



Mean-field ranking games with diffusion control

S. Ankirchner¹ · N. Kazi-Tani² · J. Wendt¹ · C. Zhou³

Received: 30 June 2023 / Accepted: 9 January 2024
© The Author(s) 2024

Abstract

We consider a stochastic differential game, where each player continuously controls the diffusion intensity of her own state process. The players must all choose from the same diffusion rate interval $[\sigma_1, \sigma_2]$, and have individual random time horizons that are independently drawn from the same distribution. The players whose states at their respective time horizons are among the best $p \in (0, 1)$ of all terminal states receive a fixed prize. We show that in the mean field version of the game there exists an equilibrium, where the representative player chooses the maximal diffusion rate when the state is below a given threshold, and the minimal rate else. The symmetric n -fold tuple of this threshold strategy is an approximate Nash equilibrium of the n -player game. Finally, we show that the more time a player has at her disposal, the higher her chances of winning.

Keywords Diffusion control · Game · Rank-based reward · Mean field limit · Oscillating Brownian motion

Mathematics Subject Classification Primary: 91A15; secondary: 91A06 · 91A10 · 91A16 · 93E20

✉ S. Ankirchner
s.ankirchner@uni-jena.de

N. Kazi-Tani
nabil.kazi-tani@univ-lorraine.fr

J. Wendt
julian.wendt@uni-jena.de

C. Zhou
matzc@nus.edu.sg

¹ Institute for Mathematics, University of Jena, Ernst-Abbe-Platz 2, 07743 Jena, Germany

² Institut Elie Cartan de Lorraine, Université de Lorraine, UFR MIM, 3 rue Augustin Fresnel, 57073 Metz Cedex 03, France

³ Department of Mathematics and Risk Management Institute, National University of Singapore, 10 Lower Kent Ridge Road, 119076 Singapore, Singapore

1 Introduction

We consider the following stochastic differential game: each player can control the fluctuation intensity of her own state process, up to an individual random time horizon. The controls of a player are individually defined as a set of progressively measurable processes, with respect to a filtration modeling the player's information flow, with values in a bounded interval $[\sigma_1, \sigma_2]$, where $0 < \sigma_1 < \sigma_2$ are the same for all. The players whose terminal states are among the highest $p \in (0, 1)$ receive a fixed prize, set to be equal to one. The other players do not receive anything.

The game models in stylized form competitions where only the best performing agents receive a fixed reward and where every agent can choose between risky and safe actions. The game thus allows to analyze the impact of rank-based rewards on the risk appetite of the competitors. The game has multiple interpretations. We refer to Sect. 6 for more details.

As usual, we fall back on the concept of Nash equilibria for predicting the players' behavior. Given the discontinuous rewards, it turns out to be difficult to compute explicit equilibria. Moreover, it seems already difficult to prove existence by abstract means. A way out for games with many players is to fall back on the game's mean field version and to derive an approximate equilibrium. The equilibrium of the mean field game here consists of a control and the respective state distribution at the terminal time. Under regularity conditions on the rewards and on the state equation coefficients, existence of equilibria in mean field games with diffusion control and common noise is proved in [2] using a relaxed formulation of the problem and the theory of second order BSDEs. Other contributions consider mean field games with diffusion control, such as [8] in a Principal-Agent setting, with applications to optimal energy demand management or [6, 7] in the case of extended mean field games with control interactions.

In this paper, we avoid second order BSDEs and show, using a direct argument, that there exists an equilibrium with a threshold control that consists in choosing the maximal diffusion rate σ_2 when the state is below a given threshold, and the minimal rate σ_1 else. The corresponding equilibrium distribution is the distribution of an oscillating Brownian motion at the terminal time.

The game bears similarities with the diffusion control game studied in [1]. In contrast to [1], however, we allow here the agents to be heterogeneously informed. Each player is assumed to observe her own state process, but the assumption on how much the agents know beyond this remains general. Some of the players may know, e.g., the time horizons, and some not. Whether a player knows her own time horizon or the time horizons of the opponents turns out to have no effect on the equilibrium. Indeed, in the many player game the empirical distribution of the time horizons is close to \mathcal{T} . Thus knowing \mathcal{T} is sufficient for implementing an approximate equilibrium control. Similarly, it has no effect on the equilibrium whether the agents can observe the state processes of the opponents or not. For implementing the threshold control of the equilibrium each agent needs only to observe her own process.

By defining the state dynamics in terms of solutions to controlled martingale problems and choosing controls of open loop type, we obtain a model framework that allows to cover general information structures, in particular situations where a player has only partial knowledge about the other players' states and about the individual random time horizons. In the game of [1] the state dynamics are described in terms of stochastic differential equations and the players' controls are modeled as closed loop controls.

In the appropriate mean field version of the game the representative player knows the distribution \mathcal{T} of the time horizon, but the actual time horizon arrives unpredicted. The

threshold describing the equilibrium control strongly depends on the distribution \mathcal{T} . For the cases where \mathcal{T} is an exponential distribution or a uniform distribution, we characterize the threshold of the equilibrium control as the unique root of a simple equation.

The article is organized as follows: in Sect. 2, we introduce the game model in more detail. For deriving a candidate for an approximate Nash equilibrium, we study the corresponding mean field game in Sect. 3 and show that an equilibrium control is given by a threshold control. In Sect. 4, we show that the n -tuple consisting of the mean field equilibrium controls is an $\mathcal{O}(n^{-1/2})$ -Nash equilibrium of the n -player game. This means, in the approximate Nash equilibrium players only use information about their own state and choose maximal diffusion intensity below the optimal threshold and minimal diffusion intensity above. All additional information about the opponents' states or the random times is irrelevant. In Sect. 5, we analyze the winning probability of a given player depending on the time horizon T and in Sect. 6, we provide a particular application of our model setting to online competitions. Finally, we consider in Sect. 7 an extension of the game to more general reward functions that are continuous, have exponential growth, satisfy a symmetry condition, and are convex below a certain threshold and concave above. We show that the same tuple consisting of the mean field equilibrium controls is also an approximate Nash equilibrium for these reward functions.

2 Game model

We describe the players' states by means of controlled martingale problems. To this end, let $\mathbb{R}_+ := [0, \infty)$, and $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n)$ denote the space of continuous functions $f : \mathbb{R}_+ \rightarrow \mathbb{R}^n$ equipped with the metric

$$d(\omega_1, \omega_2) := \sum_{n=1}^{\infty} \frac{1}{2^n} \frac{\sup_{t \in [0, n]} |\omega_1(t) - \omega_2(t)|}{1 + \sup_{t \in [0, n]} |\omega_1(t) - \omega_2(t)|}, \quad \omega_1, \omega_2 \in \mathcal{C}(\mathbb{R}_+, \mathbb{R}^n).$$

Let $\mathcal{B}(\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n))$ denote the corresponding Borel σ -algebra on $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n)$ and denote by $X = (X^1, \dots, X^n)$ the canonical process on $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n)$, i.e., $X_t^i(\omega) := \omega^i(t)$, $t \geq 0$, for $i = 1, \dots, n$ and $\omega \in \mathcal{C}(\mathbb{R}_+, \mathbb{R}^n)$. We refer to [18], Section 1.3, for more details on the construction of this measurable space and its properties.

Let \mathcal{T} be a probability measure on \mathbb{R}_+ , equipped with the usual Borel σ -algebra $\mathcal{B}(\mathbb{R}_+)$, that describes the distribution of the players' random times. We suppose that \mathcal{T} satisfies:

Assumption 2.1

$$\int_0^{\infty} \frac{1}{\sqrt{t}} \mathcal{T}(dt) < \infty.$$

Let $\mathcal{B}(\mathbb{R}_+^n)$ be the Borel σ -algebra on \mathbb{R}_+^n and define the measure $\mathcal{T}^n = \otimes_{i=1}^n \mathcal{T}$ on \mathbb{R}_+^n , i.e., \mathcal{T}^n is the n -fold product of \mathcal{T} . We define $\tau = (\tau_1, \dots, \tau_n)$ as canonical map on \mathbb{R}_+^n , i.e., $\tau(\omega) = (\tau_1(\omega), \dots, \tau_n(\omega)) = (\omega_1, \dots, \omega_n)$ for any $\omega \in \mathbb{R}_+^n$. Note that τ_1, \dots, τ_n are independent and identically distributed under \mathcal{T}^n by definition and the law of τ_i is given by \mathcal{T} .

We define a common measurable space by setting:

- (i) $\Omega := \mathcal{C}(\mathbb{R}_+, \mathbb{R}^n) \times \mathbb{R}_+^n$,
- (ii) $\mathcal{F} := \mathcal{B}(\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n)) \otimes \mathcal{B}(\mathbb{R}_+^n)$, i.e., the smallest σ -algebra containing all sets of the form $A \times B$ for $A \in \mathcal{B}(\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n))$, $B \in \mathcal{B}(\mathbb{R}_+^n)$.

We extend the definitions of X and τ to Ω by setting $X(\omega_1, \omega_2) := X(\omega_1)$ and $\tau(\omega_1, \omega_2) := \tau(\omega_2)$ for $(\omega_1, \omega_2) \in \Omega$. We define on (Ω, \mathcal{F}) the filtration $(\mathcal{F}_t)_{t \geq 0}$ by

$$\mathcal{F}_t = \sigma(\tau_1, \dots, \tau_n) \vee \sigma(X_s : 0 \leq s \leq t), \quad t \geq 0.$$

The filtration $(\mathcal{F}_t)_{t \geq 0}$ describes the overall information flow including the information about the values of the time horizons. Moreover, for Player i , we introduce the filtration $(\mathcal{F}_t^i)_{t \geq 0}$ describing the private information of Player i . We assume that $\sigma(X_s^i : 0 \leq s \leq t) \subseteq \mathcal{F}_t^i \subseteq \mathcal{F}_t$ for all $t \geq 0$.

Let \mathcal{A}^i denote the set of all $(\mathcal{F}_t^i)_{t \geq 0}$ -progressively measurable $\alpha : \Omega \times \mathbb{R}_+ \rightarrow [\sigma_1, \sigma_2]$. We refer to elements of \mathcal{A}^i as *admissible controls* or *strategies* and write \mathcal{A}_n for the product $\mathcal{A}^1 \times \dots \times \mathcal{A}^n$. We characterize the law of the state processes by means of martingale problems, introduced by Stroock and Varadhan. We refer to the monograph [18] for more details on martingale problems.

Definition 2.2 Let $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathcal{A}_n$. Then, a probability measure P^α on (Ω, \mathcal{F}) is called a *feasible state distribution* if

- (i) $P^\alpha \circ X_0^{-1} = \delta_0$,
- (ii) $P^\alpha(\mathcal{C}(\mathbb{R}_+, \mathbb{R}^n) \times B) = \mathcal{T}^n(B)$ for all $B \in \mathcal{B}(\mathbb{R}_+^n)$,
- (iii) for all $f \in \mathcal{C}_c^2(\mathbb{R}^n, \mathbb{R})$, i.e., for all twice continuously differentiable $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with compact support, the process M^f , defined by

$$M_s^f := f(X_s) - f(0) - \frac{1}{2} \int_0^s \sum_{j=1}^n (\alpha_r^j)^2 \partial_{jj} f(X_r) dr, \quad s \geq 0, \tag{1}$$

is an $(\mathcal{F}_t)_{t \geq 0}$ -martingale under P^α .

We denote by $\mathcal{Q}(\alpha)$ the set of all feasible state distributions P^α .

Remark 2.3 The assumptions on the tuple α in the previous definition do not exclude that $\mathcal{Q}(\alpha)$ is empty. For a tuple α to be a Nash equilibrium it is necessary, however, that $\mathcal{Q}(\alpha) \neq \emptyset$ (see Definition 2.8 below).

Remark 2.4 Note that condition (ii) implies that each random time τ_i has distribution \mathcal{T} . Moreover, condition (iii) yields that each state process X^i is a local $(\mathcal{F}_t)_{t \geq 0}$ -martingale and

$$\langle X^i, X^j \rangle_t = \begin{cases} \int_0^t (\alpha_s^i)^2 ds, & \text{if } i = j, \\ 0, & \text{else,} \end{cases} \quad t \geq 0.$$

The state processes are even true martingales that are square integrable, i.e., $E^{P^\alpha}[(X_t^i)^2] < \infty, t \geq 0$, because the controls are bounded and thus, $E^{P^\alpha}[\langle X^i, X^i \rangle_t] < \infty$ for all $t \geq 0$ (see, e.g., [17], Section II.6, Corollary 3, p.73).

Remark 2.5 The definition of the filtrations $(\mathcal{F}_t^i)_{t \geq 0}$ is rather general. Each filtration $(\mathcal{F}_t^i)_{t \geq 0}$ can contain information about the other players' states and random times. Therefore, the game setting covers the cases where players can or cannot observe each other, and have knowledge about the random time horizons.

If, e.g., $\mathcal{F}_t^i = \mathcal{F}_t$, each player can observe the state processes of the opponents and make her strategy depend on the opponents' state trajectories. Moreover, each player has prior knowledge of the random time horizons. If, however, $\mathcal{F}_t^i = \sigma(X_s^i : 0 \leq s \leq t)$, then each player can only observe her own state process. Neither information about the opponents' states nor about the random times is available.

Despite this quite general information structure, we show in Theorem 4.1 that an approximate Nash equilibrium is given by a tuple of threshold controls depending only on the position of the single player’s state process.

Remark 2.6 The state processes can be equivalently described as weak solutions to an n -dimensional SDE (or via a stochastic integral w.r.t. some Brownian motion). Indeed, for some control $\alpha \in \mathcal{A}_n$ with $\mathcal{Q}(\alpha) \neq \emptyset$ and $P \in \mathcal{Q}(\alpha)$, there exists an n -dimensional Brownian motion W on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$ such that the state processes X^1, \dots, X^n satisfy P -a.s.

$$X_t^i = \int_0^t \alpha_s^i dW_s^i, \quad t \geq 0, \quad i = 1, \dots, n, \tag{2}$$

and $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P, X, W)$ is a weak solution to (2) (see, e.g., [11], Proposition 5.4.6). We use this connection in the mean field game presented in Sect. 3, and thus, characterize the state processes via stochastic integrals.

Remark 2.7 For all measurable feedback functions $a : \mathbb{R}_+ \times \mathbb{R}^n \rightarrow [\sigma_1, \sigma_2]^n$, there exists a solution P^a to the martingale problem (1) with $(\alpha_s)_{s \geq 0} = (a(s, X_s^1, \dots, X_s^n))_{s \geq 0}$. This follows, e.g., from [13], Theorem 2.6.1, and [11], Proposition 5.4.11. If the filtration $(\mathcal{F}_t^i)_{t \geq 0}$ contains the information about all states, i.e. if $\sigma(X_s : 0 \leq s \leq t) \subset \mathcal{F}_t^i$, then the control $(a(s, X_s^1, \dots, X_s^n))_{s \geq 0}$ is contained in \mathcal{A}^i and $\{P^a\} \subseteq \mathcal{Q}(\alpha)$. This particularly holds for control tuples where each entry is a threshold control, i.e., each entry is given by the feedback function

$$m_b(x) = \begin{cases} \sigma_2, & \text{if } x \leq b, \\ \sigma_1, & \text{if } x > b, \end{cases} \tag{3}$$

where $b \in \mathbb{R}$. In this case, the solution to the martingale problem (1) is even unique: Remark 3.1 below implies that the SDE (7) with $m = m_b$ has a unique strong solution and [11], Corollary 5.4.9, then implies that uniqueness for the martingale problem (1) holds.

We suppose that each player aims at maximizing the probability of her own state at her terminal random time to be greater than the empirical $(1 - p)$ -quantile of all states at the individual random times. More precisely, let

$$\mu^n = \frac{1}{n} \sum_{i=1}^n \delta_{X_{\tau_i}^i}$$

be the empirical distribution of the players’ states at the terminal times. We define the empirical $(1 - p)$ -quantile by $q(\mu^n, 1 - p) = \inf\{r \in \mathbb{R} : \mu^n((-\infty, r]) \geq 1 - p\}$. Note that $X_{\tau_i}^i > q(\mu^n, 1 - p)$ if and only if the state of Player i is among the best $[np]$ players at the terminal times.

It is standard to predict or explain the players’ behavior in terms of (approximate) Nash equilibria, which are here defined as follows.

Definition 2.8 Let $\varepsilon \geq 0$. A tuple $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathcal{A}_n$ with $\mathcal{Q}(\alpha) \neq \emptyset$ is called ε -Nash equilibrium of the n -player game if for all $i \in \{1, \dots, n\}$ and $P^\alpha \in \mathcal{Q}(\alpha)$, we have

$$P^\alpha(X_{\tau_i}^i > q(\mu^n, 1 - p)) + \varepsilon \geq \sup_{\beta \in \mathcal{A}^i} \sup_{P \in \mathcal{Q}(\alpha_{-i}, \beta)} P(X_{\tau_i}^i > q(\mu^n, 1 - p)), \tag{4}$$

where $(\alpha_{-i}, \beta) = (\alpha_1, \dots, \alpha_{i-1}, \beta, \alpha_{i+1}, \dots, \alpha_n)$ and $\sup \emptyset = -\infty$.

Note that for $\varepsilon = 0$, the tuple α in Definition 2.8 is a Nash equilibrium in the usual sense. In the case $\varepsilon > 0$, the tuple α is also called an *approximate* Nash equilibrium.

We do not assume that the solutions to the martingale problem (1) are unique by limiting the set of controls. Hence, we require that (4) holds for all solutions to the martingale problem (1). We emphasize that for an arbitrary control tuple $\alpha \in \mathcal{A}_n$, uniqueness of solutions to the martingale problem (1) can fail. For example, if $n \geq 3$, then it was shown in [15] that there exists a diffusion coefficient, and hence, an operator such that the corresponding martingale problem does not have a unique solution. Equivalently, there exists a diffusion coefficient $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ that is uniformly elliptic and such that for the corresponding SDE uniqueness in law fails (see also [5], Example 1.24). Thus, one can interpret (4) as follows: Player i has no incentive to change her strategy from α to β , no matter which distribution from $\mathcal{Q}(\alpha_{-i}, \beta)$ is chosen. We stress, however, that for the approximate Nash equilibria derived from the corresponding mean field game, uniqueness is always satisfied.

We do not compute an exact Nash equilibrium for the n -player game. As in [1], we compute an approximate Nash equilibrium for large games by considering the mean field limit of the game. We show that a mean field equilibrium strategy is given by a threshold control.

Remark 2.9 Let $T > 0$ and $\tau_1 = \dots = \tau_n = T$, i.e., $\mathcal{T} = \delta_T$. Set $\mathcal{F}_t^i = \sigma(X_s : 0 \leq s \leq t)$, $t \geq 0$, $i = 1, \dots, n$, and consider only controls of the form $\alpha_t = a(t, X_t^1, \dots, X_t^n)$ for some $a : \mathbb{R}_+ \times \mathbb{R}^n \rightarrow [\sigma_1, \sigma_2]$. Then, the above model is equivalent to the game model presented in the article [1].

3 Mean field game

In this section we describe the mean field version of the game introduced in Sect. 2.

Let $0 < \sigma_1 < \sigma_2$ and $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$ be a complete filtered probability space satisfying the usual conditions. We suppose that $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$ supports an $(\mathcal{F}_t)_{t \geq 0}$ -Brownian motion $(W_t)_{t \geq 0}$ and an \mathbb{R}_+ -valued random variable τ . We assume that $(W_t)_{t \geq 0}$ and τ are independent and

$$E \left[\tau^{-\frac{1}{2}} \right] < \infty, \tag{5}$$

see Assumption 2.1 above. Let $\tilde{\mathcal{A}}$ be the set of all processes $\alpha : \Omega \times \mathbb{R}_+ \rightarrow [\sigma_1, \sigma_2]$ that are $(\mathcal{F}_t)_{t \geq 0}$ -progressively measurable. Given an agent chooses the control $\alpha \in \tilde{\mathcal{A}}$, the state process is defined by

$$X_t^\alpha := \int_0^t \alpha_s dW_s, \quad t \geq 0. \tag{6}$$

Remark 3.1 All feedback controls with a feedback function $m : \mathbb{R} \rightarrow [\sigma_1, \sigma_2]$ of bounded variation are contained in $\tilde{\mathcal{A}}$. Indeed, the SDE

$$dX_t = m(X_t)dB_t, \quad X_0 = 0, \tag{7}$$

has a weak solution because of Theorem 2.6.1 in [13], and pathwise uniqueness applies according to results in [16]. Hence, there exists a unique strong solution X^m to (7) (cf. Section 5.3 in [11]) and the control $(m(X_t^m))_{t \geq 0}$ is contained in $\tilde{\mathcal{A}}$.

Let $p \in (0, 1)$ and denote by $q(\mu, 1 - p)$ the $(1 - p)$ -quantile of some probability measure $\mu \in \mathcal{P}(\mathbb{R})$, i.e., for a Borel probability measures μ on \mathbb{R} , $q(\mu, 1 - p) = \inf\{r \in$

$\mathbb{R} : \mu((-\infty, r]) \geq 1 - p$. If $\mu = \text{Law}(X_\tau^\alpha)$, i.e., μ is the law of X_τ^α for a control $\alpha \in \tilde{\mathcal{A}}$, we write $q(X_\tau^\alpha, 1 - p)$ for $q(\mu, 1 - p)$. In the mean field game that corresponds to the n -player game described in Sect. 2, a single player wants to maximize the probability of being larger than the population $(1 - p)$ -quantile over all admissible controls. In the mean field game, we define equilibria in the following sense:

Definition 3.2 A tuple $(\mu^*, \alpha^*) \in \mathcal{P}(\mathbb{R}) \times \tilde{\mathcal{A}}$ is called *mean field equilibrium* if

(i) for all $\alpha \in \tilde{\mathcal{A}}$,

$$P(X_\tau^{\alpha^*} > q(\mu^*, 1 - p)) \geq P(X_\tau^\alpha > q(\mu^*, 1 - p)),$$

(ii) it holds $\mu^* = \text{Law}(X_\tau^{\alpha^*})$.

Lemma 3.3 Let $b \in \mathbb{R}$. Then

$$P(X_\tau^{m_b} > b) = \max_{\alpha \in \tilde{\mathcal{A}}} P(X_\tau^\alpha > b), \tag{8}$$

where m_b is defined in Eq. (3).

Proof The result follows from the diffusion control problem studied by McNamara [14]. In more detail, let $\alpha \in \tilde{\mathcal{A}}$. Assume, without loss of generality, that there exists a *regular conditional probability* $Q : \mathbb{R}_+ \times \mathcal{F} \rightarrow [0, 1]$ for \mathcal{F} given τ . Because τ and W are independent, W is also a Brownian motion under $Q(t, \cdot)$ for P_τ -a.e. $t \in \mathbb{R}_+$. Hence, we see that

$$Q(t, \{X_t^\alpha > b\}) = Q(t, \{X_t^\alpha > b\}) \leq Q(t, \{X_t^{m_b} > b\}), \text{ for } P_\tau\text{-a.e. } t \in \mathbb{R}_+,$$

using either [14], Remark 8, or [19], Proposition C.5. Finally,

$$P(X_\tau^\alpha > b) = \int_0^\infty Q(t, X_t^\alpha > b) P_\tau(dt) \leq \int_0^\infty Q(t, X_t^{m_b} > b) P_\tau(dt) = P(X_\tau^{m_b} > b).$$

We refer to [11], Section 5.3.C, or [10], Section 1.3, for more details on regular conditional probabilities. □

Lemma 3.3 implies that the optimal control for (8) is the threshold control m_b . In the following, we just write X^b for the state X^{m_b} . The process X^b is a so-called *oscillating Brownian motion* (OBM) with threshold b , introduced in [12]. In more detail, OBM is defined as follows.

Definition 3.4 Let $b \in \mathbb{R}$. We call the solution X^b of the SDE

$$dX_t = m_b(X_t)dB_t, \quad X_0 = 0, \tag{9}$$

oscillating Brownian motion (OBM) with threshold b and initial value 0.

There exists indeed a unique strong solution to the SDE (9) because of Remark 3.1. For OBMs, one can explicitly calculate the probability density function and the cumulative distribution function. For the reader’s convenience, we recall the following result on OBMs.

Proposition 3.5 Let $b \in \mathbb{R}$ and X^b be an OBM with threshold b and initial value 0. Then, for all $t > 0$, the random variable X_t^b has a probability density function $p(t, -b, \cdot - b)$ with

respect to the Lebesgue measure, where

$$p(t, x, y) = \begin{cases} \frac{2\sigma_1}{\sigma_2(\sigma_1+\sigma_2)} \frac{1}{\sqrt{2\pi t}} e^{-\left(\frac{x}{\sigma_1} - \frac{y}{\sigma_2}\right)^2 \frac{1}{2t}}, & \text{if } x \geq 0, y < 0, \\ \frac{2\sigma_2}{\sigma_1(\sigma_1+\sigma_2)} \frac{1}{\sqrt{2\pi t}} e^{-\left(\frac{y}{\sigma_1} - \frac{x}{\sigma_2}\right)^2 \frac{1}{2t}}, & \text{if } x < 0, y \geq 0, \\ \frac{1}{\sigma_1\sqrt{2\pi t}} \left(e^{-\frac{(y-x)^2}{2\sigma_1^2 t}} + \frac{\sigma_2 - \sigma_1}{\sigma_1 + \sigma_2} e^{-\frac{(y+x)^2}{2\sigma_1^2 t}} \right), & \text{if } x \geq 0, y \geq 0, \\ \frac{1}{\sigma_2\sqrt{2\pi t}} \left(e^{-\frac{(y-x)^2}{2\sigma_2^2 t}} + \frac{\sigma_1 - \sigma_2}{\sigma_1 + \sigma_2} e^{-\frac{(y+x)^2}{2\sigma_2^2 t}} \right), & \text{if } x < 0, y < 0, \end{cases}$$

for all $x, y \in \mathbb{R}$. Moreover, the cumulative distribution function of X_t^b is given by

$$F_t^b(x) = \begin{cases} \Phi\left(\frac{x}{\sigma_2\sqrt{t}}\right) - \frac{\sigma_2 - \sigma_1}{\sigma_1 + \sigma_2} \Phi\left(\frac{x - 2b}{\sigma_2\sqrt{t}}\right), & \text{if } x < b, b \geq 0, \\ \frac{2\sigma_2}{\sigma_1 + \sigma_2} \Phi\left(\frac{x - b\left(1 - \frac{\sigma_1}{\sigma_2}\right)}{\sigma_1\sqrt{t}}\right) - \frac{\sigma_2 - \sigma_1}{\sigma_1 + \sigma_2}, & \text{if } x \geq b, b \geq 0, \\ \frac{2\sigma_1}{\sigma_1 + \sigma_2} \Phi\left(\frac{x - b\left(1 - \frac{\sigma_2}{\sigma_1}\right)}{\sigma_2\sqrt{t}}\right), & \text{if } x < b, b < 0, \\ \Phi\left(\frac{x}{\sigma_1\sqrt{t}}\right) - \frac{\sigma_2 - \sigma_1}{\sigma_1 + \sigma_2} \Phi\left(\frac{2b - x}{\sigma_1\sqrt{t}}\right), & \text{if } x \geq b, b < 0. \end{cases} \tag{10}$$

The proof of Proposition 3.5 follows either from [12], Theorem 1, or [19], Proposition B.2 and Proposition B.4. For more details on OBMs, we refer to [12] and [19], Appendix B. Using the independence of the OBM X^b and the random time τ , one can derive from Proposition 3.5 the cumulative distribution function of X_τ^b .

Lemma 3.6 *Let $b \in \mathbb{R}$ and X^b be an OBM with threshold b and initial value 0. Then the cumulative distribution function of the random variable X_τ^b , denoted by F_τ^b , is given by*

$$F_\tau^b(x) = \int_0^\infty F_t^b(x) P_\tau(dt), \quad x \in \mathbb{R}.$$

We refer to Proposition B.8 in [19] for the proof of Lemma 3.6. One can show that the functions F_t^b and F_τ^b are Lipschitz continuous in the state variable as well as in the threshold b (see Appendix B in [19] for more details).

The standard approach to solve mean field games is to consider mappings from probability distributions to the distributions of optimally controlled states and find their fixed points, the so-called *equilibrium measures* (see, e.g., [3] and [4]). However, Lemma 3.3 allows to study the distributions of OBMs only, which can be parameterized by the real-valued threshold $b \in \mathbb{R}$. For identifying equilibria, it suffices to show that the function

$$f : \mathbb{R} \rightarrow \mathbb{R}, \quad b \mapsto q(X_\tau^b, 1 - p),$$

has a unique fixed point. Indeed, if $f(b) = b$, then $b = q(X_\tau^b, 1 - p)$. Lemma 3.3 further implies that $P(X_\tau^b > q(X_\tau^b, 1 - p)) = \max_{\beta \in \mathcal{M}} P(X_\tau^\beta > q(X_\tau^b, 1 - p))$; hence, $(\text{Law}(X_\tau^b), m_b)$ is an equilibrium. The main result of this section is the following:

Theorem 3.7 *There exists a unique $b^* \in \mathbb{R}$ such that $q(X_\tau^{b^*}, 1 - p) = b^*$. The tuple $(\text{Law}(X_\tau^{b^*}), m_{b^*})$ is a mean field equilibrium.*

Proof Recall that $F_\tau^b(x) = P(X_\tau^b \leq x)$, $x \in \mathbb{R}$. Consider the map $f(b) := F_\tau^b(b)$. Using the explicit formula for F_τ^b given by Proposition 3.5 and Lemma 3.6 (see also [19], Proposition B.4 and Proposition B.8), we observe that

$$f(b) = \begin{cases} \int_0^\infty \left(\frac{2\sigma_2}{\sigma_2 + \sigma_1} \Phi\left(\frac{b}{\sigma_2\sqrt{t}}\right) - \frac{\sigma_2 - \sigma_1}{\sigma_2 + \sigma_1} \right) P_\tau(dt), & \text{if } b \geq 0, \\ \int_0^\infty \frac{2\sigma_1}{\sigma_2 + \sigma_1} \Phi\left(\frac{b}{\sigma_1\sqrt{t}}\right) P_\tau(dt), & \text{if } b < 0. \end{cases} \tag{11}$$

The function f is strictly increasing and $f(0) = \frac{\sigma_1}{\sigma_1 + \sigma_2}$. Moreover, f is Lipschitz continuous because of the assumption (5) above, and Proposition B.10 and Proposition B.12 in [19]. Monotone convergence implies that $\lim_{b \rightarrow +\infty} f(b) = 1$ and $\lim_{b \rightarrow -\infty} f(b) = 0$. Hence, the intermediate value theorem implies that there exists a unique $b^* \in \mathbb{R}$ such that

$$f(b^*) = 1 - p \in (0, 1).$$

We have $f(0) \leq 1 - p$ if and only if $b^* \geq 0$. We conclude that $q(X_\tau^{b^*}, 1 - p) = b^*$ because $F_\tau^{b^*}(b^*) = 1 - p$. Lemma 3.3 implies that

$$\begin{aligned} P(X_\tau^{b^*} > q(X_\tau^{b^*}, 1 - p)) &= P(X_\tau^{b^*} > b^*) \\ &= \sup_{\alpha \in \tilde{\mathcal{A}}} P(X_\tau^\alpha > b^*) = \sup_{\alpha \in \tilde{\mathcal{A}}} P(X_\tau^\alpha > q(X_\tau^\alpha, 1 - p)), \end{aligned}$$

and thus, $(\text{Law}(X_\tau^{b^*}), m_{b^*})$ is a mean field equilibrium. □

Example 3.8 If τ is exponentially distributed with parameter $\lambda > 0$, then the threshold b^* in Theorem 3.7 is given by

$$b^* = \begin{cases} -\frac{\sigma_2}{\sqrt{2\lambda}} \log\left(\frac{\sigma_1 + \sigma_2}{\sigma_2} p\right), & \text{if } p < \frac{\sigma_2}{\sigma_1 + \sigma_2}, \\ \frac{\sigma_2}{\sqrt{2\lambda}} \log\left(\frac{\sigma_1 + \sigma_2}{\sigma_1} (1 - p)\right), & \text{if } p \geq \frac{\sigma_2}{\sigma_1 + \sigma_2}. \end{cases} \tag{12}$$

Note that the exponential distribution satisfies the condition (5). Proposition 3.5 and Lemma 3.6 imply that for $b \geq 0$

$$F_\tau^b(b) = 1 - \frac{2\lambda\sigma_2}{\sigma_1 + \sigma_2} \int_0^\infty e^{-\lambda t} \Phi\left(-\frac{b}{\sigma_2\sqrt{t}}\right) dt = 1 - \frac{\sigma_2}{\sigma_1 + \sigma_2} e^{-\frac{b}{\sigma_2}\sqrt{2\lambda}}.$$

The last step follows because

$$\Phi\left(-\frac{b}{\sigma_2\sqrt{t}}\right) = \frac{1}{2} \text{erfc}\left(\frac{b}{\sigma_2\sqrt{2t}}\right), \quad \text{erfc}(x) := \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-y^2} dy,$$

and the Laplace transform of the right-hand side is equal to $\frac{1}{2\lambda} e^{-\frac{b}{\sigma_2}\sqrt{2\lambda}}$ (see, e.g., [9], Chapter 8, Table 3). Similarly, we see for $b < 0$ that

$$F_\tau^b(b) = \frac{\sigma_1}{\sigma_1 + \sigma_2} e^{\frac{b}{\sigma_1}\sqrt{2\lambda}}.$$

Therefore, the function f in (11) is given by

$$f(b) = \begin{cases} 1 - \frac{\sigma_2}{\sigma_1 + \sigma_2} e^{-\frac{b}{\sigma_2}\sqrt{2\lambda}}, & \text{if } b \geq 0, \\ \frac{\sigma_1}{\sigma_1 + \sigma_2} e^{\frac{b}{\sigma_1}\sqrt{2\lambda}}, & \text{if } b < 0. \end{cases}$$

Now, it is straightforward to verify that $f(b^*) = 1 - p$ if and only if b^* satisfies Eq. (12).

Example 3.9 Suppose that τ is uniformly distributed on some interval $[t_1, t_2]$ with $0 \leq t_1 < t_2$. Notice that condition (5) is again satisfied. The cumulative distribution function of the OBM X^b at time τ and the point b is given by

$$\begin{aligned}
 F_\tau^b(b) &= \int_{t_1}^{t_2} \frac{1}{t_2 - t_1} F_t^b(b) dt \\
 &= \begin{cases} \frac{1}{t_2 - t_1} \frac{2b^2}{\sigma_2(\sigma_1 + \sigma_2)} \int_{\frac{\sigma_2^2}{b^2}t_1}^{\frac{\sigma_2^2}{b^2}t_2} \Phi\left(\frac{1}{\sqrt{t}}\right) dt - \frac{\sigma_2 - \sigma_1}{\sigma_1 + \sigma_2}, & \text{if } b > 0, \\ \frac{\sigma_1}{\sigma_1 + \sigma_2}, & \text{if } b = 0, \\ \frac{1}{t_2 - t_1} \frac{2b^2}{\sigma_1(\sigma_1 + \sigma_2)} \int_{\frac{\sigma_1^2}{b^2}t_1}^{\frac{\sigma_1^2}{b^2}t_2} \left(1 - \Phi\left(\frac{1}{\sqrt{t}}\right)\right) dt, & \text{if } b < 0, \end{cases} \\
 &= \begin{cases} \frac{1}{t_2 - t_1} \frac{2b^2}{\sigma_2(\sigma_1 + \sigma_2)} \left(G\left(\frac{\sigma_2^2 t_2}{b^2}\right) - G\left(\frac{\sigma_2^2 t_1}{b^2}\right)\right) - \frac{\sigma_2 - \sigma_1}{\sigma_1 + \sigma_2}, & \text{if } b > 0, \\ \frac{\sigma_1}{\sigma_1 + \sigma_2}, & \text{if } b = 0, \\ \frac{1}{t_2 - t_1} \frac{2b^2}{\sigma_1(\sigma_1 + \sigma_2)} \left(G\left(\frac{\sigma_1^2 t_2}{b^2}\right) - G\left(\frac{\sigma_1^2 t_1}{b^2}\right)\right), & \text{if } b < 0, \end{cases}
 \end{aligned}$$

where $G(s) := (s + 1)\Phi\left(\frac{1}{\sqrt{s}}\right) + \sqrt{s}\varphi\left(\frac{1}{\sqrt{s}}\right)$, $s > 0$, and where φ and Φ denote respectively the standard Gaussian density and cumulative distribution functions. For given parameters, one can solve the equation $F_\tau^b(b) = 1 - p$ for b numerically, but one cannot find a solution in closed form.

4 Approximate Nash equilibrium

We return to the setting of Sect. 2 and state the main result.

Theorem 4.1 *Let $b^* \in \mathbb{R}$ be the threshold of the mean field equilibrium control, given by Theorem 3.7. Let $\alpha^* = (\alpha^{1,*}, \dots, \alpha^{n,*}) \in \mathcal{A}_n$ be the tuple of strategies, defined by*

$$\alpha_t^{i,*} := m_{b^*}(X_t^i) = \begin{cases} \sigma_2, & \text{if } X_t^i \leq b^*, \\ \sigma_1, & \text{if } X_t^i > b^*, \end{cases} \tag{13}$$

Then, there exists a sequence $\varepsilon_n \geq 0$ with $\lim_{n \rightarrow \infty} \varepsilon_n = 0$ such that α^ is an ε_n -Nash equilibrium of the n -player game. We can choose $\varepsilon_n \in \mathcal{O}(n^{-1/2})$.*

Remark 4.2 The approximate equilibrium in Theorem 4.1 does not depend on the information structure of the game. Players may or may not have information about the other players' states and the random times. For implementing the strategy (13), Player i only needs to observe her own state. Therefore, the tuple of strategies (13) is also an approximate Nash equilibrium for the game version where the players cannot observe each other.

Moreover, the tuple of strategies (13) is an approximate Nash equilibrium for the game where players are only allowed to choose feedback strategies depending on all players' states as in the article [1]. In the case $\tau_1, \dots, \tau_n \equiv T$, Theorem 4.1 covers the main result of the n -player game with finite deterministic time horizon presented in article [1]. We refer to Remark 2.9 above for more details.

The proof of Theorem 4.1 is similar to the proof of Theorem 2 in [1], where the case with a common deterministic time horizon is treated. Nevertheless, we present the proof of Theorem 4.1 in a concise manner. More details can be found in [19], Theorem 3.2.16.

Proof of Theorem 4.1 Note that $\mathcal{Q}(\alpha^*) \neq \emptyset$ (see Remark 2.7). Let $i \in \{1, \dots, n\}$. We compare $\alpha^* = (\alpha^{1,*}, \dots, \alpha^{n,*})$ with the tuple where Player i deviates from $\alpha^{i,*}$ by choosing a strategy β . To this end, let $\beta \in \mathcal{A}^i$ such that $\mathcal{Q}(\alpha^{-i,*}, \beta) \neq \emptyset$, and let $P \in \mathcal{Q}(\alpha^*)$ and $\tilde{P} \in \mathcal{Q}(\alpha^{-i,*}, \beta)$. In the proof, we use the following notation:

- $(\alpha^{-i,*}, \beta) = (\alpha^{1,*}, \dots, \alpha^{i-1,*}, \beta, \alpha^{i+1,*}, \dots, \alpha^{n,*})$,
- $X^{-i} = (X^1, \dots, X^{i-1}, X^{i+1}, \dots, X^n)$,
- $\tau_{-i} = (\tau_1, \dots, \tau_{i-1}, \tau_{i+1}, \dots, \tau_n)$,
- $\mu := P \circ (X_{\tau_1}^1)^{-1}$ (i.e., the law of an OBM with threshold b^* at time τ_1),
- $\mu^n = \frac{1}{n} \sum_{j=1}^n \delta_{X_{\tau_j}^j}$ and $\mu^{n-1} = \frac{1}{n-1} \sum_{j \neq i} \delta_{X_{\tau_j}^j}$.

We aim at showing that

$$\tilde{P} \left(X_{\tau_i}^i > q(\mu^n, 1 - p) \right) - P \left(X_{\tau_i}^i > q(\mu^n, 1 - p) \right) \leq \frac{C}{\sqrt{n}}, \tag{14}$$

for some constant $C > 0$ independent of \tilde{P} , P , and the control β . Note that

$$P \circ (X^{-i}, \tau_1, \dots, \tau_n)^{-1} = \tilde{P} \circ (X^{-i}, \tau_1, \dots, \tau_n)^{-1}. \tag{15}$$

Step 1. We first estimate $P(X_{\tau_i}^i > q(\mu^n, 1 - p))$ from below. Notice that $X_{\tau_i}^i$ and the empirical quantile $q(\mu^n, 1 - p)$ are *not independent*. Therefore, we consider the empirical measure μ^{n-1} and define

$$A(n) := q \left(\mu^{n-1}, \frac{n}{n-1}(1 - p) \right), \tag{16}$$

i.e., $A(n)$ is the empirical $\frac{n}{n-1}(1 - p)$ -quantile of the states X^{-i} . Notice that $A(n)$ is independent of $X_{\tau_i}^i$ under P and

$$A(n) \geq q(\mu^n, 1 - p).$$

Thus,

$$P \left(X_{\tau_i}^i > q(\mu^n, 1 - p) \right) \geq P \left(X_{\tau_i}^i > A(n) \right) = E \left[1 - F_{\tau}^{b^*} (A(n)) \right], \tag{17}$$

where E denotes the expectation under P and $F_{\tau}^{b^*}$ denotes the cumulative distribution function of the OBM with threshold b^* at a random time τ (see Lemma 3.6).

Step 2. Next, we estimate the first term in (14) from above: we replace in (14) the quantile $q(\mu^n, 1 - p)$ with a quantile that does not depend on β . To this end, we observe that

$$D(n) := q\left(\mu^{n-1}, 1 - \frac{n}{n-1}p\right) \leq q(\mu^n, 1 - p). \tag{18}$$

From (18), we conclude that

$$\begin{aligned} \tilde{P}\left(X_{\tau_i}^i > q(\mu^n, 1 - p)\right) &\leq \tilde{P}\left(X_{\tau_i}^i > D(n)\right) \\ &= \tilde{E}\left[\tilde{P}\left(X_{\tau_i}^i > D(n) \mid X^{-i}, \tau_1, \dots, \tau_n\right)\right], \end{aligned} \tag{19}$$

where \tilde{E} denotes the expectation w.r.t. the probability measure \tilde{P} . If Player i knew from the very beginning the value $D(n)$, then $m_{D(n)}$ would be the control maximizing the probability for Player i 's state to be greater than $D(n)$ at time τ_i . This is a consequence of the control problem studied by McNamara [14].

To be more precise, note that Ω is a complete, separable metric space and \mathcal{F} is its Borel σ -algebra. Thus, there exists a regular conditional probability $Q : \Omega \times \mathcal{F} \rightarrow [0, 1]$ for \mathcal{F} given $\sigma(X^{-i}, \tau_1, \dots, \tau_n)$. For \tilde{P} -almost every $\omega \in \Omega$, we make the following observations: we can construct a Brownian motion under $Q(\omega, \cdot)$ such that

$$X_t^i = \int_0^t \beta_s dW_s^i, \quad t \geq 0, \quad Q(\omega, \cdot)\text{-a.s.}$$

Moreover, we can work on a probability space $(\Omega, \hat{\mathcal{F}}, (\hat{\mathcal{F}}_t)_{t \geq 0}, Q(\omega, \cdot))$ satisfying the usual conditions because we can consider the completed version of \mathcal{F} , denoted by $\hat{\mathcal{F}}$, and the completed and right-continuous version of $(\mathcal{F}_t)_{t \geq 0}$, denoted by $(\hat{\mathcal{F}}_t)_{t \geq 0}$. Now, the results of McNamara (see [14], Remark 8) imply that on the probability space $(\Omega, \hat{\mathcal{F}}, (\hat{\mathcal{F}}_t)_{t \geq 0}, Q(\omega, \cdot))$ with Brownian motion W , time horizon $T = \tau_i(\omega)$, and $b = D(\omega, n)$, we have

$$Q\left(\omega, \left\{X_T^i > b\right\}\right) \leq 1 - F_T^b(b),$$

where F_T^b denotes the cumulative distribution function of an OBM with threshold b at time T (see Proposition 3.5 above). Hence,

$$\tilde{P}\left(X_{\tau_i}^i > D(n) \mid X^{-i}, \tau_1, \dots, \tau_n\right)(\omega) \leq 1 - F_T^b(b)|_{b=D(\omega, n), T=\tau_i(\omega)},$$

for \tilde{P} -almost every $\omega \in \Omega$. This implies, using the estimate (19), that

$$\tilde{P}\left(X_{\tau_i}^i > q(\mu^n, 1 - p)\right) \leq E\left[1 - F_{\tau}^b(b)|_{b=D(n)}\right], \tag{20}$$

where we use the identity (15), the fact that $D(n)$ only depends on (X^{-i}, τ^{-i}) , the independence of τ_1, \dots, τ_n and X^{-i} , and the connection of F_{τ}^b and F_t^b (see Lemma 3.6).

Step 3. We can now combine the estimates in (17) and (20) above leading to

$$\tilde{P} \left(X_{\tau_i}^i > q(\mu^n, 1 - p) \right) - P \left(X_{\tau_i}^i > q(\mu^n, 1 - p) \right)$$

$$\leq E[F_{\tau}^{b^*}(A(n))] - E[F_{\tau}^{D(n)}(D(n))]$$

$$\leq E \left| F_{\tau}^{b^*}(A(n)) - F_{\tau}^{b^*} \left(q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right) \right| \tag{21}$$

$$+ \left| F_{\tau}^{b^*} \left(q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right) - F_{\tau}^{b^*} \left(q \left(\mu, 1 - \frac{n}{n-1}p \right) \right) \right| \tag{22}$$

$$+ E \left| F_{\tau}^{b^*} \left(q \left(\mu, 1 - \frac{n}{n-1}p \right) \right) - F_{\tau}^{b^*}(D(n)) \right| \tag{23}$$

$$+ E \left| F_{\tau}^{b^*}(D(n)) - F_{\tau}^{D(n)}(D(n)) \right|. \tag{24}$$

We estimate the four terms on the right-hand side separately.

- (i) The term (21) can be estimated, using the Lipschitz continuity of $F_{\tau}^{b^*}$ (see [19], Proposition B.10):

$$\begin{aligned} & E \left| F_{\tau}^{b^*}(A(n)) - F_{\tau}^{b^*} \left(q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right) \right| \\ & \leq E \left[C_1 \left| A(n) - q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right| \wedge 2 \right] \\ & = \int_0^{\infty} P \left(C_1 \left| A(n) - q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right| \wedge 2 > \varepsilon \right) d\varepsilon \\ & = C_1 \int_0^{\frac{2}{C_1}} P \left(\left| A(n) - q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right| > \varepsilon \right) d\varepsilon, \end{aligned} \tag{25}$$

for $C_1 := \frac{2}{\sigma_1 \sqrt{2\pi}} E[\tau_1^{-1/2}]$. The $(n - 1)$ -states $X_{\tau_j}^j, j \neq i$, are independent and identically distributed under P . Using Hoeffding’s inequality, one can deduce that

$$P \left(\left| A(n) - q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right| > \varepsilon \right) \leq 2e^{-2(n-1)C_2^2\varepsilon^2}, \quad 0 < \varepsilon < 2,$$

where the constant $C_2 > 0$ only depends on C_1 . We find with (25) that

$$\begin{aligned} & E \left| F_{\tau}^{b^*}(A(n)) - F_{\tau}^{b^*} \left(q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right) \right| \\ & \leq 2C_1 \int_0^{\frac{2}{C_1}} e^{-2(n-1)C_2^2\varepsilon^2} d\varepsilon \leq \frac{C_1}{C_2\sqrt{n-1}} \int_0^{\infty} e^{-\frac{x^2}{2}} dx \leq \frac{\sqrt{2\pi}C_1}{C_2} \frac{1}{\sqrt{n-1}}. \end{aligned} \tag{26}$$

- (ii) The term (22) can be rewritten as follows:

$$\begin{aligned} & \left| F_{\tau}^{b^*} \left(q \left(\mu, \frac{n}{n-1}(1 - p) \right) \right) - F_{\tau}^{b^*} \left(q \left(\mu, 1 - \frac{n}{n-1}p \right) \right) \right| \\ & = \left| \frac{n}{n-1}(1 - p) - \left(1 - \frac{n}{n-1}p \right) \right| = \frac{1}{n-1}, \end{aligned} \tag{27}$$

because $F_{\tau}^{b^*}(x) = \mu((-\infty, x]), x \in \mathbb{R}$.

- (iii) The term (23) can be estimated as the term (21), and thus,

$$E \left| F_{\tau}^{b^*} \left(q \left(\mu, 1 - \frac{n}{n-1}p \right) \right) - F_{\tau}^{b^*}(D(n)) \right| \leq \frac{\sqrt{2\pi}C_1}{C_2} \frac{1}{\sqrt{n-1}}. \tag{28}$$

(iv) For the term (24), we observe, using the Lipschitz continuity of F_τ^b in the threshold b (see [19], Proposition B.12),

$$\begin{aligned} & \left| F_\tau^{b^*}(D(n)) - F_\tau^{D(n)}(D(n)) \right| \\ & \leq (C_1 |b^* - D(n)|) \wedge 2 \\ & \leq \left(C_1 \left| b^* - q \left(\mu, 1 - \frac{n}{n-1} p \right) \right| \right) \wedge 2 + \left(C_1 \left| q \left(\mu, 1 - \frac{n}{n-1} p \right) - D(n) \right| \right) \wedge 2, \end{aligned} \tag{29}$$

with C_1 defined as above in (i). The expected value of the second term can be estimated as the term (21). For the first term, we observe that there exists a constant $C_3 > 0$ such that

$$\begin{aligned} \left| b^* - q \left(\mu, 1 - \frac{n}{n-1} p \right) \right| &= \left| (F_\tau^{b^*})^{-1}(1-p) - (F_\tau^{D(n)})^{-1} \left(1 - \frac{n}{n-1} p \right) \right| \\ &\leq C_3 \left| 1 - p - \left(1 - \frac{n}{n-1} p \right) \right| = \frac{C_3 p}{n-1}. \end{aligned}$$

We conclude that

$$E \left| F_\tau^{b^*}(D(n)) - F_\tau^{D(n)}(D(n)) \right| \leq \left(\frac{\sqrt{2\pi} C_1}{C_2} + C_3 p \right) \frac{1}{\sqrt{n-1}}. \tag{30}$$

Finally, Eqs. (26), (27), (28), and (30) imply that

$$\tilde{P} \left(X_{\tau_i}^i > q(\mu^n, 1-p) \right) - P \left(X_{\tau_i}^i > q(\mu^n, 1-p) \right) \leq \frac{C}{\sqrt{n-1}} \in \mathcal{O} \left(\frac{1}{\sqrt{n}} \right),$$

for an appropriate constant $C > 0$ independent of β , P , and \tilde{P} . Thus, the feedback strategy α^* is an $\mathcal{O}(n^{-1/2})$ -Nash equilibrium. \square

Remark 4.3 The assumption that the random times τ_1, \dots, τ_n in the n -player game are independent is crucial. Theorem 4.1 does not extend to the case where the random times are non-constant and identical, because then the corresponding mean field game consists of

1. solving the optimization problem

$$\sup_{\alpha} \int_0^\infty P(X_t^\alpha > q(\mu_t, 1-p)) P_\tau(dt),$$

and finding an optimal control $\alpha^*(\mu)$ for a given measure flow $\mu = (\mu_t)_{t \geq 0}$,

2. finding a fixed point of the map $\mu \mapsto \text{Law} \left(X^{\alpha^*(\mu)} \right)$.

This mean field game is different from the one defined in Sect. 3. The control problem depends on flows of probability measures instead of a single probability measure on \mathbb{R} . Therefore, the fixed point arguments used in the proof of Theorem 3.7 are not applicable.

5 The more time, the higher the winning probability

In this section we analyze how the winning probability of a player depends on the actual time horizon. We consider only the case where the competition is sufficiently fierce, i.e.

$$p < \frac{\sigma_2}{\sigma_1 + \sigma_2}. \tag{31}$$

Notice that (31) implies that the threshold level b^* of Theorem 3.7 is positive. We denote the mean field equilibrium distribution by μ^* . Recall that μ^* is the law of the OBM with threshold b^* at an independent random time with distribution \mathcal{T} .

Now suppose that n is large and that in the game with n players everyone controls their states with the threshold control m_{b^*} . We select one player and assume that her realized time horizon is t . Then the probability for this particular player to be among the best p at the end of the game is approximately given by

$$w(t) := P\left(X_t^{b^*} > q(\mu^*, 1 - p)\right) = 1 - F_t^{b^*}(b^*) = \frac{2\sigma_2}{\sigma_1 + \sigma_2} \left(1 - \Phi\left(\frac{b^*}{\sigma_2\sqrt{t}}\right)\right).$$

We refer to the function w as the winning probability. Note that the winning probability is continuous and increasing in t . Moreover, we have

$$\lim_{t \downarrow 0} w(t) = 0.$$

Thus, if the actual time horizon t is small, then the winning probability is close to zero. This is plausible, since the state process starts in zero and is stopped early, and hence attains the positive level b^* with a small probability only.

Next observe that

$$\lim_{t \rightarrow \infty} w(t) = \frac{\sigma_2}{\sigma_1 + \sigma_2}.$$

Thus, the winning probability is bounded by $\frac{\sigma_2}{\sigma_1 + \sigma_2}$, and the bound is almost attained for large time horizons t . The bound corresponds to the expected average time that an OBM is spending above the threshold in the long run. Indeed, irrespective of the threshold b , one can show that

$$\lim_{t \rightarrow \infty} P(X_t^b \geq b) = \frac{\sigma_2}{\sigma_1 + \sigma_2}. \tag{32}$$

The bound allows also for a control theoretical interpretation: to this end consider the ergodic control problem with target functional

$$J(\alpha) := \liminf_{T \rightarrow \infty} \frac{1}{T} E \int_0^T 1_{\{X_t^\alpha \geq b\}} dt, \quad \alpha \in \tilde{\mathcal{A}},$$

where $\tilde{\mathcal{A}}$ is defined as in Sect. 3 and X^α is defined as in (6). One can show that m_b is an optimal control (see Remark 8 in [14]) and hence, using (32), $\sup_{\alpha \in \tilde{\mathcal{A}}} J(\alpha) = \frac{\sigma_2}{\sigma_1 + \sigma_2}$.

6 An application to online competitions

Our game formulation is generic and can correspond to a variety of practical situations. For example, the game applies to managers of mutual funds striving for their funds to be among the best performing. This application is described, for homogeneous managers, in detail in Section 7 of [1].

We here provide an alternative application to online competitions in which, usually, teams have to collaborate in order to solve a problem and provide a solution within a limited time frame. *Hackatons* are examples of such competitions, but also data science and machine learning related competitions offered on some platforms such as Kaggle, DrivenData, AICrowd etc. During these competitions, a score, based on a given evaluation metric, can be calculated

by each participant and a "public leaderboard" displays the relative ranks during the whole length of the competition.

In this context, the private state X_t^i of player i is interpreted as her score displayed at time t . The number of teams involved in online competitions can reach several tens of thousands, enough to consider the mean field approximation.

Level of risk.

We interpret the diffusion control as the possibility to control the level of risk taken. In the online contest's setting, teams can indeed choose to try and use well established methods, whose robustness is already studied and for which errors can be more easily and quickly corrected. We interpret this as low risk and diffusion coefficient σ_1 . On the other hand, the choice of diffusion coefficient σ_2 is interpreted as trying new techniques, for which there is less or no experience. Our results rigorously show in this context that if subtasks are going well, players will play safe, whereas if a given team is poorly performing on the evaluation metric, it has an incentive to play risky and try less common strategies.

Observability.

Participants to online competitions can submit a solution, and in that case their current score is calculated and displayed to all participants. However, if a team obtains a solution and tests it offline on the provided data set, then the associated score is not visible during the time it is not submitted on the platform. A given team can choose to reveal its solution and score only towards the end of the submission period.

Moreover, it is possible to design online competitions with only partial observability: one could easily imagine that teams only observe the best score, or a given quantile of the scores distribution, to assess their relative performance.

Our results show that if the number of players is large, observability does not matter, at least for the particular type of discontinuous criteria that we consider, which are common in these competitions, where a fixed cash prize is offered to the best performing team. This result also holds for continuous functions of the rank satisfying the symmetry condition given in Assumption 7.1.

Terminal time.

We consider two cases. Firstly, the case where the time horizons of all players are constant equal to $T \in (0, \infty)$, interpreted as the date at which a final assessment of the evaluation metric is made by the contest organizers. Secondly, the time horizon τ_i of team i quantifies the resources it can put into the competition, e.g. the number of working hours. For example, τ_i can be set proportional to the deterministic assessment date and the number of team members. Section 5 reveals that larger teams have an advantage compared to smaller teams.

7 Extensions

In this section, we discuss more general reward criteria for the n -player game. In particular, we consider rewards at the random time horizons that are given by measurable functions $g : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ of the state and the population quantile, instead of the "all-or-nothing" payoff given by the function $(x, q) \mapsto \mathbb{1}_{(q, \infty)}(x)$ before. First, we show that the mean field equilibrium control of Sect. 3 is also an equilibrium for the reward functions g , if g satisfies a symmetry and convexity condition. Then, we prove that this equilibrium provides an approximate Nash equilibrium of the n -player game.

7.1 Mean field game

Assume that we are in the setting of Sect. 3. The whole analysis in Sect. 3 depends on the optimality of the threshold control m_b for the particular choice of the reward $\mathbb{1}_{(b, \infty)}$ (Lemma 3.3). Results of McNamara [14] imply that m_b is not only optimal for this reward but also for more general reward functions that are continuous, have exponential growth, and satisfy a convexity condition and a symmetry condition. In more detail, we can generalize Theorem 3.7 to measurable functions $g : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ satisfying:

- Assumption 7.1** (i) $g(\cdot, q)$ is continuous and has exponential growth for any $q \in \mathbb{R}$,
 (ii) $g(\cdot, q)$ is convex on $(-\infty, q]$ and concave on $[q, \infty)$ for any $q \in \mathbb{R}$,
 (iii) for all $x \geq 0$ and $q \in \mathbb{R}$ it holds

$$\sigma_2 g(\sigma_1 x + q, q) + \sigma_1 g(-\sigma_2 x + q, q) = (\sigma_1 + \sigma_2)g(q, q).$$

Proposition 7.2 Let b^* be given by Theorem 3.7 and g satisfy Assumption 7.1. Then, $(\text{Law}(X_\tau^{b^*}), m_{b^*})$ is also a mean field equilibrium for the reward function g , i.e.,

$$E \left[g(X_\tau^{b^*}, q(X_\tau^{b^*}, 1 - p)) \right] = \sup_{\alpha \in \tilde{\mathcal{A}}} E \left[g(X_\tau^\alpha, q(X_\tau^{b^*}, 1 - p)) \right].$$

Proof As in Lemma 3.3, one can show for fixed $b \in \mathbb{R}$ that

$$E \left[g(X_\tau^b, b) \right] = \sup_{\alpha \in \tilde{\mathcal{A}}} E \left[g(X_\tau^\alpha, b) \right],$$

using either [14], Theorem 6, or [19], Theorem C.5. The assumptions on g guarantee that these theorems apply. Moreover, Theorem 3.7 implies the existence of a unique fixed point b^* of the map $b \mapsto q(X_\tau^b, 1 - p)$. For this fixed point b^* , we see that

$$E \left[g(X_\tau^{b^*}, q(X_\tau^{b^*}, 1 - p)) \right] = \sup_{\alpha \in \tilde{\mathcal{A}}} E \left[g(X_\tau^\alpha, q(X_\tau^{b^*}, 1 - p)) \right],$$

i.e., $(\text{Law}(X_\tau^{b^*}), m_{b^*})$ is an equilibrium for the reward function g . □

7.2 Approximate Nash equilibrium in the n -player game

Now, in the setting of Sect. 2, we show that the n -tuple with each entry equal to the mean field equilibrium strategy provides an approximate Nash equilibrium of the n -player game with reward g .

Definition 7.3 Let $\varepsilon > 0$. A tuple $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathcal{A}_n$ with $\mathcal{Q}(\alpha) \neq \emptyset$ is called ε -Nash equilibrium of the n -player game if for all $i \in \{1, \dots, n\}$ and $P^\alpha \in \mathcal{Q}(\alpha)$

$$E^{P^\alpha} \left[g \left(X_{\tau_i}^i, q(\mu^n, 1 - p) \right) \right] + \varepsilon \geq \sup_{\beta \in \mathcal{A}^i} \sup_{P \in \mathcal{Q}(\alpha_{-i}, \beta)} E^P \left[g \left(X_{\tau_i}^i, q(\mu^n, 1 - p) \right) \right],$$

where $(\alpha_{-i}, \beta) = (\alpha_1, \dots, \alpha_{i-1}, \beta, \alpha_{i+1}, \dots, \alpha_n)$ and $\sup \emptyset = -\infty$.

With some additional assumptions on the terminal reward g , we can show:

Proposition 7.4 Let g satisfy Assumption 7.1. In addition, assume that g is uniformly bounded, continuous, and $g(x, \cdot)$ is monotonically decreasing for all $x \in \mathbb{R}$. Let $\alpha^* \in \mathcal{A}_n$ be defined as in Theorem 4.1. Then, there exists a sequence $\varepsilon_n \geq 0$ with $\lim_{n \rightarrow \infty} \varepsilon_n = 0$ such that the control tuple $\alpha^* = (\alpha^{1,*}, \dots, \alpha^{n,*})$ is an ε_n -Nash equilibrium of the n -player game.

Proof All details of the proof can be found in Proposition 3.2.25 in [19]. Let $i \in \{1, \dots, n\}$. Moreover, let $\beta \in \mathcal{A}^i$ such that $\mathcal{Q}(\alpha^{-i,*}, \beta) \neq \emptyset$ and choose $P \in \mathcal{Q}(\alpha^*)$, $\tilde{P} \in \mathcal{Q}(\alpha^{-i,*}, \beta)$. Using the empirical quantiles $A(n)$ and $D(n)$ defined in (16) and (18), respectively, we find that

$$\begin{aligned} \tilde{E} \left[g(X_{\tau_i}^i, q(\mu^n, 1 - p)) \right] &\leq \tilde{E} \left[g(X_{\tau_i}^i, D(n)) \right], \\ E \left[g(X_{\tau_i}^i, q(\mu^n, 1 - p)) \right] &\geq E \left[g(X_{\tau_i}^i, A(n)) \right], \end{aligned}$$

because of the monotonicity of $g(x, \cdot)$. We use the notation E and \tilde{E} for the expectation w.r.t. P and \tilde{P} , respectively. Define the function

$$G(q) := E[g(X_{\tau_i}^i, q)] = \int_0^\infty \int_{-\infty}^\infty g(x, q) p(t, -q, x - q) dx P_\tau(dt), \quad q \in \mathbb{R},$$

where p denotes the probability density function of the OBM defined in Proposition 3.5. Note that

$$E \left[g(X_{\tau_i}^i, A(n)) \right] = E [G(A(n))],$$

because $X_{\tau_1}^1, \dots, X_{\tau_n}^n$ are independent under P . Similar to S, one can show that

$$\tilde{E} \left[g(X_{\tau_i}^i, D(n)) \right] \leq \tilde{E} [G(D(n))] = E [G(D(n))],$$

similar to Step 2 in the proof of Theorem 3.2.16 in [19]. The upper bound follows from Theorem 6 in [14]: the control maximizing the left-hand side is the threshold control with threshold $D(n)$, if conditioned on $D(n)$. We conclude that

$$\tilde{E} \left[g(X_{\tau_i}^i, q(\mu^n, 1 - p)) \right] - E \left[g(X_{\tau_i}^i, q(\mu^n, 1 - p)) \right] \leq E [G(D(n))] - E [G(A(n))].$$

Note that G is bounded and continuous, and the right-hand side only depends on the distribution of $n - 1$ independent OBMs with threshold b^* . Moreover, $A(n)$ and $D(n)$ converge to b^* in probability (see Lemma 3.2.20 in [19]) and hence, also in distribution. This means

$$\lim_{n \rightarrow \infty} |E [G(D(n))] - E [G(A(n))]| = 0.$$

Therefore, we can find a sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ with the desired properties. The sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ is independent of $i \in \{1, \dots, n\}$, $\beta \in \mathcal{A}$, $P \in \mathcal{Q}(\alpha^*)$, and $\tilde{P} \in \mathcal{Q}(\alpha^{-i,*}, \beta)$ because the distributions of $A(n)$ under P and of $D(n)$ under \tilde{P} are unique. □

Acknowledgements We thank two anonymous referees for carefully reading the article and for providing many suggestions for improvements. Support from the German Research Foundation through the project AN 1024/5-1 is gratefully acknowledged. Nabil Kazi-Tani's research is supported by the ANR project DREAMES ANR-21-CE46-0002-03.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ankirchner, S., Kazi-Tani, N., Wendt, J., Zhou, C.: Large ranking games with diffusion control. *Math. Oper. Res.* (2023) (in press)
2. Barrasso, A., Touzi, N.: Controlled diffusion mean field games with common noise and McKean–Vlasov second order backward sdes. *Theory Probab. Appl.* **66**(4), 613–639 (2022)
3. Carmona, R., Delarue, F.: Probabilistic theory of mean field games with applications. I, volume 83 of *Probability Theory and Stochastic Modelling*. Springer, Cham. Mean field FBSDEs, control, and games (2018)
4. Carmona, R., Delarue, F.: Probabilistic theory of mean field games with applications. II, volume 84 of *Probability Theory and Stochastic Modelling*. Springer, Cham. Mean field games with common noise and master equations (2018)
5. Cherny, A., Engelbert, H.-J.: Singular stochastic differential equations. *Lecture Notes in Mathematics*, vol. 1858. Springer, Berlin (2005)
6. Djete, M.F.: Extended mean field control problem: a propagation of chaos result. *Electron. J. Probab.* **27**, 1–53 (2022)
7. Djete, M.F.: Mean field games of controls: on the convergence of nash equilibria. *Ann. Appl. Probab.* **33**(4), 2824–2862 (2023)
8. Elie, R., Hubert, E., Mastrolia, T., Possamai, D.: Mean-field moral hazard for optimal energy demand response management. *Math. Financ.* **31**(1), 399–473 (2021)
9. Folland, G.B.: *Fourier Analysis and Its Applications*. Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove (1992)
10. Ikeda, N., Watanabe, S.: *Stochastic differential equations and diffusion processes*, volume 24 of *North-Holland Mathematical Library*, 2 edn. North-Holland Publishing Co., Amsterdam; Kodansha, Ltd., Tokyo (1989)
11. Karatzas, I., Shreve, S. E.: *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*, 2nd edn. Springer, New York (1991)
12. Keilson, J., Wellner, J.A.: Oscillating Brownian motion. *J. Appl. Probab.* **15**(2), 300–310 (1978)
13. Krylov, N.V.: *Controlled diffusion processes*, volume 14 of *Stochastic Modelling and Applied Probability*. Springer, Berlin. Translated from the 1977 Russian original by A. B. Aries, Reprint of the (1980) edition (2009)
14. McNamara, J.M.: Optimal control of the diffusion coefficient of a simple diffusion process. *Math. Oper. Res.* **8**(3), 373–380 (1983)
15. Nadirashvili, N.: Nonuniqueness in the martingale problem and the Dirichlet problem for uniformly elliptic operators. *Annali della Scuola Normale Superiore di Pisa. Classe di Scienze. Serie IV* **24**(3), 537–549 (1997)
16. Nakao, S.: On the pathwise uniqueness of solutions of one-dimensional stochastic differential equations. *Osaka J. Math.* **9**(3), 513–518 (1972)
17. Protter, P.E.: *Stochastic integration and differential equations*, volume 21 of *Stochastic Modelling and Applied Probability*, 2 edn. Springer, Berlin, 2005. Version 2.1, Corrected third printing (2005)
18. Stroock, D.W., Varadhan, S.R.S.: *Multidimensional diffusion processes*. In: *Classics in Mathematics*. Springer, Berlin. Reprint of the (1997) edition (2006)
19. Wendt, J.: *Diffusion control and games*. Ph.d thesis, Friedrich-Schiller-Universität Jena (2023)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.