

A Mean Field Game of Optimal Stopping

Marcel Nutz*

First version: May 30, 2016. This version: November 30, 2017

Abstract

We formulate a stochastic game of mean field type where the agents solve optimal stopping problems and interact through the proportion of players that have already stopped. Working with a continuum of agents, typical equilibria become functions of the common noise that all agents are exposed to, whereas idiosyncratic randomness can be eliminated by an Exact Law of Large Numbers. Under a structural monotonicity assumption, we can identify equilibria with solutions of a simple equation involving the distribution function of the idiosyncratic noise. Solvable examples allow us to gain insight into the uniqueness of equilibria and the dynamics in the population.

Keywords Mean Field Game; Optimal Stopping; Bank-Run
AMS 2010 Subject Classification 91A13; 60G40; 91A15; 91A55

1 Introduction

Stochastic games with a large number n of players are notoriously intractable. Mean field games were introduced by Lasry and Lions [24, 25, 26] and Huang, Malhamé, and Caines [19, 20] to study Nash equilibria in the limiting regime where n tends to infinity and the players interact symmetrically through the empirical distribution of the private states of all players. Given such a distribution μ , each player typically solves a standard control problem; that is, controls a diffusion while paying some cost of effort. On the other hand,

*Departments of Statistics and Mathematics, Columbia University, mnutz@columbia.edu. Research supported by an Alfred P. Sloan Fellowship and NSF Grant DMS-1512900. This work has enormously profited from discussions with Bruno Bouchard, René Carmona, Ioannis Karatzas, Daniel Lacker, José Scheinkman, Nizar Touzi, and detailed comments from two anonymous referees, to whom the author is most grateful.

the reward (and possibly the diffusion) depend on μ , which is in turn determined by the actions of all agents. In the analytic theory, such a system is described by a coupled system of nonlinear partial differential equations (PDEs): a Hamilton–Jacobi–Bellman equation describes the optimal control problem when μ is given, and a Kolmogorov-type equation describes the evolution of μ over time as a result of the optimal controls. One of the major difficulties is that the former equation naturally starts from a terminal condition and runs backward in time, whereas the latter runs forward to ensure the consistency of μ ; we refer to [6, 17] for background. In a probabilistic version of the theory, the stochastic maximum principle is used and the system of PDEs is replaced by a coupled forward-backward stochastic differential equation; cf. [4, 7, 8, 9]. In the simplest case, the agents are exposed to idiosyncratic i.i.d. noise (essentially, an independent Brownian motion for each diffusion equation) and thus the equilibria are formulated as deterministic. More recently, the presence of an additional common noise and stochastic equilibria have received considerable attention; see [10, 13, 16, 22, 28]. A wide range of applications from production models to population dynamics have emerged over the last decade, several of them summarized in [18]; see also [12] for a recent model of systemic risk and [13] for price impact in finance.

While mean field games were introduced as a tractable model for a large stochastic game, they are still rather complex. To the best of our knowledge, the only case that can be solved explicitly is linear-quadratic control (linear dynamics, quadratic cost). This situation has been studied in detail; see [2, 3, 5, 12]. In other cases, one generally has to settle for an abstract description by a coupled system of nonlinear equations.

The main aim of the present paper is to formulate a tractable game of mean field type where the properties of equilibria can be understood somewhat more directly. In our case, the agents will be solving optimal stopping problems rather than diffusion control¹. While in a standard mean field game the (spatial) location of the players matters, the state space here becomes binary: each player either has stopped or is still in the game, and the interaction occurs through the number of players that have already stopped. This structure seems appealing due to its simple interpretation and a wide range of possible applications from bank-run models to traffic optimization.

¹ The possible interest of such a game was first pointed out to the author by René Carmona. Section 2 of [18] can be seen as a predecessor, at least in spirit: in a toy example called “When Does the Meeting Start?” the agents indirectly control their arrival times at a prespecified location.

On the other hand, it produces an inherent discontinuity in the game: as is well-known in economics (e.g., [14, 27]), games of optimal timing may easily degenerate in that all players stop at the same time. Thus, one of the challenges is to produce a class of models where typical equilibria are non-trivial.

Specifically, we shall study a continuous-time stochastic game with a continuum of players. In equilibrium, each agent i will be solving an optimal stopping problem of the form

$$\sup_{\tau} E \left[\exp \left(\int_0^{\tau} r_s ds \right) \mathbf{1}_{\{\theta > \tau\} \cup \{\theta = \infty\}} \right];$$

it has two competing forces. The process r can be interpreted as a reward or *interest rate* that is accrued as long as the agent does not stop, thus incentivizing the agent to stay in the game. On the other hand, there is a random time θ of *default* (of the interest-paying institution): the agent will lose everything if θ happens before she leaves the game. While the default happens as a “surprise” to the agent, the distribution of θ is governed by an intensity process γ^i that is known to the agent: the larger γ^i , the more likely it is that default happens soon. More precisely, θ is modeled as the first jump time of a Cox process with intensity γ^i . This leads to a tractable solution of the single-agent optimal stopping problem—we are taking inspiration from the finance literature (e.g., [23, Chapter 5]) where it is well-known that a defaultable bond in a similar setting will be priced just like a non-defaultable one, but with an adjusted interest rate $r - \gamma^i$.

The agents are heterogeneous in their views on the distribution of the default—we think of the intensity γ^i as depending on the subjective probability used by agent i . As a result, the players face different optimal stopping problems and may stop at different times. The agents’ views on the default intensity will also be influenced by how many players have already stopped; more precisely, the proportion $\rho_t \in [0, 1]$ of players that have left the game by time t . This process is observed by all agents and creates an interaction of mean field type: if ρ_t is larger, the intensity of any player will also be larger, meaning that the perceived default will happen sooner. As in bank-run models, this represents that the default of the institution is more likely if more customers have abandoned ship.

While we defer the general formulation of the setting to Section 5, a typical model may postulate that γ^i is of the form

$$\gamma_t^i = X_t + Y_t^i + c\rho_t. \tag{1.1}$$

Here X plays the role of a common noise (the same for all agents) whereas Y^i is an idiosyncratic noise that will be i.i.d. within the population. Depending on the application, one may interpret X and Y^i as public and private signals, respectively, or see their sum as a noisy observation of the true signal X . Moreover, the constant $c \geq 0$ governs the strength of interaction; that is, how much the agents' views are affected by ρ_t . Suppose that τ^i is the stopping time chosen by agent i , and that the continuum of agents is represented by an atomless probability space $(I, \mathcal{I}, \lambda)$. Then,

$$\rho_t(\omega) = \lambda\{i : \tau^i(\omega) \leq t\} \quad (1.2)$$

is the “proportion” of players that have stopped prior to time t . This can also be seen as the cumulative distribution function (c.d.f.) at time t of the empirical measure that describes the evolution of the system on $I \times \{0, 1\}$, recording for each agent i whether stopping has occurred (1) or not (0).

If we start with a given process ρ , the intensities γ^i of the agents are determined. Let us suppose that the associated optimal stopping problems have solutions $(\tau^i)_{i \in I}$. Tacitly assuming a suitable measurability, we may then consider the process $\lambda\{i : \tau^i(\omega) \leq t\}$, and if it satisfies (1.2), we shall say that ρ and $(\tau^i)_{i \in I}$ form an equilibrium. Since we are working with a continuum of players, the decision of a single agent does not influence ρ , and hence this notion corresponds to a Nash equilibrium: given the strategies of the other players, each player is behaving optimally.

Our main result (Theorem 5.1) relates equilibria ρ_t to the solution of a finite-dimensional equation. For instance, in the case of (1.1), it reads

$$1 - u = F_t(r - x - cu), \quad u \in [0, 1], \quad (1.3)$$

where F_t is the c.d.f. of the idiosyncratic noise Y_t and r is the (constant) interest rate. If $\rho(t, x)$ is the (maximal) solution u at time t and X_t is the common noise, then $\rho(t, X_t)$ describes an equilibrium, and a converse is also established. This simple equation allows us to understand the structure and multiplicity of equilibria in some detail. Two ingredients are important for the tractability of our setting. On the one hand, we use an Exact Law of Large Numbers to completely eliminate idiosyncratic randomness (the associated mathematical setup will be discussed later on). This idea and its many incarnations are well-known in economics; see [1, 15, 21, 31] to cite but a few examples. On the other hand, we impose the structural assumption that $\gamma^i - r$ is increasing, and as we shall see, that leads to a simple solution of the single-agent problem. In comparison to the coupled forward-backward system of equations that is common in the literature on mean field games, one

may say that the Hamilton–Jacobi–Bellman part becomes irrelevant because we know the solution of the single-agent problem (given ρ) in feedback form, whereas our equation represents the Kolmogorov forward equation—indeed, we may take time derivatives in (1.3) to find a PDE for ρ , or an ODE in the case without common noise.

While the main aim of the present paper is to formulate a tractable example of a mean field game of optimal stopping, there are important aspects that are not discussed. Three major questions are the passage to the limit from a game with n players, what happens when the monotonicity condition is dropped, and the analysis of applications. Recent results on these can be found in [11], and we would like to emphasize that much of that work was carried out in parallel or before ours. In particular, it is shown that in our model of Section 5.1, equilibrium strategies from the continuum formulation are ε -equilibrium strategies in an n -player game with large n . Moreover, a general framework for mean field games of optimal stopping (or “timing”) is introduced and analyzed, and applications to bank run models are discussed.

The remainder of this paper is organized as follows. In the next section, we describe the game more rigorously and analyze the single-player problem in detail, whereas Section 3 introduces the mathematical setting that allows for an Exact Law of Large Numbers. Section 4 analyzes an insightful toy model without common noise; we discuss examples of (non-)uniqueness and the impact of noise and strength of interaction on the continuity of equilibria in time. Finally, Section 5 treats the general model with common noise.

2 Description of the Game

Let $(I, \mathcal{I}, \lambda)$ be a probability space; each $i \in I$ will correspond to an agent. Moreover, let (Ω, \mathcal{F}, P) be another probability space, to be used as the sample space. We suppose that (Ω, \mathcal{F}, P) is equipped with right-continuous filtrations $\mathbb{G}^i = (\mathcal{G}_t^i)_{t \in \mathbb{R}_+}$ and an exponentially distributed random variable \mathcal{E} which is independent of \mathbb{G}^i for all $i \in I$. We interpret \mathbb{G}^i as the information available to agent i . Finally, let r be a real-valued and locally integrable process (i.e., Lebesgue-integrable on bounded intervals) which is \mathbb{G}^i -progressively measurable for all $i \in I$; that is, observed by all agents.

2.1 Single-Agent Problem

We first consider the optimization problem for a fixed agent $i \in I$ and we denote by \mathcal{T}^i the set of all \mathbb{G}^i -stopping times. Let $\gamma^i \geq 0$ be a \mathbb{G}^i -progressively

measurable process which is locally integrable and consider the random time

$$\theta^i = \inf \left\{ t : \int_0^t \gamma_s^i ds = \mathcal{E} \right\}.$$

One may think of θ^i as the first jump time of a Cox process with intensity γ^i . The default time θ^i depends on i , which will allow us to write all optimal stopping problems under a common probability measure P . Alternately, we could deal with a single random time on a canonical space and endow the agents with subjective probabilities P^i . We have found the former solution easier to write, and they are equivalent in that the agents' decisions (and the equilibria) only depend on the distribution of the intensity. For technical reasons, we shall assume that

$$\left. \begin{array}{l} r^+ \text{ is integrable on } [0, \infty), P\text{-a.s.}, \\ \text{or} \\ \inf\{t : \gamma_t^i - r_t \geq 0\} < \infty, P\text{-a.s.} \end{array} \right\} \quad (2.1)$$

See also Remark 2.2 below for the necessity of such a condition. We then have the following result on the single-agent problem.

Lemma 2.1. *Suppose that $\gamma^i - r$ is increasing² and (2.1) holds. Then,*

$$\tau^i := \inf\{t : \gamma_t^i - r_t \geq 0\} \in \mathcal{T}^i$$

is a solution of the optimal stopping problem³

$$\sup_{\tau \in \mathcal{T}^i} E \left[\exp \left(\int_0^\tau r_s ds \right) \mathbf{1}_{\{\theta^i > \tau\} \cup \{\theta^i = \infty\}} \right]. \quad (2.2)$$

If the value of (2.2) is finite, then τ^i is minimal among all solutions, and if, in addition, $\gamma^i - r$ is strictly increasing, then τ^i is the unique solution.

Proof. Due to the increase of $\gamma^i - r$ and the right-continuity of \mathbb{G}^i , the right limit process ζ of $\gamma^i - r$ exists and is \mathbb{G}^i -progressively measurable. As $\{\tau^i \leq t\} = \{\zeta_t \geq 0\} \in \mathcal{G}_t^i$, we have $\tau^i \in \mathcal{T}^i$.

Let $\tau \in \mathcal{T}^i$ be such that r^+ is integrable on $[0, \tau)$, P -a.s. Using also the independence of \mathcal{E} and \mathbb{G}^i , we have

$$\begin{aligned} P \left[\{\theta^i > \tau\} \cup \{\theta^i = \infty\} \middle| \mathcal{G}_\tau^i \right] &= P \left[\int_0^\tau \gamma_s^i ds < \mathcal{E} \middle| \mathcal{G}_\tau^i \right] \\ &= E \left[\exp \left(- \int_0^\tau \gamma_s^i ds \right) \middle| \mathcal{G}_\tau^i \right]. \end{aligned}$$

²Increase is to be understood in the non-strict sense throughout the paper.

³We use the convention that $\int_0^\infty r_s ds := -\infty$ if $\int_0^\infty r_s^+ ds = \int_0^\infty r_s^- ds = \infty$.

Hence,

$$E \left[\exp \left(\int_0^\tau r_s ds \right) \mathbf{1}_{\{\theta^i > \tau\} \cup \{\theta^i = \infty\}} \middle| \mathcal{G}_\tau^i \right] = E \left[\exp \left(\int_0^\tau (r_s - \gamma_s^i) ds \right) \middle| \mathcal{G}_\tau^i \right]$$

and finally

$$E \left[\exp \left(\int_0^\tau r_s ds \right) \mathbf{1}_{\{\theta^i > \tau\} \cup \{\theta^i = \infty\}} \right] = E \left[\exp \left(\int_0^\tau (r_s - \gamma_s^i) ds \right) \right]. \quad (2.3)$$

If we are in the first case of (2.1), our integrability condition holds for all $\tau \in \mathcal{T}^i$ and as $r - \gamma^i$ is decreasing, the representation on the right-hand side shows that τ^i is optimal. In the second case of (2.1), as r is locally integrable, we still have (2.3) for every finite-valued $\tau \in \mathcal{T}^i$, and we deduce that τ^i is optimal among all those stopping times. If τ is a general stopping time and $N \in \mathbb{N}$, Fatou's lemma and that optimality yield

$$\begin{aligned} E \left[\exp \left(\int_0^\tau r_s ds \right) \mathbf{1}_{\{\theta^i > \tau\} \cup \{\theta^i = \infty\}} \right] &\leq \liminf_{N \rightarrow \infty} E \left[\exp \left(\int_0^{\tau \wedge N} r_s ds \right) \mathbf{1}_{\theta^i > \tau \wedge N} \right] \\ &\leq E \left[\exp \left(\int_0^{\tau^i} r_s ds \right) \mathbf{1}_{\theta^i > \tau^i} \right], \end{aligned}$$

so that τ^i is in fact optimal among all stopping times. The remaining assertions can also be inferred from (2.3). \square

We see from the proof of Lemma 2.1 how the increase of $\gamma^i - r$ leads to a simple solution of the optimal stopping problem and that will contribute greatly to the tractability of equilibria. We have little else to say in defense of that condition.

Remark 2.2. As usual in infinite-horizon stopping problems, an integrability assumption is necessary to ensure existence of an optimal stopping time. In particular, if $r > \gamma^i > 0$ are constant, then $\tau^i = \infty$ which is clearly not optimal as then $P(\{\theta^i > \tau^i\} \cup \{\theta^i = \infty\}) = 0$. If we consider the same problem with a horizon $T \in (0, \infty)$, no extra assumption is necessary.

Remark 2.3. In Lemma 2.1 and the remainder of this paper, we use strict monotonicity of $\gamma^i - r$ as a simple sufficient condition for the uniqueness of τ^i . In specific cases one may want to use a sharper condition; for instance, discrete-time problems can be embedded in our results by using piecewise constant processes, but then the notion of strict monotonicity needs to be adapted.

2.2 Interaction

While the i th agent chooses to stop at $\tau^i \in \mathcal{T}^i$, the agents will interact through the “proportion” of agents that have already stopped. Indeed, we shall specify γ^i as a functional depending on a process ρ , and then an *equilibrium* will be a collection of stopping times $\tau^i \in \mathcal{T}^i$ which solve (2.2) for λ -almost all $i \in I$ and such that

$$\rho_t = \lambda\{i : \tau^i \leq t\}. \quad (2.4)$$

When λ is atomless, the decision of a single agent does not influence this quantity and hence we indeed have a Nash equilibrium. Clearly, the process ρ will necessarily be increasing and $[0, 1]$ -valued. Moreover, we think of ρ as being observed by all agents, so ρ will be \mathbb{G}^i -adapted for all i .

In (2.4), we are tacitly assuming that the set on the right-hand side is \mathcal{I} -measurable P -a.s., which is highly nontrivial for a continuum of i.i.d. random variables. The setup that can guarantee this is discussed in the next section. Before that, however, let us illustrate the concepts introduced thus far by a very simple example where the agents do not use any signals except ρ .

Example 2.4 (Sunspot). (i) Let λ be the Lebesgue measure on $I = [0, 1]$, let $r > 0$ be constant and let X be a right-continuous, increasing process on (Ω, \mathcal{F}, P) , progressively measurable for the common, right-continuous filtration $\mathbb{G}^i = \mathbb{G}$ (i.e., the same for all agents) and such that $X_\infty > 1$. Suppose agent $i \in [0, 1]$ believes in the intensity

$$\gamma_t^i = (r - i + \rho_t) \vee 0.$$

Thus, i acts as an index of “optimism” or “risk tolerance”—agents with higher index believe that θ^i will happen later. We claim that

$$\rho_t = (X_t \wedge 1) \vee 0$$

yields an equilibrium. Indeed, the optimal stopping times are then given by

$$\tau^i = \inf\{t : \gamma_t^i - r \geq 0\} = \inf\{t : \rho_t = i\} = \inf\{t : X_t \geq i\}$$

and that results in

$$\lambda\{i : \tau^i \leq t\} = \lambda\{i : X_t \geq i\} = (X_t \wedge 1) \vee 0 = \rho_t;$$

note that the second condition of (2.1) is satisfied.

For instance, the choice $X_t = t$ gives rise to $\rho_t = t \wedge 1$ and $\tau^i = i$, showing that the agents stop at deterministic times which are uniformly distributed

over the time interval $[0, 1]$. If X_0 is strictly positive, we see that some of the agents stop instantaneously at $t = 0$, whereas if X_0 is strictly negative, it will take a while before any agents stop.

(ii) A similar equilibrium exists in a finite player game. Let $n \in \mathbb{N}$ and let λ be the normalized counting measure on $I = \{1/n, 2/n, \dots, 1\}$; this corresponds to n equally weighted agents. In the same setting as in (i), an equilibrium is described by $\tau^i = \inf\{t : X_t \geq i\}$ and

$$\rho_t = \lfloor (X_t \wedge 1) \vee 0 \rfloor,$$

where $\lfloor x \rfloor := \max\{s \in I : s \leq x\}$.

Remark 2.5. (i) In the preceding example, the process X is not part of the functional form of γ^i ; essentially, *any* process X gives rise to an equilibrium. The interpretation is that if all agents agree that some commonly observed signal X is relevant, it indeed becomes relevant—the name “sunspot equilibrium” suggests itself. We shall see in Example 4.3 that this situation is a degenerate limit of a model where uniqueness is the typical case.

(ii) If all agents are perfectly identical, the problem will degenerate since they will (typically) all stop at the same time. Thus, in the above example, the agents have been made heterogeneous by varying the risk tolerance. This is not necessary when the agents are already heterogeneous due to private signals, as in the later sections.

3 Mathematical Setting and Exact Law of Large Numbers

In this section, we introduce the setting to accommodate a continuum of agents and their private signals. Let $(I, \mathcal{I}, \lambda)$ be an atomless (hence, uncountable) probability space and let (Ω, \mathcal{F}, P) be another probability space.

Definition 3.1. A family $(f_i)_{i \in I}$ of random variables on (Ω, \mathcal{F}, P) is *essentially pairwise independent* if for λ -almost all $i \in I$, f_i is independent of f_j for λ -almost all $j \in I$. The family is *essentially pairwise i.i.d.* if, in addition, all f_i have the same distribution. Analogously, for a σ -field $\mathcal{C} \subseteq \mathcal{F}$, the family $(f_i)_{i \in I}$ is *essentially pairwise conditionally independent* given \mathcal{C} , if for λ -almost all $i \in I$, f_i is conditionally independent of f_j given \mathcal{C} for λ -almost all $j \in I$.

In what follows, we need to work on a probability space that is larger than the usual product⁴ $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$, because the latter does not

⁴ Here and below, we use the convention that the product σ -field $\mathcal{I} \otimes \mathcal{F}$ is completed.

support relevant families of i.i.d. random variables. More precisely, we have the following fact; see, e.g., [31, Proposition 2.1].

Remark 3.2. If $f : I \times \Omega \rightarrow \mathbb{R}$ is an $\mathcal{I} \otimes \mathcal{F}$ -measurable function such that $f(i, \cdot)$, $i \in I$ are essentially pairwise i.i.d., then f is constant $\lambda \otimes P$ -a.s.

Following [31], we say that a probability space $(I \times \Omega, \Sigma, \mu)$ is an *extension* of the product $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$ if Σ contains $\mathcal{I} \otimes \mathcal{F}$ and the restriction of μ to $\mathcal{I} \otimes \mathcal{F}$ coincides with $\lambda \otimes P$. It is a *Fubini extension* if, in addition, any μ -integrable⁵ function $f : I \times \Omega \rightarrow \mathbb{R}$ satisfies the assertion of Fubini's theorem⁶; that is,

- (i) for λ -almost all $i \in I$, the function $f(i, \cdot)$ is P -integrable,
- (ii) for P -almost all $\omega \in \Omega$, the function $f(\cdot, \omega)$ λ -integrable,
- (iii) $i \mapsto \int f(i, \cdot) dP$ is λ -integrable, $\omega \mapsto \int (\cdot, \omega) d\lambda$ is P -integrable, and

$$\int f d\mu = \iint f(i, \omega) P(d\omega) \lambda(di) = \iint f(i, \omega) \lambda(di) P(d\omega).$$

Let $(I \times \Omega, \Sigma, \mu)$ be a Fubini extension of $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$. Then, essentially pairwise independent families satisfy an exact version of the Law of Large Numbers. The simplest version runs as follows—more generally, an exact version of the Glivenko–Cantelli Theorem holds; cf. [31, Corollary 2.9].

Proposition 3.3 (Exact Law of Large Numbers). *Let $f : I \times \Omega \rightarrow \mathbb{R}$ be μ -integrable. If $f(i, \cdot)$, $i \in I$ are essentially pairwise i.i.d. with a distribution having mean m , then $\int f(\cdot, \omega) d\lambda = m$ for P -almost all $\omega \in \Omega$.*

We shall also need a conditional version as provided by [30, Corollary 2].

Proposition 3.4 (Conditional Exact Law of Large Numbers). *Let $\mathcal{C} \subseteq \mathcal{F}$ be a countably generated σ -field and let $f : I \times \Omega \rightarrow \mathbb{R}$ be μ -integrable. If $f(i, \cdot)$, $i \in I$ are essentially pairwise conditionally independent given \mathcal{C} , then $\int f(\cdot, \omega) d\lambda = \int E^\mu[f|\mathcal{I} \otimes \mathcal{C}](\cdot, \omega) d\lambda$ for P -almost all $\omega \in \Omega$.*

In view of Remark 3.2, it is not obvious that the preceding propositions are not vacuous—that is guaranteed by the next two results.

The space $(I \times \Omega, \Sigma, \mu)$ is called *rich* if there exists a Σ -measurable function $f : I \times \Omega \rightarrow \mathbb{R}$ such that $f(i, \cdot)$, $i \in I$ are essentially pairwise i.i.d. with

⁵ That is, f is measurable for the μ -completion of Σ and $\int |f| d\mu < \infty$.

⁶ Since Σ may be strictly larger than $\mathcal{I} \otimes \mathcal{F}$, this is not automatic.

a uniform distribution on $[0, 1]$. Like an atomless probability space supports random variables with any given distribution, a rich Fubini extension supports essentially pairwise i.i.d. families with any given distribution; cf. [31, Corollary 5.4].

Lemma 3.5. *Let $(I \times \Omega, \Sigma, \mu)$ be a rich Fubini extension of $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$, let S be a Polish space and let ν be a Borel probability measure on S . There exists a Σ -measurable function $f : I \times \Omega \rightarrow S$ such that $f(i, \cdot)$, $i \in I$ are essentially pairwise independent and $f(i, \cdot)$ has distribution ν for all $i \in I$.*

Lemma 3.6. *There exist atomless probability spaces $(I, \mathcal{I}, \lambda)$ and (Ω, \mathcal{F}, P) such that $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$ admits a rich Fubini extension.*

This is part of the assertion of [31, Proposition 5.6] which also shows that one can take $I = [0, 1]$ and $\Omega = \mathbb{R}^{[0,1]}$. The main result of [32] shows that, in addition, one can take λ to be an extension of the Lebesgue measure (but not the Lebesgue measure itself). A different construction, avoiding nonstandard analysis, is presented in [29].

4 A Toy Model

In this section, we discuss a simple setting where the agents' signals are i.i.d.; that is, pure idiosyncratic noise. While not suitable for most applications, this will allow us to explain the effect of the Exact Law of Large Numbers in our model and to discuss some finer questions of uniqueness and nondegeneracy without too many distractions.

Consider the setup introduced in Section 2 with atomless probability spaces $(I, \mathcal{I}, \lambda)$ and (Ω, \mathcal{F}, P) , and let $(I \times \Omega, \Sigma, \mu)$ be a Fubini extension of their product. For each $i \in I$, let $Y^i \geq 0$ be a right-continuous, increasing, \mathbb{G}^i -progressively measurable process. We assume that for each $t \geq 0$, $(i, \omega) \mapsto Y_t^i(\omega)$ is Σ -measurable and that Y_t^i , $i \in I$ are essentially pairwise i.i.d. Moreover, we assume that the distribution of Y_t^i has no atoms; that is, its c.d.f. $y \mapsto F_t(y) := P\{Y_t^i \leq y\}$ is continuous.

Proposition 4.1. *Let $r \in \mathbb{R}$ and $c \in \mathbb{R}_+$. The equation*

$$1 - u = F_t(r - cu), \quad u \in [0, 1] \tag{4.1}$$

has a maximal solution $\rho(t) \in [0, 1]$ for every $t \geq 0$, and $t \mapsto \rho(t)$ is right-continuous. Define also

$$\gamma_t^i = Y_t^i + c\rho(t), \quad \tau^i = \inf\{t : Y_t^i + c\rho(t) = r\},$$

and assume that (2.1) is satisfied for all i .

(i) Then, ρ and $(\tau^i)_{i \in I}$ define an equilibrium: $\tau^i \in \mathcal{T}^i$ is an optimal stopping time for agent i , the mapping $(i, \omega) \mapsto \tau^i(\omega)$ is Σ -measurable, and

$$\lambda\{i : \tau^i \leq t\} = \rho(t) \quad P\text{-a.s. for all } t \geq 0.$$

(ii) Conversely, let $\bar{\rho}$ be a right-continuous function corresponding to an equilibrium. If γ^i is strictly increasing for all i , then $\bar{\rho}(t)$ is a solution of (4.1) for every $t \geq 0$.

Sketch of Proof. The proposition is a special case of Theorem 5.1 that will be proved later on, so we shall only explain the most important steps.

(a) We first argue that ρ is well-defined, increasing and right-continuous. Let us consider, for a right-continuous and increasing function $F : \mathbb{R} \rightarrow [0, 1]$, the zeros of

$$G(u) := F(r - cu) - 1 + u, \quad u \in [0, 1].$$

We have $G(0) = F(r) - 1 \leq 0$ and $G(1) = F(r - c) \geq 0$. Moreover, G is left-continuous and its jumps satisfy $\Delta G \leq 0$. Thus, G must have at least one zero in $[0, 1]$. If $u_n \uparrow u$ is a maximizing sequence of zeros in $[0, 1]$, then $G(u) = 0$ by left-continuity and u is the maximal zero.

Next, write $G_t(u) := F_t(r - cu) - 1 + u$ and let $\rho(t)$ be the maximal zero for each $t \geq 0$. The increase and the right-continuity of $t \mapsto \rho(t)$ can be inferred from the increase of Y and the right-continuity of Y and the continuity of $y \mapsto F_t(y)$, respectively—we defer the details.

(b) Next, we verify that ρ and $(\tau^i)_{i \in I}$ determine an equilibrium. It follows from (a) that $\gamma^i = Y^i + c\rho$ is increasing and right-continuous; hence, Lemma 2.1 yields that $\tau^i \in \mathcal{T}^i$ is an optimal stopping time for all $i \in I$ and that $\{(i, \omega) : \tau^i(\omega) \leq t\} = \{(i, \omega) : Y_t^i(\omega) + c\rho(t) \geq r\} \in \Sigma$. Using the Exact Law of Large Numbers of Proposition 3.3, the continuity of F_t and the definition of $\rho(t)$, we have P -a.s. that

$$\begin{aligned} \bar{\rho}(t) &:= \lambda\{i : \tau^i \leq t\} = \lambda\{i : Y_t^i + c\rho(t) \geq r\} \\ &= \int P\{Y_t^i + c\rho(t) \geq r\} \lambda(di) = 1 - F_t(r - c\rho(t)) = \rho(t) \end{aligned}$$

for all $t \geq 0$.

(c) Let $\bar{\rho} : \mathbb{R}_+ \rightarrow \mathbb{R}$ be a right-continuous function corresponding to an equilibrium; that is, $\bar{\rho}(t) = \lambda\{i : \tau^i \leq t\}$ for some optimal $\tau^i \in \mathcal{T}^i$, $i \in I$. Then $\bar{\rho}$ is clearly increasing and $[0, 1]$ -valued. Due to the strict increase of γ^i ,

we know from Lemma 2.1 that $\tau^i = \inf\{t : \gamma_t^i \geq r\}$. Thus,

$$\begin{aligned}\bar{\rho}(t) &= \lambda\{i : \tau^i \leq t\} = \lambda\{i : Y_t^i + c\bar{\rho}(t) \geq r\} \\ &= \int P\{Y_t^i + c\bar{\rho}(t) \geq r\} \lambda(di) = 1 - F_t(r - c\bar{\rho}(t));\end{aligned}$$

that is, $\bar{\rho}(t)$ is a solution of (4.1) for all $t \geq 0$. \square

We begin our discussion with some observations about uniqueness.

Remark 4.2. (i) Equation (4.1) may have more than one solution; cf. Example 4.4. If $t \mapsto \rho(t)$ is any right-continuous, increasing solution of (4.1), not necessarily maximal, then ρ induces an equilibrium, by the same arguments as in the above proof of Proposition 4.1.

(ii) Equation (4.1) also has a minimal solution, and it is automatically increasing in t . However, it is not necessarily right-continuous; see also Example 4.3 (iii) below. Instead, it is left-continuous provided that Y is. If the solution is unique and Y is continuous, the solution is both minimal and maximal, and therefore continuous.

Next, we analyze a special case of Proposition 4.1 that is explicitly solvable and sheds some light on the impact of the constant $c \geq 0$ that parametrizes the strength of interaction. One intuition is that if the interaction between the agents is too strong, some agents' stopping will lead to a domino effect where all others end up stopping immediately after.

Example 4.3. Let $r \geq 1$ and let U^i , $i \in I$ be essentially pairwise i.i.d. with a uniform distribution on $[r - 1, r]$. Moreover, let $a : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a strictly increasing, right-continuous function with $a(0) = 0$ and $a(\infty) > 1$ —the latter will ensure that (2.1) holds. We then consider the strictly increasing process

$$Y_t^i = U^i + a(t)$$

and note that $F_t(y) = F(1 + y - a(t) - r)$, where F is the c.d.f. of the uniform distribution on $[0, 1]$. Thus, Equation (4.1) becomes

$$1 - u = F(1 - cu - a(t)), \quad u \in [0, 1]. \quad (4.2)$$

- (i) *No interaction*, $c = 0$. Clearly, the unique solution is $\rho(t) = a(t) \wedge 1$, and this is the unique equilibrium by the last part of Proposition 4.1.

- (ii) *Moderate interaction*, $c \in (0, 1)$. Then, (4.2) is easily seen to have a unique solution $\rho(t) \in [0, 1]$; namely,

$$\rho(t) = [(1 - c)^{-1}a(t)] \wedge 1,$$

and this is the unique equilibrium. In particular, the population of stopped agents evolves in a nondegenerate, continuous fashion for $c \in (0, 1)$. The larger the interaction coefficient c , the more agents stop earlier.

- (iii) *Critical interaction*, $c = 1$. Using that $a(t) > 0$ for $t > 0$, we can check that $\rho(t) = 1$ is the unique solution of (4.2) for $t > 0$. Thus, $\rho \equiv 1$ is the unique right-continuous solution; that is, all agents stop at $t = 0$. This is also the unique right-continuous equilibrium.

It is worth noting that any $u \in [0, 1]$ is a solution of (4.2) for $t = 0$; recall that $a(0) = 0$. Intuitively speaking, solutions $\rho(t) = u\mathbf{1}_{\{0\}} + \mathbf{1}_{(0, \infty)}$ that are not right-continuous correspond to equilibria where a fraction u of the agents stop at time zero whereas the rest stop “immediately after” zero. This may illustrate why we have imposed right-continuity in our results.

- (iv) *Supercritical interaction*, $c > 1$. We see directly that $\rho(t) = 1$ is the unique solution of (4.2) for all $t \geq 0$.

4.1 On the Multiplicity of Equilibria

The following is a fairly well-behaved example of non-uniqueness.

Example 4.4. We consider again the setting of Example 4.3, with $r = 1$ and $a(t) = t$, say, but we now replace the uniform distribution of U^i with a measure that assigns mass $\varepsilon \in (0, 1/4)$ uniformly to $[0, \varepsilon]$ and to $[1 - \varepsilon, 1]$, and the remaining mass $1 - 2\varepsilon$ uniformly to $[1/2 - \varepsilon, 1/2 + \varepsilon]$. For small enough $\varepsilon > 0$ and $c \in (0, 1)$, we see that Equation (4.1) has three interior solutions for t within a certain interval, whereas all solutions are at the origin for $t = 0$. We can select any of these solutions to form an increasing right-continuous process ρ that corresponds to a legitimate equilibrium.

To understand this bifurcation, let us first look at an even simpler situation where $\gamma_t^i = r - c + c\rho(t)$ for all i . At time $t = 0$, two obvious equilibria are: No agent stops, then $\rho(t) = 0$ and $\gamma_t^i = r - c < r$, so it is indeed optimal not to stop. Or, all agents stop immediately, then $\rho(t) = 1$ and $\gamma_t^i = r$, so it is indeed optimal to stop. (This coordination problem is very similar to the

phenomenon discussed e.g. in [14, 27].) When γ_t^i is random, a similar choice can arise at an intermediate time for a subset of the population corresponding to an atom in the distribution of γ^i . More generally, the bifurcation can also happen in a continuous fashion when the random variable is sufficiently concentrated (relative to the size of c) around some point rather than having atoms, and this is what was witnessed in Example 4.4.

The following observation is a different view on the same interplay.

Remark 4.5. Let $(t, y) \mapsto F_t(y)$ be C^1 and write $f_t = \partial_y F_t$ for the probability density at time t . Suppose that $\partial_u F_t(r - cu) \neq -1$; that is,

$$cf_t(r - cu) \neq 1$$

for u in a neighborhood of a solution $\rho(t)$ of (4.1). Then, the Implicit Function Theorem shows that ρ is locally unique and C^1 . For $c > 0$, this is true, in particular, if $0 \leq f_t < c^{-1}$ on $[r - c, r]$. Or, put differently: if F_t is not too concentrated or if c is small enough, local uniqueness holds.

The following provides a broader perspective on Example 2.4 and shows that uniqueness may fail even more dramatically in certain regimes.

Example 4.6. We consider again the setting of Example 4.3, except that we now take

$$a(t) = \begin{cases} 0, & t < T \\ 2, & t \geq T, \end{cases}$$

where $T \in (0, \infty)$ acts as a time horizon. Indeed, this definition implies $\gamma_t^i \geq r + 1$ for all $t \geq T$ and $i \in I$, so that all agents will stop at T , if not earlier.

Thus, we are interested in the situation on $[0, T)$, where (4.2) becomes

$$1 - u = F(1 - cu).$$

(i) For $c \in [0, 1)$, the unique solution is $u = 0$, and thus

$$\rho(t) = \mathbf{1}_{[T, \infty)}$$

is the corresponding equilibrium: at $t = 0$, only a nullset of agents stop, and that does not change until T . One can check that this equilibrium is unique, even though γ^i is not strictly increasing.

(ii) At the critical value $c = 1$, uniqueness is lost and the situation is completely different. Indeed, the equation becomes the tautological $1 - u =$

$1 - u$. Thus, *any* right-continuous, increasing function $X(t)$ with values in $[0, 1]$ determines an equilibrium via

$$\rho(t) = X(t)\mathbf{1}_{[0,T)} + \mathbf{1}_{[T,\infty)}.$$

This is the situation we have encountered in Example 2.4: in terms of the equilibrium distribution, there is equivalence between assigning risk aversion $r - i$ to agent i and sampling uniformly from $[r - 1, r]$ for every agent as in the present example. The latter basically corresponds to randomly permuting the labels of the agents in Example 2.4.

5 The General Model

In this section, we generalize the model from the previous section by specifying the intensities γ^i as a possibly nonlinear function of i.i.d. signals Y^i and a common signal X . As a result, the intensities are conditionally independent rather than independent, and the equilibrium becomes a function of X .

As above, we consider the setup introduced in Section 2 with atomless probability spaces $(I, \mathcal{I}, \lambda)$ and (Ω, \mathcal{F}, P) , and let $(I \times \Omega, \Sigma, \mu)$ be a Fubini extension of their product. For each $i \in I$, let $Y^i \geq 0$ be a right-continuous, increasing, \mathbb{G}^i -progressively measurable process. We assume that for each $t \geq 0$, $(i, \omega) \mapsto Y_t^i(\omega)$ is Σ -measurable and that Y_t^i , $i \in I$ are essentially pairwise i.i.d. Moreover, we assume that the distribution of Y_t^i has no atoms; that is, its c.d.f. $y \mapsto F_t(y) := P\{Y_t^i \leq y\}$ is continuous. In addition, let X be a d -dimensional, right-continuous, (componentwise) increasing process which is \mathbb{G}^i -progressively measurable for all i and such that X_t and Y_t^i are independent for all $t \geq 0$. Thus, X is interpreted as public information whereas Y^i is a private signal⁷ available only to agent i . (A slightly different setup and interpretation are discussed in Section 5.1.)

Let $r : \mathbb{R}_+ \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a right-continuous, decreasing function. The interest rate process will be assumed to be of the form⁸

$$r_t = r(t, X_t).$$

Finally, the intensity of agent i will be of the form $\gamma_t^i = g(t, X_t, Y_t^i, \rho_t)$, where

$$g : \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$$

⁷ Additional idiosyncratic signals could also be included. In particular, a signal at time zero can be used to assign different functional forms of γ^i to the agents, similarly as at the end of Example 4.3.

⁸ Since X can be multivariate, this entails no loss of generality relative to introducing yet another stochastic process.

is a continuous function with $g(t, X_t, Y_t^i, 0) \geq 0$, increasing in all its arguments and such that for all (t, x, u) , $y \mapsto g(t, x, y, u)$ admits an inverse $y' \mapsto g^{-1}(t, x, y', u)$ on its range which we assume to be \mathbb{R} for simplicity. We suppose that g^{-1} is again continuous.

Theorem 5.1. *The equation*

$$1 - u = F_t(g^{-1}(t, x, r, u)), \quad u \in [0, 1] \quad (5.1)$$

has a maximal solution $\rho(t, x, r) \in [0, 1]$ for every $(t, x, r) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R}$, and $\rho_t := \rho(t, X_t, r(t, X_t))$ is a right-continuous, increasing process. Define also

$$\gamma_t^i = g(t, X_t, Y_t^i, \rho_t), \quad \tau^i = \inf\{t : \gamma_t^i = r_t\}$$

and assume that (2.1) is satisfied for all i .

(i) Then, ρ and $(\tau^i)_{i \in I}$ define an equilibrium: $\tau^i \in \mathcal{T}^i$ is an optimal stopping time for agent i , the mapping $(i, \omega) \mapsto \tau^i(\omega)$ is Σ -measurable, and

$$\lambda\{i : \tau^i \leq t\} = \rho_t \quad P\text{-a.s. for all } t \geq 0.$$

More generally, this holds for any measurable solution $\rho(t, x, r)$ of (5.1) such that $\rho(t, X_t, r(t, X_t))$ is right-continuous and increasing.

(ii) Conversely, let $t \mapsto \bar{\rho}_t$ be a right-continuous process corresponding to an equilibrium and suppose that $\bar{\rho}_t = \bar{\rho}(t, X_t, r(t, X_t))$ for some measurable function $\bar{\rho}$. If γ^i is strictly increasing for all i , then for every $t \geq 0$, $\bar{\rho}(t, x, r(t, x))$ solves (5.1) for $(P \circ X_t^{-1})$ -almost all $x \in \mathbb{R}^d$.

Proof. (a) We claim that $\rho(t, x, r)$ is well-defined, increasing in (t, x) and decreasing in r , and (jointly) right-continuous in (t, x) and left-continuous in r . Indeed, fix (t, x, r) and consider the function

$$G_{t,x,r}(u) := F_t(g^{-1}(t, x, r, u)) - 1 + u, \quad u \in [0, 1].$$

Since F_t takes values in $[0, 1]$, we have $G_{t,x,r}(0) \leq 0$ and $G_{t,x,r}(1) \geq 0$. As $u \mapsto G_{t,x,r}(u)$ is continuous, it follows that there is least one zero in $[0, 1]$, and since the set of all zeros is compact, it has a maximum.

We write $\rho(t, x, r)$ for the maximal zero of $G_{t,x,r}$. As Y is increasing, the function $t \mapsto F_t(y)$ is decreasing and then so is $t \mapsto G_{t,x,r}(u)$; note that g^{-1} is decreasing in (t, x, u) and increasing in r . Hence, if $s \leq t$, the fact that $G_{t,x,r} > 0$ on $(\rho(t, x, r), 1]$ implies that $G_{s,x,r} > 0$ on $(\rho(t, x, r), 1]$ and hence that $\rho(s, x, r) \leq \rho(t, x, r)$. The monotonicity in x and r follows analogously.

Let $t_n \downarrow t$ and $x_n \downarrow x$ and $r_n \uparrow r$. Set $\rho_n = \rho(t_n, x_n, r_n)$ and $\rho_* = \rho(t, x, r)$. By the above, ρ_n is decreasing and $\rho_n \geq \rho_*$. Thus, we only need to verify

that $\rho_\infty := \lim \rho_n \leq \rho_*$. In view of the definition of ρ_* as a maximal zero, it suffices to show that ρ_∞ is a zero of $G_{t,x,r}$, and as $G_{t_n,x_n,r_n}(\rho_n) = 0$, that will follow if

$$(t, x, -r, u) \mapsto G_{t,x,r}(u) \quad \text{is jointly right-continuous.} \quad (5.2)$$

Indeed, $y \mapsto F_t(y)$ is continuous, and together with the right-continuity of Y , it follows that $t \mapsto F_t(y)$ is right-continuous. Using also the continuity of g^{-1} , we see that (5.2) holds as desired. This completes the proof of the claim on ρ .

(b) Next, we verify the equilibrium conditions. As a result of (a), the processes $t \mapsto \rho_t = \rho(t, X_t, r(t, X_t))$ and γ^i are increasing and right-continuous, and Lemma 2.1 yields that $\tau^i \in \mathcal{T}^i$ is an optimal stopping time. Note that $(i, \omega) \mapsto \gamma^i(\omega)$ is Σ -measurable, and so is $(i, \omega) \mapsto r_t(\omega)$. Thus, $\{\tau^i \leq t\} = \{\gamma^i \geq r_t\} \in \Sigma$ for all $t \geq 0$. Using the Conditional Exact Law of Large Numbers of Proposition 3.4, the continuity of $y \mapsto F_t(y)$ and the definition of ρ_t , we have P -a.s. that

$$\begin{aligned} \bar{\rho}_t &:= \lambda\{i : \tau^i \leq t\} = \lambda\{i : g(t, X_t, Y_t^i, \rho(t, X_t, r(t, X_t))) \geq r(t, X_t)\} \\ &= \int P\{g(t, X_t, Y_t^i, \rho(t, X_t, r(t, X_t))) \geq r(t, X_t) | X_t\} \lambda(di) \\ &= 1 - F_t(g^{-1}(t, X_t, r(t, X_t), \rho(t, X_t, r(t, X_t)))) \\ &= \rho(t, X_t, r(t, X_t)) = \rho_t. \end{aligned}$$

(c) Let $\bar{\rho}$ be a right-continuous process corresponding to an equilibrium; that is, $\bar{\rho}_t = \lambda\{i : \tau^i \leq t\}$ for some optimal $\tau^i \in \mathcal{T}^i$. Then $\bar{\rho}$ is clearly increasing and $[0, 1]$ -valued. Due to the strict increase of γ^i , we know from Lemma 2.1 that $\tau^i = \inf\{t : \gamma_t^i \geq r_t\}$, which also ensures that $\{\tau^i \leq t\} \in \Sigma$. Since we have assumed that $\bar{\rho}_t = \bar{\rho}(t, X_t, r(t, X_t))$, we obtain as in (b) that

$$\bar{\rho}_t = \lambda\{i : \tau^i \leq t\} = 1 - F_t(g^{-1}(t, X_t, r(t, X_t), \bar{\rho}(t, X_t, r(t, X_t)))) \quad P\text{-a.s.}$$

for all $t \geq 0$. □

Remark 5.2. The result in Theorem 5.1 (ii) assumes a priori that the equilibrium $\bar{\rho}_t$ is Markovian; that is, a deterministic function of (t, X_t) .

(i) First, let us observe that this is not automatically the case: randomized equilibria may exist. Consider the setting of Example 4.4 where X is deterministic and Equation (5.1) has several (deterministic, increasing, right-continuous) solutions; in particular, a maximal solution $\rho(t)$ and a minimal

solution $\rho'(t)$. Suppose that there is a Poisson process N which is \mathbb{G}^i -adapted and independent of Y^i for all i , and let σ be its first jump time. Then,

$$\bar{\rho}_t = \mathbf{1}_{[0,\sigma)}(t)\rho'(t) + \mathbf{1}_{[\sigma,\infty)}(t)\rho(t)$$

defines another right-continuous, increasing solution of (5.1) which determines an equilibrium. However, $\bar{\rho}_t$ is not of the mentioned Markovian form. Instead, the agents can agree to change their behavior according to the independent randomization σ . We also refer to [10, 22] for further insights on randomized (or “weak”) equilibria in the context of standard mean field games.

(ii) Second, let us show that the phenomenon mentioned in (i) cannot occur if uniqueness holds in Equation (5.1). In the setting of Theorem 5.1 (ii), even if we do not suppose a priori that $\bar{\rho}$ is a function of (t, X_t) , we have

$$\begin{aligned} \bar{\rho}_t &= \lambda\{i : \tau^i \leq t\} = \lambda\{i : g(t, X_t, Y_t^i, \bar{\rho}_t) \geq r(t, X_t)\} \\ &= \int P\{g(t, X_t, Y_t^i, \bar{\rho}_t) \geq r(t, X_t) | X_t\} \lambda(di) \\ &= 1 - F_t(g^{-1}(t, X_t, r(t, X_t), \bar{\rho}_t)). \end{aligned}$$

If $\rho(t, x, r)$ is the maximal solution of (5.1) as constructed in the theorem and uniqueness holds for (5.1), it follows that

$$\bar{\rho}_t = \rho(t, X_t, r(t, X_t)) \quad P\text{-a.s.}$$

and in that sense, $\bar{\rho}$ is necessarily of the Markovian form.

Remark 5.3. In the literature on mean field games driven by stochastic differential equations, the private states at time t are usually independent conditionally on the whole path $(X_s)_{s \leq t}$ of the common noise before time t . In the present setting, we have assumed that the intensities γ^i depend on X in a Markovian way, and hence it is sufficient to condition on the current value X_t . One could envision a similar result where γ^i depends on X in a path-dependent way, and then one would condition on the whole past of X .

5.1 Additive Model and Noisy Observations

In this section, we enhance the toy model from Proposition 4.1 by incorporating a public signal and obtain a tractable specification of the general model from Theorem 5.1. Consider the setup introduced in Section 2 with atomless probability spaces $(I, \mathcal{I}, \lambda)$ and (Ω, \mathcal{F}, P) , and let $(I \times \Omega, \Sigma, \mu)$ be

a Fubini extension of their product. For each $i \in I$, let $Y^i \geq 0$ be a right-continuous, increasing, measurable process. We assume that for each $t \geq 0$, $(i, \omega) \mapsto Y_t^i(\omega)$ is Σ -measurable and that Y_t^i , $i \in I$ are essentially pairwise independent. Moreover, we assume that the distribution of $Y_t^i(\cdot)$ has no atoms; that is, its c.d.f. F_t is continuous. Furthermore, let $X \geq 0$ be a right-continuous, increasing, measurable process such that X_t and Y_t^i are independent for all $t \geq 0$. We take $r \in \mathbb{R}$ to be constant (for simplicity) and

$$\gamma_t^i = X_t + Y_t^i + c\rho_t,$$

where $c \geq 0$ is a constant governing the strength of interaction; see also Example 4.3. For the information structure, we may consider two cases. Either we see \mathbb{G}^i as given and assume that

- X and Y^i are \mathbb{G}^i -progressively measurable for all $i \in I$,

which was the point of view taken above. Or, we model that the agents observe only $X + Y^i$ and ρ , and thus we convene that

- \mathbb{G}^i is the right-continuous filtration generated by $X + Y^i$ and ρ , for all $i \in I$.

This allows for the interpretation of X as a “true signal,” whereas agent i can only observe the noisy signal $X + Y^i$ with i.i.d. noise Y^i . Although the agents have more information in the first setting, both yield the same equilibria—the form of τ^i stated below shows that the agents only use the observation of γ^i . Indeed, Theorem 5.1 yields the following.

Corollary 5.4. *The equation*

$$1 - u = F_t(r - x - cu), \quad u \in [0, 1]$$

has a maximal solution $\rho(t, x) \in [0, 1]$ for every $(t, x) \in \mathbb{R}_+ \times \mathbb{R}$, and $\rho_t := \rho(t, X_t)$ is a right-continuous process. Define also

$$\gamma_t^i = X_t + Y_t^i + c\rho_t, \quad \tau^i = \inf\{t : X_t + Y_t^i + c\rho_t = r\}$$

and assume that (2.1) is satisfied for all i .

(i) Then, ρ and $(\tau^i)_{i \in I}$ define an equilibrium: $\tau^i \in \mathcal{T}^i$ is an optimal stopping time for agent i , the mapping $(i, \omega) \mapsto \tau^i(\omega)$ is Σ -measurable, and

$$\lambda\{i : \tau^i \leq t\} = \rho_t \quad P\text{-a.s.} \quad \text{for all } t \geq 0.$$

(ii) Conversely, let $t \mapsto \bar{\rho}_t$ be a right-continuous process corresponding to an equilibrium and suppose that $\bar{\rho}_t = \bar{\rho}(t, X_t)$ for some measurable function $\bar{\rho}$. If γ^i is strictly increasing for all i , then for every $t \geq 0$, $\bar{\rho}(t, x)$ solves (5.1) for $(P \circ X_t^{-1})$ -almost all $x \in \mathbb{R}$.

A solvable example can be constructed along the lines of Example 4.3.

Example 5.5. Let $r \geq 1$ and let U^i , $i \in I$ be essentially pairwise i.i.d. with a uniform distribution on $[r-1, r]$ and such that U^i and X_t are independent for all $t \geq 0$. Moreover, suppose that X is strictly increasing with $X_0 = 0$ and $X_\infty > 1$. For $c \in (0, 1)$, consider the intensity process

$$\gamma_t^i = X_t + U^i + c\rho_t.$$

Then, the equation has a unique solution $\rho(t, x)$, and

$$\rho(t, X_t) = [(1 - c)^{-1} X_t] \wedge 1$$

corresponds to the unique (Markovian) equilibrium. In particular, this equilibrium evolves in a nondegenerate way as long as X does.

References

- [1] R. J. Aumann. Markets with a continuum of traders. *Econometrica*, 32:39–50, 1964.
- [2] M. Bardi. Explicit solutions of some linear-quadratic mean field games. *Netw. Heterog. Media*, 7(2):243–261, 2012.
- [3] M. Bardi and F. S. Priuli. Linear-quadratic N -person and mean-field games with ergodic cost. *SIAM J. Control Optim.*, 52(5):3022–3052, 2014.
- [4] A. Bensoussan, J. Frehse, and S. C. P. Yam. *Mean field games and mean field type control theory*. Springer Briefs in Mathematics. Springer, New York, 2013.
- [5] A. Bensoussan, K. C. J. Sung, S. C. P. Yam, and S. P. Yung. Linear-Quadratic Mean Field Games. *J. Optim. Theory Appl.*, 169(2):496–529, 2016.
- [6] P. Cardaliaguet. Notes on mean field games (from P.-L. Lions’ lectures at Collège de France). 2013.
- [7] R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SIAM J. Control Optim.*, 51(4):2705–2734, 2013.
- [8] R. Carmona and F. Delarue. The master equation for large population equilibria. In *Stochastic analysis and applications 2014*, volume 100 of *Springer Proc. Math. Stat.*, pages 77–128. Springer, Cham, 2014.
- [9] R. Carmona and F. Delarue. Forward-backward stochastic differential equations and controlled McKean-Vlasov dynamics. *Ann. Probab.*, 43(5):2647–2700, 2015.

- [10] R. Carmona, F. Delarue, and D. Lacker. Mean field games with common noise. *Ann. Probab.*, 44(6):3740–3803, 2015.
- [11] R. Carmona, F. Delarue, and D. Lacker. Mean field games of timing and models for bank runs. *Appl. Math. Optim.*, 76(1):217–260, 2017.
- [12] R. Carmona, J.-P. Fouque, and L.-H Sun. Mean field games and systemic risk. *Commun. Math. Sci.*, 13(4):911–933, 2015.
- [13] R. Carmona and D. Lacker. A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.*, 25(3):1189–1231, 2015.
- [14] D. W. Diamond and P H. Dybvig. Bank runs, deposit insurance, and liquidity. *J. Polit. Econ.*, 91(3):401–419, 1983.
- [15] D. Duffie and Y. Sun. The exact law of large numbers for independent random matching. *J. Econom. Theory*, 147(3):1105–1139, 2012.
- [16] M. Fischer. On the connection between symmetric N -player games and mean field games. *Ann. Appl. Probab.*, 27(2):757–810, 2017.
- [17] D. A. Gomes and J. Saúde. Mean field games models—a brief survey. *Dyn. Games Appl.*, 4(2):110–154, 2014.
- [18] O. Guéant, J.-M. Lasry, and P.-L. Lions. Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Math.*, pages 205–266. Springer, Berlin, 2011.
- [19] M. Huang, P. E. Caines, and R. P. Malhamé. Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior and decentralized ϵ -Nash equilibria. *IEEE Trans. Automat. Control*, 52(9):1560–1571, 2007.
- [20] M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.*, 6(3):221–251, 2006.
- [21] I. Karatzas, M. Shubik, and W. D. Sudderth. Construction of stationary Markov equilibria in a strategic market game. *Math. Oper. Res.*, 19(4):975–1006, 1994.
- [22] D. Lacker. A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields*, 165(3-4):581–648, 2016.
- [23] D. Lando. *Credit Risk Modeling: Theory and Applications*. Princeton University Press, Princeton, 2009.
- [24] J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. I. Le cas stationnaire. *C. R. Math. Acad. Sci. Paris*, 343(9):619–625, 2006.
- [25] J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. II. Horizon fini et contrôle optimal. *C. R. Math. Acad. Sci. Paris*, 343(10):679–684, 2006.

- [26] J.-M. Lasry and P.-L. Lions. Mean field games. *Jpn. J. Math.*, 2(1):229–260, 2007.
- [27] S. Morris and H. S. Shin. Coordination risk and the price of debt. *Eur. Econ. Rev*, 48(1):133–153, 2004.
- [28] H. Pham. Linear quadratic optimal control of conditional McKean-Vlasov equation with random coefficients and applications. *Preprint arXiv:1604.06609v1*, 2016.
- [29] K. Podczeck. On existence of rich Fubini extensions. *Econom. Theory*, 45(1-2):1–22, 2010.
- [30] L. Qiao, Y. Sun, and Z. Zhang. Conditional exact law of large numbers and asymmetric information economies with aggregate uncertainty. *Econom. Theory*, pages 1–22, 2014.
- [31] Y. Sun. The exact law of large numbers via Fubini extension and characterization of insurable risks. *J. Econom. Theory*, 126(1):31–69, 2006.
- [32] Y. Sun and Y. Zhang. Individual risk and Lebesgue extension without aggregate uncertainty. *J. Econom. Theory*, 144(1):432–443, 2009.