assume that $\{Z^i\}_{i\in I}$ is a family of functions $Z^i : \Omega \to E$. We denote by $\mathcal{F} = \sigma\{Z^i : i \in I\}$ the smallest σ-field $\mathcal{F}$ of subsets of Ω for which Z^i is $(\mathcal{F}, \mathcal{E})$ measurable for all $i \in I$. In other words, $\mathcal{F}$ is the smallest σ-field $\mathcal{F}$ of subsets of Ω containing the sets $(Z^i)^{-1}(B)$ for all $i \in I$ and $B \in \mathcal{E}$.*

Let us now consider an arbitrary index set I and for each $i \in I$, a measurable space $(E_i, \mathcal{E}_i)$. When we discuss white noises later on, we will be specially interested in the special case $E_i = E$ and $\mathcal{E}_i = \mathcal{E}$ for a given measurable space $(E, \mathcal{E})$. The *product space* is denoted by $\prod_{i\in I} E_i$, or sometimes $\times_{i\in I} E_i$. As we already did earlier, we shall denote the coordinate functions by $(X^i)_{i\in I}$ where $X^i : \prod_{i\in I} E_i \to E$ is defined by

$$X^i(\omega) = \omega_i, \qquad \forall\, \omega = (\omega_i)_{i\in I} \in \prod_{i\in I} E_i.$$

When all the E_i are equal to the same E, the product is denoted by E^I. It can be identified to the collection of all the functions from I into E. In this case we often use the notation $\omega(i)$ for ω_i to emphasize that fact. It is plain to show the following result which we state as a lemma for the sake of definiteness.

Lemma 4.38 *The product σ-field $\bigotimes_{i\in I} \mathcal{E}_i$ is the smallest $\sigma-$field of subsets of E^I that makes the coordinate projection function X^i measurable for all $i \in I$. Namely*

$$\bigotimes_{i\in I} \mathcal{E} = \sigma\{X^i;\ i \in I\}.$$

The result of the next lemma states in a different context the same property of the product σ-field.

Lemma 4.39 *Let $(\Omega, \mathcal{F})$ be a measurable space and I be an arbitrary index set, $(E, \mathcal{E})$ a measurable space, and $(Y^i)_{i\in I}$ a family of functions from Ω into E. We define $\xi : \Omega \ni \omega \mapsto \xi(\omega) \in E^I$ such that*

$$\xi(\omega)(i) = Y^i(\omega), \qquad i \in I,\ \omega \in \Omega.$$

Then ξ is $(\mathcal{F}, \mathcal{E}^I)$-measurable if and only if for all $i \in I$, Y^i is $(\mathcal{F}, \mathcal{E})-$measurable.

In fact, $\sigma\{Y^i;\ i \in I\}$ is the smallest σ-field $\mathcal{F}$ satisfying the assumptions of the above lemma. The next result is the main property of σ-fields generated by families of functions (and hence of product σ-fields) which we use in our analysis of white noises later on.

Proposition 4.40 *Let Ω be an arbitrary set, I an arbitrary index set, $(E, \mathcal{E})$ a measurable space, and $(Y^i)_{i \in I}$ a family of functions from Ω into E. Then the following two properties hold:*

1. *For every $A \in \sigma\{Y^i;\ i \in I\}$ and every $\omega \in A$, if there is another $\omega' \in \Omega$ satisfying*
$$Y^i(\omega) = Y^i(\omega'), \qquad \forall\, i \in I,$$
then $\omega' \in A$.
2. *For any $A \in \sigma\{Y^i;\ i \in I\}$, there exists a countable subset $J \subset I$ such that*
$$A \in \sigma\{Y^j;\ j \in J\}.$$

Proof: 1. The proof of the first claim follows from the following fact. If we denote by $\mathcal{G}$ the family of elements $A \in \sigma\{Y^i;\ i \in I\}$ such that for every $\omega \in A$, if there is another $\omega' \in \Omega$ satisfying $Y^i(\omega) = Y^i(\omega')$ forall $i \in I$, then $\omega' \in A$, it is easy to show that $\mathcal{G}$ is a σ-field on its own, and by construction, $\mathcal{G} \subset \sigma\{Y^i;\ i \in I\}$. On the other hand, if i_0 is any element of I and B any element of $\mathcal{E}$, it is obvious that if $\omega \in \{Y^{i_0} \in B\}$, namely if $Y^{i_0}(\omega) \in B$, and if $Y^i(\omega) = Y^i(\omega')$ for all $i \in I$, it is also true for $i = i_0$ so that $Y^{i_0}(\omega') \in B$ which proves that $\{Y^{i_0} \in B\} \in \mathcal{G}$. So Y^{i_0} is $(\mathcal{G}, \mathcal{E})$ measurable, and $\mathcal{E} \subset \mathcal{G}$ since $\mathcal{E}$ is the smallest σ-field with this property.

2. The second claim means that
$$\mathcal{F} = \sigma\{Y^i;\ i \in I\} = \bigcup_{J \subseteq I,\, J \text{ countable}} \sigma\{Y^j;\ j \in J\}.$$

We prove this equality in the same way as above. Let denote by $\mathcal{G}$ the right hand side. It is obvious that $\mathcal{G} \subseteq \mathcal{F}$. To prove the other inclusion, we show that $\mathcal{G}$ is a $\sigma-$field, and since all the Y^i are $\mathcal{G}$-measurable, this will imply the desired inclusion. The only non-trivial property to show is the countable union property. So let $(A_n)_{n \geqslant 1}$ be a countable sequence of elements in $\mathcal{G}$. For each $n \geqslant 1$, there exists a countable subset $J_n \subset I$ such that $A_n \in \sigma\{Y^j;\ j \in J_n\}$. Let $J = \bigcup_n J_n$. Then J is also a countable subset of I. Moreover, for all $n \geqslant 1$
$$A_n \in \sigma\{Y^j;\ j \in J_n\} \subseteq \sigma\{Y^j;\ j \in J\},$$
which is a σ-field, and hence
$$\bigcup_n A_n \in \sigma\{Y^j;\ j \in J\} \subseteq \mathcal{G}$$
which is what we wanted to prove. □

4.5.3 Measurability of White Noises

We now return to the construction of white noises on product spaces as given by Kolmogorov's extension theorem.

Let $(I, \mathcal{I}, \lambda)$ be a probability space, E a polish space and let $\mu \in \mathcal{P}(E)$ be a probability measure on the Borel $\sigma-$field of E.

Kolmogorov's theorem says that (even if I is only a set without special structure) on the product space $(\Omega, \mathcal{F}) = (E^I, \mathcal{B}(E)^I = \bigotimes_{i \in I} \mathcal{B}(E))$, there exists a probability measure $\mathbb{P} = \mathbb{P}_\mu$ such that if we denote by X^i for $i \in I$ the i-th coordinate projection, namely
$$X^i : \Omega \ni \omega \mapsto \omega(i) \in E,$$

then for each $i \in I$, X^i is $(\mathcal{F}, \mathcal{B}(E))$-measurable, the law $\mathcal{L}_{\mathbb{P}}(X^i)$ of X^i, namely the push forward probability measure $\mathbb{P}_\mu \circ (X^i)^{-1}$ is equal to μ, in other words:

$$\mathcal{L}(X^i) = \mathbb{P}_\mu \circ (X^i)^{-1} = \mu,$$

and $(X^i)_{i \in I}$ are independent (i.e. any finite subset $X^{i_1}, \ldots, X^{i_k}$ are independent). The stochastic process $(X^i)_{i \in I}$ is called a *white noise* as it comprises independent and identically distributed random variables.

The following lemma, due to Doob (1953), indicates the inadequacy of the product $\sigma-$filed $\mathcal{F}$ for the analysis of white noises. We recall that the above σ-field $\mathcal{F}$ is generated by the cylinders.

Lemma 4.41 *Assume that $I = [0,1]$, $(E, \mathcal{E}) = (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, and for any function $h : [0,1] \to E$ we define the set:*

$$\mathfrak{M}_h = \{\omega \in \Omega;\ X^i(\omega) = h(i) \textit{ except for at most countably many } i \in I\}.$$

Then $\mathfrak{M}_h$ has $\mathbb{P}$-outer measure 1.

Proof: Let $A \in \mathcal{F} = \mathcal{B}(E)^{[0,1]}$. Then according to Proposition 4.40, A is determined by a countable subset $I_A \subseteq I$ in the sense that if $\omega \in \Omega$ and $\omega' \in \Omega$ satisfy $\omega(i) = \omega'(i)$ for all $i \in I_A$, then $\omega \in A$ if and only if $\omega' \in A$.

Now suppose that $\mathfrak{M}_h \subset A$, and pick any $\omega \in \Omega$. We construct $\omega' \in \Omega$ such that

$$\omega'(i) = \begin{cases} \omega(i) & i \in I_A \\ h(i) & i \notin I_A. \end{cases}$$

Since $\omega'(i) = h(i)$ except for countably many $i \in I$, we have $\omega' \in \mathfrak{M}_h$. Since $\mathfrak{M}_h \subset A$, we have $\omega' \in A$, and since $\omega'(i) = \omega(i)$ for all $i \in I_A$, then $\omega \in A$ from the property of the $\sigma-$field $\mathcal{F}$. This means that $A = \Omega$. Hence, the outer measure of $\mathfrak{M}_h$ takes the value

$$\mathbb{P}^*(\mathfrak{M}_h) = \inf_{\mathfrak{M}_h \subseteq A} \mathbb{P}(A) = \mathbb{P}(\Omega) = 1,$$

which completes the proof of the lemma. □

Remark 4.42 *The above lemma has striking consequences. Among them:*

- *The space of E-valued continuous functions $\mathcal{C}(I; E)$ is not in $\mathcal{F}$.*
- *A set $A \subseteq \Omega$ cannot belong to $\mathcal{F}$ unless there exists a countable set $I_A \subset I$ such that for any $\omega, \omega' \in \Omega$ satisfying $\omega(i) = \omega'(i)$ for all $i \in I_A$ we have $\omega \in A$ if and only if $\omega' \in A$.*
- *If we define a process of Brownian motion and a Poisson process as processes with independent increments with mean zero Gaussian and Poisson distributions respectively, and we construct them on the above product probability space using Kolmogorov's theorem, the sets of typical sample paths of these processes are not in $\mathcal{F}$, in other words, they are not measurable for this product σ-field.*

Remark 4.43 *Let us assume that $I = [0,1]$, $\mathcal{I} = \mathcal{B}(I)$ and λ is the Lebesgue measure on $[0,1]$.*

- *Let us assume that* $h : [0,1] \to \mathbb{R}$*, viewed as an element of* $\mathbb{R}^I$*, is not Lebesgue measurable (no measurable function is equals to* h λ*-almost everywhere), then all the elements of* $\mathfrak{M}_h$ *are non-Lebesgue measurable. Then if we denote by* $\mathcal{N}$ *the set of non-Lebesgue measurable functions, we have:*

$$\mathbb{P}^*(\mathcal{N}) = 1, \qquad \textit{so that} \qquad \mathbb{P}_*(\mathcal{N}^c) = 0.$$

Notice that $\mathcal{N}^c$ *is the set of Lebesgue measurable functions.*
- *Let* $h : [0,1] \to \mathbb{R}$ *be a Lebesgue measurable function, then all elements of* $\mathfrak{M}_h$ *are Lebesgue measurable, so that*

$$\mathbb{P}^*(\mathcal{N}^c) = 1.$$

Thus, we can conclude that $\mathcal{N}^c$ *is not* $\mathbb{P}-$*measurable.*

4.5.4 **Fubini Extensions**

The following material is in large part lifted from Section 3.7 of [11] which was intended to give a game oriented presentation of the results of [37].

The classical Glivenko-Cantelli form of the Law of Large Numbers (LLN) states that if F denotes the cumulative distribution function of a probability measure on $\mathbb{R}$, if $(X^n)_{n\geqslant 1}$ is an infinite sequence of independent identically distributed random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with common distribution μ, and if we use the notation:

$$F_\omega(x) = \limsup_{N\to\infty} \frac{1}{N} \#\big\{n \in \{1, \cdots, N\} : X^n(\omega) \leqslant x\big\}, \quad x \in \mathbb{R},\ \omega \in \Omega, \tag{4.28}$$

for the proportion of $X^n(\omega)$'s not greater than x, then this $\limsup$ is in fact a limit for all $x \in \mathbb{R}$ and $\mathbb{P}$-almost all $\omega \in \Omega$, and $\mathbb{P}[\{\omega \in \Omega : F_\omega(\cdot) = F\}] = 1$.

In this section, we use a fixed probability space $I, \mathcal{I}, \lambda)$ for the index set of families of random variables. Most often we shall use $I = [0,1]$, $\mathcal{I} = \mathcal{B}([0,1])$ its Borel σ-field, and the Lebesgue measure for λ. If E is a Polish space, for each probability measure $\mu \in \mathcal{P}(E)$, Kolmogorov's theorem can be used to construct on the product space $\Omega = E^I$ equipped with the product σ-field $\mathcal{F}$ of copies of the Borel σ-field of E, the product probability measure $\mathbb{P}$ for which the coordinate projections $(X^i : \Omega \ni \omega \mapsto X^i(\omega) = \omega(i) \in E)_{i\in I}$ become independent and identically distributed with common distribution μ on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. As we saw in the previous section, the sample paths $I \ni i \mapsto X^i(\omega) \in E$ are pretty rough functions since they are (for $\mathbb{P}$-almost $\omega \in \Omega$) nowhere continuous and not even measurable.

Hence, this construction of a continuum of independent identically distributed random variables leads to irregular structures lacking measurability properties. The following definition offers an alternative which keeps most of what is needed from the independence.

Definition 4.44 *If* E *is a Polish space, a family* $(X^i)_{i\in I}$ *of* E*-valued random variables is said to be essentially pairwise independent if, for* λ*-almost every* $i \in I$*, the random variable* X^i *is independent of* X^j *for* λ*-almost every* $j \in I$*. Accordingly, if these random variables are real valued and square integrable, we say that the family* $(X^i)_{i\in I}$ *is essentially pairwise uncorrelated if, for* λ*-almost every* $i \in I$*, the correlation coefficient of* X^i *with* X^j *is* 0 *for* λ*-almost every* $j \in I$*.*

One may wonder if essentially pairwise independent families $(X^i)_{i\in I}$ can be constructed on probability spaces $(\Omega, \mathcal{F}, \mathbb{P})$ so that the process $\boldsymbol{X} : I \times \Omega \ni (i,\omega) \mapsto X^i(\omega)$ satisfies relevant measurability properties. To do so, we shall construct such processes on extensions of the product space $(I\times\Omega, \mathcal{I}\otimes\mathcal{F}, \lambda\otimes\mathbb{P})$, which are called Fubini's extensions.

Definition 4.45 *If $\mathcal{I}\boxtimes\mathcal{F}$ is a σ-field containing $\mathcal{I}\otimes\mathcal{F}$ and $\lambda\boxtimes\mathbb{P}$ is a probability measure on $(I\times\Omega,\mathcal{I}\boxtimes\mathcal{F})$, then $(I\times\Omega,\mathcal{I}\boxtimes\mathcal{F},\lambda\boxtimes\mathbb{P})$ is said to be a Fubini extension of $(I\times\Omega,\mathcal{I}\otimes\mathcal{F},\lambda\otimes\mathbb{P})$ if, for every measurable and $\lambda\boxtimes\mathbb{P}$-integrable $X : I\times\Omega\ni(i,x)\mapsto X^i(\omega)\in\mathbb{R}$, we have:*

1. *for λ-a.e. $i\in I$, $\Omega\ni\omega\mapsto X^i(\omega)$ is a $\mathbb{P}$-integrable random variable, and for $\mathbb{P}$-a.e. $\omega\in\Omega$, $I\ni i\mapsto X^i(\omega)$ is measurable and λ-integrable;*
2. *$I\ni i\mapsto\int_\Omega X^i(\omega)d\mathbb{P}(\omega)$ is measurable and λ-integrable, and $\Omega\ni\omega\mapsto\int_I X^i(\omega)d\lambda(i)$ is a $\mathbb{P}$-integrable random variable, and:*

$$\begin{aligned}\int_I\left(\int_\Omega X^i(\omega)d\mathbb{P}(\omega)\right)d\lambda(i) &= \int_\Omega\left(\int_I X^i(\omega)d\lambda(i)\right)d\mathbb{P}(\omega)\\ &= \int_{I\times\Omega} X^i(\omega)d\big(\lambda\boxtimes\mathbb{P}\big)(i,\omega).\end{aligned}\tag{4.29}$$

In the sequel, we shall use the standard symbol $\mathbb{E}$ for denoting the expectation under the sole probability $\mathbb{P}$.

Measurable essentially pairwise independent processes $\boldsymbol{X}$ are first constructed in such a way that, for each $i\in I$, the law of X^i is the uniform distribution on the unit interval $[0,1]$. Then, transforming a uniform random variable into a random variable with a given distribution using its quantile function, we easily construct measurable essentially pairwise independent Euclidean-valued processes with any given prescribed marginals. So the actual problem is to construct rich product probability spaces in the sense of the following definition.

Definition 4.46 *A Fubini extension $(I\times\Omega,\mathcal{I}\boxtimes\mathcal{F},\lambda\boxtimes\mathbb{P})$ is said to be rich if there exists a real valued $\mathcal{I}\boxtimes\mathcal{F}$-measurable essentially pairwise independent process $\boldsymbol{X}$ such that the law of X^i is the uniform distribution on $[0,1]$ for every $i\in I$.*

We refer to the Notes & Complements at the end of the chapter for references to papers giving the construction of essentially pairwise independent measurable processes on Fubini extensions.

The following gives a simple property of rich Fubini extensions.

Lemma 4.47 *If the Fubini extension $(I\times\Omega,\mathcal{I}\boxtimes\mathcal{F},\lambda\boxtimes\mathbb{P})$ is rich, then λ is necessarily atomless.*

Proof: We shall argue by contradiction. If $A\in\mathcal{I}$, with $\lambda(A)>0$, is an atom of $(I,\mathcal{I},\lambda)$, then, for $\mathbb{P}$-a.e. $\omega\in\Omega$, the function $I\ni i\mapsto X^i(\omega)$ is λ-a.e. constant on A. So for $\mathbb{P}$-a.e. $\omega\in\Omega$ and λ-a.e. $i\in A$,

$$X^i(\omega)=\int_A X^j(\omega)\frac{d\lambda(j)}{\lambda(A)},$$

and using the Fubini property (4.29), we deduce that for λ-a.e. $i \in A$, the random variable $\Omega \ni \omega \mapsto X^i(\omega)$ is $\mathbb{P}$-a.e. equal to the random variable $\theta : \Omega \ni \omega \mapsto \int_A X^j(\omega) d\lambda(j)/\lambda(A)$. Also, for any event $B \in \mathcal{F}$,

$$\begin{aligned}\mathbb{P}[\theta \in B] &= \lambda \boxtimes \mathbb{P}\big[(i,\omega) \in I \times \Omega : X^i(\omega) \in B\big] \\ &= \int_I \mathbb{P}[X^i \in B] d\lambda(i) = \mathrm{Leb}_1(B),\end{aligned}$$

proving that θ, as a random variable on $(\Omega, \mathcal{F}, \mathbb{P})$, has the uniform distribution. In particular, $\mathbb{E}[\theta^2] = 1/3$.

On the other hand, we know that, for almost every $i \in I$, the function $I \times \Omega \ni (j,\omega) \mapsto X^i(\omega)X^j(\omega)$ is $\mathcal{I} \boxtimes \mathcal{F}$-measurable. Also, by the Fubini property, the function $I \ni j \mapsto \mathbb{E}[X^i X^j]$ is integrable with respect to λ and

$$\int_I \mathbb{E}[X^i X^j] d\lambda(j) = \mathbb{E}[X^i \theta]. \tag{4.30}$$

Now, we observe that the function $I \times \Omega \ni (i,\omega) \mapsto X^i(\omega)\theta(\omega)$ is also $\mathcal{I} \boxtimes \mathcal{F}$-measurable. Hence, $I \ni i \mapsto \mathbb{E}[X^i \theta]$ is integrable with respect to λ and

$$\int_I \mathbb{E}[X^i \theta] d\lambda(i) = \mathbb{E}[\theta^2] = \frac{1}{3}.$$

The contradiction comes from the fact that, for almost every $i \in I$, X^i is independent to X^j for almost every $j \in I$. In other words, the left-hand side in (4.30) is equal to:

$$\int_I \mathbb{E}[X^i X^j] d\lambda(j) = \frac{1}{\lambda A} \int_A \mathbb{E}[X^i]\mathbb{E}[X^j] d\lambda(j) = \frac{1}{4},$$

which gives the desired contradiction. □

We recall the following lemma from real analysis. A proof can be found in [11, Lemma 5.29].

Lemma 4.48 *Let E be a Polish space, $\exists \varphi : [0,1] \times \mathcal{P}(E) \to E$ measurable such that $\forall \nu \in \mathcal{P}(E)$*

$$Leb \circ \psi(\cdot,\nu)^{-1} = \nu.$$

Given the result of the above lemma, we can prove the following crucial property of rich Fubini extensions.

Proposition 4.49 *If the Fubini extension $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes \mathbb{P})$ is rich, if E is a Polish space, and if $\mu : I \mapsto \mathcal{P}(E)$ is $\mathcal{I}$-measurable, then there exists a $\mathcal{I} \boxtimes \mathcal{F}$-measurable E-valued essentially pairwise independent process $\boldsymbol{Y} : I \times \Omega \mapsto E$ such that for λ-a.e. $i \in I$, $\mathbb{P} \circ Y_i^{-1} = \mu_i$.*

Proof: Define $Y(i,\omega) = \psi(X^i(\omega), \mu_i)$ where $X = (X^i)_{i \in I}$ is an essential white noise with marginal distribution the uniform distribution over $[0,1]$. For $i \in I$ fixed, the mapping $\omega \mapsto X^i(\omega)$ has a uniform distribution on $[0,1]$, i.e. $\mathcal{L}(X^i) = Leb([0,1])$. Since $\psi(\cdot,\mu_i)$ takes the uniform distribution on $[0,1]$ into μ_i, we conclude that Y^i has distribution μ_i.

Since for $\lambda-$a.e. $i \in I$, X^i is independent of X^j for $\lambda-$a.e. $j \in I$, this implies that for $\lambda-$a.e. $i \in I$, $\psi(X^i, \mu_i)$ is independent of $\psi(X^j, \mu_j)$ for $\lambda-$a.e. $j \in I$, which shows that $\boldsymbol{Y}$ is a white noise with the desired properties. □

Remark 4.50 *It is very important to realize that Fubini extensions do not provide a panacea to all the measurability problems of continuous white noise. It is a fact that, if* $\boldsymbol{X} = (X^i)_{i\in I}$ *is essentially pairwise independent with values in a Polish space on a Fubini extension* $(I\times\Omega, \mathcal{I}\boxtimes\mathcal{F}, \lambda\boxtimes\mathbb{P})$ *and if* $I = [0,1]$ *and* $(I, \mathcal{I}, \lambda)$ *extends* $([0,1], \mathcal{B}_{[0,1]}, \lambda)$ *then the set of* $\omega \in \Omega$ *for which the sample function* $I \ni i \mapsto X^i(\omega)$ *is Lebesgue measurable has* $\mathbb{P}$*-probability* 0*.*

4.5.5 **The Exact Law of Large Numbers**

An exact law of large numbers can be proven on Fubini's extensions. In a weak form, this law can be given in the following wasy.

Theorem 4.51 *Let* $\boldsymbol{X} = (X^i)_{i\in I}$ *be a measurable square integrable process on a Fubini extension* $(I\times\Omega, \mathcal{I}\boxtimes\mathcal{F}, \lambda\boxtimes\mathbb{P})$*. The following are equivalent:*

(i) The random variable $(X^i)_{i\in I}$ *are essentially pairwise uncorrelated;*

(ii) For every $A \in \mathcal{I}$ *with* $\lambda(A) > 0$*, one has for* $\mathbb{P}$*-almost surely in* $\omega \in \Omega$*:*

$$\int_A X^i(\omega) d\lambda(i) = \int_A \mathbb{E}[X^i] d\lambda(i).$$

Proof: First Step: We first check that if $\boldsymbol{Y} = (Y^i)_{i\in I}$ and $\boldsymbol{Z} = (Z^i)_{i\in I}$ are measurable and square integrable processes on the Fubini extension $(I\times\Omega, \mathcal{I}\boxtimes\mathcal{F}, \lambda\boxtimes\mathbb{P})$, and if we set $\tilde{X}^{i,j}(\omega) = Y^i(\omega)Z^j(\omega)$ for $i,j \in I$ and $\omega \in \Omega$, then $\Omega \ni \omega \mapsto \tilde{X}^{i,j}$ is $\mathbb{P}$-integrable for λ-a.e. $i \in I$ and $j \in I$. Now, proceeding as in the proof of Lemma 4.47 and using the Fubini property of the space, we easily check that, for λ-a.e. $i \in I$, the function $I \ni j \mapsto \mathbb{E}[\tilde{X}^{i,j}]$ is λ-integrable, that the function $I \ni i \mapsto \int_I \mathbb{E}[\tilde{X}^{i,j}]d\lambda(j) = \mathbb{E}[Y^i \int_I Z^j d\lambda(j)]$ is λ-integrable, that the function $\Omega \ni \omega \mapsto (\int_I Y^i(\omega)d\lambda(i))(\int_I Z^j(\omega)d\lambda(j))$ is $\mathbb{P}$-integrable and that:

$$\mathbb{E}\Big[\Big(\int_I Y^i(\omega)d\lambda(i)\Big)\Big(\int_I Z^j(\omega)d\lambda(j)\Big)\Big] = \int_I\Big(\int_I \mathbb{E}[\tilde{X}^{i,j}]d\lambda(i)\Big)d\lambda(j). \tag{4.31}$$

Second Step: Let $A, B \in \mathcal{I}$, and let us define the processes $\boldsymbol{Y} = (Y^i)_{i\in I}$ and $\boldsymbol{Z} = (Z^i)_{i\in I}$ by $(Y^i = \mathbf{1}_A(i)(X^i - \mathbb{E}[X^i]))_{i\in I}$ and $(Z^i = \mathbf{1}_B(i)(X^i - \mathbb{E}[X^i]))_{i\in I}$ respectively. Applying (4.31) from the first step we get:

$$\begin{aligned}&\int_A\int_B \mathbb{E}\Big[\Big(X^i - \mathbb{E}[X^i]\Big)\Big(X^j - \mathbb{E}[X^j]\Big)\Big]d\lambda(i)d\lambda(j)\\ &\qquad\qquad = \mathbb{E}\Big[\int_A\Big(X^i - \mathbb{E}[X^i]\Big)d\lambda(i)\int_B\Big(X^j - \mathbb{E}[X^j]\Big)d\lambda(j)\Big],\end{aligned} \tag{4.32}$$

and the implication *(i)* ⇒ *(ii)* follows by taking $B = A$. On the other hand, if we assume that *(ii)* holds, equation (4.32) implies that:

$$\int_A\int_B \mathbb{E}\Big[\Big(X^i - \mathbb{E}[X^i]\Big)\Big(X^j - \mathbb{E}[X^j]\Big)\Big]d\lambda(i)d\lambda(j) = 0$$

for all $A, B \in \mathcal{I}$. The set $A \in \mathcal{I}$ being arbitrary, we conclude that:

$$\int_B \mathbb{E}\Big[\Big(X^i - \mathbb{E}[X^i]\Big)\Big(X^j - \mathbb{E}[X^j]\Big)\Big]d\lambda(j) = 0,$$

for λ-a.e. $i \in I$. So for λ-a.e. $i \in I$, $B \in \mathcal{I}$ being arbitrary, we conclude that:

$$\mathbb{E}\Big[\Big(X^i - \mathbb{E}[X^i]\Big)\Big(X^j - \mathbb{E}[X^j]\Big)\Big] = 0$$

for λ-a.e. $j \in I$ which completes the proof. □

Theorem 4.51 provides a form of the weak law of large numbers for essentially pairwise uncorrelated uncountable families of random variables. Here is a stronger form for essentially pairwise independent families of random variables.

Theorem 4.52 *Let E be a Polish space and $\boldsymbol{X} = (X^i)_{i\in I}$ be a measurable E-valued process on a Fubini extension $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes \mathbb{P})$ such that the random variables $(X^i)_{i\in I}$ are essentially pairwise independent. Then, for $\mathbb{P}$ almost every $\omega \in \Omega$ and for any B in the Borel σ-field $\mathcal{B}(E)$,*

$$\lambda\big[\{i \in I : X^i(\omega) \in B\}\big] = \int_I \mathbb{P}\big[X^i \in B\big] d\lambda(i).$$

Of course, we may choose E as a Euclidean space, in which case we get a strong form of the exact law of large numbers for essentially pairwise independent families of random variables with values in $\mathbb{R}^d$, for some $d \geqslant 1$. By choosing E as a functional space, the same holds true for a continuum of essentially pairwise independent random processes.

Finally, we can also derive conditional versions of these exact laws. We do not give the details here because we want to keep the presentation to a rather non-technical level since our motivation is merely to connect our approach to mean field games to the existing literature on games with a continuum of players. The interested reader is referred to the Notes & Complements at the end of the chapter for references.

4.5.6 More Unexpected Behavior and finally, a Positive Existence Result

In this final section, we assume that the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is the product probability space constructed via Kolmogorov's existence theorem from a Polish space E equipped with its Borel σ-field $\mathcal{B}_E$ and a fixed probability measure μ, the index set of the product being the unit interval $I = [0, 1]$. Recall that using a notation introduced earlier, $\mathbb{P} = \mathbb{P}_\mu$.

First we address the following natural question: *Can the product probability space $(\Omega, \mathcal{F}, \mathbb{P}_\mu)$ be a component of a Fubini extension together with the unit interval $I = [0, 1]$?* The following result gives a negative answer to this question.

Proposition 4.53 *For $I = [0, 1]$, there is no atomless probability space $(I, \mathcal{I}, \lambda)$ satisfying*

- *$(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes \mathbb{P})$ is a Fubini extension with marginals $(I, \mathcal{I}, \lambda)$ and $(\Omega, \mathcal{F}, \mathbb{P})$ with $\mathbb{P} = \mathbb{P}_\mu$.*
- *$\boldsymbol{X} = (X^i)_{i\in I}$ is $\mathcal{I} \boxtimes \mathcal{F}$ measurable.*

Proof: Suppose that such a Fubini extension exists and try to reach a contradiction. Take $E = \mathbb{R}$, B a non-empty interval in E such that $\mu(B) < 1$ and $c \in B$. We define two functions $h : I \ni i \mapsto h(i) = c \in E$ and $g = \mathbf{1}_B \circ \boldsymbol{X} : I \times \Omega \to E$, so that for every $(i, \omega) \in I \times \Omega$,

$$g(i, \omega) = \begin{cases} 1 \; if \quad X^i(\omega) \in B \\ 0 \; otherwise. \end{cases}$$

We prove that $\mathbb{P}(\mathfrak{M}_h) = 0$, contradicting the result of Lemma 4.41.

Notice that g is bounded and $\mathcal{I} \boxtimes \mathcal{F}$ measurable since $\boldsymbol{X}$ is, and $(g(i, \cdot))_{i \in I}$ inherits the independence properties of $\boldsymbol{X}$. So by the exact law of large numbers,

$$\lambda \circ g_\omega^{-1} = (\lambda \boxtimes \mathbb{P}) \circ g^{-1}$$

for $\mathbb{P}$-almost every $\omega \in \Omega$, where $g_\omega = g(\cdot, \omega)$ for $\omega \in \Omega$. Evaluating both sides on B, we get

$$\lambda(g_\omega^{-1}(B)) = (\lambda \boxtimes \mathbb{P})(g^{-1}(B))$$

for $\mathbb{P}$-almost every $\omega \in \Omega$, or equivalently

$$\lambda(\{i \in I;\ X^i(\omega) \in B\}) = \int \lambda(di) g(i, \omega) = \int \lambda(di) \mathbb{P}(d\omega) g(i, \omega) = \int \lambda(di) \mu(B) = \mu(B) < 1.$$

So for $\mathbb{P}$-a.e. $\omega \in \Omega$, $X^{\cdot}(\omega) = \omega(\cdot) \notin \mathfrak{M}_h$. This is because if $\omega \in \Omega$ such that $X^{\cdot}\imath \mathcal{M}_h$, then $X^i(\omega) = h(i) = c$ except for countably many $i \in I$, hence for λ-almodt every $i \in I$, and since $h \equiv c$ and $c \in B$, we conclude that $\lambda(X(\omega)^{-1}(B)) = 1$. $X^i(\omega) \in B$. If that were to happen except for countably many i, then $\lambda-$measure of those i's would be 1 since λ is atomless. Hence $\mathbb{P}(\mathfrak{M}_h) = 0$ which gives the desired contradiction. ▫

Still, the following result comes to *save the day.*

Proposition 4.54 *There exist an atomless probability space $(I, \tilde{\mathcal{I}}, \tilde{\lambda})$ and a probability space $(\Omega, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ extending $(\Omega, \mathcal{F}, \mathbb{P}_\mu)$such that*

(i) There is a rich Fubini extension $(I \times \Omega, \tilde{\mathcal{I}} \boxtimes \tilde{\mathcal{F}}, \tilde{\lambda} \boxtimes \tilde{\mathbb{P}})$;

(ii) The coordinate process $\boldsymbol{X} = (X^i)_{i \in I}$ is $\mathcal{I} \boxtimes \tilde{\mathcal{F}}$-measurable.

5

Signals and Correlated Equilibria

Summary. We first review the abstract set-up of the search for mixed Nash equilibrium and we extend this notion to correlated equilibrium. This generalized form of equilibrium seems to have been introduced by R. Aumann in [**?**].

5.1 Mixed Strategies, Signals, and Correlated Equilibrium

As in the previous chapters, we consider static games first.

5.1.1 Correlated Equilibria

Consider N players and let us denote their set by $[N] = \{1, \ldots, N\}$. Let A^i be the set of actions available to player $i \in [N]$, $\mathcal{A}^i$ be a σ-field on A^i and let $A = A^1 \times \cdots \times A^N$ be the set of strategy profiles and $\mathcal{A} = \mathcal{A}^1 \otimes \cdots \otimes \mathcal{A}^N$ the product $\sigma-$field on A. Next we consider a cost function for each player $J : A \to \mathbb{R}^N$ where $J = (J^i)_{i=1,\ldots,N}$.

We denote by $\mathcal{G} = ([N], A, J)$ the game defined by the above quantities.

Let $\Sigma^i = \mathcal{P}(A^i)$ be the set of mixed strategies for player i. This is the set of probability measures on the measurable space $(A^i, \mathcal{A}^i)$. If A^i is a finite set, the $\Sigma^i = \Delta(A^i)$ is just a simplex. Let $\Sigma = \Sigma^1 \times \cdots \times \Sigma^N$ be the set of mixed strategy profiles. If $\pi \in \Sigma$, then there exists $\pi^1, \ldots, \pi^N$ with $\pi^i \in \mathcal{P}(A^i)$ such that $\pi = (\pi^1, \ldots, \pi^N)$. Since they are in one-to-one correspondence, we identify the mixed strategy profiles $\pi = (\pi^1, \ldots, \pi^N)$ with the product measures $\pi^1 \otimes \ldots \otimes \pi^N \in \mathcal{P}(A)$. Thus, Σ can be viewed as the set of product probability measures on $(A, \mathcal{A})$ and it is a subset of the set of all the probability distributions on $(A, \mathcal{A})$, namely $\Sigma \subset \mathcal{P}(A)$.

Remark 5.1 *It is very important to emphasize that for every mixed strategy profile $\pi \in \Sigma$, the individual strategies of the N players are sampled independently according to their marginal distributions $(\pi^i)_{i=1,\ldots,N}$. However, for any $\pi \in \mathcal{P}(A)$, one can imagine a (random) strategy profile $(\alpha^1, \ldots, \alpha^N) \sim \pi$ for which the individual actions $(\alpha^i)_{i \in N}$ are not sampled independently according to the marginal distribution $\pi^i = \int_{A^{-i}} \pi(\cdot, d\boldsymbol{\alpha}^{-i})$ on $(A^i, \mathcal{A}^i)$, but instead, sampled jointly according to the joint distribution π. This last point is at the root of the notion of correlated equilibrium which we study next.*

To emphasize the specific nature of a correlated equilibrium, we introduce the notion in parallel with the notion of mixed equilibrium in order to highlight the differences between the two notions. For this reason, we recall the definition and the notations used when we first introduced mixed equilibria. As always when we use mixed strategies, we actually work with the extended game introduced early. In order to avoid using too many tildes ˜ complicating the notation, we redefine the components of the extended game from scratch. We extend the cost function J form A to the set $\Sigma \subset \mathcal{P}(A)$ in a natural way. More precisely, we consider $J : \Sigma \to \mathbb{R}^N$ such that for every $i \in [N]$,

$$J^i(\pi) = \int J^i(\alpha^1, \ldots, \alpha^N)\pi^1(d\alpha^1)\ldots\pi^N(d\alpha^N). \tag{5.1}$$

In other words, $J^i(\pi)$ is the integral of the function $\boldsymbol{\alpha} \mapsto J^i(\boldsymbol{\alpha})$ with respect to the product measure $\pi^1 \otimes \ldots \otimes \pi^N$. If A^i are finite for all $i \in [N]$, then the above extension of the cost function reads

$$J^i(\pi) = \sum_{\boldsymbol{\alpha}=(\alpha^1,\ldots,\alpha^N)\in A} J^i(\alpha^1, \ldots, \alpha^N)\prod_{j=1}^{N}\pi^j(\alpha^j).$$

Definition 5.2 *An element $\pi^i \in \Sigma^i$ is called a mixed strategy for player $i \in [N]$. An element $\pi \in \Sigma$ is called a mixed strategy profile. We define the extended game with mixed strategy by $\overline{\mathcal{G}} = ([N], \Sigma, J)$.*
A mixed strategy profile $\pi^* = (\pi^{1,*}, \ldots, \pi^{N,*})$ is said to be a mixed Nash equilibrium for the original game $\mathcal{G} = ([N], A, J)$ if it is a Nash equilibrium for the extended game $\overline{\mathcal{G}} = ([N], \Sigma, J)$.

We now introduce the important concept of correlated equilibrium.

Definition 5.3 *A probability measure $\pi \in \mathcal{P}(A)$ (which is not necessarily a product distribution) is said to be a correlated equilibrium (cNE) if for π^i-a.e. $\alpha^i \in A^i$ and π^i-a.e. $\tilde{\alpha}^i \in A^i$, we have*

$$\int_{A^{-i}} J^i(\alpha^i, \boldsymbol{\alpha}^{-i})\pi(d\boldsymbol{\alpha}^{-i}|\alpha^i) \leqslant \int_{A^{-i}} J^i(\tilde{\alpha}^i, \boldsymbol{\alpha}^{-i})\pi(d\boldsymbol{\alpha}^{-i}|\alpha^i). \tag{5.2}$$

where π^i is the marginal distribution of π on the space $(A^i, \mathcal{A}^i)$.

In words, for a particular player, say $i \in [N]$, and for a given signal/recommendation, say $\alpha^i \in A^i$, then conditional on this knowledge, the signal α^i is the best response of player i against the other players playing according to the joint probability distribution $\pi \in \mathcal{P}(A)$.

Remark 5.4 *If the sets A^i are finite for all $i \in [N]$, then for every function φ on A and $\alpha^i \in A^i$,*

$$\int_{A^{-i}} \varphi(\boldsymbol{\alpha})\pi(d\boldsymbol{\alpha}^{-i}|\alpha^i) = \sum_{(\alpha^1,\ldots,\alpha^{i-1},\alpha^{i+1},\ldots,\alpha^N)} \varphi(\alpha^1, \ldots, \alpha^N)\frac{\pi(\alpha^1, \ldots, \alpha^N)}{\sum_{\boldsymbol{\alpha}^{-i}}\pi(\boldsymbol{\alpha})}.$$

So the above condition for (cNE) rewrites

$$\sum_{\boldsymbol{\alpha}^{-i}} J(\alpha^i, \boldsymbol{\alpha}^{-i})\pi(\alpha^i, \boldsymbol{\alpha}^{-i}) \leqslant \sum_{\boldsymbol{\alpha}^{-i}} J(\tilde{\alpha}^i, \boldsymbol{\alpha}^{-i})\pi(\alpha^i, \boldsymbol{\alpha}^{-i}).$$

Given the above definition, the following proposition is straightforward. We state it without proof because of its importance.

Proposition 5.5

1. *If $(\pi^1, \ldots, \pi^N)$ is a mixed NE, then $\pi = \pi^1 \otimes \cdots \otimes \pi^N$ is a correlated equilibrium.*
2. *If $(\pi_k)_{k=1,\ldots,K} \in (\mathcal{P}(A))^K$ are correlated equilibria and if $(p_k)_{k=1,\ldots,K} \in [0,1]^K$ is such that $\sum_{k=1}^K p_k = 1$, then $\pi = \sum_k p_k \pi_k$ is also a correlated equilibrium. In other words the set (cNE) of correlated Nash equilibria is a convex set.*

Remark 5.6 *Notice that it is NOT true in general that a convex combination of mixed Nash equilibria is a mixed Nash equilibrium.*

Remark 5.7 *While the above sobering remark says that the set of values of the mixed Nash equilibria may not be convex, the structure of its convex hull is not too complex. Indeed, we can extend the definition* (5.1) *of the cost function to accommodate general correlated random strategies. Using the same notation:*

$$J^i(\pi) = \int J^i(\alpha^1, \ldots, \alpha^N)\pi(d\alpha^1, \ldots, d\alpha^N). \tag{5.3}$$

for $\pi \in \mathcal{P}(A)$ defining in this way a cost function $J : \mathcal{P}(A) \to \mathbb{R}^N$. The set of values of the mixed Nash equilibria is given by:

$$J_{mNE} = \{J(\pi), \pi \in \Sigma \text{ mixed Nash equilibrium}\} \subset \mathbb{R}^N.$$

We denote by coJ_{mNE} its convex hull. By Carathéodory theorem, for any $J \in coJ_{mNE}$, there exists at most $N+1$ mixed Nash equilibria $\pi_1, \ldots, \pi_{N+1}$ and $N+1$ non-negative real values $p_1, \ldots, p_{N+1} \in [0,1]$ with $\sum_{k=1}^{N+1} p_k = 1$ such that

$$J = \sum_{k=1}^{N+1} p_k J(\pi_k).$$

We shall use this fact later on in these lectures.

The numerical computation of correlated equilibria is a difficult challenge, even when the action spaces A^i are finite. For practical algorithms for the computation of correlated equilibria in polynomial time, the reader is referred to [**?**].

5.1.2 **Signals**

All the random quantities used in this section are assumed to be defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The idea behind the discussion of this section is to sample from the state of the world $\omega \in \Omega$, and then provide a signal or a recommendation to each player based on that sample. Let ξ be an A-valued random variable on $(\Omega, \mathcal{F}, \mathbb{P})$, that is $\xi : \Omega \to A$ is $(\mathcal{F}, \mathcal{A})$-measurable. Since $A = A^1 \times \ldots \times A^N$, we can write $\xi = (\xi^1, \ldots, \xi^N)$ in which case $\xi^i \in A^i$ is interpreted as the signal or recommendation for player i.

Let $\pi = \mathcal{L}(\xi) = \mathbb{P} \circ \xi^{-1} \in \mathcal{P}(A)$, i.e. $\pi(B) = \mathbb{P}(\xi \in B)$ for all $B \in \mathcal{A}$ be the push forward of the probability $\mathbb{P}$, in other words, the law of the random quantity ξ. We notice that π is not necessarily a product measure, in other words, the N signals are not necessarily independent since they can depend upon each other.

In general, for every $i \in [N]$, we denote by $\pi^i = \mathcal{L}(\xi^i) \in \mathcal{P}(A^i)$ the marginal law of ξ^i on $(A^i, \mathcal{A}^i)$, namely the push forward of $\mathbb{P}$ by ξ^i, and by $\pi(\cdot|\xi^i = \alpha^i) \in \mathcal{P}(A^{-i})$ the conditional law of ξ^{-i} on $(A^{-i}, \mathcal{A}^{-i})$ given that $\xi^i = \alpha^i$. So for any measurable and bounded function f on A we have:

$$\begin{aligned}\int f(\boldsymbol{\alpha})\pi(d\boldsymbol{\alpha}) &= \mathbb{E}[f(\xi^1, \dots, \xi^N)] = \mathbb{E}\left[\mathbb{E}[f(\xi^1, \dots, \xi^N)|\xi^i]\right] \\ &= \int_{A^i} \Big(\int_{A^{-i}} f(\alpha^i, \boldsymbol{\alpha}^{-i})\pi(d\boldsymbol{\alpha}^{-i}|\xi^i = \alpha^i) \Big) \pi^i(d\alpha^i)\end{aligned}$$

Accordingly, saying that the law of ξ is a correlated Nash Equilibrium (cNE) is equivalent to saying that $\forall i \in [N], \forall \tilde{\alpha}^i \in A^i$,

$$\mathbb{E}\left[J^i(\xi^i, \xi^{-i})|\xi^i\right] \leqslant \mathbb{E}\left[J^i(\tilde{\alpha}^i, \xi^{-i})|\xi^i\right]. \qquad \mathbb{P} - a.s.,$$

or in other words, for every $\tilde{\alpha}^i \in A^i$,

$$\mathbb{E}\left[J^i(\alpha^i, \xi^{-i})|\xi^i = \alpha^i\right] \leqslant \mathbb{E}\left[J^i(\tilde{\alpha}^i, \xi^{-i})|\xi^i = \alpha^i\right], \qquad \pi^i - a.e \quad \alpha^i \in A^i.$$

Definition 5.8 (Information structure) *An information structure is a tuple*

$$(\Omega, (\mathcal{F}^i)_{i \in [N]}, (\mathbb{P}^i)_{i \in [N]}),$$

where Ω is the state of the world, and for each $i \in [N]$, $\mathcal{F}^i$ is the set of events known to player i, and $\mathbb{P}^i$ is a probability measure on the measurable space $(\Omega, \mathcal{F}^i)$.

The measure $\mathbb{P}^i$ is interpreted as the probability measure of player i. For example, given $(\Omega, \mathcal{F}, \mathbb{P})$ and a signal $\xi = (\xi^1, \dots, \xi^N)$, we can choose $\mathcal{F} \supset \mathcal{F}^i = \sigma\{\xi^i\} = \{\{\xi^i \in B\}; B \in \mathcal{A}^i\}$ for the set of events which depend only upon the signal ξ^i of player i, and $\mathbb{P}^i = \mathbb{P}_{|\mathcal{F}^i}$ the restriction of the probability measure $\mathbb{P}$ to this class of events.

5.1.3 **Examples**

Example 1.

We first consider the classical game model known under the name of *Battle of the Sexes*. A couple, wife and husband, have to choose between going to the opera or a soccer game. The rewards for the different choices are given in Table 5.1. Uncharacteristically given how we proceeded so far, we use rewards instead of costs, so in this example, the players (husband and wife) try to *maximize* their rewards.

- In this game, there are two Nash equilibria in pure strategies: $\alpha^1 = \alpha^2 =$ Opera, and $\alpha^1 = \alpha^2 =$ Soccer.
- If one uses mixed strategies given by probability distributions as in Table 5.2, one can see that there is also one (non-pure) Nash equilibrium in mixed strategies: wife chooses opera with probability $p = 4/5$ and husband chooses opera with probability $q = 1/5$.

Wife \ Husband	Opera	Soccer
Opera	$(4,1)$	$(0,0)$
Soccer	$(0,0)$	$(1,4)$

Table 5.1: *Battle of sexes: reward table*

Wife \ Husband	Opera	Soccer
Opera	pq	$p(1-q)$
Soccer	$(1-p)q$	$(1-p)(1-q)$

Table 5.2: *Battle of sexes: typical probability distribution $\pi \in \Sigma$ giving a mixed strategy profile. Here $\pi_{wife}(Opera) = p$ and $\pi_{husband}(Opera) = q$.*

- However, one can see by inspection that there are many correlated equilibria, case in point:
$$\pi_{op,op} = 0.5 = \pi_{soc,soc}, \qquad \pi_{op,soc} = \pi_{soc,op} = 0,$$
is one of them.

Example 2.

We now consider the classical game model known as the *traffic light dilemma*. The costs to the two players (driver of car 1 and driver of car 2) are given in Table 5.3 below. While the possible actions are stated as **Stop** and **Go** in this table, we shall abbreviate them to S and G in the discussion and the computations below. Note that in this example, both drivers want to minimize their cost.

Car 1 \ Car 2	Stop	Go
Stop	$(0,0)$	$(0,-1)$
Go	$(-1,0)$	$(100,100)$

Table 5.3: *Traffic light game: costs.*

- **Pure Nash equilibria:**
  - $(\alpha^1, \alpha^2) = (G, S)$ is a pure Nash equilibrium. Indeed, $J^1(G,S) = -1 < 0 = J^1(S,S)$, and $J^2(G,S) = 0 < 100 = J^2(G,G)$.
  - Similarly, $(\alpha^1, \alpha^2) = (S, G)$ is also a Nash equilibrium in pure strategies.
- **Mixed Nash equilibria:** to compute them all, let us denote by $\pi = (\pi^1, \pi^2)$ a generic strategy profile in mixed strategies, and let $p = \pi^1(S)$ and $q = \pi^2(S)$ so that $1 - p = \pi^1(G)$, and $1 - q = \pi^1(G)$. The cost to player 1 is

$$
\begin{aligned}
J^1(\pi) &= J^1(S,S)\pi^1(S)\pi^2(S) + J^1(S,G)\pi^1(S)\pi^2(G) \\
&\qquad\qquad + J^1(G,S)\pi^1(G)\pi^2(S) + J^1(G,G)\pi^1(G)\pi^2(G) \\
&= -q(1-p) + 100(1-p)(1-q) \\
&= (1-p)(100-101q)
\end{aligned}
$$

and similarly the cost for the second player is

$$J^2(\pi) = (1-q)(100-101p).$$

We first compute the best response π^{1*} of the first player to the mixed strategy π^2 of the second player. We identify the search for π^{1*} to the search for a probability p^* and we use the fact that π^2 is given by the probability q:

– if $q < 100/101$, $J^1(\pi)$ is minimum when $p^*(q) = 1$;
– if $q = 100/101$, $p^*(q)$ can be anything;
– if $q > 100/101$, $J^1(\pi)$ is minimum when $p^*(q) = 0$;

Similarly, we determine the best response q^* of the second player when the first player acts with probability p:

– if $p < 100/101$, then $q^*(p) = 1$;
– if $p = 100/101$, then $q^*(p)$ can be anything;
– if $p > 100/101$, then $q^*(p) = 0$.

By definition the equilibria are the mixed strategies π given by couples $(p,q) \in [0,1]^2$ such that $p^*(q) = p$ and $q^*(p) = q$. We find

– $(p,q) = (0,1)$, which is a pure Nash equilibrium, hence a mixed Nash equilibrium as well;
– $(p,q) = (1,0)$ which is a pure Nash equilibrium, hence a mixed Nash equilibrium as well;
– $(p,q) = (100/101, 100/101)$: indeed if $p = 100/101$, then $q^*(p)$ can a priori be anything. But due to the fixed point condition, we have necessarily $q = q^*(p) = 100/101$. Otherwise $p = p^*(q) = p^*(q^*(p)) \in \{1,0\}$, contradiction. This proves that there is a unique Nash equilibrium in non-pure mixed strategies.

To complete the analysis of these equilibria, we compute the **social cost**:

$$J(\pi) = \sum_{(a,b)\in\{S,G\}^2} (J^1(a,b) + J^2(a,b))\pi_{a,b}.$$

For the Nash equilibria in pure strategies, we find $J = -1$ in both cases. However, for the Nash equilibrium in mixed strategies we find:

$$J = \frac{100}{101^2} - \frac{100}{101^2} + \frac{200}{101^2} = 0.$$

So the latter is worse than the NE in pure strategies.

- **Correlated equilibria:** We now turn our attention to correlated equilibria. Such an equilibrium, if it exists, is given by a strategy profile which is a probability distribution π over $\{S,G\}^2$ which is not necessarily the product of two probability distributions over $\{S,G\}$. We shall use the fact that for any such π, we have $\pi_{S,S} = 1 - \pi_{G,S} - \pi_{S,G} - \pi_{G,G}$. Recall that the definition of a correlated equilibrium says that for player $i \in \{1,2\}$, and for every $\alpha^i \in \{S,G\}$ and $\hat{\alpha}^i \in \{S,G\}$,

$$\int J^i(\alpha^i, \alpha^{-i})\pi(d\alpha^{-i}|\alpha^i) \leqslant \int J^i(\hat{\alpha}^i, \alpha^{-i})\pi(d\alpha^{-i}|\alpha^i).$$

In the current example, we have only two players and two actions per player. The inequality holds automatically when $\hat{\alpha}^i = \alpha^i \in \{S, G\}$, so we have only two inequalities to check for each player.

We first consider the case of player $i = 1$.

– Consider $\alpha^1 = S$, $\hat{\alpha}^1 = G$. We want to have

$$J^1(S,S)\pi_{S,S} + J^1(S,G)\pi_{S,G} \leqslant J^1(G,S)\pi_{S,S} + J^1(G,G)\pi_{S,G}$$

which can be rewritten as

$$0 \leqslant -\pi_{S,S} + 100\pi_{S,G}.$$

Using the fact that π is a probability distribution, the above inequality rewrites

$$101\pi_{S,G} + \pi_{G,S} + \pi_{G,G} \geqslant 1. \tag{5.4}$$

– Consider now $\alpha^1 = G, \hat{\alpha}^1 = S$. We want to have

$$J^1(G,S)\pi_{G,S} + J^1(G,G)\pi_{G,G} \leqslant J^1(S,S)\pi_{G,S} + J^1(S,G)\pi_{G,G}$$

which rewrites

$$-\pi_{G,S} + 100\pi_{G,G} \leqslant 0,$$

or equivalently

$$100\pi_{G,G} \leqslant \pi_{G,S}. \tag{5.5}$$

Likewise, for player $i = 2$,

– Consider $\alpha^2 = S, \hat{\alpha}^2 = G$. We want to check that

$$J^2(S,S)\pi_{S,S} + J^2(G,S)\pi_{G,S} \leqslant J^2(S,G)\pi_{S,S} + J^2(G,G)\pi_{G,S}$$

which rewrites

$$0 \leqslant -\pi_{S,S} + 100\pi_{G,S}.$$

Using the fact that π is a probability distribution, the above inequality rewrites

$$101\pi_{G,S} + \pi_{S,G} + \pi_{G,G} \geqslant 1. \tag{5.6}$$

– If now $\alpha^2 = G$ and $\hat{\alpha}^2 = S$, we want

$$J^2(S,G)\pi_{S,G} + J^2(G,G)\pi_{G,G} \leqslant J^2(S,S)\pi_{S,G} + J^2(G,S)\pi_{G,G}$$

which rewrites

$$-\pi_{S,G} + 100\pi_{G,G} \leqslant 0,$$

or equivalently

$$100\pi_{G,G} \leqslant \pi_{S,G}. \tag{5.7}$$

So a probability distribution π on $\{S,G\}^2$ is a correlated equilibrium if it satisfies the above four inequalities (5.4), (5.5), (5.6), (5.7). As expected, one easily checks that the mixed Nash equilibrium found earlier satisfies the 4 inequalities. Indeed,

- for $(p,q)=(1,0)$ we have $\pi_{S,S}=\pi_{G,G}=\pi_{G,S}=0$ and $\pi_{S,G}=1$.
- for $(p,q)=(0,1)$ we have $\pi_{S,S}=\pi_{G,G}=\pi_{S,G}=0$ and $\pi_{G,S}=1$.
- for $(p,q)=(100/101,100/101)$, we have $\pi_{S,S}=\frac{100^2}{101^2}$, $\pi_{S,G}=\pi_{G,S}=\frac{100}{101^2}$ and $\pi_{G,G}=\frac{1}{101^2}$.

To limit the search for correlated equilibria we restrict ourselves to symmetric equilibria for which $\pi_{G,S}=\pi_{S,G}$. With this extra demand, a joint probability distribution $\pi\in\mathcal{P}(\{S,G\}^2)$ is a symmetric correlated equilibrium if it satisfies the following two inequalities:

$$100\pi_{G,G}\leqslant\pi_{G,S},\qquad 102\pi_{G,S}+\pi_{G,G}\geqslant 1. \tag{5.8}$$

If we compute the social cost associated to such a joint (symmetric) probability distribution $\pi\in\mathcal{P}(\{S,G\}^2)$, we see that the social cost which should be of the form

$$J=-\pi_{S,G}-\pi_{G,S}+200\pi_{G,G},$$

becomes

$$J=2(-\pi_{G,S}+100\pi_{G,G}),$$

and if we want to look for symmetric correlated equilibria with small social costs, we see that $\pi_{S,G}=\pi_{G,S}=1/2$ and $\pi_{G,G}=\pi_{S,S}=0$ is a symmetric correlated equilibrium since it satisfies the constraints given by the inequalities (5.8) and its social cost is $J=-1$.

So while the non-pure equilibrium in mixed strategy had a worse social cost than the Nash equilibria in pure strategies, for the particular cost functions given in Table 5.3, symmetric correlated equilibria can achieve the same social cost than the equilibria in pure strategies, but they cannot do better.

5.1.4 Coarse Equilibria and Regret

There is another popular notion of equilibrium which is even *weaker* than the notion of correlated equilibrium. It is called *coarse correlated equilibrium*. $\pi\in\mathcal{P}(A)$ is such an equilibrium if for all $i\in[N]$, for all $\hat{\alpha}^i\in A^i$

$$\int J^i(\boldsymbol{\alpha})\pi(d\boldsymbol{\alpha})\leqslant\int J^i(\hat{\alpha}^i,\boldsymbol{\alpha}^{-i})\pi(d\boldsymbol{\alpha}).$$

Intuitively, the notion of the best response is understood in an average sense for this notion of equilibrium.

Conveniently associated to this notion of equilibrium is the notion of regret.

Definition 5.9 *For every $\boldsymbol{\alpha}\in A$, for every $i\in[N]$, and for any measurable function $\varphi^i:A^i\to A^i$, we define the regret of player i with respect to $(\boldsymbol{\alpha},\varphi^i)$ by*

$$R_i(\boldsymbol{\alpha},\varphi^i)=J^i(\varphi^i(\alpha^i),\boldsymbol{\alpha}^{-i})-J^i(\boldsymbol{\alpha}).$$

The function φ^i is called a strategy modification.

The condition for a coarse correlated equilibrium (ccNE) can be written as follows: for all $i \in [N]$, for all constant strategy modification φ^i, i.e. $\varphi^i(\cdot) = \hat{\alpha}^i$ with some $\hat{\alpha}^i \in A^i$, we have

$$\int R_i(\boldsymbol{\alpha}, \varphi^i)\pi(d\boldsymbol{\alpha}) \geqslant 0.$$

Proposition 5.10 *A joint probability distribution $\pi \in \mathcal{P}(A)$ is a coarse correlated equilibrium if and only if for every $i \in [N]$ and for every strategy modification φ^i,*

$$\int R_i(\boldsymbol{\alpha}, \varphi^i)\pi(d\boldsymbol{\alpha}) \geqslant 0.$$

Proof: One way to think about this is to view $\pi = \mathcal{L}(\xi^1, \ldots, \xi^N)$ as a set of strategy modification. Specifically, the strategy modification φ^i changes π to $\mathcal{L}(\varphi^i(\xi^i), \xi^{-i}) \in \mathcal{P}(A)$. □

Later on in these lectures, we move on to multiple stages and dynamic games. For a quick summary capturing many of the results seen so far for one stage game, we have

$$NE \subset mNE \subset cNE \subset ccNE$$

and while we do not always have existence of NEs, generically, we have existence of mixed NEs (and hence cNE and ccNE).

5.2 Correlated Equilibria for Graphon Games

Let $w : [0,1] \times [0,1] \to [0,1]$ be a graphon kernel and $A^w : L^2([0,1]; \mathbb{R}^k) \to L^2([0,1]; \mathbb{R}^k)$ the corresponding graphon operator

$$[A^w f](x) = \int f(y)w(x,y)dy.$$

Recall that A^w is a compact operator on $L^2([0,1]; \mathbb{R}^k)$. Let $J : \mathbb{R}^k \times \mathbb{R}^k \ni (\alpha, z) \mapsto J(\alpha, z) \in \mathbb{R}$ be the cost function of the game and recall that a Nash equilibrium is a strategy profile $\boldsymbol{\alpha} : [0,1] \to \mathbb{R}^k$ such that for almost every $x \in [0,1]$ and $\alpha' \in \mathbb{R}^k$,

$$J(\boldsymbol{\alpha}(x), z(x|\boldsymbol{\alpha})) \leqslant J(\alpha', z(x|\boldsymbol{\alpha})),$$

where $z(x|\boldsymbol{\alpha}) = [A^w\boldsymbol{\alpha}](x) = \int_{[0,1]} w(x,y)\boldsymbol{\alpha}(y)dy$. Intuitively, $J(\alpha, z)$ should be understood as the cost to any player taking action α when the graphon aggregate of the other players' actions is z. One more time, we stress that J is defined everywhere but the requirement for the Nash equilibrium is only almost everywhere.

In order to consider correlated equilibria, we need first to introduce the notion of mixed strategy in the context of graphon games. Recall that for each player x, the set of admissible actions is A^x, a Borel subset of a closed convex set $A_0 \subset \mathbb{R}^k$. Recall that we use the notation $I = [0,1]$ when we want to emphasize that $[0,1]$ is the set of players in the graphon game.

Definition 5.11 *A mixed strategy for player $x \in I$, denoted by π^x, is a probability measure on $\mathbb{R}^k$ concentrated on A^x, namely $\pi^x \in \mathcal{P}(\mathbb{R}^k)$ such that $\pi^x(A^x) = 1$.*

A mixed strategy profile, denoted by $(\pi^x)_{x\in I}$, is a collection of mixed strategies for players indexed by $I = [0,1]$ such that the mapping $I \ni x \mapsto \pi^x \in \mathcal{P}(\mathbb{R}^k)$ is measurable.

The notion of measurability of $x \mapsto \pi^x$ can be defined in various ways:

- for all $B \in \mathcal{B}(\mathbb{R}^k)$, the map $x \mapsto \pi^x(B)$ is measurable;
- for all $f : \mathbb{R}^k \to \mathbb{R}$ continuous with compact support, the function $x \mapsto \int_{\mathbb{R}^k} f(\alpha)\pi^x(d\alpha)$ is measurable.

Classical monotone class arguments from measure theory can be used to prove that these two statements are equivalent.

Recall that in the case of an N-player game, say $\mathcal{G} = ([N], A = A^1 \times \cdots \times A^N, J = (J^i)_{i\in[N]})$, we defined the notion of Nash equilibrium in mixed strategies as a Nash equilibrium for the corresponding extended game $\tilde{\mathcal{G}}$. In the present context, the set of players $[N]$ should be replaced by $I = [0,1]$, and the set A, when viewed as a product space, should be replaced by the continuum product $A = \prod_{x\in I} A^x$. Recall that the space of *admissible strategy profiles* was defined as a smaller space, namely the subset $\mathbb{A}$ of this product space A comprising the (equivalent classes of) measurable and square integrable functions $\boldsymbol{\alpha}$ from I into $\mathbb{R}^k$ such that $\boldsymbol{\alpha}(x) \in A^x$ for almost every $x \in I$. The natural extension $\tilde{\mathbb{A}}$ of the space $\mathbb{A}$ of admissible strategy profiles should be defined from the product $\tilde{A} = \prod_{x\in I} \mathcal{P}(A^x)$ as the subset of the families $(\pi^x)_{x\in I} \in \tilde{A}$ which are measurable in x in any of the senses given above.

Remark 5.12 *1. To generalize the notion of mixed Nash equilibrium introduced for N-player games to the context of graphon games, it is natural to consider the product measure*

$$\mathbb{P}_0 = \bigotimes_{x\in I} \pi^x(d\alpha^x) \quad \text{on the product space} \quad A = \prod_{x\in I} A^x.$$

2. Recall that for N players, we found it convenient to approach the mixed strategies as laws of random variables $(\xi^i)_{i\in[N]}$ such that $\mathcal{L}(\xi^i) = \pi^i$ and which can be viewed as recommendations. The difference between Nash equilibria in mixed strategies and correlated equilibria could then be read off the independence of the ξ^i, or their lack thereof. Here we would like to use the same idea with a continuum of random variables $(\xi^x)_{x\in I}$. Obviously, we will have to work with a Fubini extension to define these objects rigorously if and when we want the ξ^x to be independent and depend measurably in $x \in I$.

3. As an aside, let us mention an analogy with stochastic optimal control in continuous time. Let us assume that for $t \in [0,1]$,

$$dX_t = \alpha_t dt + dW_t$$

where X_t is the state of the system at time t, α_t is the control exerted by the controller, and the goal is to minimize the cost functional

$$J(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^1 \frac{1}{2}|\alpha_t|^2 dt + \ \ldots\right].$$

We seek to minimize J over some set $\mathbb{A}$ of admissible controls which consists of stochastic processes $\alpha : [0,1] \times \Omega \to \mathbb{R}^k$ satisfying some extra conditions. Since this optimization problem may not have a solution, or since conditions for existence of an optimal control may be hard to come by, the notion of relaxed control has been introduced to alleviate some of these difficulties. Relaxed controls are measure-valued stochastic processes of the form $\pi : [0,1] \times \Omega \to \mathcal{P}(\mathbb{R}^k)$ with $\pi^t(\omega) \in \mathcal{P}(\mathbb{R}^k)$ for $t \in [0,1]$ and they bear strong relationships to the mixed controls we want to introduce to solve graphon games. Indeed, if we forget the stochastic aspect (and remove Ω from the set-up), these relaxed controls can be viewed as probability measures $\pi \in \mathcal{P}([0,1] \times \mathbb{R}^k)$ for which the first marginal is the Lebesgue measure. In other words, if we use their disintegration into conditional probabilities with respect to their first marginals,

$$\pi(dt, d\alpha) = \pi^t(d\alpha)dt,$$

or equivalently,

$$\int f(t,\alpha)\pi(dt, d\alpha) = \int_0^1 dt \left(\int f(t,\alpha)\pi^t(d\alpha) \right),$$

for every bounded measurable function f. Clearly, the kernel $[0,1] \ni t \to \pi^t$ look exactly like the mixed strategy profiles we want to analyze in the case of graphon games if we swap the time variable t for the player index x.

Back to our graphon games. Using the same type of cost function $(\alpha, z) \mapsto J(\alpha, z)$ as before, we would like to define the cost to player $x \in I$ if all the players use the mixed strategy profile $\pi = (\pi^x)_{x \in I}$, via the value $\tilde{J}^x(\pi)$ of a function $\tilde{J}^x$ defined for each mixed strategy profile π as an extension of the original cost function by the formula:

$$\begin{aligned}
\tilde{J}^x(\pi) &= \int\int\int J\Big(\alpha, \int_I w(x,y)\alpha^y dy\Big)\pi^x(d\alpha) \prod_{y \neq x} \pi^y(d\alpha^y) \\
&= \int\int J(\alpha, z)\pi^x(d\alpha)\pi^{w,x}(dz)
\end{aligned}$$

if we denote by $\pi^{w,x}$ the push forward of the measure $\mathbb{P}_0 = \bigotimes_{x \in [0,1]} \pi^x(d\alpha^x)$ by the map $(\alpha^y)_{y \in I} \mapsto \int_I w(x,y)\alpha^y dy$. This push forward measure has a simple expression when $\mathbb{P}_0$ is a product measure, or to be more rigorous, when the $(\alpha^y)_{y \in I}$ are essentially pairwise independent. Indeed, in this case, the exact law of large numbers applied for each fixed $x \in I$ to the essentially pairwise independent random variables $X^y(\omega) = w(x,y)\omega(y)$ says that for $\mathbb{P}_0$-almpst surely in ω,

$$\int_I X^y(\omega)dy = \int_I \mathbb{E}^{\mathbb{P}_0}[X^y]dy$$

or in other words

$$\int_I w(x,y)\omega(y)dy = \int_I w(x,y)\Big(\int_{\mathbb{R}^k} \alpha\pi^y(d\alpha)\Big)dy = \int_{\mathbb{R}^k} \alpha\Big(\int_I w(x,y)\pi^y(d\alpha)dy\Big)$$

which is not random, identifying the push forward measure as the Dirac point mass at the mean of the measure

$$\pi^{w,x}(d\alpha) = \int_I w(x,y)\pi^y(d\alpha)dy \in \mathcal{P}(\mathbb{R}^k) \tag{5.9}$$

which is the aggregate (weigted by the graphon weights $w(x,y)$) of the mixed strategies $\pi^y(d\alpha)$ of all the players viewed from the point of view of player $x \in I$, and which is in full analogy with the definition of the pure aggregate $z(x|\boldsymbol{\alpha})$.

Definition 5.13 *The mixed strategy profile $\pi = (\pi^x)_{x\in I}$ is said to be a mixed Nash Equilibrium (mNE) for the graphon game with interaction graphon w if for almost every $x \in I$ and for all $\theta \in \mathcal{P}(\mathbb{R}^k)$ such that $\theta(A^x) = 1$ we have:*

$$\tilde{J}^x(\pi) \leqslant \int J(\alpha, z)\theta(d\alpha)\pi^{w,x}(dz).$$

where $\pi^{w,x}$ is defined by (5.9).

Given a mixed strategy profile $\pi = (\pi^x)_{x\in I}$, Proposition 4.49 says that there exists a white noise $(\xi^x)_{x\in I}$ of essentially pairwise independent A_0-valued random variables satisfying $\mathcal{L}(\xi^x) = \pi^x$ for all $x \in I$. Note that the above definition of the cost $\tilde{J}^x(\pi)$ can be rewritten as:

$$\tilde{J}^x(\pi) = \mathbb{E}\Big[J\Big(\xi^x, \int_I w(x,y)\xi^y dy\Big)\Big]$$

and $\pi = (\pi^x)_{x\in I}$ is a Nash equilibrium in mixed strategies if for almost every $x \in I$ and for any random variable θ^x with values in A^x we have:

$$\tilde{J}^x(\pi) \leqslant \mathbb{E}\Big[J\Big(\theta^x, \int_I w(x,y)\xi^y dy\Big)\Big].$$

Now, take any measurable A_0-valued process $(\xi^x)_{x\in I}$ (not necessarily of essentially pairwise independent random variables) on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. By Kolmogorov's theorem we know that, up to the measurability requirement which is crucial for us for the formulas below to make sense, such a process could be defined through a probability measure $\mathbb{P}$ on the product space $(\Omega, \mathcal{F}) = ((\mathbb{R}^k)^I(\mathcal{B}_k)^I)$ such that the marginal distribution at x coincides with that of ξ^x. In any case, we set $\pi^x = \mathbb{P}\circ(\xi^x)^{-1}$ for these marginal distributions. Denoting by $\mathbb{E}$ the expectation under this probability $\mathbb{P}$, we have

$$J^x(\pi) = \mathbb{E}\left[J\left(\xi^x, \int_0^1 w(x,y)\xi^y dy\right)\right],$$

and if we define the process $\boldsymbol{z} = (z(x|\xi))_{x\in[0,1]}$ by

$$z(x|\xi) = \int w(x,y)\xi^y dy$$

the latter can be interpreted as the graphon weighted aggregate signal and it is natural to introduce the following definition. Note that the fact that for each $\omega \in \Omega$ fixed the function $y \mapsto \xi^y(\omega)$ is measurable is crucial for the above integrals to make sense.

Definition 5.14 *For a given mixed strategy profile $\pi = (\pi^x)_{x \in I}$, we say that a probability distribution $\mathbb{P}$ satisfies $\pi^x = \mathbb{P} \circ (\xi^x)^{-1}$ for every $x \in I$ is a correlated equilibrium if*

$$J^x(\pi) \leqslant \mathbb{E}\left[J\left(\zeta, \int_0^1 w(x,y)\xi^y dy\right)\right],$$

for almost every $x \in I$ and for all A^x-valued random variable ζ.

With this definition of correlated equilibrium, the notion of mixed Nash equilibrium (mNE) can be reformulated in the following way: $\pi = (\pi^x)_{x \in [0,1]}$ *is a mixed Nash equilibrium if and only if the product measure $\mathbb{P}_0$ is a correlated equilibrium in the sense of the above definition.*

5.3 Large Bayesian Games and Cournot Competition

We first describe mathematically how under some very specific conditions, one can take the limit when the number N of players tends to ∞ and obtain limiting objects of the anonymous game type.

5.3.1 Anonymous Games as Limits

In an N-player game, we denote by Θ^i the set of possible types of player $i \in [N]$, and the set of type profiles is denoted by $\Theta = \Theta^1 \times \cdots \times \Theta^N$. We also denote the set of pure strategy profiles by $A = A^1 \times \cdots \times A^N$. Here, we consider game models which have some kind of symmetry among the players. Previously, the cost to to player $i \in [N]$ depended on the type profile and the strategy profile $(\boldsymbol{\theta}, \boldsymbol{\alpha})$ of all players and we would write it as $J^i(\boldsymbol{\theta}, \boldsymbol{\alpha})$. Now we consider costs which take the form:

$$J^i(\boldsymbol{\theta}, \boldsymbol{\alpha}) = J_0\left(\theta^i, \alpha^i, \frac{1}{N-1}\sum_{j=1, j\neq i}^{N} \delta_{(\theta^i, \alpha^i)}\right)$$

for some common function J_0. In particular,

$$J^i(\theta^i, \alpha^i, (\theta^j, \alpha^j)_{j\neq i}) = J^i(\theta^i, \alpha^i, (\theta^{\sigma(j)}, \alpha^{\sigma(j)})_{j \neq i})$$

Furthermore, we assume that all the type sets Θ^j for $j \in [N]$ are identical, and we denote their common value by Θ_0. Similarly, we assume that all the feasible control sets A^j for $j \in [N]$ are the same and we denote their common value by A_0. So the function J_0 used above to define the individual casts is of the form:

$$J_0 : \Theta_0 \times A_0 \times \mathcal{P}(\Theta_0 \times A_0) \to \mathbb{R}.$$

The following proposition (borrowed from [10]) is a rigorous justification of the intuitive fact that the symmetry among the players implies that for large games, the cost function J_0 can indeed be derived from a function of measures.

Lemma 5.15 *Let E be a compact metric space, and for each integer N, let $u^N : E^N \to \mathbb{R}$ be symmetric, i.e.,*

$$u^N((x_i)_{i\in[N]}) = u^N((x_{\sigma(i)})_{i\in[N]})$$

for any permutation $\sigma : [N] \to [N]$ of the indices $\{1, \ldots, N\}$. Furthermore we assume:

1. *uniform boundedness:*

$$\sup_{N \geqslant 1} \sup_{(x_i)_{i\in[N]}\in E^N} \left|u^N\left((x_i)_{i\in[N]}\right)\right| < +\infty;$$

2. *uniform Lipschitz continuity: there exists a constant $c > 0$ such that for all $N \geqslant 1$ and for all $X, Y \in E^N$,*

$$\left|u^N(X) - u^N(Y)\right| \leqslant c\rho\left(\bar{\mu}^N_X, \bar{\mu}^N_Y\right)$$

where $\bar{\mu}^N_X = \frac{1}{N}\sum_{j=1}^N \delta_{X_j}$ is the empirical measure, and ρ is a metric on $\mathcal{P}(E)$ for weak convergence.

Then there is a subsequence $(u^{N_k})_{k\geqslant 1}$ and a Lipschitz continuous map $U : \mathcal{P}(E) \to \mathbb{R}$ such that

$$\lim_{k\to+\infty} \sup_{X\in E^{N_k}} |u^{N_k}(X) - U(\bar{\mu}^{N_k}_X)| = 0.$$

Intuitively, the above result says that a symmetric function of a large number of variables can be approximated by a function of the empirical distribution when the number of variables is large enough.

The relevance of this result to our current concerns is to allow us to apply it to the set $E^{N-1} = (\Theta_0 \times A_0)^{N-1}$ for $u^{N-1}(\cdot) = J^i(\theta^i, \alpha^i, \cdot)$ to be able to say

$$J^i(\theta^i, \alpha^i, (\theta^j, \alpha^j)_{j\neq i}) \approx J_0\left(\theta^i, \alpha^i, \frac{1}{N-1}\sum_{j\neq i} \delta_{(\theta^j,\alpha^j)}\right).$$

In fact, we do not really need to explicitly assume that u^N is symmetric because what is needed in the proof can be inferred from the uniform Lipschitz continuity condition.

Proof: We could repeat mutatis mutandis the proof given in [11, Lemma 1.2]. We only outline the main steps. For each $n \geqslant 1$, we define the function $U^n : \mathcal{P}(E) \to \mathbb{R}$ by

$$U^n(\mu) = \inf_{X\in E^n} \left[u^n(X) + c\rho(\overline{\mu}^n_X, \mu)\right], \qquad \mu \in \mathcal{P}(E)$$

We can split the proof into 4 steps:

- First we show that U^n is uniformly bounded.
- Next we show that U^n extends u^n in the sense that for any $Y \in E^n$ and $n \geqslant 1$,

$$u^n(Y) = U^n(\overline{\mu}^n_Y).$$

- Then we argue that U^n is uniformly $c-$Lipschitz on $\mathcal{P}(E)$:

$$|U^n(\mu) - U^n(\nu)| \leqslant c\rho(\mu, \nu)$$

  for all $\mu, \nu \in \mathcal{P}(E)$.
- Finally, using Arzela-Ascoli's theorem and the compactness of the metric space $\mathcal{P}(E)$, we see that there exists a subsequence $(n_k)_{k\geqslant 1}$ for which U^{n_k} converges uniformly toward a limit U so that

$$\limsup_{k\to\infty} \sup_{X\in E^{n_k}} |u^{n_k}(X) - U(\overline{\mu}^{n_k}_X)| \leqslant \lim_{k\to\infty} \sup_{\mu\in\mathcal{P}(E)} |U^{n_k}(\mu) - U(\mu)| = 0$$

This concludes the proof. □

5.3.2 **Large Bayesian Games**

Recall that the basic building blocks of a Bayesian game are:

- An action space A_0.
- The set $[N] = \{1, \dots, N\}$ of players.
- For each $i \in [N]$, the set $A^i \subset A_0$ of feasible actions for player i. We assume that $A_0 \subset \mathbb{R}^k$ for some integer $k \geqslant 1$ and that A^i is closed.
- $\Theta = \Theta^1 \times \cdots \times \Theta^N$ be the space of type profiles, where $\Theta^i \subset \Theta_0$.
- The costs functions $J = (J^i)_{i \in [N]}$, for which $J^i(\boldsymbol{\theta}, \boldsymbol{\alpha}) = J^i(\theta^1, \dots, \theta^N, \alpha^1, \dots, \alpha^N)$.

We now add a symmetry assumption which will allow us to use Lemma 5.15. To be specific, we assume that for each $i \in [N]$ and $\alpha^i \in A^i$ and $\theta^i \in \Theta^i$ the function:

$$(\theta^j, \alpha^j)_{j \neq i} \mapsto J^i(\boldsymbol{\theta}, \boldsymbol{\alpha})$$

is symmetric and satisfies the assumptions of Lemma 5.15. As a result, if we want to analyze models of large games we can assume that

$$J^i(\boldsymbol{\theta}, \boldsymbol{\alpha}) = J^i(\text{ symmetric function of } (\theta^j, \alpha^j)_{j \neq i}, \theta^i, \alpha^i).$$

for some function $J_0 : \mathcal{P}(\Theta_0 \times A_0) \times \Theta_0 \times A_0 \to \mathbb{R}$.

In this set up, a pure strategy for player i is a measurable function $\alpha^i : \Theta^i \to A^i$, while a mixed strategy is a function $\pi^i : \Theta^i \to \mathcal{P}(A^i)$. We restrict ourselves to the case when Θ^i are compact subset of Θ_0 and because of the symmetry assumption we shall also assume that $\Theta^i = \Theta_0$, $A^i = A_0$ for al $i \in [N]$.

5.3.3 **Cournot Competition**

Definition 5.16 *Given a probability measure $\lambda \in \mathcal{P}(\Theta_0)$ over types, we define the set $\mathcal{C}(\lambda)$ of Cournot-Nash equilibria as*

$$\mathcal{C}(\lambda) = \Big\{ \mu \in \mathcal{P}(\Theta_0 \times A_0); \;\; \Pi^1(\mu) = \lambda \textit{ and } \\ \mu\big(\big\{(\theta, \alpha) \in \Theta_0 \times A_0; \alpha \in A(\theta), J_0(\mu, \theta, \alpha) = \inf_{\beta \in A(\theta)} J_0(\mu, \theta, \beta)\big\}\big) = 1 \Big\}$$

where Π^1 denotes the first marginal, namely $\mu(B \times A_0) = \lambda(B)$ for all $B \in \mathcal{B}()$.

Remark 5.17 *1. In the above definition, we see that the types are revealed first, according to the distribution λ, and then players take actions. So λ should be viewed as a prior distribution of the types.*
2. The condition $\mu(\cdots) = 1$ means that the measure μ is concentrated on the set of type-action pairs (θ, α) such that for each type θ, the action α is the best response.
3. By disintegration of the measure μ (i.e. using a regular version of the conditional probability given the first component θ), we have $\mu(d\theta, d\alpha) = \lambda(d\theta)\pi^\theta(d\alpha)$ where $\pi = (\pi^\theta)_{\theta \in \Theta_0}$ is measurable in θ, in other words, is a probability transition kernel form $(\Theta_0, \mathcal{B}(\Theta_0))$ to $(A_0, \mathcal{B}(A_0))$.

4. *If there is a unique best response $\alpha(\theta)$ for each type θ, $\pi^\theta(d\alpha) = \delta_{\alpha(\theta)}(d\alpha)$ for some $\alpha(\theta) \in A_0$, then the mapping defined by $\Theta_0 \ni \theta \mapsto \alpha(\theta) \in A_0$ is a pure strategy.*

For the next result we shall make the following assumptions.

Assumption 2.

A1. *A_0 is a compact metric space;*
A2. *Θ_0 is a Polish space (complete separable metric space);*
A3. *The constraint correspondence $A : \Theta_0 \twoheadrightarrow A_0$ satisfies*
- *$\forall \theta \in \Theta_0$, $A(\theta)$ is a non-empty closed subset of A_0.*
- *A is continuous, or equivalently,*
  - *the graph $gr(A) = \{(\theta, \alpha) \in \Theta_0 \times A_0 : \alpha \in A(\theta)\}$ is closed (which implies that A is upper hemi-continuous since A_0 is compact);*
  - *for every sequence $(\theta_n)_n$ and $\theta \in \Theta_0$, $\alpha \in A(\theta)$ with $\theta_n \to \theta$, there exists a subsequence $(n_k)_k$ and $\alpha_{n_k} \in A(\theta_{n_k})$ such that $\alpha_{n_k} \to \alpha$ as $k \to \infty$ (A is lower hemi-continuous).*

A4. *$J_0 : \mathcal{P}(\Theta_0 \times A_0) \times \Theta_0 \times A_0 \to \mathbb{R}$ is continuous and bounded.*

The following proposition is borrowed from [23] where the authors also study the corresponding law of large numbers and the central limit theorems appropriate for these models.

Proposition 5.18 (Mas-Colell) *Under assumption 2, for each $\lambda \in \mathcal{P}(\Theta_0)$, the set $\mathcal{C}(\lambda)$ of Cournot-Nash equilibria is not empty.*

Proof: For $\lambda \in \mathcal{P}(\Theta_0)$. Set

$$\mathcal{M}(\lambda) = \{\mu \in \mathcal{P}(\Theta_0 \times A_0); \Pi^1(\mu) = \lambda, \mu(\{(\theta, \alpha); \alpha \in A(\theta)\}) = 1\}.$$

1. Step 1. $\mathcal{M}(\lambda) \subset \mathcal{P}(\Theta_0 \times A_0)$ is closed.
   Indeed, if $\mu_n \in \mathcal{M}(\lambda)$ is such that $\mu_n \Rightarrow \mu$ weakly we can show that $\mu \in \mathcal{M}(\lambda)$. By definition of weak convergence, for all $f : \Theta_0 \times A_0 \to \mathbb{R}$ continuous, $\int f d\mu_n \to \int f d\mu$. In particular for $f(\theta, \alpha) = f(\theta)$ independent of the second variable, we have $\int f(\theta) \Pi^1(\mu_n)(d\theta) \xrightarrow[n\to\infty]{} \int f(\theta) \Pi^1(\mu)(d\theta)$. So that $\Pi^1(\mu)(d\theta) = \lambda(d\theta)$.
2. Step 2. $\mathcal{M}(\lambda)$ is tight, i.e., for all $\epsilon > 0$, there exists a compact $K_\epsilon \subset \Theta_0 \times A_0$ such that for all $\mu \in \mathcal{M}(\lambda)$, $\mu(K_\epsilon) \geqslant 1 - \epsilon$. We should notice that Θ_0 is not necessarily compact. Since Θ_0 is a Polish space, for $\epsilon > 0$ there exists a compact set $\tilde{K}_\epsilon \subset \Theta_0$ such that $\lambda(\tilde{K}_\epsilon) \geqslant 1 - \epsilon$. Since the first marginal of any $\mu \in \mathcal{M}(\lambda)$ is equal to λ it is obvious that $\mu(K_\epsilon) \geqslant 1 - \epsilon$ if we set $K_\epsilon = \tilde{K}_\epsilon \times A_0$ which is obviously compact.
3. Step 3. We define the correspondence $\Phi : \mathcal{M}(\lambda) \twoheadrightarrow \mathcal{M}(\lambda)$ by

$$\Phi(\mu) = \left\{\tilde{\mu} \in \mathcal{M}(\lambda);\ \int_{\Theta_0 \times A_0} G(\mu, \theta, \alpha)\tilde{\mu}(d\theta, d\alpha) \leqslant 0\right\}$$

   where

$$G(\mu, \theta, \alpha) = J_0(\mu, \theta, \alpha) - \inf_{\beta \in A(\theta)} J_0(\mu, \theta, \beta).$$

   Notice that a fixed point for Φ is a Cournot-Nash equilibrium. So our goal is to conclude the proof using Kakutani's fixed point theorem (see Theorem 1.29 in Chapter 1), so we proceed to check the assumptions of this theorem. G is bounded because J_0 is bounded. Since J_0 is continuous, the assumptions of Berge's minimum theorem recalled in Chapter 1 as Theorem 1.30 are satisfied and we learn that G is bounded and continuous.

- Since the correspondence A is continuous, it is measurable with respect to the Borel $\sigma-$algebras on Θ_0 and A_0, so that A is also weakly measurable (since both Θ_0 and A_0 are metrizable, see [1, Lemma 18.2]). Moreover, since J_0 is continuous, it is a Carathéodory function and we can use Theorem 1.32 recalled in the appendix of Chapter 1 to conclude the existence of $\hat{\alpha} : \mathcal{M}(\lambda) \times \Theta_0 \to A_0$ such that
$$J_0(\mu, \theta, \hat{\alpha}(\mu, \theta)) = \inf_{\beta \in A(\theta)} J_0(\mu, \theta, \beta).$$
So if we define the measure $\tilde{\mu}$ by
$$\tilde{\mu}(d\theta, d\alpha) = \lambda(d\theta)\delta_{\hat{\alpha}(\mu,\theta)}(d\alpha)$$
we immediately see that $\tilde{\mu} \in \Phi(\mu)$ proving that $\Phi(\mu)$ is not empty.
- $\Phi(\mu)$ is convex, because $\mathcal{M}(\lambda)$ is convex.
- We now check that the graph of Φ, namely
$$gr(\Phi) = \{(\mu, \tilde{\mu}) \in \mathcal{M}(\lambda) \times \mathcal{M}(\lambda) : \quad \tilde{\mu} \in \Phi(\mu)\},$$
is closed. Remember that we only assume that Θ_0 is a Polish space. Let us assume that $(\mu_n, \tilde{\mu}_n)_{n \geqslant 1}$ is a sequence with $\tilde{\mu}_n \in \Phi(\mu_n)$ for each $n \geqslant 1$, and such that there exists $\mu, \tilde{\mu} \in \mathcal{M}(\lambda)$ such that $(\mu_n, \tilde{\mu}_n) \Rightarrow (\mu, \tilde{\mu})$ weakly in $\mathcal{P}(\Theta_0 \times A_0)$. We want to show that $(\mu, \tilde{\mu}) \in gr(\Phi)$ to conclude that $gr(\Phi)$ is closed.
To do so, let us pick $\epsilon > 0$. Because a convergent sequence of measures is tight, there exists a compact set $K_\epsilon \subset \Theta_0 \times A_0$ such that for every $n \geqslant 1$
$$\tilde{\mu}_n(K_\epsilon^c) < \frac{\epsilon}{4\|G\|_\infty}.$$
Since G is continuous, it is uniformly continuous on $K \times K_\epsilon$ where K is a compact subset of $\mathcal{P}(\Theta_0 \times A_0)$ containing $\{\mu_n, n \geqslant 1\} \cup \{\mu\} \subset K$. Then, there exists an n_ϵ such that for all $n \geqslant n_\epsilon$ and for all $(\theta, \alpha) \in K_\epsilon$,
$$|G(\mu_n, \theta, \alpha) - G(\mu, \theta, \alpha)| \leqslant \frac{\epsilon}{2}$$
So, for $n \geqslant n_\epsilon$,
$$\begin{aligned}
&\int |G(\mu_n, \theta, \alpha) - G(\mu, \theta, \alpha)| \, \tilde{\mu}_n(d\theta, d\alpha) \\
&\qquad = \left(\int_{K_\epsilon^c} + \int_{K_\epsilon} \right) |G(\mu_n, \theta, \alpha) - G(\mu, \theta, \alpha)| \, \tilde{\mu}_n(d\theta, d\alpha) \\
&\qquad \leqslant 2\tilde{\mu}_n(K_\epsilon^c)\|G\|_\infty + \frac{\epsilon}{2} \\
&\qquad \leqslant \epsilon.
\end{aligned}$$
Thus, for n large enough we have:
$$\begin{aligned}
&\int G(\mu, \theta, \alpha)\tilde{\mu}(d\theta, d\alpha) \\
&\quad = \int G(\mu, \theta, \alpha)[\tilde{\mu} - \tilde{\mu}_n](d\theta, d\alpha) + \int [G(\mu, \theta, \alpha) - G(\mu_n, \theta, \alpha)]\tilde{\mu}_n(d\theta, d\alpha) \\
&\qquad\qquad + \int G(\mu_n, \theta, \alpha)\tilde{\mu}_n(d\theta, d\alpha) \\
&\quad \leqslant \epsilon + \epsilon + 0 \\
&\qquad\qquad \leqslant 2\epsilon
\end{aligned}$$

because the first integral tends to 0 because G is bounded and continuous and μ_n convergences weakly towards μ, the second integral is smaller than ϵ for all $n \geqslant n_\epsilon$, and the third integral is not greater than zeros because $(\mu, \tilde{\mu}_n) \in gr(\Phi)$. Since $\epsilon > 0$ was arbitraty, this proves that $(\mu, \tilde{\mu}) \in gr(\Phi)$. Therefore, we can apply Kakutani's fixed point theorem and conclude the existence of a Cournot-Nash equilibrium.

This concludes the proof. □

Part II

MULTI-STAGE & DYNAMIC GAMES

6

Multi-Stage & Dynamic Games

6.1 Multi-Stage Games

A multi-stage game consists of the following components:

- A sequence $(0, \ldots, K)$ which represents the stages of the game;
- N players who play repeatedly and take action simultaneously at each stage of the game;

At each stage, each player is cognizant of the actions of all the N players in all the previous stages. For each stage $k \in \{0, \ldots, K\}$, the action profile at this stage is denoted by:

$$\boldsymbol{\alpha}_k = (\alpha_k^1, \ldots, \alpha_k^N)$$

where α_k^i is the action chosen by player i at stage k. The available information before this action is taken at stage k is called the *history*. It is defined as:

$$h^k = (\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_{k-1}),$$

for $k \geqslant 1$, and $h^0 = \varnothing$ by convention. At stage k, the feasible set of actions for player i is denoted by $A^i(h^k)$. It depends on the information contained in the history prior to stage k. For example, the resources that have been used in the past can affect the current actions of players. So, at stage k, the action chosen by player i can be viewed as a function of history h^k;

We include *"no action"* as part of $A^i(h^k)$ for every $i = 1, \ldots, N$ and for every $k = 0, \ldots, K$. This will gave flexibility to the model, allowing players to be inactive at certain stages. This will be the case when player i choses "no action".

Here are some examples for multi-stage games:

1. $K = 0$, i.e. with one single stage, so no repetition, the **Cournot competition** discussed in Subsection 1.2.2 is an example.
2. $K = 1$, i.e. with two stages, our model of repeated games can capture the so-called **Stackelberg game** models if we choose $N = 2$ to consider only two players. One player is called the *leader*, and the other is called the *follower*. Then
   - at stage $k = 0$, the leader chooses an action $\alpha_0 \in A^{leader}(\varnothing)$. The follower is inactive, namely $A^{follower}(\varnothing) = \{\text{“}no\,action\text{”}\}$.

- at stage $k = 1$, the leader is inactive which we force by setting $A^{leader}(h^1) = \{\text{“}no\,action\text{”}\}$, and the follower chooses an action $\alpha^{follower} \in A^{follower}(h^1)$ where $h^1 = ((\alpha_0, \text{“}no\,action\text{”}))$.

3. $K = 2$, i.e. with two stages, our formalism can be used to include **entry-deterrence** games between two players, an *incombent* and an *entrant*.
   - at stage $k = 0$, the incumbent raises and amount α_0 of money, and the entrant is forced to be inactive;
   - at stage $k = 1$, the incumbent is inactive (can only choose “no action”), while the entrant choses between “enter” or “do not enter”;
   - at stage $k = 2$, if the entrant chose to enter at stage $k = 1$, the game is played between the incumbent and the entrant as a Cournot competition.

Definition 6.1 *A pure strategy for player $i \in \{1, \ldots, N\}$, is a sequence $\hat{\alpha}^i = (\hat{\alpha}^i_k)_{k=0,\cdots,K}$ of actions, one for each stage. A pure strategy profile, denoted by $\hat{\boldsymbol{\alpha}}$, is a collection of pure strategies, one for each player:*

$$\hat{\boldsymbol{\alpha}} = (\hat{\alpha}^i)_{i=1,\ldots,N} = (\hat{\alpha}^i_k)_{i=1,\ldots,N,\,k=0,\ldots,K}.$$

It is said to be an admissible pure strategy profile if $\hat{\alpha}^i_k \in A^i(h^k)$ for all $k = 0, \ldots, K$ and $i = 1, \ldots, N$.

The set of histories h^k is denoted by $\mathcal{H}(k)$. So $\mathcal{H}(K+1)$ represents the set of all strategy profiles throughout the game, namely the set of all the actions taken by all the players throughout the game.We complete the definition of the repeated game by defining the costs to the players. At the last stage $k = K$, we collect all the histories and we decide of the rewards or costs to the individual players.

Definition 6.2 *To each player $i \in \{1, \ldots, N\}$, we assign a cost function $J^i : \mathcal{H}(K+1) \to \mathbb{R}$.*

We shall often denote by $[N]$ the set $\{1, \cdots, N\}$ of the N players. Also, the cost function J^i is often assumed to be distributed and addictive, i.e. of the form

$$J^i(h^K) = \sum_{k=0}^{K} c_k f^i(\alpha^i_k).$$

The coefficients $(c_k)_{k=0,\ldots,K}$ are most often used as discount factors. For example, when they are of the form $c_k = c_K \delta^k$ for some $c_K \in \mathbb{R}$ and $\delta \in [0, 1]$, $J^i(h^K)$ can represent the *present value* of the aggregate of the one-stage costs $f^i(\alpha^i_k)$ accumulated over the length of the game. They can also be used to model the patience or the impatience of the players. For infinitely repeated games, i.e. when $K = \infty$, we often assume that the functions f^i are uniformly bounded and $\delta < 1$ to make sure that the (infinite) summation defining $J^i(h^K)$ makes sense.

Definition 6.3 *A pure strategy profile $\hat{\boldsymbol{\alpha}}$ is said to be a pure strategy Nash equilibrium if for every player $i \in \{1, \ldots, N\}$ and every strategy (i.e. sequence of actions $\alpha^i = (\alpha^i_0, \ldots, \alpha^i_K)$ for player i), such that the strategy profile $(\alpha^i, \boldsymbol{\alpha}^{-i})$ is still admissible, we have:*

$$J^i(\hat{\boldsymbol{\alpha}}) \leqslant J^i((\alpha^i, \hat{\boldsymbol{\alpha}}^{-i}))$$

where as we did up to now, we use the notation $(\alpha^i, \hat{\boldsymbol{\alpha}}^{-i}) = (\alpha^i_k, \hat{\boldsymbol{\alpha}}^{-i}_k)_{k=0,\ldots,K-1}$ *with* $(\alpha^i_k, \hat{\boldsymbol{\alpha}}^{-i}_k) = (\hat{\alpha}^1_k, \ldots, \hat{\alpha}^{i-1}_k, \alpha^i_k, \hat{\alpha}^{i+1}_k, \ldots, \hat{\alpha}^N_k)$ *is the action profile at stage* k*.*

Remark 6.4 *Notice that the admissibility of the strategy profile* $(\alpha^i, \boldsymbol{\alpha}^{-i})$ *may be quite difficult to check in some cases. Indeed, having player* i *change his strategy* $\hat{\alpha}^i$ *may change all the histories* h^k *which may make the admissibility condition* $\alpha^j_k \in A^j(h^k)$ *difficult to verify.*

We now introduce a concept which plays an important role in the analysis of repeated and dynamic games. We first define the notion of sub-game.

Definition 6.5 *For each stage* $k = 0, \ldots, K$ *and for each history* $h^k \in \mathcal{H}(k)$*, we define the game* $\mathcal{G}(h^k)$ *as the new game, depending on the history* $h^k = (\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_{k-1})$ *in the following way:*

- *The new game has still the same* N *players;*
- *The new pure strategy profiles for the game* $G(h^k)$ *are the*

$$\boldsymbol{\alpha}_{h^k} = (\alpha^i_\ell)_{i=1,\ldots,N,\, \ell=k,k+1,\ldots,K};$$

- *The new histories* $h^{k,\ell} \in \mathcal{H}^k(\ell)$ *are given by* $h^{k,\ell} = (h^k, \boldsymbol{\alpha}_k, \ldots, \boldsymbol{\alpha}_{\ell-1})$ *for* $\ell = k, \cdots, K+1$*;*
- *The new cost functions are the functions* $\mathcal{H}^k(K+1) \ni h^{k,K+1} \mapsto J^i(h^{k,K+1})$*.*

We now define the notion of *sub-game perfect Nash equilibrium*.

Definition 6.6 *A strategy profile* $\hat{\boldsymbol{\alpha}}$ *is said to be a sub-game perfect Nash equilibrium if for every stage* $k \in \{0, \ldots, K\}$ *and every history* $h^k \in \mathcal{H}(k)$*, the strategy profile* $(\hat{\alpha}^i_\ell)_{i=1,\ldots,N,\, \ell=k,\ldots,K-1}$ *is a Nash equilibrium for the sub game* $\mathcal{G}(h^k)$*.*

In words, at every stage k, whatever the history, if players play according to the strategy profile $(\hat{\alpha}^i_\ell)_{i=1,\ldots,N,\, \ell=k,\ldots,K-1}$ afterwards, then they are in a Nash equilibrium of the new game.

Remark 6.7 *We now introduce the notions of open and closed loop strategies. Let* α^i_k *be the action taken by player* i *at stage* k*. For the sake of simplicity, we assume that* $A^i(h^k) = \mathbb{R}$ *for every* $h^k \in H(k)$ *and every* $k = 0, \ldots, K$*.*

- $\boldsymbol{\alpha} = (\alpha^i_k)_{i,k}$ *is said to be an open loop strategy profile if there exists a function* $\varphi^i : \{0, \ldots, K\} \to \mathbb{R}$ *such that* $\alpha^i_k = \varphi^i(k)$ *for every* $k = 0, \ldots K$*, the functions* φ^i *being chosen before the game begins.*
- $\boldsymbol{\alpha} = (\alpha^i_k)_{i,k}$ *is said to be a closed loop strategy profile if for each* $i \in [N]$ *and* $k \in \{0, \cdots, K\}$*, there exists a function* $\varphi^i_k : \mathcal{H}(k) \to \mathbb{R}$ *such that*

$$\alpha^i_k = \varphi^i_k(\boldsymbol{\alpha}_0, \cdots, \boldsymbol{\alpha}_{k-1}).$$

We shall revisit these definitions in the case of dynamic stochastic games.

6.1.1 The One-Stage Prisoner's Dilemma Game

Assume that there are only two players $N = 2$, and that the feasible sets of actions are $A_1 = A_2 = \{C, D\}$ for every stage k of the game. Here C stands for cooperation and D stands for defection.

The costs to the two players are defined in terms of four real numbers $T, R, P, S \in \mathbb{R}$ which satisfy $T < R < P < S$. The costs are labeled after letters which have an intuitive meaning, so we kept the letters used in the classical literature on game theory: T stands for *temptation*, R for *reward* , P for *punishment*, and S for *sucker*.

In the one-stage game, the costs of the two players are given by the following two functions:

$$J^1(\alpha^1, \alpha^2) = \mathbf{1}_{\alpha^1=C}\Big(R\mathbf{1}_{\alpha^2=C} + T\mathbf{1}_{\alpha^2=D}\Big) + \mathbf{1}_{\alpha^1=D}\Big(S\mathbf{1}_{\alpha^2=C} + P\mathbf{1}_{\alpha^2=D}\Big),$$

and

$$J^2(\alpha^1, \alpha^2) = \mathbf{1}_{\alpha^2=C}\Big(R\mathbf{1}_{\alpha^1=C} + T\mathbf{1}_{\alpha^1=D}\Big) + \mathbf{1}_{\alpha^2=D}\Big(S\mathbf{1}_{\alpha^1=C} + P\mathbf{1}_{\alpha^1=D}\Big).$$

To conform with the standard practice which gives the costs in a matrix format, we restate the definition of the costs in table 6.4):

P1\P2	C	D
C	(R, R)	(S, T)
D	(T, S)	(P, P)

Table 6.4: *Cost functions of the one-stage game for the two players in a matrix form. The entries of the matrix give the values of the costs $(J^1(\cdot), J^2(\cdot))$ representing the costs of player 1 and player 2 in different situation. The column names represent the strategies chosen by player 2 and the row names are for player 1.*

The objective of the players is to minimize their own costs. Notice that the inequalities $T < R$ and $P < S$ imply that defection is a preferable strategy for both players. Still the unique Nash equilibrium is different.

Proposition 6.8 *The strategy (D, D) is a Nash equilibrium.*

Proof: The proof is done by inspection.

1. If players choose the strategy $(\alpha^1, \alpha^2) = (D, D)$, then
$$J^1(D, D) = P < S = J^1(C, D), \qquad J^2(D, D) = P < S = J^2(D, C).$$
Thus, (D, D) is a Nash equilibrium.
2. If $(\alpha^1, \alpha^2) = (D, C)$, then
$$J^1(D, C) = T < R = J^1(C, C), \qquad J^2(D, C) = S > P = J^2(D, D).$$
Thus, (D, C) is not a Nash equilibrium.

3. If $(\alpha^1, \alpha^2) = (C, D)$, then

$$J^1(C, D) = S > P = J^1(D, D), \qquad J^2(C, D) = T < R = J^2(C, C).$$

Thus, (C, D) is not a Nash equilibrium.

4. If $(\alpha^1, \alpha^2) = (C, C)$, then

$$J^1(C, C) = R > T = J^1(D, C), \qquad J^2(C, C) = R > T = J^2(C, D).$$

Thus, (C, D) is not a Nash equilibrium,

which completes the proof. □

6.1.2 **Repeated Games**

Let us denote the cost function in the static one-stage game by

$$\tilde{J}^i(\boldsymbol{\alpha}) = J^i(\alpha^1, \alpha^2), \qquad with \quad \boldsymbol{\alpha} = (\alpha^1, \alpha^2), \quad i = 1, 2,$$

Finitely Repeated Games

The cost for a $K + 1$-stage game for player i can be expressed by

$$J^i(\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_K) = \frac{1-\delta}{1-\delta^{K+1}} \sum_{j=0}^{K} \delta^j \, \tilde{J}^i(\boldsymbol{\alpha}_j)$$

where the constant $0 \leqslant \delta < 1$ captures the fact that the players are impatient and want to weigh the recent results more heavily. The constant in front of the summation sign does not play any particular role in the search for equilibria. It is there as a normalization factor, making sure that the sum of the discount factors is 1. The following result, known as the *one stage deviation principle* is useful to identify sub-game perfect Nash equilibria.

Proposition 6.9 *For finitely repeated games, a pure strategy profile $\hat{\boldsymbol{\alpha}} = (\hat{\alpha}^i_k)^{i=1,\ldots,N}_{k=0,\ldots,K}$ is a sub-game perfect Nash equilibrium if and only if*

(*) there is no player i and no strategy $\hat{\boldsymbol{\beta}} = (\hat{\beta}^i_k)^{i=1,\ldots,N}_{k=0,\ldots,K}$ that agrees with $\boldsymbol{\alpha}^i$ except for one single stage k_0 and such that $\boldsymbol{\beta}^i$ is a better response to $\hat{\boldsymbol{\alpha}}^{-i}$ tahn $\boldsymbol{\alpha}^i$ conditional on the history h^{k_0}.

Proof: The *Only if* part follows directly from the definition of sub-game perfection. So we only argue the *If* part. If $\hat{\boldsymbol{\alpha}}$ is not a sub-game perfect Nash equilibrium, there exists a stage k, a history h^k a player i and a strategy $\boldsymbol{\beta}^i$ which is a better response to $\hat{\boldsymbol{\alpha}}^{-i}$ than $\boldsymbol{\alpha}^i$ for the game startig at stage k with history h^k as conditioning. If we define

$$\hat{k} = \max\{k';\ \exists h^{k'}, \beta^i(h^{k'}) \neq \hat{\alpha}^i(h^{k'})\},$$

then assumption $*$ implies that $\hat{k} > k$. Now, let us consider the strategy $\tilde{\boldsymbol{\alpha}}^i$ defined by

$$\begin{cases} \tilde{\alpha}^i_\ell = \hat{\alpha}^i_\ell & for \ell < \hat{k}; \\ \tilde{\alpha}^i_\ell = \beta^i_\ell & for \ell \geqslant \hat{k}. \end{cases}$$

Since $\tilde{\boldsymbol{\alpha}}^i$ agrees with $\boldsymbol{\beta}^i$ from stage k onward, the one stage deviation principle which we assume implies that $\tilde{\boldsymbol{\alpha}}^i$ is as good a response as $\boldsymbol{\beta}^i$ in every subgame starting from $\hat{k}$, so $\tilde{\boldsymbol{\alpha}}^i$ is as good a response as $\boldsymbol{\beta}^i$ in the subgame starting at stage k with history h^k.

If $\hat{k} = k + 1$, then $\tilde{\boldsymbol{\alpha}}^i = \hat{\boldsymbol{\alpha}}^i$ which contradicts the hypothesis that $\boldsymbol{\beta}^i$ improves on $\hat{\boldsymbol{\alpha}}^i$. Now if $\hat{k} > k + 1$, we can construct a strategy that agrees with $\boldsymbol{\beta}^i$ until stage $k - 2$ and argues that it is as good as $\boldsymbol{\beta}^i$. Repeating the procedure over and over we end up with a contradiction which concludes the proof. □

Remark 6.10 *Note that finitely repeated games can have multiple Nash equilibria.*

Infinitely Repeated Games

When the one stage game is repeated infinitely many times, i.e. when $K = \infty$, the cost to player i can be expressed by

$$J^i(\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_K) = (1 - \delta) \sum_{j=0}^{\infty} \delta^j \, \tilde{J}^i(\boldsymbol{\alpha}_j).$$

These costs can be interpreted in the following way. If we denote by τ a geometric random variable with parameter $1 - \delta \in (0, 1]$,

$$\mathbb{P}(\tau = k) = \delta^k (1 - \delta), \quad for\, k = 0, 1, \ldots,$$

then for a sequence of strategy profiles $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_k)_{k=0,1,\ldots}$, the cost to player i can be interpreted as the expectation of the cost of the one-stage game at the random stage τ. Indeed:

$$\mathbb{E}[\tilde{J}^i(\boldsymbol{\alpha}_\tau)] = (1 - \delta) \sum_{j=0}^{\infty} \delta^j \, \tilde{J}^i(\boldsymbol{\alpha}_j).$$

Remark 6.11 *A result similar (though under some extra assumption) to the one-stage deviation principle stated and proved as Proposition 6.9 above can also be proven for infinite horizon repeated games. We shall not state it here because of its more technical nature.*

6.1.3 **The Repeated Prisoner's Dilemma Game**

Proposition 6.12 *If the prisoner dilemma game is repeated finitely many times, the strategy profile $((D, D), \ldots, (D, D))$ is the unique sub-game perfect Nash equilibrium.*

Proof: While we do not give a complete proof of this standard result, we highlight several important steps.

Step 1: We first check that if we are looking for a sub-game perfect Nash equilibrium, the last choice should necessarily be (D, D) whatever the history is.

Step 2: We complete the proof by an induction argument backward in time. The so-called *one stage deviation principle* for sub-game perfect Nash equilibria (stated and proved as Proposition 6.9) provides the induction part of the proof. It imples that when $K < \infty$, the strategy profile $\hat{\boldsymbol{\alpha}} = (\hat{\alpha}^i_k)_{i=1,2,\ k=0,\ldots,K-1}$ where $\boldsymbol{\alpha}_k = (\alpha^1_k, \alpha^2_k) = (D, D)$ for every $k = 0, 1, \ldots, K - 1$ is the only sub-game perfect Nash equilibrium. □

Remark 6.13 *The result is different when the number of stages is infinite, i.e. $K = \infty$. Indeed, using an appropriate version of the one-stage deviation principle it is possible to prove that when $K = \infty$ and $\delta > \frac{1}{2}$ the following strategy profile is a sub-game perfect Nash equilibrium. Each player follows the strategy described below:*

1. *she starts playing with action C;*
2. *she keeps on playing with action C until she sees an action D of the other player;*
3. *she switches to play with action D and keeps on playing with action D afterwards.*

So if the game has to be played infinitely many times, then cooperation *may be observed for a long period of time.*

6.2 Stochastic Games

The concept of a stochastic game can be viewed as a generalization of the one of repeated game, as well as a generalization of the notion of Markov Decision Processes (MDP for short) to the case of multiple controllers.

We switch from the terminology of *stage* to *time* and accordingly, we use the notation t instead of k. We consider first the case of discrete time $t = 0, 1, \ldots, T$, and we focus on three different settings:

- the finite horizon problem with $T < \infty$;
- the infinite horizon problem ($T = \infty$) with a discount factor $0 \leqslant \delta < 1$;
- the ergodic problem ($T \to \infty$).

Typically, a stochastic game consists of the following components:

1. A finite number of players, say N, and we denote by $[N] = \{1, \ldots, N\}$ the set of players.
2. A state space which we take as a measurable space $(X, \mathcal{X})$. States $x \in X$ capture information from the history and the environment. We also want the set of feasible actions to depend upon the current state. Also, the costs depend upon the state and the actions of the players.
3. For every player $i \in [N]$ we denote by $A^i(x)$ the set of feasible actions for player i when the state is $x \in X$. We assume that $A^i(x)$ is a measurable subset of a measurable space $(A^i, \mathcal{A}^i)$. In many applications, A^i is supposed to be a compact space for mathematical convenience.
   For every $x \in X$, we denote by $A(x) = A^1(x) \times \ldots \times A^N(x)$ be the set of strategy profiles $\boldsymbol{\alpha} = (\alpha^1, \ldots, \alpha^N)$ when the state is $x \in X$. We also denote by $A = A^1 \times \ldots \times A^N$. and we equip it with the product σ-field $\mathcal{A} = \mathcal{A}^1 \times \ldots \times \mathcal{A}^N$.
4. For every $i \in [N]$, the running cost function for player i in state $x \in X$ and for a strategy profile $\boldsymbol{\alpha} \in A(x)$ is given by a function
   $$f^i : X \times A \ni (x, \alpha^1, \ldots, \alpha^N) \mapsto f^i(x, \boldsymbol{\alpha}) \in \mathbb{R}$$
   In most cases, we shall assume that f^i is bounded for every $i \in [N]$.

For finite horizon problems, we can also define a terminal cost function for player i, denoted by $g^i : X \to \mathbb{R}$ for example. And again for the sake of convenience, we shall also assume that g^i is bounded for every $i \in [N]$.

5. The time evolution of the state given by a process of random variables $(x_t)_{t \geqslant 0}$, x_0 having a prescribed distribution, say $\mu_0 \in \mathcal{P}(X, \mathcal{X})$, and for every $t \geqslant 0$, x_{t+1} stands for the new state. It has a distribution depending only on x_t and the actions $\boldsymbol{\alpha}_t = (\alpha_t^1, \ldots, \alpha_t^N) \in A(x_t)$ taken by all the players at time t. The stochastic mechanism driving the evolution of the state is given by a transition probability function $P : X \times A^1 \times \ldots \times A^n \to \mathcal{P}(X)$:

$$P(x, \boldsymbol{\alpha}, dx') = P(x, \alpha^1, \ldots, \alpha^N, dx') := \mathbb{P}[dx' | x, \alpha^1, \ldots, \alpha^N]$$

giving the probability that the next state is in dx' conditioned by the fact that the current state is x and the players chose the actions $\boldsymbol{\alpha}$.

We need to define precisely at each time, the kind of information each player can access to take action. This should make clear what $\boldsymbol{\alpha}$ depends upon, and help us decide whether or not the process $(x_t)_{t \geqslant 0}$ is Markovian. But first, we adapt the notion of history to the present set-up.

Definition 6.14 *We define the set of histories up to time t, denoted by $\mathcal{H}^t$, as the collection of state action profiles up to time t. More precisely,*

$$\mathcal{H}^t = \{(x_0, \boldsymbol{\alpha}_0, x_1, \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_{t-1}, x_t)\}$$

where $x_s \in X$ for $s \in 0, 1, \ldots, t$ and $\boldsymbol{\alpha}_s \in A(x_s)$ for $s = 1, \ldots, t-1$.

We can now define the types of strategies used by the players in a stochastic game.

Definition 6.15 *A pure strategy for player $i \in [N]$ is a function $\pi^i : \mathbb{N} \times \mathcal{H}^\infty \to A^i$ such that:*

1. *π^i is non-anticipative in the sense that $\pi^i(t, h)$ depends only on $h^t = (x_0, \boldsymbol{\alpha}_0, x_1, \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_{t-1}, x_t)$ for every $(t, h) \in \mathbb{N} \times H^\infty$.*
2. *$\pi^i(t, h) \in A^i(x_t)$.*

A mixed strategy for player i is a function $\pi^i : \mathbb{N} \times \mathcal{H}^\infty \to \mathcal{P}(A^i)$ such that:

1. *π^i is non-anticipative;*
2. *$\pi^i(t, h)$ with $h^t = (x_0, \boldsymbol{\alpha}_0, x_1, \ldots, \boldsymbol{\alpha}_{t-1}, x_t)$ is concentrated on $A^i(x_t)$ in the sense that*

$$\pi^i(t, h)(A^i(x_t)) = 1.$$

A pure stationary strategy for player i is a function $\pi^i : X \to A^i$ such that $\pi^i(x) \in A^i(x)$ for all $x \in X$. Similarly, a mixed stationary strategy is a function $\pi^i : X \to \mathcal{P}(A^i)$ such that for every $x \in X$, $\pi^i(x)[A^i(x)] = 1$.

The above is a pileup of notations and definitions. Next we show that they are not unreasonable in the sense that mathematical objects satisfying all these requirements do exist.

6.2.1 Set-up for Pure Strategies

Given a pure strategy profile $\pi = (\pi_t)_{t \geqslant 0}$ as defined above, where $\pi_t = (\pi_t^1, \ldots \pi_t^N)$ and $\pi_t^i := \pi^i(t, \cdot)$ for every $i \in [N]$ and $t \geqslant 0$, we would like to know if there exists a stochastic process $(x_t)_{t \geqslant 0}$ such that for every $t \geqslant 0$, we have:

$$x_{t+1} \sim P(x_t, \boldsymbol{\alpha}_t)$$

where $P : X \times A^1 \times \ldots \times A^n \to \mathcal{P}(X)$ is the prescribed transition probability function. and $\boldsymbol{\alpha}_t = (\alpha_t^1, \ldots, \alpha_t^n)$ with $\alpha_t^i = \pi_t^i(h^t)$ for every $i \in [N]$ and $t \geqslant 0$.

One way to do just that would be, given $\pi = (\pi_t)_{t \geqslant 0}$ and $P : X \times A \to \mathcal{P}(X)$, to construct a probability distribution denoted by $\mathbb{P}$ for example, on the product space $X^{\mathbb{N}}$ equipped with its product σ-filed $\mathcal{X}^{\mathbb{N}}$, in such a way that the coordinate process $(x_t)_{t \geqslant 0}$ satisfies the desired properties. In particular, it should be such that for every $t \geqslant 0$, conditioned on $(x_t, \pi_t(h^t))$, the marginal distribution of x_{t+1} is given by the probability distribution $P(x_t, \pi_t(h^t))(\cdot)$, namely

$$\mathbb{P}[\cdot\ | x_t, \pi_t(h^t)] = P(x_t, \pi_t(h^t))(\cdot) \in \mathcal{P}(X).$$

To achieve this, we start with an initial distribution $\mu_0 \in \mathcal{P}(X)$ (for example $\mu_0 = \delta_{x_0}$).

- For $t = 0$, we have $x_0 \sim \mu_0$ and we denote it by $\mathbb{P}_{x_0}$.
- For $t = 1$, we define a distribution $\mathbb{P}_{(x_0,x_1)} \in \mathcal{P}(X \times X)$ by:

  $$\mathbb{P}_{(x_0,x_1)}(dx_0, dx_1) = \mu_0(dx_0) P(x_0, \pi_0(x_0))(dx_1).$$

  This means that, for every measurable sets $B_0, B_1 \in \mathcal{X}$, we have

  $$\mathbb{P}_{(x_0,x_1)}(B_0 \times B_1) = \int_{B_0} \mu_0(dx_0) \int_{B_1} P(x_0, \pi_0(x_0))(dx_1).$$

- For $t = 2$, we define the distribution $\mathbb{P}_{(x_0,x_1,x_2)} \in \mathcal{P}(X \times X \times X)$ by:

  $$\mathbb{P}_{(x_0,x_1,x_2)}(dx_0, dx_1, dx_2) = \mu_0(dx_0) P(x_0, \pi_0(x_0))(dx_1) P(x_1, \pi_1(x_0, \pi_0(x_0), x_1))(dx_2).$$

- For $t = 3$, we define a distribution $\mathbb{P}_{(x_0,x_1,x_2,x_3)} \in \mathcal{P}(X \times X \times X \times X)$ by:

  $$\mathbb{P}_{(x_0,x_1,x_2,x_3)}(dx_0, dx_1, dx_2, dx_3) = \mathbb{P}_{(x_0,x_1,x_2)}(dx_0, dx_1, dx_2) \cdot P(x_2, \pi_2(h^2))(dx_3)$$

  where $h^2 = (x_0, \pi_0(x_0), x_1, \pi_1(x_0, \pi_0(x_0), x_1), x_2)$.
- and so on, and so on, for $t \geqslant 4$.

We can check that the collection of probability distributions so constructed is consistent. Hence, we can apply Kolmogorov's extension theorem and conclude that there exists a probability measure $\mathbb{P}$ on the product space $X^{\mathbb{N}}$ equipped with its product σ-field $\mathcal{X}^{\mathbb{N}}$ such that the coordinate process $(x_t)_{t \geqslant 0}$ has the desired properties.

Remark 6.16 *If the pure strategy profile π is stationary in the sense that for every $t \geqslant 0$ it is such that*

$$\pi_t(h^t) = \pi_t(x_t) = (\pi_t^1(x_t), \ldots, \pi_t^N(x_t)),$$

then $\pi_1(x_0, \pi_0(x_0), x_1) = \pi_1(x_1)$ and $\pi_2(h^2) = \pi_2(x_2)$, etc, and the state process $(x_t)_{t \geqslant 0}$ is Markovian.

6.2.2 Set-up for Mixed Strategies

We address the existence of the state process similarly. For any mixed strategy profile $\pi = (\pi_t)_{t \geqslant 0}$ where $\pi^i_t(\cdot) = \pi^i(t, \cdot) : \mathcal{H}^t \to \mathcal{P}(A^i)$ for every $i \in [N]$ and $t \geqslant 0$, we can construct a probability distribution $\mathbb{P}$ on the product space $(X \times A)^{\mathbb{N}}$ such that the coordinate process $(x_t, \boldsymbol{\alpha}_t)_{t \geqslant 0}$ has the properties required to describe the time evolution of the state and the decisions of the players.

1. For $t = 0$, we start from an initial distribution $\mu_0 \in \mathcal{P}(X)$ such that $x_0 \sim \mu_0$ and we define a distribution in $\mathcal{P}(X \times A)$ by:
$$\mu_0(dx_0)\pi_0(x_0)(d\boldsymbol{\alpha}_0),$$
where $\pi_0(x)(d\boldsymbol{\alpha}) = \prod_{j=1}^{n} \pi_0^j(x)(d\alpha^j)$ for every $x \in X$. Notice the product nature of this last measure. This is required by the fact that in models with mixed strategies, the players randomized their actions *independently of each other*.
2. For $t = 1$, we define a probability distribution in $\mathcal{P}(X \times A \times X)$ by:
$$\mu_0(dx_0)\pi_0(x_0)(d\boldsymbol{\alpha}_0)P(x_0, \boldsymbol{\alpha}_0)(dx_1)$$
And we also define a probability distribution in $\mathcal{P}(X \times A \times X \times A)$ by:
$$\mu_0(dx_0)\pi_0(x_0)(d\boldsymbol{\alpha}_0)P(x_0, \boldsymbol{\alpha}_0)(dx_1)\pi_1(x_0, \alpha_0, x_1).$$
3. For $t = 2$, we define a probability distribution in $\mathcal{P}(X \times A \times X \times A \times X)$ by:
$$\mu_0(dx_0)\pi_0(x_0)(d\boldsymbol{\alpha}_0)P(x_0, \boldsymbol{\alpha}_0)(dx_1)\pi_1(x_0, \boldsymbol{\alpha}_0, x_1)(d\boldsymbol{\alpha}_1)P(x_1, \boldsymbol{\alpha}_1)(dx_2)$$
etc, and as before, we use Kolmogorov's extension theorem to conclude the existence of the desired probability measure $\mathbb{P}$.

Remark 6.17 *In some applications, the state space naturally appears as a product $X = X^1 \times \ldots \times X^N$, so that the states are of the form $x = (x^1, \ldots, x^N)$. In this case, the component x^i can be interpreted as the private state of player i.*

6.2.3 The Cost Functions

Whether the model requires the use of pure strategies or mixed strategies, the cost functions are defined as follows.

Definition 6.18 • ***Finite horizon*** *($T < \infty$): For every $i \in [N]$, the finite horizon cost function for player i is given by a function J^i depending on the initial distribution $\mu_0 \in \mathcal{P}(X)$ and a strategy profile $\pi = (\pi^1, \ldots, \pi^N)$ with $\pi^i : \mathbb{N} \times \mathcal{H}^T \to A^i$ such that*
$$J^i(\mu_0, \pi) = \mathbb{E}\left[\sum_{t=0}^{T-1} f^i(x_t, \boldsymbol{\alpha}_t) + g^i(x_T)\right]$$
where the process $(x_t, \boldsymbol{\alpha}_t)_{t \geqslant 0}$ is the process whose existence was argued earlier. Here, $g^i : X \to \mathbb{R}$ is a bounded function giving the terminal cost for player i.

- ***Infinite horizon** ($T = \infty$) **with discount** ($\delta \in [0,1)$): For every $i \in [N]$, the infinite horizon discounted cost function for player i associated to an initial distribution μ_0 and a strategy profile $\pi = (\pi^1, \ldots, \pi^n)$ is defined by*

$$J^i(\mu_0, \pi) = (1-\delta)\mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t f^i(x_t, \boldsymbol{\alpha}_t)\right].$$

- ***Ergodic**: For every $i \in [N]$, the ergodic cost function for player i associated to an initial distribution μ_0 and a stationary strategy profile $\pi = (\pi^1, \ldots, \pi^N)$ with $\pi^i : X \to A^i$ (or a mixed stationary strategy profile π with $\pi^i : X \to \mathcal{P}(A^i)$) is defined by*

$$J^i(\mu_0, \pi) = \liminf_{T \to \infty} \frac{1}{T}\mathbb{E}\left[\sum_{t=0}^{T-1} f^i(x_t, \boldsymbol{\alpha}_t)\right]$$

for the process $(x_t, \boldsymbol{\alpha}_t)_{t \geqslant 0}$ constructed earlier.

Remark 6.19 *1. In the infinite horizon discounted case, the boundedness of the function f^i guarantees the convergence of the infinite series and the finiteness of the expectation as well.*

2. In the ergodic setting, we wish that the cost could be defined as:

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} f^i(x_t, \boldsymbol{\alpha}_t)$$

but there is no guarantee that this limit exists, and even if it did, it could very well be random. So we not only take the expectation to remove the randomness in the cost, but we also use $\liminf$ to be sure that the cost is well defined.

7

Stochastic Differential Equations

The purpose of this chapter is to introduce enough of the classical theory of Wasserstein space and stochastic differential equations to prepare for the analysis of anonymous stochastic differential games with mean field interactions.

7.1 Notation and First Definitions

We assume that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a stochastic basis where the filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leqslant t \leqslant T}$ supports an m-dimensional $\mathbb{F}$-Brownian motion $\mathbf{W} = (W_t)_{0 \leqslant t \leqslant T}$ in $\mathbb{R}^m$. For each integer $k \geqslant 1$, we denote by $\mathbb{H}^{0,k}$ the collection of all $\mathbb{R}^k$-valued progressively measurable processes on $[0, T] \times \mathbb{R}$, and we introduce the subspaces:

$$\mathbb{H}^{2,k} := \Big\{ Z \in \mathbb{H}^{0,k};\ \mathbb{E} \int_0^T |Z_s|^2 ds < \infty \Big\}$$

and using the square root of the expectation appearing in the above definition as a norm, the space $\mathbb{H}^{2,k}$: becomes a Hilbert space, the inner product being obtained from this norm by polarization. We shalll also use the following space

$$\mathbb{S}^2 := \Big\{ Y \in \mathbb{H}^{0,k};\ \mathbb{E} \sup_{0 \leqslant t \leqslant T} |Y_s|^2 < \infty \Big\}.$$

Equipped with the norm given by the square root of the expectation appearing in its definition, and by the scalar product obtained by thepolarization identity, the space $\mathbb{H}^{2,k}$ becomes a separable Hilbert space. Similarly, $\mathbb{S}$ becomes a Banach space when equipped with the norm given by the square root of the expectation appearing in its definition. We may also use the notation $\mathbb{B}^k$ for the subspace of bounded processes, namely:

$$\mathbb{B}^k := \Big\{ Z \in \mathbb{H}^{0,k};\ ,\ \sup_{0 \leqslant t \leqslant T} |Z_t| < \infty,\ \mathbb{P} - a.s. \Big\}.$$

In the case of scalar processes, when $k = 1$, we skip the exponent k from our notation. We are interested in stochastic differential equations (SDEs) of the form

$$dX_t = b(t, X_t)dt + \sigma(t, X_t)dW_t \tag{7.1}$$

where the coefficients b and σ

$$(b, \sigma) :\ [0, T] \times \Omega \times \mathbb{R}^d \to \mathbb{R}^d \times \mathbb{R}^{d \times m}$$

satisfy the following assumptions.

(A1) For each $x \in \mathbb{R}^d$, the processes $(b(t,x))_{0\leqslant t\leqslant T}$ and $(\sigma(t,x))_{0\leqslant t\leqslant T}$ are in $\mathbb{H}^{2,d}$ and $\mathbb{H}^{2,dm}$ respectively;

(A2) $\exists c > 0$, $\forall t \in [0,T]$, $\forall \omega \in \Omega$, $\forall x, x' \in \mathbb{R}^d$,

$$|(b,\sigma)(t,\omega,x) - (b,\sigma)(t,\omega,x')| \leqslant c|x - x'|.$$

As most probabilists do, we shall refrain from making the dependence upon $\omega \in \Omega$ explicit whenever possible.

Definition 7.1 *We say that an $\mathbb{F}$-progressively measurable process $\mathbf{X} = (X_t)_{0\leqslant t\leqslant T}$ is a strong solution of the SDE* (7.1) *if*

- $\int_0^T (|b(t,X_t)| + |\sigma(t,X_t)|^2)dt < \infty$ *$\mathbb{P}$-almost surely,*
- $X_t = X_0 + \int_0^t b(s,X_s)ds + \int_0^T \sigma(s,X_s)dW_s, \quad 0 \leqslant t \leqslant T.$

7.2 Existence and Uniqueness of Strong Solutions: the Lipschitz Case

Theorem 7.2 *Let us assume that $X_0 \in L^2$ is independent of $\mathbf{W}$ and that the coefficients b and σ satisfy the assumptions **(A1)** and **(A2)** above. Then, there exists a unique solution of equation* (7.9) *in $\mathbb{H}^{2,d}$, and for some $c > 0$ depending only upon T and the Lipschitz constant of b and σ, this solution satisfies*

$$\mathbb{E} \sup_{0\leqslant t\leqslant T} |X_t|^2 \leqslant c(1 + \mathbb{E}|X_0|^2)e^{cT}. \tag{7.2}$$

Here and in the following, we use the letter c for a generic constant which can change from line to line.

Proof: For each $\mathbf{X} \in \mathbb{H}^{2,d}$ the space of square integrable progressively measurable processes, we define the process $U(\mathbf{X})$ by:

$$U(\mathbf{X})_t = X_0 + \int_0^t b(s,X_s)ds + \int_0^t \sigma(s,X_s)dW_s. \tag{7.3}$$

First we prove that $U(\mathbf{X}) \in \mathbb{H}^{2,d}$, and then, since $\mathbf{X}$ is a solution of the SDE (7.1) if and only if $U(\mathbf{X}) = \mathbf{X}$, we prove that U is a strict contraction in the Hilbert space $\mathbb{H}^{2,d}$. By definition of the norm of $\mathbb{H}^{2,d}$ we have

$$\|U(\mathbf{X})\|^2 \leqslant (i) + (ii) + (iii)$$

with

$$(i) = 3T\mathbb{E}|X_0|^2 < \infty,$$

and (ii) and (iii) defined below. Using the fact

$$|b(t,x)|^2 \leqslant c(1 + |b(t,0)|^2 + |x|^2)$$

implied by our Lipschitz assumption, we have:

$$
\begin{aligned}
(ii) &= 3\mathbb{E}\int_0^T \left|\int_0^t b(s,X_s)ds\right|^2 dt \\
&\leqslant 3\mathbb{E}\int_0^T t\left(\int_0^t |b(s,X_s)|^2 ds\right)dt \\
&\leqslant 3c\mathbb{E}\int_0^T t\left(\int_0^t (1+|b(s,0)|^2+|X_s|^2)ds\right)dt \\
&\leqslant 3cT^2\left(1+\|b(\,\cdot\,,0)\|^2+\mathbb{E}\int_0^T |X_t|^2\right) \\
&< \infty
\end{aligned}
$$

if we use assumption **(A1)** which says that the norm of $(b(t,x))_{0\leqslant t\leqslant T}$ in $\mathbb{H}^{2,d}$ is finite and the assumption that $\boldsymbol{X} = (X_t)_{0\leqslant t\leqslant T}$ is in $\mathbb{H}^{2,d}$. Finally, in order to estimate

$$
(iii) = 3\mathbb{E}\int_0^T \left|\int_0^t \sigma(s,X_s)dW_s\right|^2 dt
$$

we use Doob's maximal inequality (or its generalization in the form of the Burkhölder-Davis-Gundy (BDG) inequality which we recall below) together with the fact that

$$
|\sigma(t,x)|^2 \leqslant c(1+|\sigma(t,0)|^2+|x|^2)
$$

implied by our Lipschitz assumption on σ. Doing so, we have:

$$
\begin{aligned}
(iii) &\leqslant 3T\mathbb{E}\sup_{0\leqslant t\leqslant T}\left|\int_0^t \sigma(s,X_s)dW_s|^2\right| dt \\
&\leqslant 12T\mathbb{E}\int_0^T |\sigma(s,X_s)|^2 ds \\
&\leqslant 12Tc\mathbb{E}\int_0^T (1+|\sigma(s,0)|^2+|X_s|^2)ds \\
&< \infty
\end{aligned}
$$

for the same reasons as above. Again, remember that the value of the generic constant c can change from line to line. Now that we know that U maps $\mathbb{H}^{2,d}$ into itself, we prove that it is a strict contraction. In order to do so, we find convenient to change the norm of the Hilbert space $\mathbb{H}^{2,d}$ to an equivalent norm. For each $\alpha > 0$ we define a norm on the space $\mathbb{H}^{2,d}$ by:

$$
\|\boldsymbol{\xi}\|_\alpha^2 = \mathbb{E}\int_0^T e^{-\alpha t}|\xi_t|^2 dt.
$$

The norm $\|\cdot\|_\alpha$ and the original norm $\|\cdot\|$ (which correspond to $\alpha = 0$) are equivalent and define the same topology. If $\mathbf{X}$ and $\mathbf{Y}$ are generic elements of $\mathbb{H}^{2,d}$ with $X_0 = Y_0$, we have

$$
\begin{aligned}
&\mathbb{E}|U(X)_t - U(Y)_t|^2 \\
&\quad\leqslant 2\mathbb{E}\left|\int_0^t [b(s,X_s)-b(s,Y_s)]ds\right|^2 + 2\mathbb{E}\left|\int_0^t [\sigma(s,X_s)-\sigma(s,Y_s)]dW_s\right|^2 \\
&\quad\leqslant 2t\mathbb{E}\int_0^t |b(s,X_s)-b(s,Y_s)|^2 ds + 8\mathbb{E}\int_0^t |\sigma(s,X_s)-\sigma(s,Y_s)|^2 ds \\
&\quad\leqslant ct\int_0^t \mathbb{E}|X_s-Y_s|^2 ds
\end{aligned}
$$

if we use the Lipschitz property of the coefficients. Consequently:

$$\begin{aligned}\|U(\mathbf{X}) - U(\mathbf{Y})\|_\alpha^2 &= \int_0^T e^{-\alpha t}\mathbb{E}|U(\mathbf{X})_t - U(\mathbf{Y})_t|^2 dt \\ &\leqslant cT \int_0^T e^{-\alpha t} \int_0^t \mathbb{E}|X_s - Y_s|^2 ds\, dt \\ &\leqslant cT \int_0^T \mathbb{E}|X_s - Y_s|^2 ds \int_t^T e^{-\alpha t} dt \\ &\leqslant \frac{cT}{\alpha}\|\mathbf{X} - \mathbf{Y}\|_\alpha^2,\end{aligned}$$

and U is indeed a strict contraction if $\alpha > cT$ is large enough! Finally, we prove the estimate (7.2) for the solution. For $t \in [0, T]$ fixed we have:

$$\begin{aligned}\mathbb{E}\sup_{0\leqslant s\leqslant t}|X_s|^2 &= \mathbb{E}\sup_{0\leqslant s\leqslant t}\left|X_0 + \int_0^s b(r, X_r)dr + \int_0^s \sigma(r, X_r)dW_r\right|^2 \\ &\leqslant 3\left(\mathbb{E}|X_0|^2 + t\mathbb{E}\int_0^t |b(s, X_s)|^2 ds + 4\mathbb{E}\int_0^t |\sigma(s, X_s)|^2 ds\right) \\ &\leqslant c\left(1 + \mathbb{E}|X_0|^2 + \int_0^t \mathbb{E}\sup_{0\leqslant r\leqslant s}|X_r|^2 dr\right)\end{aligned}$$

where the constant c depends only upon T, $\|b(\,\cdot\,, 0)\|^2$ and $\|\sigma(\,\cdot\,, 0)\|^2$, and finally we conclude using Gronwall inequality. □

For the sake of completeness we state the version of Gronwall and BDG inequalities which we use throughout.

Remark 7.3 ***Gronwall's inequality***. *Since we will use Gronwall's inequality repeatedly in the sequel, we state it for later reference:*

$$\varphi(t) \leqslant \alpha + \int_0^t \beta(s)\varphi(s)ds \quad \Longrightarrow \quad \varphi(t) \leqslant \alpha e^{\int_0^t \beta(s)ds}. \tag{7.4}$$

Remark 7.4 ***BDG inequality***. *For each $p \in (0, \infty)$ there exist (universal) constants c_p and C_p such that, for each continuous time martingale $\boldsymbol{M} = (M_t)_{0\leqslant t\leqslant T}$ such that $M_0 = 0$ we have*

$$c_p\mathbb{E}[< M >_T^{p/2}] \leqslant \mathbb{E}[\sup_{0\leqslant t\leqslant T}|M_t|^p] \leqslant C_p\mathbb{E}[< M >_T^{p/2}]. \tag{7.5}$$

We shall use this inequality when $M_t = \int_0^t \xi_s dW_s$ for a adapted square integrable integrand, in which case $< M >_T = \int_0^T |\xi_t|^2 dt$.

7.3 Spaces of Probability Measures

Our goal is to study stochastic differential equations with coefficients which can depend upon probability measures, and in particular, the marginal distributions of the solutions. As we show later on, this type of nonlinearity occurs in the asymptotic regime of large stochastic systems with mean field interactions.

7.3.1 **Notation**

For any given measurable space $(E, \mathcal{E})$, we denote by $\mathcal{P}(E)$ the set of probability measures on $(E, \mathcal{E})$. In the sequel, we shall most often assume that E is a complete separable metric space, e.g. $E = \mathbb{R}^d$ or a separable Banach space with norm $\| \cdot \|$. In such a case, the σ fiedl $\mathcal{E}$ is taken to be the Borel σ-field of E. For $p \geqslant 1$, we denote by $\mathcal{P}_p(E)$ the set of probability measures with finite second order moment (i.e. the set of those $\mu \in \mathcal{P}(E)$ for which $\int_E d(x, x_0)^p \mu(dx) < \infty$) for some (or equivalently any) $x_0 \in E$. Obviously, we use the notation d for the distance of the metric space E. For μ and ν in $\mathcal{P}_p(E)$, the p-Wasserstein distance between μ and ν is defined by the formula

$$W^{(p)}(\mu, \nu) = \inf \left\{ \left[\int d(x, y)^p \pi(dx, dy) \right]^{1/p} ;\ \pi \in \mathcal{P}(E \times E) \text{ with marginals } \mu \text{ and } \nu \right\}. \tag{7.6}$$

Any probability measure $\pi \in \mathcal{P}(E \times E)$ with marginals μ and ν is called a coupling of μ and ν. When E is a complete separable metric space (one often says E is a Polish space) the space $\mathcal{P}(E)$ is typically equipped with the topology of the *weak convergence* for which $\mathcal{P}(E)$ is also a Polish space. This topology is in fact the topology of the convergence in the sense of the modified Wasserstein distance $W^{(0)}$ on $\mathcal{P}(E)$ defined as

$$W^{(0)}(\mu, \nu) = \inf \left\{ \int 1 \wedge d(x, y)\ \pi(dx, dy);\ \pi \in \mathcal{P}(E \times E) \text{ with marginals } \mu \text{ and } \nu \right\}. \tag{7.7}$$

Notice that the distance $1 \wedge d(x, y)$ is bounded and cannot feel if the probability measures have finite moments. In fact, the convergence in the sense of the p-Wasserstein distances $W^{(p)}$ is equivalent to the weak convergence of measures together with the convergence of all moments of order up to p.

In what follows, we shall almost exclusively use the distance $W^{(2)}$ which we call the Wasserstein distance on the space $\mathcal{P}_2(E)$ which we call the Wasserstein space.

In order to formulate the results of this part of the chapter we need to give a meaning to the continuity and Lipschitz property of functions of probability measures on the space $E = C([0, T]; \mathbb{R}^d)$ of $\mathbb{R}^d$-valued bounded continuous functions on $[0, T]$ equipped with the norm of the uniform convergence. A notion of distance on the space of such measures is all we need for now. For the sake of definiteness, we rewrite the definitions of the Wasserstein distances $W^{(0)}$ and $W^{(2)}$ in the particular case of the complete metric space $E = C([0, T]; \mathbb{R}^d)$.

$$\begin{aligned} W^{(0)}(m_1, m_2) = \inf \Big\{ \int \sup_{0 \leqslant t \leqslant T} 1 \wedge |X_t(w_1) - X_t(w_2)|\ m(dw_1, dw_2); \\ m \in \mathcal{P}(C([0, T]; \mathbb{R}^d) \times C([0, T]; \mathbb{R}^d)),\ \text{with marginals } m_1 \text{ and } m_2 \Big\} \end{aligned}$$

where $(X_t)_{0 \leqslant t \leqslant T}$ is the canonical process $w \hookrightarrow X_t(w) = w(t)$ on $\mathcal{P}(C([0, T]; \mathbb{R}^d))$. $W^{(0)}$ is a distance on $\mathcal{P}(C([0, T]; \mathbb{R}^d))$ which defines a notion of convergence equivalent to the weak convergence of probability measures. So $\mathcal{P}(C([0, T]; \mathbb{R}^d))$ is a complete metric space for $W^{(0)}$. Similarly, we define $\mathcal{P}_2(C([0, T]; \mathbb{R}^d)$ as the subset of $\mathcal{P}(C([0, T]; \mathbb{R}^d)$ of the measures m with a finite moment of order 2, namely satisfying

$$\int \sup_{0\leqslant t\leqslant T} |X_t(w)|^2 m(dw) < \infty.$$

The corresponding Wasserstein distance $W^{(2)}$ can be defined as:

$$W^{(2)}(m_1,m_2) = \inf\Bigg\{ \Bigg[\int \sup_{0\leqslant t\leqslant T} |X_t(w_1) - X_t(w_2)|^2 \; m(dw_1,dw_2) \Bigg]^{1/2};$$
$$m \in \mathcal{P}_2(C([0,T];\mathbb{R}^d) \times C([0,T];\mathbb{R}^d)), \text{ with marginals } m_1 \text{ and } m_2 \Bigg\}$$

and for later convenience we introduce the notation

$$W_t^{(2)}(m_1,m_2) = \inf\Bigg\{ \Bigg(\int \sup_{0\leqslant s\leqslant t} |X_s(w_1) - X_s(w_2)|^2 \; m(dw_1,dw_2) \Bigg)^{1/2};$$
$$m \in \mathcal{P}_2(C([0,T];\mathbb{R}^d) \times C([0,T];\mathbb{R}^d)), \text{ with marginals } m_1 \text{ and } m_2 \Bigg\},$$

so that $W^{(2)}(m_1,m_2) = W_T^{(2)}(m_1,m_2)$.

7.3.2 Example: Stochastic System with Mean Field Interactions

Let us consider the system of N stochastic differential equations with mean field interactions:

$$dX_t^i = b^i\Big(t, X_t^i, \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}\Big)dt + \sigma^i\Big(t, X_t^i, \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}\Big)dW_t^i, \qquad i = 1,\cdots,N \tag{7.8}$$

where the $\boldsymbol{W}^i = (W_t^i)_{0\leqslant t\leqslant T}$ are independent Wiener processes for $i = 1,\cdots,N$, and the $(b^i,\sigma^i) : [0,T]\times\Omega\times\mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}^d\times\mathbb{R}^{d\times d}$ satisfy assumption **(A1)** of the previous section and are uniformly Lipschitz in the sense that:

$$\|b^i(t,x,\mu) - b^i(t,x',\mu')\|^2_{\mathbb{R}^d} + \|\sigma^i(t,x,\mu) - \sigma^i(t,x',\mu')\|^2_{\mathbb{R}^{d\times d}}$$
$$\leqslant \Big(\|x-x'\|^2_{\mathbb{R}^d} + W^{(2)}(\mu,\mu')^2\Big), \qquad i = 1,\cdots,N,$$

for a constant $c > 0$ independent of $t \in [0,T]$ and $\omega \in \Omega$. In order to apply the existence and uniqueness result proven earlier, we recast the system as an SDE in $\mathbb{R}^{Nd}$ by setting

$$\boldsymbol{X}_t = \begin{bmatrix} X_t^1 \\ \vdots \\ X_t^N \end{bmatrix}, \qquad B(t,\boldsymbol{x}) = \begin{bmatrix} b^1\Big(t, x^1, \frac{1}{N}\sum_{j=1}^N \delta_{x^j}\Big) \\ \vdots \\ b^N\Big(t, x^N, \frac{1}{N}\sum_{j=1}^N \delta_{x^j}\Big) \end{bmatrix}, \qquad \boldsymbol{W}_t = \begin{bmatrix} W_t^1 \\ \vdots \\ W_t^N \end{bmatrix},$$

and $\Sigma(t,\boldsymbol{x})$ equal to the $(Nd)\times(Nd)$ block diagonal matrix with i-th diagonal block given by $\sigma^i\Big(t, X_t^i, \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}\Big)$. Since the system (7.8) can be rewritten as the SDE

$$d\boldsymbol{X}_t = B(t,\boldsymbol{X}_t)dt + \Sigma(t,\boldsymbol{X}_t)d\boldsymbol{W}_t,$$

in order to prove existence and uniqueness of a strong solution we only need to prove the Lipschitz condition **(A2)**. The latter is an immediate consequence of the Lipschitz assumption on the coefficients b^i and σ^i and the fact that by definition

$$W^{(2)}\Big(\frac{1}{N}\sum_{j=1}^{N}\delta_{x^j}, \frac{1}{N}\sum_{j=1}^{N}\delta_{x'^j}\Big)^2 \leqslant \frac{1}{N}\sum_{j=1}^{N}\|x^j - x'^j\|^2$$

as shown by using the coupling $\pi = \frac{1}{N}\sum_{j=1}^{N}\delta_{(x^j, x'^j)}$.

7.4 SDEs of McKean-Vlasov Type

We now introduce the assumptions on the coefficients of the *non-linear* stochastic differential equations which we study in this chapter. As in the classical case studied earlier, we assume that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a stochastic basis where the filtration $\mathbb{F} = (\mathcal{F}_t)_{0\leqslant t\leqslant T}$ supports a $\mathbb{F}$-Brownian motion $\mathbf{W} = (W_t)_{0\leqslant t\leqslant T}$ in $\mathbb{R}^d$. We are interested in stochastic differential equations of the form

$$dX_t = b(t, X_t, \mathcal{L}(X_t))dt + \sigma(t, X_t, \mathcal{L}(X_t))dW_t \tag{7.9}$$

where the coefficients b and σ

$$(b, \sigma) : [0, T] \times \Omega \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}^d \times \mathbb{R}^{d\times m}$$

satisfy the following assumptions.

(A1) For each $x \in \mathbb{R}^d$ and $\mu \in \mathcal{P}(\mathbb{R}^d)$, the processes $(b(t, x, \mu))_{0\leqslant t\leqslant T}$ and $(\sigma(t, x, \mu))_{0\leqslant t\leqslant T}$ are in $\mathbb{H}^{2,d}$ and $\mathbb{H}^{2,dm}$ respectively;
(A2) $\exists c > 0$, $\forall t \in [0, T]$, $\forall \omega \in \Omega$, $\forall x, x' \in \mathbb{R}^d$, $\forall \mu, \mu' \in \mathcal{P}_2(\mathbb{R}^d)$,

$$|(b, \sigma)(t, \omega, x, \mu) - (b, \sigma)(t, \omega, x', \mu')| \leqslant c(|x - x'| + W^{(2)}(\mu, \mu'))$$

Here and in the following, we use the notation $\mathbb{P}_X$ or $\mathcal{L}(X)$ for the distribution or law of the random element X.

7.4.1 Examples of Mean Field Interactions

In practical applications, interactions through the marginal distribution of the process as is the case of the SDE of McKean-Vlasov type (7.9) come in various forms of complexity.

In the simplest case, which we shall call *mean field interaction of scalar type*, the dependence upon the distribution degenerates into a dependence upon some moments of this distribution. To be more specific, in the case of interactions of scalar type we have

$$b(t, \omega, x, \mu) = \tilde{b}(t, \omega, x, \langle\varphi, \mu\rangle) \tag{7.10}$$

for some scalar function φ defined on $\mathbb{R}^d$ and some function $\tilde{b}$ defined on $[0, T] \times \Omega \times \mathbb{R}^d \times \mathbb{R}$. As before, we use the angular bracket notation

$$\langle \varphi, \mu \rangle = \int \varphi(x') d\mu(x')$$

for the integral of a function with respect to a measure.

Remark 7.5 *It is clear that scalar interactions can include functions of* several*, say n, moments of the measure. Still using the same notation and framework, it amounts to considering vector functions φ taking values in a Euclidean space $\mathbb{R}^n$ and having the function $\tilde{b}$ be defined on $[0,T] \times \Omega \times \mathbb{R}^d \times \mathbb{R}^n$.*

We shall also encounter applications in which the dependence upon the distribution is given by means of an auxiliary function $\tilde{b}$ defined on $[0,T] \times \Omega \times \mathbb{R}^d \times \mathbb{R}^d$, the interaction taking the form:

$$b(t, \omega, x, \mu) = \int_{\mathbb{R}^d} \tilde{b}(t, \omega, x, x') \mu(dx'). \tag{7.11}$$

This type of mean field interaction will be called *interaction of order 1*. It is linear in μ. Similarly, one could define mean field interactions of higher orders. For example, a mean field interaction of order 2 (which is quadratic in μ) should be of the form

$$b(t, \omega, x, \mu) = \int\int_{\mathbb{R}^d \times \mathbb{R}^d} \tilde{b}(t, \omega, x, x', x'') \mu(dx') \mu(dx''). \tag{7.12}$$

7.4.2 **Existence and Uniqueness of Solutions: the Lipschitz Case**

Theorem 7.6 *Let us assume that $X_0 \in L^2$ is independent of $\mathbf{W}$, and that the coefficients b and σ satisfy the assumptions **(A1)** and **(A2)** stated above. Then, there exists a unique solution to the equation* (7.9) *in $\mathbb{H}^{2,d}$, and for some $c > 0$ depending only upon T and the Lipschitz constant of b and σ, this solution satisfies*

$$\mathbb{E} \sup_{0 \leqslant t \leqslant T} |X_t|^2 \leqslant c(1 + \mathbb{E}|X_0|^2) e^{cT}. \tag{7.13}$$

Proof: Let $m \in \mathcal{P}_2(C([0,T];\mathbb{R}^d))$ be temporarily fixed, and let us denote by m_t its time marginals, i.e. the push-forward image of the measure m by X_t viewed as a map from $C([0,T];\mathbb{R}^d)$ into $\mathbb{R}^d$. By Lebesgue's dominated convergence theorem, the inequality

$$W^{(2)}(m_s, m_t)^2 \leqslant \int |X_s(\omega) - X_t(\omega)|^2 \, m(d\omega)$$

implies that the map $[0,T] \ni t \hookrightarrow m_t \in \mathcal{P}_2(\mathbb{R}^d)$ is continuous for the Wasserstein distance $W^{(2)}$. Hence, substituting momentarily m_t for $\mathcal{L}(X_t)$ for all $t \in [0,T]$ in (7.9), since X_0 is given, Theorem 7.2 gives existence and uniqueness of a strong solution of the classical stochastic differential equation

$$dX_t = b(t, X_t, m_t)dt + \sigma(t, X_t, m_t)dW_t \tag{7.14}$$

with random coefficients, and we denote its solution by $\mathbf{X}^m = (X_t^m)_{0 \leqslant t \leqslant T}$. We first notice that, because of the upper bound proven in Theorem 7.2, the law of $\mathbf{X}^m$ is of order 2. We then define the mapping $\Phi : \mathcal{P}_2(C([0,T];\mathbb{R}^d)) \ni m \hookrightarrow \Phi(m) = \mathcal{L}(\mathbf{X}^m) = \mathbb{P}_{\mathbf{X}^m} \in \mathcal{P}_2(C([0,T];\mathbb{R}^d))$. Since a process $\mathbf{X} = (X_t)_{0 \leqslant t \leqslant T}$ satisfying $\mathbb{E} \sup_{0 \leqslant t \leqslant T} |X_t|^2 < \infty$ is a solution of (7.9) if and only if its law is a fixed point of Φ, we prove the existence and uniqueness result of the theorem by proving that the mapping Φ has a unique fixed point. Let us choose m and m' in $\mathcal{P}_2(C([0,T];\mathbb{R}^d))$. Then since $\mathbf{X}^m$ and $\mathbf{X}^{m'}$ have the same initial conditions, for each $t \in [0,T]$, using Doob's maximal inequality and the Lipschitz assumption we have:

$$\begin{aligned}
\mathbb{E} \sup_{0\leqslant s\leqslant t} |X_s^m - X_s^{m'}|^2 \leqslant 2\mathbb{E} \sup_{0\leqslant s\leqslant t} & \left| \int_0^s [b(r, X_r^m, m_r) - b(r, X_r^{m'}, m'_r)] dr \right|^2 \\
& + \mathbb{E} \sup_{0\leqslant s\leqslant t} \left| \int_0^s [\sigma(r, X_r^m, m_r) - \sigma(r, X_r^{m'}, m'_r)] dW_r \right|^2 \\
& \leqslant cT \bigg(\int_0^t \mathbb{E} \sup_{0\leqslant r\leqslant s} |X_r^m - X_r^{m'}|^2 ds + \int_0^t W^{(2)}(m_s, m'_s) ds \\
& \qquad + \mathbb{E} \int_0^t |\sigma(r, X_r^m, m_r) - \sigma(r, X_r^{m'}, m'_r)|^2 dr \bigg) \\
& \leqslant cT \left(\int_0^t \mathbb{E} \sup_{0\leqslant r\leqslant s} |X_r^m - X_r^{m'}|^2 ds + \int_0^t W^{(2)}(m_s, m'_s) ds \right).
\end{aligned}$$

As usual, and except for the dependence upon T which we keep track of, we use the same notation even though the value of the constant c can change from line to line. Using Gronwall's inequality one concludes that

$$\mathbb{E} \sup_{0\leqslant s\leqslant t} |X_s^m - X_s^{m'}|^2 \leqslant c(T) \int_0^t W^{(2)}(m_s, m'_s)^2 \, ds. \tag{7.15}$$

with $c(T) = cTe^{cT}$. Notice that

$$W_t^{(2)}(\Phi(m), \Phi(m')) \leqslant \mathbb{E} \sup_{0\leqslant s\leqslant t} |X_s^m - X_s^{m'}|^2$$

because the joint law $\mathcal{L}(\boldsymbol{X}^m, \boldsymbol{X}^{m'})$ of $\boldsymbol{X}^m$ and $\boldsymbol{X}^{m'}$ is obviously a coupling of $\phi(m)$ and $\phi(m')$. Moreover, since $W^{(2)}(m_s, m'_s) \leqslant W_s^{(2)}(m, m')$ so that (7.15) implies

$$W_t^{(2)}(\Phi(m), \Phi(m')) \leqslant c(T) \int_0^t W_s^{(2)}(m, m')^2 \, ds.$$

Iterating this inequality and denoting by Φ^k the k-th composition of the mapping Φ with itself we get that for any integer $k > 1$

$$\begin{aligned}
W_T^{(2)}(\Phi^k(m), \Phi^k(m')) &\leqslant c(T) \int_0^T W_{t_k}^{(2)}(\Phi^{k-1}(m), \Phi^{k-1}(m'))^2 \, dt_k \\
&\leqslant c(T) \int_0^T \int_0^{t_k} W_{t_{k-1}}^{(2)}(\Phi^{k-2}(m), \Phi^{k-2}(m'))^2 \, dt_{k-1} dt_k \\
&\leqslant \cdots\cdots\cdots \\
&\leqslant c(T) \int_0^T \int_0^{t_k} \cdots \int_0^{t_2} W_{t_1}^{(2)}(m, m')^2 \, dt_1 \cdots dt_{k-1} dt_k \\
&= c(T)^k \int_0^T \frac{(T-s)^{k-1}}{(k-1)!} W_s^{(2)}(m, m')^2 \, ds \\
&\leqslant \frac{c^k T^k}{k!} W_T^{(2)}(m, m')^2,
\end{aligned}$$

which shows that for k large enough, Φ^k is a strict contraction and hence, Φ admits a unique fixed point. $\square$

Remark 7.7 *Let us fix $\varphi \in C_b^2(\mathbb{R}^d)$ and let us apply Itô's formula to $\varphi(X_t)$. We get*

$$\varphi(X_t) = \varphi(X_0) + \int_0^t [\frac{1}{2}\mathit{trace}[\sigma(s, X_s, \mathcal{L}(X_s))^\dagger \sigma(s, X_s, \mathcal{L}(X_s)) D^2\varphi(X_s)] \\ + b(s, X_s, \mathcal{L}(X_s)) D\varphi(X_s) ds + \int_0^t D\varphi(X_s)\sigma(s, X_s, \mathcal{L}(X_s))\, dW_s.$$

Taking expectations of both sides and using the notation $\mu_t = \mathcal{L}(X_t)$ we get:

$$\langle \varphi, \mu_t\rangle = \langle \varphi, \mu_0\rangle + \int_0^t [\frac{1}{2}\langle \mathit{trace}[\sigma(s, \cdot, \mu_s)^\dagger \sigma(s, \cdot, \mu_s) D^2\varphi(\cdot)], \mu_s\rangle + \langle b(s, \cdot, \mu_s) D\varphi(\cdot), \mu_s\rangle] ds$$

and after integration by parts

$$\langle \varphi, \mu_t\rangle = \langle \varphi, \mu_0\rangle + \int_0^t \langle \frac{1}{2}\mathit{trace}[\sigma(s, \cdot, \mu_s)^\dagger \sigma(s, \cdot, \mu_s) D^2\mu_s - \mathit{div}(b(s, \cdot, \mu_s)\mu_s, \varphi\rangle ds$$

or in differential form

$$\partial_t \mu_t = L(\mu_t)\mu_t$$

where for each $t \in [0, T]$ and $\mu \in \mathcal{P}(\mathbb{R}^d)$, the second order partial differential operator $L_t(\mu)$ is defined by:

$$L_t(\mu) f = \frac{1}{2}\mathit{trace}[\sigma(t, \cdot, \mu)^\dagger \sigma(t, \cdot, \mu) D^2 f - \mathit{div}[b(t, \cdot, \mu) f].$$

This is a form of (non-linear) Kolmogorov's equation for $\boldsymbol{\mu} = (\mu_t)_{0 \leqslant t \leqslant T}$.

7.4.3 Particle Approximations and Propagation of Chaos

Our next step is to study pathwise particle approximations of the solution of the McKean-Vlasov SDE (7.9). In particular, this will provide a proof of the original propagation of chaos result which was stated in terms of convergence of laws instead of a pathwise behavior. Let $((X_0^i, \mathbf{W}^i))_{i \geqslant 1}$ be a sequence of independent copies of $(X_0, \mathbf{W})$. For each $i \geqslant 1$, we let $\mathbf{X}^i = (X_t^i)_{0 \leqslant t \leqslant T}$ denote the solution of (7.9) constructed in Theorem 7.6 starting from X_0^i and driven by the Wiener process $\mathbf{W}^i$. It satisfies:

$$X_t^i = X_0^i + \int_0^t b(s, X_s^i, \mathcal{L}(X_s^i)) ds + \int_0^t \sigma(s, X_s^i, \mathcal{L}(X_s^i)) dW_s^i. \tag{7.16}$$

Notice that the probability measures $\mathcal{L}(X_s^i)$ do not depend upon i. Clearly, all the processes $\mathbf{X}^i$ are independent by construction. We show that they can be approximated by finite systems of classical Itô processes (which we often call particles) depending upon each other through specific interactions. For each integer $N \geqslant 1$, we consider the *particle processes* $\mathbf{X}^{i,N}$ for $i = 1, \cdots, N$ solving the system of standard SDEs:

$$X_t^{i,N} = X_0^i + \int_0^t b\big(s, X_s^{i,N}, \frac{1}{N}\sum_{i=1}^N \delta_{X_s^{i,N}}\big) ds + \int_0^t \sigma\big(s, X_s^{i,N}, \frac{1}{N}\sum_{i=1}^N \delta_{X_s^{i,N}}\big) dW_s^i, \tag{7.17}$$

for $i = 1, \cdots, N$, where we use the standard notation δ_x for the unit (Dirac) point mass at x. Notice that the coupling between these N SDEs is obtained by replacing in the McKean-Vlasov dynamics (7.9), the distributions $\mathcal{L}(X_t^i)$ whose presence creates the nonlinearity in the form of a self-interaction, by the empirical distributions of the particles $X_t^{1,N}, \cdots, X_t^{N,N}$. The hope is that a form of Law of Large Numbers will prove that the impact of this substitution on the solutions will be minimal. A simple application of the definition of the Wasserstein distance to the case of point measures shows that if $\mathbf{x} = (x_1, \cdots, x_N)$, and $\mathbf{y} = (y_1, \cdots, y_N)$ are generic elements of $\mathbb{R}^{dN}$, one has:

$$W^{(2)}\big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i}, \frac{1}{N}\sum_{i=1}^N \delta_{y_i}\big) \leqslant \left(\frac{1}{N}\sum_{i=1}^N |x_i - y_i|^2\right)^{1/2} = \frac{1}{\sqrt{N}}|\mathbf{x} - \mathbf{y}|. \tag{7.18}$$

This implies for fixed N, the uniform Lipschitz property for the coefficients of the system (7.17) and in turn, existence and uniqueness of a strong solution. The main result of this section establishes pathwise propagation of chaos for the interacting particle system (7.17). The following lemma will be useful in the control of the particle approximation.

Lemma 7.8 *Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, let $(\xi_i)_{i\geqslant 1}$ be a sequence of independent random variables with common law μ, and for each integer $N \geqslant 1$, let μ^N denote the empirical distribution of $\xi_1, \cdots, \xi_N$, namely $\mu^N = \frac{1}{N}\sum_{i=1}^N \delta_{\xi_i}$. Then for each $N \geqslant 1$ we have:*

$$\mathbb{E}W^{(2)}(\mu^N, \mu)^2 \leqslant \int_{\mathbb{R}^d} |x|^2 \mu(dx), \quad \textit{and} \quad \lim_{N\to\infty} \mathbb{E}W^{(2)}(\mu^N, \mu)^2 = 0.$$

Proof: By the strong law of large numbers, μ^N converges weakly toward μ almost surely as $N \to \infty$. Similarly, for each $1 \leqslant i, j \leqslant d$,

$$\lim_{N\to\infty} \int_{\mathbb{R}^d} x_i \mu^N(dx) = \int_{\mathbb{R}^d} x_i \mu(dx), \quad \text{and} \quad \lim_{N\to\infty} \int_{\mathbb{R}^d} x_i x_j \mu^N(dx) = \int_{\mathbb{R}^d} x_i x_j \mu(dx)$$

almost surely. Since the Wasserstein distance $W^{(2)}$ induces the topology of weak convergence together with the convergence of all the moments up to order 2, one concludes that $W^{(2)}(\mu^N, \mu)$ converges toward 0 almost surely as $N \to \infty$. Moreover, the sequence $(W^{(2)}(\mu^N, \mu)^2)_{N\geqslant 1}$ of random variables is uniformly integrable. Indeed, for any coupling π of μ^N and μ,

$$\begin{aligned}
W^{(2)}(\mu^N, \mu)^2 &\leqslant \int_{\mathbb{R}^d\times\mathbb{R}^d} |x - y|^2 \,\pi(dx, dy) \\
&\leqslant \int_{\mathbb{R}^d\times\mathbb{R}^d} (|x|^2 + |y|^2)\,\pi(dx, dy) \\
&= 2\left(\int_{\mathbb{R}^d} (|x|^2\,\mu^N(dx) + \int_{\mathbb{R}^d} |y|^2\,\mu(dy)\right) \\
&= \frac{2}{N}\sum_{i=1}^N |\xi_i|^2 + 2\int_{\mathbb{R}^d} |x|^2 \mu(dx)
\end{aligned}$$

which is nonnegative and converges almost surely toward $4\int_{\mathbb{R}^d} |x|^2\mu(dx)$ which is a finite constant. Since the limit is a constant, the convergence is also in the sense of L^1 from which we conclude. □

Theorem 7.9 *Under the above assumptions for the processes defined by* (7.16) *and* (7.17) *we have:*

$$\lim_{N\to\infty} \sup_{1\leqslant i\leqslant N} \mathbb{E} \sup_{0\leqslant t\leqslant T} |X_t^{i,N} - X_t^i|^2 = 0. \tag{7.19}$$

Usually, propagation of chaos is a statement about distributions rather than a pathwise statement. It says that for fixed $k \geqslant 1$, if we let $N \nearrow \infty$, then the law of $(X_t^{i,N})_{t\in[0,T]}^{i=1,\cdots,k}$ converges toward the probability distribution of $(X_t^i)_{t\in[0,T]}^{i=1,\cdots,k}$ implying that the k particles $\mathbf{X}^{1,N}, \cdots, \mathbf{X}^{k,N}$ become independent and acquire the same distribution given by the solution of the McKean-Vlasov SDE (7.9). Theorem 7.9, and especially statement (7.19), give a pathwise form of this statement by constructing the finite particle systems and their chaotic limits on the same probability space and proving pathwise convergence via a mean field coupling.

Proof: For any fixed $t \in [0,T]$, $N \geqslant 1$ and $i \in \{1, \cdots, N\}$, we have:

$$\begin{aligned}
\mathbb{E} \sup_{0\leqslant s\leqslant t} |X_s^{i,N} - X_s^i|^2 \leqslant c\int_0^t & \mathbb{E}|b\big(s, X_s^{i,N}, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^{i,N}}\big) - b\big(s, X_s^i, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^i}\big)|^2 ds \\
& c\int_0^t \mathbb{E}|b\big(s, X_s^i, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^i}\big) - b(s, X_s^i, \mathcal{L}(X_s^i)|^2 ds \\
& + c\int_0^t \mathbb{E}|\sigma\big(s, X_s^{i,N}, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^{i,N}}\big) - \sigma\big(s, X_s^i, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^i}\big)|^2 ds \\
& c\int_0^t \mathbb{E}|\sigma\big(s, X_s^i, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^i}\big) - \sigma(s, X_s^i, \mathcal{L}(X_s^i)|^2 ds.
\end{aligned}$$

Using the Lipschitz property of the coefficients b and σ, the elementary estimate (7.18) controlling the distance between two empirical measures, and the exchangeability of the couples $(X^i, X^{i,N})_{1\leqslant i\leqslant N}$, one sees that the first and third terms of the above right are bounded from above by $c\int_0^t \mathbb{E}\sup_{0\leqslant r\leqslant s} |X_r^{i,N} - X_r^i|^2 ds$. Using GronwallÕs inequality and once more the Lipschitz property of the coefficients b and σ, we get

$$\begin{aligned}
\mathbb{E} \sup_{0\leqslant s\leqslant t} |X_x^{i,N} - X_x^i|^2 \leqslant c\int_0^t & \mathbb{E}|b\big(s, X_s^i, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^{k,N}}\big) - b(s, X_s^i, \mathcal{L}(X_s^i)|^2 ds \\
& + c\int_0^t \mathbb{E}|\sigma\big(s, X_s^i, \frac{1}{N}\sum_{k=1}^N \delta_{X_s^{k,N}}\big) - \sigma(s, X_s^i, \mathcal{L}(X_s^i)|^2 ds \\
\leqslant c\int_0^t & \mathbb{E}W^{(2)}(\frac{1}{N}\sum_{k=1}^N \delta_{X_s^{k,N}}, \mathcal{L}(X_s^i))^2 ds.
\end{aligned}$$

The claim of the theorem now follows from Lemma 7.8 and Lebesgue's dominated convergence theorem. □

7.5 Some Important Examples

7.5.1 The Kuramoto Model

The Kuramoto model (KM) of coupled phase oscillators is a very popular example of dynamical system, its success being mostly due to its analytical simplicity and universality of the dynamical mechanisms that it helped to reveal. It describes the evolution of a system of N interconnected phase oscillators $\theta^{N,i} : [0,\infty) mapto \mathbb{R}/2\pi\mathbb{Z}$ with intrinsic frequencies $\omega^{N,i}$ for $i = 1, \cdots, N$. Their dynamics are given by the following system of Ordinary Differential Equations (ODE) :

$$\dot{\theta}_t^{N,i} = \omega^{N,i} + \frac{\kappa}{N} \sum_{j=1}^{N} a_{i,j}^N \sin(\theta_t^{N,i} - \theta_t^{N,i} + \alpha), \qquad i \in [N],$$

where κ is the strength of the coupling, $\alpha \in [0, 2\pi)$specifies the type of interaction (attractive or repulsive), and the symmeric matrix $\mathbf{a}^N = [a_{i,j}^N]_{i,j=1,\cdots,N}$ defines the graph underpinning the network of interactions between the oscillators. One of the most interesting properties of this dynamical system is the existence of a critical value κ_c separating incoherent dynamics from perfect synchronization.

The sum on the right–hand side models the interactions between the oscillators, $\alpha \in [0, 2\pi)$ determines the type of interactions (attractive vs repulsive), and K is the strength of coupling. The spatial structure of interconnections is encoded in the adjacency matrix (a_n^{ij}). The KM plays an important role in the theory of synchronization. We mention two major contributions that are especially relevant to the present study. First, it reveals a universal mechanism for the transition to synchronization in systems of coupled oscillators with random intrinsic frequencies. The analysis of the KM shows that there is a critical value of the coupling strength Kc separating the incoherent (mixing) dynamics (Fig. 1a) from synchronization (Fig. 1b) [19, 7, 8]. Second, studies of the KM led to the discovery of chimera states, patterns combining regions of coherent and incoherent dynamics

8

Stochastic Differential Games

This chapter is devoted to the analysis of stochastic differential games with a strong emphasis on several important features which may not always be compatible. First, we extend to stochastic differential games the notion of anonymous game, especially games with mean field interactions leading to what is now known as Mean Field Games (MFGs). Concurrently, we consider games for which the interactions between players are underpinned by a graph structure, possibly breaking the symmetry of the interactions found in MFGs. In both cases, we focus on a probabilistic approach (as opposed to those based on PDEs). To this end, we present a version of the stochastic maximum principle and we apply it to the solution of a simple Linear Quadratic (LQ) model chosen in hope that its limit when the number of players go to infinity can be analyzed.

8.1 Introduction and First Definitions

The purpose of this chapter is to introduce and develop the mathematical analysis of competitive stochastic differential games with finitely many players. We denote by N the number of players. We label them by the integers $1, \cdots, N$, using frequently the notation $[N] = \{1, \cdots, N\}$. At each time t, the players act on a system whose state X_t they influence through their actions, the dynamics of X_t being given by a stochastic differential equation of the Itô's type. The Itô process giving the state dynamics is driven by a m-dimensional Wiener process $\boldsymbol{W} = (W_t)_{0 \leqslant t \leqslant T}$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, the filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leqslant t \leqslant T}$ being assumed to be most of the time its natural filtration.

We denote by $A^1, \cdots, A^N$ the sets of actions that players $1, \cdots, N$ can take at any point in time. A^i is typically a compact metric space or a subset of a closed convex subset of a Euclidean space, say $A^i \subset \mathbb{R}^{k_i}$, and we denote by $\mathcal{A}^i$ its Borel σ-field. We will use the notation $A = A^1 \times \cdots \times A^N$ for the set of actions $\alpha = (\alpha^1, \cdots, \alpha^N)$ available to players $1, \cdots, N$ at any given time. Also, in order to emphasize the multivariate nature of the game models (as opposed to the single agent nature of control problems) we use the term strategy profile. We use the notation $\mathbb{A}$ for the set of admissible strategy profiles. The elements $\boldsymbol{\alpha}$ of $\mathbb{A}$ are N-tuples $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^1, \cdots, \boldsymbol{\alpha}^N)$ where each $\boldsymbol{\alpha}^i = (\alpha^i_t)_{0 \leqslant t \leqslant T}$ is a A^i-valued adapted process. Moreover, these individual strategies will have to satisfy extra conditions (e.g. measurability and integrability constraints) which change from one application to another. In most of the cases considered here, we shall assume that these constraints can be defined player by player, independently of each other. To be more specific, we shall often assume that $\mathbb{A} = \mathbb{A}^1 \times \cdots \times \mathbb{A}^N$ where for each $i \in \{1, \cdots, N\}$, $\mathbb{A}^i$ is the space of controls / strategies which are deemed admissible to player i, irrespective of what the other players do. Typically, $\mathbb{A}^i$ will be a space of A^i-valued, progressively measurable processes $\boldsymbol{\alpha}^i = (\alpha^i_t)_{0\leqslant t\leqslant T}$ being either bounded, or satisfying an integrability condition such as $\mathbb{E}\int_0^T |\alpha^i_t|^2 dt < \infty$.

A Convenient Notation. If $\alpha = (\alpha^1, \cdots, \alpha^N) \in A$, $i \in \{1, \cdots, N\}$, and $\beta^i \in A^i$, we will denote by (α^{-i}, β^i) the collective set of actions where all players except player i keep the same actions, while player i switches from action α^i to β^i. Similarly, if $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^1, \cdots, \boldsymbol{\alpha}^N) \in \mathbb{A}$ is a set of admissible strategies for the N players, and $\boldsymbol{\beta}^i \in \mathbb{A}^i$ is an admissible strategy for player i, then we denote by $(\boldsymbol{\alpha}^{-i}, \boldsymbol{\beta}^i)$ the new set of strategies where at each time t all players $j \neq i$ keep the same action α^j_t while player i switches from action α^i_t to β^i_t. In other words, $(\boldsymbol{\alpha}^{-i}, \boldsymbol{\beta}^i)^j_t$ is equal to α^j_t if $j \neq i$ and β^i_t otherwise.

For each choice of admissible strategy profile $\boldsymbol{\alpha} = (\alpha_t)_{0\leqslant t\leqslant T} \in \mathbb{A}$, it is assumed that the time evolution of the state $\boldsymbol{X} = \boldsymbol{X}^{\boldsymbol{\alpha}}$ of the system satisfies:

$$\begin{cases} dX_t = b(t, X_t, \alpha_t)dt + \sigma(t, X_t, \alpha_t)dW_t & 0 \leqslant t \leqslant T, \\ X_0 = x. \end{cases} \tag{8.1}$$

where

$$(b, \sigma) : [0.T] \times \Omega \times \mathbb{R}^d \times A \hookrightarrow \mathbb{R}^d \times \mathbb{R}^{d\times m}$$

satisfies

(a1) $\forall x \in \mathbb{R}^d, \forall \alpha \in A, (b(t, x, \alpha))_{0\leqslant t\leqslant T}$ and $(\sigma(t, x, \alpha))_{0\leqslant t\leqslant T}$ are progressively measurable processes with values in $\mathbb{R}^d$ and $\mathbb{R}^{d\times m}$ respectively;

(a2) $\exists c > 0, \forall t \in [0, T], \forall \alpha \in A, \forall \omega \in \Omega, \forall x, x' \in \mathbb{R}^d$,
$|b(t, \omega, x, \alpha) - b(t, \omega, x', \alpha)| + |\sigma(t, \omega, x, \alpha) - \sigma(t, \omega, x', \alpha)| \leqslant c|x - x'|$.

As usual, we omit ω from the notation whenever possible.

Unless the α^i_t are in Markovian feedback form (i.e. of the form $\alpha^i_t = \phi(t, X_t)$ for some deterministic function $(t, x) \mapsto \phi(t, x)$) the dynamics given by (8.1) are given by a stochastic differential equation whose coefficients depend upon the past.

8.1.1 A Frequently Encountered Special Case

It happens often that the state of the system is the aggregation of private states of individual players, say $X_t = (X^1_t, \cdots, X^N_t)$ where $X^i_t \in \mathbb{R}^{d_i}$ can be interpreted as the private state of player $i \in \{1, \cdots, N\}$. Here $d = d_1 + \cdots + d_N$ and consequently, $\mathbb{R}^d = \mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_N}$. The main feature of such a special case is that we usually assume that the dynamics of the private states are given by stochastic differential equations driven by separate Wiener processes $\boldsymbol{W}^i = (W^i_t)_{0\leqslant t\leqslant T}$ which are most often assumed to be independent of each other. See nevertheless our discussion of the *common noise case* later on. So typically we assume that

$$dX^i_t = b^i(t, X_t, \alpha_t)dt + \sigma^i(t, X_t, \alpha_t)dW^i_t \qquad 0 \leqslant t \leqslant T,\ i = 1, \cdots, N \tag{8.2}$$

where the $\boldsymbol{W}^i = (W^i_t)_{0\leqslant t\leqslant T}$ are m_i-dimensional independent Wiener processes giving the components of $\boldsymbol{W} = (W_t)_{0\leqslant t\leqslant T}$, and where the functions

$$(b^i, \sigma^i) : [0.T] \times \Omega \times \mathbb{R}^d \times A \hookrightarrow \mathbb{R}^{d_i} \times \mathbb{R}^{d_i \times m_i}$$

satisfy the same assumptions as before. It is important to notice that these N dynamical equations are coupled by the fact that all the private states and all the actions enter into the coefficients of these N equations. We can use vector/matrix notation, and rewrite the system of N stochastic dynamical equations in Euclidean spaces of dimensions $d_1, \cdots, d_N$ respectively as a stochastic differential equation giving the dynamics of the d-dimensional state X_t. So if we set

$$X_t = \begin{bmatrix} X_t^1 \\ X_t^2 \\ \vdots \\ X_t^N \end{bmatrix}, \quad b(t,x,\alpha) = \begin{bmatrix} b^1(t,x,\alpha) \\ b^2(t,x,\alpha) \\ \vdots \\ b^N(t,x,\alpha) \end{bmatrix}, \quad \text{and} \quad W_t = \begin{bmatrix} W_t^1 \\ W_t^2 \\ \vdots \\ W_t^N \end{bmatrix}, \tag{8.3}$$

with $m = m_1 + \cdots + m_N$ and

$$\sigma(t,x,\alpha) = \begin{bmatrix} \sigma^1(t,x,\alpha) & 0 & \cdots & 0 \\ 0 & \sigma^2(t,x,\alpha) & \cdots & 0 \\ \ddots & \ddots & \ddots & \ddots \\ 0 & \cdots & 0 & \sigma^N(t,x,\alpha) \end{bmatrix}.$$

we recover the dynamics of the state of the system given by (8.1). However, it is important to emphasize the special block-diagonal structure of the volatility matrix σ, and most importantly, the fact that when the Wiener processes $\boldsymbol{W}^i$ are not independent, the components of the Wiener process $\boldsymbol{W}$ are not independent which is an assumption which we often make when we use dynamics of the form (8.1).

The popularity of this formulation is due to the ease with which we can define the information structures and admissible strategy profiles of some specific games of interest. For example, in a game where each player can only use the information of the state of the system at time t when making a strategic decision at that time, the admissible strategy profiles will be of the form $\alpha_t^i = \phi^i(t, X_t)$ for some deterministic function ϕ^i. These strategies are said to be closed loop in feedback form, or Markovian. Moreover, if the information which can be used by player i at time t can only depend upon his/her own private state at time t, then the admissible strategy profiles will be of the form $\alpha_t^i = \phi^i(t, X_t^i)$. Such strategies are usually called distributed.

8.1.2 Cost Functionals and Notions of Optimality

In full analogy with the case of stochastic control involving only one controller, we assume that each player (controller) faces instantaneous and running costs. So for each $i \in \{1, \cdots, N\}$ we assume that we have

- an $\mathcal{F}_T$-measurable square integrable random variable $\xi^i \in L^2(\Omega, \mathcal{F}_T, \mathbb{P})$ usually called the *terminal cost*. Most often, ξ^i will be of the form $\xi^i = g^i(X_T)$ for some $\mathcal{F}_T \times \mathcal{B}_{\mathbb{R}^d}$-measurable function $g^i : \Omega \times \mathbb{R}^d \hookrightarrow \mathbb{R}$ which is assumed to grow at most quadratically;

- a function $f^i : [0,T] \times \Omega \times \mathbb{R}^d \times A \hookrightarrow \mathbb{R}$ called the *running cost* satisfying the same assumption as the drift b;

from which we define the overall expected cost to player i

- (*cost functional*) If the N players use the strategy profile $\boldsymbol{\alpha} \in \mathbb{A}$, the expected total cost to player i is defined as

$$J^i(\boldsymbol{\alpha}) = \mathbb{E}\Big[\int_0^T f^i(s, X_s, \alpha_s)ds \,+\, \xi^i\Big\}, \qquad \boldsymbol{\alpha} = (\boldsymbol{\alpha}^1, \cdots, \boldsymbol{\alpha}^N) \in \mathbb{A}, \quad (8.4)$$

  where $\boldsymbol{X}$ is the state of the system whose dynamics are given by equation (8.1).

Notice that, in the general situation considered here, the cost to a given player depends upon the strategies used by the other players indirectly through the values of the state X_t over time, but also directly as the specific actions α_t^j taken by the other players may appear explicitly in the expression of the running const f^i of player i.

Each player i attempts to minimize his/her total expected cost J^i. If we introduce the notation

$$J(\boldsymbol{\alpha}) = (J^1(\boldsymbol{\alpha}), \cdots, J^N(\boldsymbol{\alpha})), \qquad \boldsymbol{\alpha} \in \mathbb{A},$$

heuristically speaking, finding a solution to the game amounts to searching for a solution to the stochastic optimization problem for the functional J over the set $\mathbb{A}$ of admissible strategy profiles. The major difficulty is that J is taking values in a multidimensional Euclidean space which is not totally ordered in a natural fashion. We need to specify how we compare the costs of different strategy profiles in order to clearly define the notion of optimality. We shall use the notion of Nash equilibrium.

Definition 8.1 *A set of admissible strategies $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^{*1}, \cdots, \boldsymbol{\alpha}^{*N}) \in \mathbb{A}$ is said to be a Nash equilibrium for the game if*

$$\forall i \in \{1, \cdots, N\}, \;\; \forall \boldsymbol{\alpha}^i \in \mathbb{A}^i, \qquad J^i(\boldsymbol{\alpha}^*) \leqslant J^i(\boldsymbol{\alpha}^{*-i}, \boldsymbol{\alpha}^i).$$

NB: Nash equilibriums are not Pareto efficient in general!

The existence and uniqueness (or lack thereof) of equilibriums, as well as the properties of the corresponding optimal strategy profiles strongly depend upon the information structures available to the players, and the types of actions they are allowed to take. So rather than referring to a single game with several information structures and admissible strategy profiles for the players, we choose to talk about models, e.g. the open loop model for the game or the closed loop model, or even the Markovian model for the game. We give precise definitions below. The published literature on stochastic differential games, at least the part addressing terminology issues, is rather limited, and there is no clear consensus on the names to give to the many notions of admissibility for strategy profiles. We warn the reader that the definitions we use reflect our own personal biases and, this disclaimer being out of the way, the best we can do is to pledge consistency with our choices.

Definition 8.2 *If the strategy profile $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^{*1}, \cdots, \boldsymbol{\alpha}^{*N}) \in \mathbb{A}$ satisfies the conditions of Definition 8.1 without further restriction on the strategies $\boldsymbol{\alpha}^{*i}$ and $\boldsymbol{\alpha}^i$, we say that $\boldsymbol{\alpha}^*$ is an open loop Nash equilibrium (OLNE for short) for the game, or equivalently, a Nash equilibrium for the open loop game model.*

If the filtration $\mathbb{F}$ is generated by the Wiener process $\boldsymbol{W}$, except possibly for the presence of independent events in $\mathcal{F}_0$, the strategy profiles used in an open loop game model can be viewed as given by controls of the form

$$\alpha_t^i = \varphi^i(t, X_0, W_{[0,t]})$$

for some deterministic functions $\varphi^1, \cdots, \varphi^N$ where we used the notation $W_{[0,t]}$ for the path of the Wiener process between time $t = 0$ and t. The definition of open loop equilibrium warrants some caution. This definition is very natural and very convenient from a mathematical point of view, and we shall see that powerful existence results can be proved for these game models. However, it is rather unrealistic from a practical point of view. Indeed, it is very difficult to imagine situations in which the whole trajectory $W_{[0,t]}$ of the random shocks can be observed. So using functions of this trajectory as strategies does not seem very constructive as an approach to the search for an equilibrium !

Definition 8.3 *If the strategy profile $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^{*1}, \cdots, \boldsymbol{\alpha}^{*N}) \in \mathbb{A}$ satisfies the conditions of Definition 8.1 with the restriction that the strategies $\boldsymbol{\alpha}^{*i}$ and $\boldsymbol{\alpha}^i$ are deterministic functions of time and the initial state, we say that $\boldsymbol{\alpha}^*$ is a deterministic Nash equilibrium (DNE for short) for the game.*

The strategy profiles used in the search for a deterministic equilibrium are given by controls of the form

$$\alpha_t^i = \varphi^i(t, X_0)$$

for some deterministic functions $\varphi^1, \cdots, \varphi^N$ of the time variable t and the state of the system x. Given the fact that the Wiener process $\boldsymbol{W}$ is not present in deterministic game models, the above definitions are consistent with the standard terminology used in the classical analysis of deterministic games.

Definition 8.4 *If the strategy profile $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^{*1}, \cdots, \boldsymbol{\alpha}^{*N}) \in \mathbb{A}$ satisfies the conditions of Definition 8.1 with the restriction that the strategies $\boldsymbol{\alpha}^{*i}$ and $\boldsymbol{\alpha}^i$ are deterministic functions of time and the trajectory of the state between time $t = 0$ and t, we say that $\boldsymbol{\alpha}^*$ is a closed loop Nash equilibrium (CLNE for short) for the game, or equivalently, a Nash equilibrium for the closed loop game model.*

The strategy profiles used in the search for a closed loop equilibrium are given by controls of the form

$$\alpha_t^i = \varphi^i(t, X_{[0,t]})$$

for some deterministic functions $\varphi^1, \cdots, \varphi^N$. Finally,

Definition 8.5 *If the strategy profile $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^{*1}, \cdots, \boldsymbol{\alpha}^{*N}) \in \mathbb{A}$ satisfies the conditions of Definition 8.1 with the restriction that the strategies $\boldsymbol{\alpha}^{*i}$ and $\boldsymbol{\alpha}^i$ are deterministic functions of time, the initial state and the trajectory of the state at time t, we say that $\boldsymbol{\alpha}^*$ is a closed loop Nash equilibrium in feedback form (CLFFNE for short) for the game.*

The strategy profiles used in the search for a closed loop equilibrium in feedback form are given by controls of the form

$$\alpha_t^i = \varphi^i(t, X_0, X_t)$$

for some deterministic functions $\varphi^1, \cdots, \varphi^N$. The most important example of application of this notion of equilibrium concerns the case of deterministic drift and volatility coefficients b and σ, and cost functions f and g. This case will be discussed in Subsection 8.1.4 treating Markovian diffusions and we shall strengthen the notion of closed loop equilibrium in feedback form into the notion of Markovian equilibriumm.

Remark 8.6 *While we went out on a limb in our choice for the definition of open loop models, the following remarks are consistent with the terminology used in most write ups on deterministic games. Typically, in the open loop model, players cannot observe the play of their opponents while in the closed loop model, at each time, all past play is common knowledge. From a mathematical stand point, open loop equilibriums are more tractable than closed loop equilibriums because players need not consider how their opponents would react to deviations from the equilibrium path. With this in mind. one should expect that when the impact of players on their opponents' costs/rewards is small, open loop and closed loop equilibriums should be the same. We shall see instances of this intuition in the large N limit of the Linear Quadratic (LQ) model of systemic risk analyzed in Section* **??** *and in our discussion of Mean Field Games (MFG) in Chapter* **??**.

8.1.3 Toward a Version of the Pontryagin Stochastic Maximum Principle

Players' Hamiltonians

For each player $i \in [N]$, we define their Hamiltonian as the function H^i:

$$[0,T] \times \Omega \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^{d\times m} \times A \ni (t,x,y,z,\alpha) \hookrightarrow H^i(t,x,y,z,\alpha) \in \mathbb{R} \tag{8.5}$$

defined by

$$H^i(t,x,y,z,\alpha) = \underbrace{b(t,x,\alpha)\cdot y}_{\substack{\text{inner product of}\\ \text{state drift } b \text{ and}\\ \text{covariable } y}} + \underbrace{\text{trace}\,[\sigma(t,x,\alpha)^\dagger z]}_{\substack{\text{inner product of}\\ \text{state volatility } \sigma\\ \text{and covariable } z}} + \underbrace{f^i(t,x,\alpha)}_{\substack{\text{running cost}\\ \text{of player } i}}.$$

When the actions of the players do not appear in the volatility of the state (i.e. when the volatility is not controlled), we use the reduced Hamiltonians

$$\tilde{H}^i(t,x,y,\alpha) = b(t,x,\alpha)\cdot y + f^i(t,x,\alpha).$$

whose minimum in the variable α^i is attained for the same value as for the full Hamiltonian. We explain below the relevance of this remark.

Generalized Isaacs (MinMax) Condition

The following definition is motivated by the generalization to stochastic differential games of the necessary part of the stochastic maximum principle which we give below.

Definition 8.7 *We say that the generalized Isaacs (minmax) condition holds if there exists a function*

$$\hat{\alpha} : [0,T] \times \mathbb{R}^d \times (\mathbb{R}^d)^N \times (\mathbb{R}^{d\times m})^N \ni (t,x,y,z) \hookrightarrow \hat{\alpha}(t,x,y,z) \in A$$

satisfying, for every $i \in \{1,\cdots,N\}$*,* $t \in [0,T]$*,* $x \in \mathbb{R}^d$*,* $y = (y^1,\cdots,y^N) \in (\mathbb{R}^d)^N$ *and* $z = (z^1,\cdots,z^N) \in (\mathbb{R}^{d\times m})^N$

$$H^i(t,x,y^i,z^i,\hat{\alpha}(t,x,y,z)) \leqslant H^i(t,x,y^i,z^i,(\hat{\alpha}(t,x,y,z)^{-i},\alpha^i)) \quad \textit{for all } \alpha^i \in A^i. \tag{8.6}$$

Notice that in this definition, the function $\hat{\alpha}$ could depend upon the random scenario $\omega \in \Omega$ if the Hamiltonians H^i do. In words, this definition says that for each set of dual variables $y = (y^1,\cdots,y^N)$ and $z = (z^1,\cdots,z^N)$, for each time t and state x at time t, and possibly random scenario ω, one can find a set of actions $\hat{\alpha} = (\hat{\alpha}^1,\cdots,\hat{\alpha}^N)$ depending on these quantities, and such that if we fix $N-1$ of these actions, say $\hat{\alpha}^{-i}$, then the remaining one $\hat{\alpha}^i$ minimizes the i-th Hamiltonian in the sense that:

$$\hat{\alpha}^i \in \arg \inf_{\alpha^i \in A^i} H^i(t,x,y^i,z^i,(\hat{\alpha}^{-i},\alpha^i)), \qquad \text{for all } i \in \{1,\cdots,N\}. \tag{8.7}$$

The notation can be lightened slightly when the volatility is not controlled. Indeed, as we explained above, minimizing the Hamiltonian gives the same $\hat{\alpha}$ as minimizing the reduced Hamiltonian. Note also that in this case, the argument $\hat{\alpha}$ of the minimization is independent of z. So when the volatility is not controlled we say that the generalized Isaacs (minmax) condition holds if there exists a function

$$\hat{\alpha} : [0,T] \times \mathbb{R}^d \times (\mathbb{R}^d)^N \ni (t,x,y,z) \hookrightarrow \alpha^*(t,x,y) \in A$$

satisfying

$$\begin{aligned} &\forall i \in \{1,\cdots,N\}, \forall t \in [0,T], \forall x \in \mathbb{R}^d, \forall y = (y^1,\cdots,y^N) \in (\mathbb{R}^d)^N,\\ &\qquad \tilde{H}^i(t,x,y^i,\hat{\alpha}(t,x,y)) \leqslant \tilde{H}^i(t,x,y^i,(\hat{\alpha}(t,x,y)^{-i},\alpha^i)) \quad \text{for all } \alpha^i \in A^i. \end{aligned} \tag{8.8}$$

Rationale: the fact that individual players' Hamiltonians should be minimized is suggested by the construction of the best response function as a solutions of N stochastic control problems and the necessary part of the Pontryagin stochastic maximum principle for stochastic control problems. Accordingly, the fact that the same function $\hat{\alpha}$ minimizes ALL the Hamiltonians simultaneously is dictated by search for a fixed point of the best response function.

8.1.4 The Particular Case of Markovian / Diffusion Dynamics

In most applications for which actual numerical computations are possible, the coefficients of the state dynamics (8.1) depend only upon the *present value* X_t of the state instead of the entire past $X_{[0,t]}$ of the state of the system, or of the WIener process driving the evolution of the state. In this case, the dynamics of the state are given by a diffusion-like equation

$$dX_t = b(t,X_t,\alpha_t)dt + \sigma(t,X_t,\alpha_t)dW_t \qquad 0 \leqslant t \leqslant T, \tag{8.9}$$

with initial condition $X_0 = x$, and for deterministic drift and volatility functions

$$(b, \sigma) : [0, T] \times \mathbb{R}^d \times A \ni (t, x, \alpha) \hookrightarrow (b(t, x, \alpha), \sigma(t, x, \alpha)) \in \mathbb{R}^d \times \mathbb{R}^{d \times m}.$$

So except for the fact that the strategy profile $\boldsymbol{\alpha}$ may depend upon the past (feature that we are about to discard) the solution of (8.9) should be like a Markov diffusion. For this reason, we shall concentrate on strategy profiles which are deterministic functions of time and the current value of the state to indeed force the controlled state process to be a Markov diffusion. In fact, we shall also assume that the running cost functions f^i and the terminal cost random variables ξ^i are Markovian in the sense that, like b and σ, f^i does not depend upon the random scenario $\omega \in \Omega$, but only upon the current values of the state and the actions taken by the players, so that we can have $f^i : [0, T] \times \mathbb{R}^d \times A \ni (t, x, \alpha) \hookrightarrow f^i(t, x, \alpha) \in \mathbb{R}$, and ξ^i is of the form $\xi^i = g^i(X_T)$ for some measurable function $g^i : \mathbb{R}^d \ni x \hookrightarrow g^i(x) \in \mathbb{R}$ with (at most) quadratic growth. So in the case of Markovian / diffusion dynamics, the cost functional of player i is of the form:

$$J^i(\boldsymbol{\alpha}) = \mathbb{E}\{\int_0^T f^i(t, X_t, \alpha_t)dt \ + \ g^i(X_T)\}, \qquad \boldsymbol{\alpha} \in \mathbb{A}, \tag{8.10}$$

and we tailor the notion of equilibrium to this situation by considering closed loop strategy profiles in feedback forms which provide *simultaneously* Nash equilibriums for all the games starting at times $t \in [0, T]$ (i.e. over the time periods $[t, T]$) and all the possible initial conditions $X_t = x$ as long as they share the same drift and volatility coefficients b and σ, and cost functions f^i and g^i.

Markov Nash Equilibriums

Inspired by the notion of *sub-game perfect* equilibriums, we introduce the strongest notion yet of Nash equilibrium.

Definition 8.8 *A set $\phi = (\varphi^1, \cdots, \varphi^N)$ of N deterministic functions $\varphi^i : [0, T] \times \mathbb{R}^d \hookrightarrow \mathbb{R}^k$ for $i = 1, \cdots, N$ is said to be a Markov Nash equilibrium (MNE for short), or a Nash equilibrium for the Markovian game model if for each $(t, x) \in [0, T] \times \mathbb{R}^d$, the strategy profile $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^{*1}, \cdots, \boldsymbol{\alpha}^{*N}) \in \mathbb{A}$ defined for $s \in [t, T]$ by $\alpha_s^{*i} = \varphi^i(s, X_s^{t,x})$ where $\boldsymbol{X}^{t,x}$ is the unique solution of the stochastic differential equation*

$$dX_s = b(s, X_s, \phi(s, X_s))ds + \sigma(s, X_s, \phi(s, X_s))dW_s, \qquad t \leqslant s \leqslant T$$

with initial condition $X_t = x$, satisfies the conditions of Definition 8.1 with the restriction that the strategy $\boldsymbol{\alpha}^i$ is also given by a deterministic function φ on $[t, T] \times \mathbb{R}^d$.

It goes without saying that regularity assumptions on the functions φ^i are needed for the stochastic differential equations giving the dynamics of the controlled state to have a unique strong solution. Typically, we assume that

the coefficients b and σ are Lipschitz in (x, α) uniformly in $t \in [0, T]$.

The strategy profiles used in the above definition are called Markovian strategy profiles. Obviously, they are close loop in feedback form.

8.2 Game Version of the Stochastic Maximum Principle

Proving generalizations of the Pontryagin maximum principle to stochastic games is not as straightforward as one would like. In these lectures, we limit ourselves to open loop equilibriums for the sake of simplicity. Throughout this section, we assume that the drift and volatility functions, as well as the running and terminal cost functions are deterministic functions which are differentiable with respect to the variable x, and that the partial derivatives $\partial_x b$, $\partial_x \sigma$, $\partial_x f^i$ and $\partial_x g^i$ for $i = 1, \cdots, N$ are uniformly bounded. Notice that since b takes values in $\mathbb{R}^d$ and $x \in \mathbb{R}^d$, $\partial_x b$ is an element of $\mathbb{R}^{d\times d}$, in other words a $d \times d$ matrix whose entries are the partial derivatives of the components b^j of b with respect to the components x^i of x. Analog statements can be made concerning $\partial_x \sigma$ which has the interpretation of a tensor.

8.2.1 Open Loop Equilibriums

The generalization of the stochastic Pontryagin maximum principle to open loop stochastic games can be approached in a very natural way, and forms of the open loop sufficient condition for the existence and identification of a Nash equilibrium have been used in the case of linear quadratic models. See the Notes & Complements at the end of the chapter for references.

Definition 8.9 *Given an open loop admissible strategy profile $\boldsymbol{\alpha} \in \mathbb{A}$ and the corresponding evolution $\boldsymbol{X} = \boldsymbol{X}^{\boldsymbol{\alpha}}$ of the state of the system, a set of N couples $(\boldsymbol{Y}^{i,\boldsymbol{\alpha}}, \boldsymbol{Z}^{i,\boldsymbol{\alpha}}) = (Y_t^{i,\boldsymbol{\alpha}}, Z_t^{i,\boldsymbol{\alpha}})_{t\in[0,T]}$ of processes in $\mathbb{S}^{2,d}$ and $\mathbb{H}^{2,d\times m}$ respectively for $i = 1, \cdots, N$, is said to be a set of adjoint processes associated with $\boldsymbol{\alpha} \in \mathbb{A}$ if for each player $i \in \{1, \cdots, N\}$ they satisfy the BSDEs*

$$\begin{cases} dY_t^{i,\boldsymbol{\alpha}} = -\partial_x H^i(t, X_t, Y_t^{i,\boldsymbol{\alpha}}, Z_t^{i,\boldsymbol{\alpha}}, \alpha_t)dt + Z_t^{i,\boldsymbol{\alpha}} dW_t \\ Y_T^{i,\boldsymbol{\alpha}} = \partial_x g^i(X_T). \end{cases} \tag{8.11}$$

We shall not argue the existence and uniqueness of the adjoint processes here. Even though we did not present the theory of Backward Stochastic Differential Equations (BSDEs) here, existence and uniqueness do not represent an issue under the present hypotheses. Indeed, given $\boldsymbol{\alpha} \in \mathbb{A}$ and the corresponding state evolution $\boldsymbol{X} = \boldsymbol{X}^{\boldsymbol{\alpha}}$, equation (8.11) can be viewed as a BSDE with random coefficients, terminal condition in L^2, and driver:

$$\begin{aligned} \psi(t, \omega, y, z) = -\partial_x b(t, X_t(\omega), \alpha_t(\omega)) \cdot y - \partial_x \sigma(t, X_t(\omega), \alpha_t(\omega)) \cdot z \\ - \partial_x f^i(t, X_t(\omega), \alpha_t(\omega)) \end{aligned}$$

which is an affine function of y and z with uniformly bounded random coefficients and an L^2 intercept. So for each $i \in \{1, \cdots, N\}$, existence and uniqueness of a solution $(Y_t^{i,\boldsymbol{\alpha}}, Z_t^{i,\boldsymbol{\alpha}})_{0\leqslant t\leqslant T}$ follows from standard existence results. See for example Theorem **??**.

The following result is the open loop *game version* of the necessary part of the Pontryagin maximum principle. Its proof can be conducted along the lines of the corresponding proof in the case of stochastic control. We do not give it as we only use this result as a rationale for the search for a function satisfying the min-max Isaacs condition.

Theorem 8.10 *Under the above conditions, if $\boldsymbol{\alpha}^* \in \mathbb{A}$ is a Nash equilibrium for the open loop game, and if we denote by $\boldsymbol{X}^* = (X_t^*)_{0 \leqslant t \leqslant T}$ the corresponding controlled state of the system, and by $(\boldsymbol{Y}^*, \boldsymbol{Z}^*) = \big((\boldsymbol{Y}^{*1}, \cdots, \boldsymbol{Y}^{*N}), (\boldsymbol{Z}^{*1}, \cdots, \hat{Z}^{*N})\big)$ a set of adjoint processes, then the generalized min-max Isaacs conditions hold along the optimal paths in the sense that for each $i \in \{1, \cdots, N\}$:*

$$H^i(t, X_t^*, Y_t^{*i}, Z_t^{*i}, \alpha_t^*) = \inf_{\alpha^i \in A^i} H^i(t, X_t^*, Y_t^{*i}, Z_t^{*i}, (\alpha^{*-i}, \alpha^i)), \quad dt \otimes d\mathbb{P} \text{ a.s.} \tag{8.12}$$

We now state and prove the sufficient condition which we will use in the applications we consider in these lectures.

Theorem 8.11 *Assuming that the functions b, σ and f^i are twice continuously differentiable in $(x, \alpha) \in \mathbb{R}^d \times A$ and g^i are twice continuously differentiable in $x \in \mathbb{R}^d$, all with bounded partial derivatives, if $\hat{\boldsymbol{\alpha}} \in \mathbb{A}$ is an admissible adapted (open loop) strategy profile, $\hat{\boldsymbol{X}} = (\hat{X}_t)_{0 \leqslant t \leqslant T}$ the corresponding controlled state, and $(\hat{\boldsymbol{Y}}, \hat{\boldsymbol{Z}}) = \big((\hat{\boldsymbol{Y}}^1, \cdots, \hat{\boldsymbol{Y}}^N), (\hat{\boldsymbol{Z}}^1, \cdots, \hat{\boldsymbol{Z}}^N)\big)$ a set of corresponding adjoint processes such that for each $i \in \{1, \cdots, N\}$:*

1. *$(x, \alpha) \hookrightarrow H^i(t, x, \hat{Y}_t^i, \hat{Z}_t^i, \alpha)$ is a convex function , $\quad dt \otimes d\mathbb{P}$ a.s.;*
2. *g^i is convex,*

and if moreover, for every $i \in \{1, \cdots, N\}$ we have:

$$H^i(t, \hat{X}_t, \hat{Y}_t^i, \hat{Z}_t^i, \hat{\alpha}_t) = \inf_{\alpha^i \in A^i} H^i(t, \hat{X}_t, \hat{Y}_t^i, \hat{Z}_t^i, (\hat{\alpha}^{-i}, \alpha^i)), \quad dt \otimes d\mathbb{P} \text{ a.s.}, \tag{8.13}$$

then $\hat{\boldsymbol{\alpha}}$ is a Nash equilibrium for the open loop game.

Proof: We fix $i \in \{1, \cdots, N\}$, a generic (adapted) $\boldsymbol{\alpha}^i \in \mathbb{A}^i$, and for the sake of simplicity, we denote by $\boldsymbol{X}$ the state $\boldsymbol{X}^{(\hat{\boldsymbol{\alpha}}^{-i}, \boldsymbol{\alpha}^i)}$ controlled by the strategies $(\hat{\boldsymbol{\alpha}}^{-i}, \boldsymbol{\alpha}^i)$. The function g^i being convex almost surely, we have:

$$\begin{aligned}
&g^i(\hat{X}_T) - g^i(X_T) \\
&\quad \leqslant (\hat{X}_T - X_T)\partial_x g^i(\hat{X}_T) \\
&\quad = (\hat{X}_T - X_T)\hat{Y}_T^i \\
&\quad = \int_0^T (\hat{X}_t - X_t)\, d\hat{Y}_t^i + \int_0^T \hat{Y}_t^i\, d(\hat{X}_t - X_t) + \int_0^T [\sigma(t, \hat{X}_t, \hat{\alpha}_t) - \sigma(t, X_t, (\hat{\alpha}^{-i}, \alpha^i))] \cdot \hat{Z}_t^i\, dt \\
&\quad = -\int_0^T (\hat{X}_t - X_t)\partial_x H^i(t, \hat{X}_t, \hat{Y}_t^i, \hat{Z}_t^i, \hat{\alpha}_t)\, dt + \int_0^T \hat{Y}_t^i [b(t, \hat{X}_t, \hat{\alpha}_t) - b(t, X_t, (\hat{\alpha}^{-i}, \alpha^i))]\, dt \\
&\qquad\qquad + \int_0^T [\sigma(t, \hat{X}_t, \hat{a}_t) - \sigma(t, X_t, (\hat{\alpha}^{-i}, \alpha^i))] \cdot \hat{Z}_t^i\, dt + \text{ martingale}
\end{aligned} \tag{8.14}$$

So that, taking expectations of both sides and plugging the result into

$$J^i(\hat{\boldsymbol{\alpha}}) - J^i((\hat{\boldsymbol{\alpha}}^{-i}, \boldsymbol{\alpha}^i)) = \mathbb{E}\{\int_0^T [f^i(t, \hat{X}_t, \hat{\alpha}_t) - f^i(t, X_t, (\hat{\alpha}^{-i}, \alpha^i))]dt\} + \mathbb{E}\{g^i(\hat{X}_T) - g^i(X_T)\}$$

we get:

$$
\begin{aligned}
J^i(\hat{\boldsymbol{\alpha}}) - J^i((\hat{\boldsymbol{\alpha}}^{-i}, \alpha^i)) &= \mathbb{E}\{\int_0^T [H^i(t, \hat{X}_t, \hat{Y}_t^i, \hat{Z}_t^i, \hat{\alpha}_t) - H^i(t, X_t, \hat{Y}_t^i, \hat{Z}_t^i, (\hat{\alpha}^{-i}, \alpha^i))]dt\} \\
&\qquad - \mathbb{E}\{\int_0^T \hat{Y}_t^i [b(t, \hat{X}_t, \hat{\alpha}_t) - b(t, X_t, (\hat{\alpha}^{-i}, \alpha^i))]\, dt\} \\
&\qquad - \mathbb{E}\{\int_0^T [\sigma(t, \hat{X}_t, \hat{\alpha}_t) - \sigma(t, X_t, (\hat{\alpha}^{-i}, \alpha^i))] \cdot \hat{Z}_t^i\, dt\} \\
&\qquad + \mathbb{E}\{g^i(\hat{X}_T) - g^i(X_T)\} \\
&\leqslant \mathbb{E}\{\int_0^T [H^i(t, \hat{X}_t, \hat{Y}_t^i, \hat{Z}_t^i, \hat{\alpha}_t) - H^i(t, X_t, \hat{Y}_t^i, \hat{Z}_t^i, (\hat{\alpha}^{-i}, \alpha^i)) \\
&\qquad - (\hat{X}_t - X_t)\partial_x H^i(t, \hat{X}_t, \hat{Y}_t^i, \hat{Z}_t^i, \hat{\alpha}_t)]\, dt\} \\
&\leqslant 0,
\end{aligned} \tag{8.15}
$$

because the above integrand is non-positive for $dt \otimes d\mathbb{P}$ almost all $(t, \omega) \in [0, T] \times \Omega$. Indeed, this is easily seen by a second order Taylor expansion as a function of (x, α), using the convexity assumption and the fact that $\hat{\alpha}$ is a critical point (where the first order derivative vanishes) because it satisfies the generalized Isaacs condition by assumption. □

Implementation Strategy

We shall try to use this sufficient condition in the following manner. When the coefficients of the model are differentiable with respect to the state variable x with bounded derivatives, if the convexity assumptions 1. and 2. of the above theorem are satisfied, we shall search for a deterministic function $\hat{\alpha}$

$$(t, x, (y^1, \cdots, y^N), (z^1, \cdots, z^N)) \hookrightarrow \hat{\alpha}(t, x, (y^1, \cdots, y^N), (z^1, \cdots, z^N)) \in A$$

on $[0, T] \times \mathbb{R}^d \times \mathbb{R}^{dN} \times \mathbb{R}^{dmN}$ which satisfies Isaacs conditions. Next, we replace the *adapted* controls $\boldsymbol{\alpha}$ in the forward dynamics of the state as well as in the adjoint BSDEs by

$$\hat{\alpha}(t, X_t, (Y_t^1, \cdots, Y_t^N), (Z_t^1, \cdots, Z_t^N)).$$

This creates a large FBSDE comprising a forward equation in dimension d and N backward equations in dimension d. The couplings between these equations may be highly nonlinear, and this system may be very difficult to solve. However, if we find processes $\boldsymbol{X}$, $(\boldsymbol{Y}^1, \cdots, \boldsymbol{Y}^N)$, $(\boldsymbol{Z}^1, \cdots, \boldsymbol{Z}^N)$ solving this FBSDE

$$
\begin{cases}
dX_t = b(t, X_t, \hat{\alpha}(t, X_t, (Y_t^1, \cdots, Y_t^N), (Z_t^1, \cdots, Z_t^N)))dt \\
\qquad\qquad + \sigma(t, X_t, \hat{\alpha}(t, X_t, (Y_t^1, \cdots, Y_t^N), (Z_t^1, \cdots, Z_t^N)))dW_t, \\
dY_t^1 = -\partial_x H^1(t, X_t, Y_t^1, Z_t^1, \hat{\alpha}(t, X_t, (Y_t^1, \cdots, Y_t^N), (Z_t^1, \cdots, Z_t^N)))dt + Z_t^1 dW_t, \\
\cdots \;=\; \cdots \\
dY_t^N = -\partial_x H^N(t, X_t, Y_t^N, Z_t^N, \hat{\alpha}(t, X_t, (Y_t^1, \cdots, Y_t^N), (Z_t^1, \cdots, Z_t^N)))dt + Z_t^N dW_t,
\end{cases} \tag{8.16}
$$

with initial condition $X_0 = x$ for the forward equation, and for each $i \in \{1, \cdots, N\}$, terminal condition $Y_T^i = \partial_x g^i(X_T)$ for the backward equation, the above sufficient condition says that the strategy profile $\hat{\boldsymbol{\alpha}}$ defined by $\hat{\alpha}_t = \hat{\alpha}(t, X_t, (Y_t^1, \cdots, Y_t^N), (Z_t^1, \cdots, Z_t^N))$ is an open loop Nash equilibrium.

8.3 First Application to a Simple Network Game

We first describe the model.

8.3.1 The Finite Player Game Model

For each N, positive integer denoting the number of players, let $W^N = [w^N_{i,j}]_{i,j=1,\cdots,N}$ be an $N \times N$ symmetric matrix of real numbers. We want to think of the integers in $\{1, \ldots, N\}$ as players, and $W^N_{i,j}$ as a weight quantifying the strength of the interaction between players i and j. When these weights only take values 0 or 1, the matrix W is merely the adjacency matrix of the graph of players, to be understood as a way to indicate the couples of players in interaction. For the sake of definiteness, we shall assume that $w^N_{ii} = 0$. So players do not interact with themselves. We carry the superscript N to emphasize the number of players as we intend to eventually take the limit $N \to \infty$ to analyze the situation for *large network games*.

We denote by A the set of actions which are admissible to each player. A will typically be a closed convex subset of a Euclidean space $\mathbb{R}^d$.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ a probability space on which is defined a sequence $(\mathbf{B}^n)_{n\geqslant 1}$ of independent processes of Brownian motion, $\mathbf{B}^n = (B^n_t)_{t\geqslant 0}$, and we denote by $\mathbb{F} = (\mathcal{F}_t)_{t\geqslant 0}$ the filtration $\mathcal{F}_t = \sigma\{B^n_s :\ 0 \leqslant s \leqslant t,\ n \geqslant 1\}$. We also assume the existence of a sequence $(X^n_0)_{n\geqslant 1}$ of independent random variables in $\mathbb{R}^d$ which is independent of the sequence of Brownian motions.

For each $N \geqslant 1$, we denote by $(\boldsymbol{X}^{N,i})_{i=1,\cdots,N}$ the Itô stochastic processes with initial conditions $X^{N,i}_0 = X^i_0$ and stochastic differentials:

$$dX^{N,i}_t = \alpha^{N,i}_t dt + \sigma_N dB^i_t, \qquad i = 1, \cdots, N.$$

where for each $i \in [N]$, $\boldsymbol{\alpha}^{N,i} = (\alpha^{N,i}_t)_{0\leqslant t\leqslant T}$ is an $\mathbb{R}^k$-valued process progressively measurable with respect to the filtration generated by the Brownian motion $\mathbf{B}^{(N)} = (\mathbf{B}^1, \cdots, \mathbf{B}^N)$. $X^{N,i}_t$ represents the (private) state of player i in the N player version of our game, and $(\alpha^{N,i}_t)_{0\leqslant t\leqslant T}$ the control process giving their strategy. For the sake of simplicity we assume that the dimensions $d_1, \cdots, d_N$ of all the private states are equal to 1 so that $d = N$. As a result all the k_i are also equal to 1 and $\alpha^{N,i}_t \in \mathbb{R}$.

We complete the definition of the game by introducing the costs incurred by the players. The cost to player $i \in [N]$ is given by:

$$J^i(\boldsymbol{\alpha}^{N,1}, \cdots, \boldsymbol{\alpha}^{N,N}) = \mathbb{E}\Big[\int_0^T f^i(t, (X^{N,1}_t, \cdots, X^{N,N}_t), \alpha^{N,i}_t) dt\Big]$$

with

$$f^i(t, (x^1, \cdots, x^N), \alpha) = \frac{1}{2}|\alpha|^2 + \frac{\kappa^2}{2}|x^i - \frac{1}{N}\sum_{j=1}^N w^N_{ij} x^j|^2.$$

This quadratic cost has two components. First, a penalization for the choice of the action $\alpha \in A$, and second a penalty of the state x^i of player i to be far from the aggregate of the states of the players they are interacting with. Note that this aggregate involves the graph structure underpinning the interactions in the network, and that it changes with i. As a result, while this interaction is mean field when $w_{ij}^N \equiv 1$, in general, this type of interaction is not mean field because the symmetry between the players is broken by the particular structure of the graph.

8.3.2 **Construction of Nash Equilibriums**

We use the sufficient condition of the Pontryagin stochastic maximum principle to identify a Nash equilibrium for this game. Notice that for each $i \in [N]$, the terminal cost is zero which is obviously convex. Moreover, the Hamiltonian of player i reads:

$$H^i(t, \boldsymbol{x}, \boldsymbol{y}^i, \boldsymbol{z}^i, \boldsymbol{\alpha}) = \sum_{j=1}^N \alpha^j y^{i,j} + \sigma \sum_{j=1}^N z^{ijj} + \frac{1}{2}|\alpha^i|^2 + \frac{\kappa^2}{2}|x^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N x^j|^2$$

which is clearly convex in $(\boldsymbol{x}, \boldsymbol{\alpha})$ so that conditions 1. and 2. of Theorem 8.11 are satisfied. Also, the minimization of H^i with respect to α^i can easily be achieved since:

$$\frac{\partial H^i}{\partial \alpha^i} = 0 \quad \Leftrightarrow \quad y^{ii} + \alpha^i = 0 \quad \Leftrightarrow \quad \alpha^i = -y^{ii}.$$

So using the notation of the above *Implementation Strategy* subsection, the function $\hat{\alpha}$ satisfying the Isaacs conditions is given by:

$$\hat{\alpha}(t, x, (y^1, \cdots, y^N), (z^1, \cdots, z^N)) = (-y^{11}, \cdots, -y^{NN}).$$

Since

$$\frac{\partial H^i(\cdots)}{\partial x^j} = -\kappa^2 \Big(x^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N x^j\Big)\Big(\delta_{ij} - \frac{1}{N} w_{ij}^N\Big)$$

the forward/backward system (8.16) becomes (dropping momentarily the superscript N for convenience):

$$\begin{cases} dX_t^i = -Y_t^{ii} dt + \sigma dB_t^i, \qquad i = 1, \cdots, N \\ dY_t^{ij} = -\kappa^2 \Big(X_t^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N X_t^j\Big)\Big(\delta_{ij} - \frac{1}{N} w_{ij}^N\Big) dt + \sum_{k=1}^N Z_t^{ijk} dB_t^k, \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad i, j = 1, \cdots, N \end{cases} \tag{8.17}$$

with initial conditions X_0^i for the N forward equations and $Y_T^{ij} = 0$ for the N^2 backward equations. Below, we first solve the system

$$\begin{cases} dX_t^i = -Y_t^{ii} dt + \sigma dB_t^i, \qquad i = 1, \cdots, N \\ dY_t^{ii} = -\kappa^2 \Big(X_t^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N X_t^j\Big) dt + \sum_{k=1}^N Z_t^{iik} dB_t^k, \quad i, j = 1, \cdots, N, \end{cases} \tag{8.18}$$

recall that we assumed $w_{ii}^N = 0$, and then we shall set

$$Y_t^{ij} = \kappa^2 \mathbb{E}\Big[\int_t^T \Big(X_s^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N X_s^j\Big) ds \Big| \mathcal{F}_t\Big], \qquad \text{for } j \neq i.$$

This is indeed the desired solution because

$$Y_t^{ij} = -\kappa^2 \int_0^t \Big(X_s^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N X_s^j\Big)ds + \kappa^2 \mathbb{E}\Big[\int_0^T \Big(X_s^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N X_s^j\Big)ds \Big| \mathcal{F}_t\Big],$$

and

$$M_t = \kappa^2 \mathbb{E}\Big[\int_0^T \Big(X_s^i - \frac{1}{N}\sum_{j=1}^N w_{ij}^N X_s^j\Big)ds \Big| \mathcal{F}_t\Big]$$

being a martingale in the Brownian filtration $(\mathcal{F}_t)_{t\geqslant 0}$, it can be represented as a stochastic integral with respect to the Brownian motion processes generating the filtration. In other words, there exist square integrable processes $(\boldsymbol{Z}^{ijk})_{k=1,\cdots,N}$ such that:

$$M_t = \sum_{k=1}^N \int_0^t Z_s^{ijk} dB_s^k, \qquad 0 \leqslant t \leqslant T,$$

which is exactly what was needed to complete the solution of the adjoint equations. See for example [21, Theorem 4.15 p. 182]. In order to streamline the notations we set:

$$X_t = \begin{bmatrix} X_t^1 \\ X_t^2 \\ \vdots \\ X_t^N \end{bmatrix}, \; Y_t = \begin{bmatrix} Y_t^{11} \\ Y_t^{22} \\ \vdots \\ Y_t^{NN} \end{bmatrix}, \; Z_t = \begin{bmatrix} Z_t^{111} & \cdots & Z_t^{11N} \\ Z_t^{221} & \cdots & Z_t^{22N} \\ \vdots & \cdots & \vdots \\ Z_t^{NN1} & \cdots & Z_t^{NNN} \end{bmatrix}, \text{ and } B_t = \begin{bmatrix} B_t^1 \\ B_t^2 \\ \vdots \\ B_t^N \end{bmatrix}, \tag{8.19}$$

so we can rewrite the forward backward system (8.18) in the condensed manner:

$$\begin{cases} dX_t = -Y_t dt + \sigma dB_t, \qquad X_0 = X_0 \\ dY_t = -\kappa^2\Big(I - \frac{1}{N}W\Big)X_t dt + Z_t dB_t, \qquad Y_T = 0. \end{cases} \tag{8.20}$$

The system being linear, it is not unreasonable to expect that the solution of the backward equation be a linear function of the solution of the forward equation. There are sound reasons for that, but we shall not give them because they are beyond the scope of these lectures. Still, based on this intuition, we make the ansatz $Y_t = \eta_t X_t$ for a differentiable, $N \times N$ matrix valued (deterministic) function $t \mapsto \eta_t$ which we try to determine. Again, for the sake of notation we set $A = \kappa^2\Big(I - \frac{1}{N}W\Big)$ so that

$$dY_t = -AX_t dt + Z_t dB_t, \tag{8.21}$$

and computing dY_t from the ansatz and using the expression of dX_t given by the forward equation we get:

$$\begin{aligned} dY_t &= \dot{\eta}_t X_t dt + \eta_t dX_t \\ &= \dot{\eta}_t X_t dt + \eta_t(-Y_t dt + \sigma dB_t) \\ &= (\dot{\eta}_t - \eta_t^2)X_t dt + \sigma\eta_t dB_t \end{aligned} \tag{8.22}$$

and identifying with the expression (8.21) we get

$$\begin{cases} \dot{\eta}_t - \eta_t^2 = -A \\ Z_t = \sigma \eta_t. \end{cases} \tag{8.23}$$

The first relation is a matrix Riccati equation. If we assume that the largest eigenvalue of W is not greater than N, in which case $A \geqslant 0$ in the sense of inequality between matrices, then given the terminal condition $\eta_T = 0$, this Riccati equation admits a unique (non-exploding) solution which is given explicitly by:

$$\eta_t = \sqrt{A}\big(e^{2\sqrt{A}(T-t)} - I\big)\big(e^{2\sqrt{A}(T-t)} + I\big)^{-1} \tag{8.24}$$

which can also be expressed as

$$\eta_t = \sqrt{A}\Big(I - 2\big(e^{2\sqrt{A}(T-t)} + I\big)^{-1}\Big). \tag{8.25}$$

The fact that the matrix valued function η_t defined by either of these formulas is the solution of the Riccati equation in (8.23) can be checked by inspection. Using this expression for η_t in the ansatz and the dynamics of the state as given in the first equation of (8.20), we see that

$$dX_t = -\eta_t X_t dt + \sigma dB_t$$

which shows that in equilibrium, at least when the initial condition is Gaussian, the state process is Gaussian and the equilibrium strategy profile is given by

$$\hat{\alpha}_t = -\eta_t X_t$$

which is also Gaussian whenever X_t is. In fact, if for each $s \in [0, T]$, we denote by $U(t, s)$ the solution of the ordinary (matrix) differential equation

$$\frac{d}{dt} U(t, s) = -\eta_t U(t, s), \tag{8.26}$$

over the interval $[s, T]$ with the initial condition $U(s, s) = I$, then we have:

$$X_t = U(t, 0)X_0 + \sigma \int_0^t U(t, s) dB_s \tag{8.27}$$

from which one easily computes the covariance of the process $(X_t)_{0 \leqslant t \leqslant T}$.

$$\text{cov}(X_s, X_t) = U(t, 0)\text{cov}(X_0)U(t, 0)^\dagger + \sigma^2 \int_0^s \text{trace}[U(s, u)U(t, u)^\dagger] du$$

whenever $0 \leqslant s \leqslant t \leqslant T$. We use the exponent † to denote the transpose of a matrix. Note that the first term is not present when the initial condition X_0 is deterministic.

8.3.3 A Glance at the Convergence of Large Games when $N \to \infty$

The above calculation shows the important role played by the matrix $A = \kappa^2(I - N^{-1}W)$. As we are about to see, in order for this quantity not to blow up in the limit $N \to \infty$, we need to rescale the interaction part of the cost by choosing $\kappa = \sqrt{N}$, in which case $A = NI - W$.

Moreover, in order to analyze the limit of the solution of the network game considered above as the size of the game increases without bound, we rename the important quantities used above by introducing the dependence upon the number N of players. In particular:

$$A^N = NI^N - W^N, \quad \text{and} \quad \eta_t^N = \sqrt{A^N}(e^{2\sqrt{A^N}(T-t)} - I^N)(e^{2\sqrt{A^N}(T-t)} + I^N)^{-1},$$

where we denote by I^N the $N \times N$ identity matrix. The spaces to which these matrices belong, namely $\mathbb{R}^{N\times N}$ change with N, so for us to be able to control convergence when $N \to \infty$, we embed all these matrices into a common space. We choose this common space to be the Hilbert space $L^2(I \times I, \mathcal{B}_{I\times I}, \lambda_I \otimes \lambda_I)$ of (equivalence classes of) square integrable functions on the unit square $I \times I$. In order to shorten the notation, we shall often denote this space by $L^2(I \times I)$. A generic element, say w, of this Hilbert space is a function $I \times I \ni (x, y) \mapsto w(x, y) \in \mathbb{R}$ uniquely defined for λ_I almost every $x \in I$ and $y \in I$, and its value $w(x, y)$ can be interpreted as the strength of the interaction between x and y if the latter are understood as players in a game with a continuum of players. In graph theory, these functions w are often called *graphons*. Also, to each graphon w we associate an operator $\boldsymbol{W}$ on $L^2(I, \mathcal{B}_I, \lambda_I)$ defined by:

$$[\boldsymbol{W}f](x) = \int_I w(x, y)f(y)dy, \qquad f \in L^2(I, \mathcal{B}_I, \lambda_I) \tag{8.28}$$

and because this operator has a square integrable kernel, it is a Hilbert-Schmidt operator. Recall the discussion of Section 4.2 of Chapter 4.

So associated to the matrix A^N defined above, we introduce the piecewise constant graphon $\bar{A}^N$ which is equal to $A^N_{ij} = N\delta_{ij} - w^N_{ij}$ on the plaquette $I_i \times I_j$. Recall that for each integer $N \geqslant 1$, we denote by $\mathcal{P}^{(N)} = \{I_1, \cdots, I_N\}$ the partition of $[0, 1)$ comprising the N equal length intervals $I_j = [(j-1)N, j/N)$ for $j = 1, \cdots, N$. We often use the dyadic partitions corresponding to $N = 2^n$ for some integer $n \geqslant 0$.

At this stage, we have a sequence $(\bar{A}^N)_{N\geqslant 1}$ of (piecewise constant) elements of $L^2(I \times I)$. It makes sense to study the convergence of this sequence in this space. In terms of the corresponding sequence of operators $(\bar{\mathbf{A}}^N)_{N\geqslant 1}$ obtained through the definition (8.28), convergence of the *kernels* $\bar{A}^N$ corresponds to the convergence of the operators $\bar{\mathbf{A}}^N$ in the Hilbert-Schmidt norm, which is stronger than the convergence in the operator norm. Unfortunately, even the convergence in the sense of the operator norm is too much to expect in many cases. Indeed, as seen in Lemma 9.5, strong convergence is too often the best one can expect.

9

Stochastic Differential Graphon Games

In this chapter, we introduce a new class of dynamic stochastic differential games with a continuum of players. Our goal is to generalize the section of graphon games contained in Chapter 4 to the dynamic setting of controlled states given by the solutions of stochastic differential equations.

9.1 GAMES WITH A CONTINUUM OF PLAYERS

In hope to identify, and possibly analyze, game models which could appear as limits when $N \to \infty$ of network games of the type studied in the previous chapter, we introduce a model with a continuum of players with interactions given by a graphon.

9.1.1 Measure Theoretic Background

Let $I = [0,1]$, $\mathcal{B}_I$ be its Borel σ-field and λ_I the Lebesgue measure on I. We should think of each $x \in I$ as a player in the game we are about to introduce.

Next for technical reasons which will become clear later on, we consider a rich Fubini extension $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes \mathbb{P})$ as we did in Chapter 4 for the static case. $\mathcal{I}$ is a σ-field containing $\mathcal{B}_I$. It is not countably generated. λ is a probability measure extending the Lebesgue measure λ_I. So the measure space $(I, \mathcal{B}, \lambda)$ is an extension of the standard Lebesgue space $(I, \mathcal{B}_I, \lambda_I)$. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space which we should think of as the sample space. $\mathcal{F}$ is not countably generated either, so the classical Lebesgue spaces $L^p(\Omega, \mathcal{F}, \mathbb{P})$ will not be separable. The σ-field $\mathcal{I} \boxtimes \mathcal{F}$ contains the product σ-field $\mathcal{I} \otimes \mathcal{F}$ and the probability measure $\lambda \boxtimes \mathbb{P}$ is an extension of the product measure $\lambda \otimes \mathbb{P}$. See Chapter **??** for more on rich Fubini extensions.

The theory of Fubini extensions guarantees the existence of a (measurable) essentially pairwise independent process $(\xi_x)_{x \in I}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ with values in the Polish space $E = C([0,T])$ such that for each $x \in I$, the distribution of the random variable ξ_x is the Wiener measure μ_0 on E equipped with the Borel σ-field $\mathcal{B}_E$ of the topology defined by the sup-norm. For each $t \in [0,T]$ we denote by C_t the coordinate map $E \ni \omega \mapsto C_t(\omega) = \omega(t)$. With this definition $(C_t)_{0 \leqslant t \leqslant T}$ is a process of Brownian motion on the probability space $(E, \mathcal{B}_E, \mu_0)$ and as a result, for each $x \in I$ the process $\mathbf{B}^x = (B_t^x)_{0 \leqslant t \leqslant T}$ is a also a process of Brownian motion on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Moreover, the processes $(\mathbf{B}^x)_{x \in I}$ are essentially pairwise independent in the sense that for λ-almost every $x \in I$ the processes $\mathbf{B}^x$ is independent of the process $\mathbf{B}^y$ for λ-almost every $y \in I$. We should think of dB_t^x as the idiosyncratic shocks affecting the state of player x at time t.

For each $x \in I$, we denote by $\mathcal{F}_t^x = \sigma\{B_s^x;\ 0 \leqslant s \leqslant t\}$ the complete σ-field generated by the random variables B_s^x for $0 \leqslant s \leqslant t$. Now we define an admissible strategy profile $\boldsymbol{\alpha} = (\alpha_t^x)_{(x,t)\in I\times[0,T]}$ as a measurable and square integrable process on $(I \times \Omega, \mathcal{B}_I \otimes \mathcal{F}, \lambda_I \otimes \mathbb{P})$ with values in $L^2([0,T], dt)$ which is distributed in the sense that for each $x \in I$, the strategy $\boldsymbol{\alpha}^x = (\alpha_t^x)_{0\leqslant t\leqslant T}$ of player $x \in I$ is adapted to the filtration $\mathbb{F}^x = (\mathcal{F}_t^x)_{0\leqslant\leqslant T}$ generated by the Wiener process $\mathbf{B}^x = (B_t^x)_{t\geqslant 0}$.

Let w be a graphon whose associated operator on $L^2(I, \mathcal{B}_I, \lambda_I)$ we denote by $\boldsymbol{W}$. This operator can be extended to the Hilbert space $L^2(I, \mathcal{I}, \lambda)$ since for $f \in L^2(I, \mathcal{I}, \lambda)$ and for λ-almost every $x \in I$, the integral $\int_I w(x,y)f(y)\lambda(dy)$ makes sense. In fact, even though $f \in L^2(I, \mathcal{I}, \lambda)$ may not be measurable with respect to $\mathcal{B}_I$, the function $I \ni x \mapsto [\boldsymbol{W}f](x) = \int_I w(x,y)f(y)\lambda(dy)$ is as can be seen using the eigenfunction expansion of the operator $\boldsymbol{W}$ in $L^2(I, \mathcal{B}_I, \lambda_I)$. Indeed:

$$[\boldsymbol{W}f](x) = \sum_{k\geqslant 1} \lambda_k \varphi_k(x) \int_I \varphi_k(y)f(y)\lambda(dy)$$

which shows the desired measurability since all the eigenfunctions φ_k are $\mathcal{B}_I$-measurable. Note that the infinite series in the right hand side converges in $L^2(I)$-sense because the $\|\varphi_k\|_{L^2(I)} = 1$, $|\int_I \varphi_k(y)f(y)\lambda(dy)| \leqslant \left[\int_I |f(y)|^2\lambda(dy)\right]^{1/2}$, and the λ_k are square summable.

Even though we are not able at this stage to define rigorously the dynamics of the states of all the players simultaneously, we denote for each player $x \in I$, their state at time t by X_t^x and we assume that their overall expected cost is given by a quantity of the form:

$$J^x\big((\alpha_t)_{0\leqslant t\leqslant T}, \boldsymbol{\alpha}\big) = \mathbb{E}\Big[\int_0^T f(t, X_t^x, \alpha_t, Z_t^{\boldsymbol{\alpha},x})dt\Big] \tag{9.1}$$

when player x uses the strategy $(\alpha_t)_{0\leqslant t\leqslant T}$ and all the other players use the admissible strategy profile $\boldsymbol{\alpha} = (\alpha_t^x)_{(x,t)\in I\times[0,T]}$. In so doing, the ensemble of players creates at each time $t \in [0,T]$, an aggregate state:

$$Z_t^{\boldsymbol{\alpha},x} = \int_I w(x,y)X_t^y\ \lambda(dy) \tag{9.2}$$

which represents the state of the system faced by player $x \in I$. Note that if the graphon is not constant, this aggregate is likely to depend upon x, breaking the symmetry of anonymous and mean field games. Moreover, its definition as an integral requires the measurability of the map $y \mapsto X_t^y$ which is the reason we are setting up the game model with the help of Fubini extensions.

The definition X_t^x of the state of player x at time t, as well as the definition of its dynamics, will be chosen on a case by case basis. In any case, for the sake of notation we shall often denote the aggregate (9.2) by $[\boldsymbol{W}X_t^{\cdot}](x)$, implicitly using the extension of the operator $\boldsymbol{W}$ to $L^2(I, \mathcal{I}, \lambda)$. Also, we need to keep in mind the fact that the *function* $y \mapsto X_t^y$ depends upon the choice of the admissible strategy profile $\boldsymbol{\alpha}$.

9.1.2 **A First Class of Models**

In this subsection, we assume that $b : A \times \mathbb{R} \mapsto \mathbb{R}$ is a function which is Lipschitz in both variable, with a Lipschitz constant in the second variable being strictly smaller than 1. To be specific, we assume that there exist positive constants c_α and c_z such that $c_z < 1$ and

$$|b(\alpha, z) - b(\alpha', z')|^2 \leqslant c_\alpha |\alpha - \alpha'|^2 + c_z |z - z'|^2, \qquad (\alpha, z), (\alpha', z') \in A \times \mathbb{R}. \quad (9.3)$$

Our goal is to define the state of the system as a solution of a system of a continuum of coupled stochastic differential equations of the form

$$dX_t^x = b(\alpha_t^x, Z_t^{\boldsymbol{\alpha},x})dt + dB_t^x, \qquad t \geqslant 0, x \in I. \quad (9.4)$$

The meaning of this differential form is better expressed in integral form:

$$X_t^x = X_0^x + \int_0^t b(\alpha_s^x, Z_s^{\boldsymbol{\alpha},x})ds + B_t^x, \qquad t \geqslant 0, x \in I.$$

We assume that the initial condition is square integrable in the sense that $(X_0^x)_{x \in I} \in L^2(\Omega \times I, \mathcal{F} \otimes \mathcal{I}, \mathbb{P} \otimes \lambda)$ and that the random variables X_0^x are essentially pairwise independent. If we can actually construct such a state, we should have:

$$\mathbb{E}[X_t^x] = \mathbb{E}[X_0^x] + \int_0^t \mathbb{E}[b(\alpha_s^x, Z_s^{\boldsymbol{\alpha},x})]ds, \qquad t \geqslant 0, x \in I, \quad (9.5)$$

The following notation will become handy in the discussion below. For each $x_0 \in \mathbb{R}$, function $t \mapsto u_t \in A$, and real valued functions z and ξ of t, $t \mapsto z_t$ and $t \mapsto \xi_t$, we define the function $t \mapsto X_{x_0,u,z,\xi}(t)$ by

$$X_{x_0,u,z,\xi}(t) = x_0 + \int_0^t b(\alpha_s, z_s)ds + \xi(t). \quad (9.6)$$

Notice that $\mathbb{E}[X_{X_0^x, \alpha_\cdot^x, Z_\cdot^{\boldsymbol{\alpha},x}, B_\cdot^x}(t)]$ is exactly the expectation of the state X_t^x as considered in (9.5).

Lemma 9.1 *If the Hilbert-Schmidt norm of the graphon operator satisfies*

$$c_z T^2 \|W\|_2^2 < 2, \quad (9.7)$$

then for any admissible strategy profile $\boldsymbol{\alpha}$, there exists a unique element $\Omega \times [0,T] \times I \ni (\omega, t, x) \mapsto Z_t^{\boldsymbol{\alpha},x}(\omega)$ in $L^2(\Omega \times [0,T] \times I, \mathcal{F} \otimes \mathcal{B}_{[0,T]} \otimes \mathcal{B}_I, \mathbb{P} \otimes \lambda_{[0,T]} \otimes \lambda_I)$ satisfying:

$$Z_t^{\boldsymbol{\alpha},x} = \int_I w(x,y) X_{X_0^y, \alpha_\cdot^y, Z_\cdot^{\boldsymbol{\alpha},y}, B^y}(t)\lambda(dy) \quad (9.8)$$

for $\mathbb{P} \otimes \lambda_{[0,T]} \otimes \lambda_I$ almost every $(\omega, t, x) \in \Omega \times [0,T] \times I$, showing existence and uniqueness of a state process satisfying (9.4) *where the aggregate $Z_t^{\boldsymbol{\alpha},x}$ is given by* (9.2).

Proof: Given $\boldsymbol{\alpha}$, we define the map U from $L^2(\Omega \times [0,T] \times I, \mathcal{F} \otimes \mathcal{B}_{[0,T]} \otimes \mathcal{B}_I, \mathbb{P} \otimes \lambda_{[0,T]} \otimes \lambda_I)$ which we denote by $L^2(\Omega \times [0,T] \times I)$ throughout the proof, into itself by:

$$[Uz](t,x) = \int_I w(x,y) X_{X_0^y, \alpha_\cdot^y, z(\cdot,y), B^y}(t) \lambda(dy)$$

and we prove that U is a strict contraction of $L^2(\Omega \times [0,T] \times I)$. This will prove the claim of the lemma. Note that by definition of a Fubini extension, the integral with respect to $\lambda(dy)$ of

$$X_{X_0^y, \alpha_\cdot^y, z(\cdot,y), B^y}(t) = X_0^y(\omega) + \int_0^t b(\alpha_s^y(\omega), z(\omega,s,y))ds + B_t^y(\omega)$$

in $\mathcal{F}$-measurable in $\omega \in \Omega$ and since it is continuous in t, it is jointly measurable with respect to $\mathcal{F} \otimes \mathcal{B}_I$. Consequently, using again the eigenfunction expansion of the graphon function w as we did earlier, we see that $[Uz](\omega, t, x)$ as defined above is $\mathcal{F} \otimes \mathcal{B}_{[0,T]} \otimes \mathcal{B}_I$-jointly measurable.

As usual, we shall try not to write the dependence upon the variable $\omega \in \Omega$ whenever we can. We first prove that $Uz \in L^2(\Omega \times [0,T] \times I)$ whenever z does.

$$\begin{aligned}
&\mathbb{E}\int_0^T \int_I [Uz](t,x)^2 \, dtdx \\
&\quad = \mathbb{E}\int_0^T \int_I \Big| \int_I w(x,y) X_{X_0^y, \alpha_\cdot^y, z(\cdot,y), B_\cdot^y}(t) \, \lambda(dy) \Big|^2 \, dtdx \\
&\quad \leqslant \mathbb{E}\int_0^T \int_I \Big(\int_I w(x,y)^2 dy \Big) \Big(\int_I \Big| X_0^y + \int_0^t b(\alpha_s^y, z(s,y))ds + B_t^y \Big|^2 \lambda(dy) \Big) dtdx \\
&\quad \leqslant C\|W\|_2^2 \mathbb{E}\int_0^T \Big(\int_I |X_0^y|^2 dy + \int_I \Big| \int_0^t b(\alpha_s^y, z(s,y))ds \Big|^2 dy + \int_I |B_t^y|^2 \lambda(dy) \Big) dt \\
&\quad \leqslant C\|W\|_2^2 \Big(T\mathbb{E}\int_I |X_0^y|^2 dy + \frac{2}{3}T^3 |b(0,0)|^2 + \int_I \Big| \int_0^t (c_\alpha^2 |\alpha_s^y| + c_z^2 |z(s,y)|^2) dsdy + \frac{1}{2}T^2 \Big)
\end{aligned}$$

which is finite because $\boldsymbol{\alpha}$ and z are in $L^2(\Omega \times [0,T] \times I)$. We now prove the strict contraction property. Let z and $\tilde{z}$ be elements of $L^2(\Omega \times [0,T] \times I)$. Then:

$$\begin{aligned}
&\|Uz - U\tilde{z}\|^2_{L^2(\Omega \times [0,T] \times I)} \\
&\quad = \mathbb{E}\int_0^T \int_I \Big| \int_I w(x,y) \Big(X_{X_0^y, \alpha_\cdot^y, z(\cdot,y), B_\cdot^y}(t) - X_{X_0^y, \alpha_\cdot^y, \tilde{z}(\cdot,y), B_\cdot^y}(t) \Big) \, \lambda(dy) \Big|^2 \, dtdx \\
&\quad \leqslant \mathbb{E}\int_0^T \int_I \Big(\int_I w(x,y)^2 dy \Big) \Big(\int_I \Big| \int_0^t \Big(b(\alpha_s^y, z(s,y)) - b(\alpha_s^y, \tilde{z}(s,y)) \Big) ds \Big|^2 dy \Big) dtdx \\
&\quad \leqslant \|W\|_2^2 \, \mathbb{E}\int_0^T \int_I \int_0^t \big(b(\alpha_s^y, z(s,y)) - b(\alpha_s^y, \tilde{z}(s,y)) \big)^2 ds \, dy \, dt \\
&\quad \leqslant \frac{T^2}{2} c_z \, \|W\|_2^2 \, \mathbb{E}\int_0^T \int_I |z(s,y)) - \tilde{z}(s,y)|^2 dsdy
\end{aligned}$$

which conclude the proof since $c_z T^2 \, \|W\|_2^2 < 2$ by assumption. □

Our goal is now to prove that the aggregate constructed above is in fact deterministic. In preparation, we state without proof the following result which can be obtained with exactly the same proof as above.

Lemma 9.2 *Under assumption* (9.7)*, for any admissible strategy profile* $\boldsymbol{\alpha}$*, there exists a unique element* $[0,T] \times I \ni (t,x) \mapsto q^x(t)$ *in* $L^2([0,T] \times I, \mathcal{B}_{[0,T]} \otimes \mathcal{B}_I, \lambda_{[0,T]} \otimes \lambda_I)$ *satisfying:*

$$q^x(t) = \int_I w(x,y)\mathbb{E}\big[X_{X_0^y,\alpha_\cdot^y,q^y,B^y}(t)\big]dy \tag{9.9}$$

for $\lambda_{[0,T]} \otimes \lambda_I$ *almost every* $(t,x) \in [0,T] \times I$*. We shall denote by* $q^{\boldsymbol{\alpha},x}(t)$*.*

Proposition 9.3 *Under assumption* (9.7)*, for each admissible strategy profile* $\boldsymbol{\alpha}$*, the aggregate* $Z_t^{\boldsymbol{\alpha},x}$ *is deterministic.*

Proof: Because of our definition of admissibility of a strategy profile, because the initial conditions are assumed to be pairwise essentially independent, the fact that $q^{\boldsymbol{\alpha},x}(t)$ is determinist implies that for each fixed t, the random variables $X_{X_0^y,\alpha_\cdot^y,q^{\boldsymbol{\alpha},y},B^y}(t)$ are essentially pairwise independent. Consequently, the exact law of large numbers implies that:

$$\begin{aligned} q^{\boldsymbol{\alpha},x}(t) &= \int_I w(x,y)\mathbb{E}\big[X_{X_0^y,\alpha_\cdot^y,q^{\boldsymbol{\alpha},y},B^y}(t)\big]dy \\ &= \int_I w(x,y)X_{X_0^y,\alpha_\cdot^y,q^{\boldsymbol{\alpha},y},B^y}(t)\lambda(dy) \end{aligned} \tag{9.10}$$

which shows that $q^{\boldsymbol{\alpha},x}(t)$ satisfies the identity (9.8), and by uniqueness of the fixed point constructed in the proof of Lemma 9.1, we conclude that $q^{\boldsymbol{\alpha},x}(t) = Z_t^{\boldsymbol{\alpha},x}$ proving that the latter is deterministic. □

9.1.3 A Model Inspired by the Previous Finite Player Network Game

As a warm-up, we consider a natural generalization of the finite player network game studied earlier to the set-up of games with a continuum of players. We assume that for each player $x \in I$, the dynamics of their state are given by the Itô process

$$dX_t^x = \alpha_t^x\,dt + \sigma dB_t^x, \tag{9.11}$$

where we assume that $\boldsymbol{\alpha} = (\alpha_t^x)_{(x,t)\in I\times[0,T]}$ is an admissible strategy profile. We also assume that the initial conditions are given by a process $\boldsymbol{X}_0 = (X_0^x)_{x\in I}$ for which $I \times \Omega \ni (x,\omega) \mapsto X_0^x(\omega) \in L^2(\Omega \times I, \mathcal{F} \otimes \mathcal{B}_I, \mathbb{P} \otimes \lambda_I)$. This model is clearly a particular case of what was discussed in the previous subsection, as one can see by choosing $b(\alpha, z) = \alpha$ which satisfies assumption (9.3) with $c_\alpha = 1$ and $c_z = 0$. Accordingly, Proposition 9.3 says that for each admissible strategy profile, the state aggregates $Z_t^{\boldsymbol{\alpha},x}$ are deterministic. This is clear in the present situation since:

$$Z_t^{\boldsymbol{\alpha},x} = [\boldsymbol{W}X_0^{\cdot}](x) + \int_0^t \int_I w(x,y)\mathbb{E}[\alpha_s^y]\,dyds \tag{9.12}$$

Wishful Thinking: a Formal Computation

Let us assume that in equilibrium, the state trajectories are given by the solutions of a continuum of stochastic differential equations:

$$dX_t^x = -[\eta_t X_t^{\cdot}]^x dt + dB_t^x, \qquad x \in I, \tag{9.13}$$

for some continuous function $[0,T] \ni t \hookrightarrow \eta_t$ in the space of bounded operators on $L^2(I)$. Denoting by $(U(t,s))_{0 \leqslant s \leqslant t \leqslant T}$ the fundamental solution of the equation $\dot{f}_t = -\eta_t f_t$ as before, and assuming that these operators have kernels, one should expect that the solution of (9.13) be given by the formula:

$$X_t^x = [U(t,0)X_0^{\cdot}]^x + \int_0^t [U(t,s)dB_s^{\cdot}]^x$$

where the stochastic integral should be given by

$$\int_0^t [U(t,s)dB_s^{\cdot}]^x = \int_0^t \int_I U(t,s)(x,y)dB_s^y \lambda(dy)$$

where we use the notation $U(t,s)(x,y)$ for the kernel of the bounded operator $U(t,s)$, assuming that such a kernel exists. Assuming that these *formal* calculations can be justified, for $0 \leqslant s \leqslant t \leqslant T$ one has

$$\begin{aligned}
&\mathrm{cov}(X_t^x X_s^y) \\
&\quad = \mathbb{E}\Big[\Big(\int_0^t \int_I U(t,u)(x,x')dB_u^{x'}\lambda(dx')du\Big)\Big(\int_0^s \int_I U(s,u')(y,y')dB_{u'}^{y'}\lambda(dy')du'\Big)\Big] \\
&\quad = \mathbb{E}\Big[\Big(\int_0^s \int_I U(t,u)(x,x')dB_u^{x'}\lambda(dx')du\Big)\Big(\int_0^s \int_I U(s,u')(y,y')dB_{u'}^{y'}\lambda(dy')du'\Big)\Big] \\
&\quad = \int_0^s \int_0^s \int_I \int_I U(t,u)(x,x')U(s,u')(y,y')\mathbb{E}[dB_u^{x'}dB_{u'}^{y'}]\lambda(dx')\lambda(dy')dudu' \\
&\quad = \int_0^s \int_I U(t,u)(x,x')U(s,u)(y,x')\lambda(dx')du \\
&\quad = \int_0^s \int_I U(t,u)(x,x')U(s,u)(y,x')dx'du \\
&\quad = \int_0^s [U(t,u)U(s,u)^{\dagger}](x,y)du
\end{aligned} \tag{9.14}$$

where we used successively, the independence of the increments of the Brownian motions, the essential pairwise independence of the Brownian motions, the fact that the kernels of the operators U are $\mathcal{B}_I$ measurable so we can use dx' instead of $\lambda(dx')$, and finally the definition of the transpose and the product of kernel operators.

9.2 Convergence and Approximation Analyses

In this section, instead of starting from finite systems and analyzing their possible limits when their sizes grow without bound, we start from an infinite system, based on a model for the interactions between a continuum of players, and we construct finite systems which converge in a certain sense toward the infinite system we started from.

9.2.1 Approximations by Finite Random Graphs

9.2.2 Deterministic Approximation Analysis

As usual, we denote by $I = [0,1]$ the unit interval, $\mathcal{B}_I$ its Borel σ-field and by λ_I the Lebesgue measure. To shorten the notations, we shall use $L^2(I)$ for $L^2(I, \mathcal{B}_I, \lambda_I)$ and $L^2(I \times I)$ for $L^2(I \times I, \mathcal{B}_I \otimes \mathcal{B}_I, \lambda_I \otimes \lambda_I)$.

Function Approximation Preliminaries

For each integer $N \geqslant 1$ we denote by $\mathcal{P}^{(N)} = \{I_1, \cdots, I_N\}$ the partition of $[0,1)$ given by the N equal length intervals $I_j = [(j-1)N, j/N)$ for $j = 1, \cdots, N$, and by $L^{(N)}$ the subspace of $L^2(I)$ of the (equivalent classes of) functions which are constant over the intervals I_j. We emphasize the fact that the vector space $L^{(N)}$ is isometric (up to the normalizing scaling factor $\sqrt{N}$) to the Euclidean space $\mathbb{R}^N$. Indeed, the following calculation shows that this isometry $\boldsymbol{\pi}^N$ is given by the map from $L^{(N)}$ onto $\mathbb{R}^N$ which associates to a function f taking the values f_j on the intervals I_j, the element of $\mathbb{R}^N$ with j-th component f_j. If the functions f and g in $L^{(N)}$ take the values f_j and g_j on the intervals I_j, then:

$$< f, g >_{L^{(N)}} = \sum_{j=1}^{N} \int_{I_j} f(x)g(x)dx = \frac{1}{N} \sum_{j=1}^{N} f_j g_j = \frac{1}{N} < \boldsymbol{\pi}^N f, \boldsymbol{\pi}^N g >_{\mathbb{R}^N} . \tag{9.15}$$

Here and in the following, we denote by $< f, g >_L$ the inner product of f and g in the space L, which we skip in the case of $L^2(I)$.

We now identify the orthogonal projection $\boldsymbol{\Pi}^N$ of $L^2(I)$ onto the subspace $L^{(N)}$. If $f \in L^2(I)$ and $N \geqslant 1$, we denote by $\mu^N(f)$ the element of $\mathbb{R}^N$ defined by

$$\mu^N(f)_j = N \int_{I_j} f(y)dy, \qquad j = 1, \cdots, N.$$

The entries of this vector are the averages of the function f over the intervals I_j of the partition. For each $N \geqslant 1$ and $f \in L^2(I)$ we denote by $\bar{f}^N \in L^2(I, \mathcal{B}_I, \lambda_I)$ the piecewise constant function on I which is equal to $\mu^N(f)_j$ on I_j for $j = 1, \cdots, N$. This function has a nice probabilistic interpretation. Indeed, if for each $N \geqslant 1$ we denote by $\mathcal{B}_I^N$ the sub σ-field of $\mathcal{B}_I$ generated by the partition $\mathcal{P}^{(N)}$, then $\bar{f}^N$ is nothing but the conditional expectation of f with respect to $\mathcal{B}_I^N$, in notation

$$\bar{f}^N = \mathbb{E}^{\lambda_I}[f|\mathcal{B}_I^N]. \tag{9.16}$$

Clearly, $\bar{f}^N \in L^{(N)}$, and in fact, $\bar{f}^N = \boldsymbol{\Pi}^N f$. Indeed, for every $f \in L^2(I)$ and for every $\varphi \in L^{(N)}$, if we denote by φ_j the value of φ on I_j, we have:

$$
\begin{aligned}
< f, \varphi >= \sum_{j=1}^{N} \int_{I_j} f(x)\varphi(x)dx &= \sum_{j=1}^{N} \varphi_j \int_{I_j} f(x)dx \\
&= \sum_{j=1}^{N} \frac{1}{N} \varphi_j \mu^N(f)_j \\
&= \sum_{j=1}^{N} \int_{I_j} \varphi(x) \bar{f}^N(x)dx =< \bar{f}^N, \varphi > .
\end{aligned}
\tag{9.17}
$$

Similarly, if $w \in L^2(I \times I)$ is symmetric in the sense that $w(x,y) = w(y,x)$ for almost every x and y in I, and if $N \geqslant 1$, we denote by $w^N \in \mathbb{R}^{N \times N}$ the $N \times N$ symmetric matrix defined by its entries w_{ij}^N given by:

$$
w_{ij}^N = N^2 \int_{I_i} \int_{I_j} w(x,y)dxdy, \qquad i,j = 1, \cdots, N,
$$

and by $\bar{w}^N$ the piecewise constant function in $L^2(I \times I)$ which is equal to w_{ij}^N over the square $I_i \times I_j$. Again, $\bar{w}^N$ can be viewed as a conditional expectation:

$$
\bar{w}^N = \mathbb{E}^{\lambda_I \otimes \lambda_I}[w | \mathcal{B}_{I \times I}^N]. \tag{9.18}
$$

where $\mathcal{B}_{I \times I}^N$ is the σ-field generated by the partition of $I \times I$ into the plaquettes $I_i \times I_j$ for $i,j = 1, \cdots, N$.

Lemma 9.4 *If $f \in L^2(I)$ and $w \in L^2(I \times I)$, then*

$$
\lim_{N \to \infty} \bar{f}^N = f, \qquad \lim_{N \to \infty} \bar{w}^N = w
$$

almost surely and in the L^2-sense for the respective L^2 spaces.

Proof: For the sake of simplicity, we prove the result in the case of the dyadic subsequence $N = 2^n$ with a simple martingale argument. Notice that for each $n \geqslant 0$, $\mathcal{B}_I^{2^n} \subset \mathcal{B}_I^{2^{n+1}}$ which shows that $(\bar{f}^{2^n})_{n \geqslant 0}$ is a martingale closed by f, hence the convergence almost sure and in L^2. The argument is exactly the same for w. □

Piecewise Constant Graphons and Associated Operators

If $w \in L^2(I \times I)$, we denote by $\boldsymbol{W}$ the operator with kernel w, i.e. the operator defined on $L^2(I)$ by:

$$
[\boldsymbol{W} f](x) = \int_I w(x,y) f(y)\, dy, \qquad f \in L^2(I).
$$

The operator $\boldsymbol{W}$ is not only a bounded operator, but because it has a square integrable kernel, it is a Hilbert-Schmidt operator. Its Hilbert-Schmidt norm $\|\boldsymbol{W}\|_2$ is given by

$$
\|\boldsymbol{W}\|_2^2 = \int\int_{I \times I} w(x,y)^2 dxdy.
$$

In fact there exists an complete orthonormal system (CONS) $\{\varphi_k; k \geqslant 1\}$ of $L^2(I)$ of eigenfunctions of $\boldsymbol{W}$. Indeed there exist a square summable sequence of real numbers $\{\lambda_k; k \geqslant 1\}$ satisfying:

$$\boldsymbol{W}\varphi_k = \lambda_k\varphi_k, \qquad k \geqslant 1.$$

The square of the Hilbert-Schmidt norm also equals the sum of the squares of the eigenvalues, so

$$\|\boldsymbol{W}\|_2^2 = \sum_{k\geqslant 1}\lambda_k^2.$$

Given an $N\times N$ symmetric matrix η, we denote by $\bar{\eta}^N$ the function on $I\times I$ which is equal to η_{ij} over the square $I_i\times I_j$, and by $\boldsymbol{\eta}^N$ the corresponding kernel operator on $L^2(I)$ defined by:

$$[\boldsymbol{\eta}^N f](x) = \int_I \eta^N(x,y)f(y)dy, \qquad x\in I.$$

Notice that if $f\in L^{(N)}$ is equal to f_j over the interval I_j for $j=1,\cdots,N$, then if $x\in I_i$,

$$[\boldsymbol{\eta}^N f](x) = \sum_{j=1}^N \int_{I_j}\eta^N(x,y)f(y)dy = \frac{1}{N}\sum_{j=1}^N \eta_{ij}f_j = \frac{1}{N}[\eta(\boldsymbol{\pi}^N f)]_i.$$

In particular, the function $\boldsymbol{\eta}^N f$ is constant over the intervals I_j, so it is an element of $L^{(N)}$, showing that the operator $\boldsymbol{\eta}^N$ leaves the space $L^{(N)}$ invariant. Moreover, the vector of its values, namely $\boldsymbol{\pi}^N(\boldsymbol{\eta}^N f)$ is given up to the factor $1/N$ by the product of the original matrix η by the vector $\boldsymbol{\pi}^N f$. To be specific:

$$\boldsymbol{\pi}^N\boldsymbol{\eta}^N f = \frac{1}{N}\eta(\boldsymbol{\pi}^N f) \tag{9.19}$$

for $f\in L^{(N)}$. Notice also that the operator $\boldsymbol{\eta}^N$ associated with the matrix η not only leaves invariant the subspace $L^{(N)}$, but it is identically equal to 0 on the orthogonal complement $L^{(N)\perp}$. Indeed, if $f\in L^{(N)\perp}$, then

$$[\boldsymbol{\eta}^N f](x) = \int_I \eta^N(x,y)f(y)dy = <\eta^N(x,\cdot), f> = 0$$

because for each $x\in I$, the function $I\ni y\mapsto \eta^N(x,y)$ is constant over the intervals I_j, hence belongs to $L^{(N)}$, implying that the inner product is 0 since f is orthogonal to all the functions of $L^{(N)}$. So since $\mathbf{I}-\boldsymbol{\Pi}^N$ is the orthogonal projection of $L^2(I)$ onto $L^{(N)\perp}$, we see that for every $f\in L^2(I)$,

$$\boldsymbol{\eta}^N f = \boldsymbol{\eta}^N(\boldsymbol{\Pi}^N f) + \boldsymbol{\eta}^N([\mathbf{I}-\boldsymbol{\Pi}^N]f) = \boldsymbol{\eta}^N(\boldsymbol{\Pi}^N f)$$

and applying (9.20) to $\boldsymbol{\Pi}^N f$ we get

$$\boldsymbol{\pi}^N\boldsymbol{\eta}^N\boldsymbol{\Pi}^N f = \frac{1}{N}\eta(\boldsymbol{\pi}^N\boldsymbol{\Pi}^N f) \tag{9.20}$$

for $f\in L^2(I)$. So, using the facts that $\boldsymbol{\eta}^N\boldsymbol{\Pi}^N = \boldsymbol{\eta}^N$ and $\boldsymbol{\pi}^N\boldsymbol{\Pi}^N = \mu^N$ we get the formula:

$$\boldsymbol{\eta}^N = \frac{1}{N}(\boldsymbol{\pi}^N)^{-1}\eta\mu^N. \tag{9.21}$$

The final remark on this topic is contained in the following formula. For each bounded continuous function u from $\mathbb{R}$ into $\mathbb{R}$ such that $u(0) = 0$, we have the following equality between bounded operators on $L^2(I)$:

$$\boxed{u(\boldsymbol{\eta}^N) = \frac{1}{N}(\boldsymbol{\pi}^N)^{-1}u(\eta)\mu^N. \qquad (9.22)}$$

The easiest way to prove this equality is to start from the eigenvalues $\lambda_1 \leqslant \lambda_2 \leqslant \cdots \leqslant \lambda_N$ of the symmetric matrix η and the corresponding eigenvectors $f_1, f_2, \cdots, f_N$ forming an orthonormal basis of $\mathbb{R}^N$. By definition, the matrix $u(\eta)$ is diagonalized in the same orthonormal basis $\{f_1, f_2, \cdots, f_N\}$ and its eigenvalues are $u(\lambda_1), u(\lambda_2), \cdots, u(\lambda_N)$. Formula (9.21) implies that the eigenvalues of the operator $\boldsymbol{\eta}^N$ are 0 and $\lambda_1 \leqslant \lambda_2 \leqslant \cdots \leqslant \lambda_N$, an orthonormal basis of eigenvectors being given by any orthonormal basis of $L^{(N)\perp}$, say $\{g_k\}_{k\geqslant 1}$ (all the g_k associated with the eigenvalue 0), and the orthonormal basis $\{\sqrt{N}(\boldsymbol{\pi}^N)^{-1}f_1, \sqrt{N}(\boldsymbol{\pi}^N)^{-1}f_2, \cdots, \sqrt{N}(\boldsymbol{\pi}^N)^{-1}f_N\}$ of $L^{(N)}$. Recall (9.15). Consequently, because of the way functions of a bounded symmetric (self-adjoint) operator are defined, the eigenvalues of the operator $u(\boldsymbol{\eta}^N)$ are 0 (recall that $u(0) = 0$ by assumption) and $u(\lambda_1), u(\lambda_2), \cdots, u(\lambda_N)$, with the same orthonormal basis of functions. Using again formula (9.21), we see that the operator described this way must be $\frac{1}{N}(\boldsymbol{\pi}^N)^{-1}u(\eta)\mu^N$, proving formula (9.22).

Stability Results

The nature and the statements of the following approximation results for operators on $L^2(I)$ are motivated by the analysis of the finite player network game introduced in the previous chapter.

Lemma 9.5 *If $w \in L^2(I \times I)$, and if for every integer $N \geqslant 1$ we denote by $\mathbf{I}^N$ the operator whose kernel is given by the piecewise constant function equal to δ_{ij} over the plaquette $I_i \times I_j$ of the partition of $I \times I$ determined by the partition $\mathcal{P}^{(N)}$ of I, and $\boldsymbol{W}^N$ the operator associated with the piecewise constant function $\bar{w}^N$ defined in* (9.18)*, then the operator $N\mathbf{I}^N$ converges strongly towards the identity operator $\mathbf{I}$ of $L^2(I)$ and the operator $\boldsymbol{W}^N$ converges strongly towards the operator $\boldsymbol{W}$ associated with the graphon w, i.e. for every $f \in L^2(I)$*

$$\lim_{N\to\infty}[N\mathbf{I}^N]f = f, \qquad \textit{and} \qquad \lim_{N\to\infty}\boldsymbol{W}^N f = \boldsymbol{W}f.$$

Proof: If $f \in L^2(I)$, then if $x \in I_i$,

$$\begin{aligned}
[(N\mathbf{I}^N - \boldsymbol{W})f](x) &= \sum_{j=1}^N \int_{I_j}[NI^N(x,y) - w^N(x,y)]f(y)dy \\
&= \sum_{j=1}^N [N\delta_{ij} - w_{ij}^N]\int_{I_j} f(y)dy \qquad (9.23)\\
&= N\int_{I_i} f(y)dy - N^2\sum_{j=1}^N \int_{I_i}\int_{I_j} w(x',y')dx'dy' \int_{I_j} f(y)dy
\end{aligned}$$

the first term converging toward f in $L^2(I)$. As for the second term, we rewrite it as:

$$\begin{aligned} &-N\int_{I_i}\Big(\int_I w(x',y')f(y')dy'\Big)dx' - N\sum_{j=1}^N\int_{I_i}\int_{I_j} w(x',y')\Big[N\int_{I_j} f(y)dy - f(y')\Big]dx'dy' \\ &= -N\int_{I_i}[Wf](x')dx' - N\int_{I_i}\Big(\sum_{j=1}^N\int_{I_j} w(x',y')\Big[N\int_{I_j} f(y)dy - f(y')\Big]dy'\Big)dx'. \end{aligned} \tag{9.24}$$

The first term converges in $L^2(I)$ toward the desired function $\boldsymbol{W}f$, so it remains to show that the second term converges in $L^2(I)$ toward 0. But using the notation (9.16), we can rewrite the negative of this second term as

$$N\int_{I_i}\Big(\sum_{j=1}^N\int_{I_j} w(x',y')\Big[\bar f^N(y') - f(y')\Big]dy'\Big)dx' = N\int_{I_i}[W(\bar f^N - f)](x')dx'$$

Notice that the right hand side is the value at $x \in I_i$ of the piecewise constant function equal to the average of the function $\boldsymbol{W}(\bar f^N - f)$ over each interval of the partition $\mathcal{P}^{(N)}$. By Jensen's inequality the $L^2(I)$ - norm of this function is less than or equal to the $L^2(I)$ - norm of the function $\boldsymbol{W}(\bar f^N - f)$ which goes to 0 because of the result of Lemma 9.4.

In fact the convergence of the second term in (9.23) can be argued with a high level argument. Indeed, Lemma 9.4 says that the kernels of the operators $\boldsymbol{W}^N$ converge toward the kernel of the operator $\boldsymbol{W}$ in $L^2(I \times I)$, which implies the convergence of the operators in the Hilbert-Schmidt norm, hence in the operator norm, hence in the sense of the strong convergence. □

The following result is tailored to the analysis of the finite player network game investigated earlier.

Lemma 9.6 *Let* $(\mathbf{A}^N)_{N\geqslant 1}$ *be a sequence of bounded self-adjoint operators on* $L^2(I)$ *which satisfy the following conditions:*

(i) the spectra of all the $\mathbf{A}^N$ *are contained in* $[0,c]$ *for some* $c > 0$*;*

(ii) $\mathbf{A}^N$ *converges strongly toward a bounded self-adjoint operator* $\mathbf{A}$.

So if for each $N \geqslant 1$ *and* $t \in (0,T)$ *we define:*

$$\boldsymbol{\eta}_t^N = \sqrt{\mathbf{A}^N}(e^{2\sqrt{\mathbf{A}^N}(T-t)} - I)(e^{2\sqrt{\mathbf{A}^N}(T-t)} + I)^{-1}$$

then we have

$$\lim_{N\to\infty}\boldsymbol{\eta}_t^N = \boldsymbol{\eta}_t \qquad \textit{strongly}$$

where $\boldsymbol{\eta}_t = \sqrt{\mathbf{A}}(e^{2\sqrt{\mathbf{A}}(T-t)} - I)(e^{2\sqrt{\mathbf{A}}(T-t)} + I)^{-1}$.

Proof: This is a direct consequence of the fact that for $t \in (0,T)$ fixed, the function $\alpha \mapsto \sqrt{\alpha}(e^{2\sqrt{\alpha}(T-t)} - 1)(e^{2\sqrt{\alpha}(T-t)} + 1)^{-1}$ is continuous and bounded on $[0,c]$. □

Lemma 9.7 *Let* $(\boldsymbol{\eta}^N)_{N\geqslant 1}$ *be a sequence of continuous functions from* $[0,T]$ *into the space of bounded self-adjoint operators on* $L^2(I)$ *which satisfy the following conditions:*

(i) for each $t \in (0,T)$ *the operator* $\boldsymbol{\eta}_t^N$ *converges strongly toward a bounded self-adjoint operator* $\boldsymbol{\eta}_t$*;*

(ii) the function $(0,T) \ni t \mapsto \boldsymbol{\eta}_t$ *is strongly continuous.*

Then for each $N \geqslant 1$ and $s, t \in (0, T)$ the fundamental solution $(U^N(t,s))_{0\leqslant s\leqslant t\leqslant T}$ of the equation

$$\frac{d}{dt}U^N(t,s) = \boldsymbol{\eta}_t^N U^N(t,s) \qquad t \geqslant s$$

with the initial condition $U^N(s,s) = \mathbf{I}$, converges strongly as $N \to \infty$ toward he fundamental solution $(U(t,s))_{0\leqslant s\leqslant t\leqslant T}$ of the equation

$$\frac{d}{dt}U(t,s) = \boldsymbol{\eta}_t U(t,s) \qquad t \geqslant s$$

with the same initial condition $U(s,s) = \mathbf{I}$

Proof: For each $t \in [0, T]$, the operators $\boldsymbol{\eta}_t^N$ converge strongly toward $\boldsymbol{\eta}_t$, both being given by formula (8.24) applied to A^N and A respectively. The uniform boundedness principle (see for example [32, Theorem III.9] or [39, Theorem 4.22]) implies that

$$\sup_{0\leqslant t\leqslant T}\sup_{N\geqslant 1}\|\boldsymbol{\eta}_t^N\| < \infty \quad \text{and} \quad \sup_{0\leqslant t\leqslant T}\|\boldsymbol{\eta}_t\| < \infty. \tag{9.25}$$

As a result, the fundamental solutions of the operator linear differential equation (8.26) are also uniformly bounded in norm. Indeed:

$$\sup_{0\leqslant s\leqslant t\leqslant T}\sup_{N\geqslant 1}\|U^N(t,s)\| \leqslant \sup_{0\leqslant s\leqslant t\leqslant T}\sup_{N\geqslant 0} e^{\int_0^T \|\boldsymbol{\eta}_t^N\|dt} < \infty. \tag{9.26}$$

Next we prove that for $0 \leqslant s \leqslant t \leqslant T$ fixed, $U^N(t,s)$ converges strongly toward $U(t,s)$. Indeed, if $f \in L^2(I)$ is fixed, $U^N(t,s)f$ and $U(t,s)f$ are the solutions of the ordinary differential equations

$$\dot{f}_t^N = -\boldsymbol{\eta}_t^N f_t^N, \qquad \text{and} \qquad \dot{f}_t = -\boldsymbol{\eta}_t f_t,$$

over the interval $t \in [s, T]$ with the same initial conditions $f_s^N = f_s = f$. Notice that

$$\dot{f}_t = -\boldsymbol{\eta}_t f_t = -\boldsymbol{\eta}_t^N f_t + [\boldsymbol{\eta}_t^N - \boldsymbol{\eta}_t]f_t$$

so that

$$f_t = f_t^N + \int_s^t U^N(t,u)[\boldsymbol{\eta}_u^N - \boldsymbol{\eta}_u]f_u \, du.$$

For each fixed $u \in [s, t]$, the strong convergence of $\boldsymbol{\eta}_u^N$ toward $\boldsymbol{\eta}_u$ and the uniform bound (9.26) imply that $\lim_{N\to\infty}\|U^N(t,u)[\boldsymbol{\eta}_u^N - \boldsymbol{\eta}_u]f_u\| = 0$. Moreover, the uniform bounds (9.25) and (9.26) make it possible to use Lebesgue's dominated convergence theorem to conclude that f_t^N converges toward f_t in $L^2(I)$. This proves that for each $0 \leqslant s \leqslant t \leqslant T$, $U^N(t,s)$ converges strongly toward $U(t,s)$. □

Examples

- **Constant graphon** $w(x, y) = p$ for some $p \in (0, 1)$. In this case, for each $N \geqslant 1$, $w_{ij}^N = p$ which says that each player interacts with all the other players equally, the aggregate being, up to the factor p, the sample average of all the states. We recover the Mean Field Game models.

- **Simple threshold graphon** $w(x,y) = \mathbf{1}_{x+y\leqslant 1}$. In this case,

$$w_{ij}^N = \begin{cases} 1 & \text{if } i+j \leqslant N \\ 0.5 & \text{if } N+1 \leqslant i+j \leqslant N+2 \\ 0 & \text{if } i+j > N+2 \end{cases}$$

so a given player $i \in [N]$ will have the same unit strength interaction with all the players $j \in \{1, \cdots, N-i\}$, interaction of strength 0.5 with players $j \in \{N-i+1, N \wedge (N-i+2)\}$ and no interaction with the players $j \in \{N \wedge (N-i+3), N\}$. Recall that the notation $x \wedge y$ stands for the minimum of x and y.

- **Min-max graphon** $w(x,y) = x \wedge y(1 - x \vee y)$. Recall that the notation $x \vee y$ stands for the maximum of x and y. In this case,

$$w_{ij}^N = \begin{cases} \frac{1}{N^2}(i-\frac{1}{2})(N-(j-\frac{1}{2})) & \text{if } 1 \leqslant i \leqslant j-1 \\ \frac{i^3-i+1/3}{N^3} - \frac{i^3-(3/2)i^2+i-1/4}{N^4} & \text{if } i=j \\ \frac{1}{N^2}(j-\frac{1}{2})(N-(i-\frac{1}{2})) & \text{if } 1 \leqslant j \leqslant i-1 \end{cases}$$

so the interaction of a given player $i \in [N]$ with player j decreases when the label j of the other player gets further from the label i.

- **Power law graphon** $w(x,y) = (x \wedge y(1 - x \vee y))xy)^\gamma$ for some $\gamma \in (0, dir0o/3)$. While this graphon function is not bounded, it is still square integrable. In this case,

$$w_{ij}^N = \frac{1}{(1-\gamma)^2 N^{2(1-\gamma)}}[i^{1-\gamma} - (i-1)^{1-\gamma}][j^{1-\gamma} - (j-1)^{1-\gamma}]$$

showing that the strength of the interaction of a given player $i \in [N]$ with player j decreases when the label j of the other increases.

9.2.3 **Back to the Stochastic Differential Graphon Game**

Recall that we started from a graphon w to underpin the interaction between the players and we proved that for each admissible strategy profile $\boldsymbol{\alpha} = (\alpha_t^x)_{x\in I, t\in[0,T]}$, the aggregate states $Z_t^{\boldsymbol{\alpha},x}$ felt by the individual players were deterministic.

Notion of Nash Equilibrium

Definition 9.8 *An admissible strategy profile $\boldsymbol{\alpha} = (\alpha_t^x)_{x\in I, t\in[0,T]}$ is said to be a Nash equilibrium for the game if there exists $\boldsymbol{Z} = (Z_t^x)_{x\in I, t\in[0,T]} \in L^2(I \times [0,T])$ such that*

(i) for almost every $x \in I$

$$(\alpha_t^x)_{0\leqslant t\leqslant T} \in \arg \inf_{(\alpha_t)_{0\leqslant t\leqslant T}} \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\alpha_t|^2 + (X_t - Z_t^x)^2\Big)^2 dt\Big]$$

under the constraint $dX_t = \alpha_t dt + dB_t$ for some Brownian motion $\mathbf{B} = (B_t)_{t\geqslant 0}$ where the drift $(\alpha_t)_{0\leqslant t\leqslant T}$ is adapted to the filtration of this Brownian motion;

(ii) $\boldsymbol{Z} = (Z_t^x)_{x\in I, t\in[0,T]}$ is the deterministic aggregate associated with $\boldsymbol{\alpha} = (\alpha_t^x)_{x\in I, t\in[0,T]}$ via Proposition 9.3.

Bullet point (i) in the statement of the definition states the optimal control problem (almost every) each player $x \in I$ has to solve to find their best response to the aggregate state $\boldsymbol{Z}$. If one uses Pontryagin stochastic maximum principle to solve this problem we end up having to solve the following forward-backward system of stochastic differential equations:

$$\begin{cases} dX_t^x &= -Y_t^x dt + dB_t^x \\ dY_t^x &= -(X_t^x - Z_t^x)dt + \tilde{Z}_t dB_t^x. \end{cases}$$

Unfortunately, this forward-backward system cannot be attacked x by x because of the coupling contained in the candidate $\boldsymbol{Z}$ for the actual state aggregate. Including the fixed point element contained in the second bullet point of the above definition implies that this $\boldsymbol{Z}$ (not to be confused with the $\tilde{Z}_t^x$ integrand processes) should be the actual aggregate in equilibrium. So if we combine all these coupled forward backward systems in one, one obtain (at least formally) the infinite dimensional forward backward system

$$\begin{cases} dX_t^\cdot &= -Y_t^\cdot dt + dB_t^\cdot \\ dY_t^\cdot &= -AX_t^\cdot dt + \tilde{Z}_t dB_t^x. \end{cases}$$

where A is the operator $\mathbf{I} - \boldsymbol{W}$ and $\tilde{Z}_t$ is now an operator since it should depend upon couples (x, y) of players. This coupled forward backward system being linear, it can be solved by first making an ansatz $Y_t^\cdot = \boldsymbol{\eta}_t X_t^\cdot$ for some deterministic function $t \mapsto \boldsymbol{\eta}_t$ with values in the space of bounded operators on $L^2(I)$, and differentiating the ansatz and using the differential given by the above system helps us identify $\boldsymbol{\eta}_t$ as the solution of the Riccati equation

$$\dot{\boldsymbol{\eta}}_t - \boldsymbol{\eta}_t^2 = -A$$

whose solution is given by $\boldsymbol{\eta}_t = \sqrt{A}(e^{2\sqrt{A}(T-t)} - \mathbf{I})(e^{2\sqrt{A}(T-t)} + \mathbf{I})^{-1}$. Consquently, we can conclude that the dynamics of the states in equilibrium are given by the Gaussian process whose covariance was computed in the subsection touted as *Wishful Thinking*, and that this equilibrium is indeed the limit as $N \to \infty$ of the Gaussian processes giving the states evolutions for the equilibria of the finite player network games which we constructed for the purpose of approximation.

9.3 Where Do We Go From Here?

As a conclusion to this chapter, I list informally, in no particular order, a set of questions raised by the limitations of the results and proofs given above, and a few desirable extensions which could be approachable with the tools developed presented in these lectures.

- For the very few models for which we can compute Nash equilibria, could we also solve the central planner optimization problem and as a result, estimate the price of anarchy for these large stochastic networks.
- We chose to use a version of the Pontryagin stochastic maximum principle to solve the network game model analyzed in this chapter. A more analytic approach would

be based on the use of the dynamic programming principle and the solution of a system of Hamilton-Jacobi-Bellman equations coupled with a Kolmogorov-Fokker-Planck equation. Could such an approach be implemented to control the limits of finite player network games and recover the systems touted by Peter Caines and his collaborators in a series of recent notes and conference proceedings? See for example [14].

- The analysis of static graphon games lead to expressions for Nash equilibria involving geometric properties of the underlying graph structure, such as centrality indexes for example. The form of the equilibrium strategy profiles we found in the models of dynamic stochastic graphon games we managed to solve in this chapter involved the operator $\mathbf{I} - \boldsymbol{W}$ in a form which was not explicit enough to identify the role of the geometry of the graph in the expression of the player strategies. Making this connection explicit will be helpful to understand how the nature of the graph influences the equilibria.
- We approximated the limiting continuum graph structure by deterministic analogs based on piecewise constant graphons, and the corresponding finite player games were based on these approximations. Graphons are traditionally introduced as limits of random graphs. Some of these random graphs are obtained by sampling methods, the most popular ones being described in Chapter 4. Is it possible to approach the analysis of network games with a continuum of players as limits of finite player games interacting through these random graph approximations? How different would the analysis be? Could we get a better insight using these types of approximations?
- Graphons are limits of dense graphs. In game models underpinned by such graphs, each player interacts with a significant proportion of the other players. Remember that in a mean field game, each player interacts with all the players, even if only through average quantities like means, ... Mathematical structures have been introduced to capture the limits of sparse random graphs (e.g. bounded degree graphs, Erdos-Renyi graphs with small connection probabilities, graphs with power law degree, ...) whose limits would naturally lead to the zero graphon. Could we use these to analyze the limits of network games with sparse interactions?

References

1. C.D. Aliprantis and K. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Verlag, 2006.
2. R. Aumann. Markets with a continuum of traders. *Econometrica*, 32:39 – 50, 1964.
3. R. J. Aumann and L. S. Shapley. *Values of Non-Atomic Games*. Princeton University Press, 2015.
4. M. Barlo and G. Carmona. Strategic behavior in non-atomic games. Technical report, 2015.
5. P. Billingsley. *Convergence of Probability Measures. Third edition.* John Wiley & Sons, Inc., 1995.
6. P. Bonacich. Power and centrality: A family of measures. *American Journal of Sociology*, 92(5):1170–1182, 1987.
7. Phillip Bonacich. Factoring and weighting approaches to status scores and clique identification. *Journal of mathematical sociology*, 2(1):113–120, 1972.
8. Phillip Bonacich. Technique for analyzing overlapping memberships. *Sociological methodology*, 4:176–185, 1972.
9. H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Universitext. Springer Verlag, 2011.
10. P. Cardaliaguet. Notes from P.L. Lions' lectures at the Collège de France. Technical report, `https://www.ceremade.dauphine.fr/~cardalia/MFG100629.pdf`, 2012.
11. R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games: vol. I, Mean Field FBSDEs, Control, and Games*. Stochastic Analysis and Applications. Springer Verlag, 2017.
12. P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. *Rendiconti di Matematica e delle sue Applicazioni*, 28:33 – 61, 2008.
13. P. Erdös and A. Rényi. On random graphs. i. *Publicationes Mathematicae*, 6:290 – 297, 1959.
14. S. Gao and P.Caines. Graphon control of large-scale networks of linear systems. Technical report, arXiv:1807.03412v5.
15. E.N. Gilbert. Random graphs. *Annals of Mathematical Statistics*, 30:1141 – 1144, 1959.
16. J. Harsanyi. Games with incomplete information played by bayesian players. *Management Science*, pages 159 – 182, 320 – 334, 486 – 502, 1967-68.
17. S. Hart and D. Monderer. Potentials and weighted values of non-atomic games. *Mathematics of Operations Research*, 22:619 – 630, 1997.
18. P.W. Holland, K.B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5:109 – 137, 1983.
19. Svante Janson. Graphons, cut norm and distance, couplings and rearrangements. Technical report, Department of Mathematics, Uppsala University, 2010.
20. S. Kakutani. A generalization of brouwer's fixed point theorem. *Duke Mathematical Journal*, 8:457 – 459, 1941.
21. I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus, 2nd edition*. Springer Verlag, 1991.
22. Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953.

23. D. Lacker and K. Ramanan. Rare nash equilibria and the price of anarchy in large static games. Technical report.
24. László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Society, 2012.
25. A. Mas-Colell. On a theorem by schmeidler. *Journal of Mathematical Economics*, 13:201 – 206, 1984.
26. D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14:124 – 143, 1996.
27. D.P. Myatt and C. Wallace. Endogenous information acquisition in coordination games. *Review of Economic Studies*, 79:340 – 374, 2012.
28. J. Nash. Equilibrium points in n-person games. *Proceedings National Academy of Sciences USA*, 36:48–49, 1950.
29. Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
30. F. Parise and A. Ozdaglar. Graphon games. Technical report, MIT, 2018.
31. M. Reed and B. Simon. *Methods of modern mathematical physics: I Functional Analysis*. Academic Press, 1980.
32. M. Reed and B. Simon. *Methods of modern mathematical physics. I. Functional analysis*. Academic Press, Inc. San Diego, 1980.
33. R.W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:64 – 67, 1973.
34. W. Rudin. *Real and complex analysis*. McGraw-Hill, New York, 1966.
35. D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7:295 – 300, 1973.
36. G. Scutari, D. P. Palomar, F. Facchinei, and J.S. Pang. Convex optimization, game theory, and variational inequality theory,.
37. Y. Sun. The exact law of large numbers via fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126:31–69, 2006.
38. A. Tarski. A lattice theoretic fixed point theorem. *Pacific Journal of Mathematics*, 5:285–309, 1955.
39. J. Weidmann. *Linear Operators in Hilbert Spaces*. Springer Verlag, 1980.

NOTATION INDEX

AUTHOR INDEX

SUBJECT INDEX