Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Mean Field Contest with Singularity

Marcel Nutz, Yuchong Zhang

To cite this article:
Marcel Nutz, Yuchong Zhang (2023) Mean Field Contest with Singularity. Mathematics of Operations Research 48(2):1095-1118. https://doi.org/10.1287/moor.2022.1297

informs
https://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Vol. 48, No. 2, May 2023, pp. 1095–1118
ISSN 0364-765X (print), ISSN 1526-5471 (online)

Mean Field Contest with Singularity

Marcel Nutz,[a] **Yuchong Zhang**[b,*]

[a] Departments of Statistics and Mathematics, Columbia University, New York, New York 10027; [b] Department of Statistical Sciences, University of Toronto, Toronto, Ontario M5G1Z5, Canada
*Corresponding author

Contact: mnutz@columbia.edu, https://orcid.org/0000-0003-2936-2315 (MN); yuchong.zhang@utoronto.ca, https://orcid.org/0000-0002-7687-2783 (YZ)

Received: March 9, 2021
Revised: February 21, 2022
Accepted: June 5, 2022
Published Online in Articles in Advance: August 31, 2022

MSC2020 Subject Classification: Primary: 91A13; secondary: 91A65, 91A15

https://doi.org/10.1287/moor.2022.1297

Abstract. We formulate a mean field game where each player stops a privately observed Brownian motion with absorption. Players are ranked according to their level of stopping and rewarded as a function of their relative rank. There is a unique mean field equilibrium, and it is shown to be the limit of associated n-player games. Conversely, the mean field strategy induces n-player ε-Nash equilibria for any continuous reward function—but not for discontinuous ones. In a second part, we study the problem of a principal who can choose how to distribute a reward budget over the ranks and aims to maximize the performance of the median player. The optimal reward design (contract) is found in closed form, complementing the merely partial results available in the n-player case. We then analyze the quality of the mean field design when used as a proxy for the optimizer in the n-player game. Surprisingly, the quality deteriorates dramatically as n grows. We explain this with an asymptotic singularity in the induced n-player equilibrium distributions.

Funding: M. Nutz is supported by an Alfred P. Sloan Fellowship and the Division of Mathematical Sciences of the National Science Foundation [Grants DMS-1812661 and DMS-2106056]. Y. Zhang is supported in part by the Natural Sciences and Engineering Research Council of Canada [NSERC Discovery Grant RGPIN-2020-06290].

Keywords: mean field game • stochastic contest • optimal contract • Stackelberg game

1. Introduction

We formulate a mean field game where each player stops a privately observed Brownian motion with drift and absorption at the origin. Players are ranked according to their level of stopping and paid a reward, which is a decreasing function of the rank. This is an infinite-player version of the n-player game studied in Nutz and Zhang [42], which, in turn, extends the Seel–Strack model (Seel and Strack [46]), where only the top-ranked player receives a reward. First, we establish the existence and uniqueness of a mean field equilibrium for any given reward function. Second, we solve the problem of optimal reward design (optimal contract) for a principal who can choose how to distribute a given reward budget over the ranks and aims to maximize the performance (i.e., stopping level) at a given rank—for instance, the median performance among the players. An analogous problem was studied for the n-player case in Nutz and Zhang [42], but only a partial characterization of the optimal design is available. Here, taking the mean field limit enables a clear-cut answer.

The present work also serves as a case study: from the perspective of mean field analysis, a particular feature of this game is to be tractable without necessarily being smooth. Atoms occur naturally in the equilibrium distribution, but the mean field game nevertheless admits an equilibrium that can be described in closed form, and we can prove analytically that the equilibrium is unique. As the n-player equilibrium can also be described in detail, we can observe the quality of the mean field approximation, not only for the mean field game (with fixed reward function) but also for the reward design problem—which is a Stackelberg game between the principal and a continuum of players. It turns out that this case study offers a cautionary tale.

In the n-player setting, the Seel–Strack model has been generalized and varied in several directions: more general diffusion processes (Feng and Hobson [23]), random initial laws (Feng and Hobson [24]), heterogeneous loss constraints (Seel [45]), and behavioral players (Feng and Hobson [25]), among others. See also Fang et al. [22] and Nutz and Zhang [42] for references to other models on risk taking under relative performance pay, and see Vojnović [48] for an introduction to rank-order prize allocation. The novelty in the present work is the analysis of a mean field model along the lines of Seel–Strack and its optimal design problem; we focus on the original Brownian dynamics. For the theory and applications of mean field games, the monographs by Bensoussan et al. [4] and Carmona and Delarue [8, 9] provide an excellent overview and references. The very recent mean field model

Ankirchner et al. [1] can be related to the first part of this work. In Ankirchner et al. [1], players control the volatility of a Brownian motion up to an independent exponential time and are then ranked. The reward is 1 above a certain rank and 0 below. As the horizon is exponential and volatilities can only be chosen within an interval that is bounded and bounded away from 0, the model has a smooth equilibrium, and the distinct features of the present work do not appear. (Questions of optimal design, or reward functions other than the binary one, are not studied.) Contracts between a principal and n agents have been analyzed in Demski and Sappington [16], Green and Stokey [27], Harris et al. [29], Holmstrom [30], Mookherjee [38], and Nalebuff and Stiglitz [39], among many other studies. Closer to the present work, Carmona and Wang [11] and Elie et al. [20] study optimal contracts between a principal and infinitely many agents using the theory of mean field games for diffusion control.

1.1. Mean Field Equilibrium

The mean field game as formalized in Section 2 admits an equilibrium as soon as the reward function R is right continuous and a natural integrability condition on the drift parameter holds (the latter is also present in the n-player game). The equilibrium stopping distribution can be described in closed form using the right-continuous inverse of $y \longmapsto R(1-y)$ (see Theorem 1). Once the correct ansatz is guessed, the existence result is reduced to a verification proof following a direct martingale argument. The uniqueness result is more involved, in part because—in contrast to the n-player game and many other mean field models—atoms in the equilibrium cannot be excluded; in fact, the closed-form solution already indicates that atoms will arise unless R is strictly monotone. The first part of the proof (Section 3.1) relates flat stretches in reward to atoms in any potential equilibrium distribution. The basic idea is to show, a priori, that equal pay must correspond to equal performance: ranks with the same reward are occupied by players that stop at the same level, and ranks with different rewards are occupied by players that stop at different levels. On the other hand, jumps in reward are related to gaps in the support of the equilibrium distribution. The second part of the uniqueness proof (Section 3.2) is based on the idea that in any equilibrium, the opposing players collectively act such as to minimize the value function of a given representative player. This approach enables an analytic proof using optimal stopping theory and dynamic programming arguments: using the additional constraints shown in the first part, the minimization is shown to have a unique solution, proving the uniqueness of the equilibrium. This analysis is complicated by the presence of atoms. We remark that a similar uniqueness proof could be given for the n-player game, in which case it would simplify substantially because atoms can be excluded a priori (however, a different proof is already available).

1.2. Optimal Reward Design

In the n-player game, Nutz and Zhang [42] study the design problem for a principal maximizing the performance at the kth rank—for example, maximizing the revenue in a second-best auction or the median performance among employees, customers, students, and the like. To consider the analogue in the mean field limit, we replace the kth rank by the quantile $\alpha = k/n$—for instance, $\alpha = 0.5$ for the median player. A reasonable guess for the optimal reward design is to (a) pay nothing to the ranks below the target α and (b) distribute the reward budget uniformly over the ranks above. We show in Theorem 2 that this guess is correct for any value of the drift parameter. By contrast, the guess is wrong in the finite player game: for nonnegative drift, the general shape is correct, but the optimal cutoff point can be at a rank strictly below the target rank k. For negative drift, the sharp cutoff can be replaced by a smoothed shape that also pays a small number of rewards of different sizes. (A full characterization of the optimal reward is only available for zero drift; see Nutz and Zhang [42].) Again, the mean field limit proves useful in allowing for a fuller description and a clearer result. On the other hand, knowing only the mean field limit may suggest an oversimplified picture for the finite player game. One previous model where the optimal reward design problem was solved completely for both the n-player and mean field setting is the Poissonian game of Nutz and Zhang [41], where players control the jump intensity and are ranked according to their jump times. There, the optimal designs are more similar between the two settings; part (a) of the abovementioned guess is always correct—the optimal reward has a sharp cutoff exactly at the target rank—though the shape over the ranks above the target is concave rather than being flat as in part (b). A related mean field game is considered in Bayraktar et al. [3], with diffusion instead of Poissonian dynamics. Both (a) and (b) turn out to be correct in the mean field setting. The n-player game is not tractable, and its optimal design was not studied. In the light of the present work, one should not take for granted that the shape is analogous to the mean field limit.

1.3. Mean Field Approximation

In the preceding discussion, the mean field model is formulated directly as a game with infinitely many players. To connect this model rigorously with the n-player game, we show in Theorem 3 that for any given reward function, the unique n-player equilibrium for the induced reward converges to the mean field counterpart. Moreover, the value function of a player in the n-player game converges uniformly to the value function in the mean field game. This way of connecting the two models is classical in the mean field game literature starting with Lasry and Lions [36]; see, in particular, Bardi [2], Cardaliaguet et al. [6], Carmona et al. [12], Fischer [26], and Lacker [33, 34] and the recent works by Djete [17, 18], Iseri and Zhang [32], and Lacker and Le Flem [35]. Another way of connecting the two models, going back to Huang et al. [31], is to fix the optimal strategy from the mean field equilibrium and consider it in the n-player game for large n. Consistent with a broad literature (e.g., Campi and Fischer [5], Carmona and Delarue [7], Carmona and Lacker [10], Carmona et al. [12], and Cecchin and Fischer [13]), we show in Theorem 4 that for any continuous reward function, this control induces an ε-Nash equilibrium for large n; that is, players cannot improve their expected performance by more than ε through unilateral deviations from the mean field strategy. Surprisingly, continuity is necessary: any discontinuity in the reward function is shown to rule out the ε-Nash equilibrium property for large n. A discontinuity in reward leads to a gap in the support of the mean field equilibrium distribution. Because of a knife-edge phenomenon in the sampling for large but finite n, a player can improve substantially by unilaterally stopping inside the gap with a well-chosen distribution. In the study of diffusive mean field games with absorption, Campi and Fischer [5, section 7] describe an example with degenerate volatility where the ε-Nash equilibrium property fails. The degeneracy is exogenously chosen so that absorption cannot occur in the n-player game but will occur in the mean field game, therefore creating a disconnect between the two. In Ankirchner et al. [1], on the other hand, the ε-Nash equilibrium property always holds, despite the reward being discontinuous, because the dynamics of the game itself (nondegenerate volatility) guarantee a smooth equilibrium. The models of Cecchin et al. [14], Delarue and Tchuendom [15], Martin et al. [37], and Nutz [40], highlight a different type of discrepancy where some mean field equilibria can fail to be limits of n-player equilibria. Those examples arise as a result of nonuniqueness of mean field equilibria and thus are orthogonal to the issues in the present work.

Next, we discuss the quality of the mean field approximation for the reward design problem (Section 5.3); here, using the mean field proxy seems particularly attractive because the optimal n-player design was fully solved only for zero drift. Our numerical discussion uses the zero drift case for that same reason. We observe that the optimal design for the n-player problem converges to the mean field counterpart. Moreover, the induced performance of the former in the n-player game converges to the performance of the latter in the mean field game. This is consistent with Elie et al. [20], where the authors prove convergence of the optimal designs and induced performances for an example of their diffusive game—which is much more complex, yet smoother, than ours. But more important, and maybe surprising, the *quality of the mean field proxy* from the point of view of the principal is strikingly poor in the present model (this aspect was not studied in Elie et al. [20]): for moderate n, the performance induced by the mean field optimizer is significantly inferior to the exact n-player optimal design. For large n, the performance *deteriorates* even further, eventually achieving only 50% of the optimum. The tractability of the present model allows us to explain the reason for this phenomenon in detail. In the literature, mean field approximations are often applied in finite-player games without further analysis. The present study may offer the message that the quality of the approximation warrants consideration, especially when smoothness is not guaranteed, and that the mean field model can yield an oversimplified picture of the n-player game in some cases.

2. Mean Field Equilibrium

In this section we define the mean field contest as a game with a continuum of players and prove that there exists a unique Nash equilibrium. It will be shown in Section 5 that this equilibrium is indeed the limit of associated n-player games as $n \to \infty$. Throughout, we fix a *reward function,* defined as a right-continuous and decreasing[1] function $R : [0,1] \to \mathbb{R}_+$ satisfying $R(0) > R(1-) = R(1)$. It will be shown in Remarks 2 and 5, respectively, that left continuity at the last rank is essential for the uniqueness of the equilibrium and right continuity is essential for its existence.

Each infinitesimal player i privately observes a drifted Brownian motion $X_t^i = x_0 + \mu t + \sigma W_t^i$ with absorption at $x = 0$ and chooses a (possibly randomized) stopping time $\tau_i < \infty$. The initial value $x_0 \in (0, \infty)$, drift $\mu \in \mathbb{R}$, and dispersion $\sigma \in (0, \infty)$ are identical across players, whereas the Brownian motions are independent. Let $Y_i = X^i_{\tau_i}$ be the position at stopping. If Y_i are independent and identically distributed (i.i.d.) across players with law F, the empirical distribution of (Y_i) is almost surely (a.s.) equal to F, by the exact law of large numbers (see Remark 3). That is, if all players choose the same stopping distribution, then the collection of players (deterministically) ends up distributed accordingly. Hence the rank of player i if she stops at x while all other players stop according to distribution F is defined as $1 - F(x)$.[2] If F does not have an atom at x, meaning that there are no ties at this rank, she receives the reward $R(1 - F(x))$. Otherwise, she receives the average of $R(1 - y)$ over $y \in [F(x-), F(x)]$, which is equivalent to splitting ties uniformly at random. Thus, writing

$$g(y) := R(1-y),$$

the payoff $\xi^F(x)$ for stopping at x if all other players use F is

$$\xi^F(x) = \begin{cases} g(F(x)) & \text{if } F(x) = F(x-), \\ \dfrac{1}{F(x)-F(x-)}\displaystyle\int_{F(x-)}^{F(x)} g(y)dy & \text{if } F(x) > F(x-). \end{cases} \tag{1}$$

The set $\mathcal{F}$ of distributions that are *feasible* (i.e., can be attained by stopping X^i with a randomized stopping time) is characterized through Skorokhod's embedding theorem and the scale function h of $X := X^i$, as observed in Seel and Strack [46].

Lemma 1. *The set $\mathcal{F}$ consists of all distributions F on $[0,\infty)$ satisfying $\int h dF = 1$ if $\mu > 0$ and $\int h dF \le 1$ if $\mu \le 0$, where*

$$h(x) = h_{x_0}(x) = \begin{cases} \dfrac{\exp\left(\dfrac{-2\mu x}{\sigma^2}\right) - 1}{\exp\left(\dfrac{-2\mu x_0}{\sigma^2}\right) - 1} & \text{if } \mu \ne 0, \\ \dfrac{x}{x_0} & \text{if } \mu = 0. \end{cases} \tag{2}$$

This result is due to Hall [28].[3] We say that $F \in \mathcal{F}$ is a mean field *equilibrium* if no player is incentivized to deviate from F; that is, $\int \xi^F dF \ge \int \xi^F d\tilde{F}$ for all $\tilde{F} \in \mathcal{F}$. The associated value function $u(x)$ is defined as the supremum expected reward achievable for a player starting at level x (instead of x_0) if all others use F. Denote the average reward by $\bar{R} = \int_0^1 R(r)dr$ and set

$$\bar{\mu}_\infty := \frac{\sigma^2}{2x_0}\log\left(\frac{R(0)-R(1)}{R(0)-\bar{R}}\right) > 0.$$

In the following result, g^{-1} denotes the *right*-continuous inverse:

$$g^{-1}(z) = \inf\{y \in [0,1] : g(y) > z\} \text{ for } z \in [R(1), R(0)), \qquad g^{-1}(R(0)) := 1.$$

Theorem 1. *Let $\mu < \bar{\mu}_\infty$. There exists a unique equilibrium. Its cumulative distribution function (cdf) is*

$$F^*(x) = g^{-1}([R(1) + (\bar{R} - R(1))h(x)] \wedge R(0)), \tag{3}$$

where g^{-1} is the right-continuous inverse of g, and the equilibrium value function is

$$u^*(x) = [R(1) + (\bar{R} - R(1))h(x)] \wedge R(0).$$

In particular, the equilibrium has compact support $[0,\bar{x}]$ for $\bar{x} = h^{-1}\left(\frac{R(0)-R(1)}{\bar{R}-R(1)}\right)$, and its atoms are in one-to-one correspondence with intervals where the reward R is constant.

Remark 1. The equilibrium distribution (3) is invariant under affine transformations of the reward R. In particular, we may normalize the reward to satisfy $R(1) = 0$ and $\bar{R} = 1$ without loss of generality.

Remark 2. The condition $R(1-) = R(1)$ is not necessary for the existence result in Theorem 1 (it is not used in the proof), but it is crucial for uniqueness. Indeed, we claim that infinitely many equilibria arise whenever $R(1-) > R(1)$. To see this, fix a constant $R(1-) \ge \beta > R(1)$ and define the new reward $\tilde{R}$ by $\tilde{R}(1) = \beta$ and $\tilde{R}(r) = R(r)$ for $r < 1$. We assume that β is so that the constant $\bar{\mu}_\infty$ associated with $\tilde{R}$ still satisfies $\mu < \bar{\mu}_\infty$. As mentioned previously, the reward $\tilde{R}$ admits an equilibrium $\tilde{F}^*$ as described Theorem 1, and inspection of the formula shows that $\tilde{F}^*$ differs from the equilibrium F^* corresponding to R. More generally, $\tilde{F}^*$ is different for any two choices of β. To prove the claim, we argue that $\tilde{F}^*$ is also an equilibrium for R. If all players use the same stopping distribution, their value functions are the same under both rewards because achieving the last rank is a null set for any player. However, the rewards differ in the analysis of unilateral deviations: the inequality $\tilde{R}(1) > R(1)$ implies that if a player is not incentivized to deviate under $\tilde{R}$, the same holds under R. In particular, $\tilde{F}^*$ is also an equilibrium under R, proving the claim. Conversely, F^* need not be an equilibrium under $\tilde{R}$, as can be seen from our uniqueness result for $\beta = R(1-)$.

Remark 3. The framework of Sun [47] allows for the rigorous construction of a continuum of (almost everywhere (a.e.)) independent processes satisfying an exact law of large numbers. A short summary of the pertinent results can be found, for example, in Nutz [40, section 3]. As an alternative to explicitly formulating the game with a continuum of players, one can also directly analyze the problem of a "representative" player facing a distribution, as is sometimes done in the literature on mean field games—this corresponds to taking the exact law of large numbers as a given.

3. Proof of Theorem 1

We first show by a direct verification argument that the stated distribution is indeed an equilibrium. The proof of uniqueness occupies the remainder of the section.

Proof of Theorem 1—Existence. In view of

$$\begin{aligned}\int h dF^* &= \int_0^{h(\infty)} (1 - F^* \circ h^{-1}(w))dw = \int_0^{\frac{g(1)-g(0)}{\bar{R}-g(0)}} (1 - g^{-1}(g(0) + (\bar{R} - g(0))w))dw \\ &= \frac{1}{\bar{R} - g(0)} \int_{g(0)}^{g(1)} (1 - g^{-1}(z))dz = \frac{\bar{R} - g(0)}{\bar{R} - g(0)} = 1,\end{aligned}$$

Lemma 1 yields that $F^* \in \mathcal{F}$. To see that F^* is an equilibrium, fix some player i and suppose that all other players stop according to F^*. Using the property $g(g^{-1}(z)) \le z$ of the right-continuous inverse of the left-continuous function $g(y) = R(1-y)$, we have

$$\xi^{F^*}(x) \le g(F^*(x)) \le [R(1) + (\bar{R} - R(1))h(x)] \wedge R(0) =: \varphi(x).$$

By Itô's formula and Jensen's inequality, $\varphi(X_t)$ is a bounded supermartingale. Hence, optional sampling implies that for any finite stopping time τ,

$$E[\xi^{F^*}(X_\tau)] \le E[\varphi(X_\tau)] \le \varphi(x_0) = \bar{R}.$$

On the other hand, player i can attain $\bar{R}$ by choosing F^*, by symmetry. This shows that F^* is optimal for player i and hence that F^* is an equilibrium. □

3.1. Relating Constant Rewards to Atoms and Jumps in Reward to Gaps in Support

We first relate atoms in equilibrium distributions to intervals of constancy of the reward (and hence of g). Technical details aside, the message is that, in equilibrium, equal pay must correspond to equal performance: ranks with the same reward are occupied by players that stop at the same level, and ranks with different rewards are occupied by players that stop at different levels.

We do not yet impose the continuity properties of R, which will allow us to prove that they are important for the existence of equilibria. Instead, R is any decreasing function in this subsection, which, of course, implies that its discontinuities are of the jump type.

Lemma 2. *Let $F \in \mathcal{F}$ be a mean field equilibrium. If F has an atom at $x_1 \in \mathbb{R}_+$, then g is constant on $(F(x_1-), F(x_1)]$. As a result, we have $\xi^F(x) = g(F(x))$ for all $x \in \mathbb{R}_+$.*

Proof. Set $y_0 := F(x_1-)$ and $y_1 := F(x_1)$. Let ν be the measure associated with F. Consider for each $\varepsilon > 0$ the perturbed measure

$$\nu_\varepsilon := \lambda_\varepsilon(\nu + (y_1 - y_0)(\delta_{x_1+\varepsilon} - \delta_{x_1})) + (1 - \lambda_\varepsilon)\delta_0,$$

where $\lambda_\varepsilon \in (0,1)$ is chosen so that

$$\int h(x)d\nu_\varepsilon(x) = \lambda_\varepsilon\left(\int h(x)d\nu(x) + (y_1 - y_0)[h(x_1+\varepsilon) - h(x_1)]\right) = \int h(x)d\nu(x).$$

This ensures that $\nu_\varepsilon \in \mathcal{F}$. Suppose that g is not constant on $(y_0, y_1]$; then

$$\xi^F(x_1+\varepsilon) - \xi^F(x_1) \ge g(y_1) - \frac{1}{y_1 - y_0}\int_{y_0}^{y_1} g(y)dy =: \eta > 0.$$

where η is clearly independent of ε. This implies

$$\begin{aligned}
&\int \xi^F(x)d\nu_\varepsilon(x) - \int \xi^F(x)d\nu(x) \\
&= \lambda_\varepsilon\left(\int \xi^F(x)d\nu(x) + (y_1 - y_0)[\xi^F(x_1+\varepsilon) - \xi^F(x_1)]\right) + (1-\lambda_\varepsilon)\xi(0) - \int \xi^F(x)d\nu(x) \\
&= (1-\lambda_\varepsilon)\left(\xi^F(0) - \int \xi^F(x)d\nu(x)\right) + \lambda_\varepsilon(y_1-y_0)[\xi^F(x_1+\varepsilon) - \xi^F(x_1)] \\
&\geq (1-\lambda_\varepsilon)\left(\xi^F(0) - \int \xi^F(x)d\nu(x)\right) + \lambda_\varepsilon(y_1-y_0)\eta.
\end{aligned}$$

Using $\lim_{\varepsilon\to 0+}\lambda_\varepsilon = 1$, $y_1 - y_0 > 0$ and $\eta > 0$, we obtain that $\int \xi^F dF_\varepsilon - \int \xi^F dF > 0$ for ε sufficiently small, contradicting the assumption that F is an equilibrium. Finally, if g is constant on $(y_0, y_1]$, it is clear that $\xi^F(x_1) = g(y_1) = g(F(x_1))$. □

As g is increasing, each level set $\{g = c\}$ is an interval. If the interval has positive length, we say that g has a *flat segment* at level c. In all that follows, we denote by $F^{-1}(y) = \inf\{x : F(x) \geq y\}$ the *left*-continuous inverse (or quantile function) of F.

Lemma 3. *Suppose g has a flat segment at level c, so that $y_1 = \inf\{y \in [0,1] : g(y) = c\}$ and $y_2 = \sup\{y \in [0,1] : g(y) = c\}$ satisfy $y_1 < y_2$. Suppose $F \in \mathcal{F}$ is a mean field equilibrium; define $x_1^+ = F^{-1}(y_1+)$ and $x_2 = F^{-1}(y_2)$. Then we must have $x_1^+ = x_2$. Moreover, $F(x_1^+-) = y_1$ and $F(x_1^+) = y_2$.*

Proof. Let ν be the measure associated with F and $x_1^\varepsilon = F^{-1}(y_1 + \varepsilon)$. Suppose, on the contrary, that $x_2 > x_1^\varepsilon$ for some $\varepsilon \in (0, y_2 - y_1)$. Then $y_1 < F(x_1^\varepsilon) < y_2$, and thus $\xi^F(x_1^\varepsilon) = c$. Consider the measure

$$\zeta := \nu|_{(x_1^\varepsilon, x_2)} + (y_2 - F(x_2-))\delta_{x_2}$$

with total mass $|\zeta| = y_2 - F(x_1^\varepsilon) > 0$. We distinguish two cases:

(i) *Case* $y_2 < 1$: In this case, let $\beta \in (y_2, 1)$ and $x_\beta := F^{-1}(\beta) \in [x_2, \infty)$. For some $\lambda \in [0,1]$ to be determined later, define the measure

$$\nu_\lambda = \nu - \zeta + |\zeta|(\lambda\delta_{x_1^\varepsilon} + (1-\lambda)\delta_{x_\beta}).$$

In other words, ν_λ is obtained from ν by removing all mass on (x_1^ε, x_2), plus possibly an additional atom at x_2, so that the total removed mass is $y_2 - F(x_1^\varepsilon)$ and then moving this mass to atoms at x_1^ε and x_β according to weights λ and $1-\lambda$. Clearly, ν_λ is a probability measure supported on $\mathbb{R}_+$, and we have

$$\begin{aligned}
\int h d\nu_\lambda - \int h d\nu &= -\int h d\zeta + |\zeta|(\lambda h(x_1^\varepsilon) + (1-\lambda)h(x_\beta)) \\
&= \lambda \int \left[h(x_1^\varepsilon) - h\right] d\zeta + (1-\lambda)\int \left[h(x_\beta) - h\right] d\zeta.
\end{aligned}$$

In view of $\int [h(x_1^\varepsilon) - h] d\zeta < 0$ and $\int [h(x_\beta) - h] d\zeta \geq 0$, we can choose $\lambda \in [0,1)$ so that $\int h d\nu_\lambda = \int h d\nu$. We then have $\nu_\lambda \in \mathcal{F}$ by Lemma 1. Using the optimality of F and the fact that $\xi^F(x) = c$ for all $x \in [x_1^\varepsilon, x_2)$, we obtain

$$\begin{aligned}
0 \geq \int \xi^F d\nu_\lambda - \int \xi^F d\nu &= \lambda \int \left[\xi^F(x_1^\varepsilon) - \xi^F\right] d\zeta + (1-\lambda)\int \left[\xi^F(x_\beta) - \xi^F\right] d\zeta \\
&= \lambda(c - \xi^F(x_2))\zeta\{x_2\} + (1-\lambda)\left[\xi^F(x_\beta) - c\right]\nu(x_1^\varepsilon, x_2) + (1-\lambda)(\xi^F(x_\beta) - \xi^F(x_2))\zeta\{x_2\}.
\end{aligned}$$

Lemma 2 rules out the possibility that $F(x_2-) < y_2 < F(x_2)$, so we must be in one of the following two subcases:

(a) *Case* $F(x_2-) = y_2$: In this case, $\zeta\{x_2\} = 0$, and

$$\int \xi^F d(\nu_\lambda - \nu) = (1-\lambda)\left[\xi^F(x_\beta) - c\right]|\zeta|.$$

Using $|\zeta| > 0$ and $\xi^F(x_\beta) \geq g(\beta) > c$ and $\lambda < 1$, we obtain the contradiction that $\int \xi^F d(\nu_\lambda - \nu) > 0$.

(b) *Case* $F(x_2-) < y_2 = F(x_2)$: In this case, Lemma 2 implies $\xi^F(x_2) = g(F(x_2)) = g(F(x_2-)+) = c$, and we reach the same contradiction:

$$\int \xi^F d(\nu_\lambda - \nu) = (1-\lambda)\big[\xi^F(x_\beta) - c\big]\nu(x_1^\varepsilon, x_2) + (1-\lambda)(\xi^F(x_\beta) - c)\zeta\{x_2\}$$
$$= (1-\lambda)\big[\xi^F(x_\beta) - c\big]|\zeta| > 0.$$

(ii) *Case* $y_2 = 1$: Then $y_1 > 0$ as g is not a.e. constant. Let $x_1 := F^{-1}(y_1) \le x_1^\varepsilon$. We note that $\nu[0, x_1] \ge y_1 > 0$ and consider the measure

$$\nu_\lambda = \nu - \lambda\zeta - (1-\lambda)\nu|_{[0,x_1]} + \{\lambda|\zeta| + (1-\lambda)\nu[0,x_1]\}\delta_{x_1^\varepsilon},$$

where $\lambda \in [0,1)$ is again chosen so that $\int h d\nu_\lambda = \int h d\nu$. We have

$$0 \ge \int \xi^F d\nu_\lambda - \int \xi^F d\nu = \lambda \int \big[\xi^F(x_1^\varepsilon) - \xi^F\big] d\zeta + (1-\lambda)\int \big[\xi^F(x_1^\varepsilon) - \xi^F\big] d\nu|_{[0,x_1]}$$
$$= \lambda(c - \xi^F(x_2))\zeta\{x_2\} + (1-\lambda)\int \big[c - \xi^F\big] d\nu|_{[0,x_1]}.$$

Similarly as in case (i), one can show that either $\zeta\{x_2\} = 0$ or $\xi^F(x_2) = c$, both of which lead to

$$0 \ge (1-\lambda)\int \big[c - \xi^F\big] d\nu|_{[0,x_1]} \ge 0,$$

and thus $\int [c - \xi^F(x)] d\nu|_{[0,x_1]} = 0$. It follows that $\xi^F(x) = c$ for ν-a.e. $x \in [0, x_1]$. On the other hand, the definitions of x_1 and y_1 imply that $F(x) < y_1$ and $\xi^F(x) < c$ for all $x < x_1$. So it must hold that either $x_1 = 0$ or $\nu[0, x_1) = 0$. Both cases lead to $F(x_1-) = 0$ and $\nu\{x_1\} \ge y_1 > 0$. But $F(x_1-) = 0$ yields, by Lemma 2, that $\xi^F(x_1) = g(F(x_1-)+) = g(0+) < c$, whereas $\nu\{x_1\} > 0$ implies $\xi^F(x_1) = c$—a contradiction. This completes the proof that $F^{-1}(y_1+) = F^{-1}(y_2)$.

Finally, let $x_1^+ := F^{-1}(y_1+)$. Clearly, $F(x_1^+) = F(F^{-1}(y_2)) \ge y_2$. For any $x < x_1^+$ and $\varepsilon > 0$, we have $x < x_1^\varepsilon$ and $F(x) < y_1 + \varepsilon$. Passing to the limit then yields $F(x_1^+-) \le y_1 < y_2$. As g is constant on $(F(x_1^+-), F(x_1^+)]$, we must have $F(x_1^+-) = y_1$ and $F(x_1^+) = y_2$. □

Remark 4. If in Lemma 3 we also have $g(y_1) = c$, then the proof goes through with x_1^ε replaced by $x_1 = F^{-1}(y_1) \vee 0$. (The reason for using x_1^ε is to have $\xi^F(x_1^\varepsilon) = g(F(x_1^\varepsilon)) = c$.) As a result, we have $x_1 = x_1^+ = x_2$. In particular, if $y_1 = 0$, then $F(0) = y_2$.

Remark 5. Lemmas 2 and 3 imply that for the existence of a mean field equilibrium, it is necessary that g be left continuous at any level for which it contains a flat segment. In particular, if the reward function R is piecewise constant, *a mean field equilibrium can only exist if R is right continuous.*

The feasibility constraint yields one equation to pin down the equilibrium. The best way to illustrate this is to go through a particular case of Theorem 1 where the reward function is of the cutoff type. That is the purpose of the next proposition—here, the feasibility constraint and the preceding results on atoms are already sufficient to uniquely identify the equilibrium.

Proposition 1. *Let $R(r) = (1/\alpha)1_{[0,\alpha)}(r)$ for some $\alpha \in (0,1)$. Then the unique mean field equilibrium is given by the two-point distribution $\nu_\alpha := (1-\alpha)\delta_0 + \alpha\delta_{x_1}$, where x_1 is the unique point in (x_0, ∞) with $h(x_1) = 1/\alpha$.*

Proof. We first derive a necessary condition for $F \in \mathcal{F}$ to be an equilibrium. Let $x_\alpha = F^{-1}((1-\alpha)+)$. By Lemma 3 and Remark 4, we have $F(0) = 1-\alpha$, $F(x_\alpha-) = 1-\alpha$ and $F(x_\alpha) = 1$. That is, the measure associated with F must take the form $\nu = (1-\alpha)\delta_0 + \alpha\delta_{x_\alpha}$. To determine x_α, we first note that $x_\alpha \le x_1$, for otherwise, $\int h d\nu > \int h d\nu_\alpha = 1$, contradicting $\nu \in \mathcal{F}$. Suppose $x_\alpha < x_1$ (which is only feasible if $\mu \le 0$); then there exists $\beta \in (\alpha, 1)$ such that $\nu' = (1-\beta)\delta_0 + \beta\delta_{x_\alpha}$ is feasible. In view of $\xi^F(x_\alpha) = g(F(x_\alpha)) = g(1) > g(1-\alpha) = \xi^F(0)$, the distribution ν' is strictly preferable to ν when the other players choose ν. As a result, $x_\alpha = x_1$, which uniquely identifies F. To check that F is indeed an equilibrium, we argue as in the beginning of Section 3. □

This proof does not generalize to piecewise constant reward functions with multiple jumps: although the feasibility constraint still yields one equation, there are now multiple unknowns (the locations of the atoms). To determine mean field equilibria for general reward functions, it is necessary to analyze the effect of jumps in some detail. Let

$$J(g) := \{y \in (0,1) : g(y-) < g(y+)\}$$

be the set of interior jump points of g. The next lemma says that any jump of g—or, equivalently, of R—induces a flat segment in any equilibrium distribution. (The reasoning in Remark 2 shows that this assertion fails at $y = 0$; whence the definition of $J(g)$ considers only interior jumps.)

Lemma 4. *Let F be a mean field equilibrium. For each $y \in J(g)$, the interval $\{x \geq 0 : F(x) = y\}$ has positive length.*

Proof. Let $y_1 \in J(g)$; then $x_1 := F^{-1}(y_1) \in [0,\infty)$ as $y_1 \in (0,1)$. Suppose, for contradiction, that $\{x \geq 0 : F(x) = y_1\}$ has zero length; then $F(x) > y_1$ for all $x > x_1$. Let ν be the measure associated with F. In the remainder of the proof we construct a feasible distribution ν' that is strictly better than ν. By Lemma 2, we have either $F(x_1-) = y_1$ or $F(x_1-) < y_1 = F(x_1)$.

(i) *Case* $F(x_1-) = y_1$: In this case, $x_1 > 0$, and F is nonconstant in any left neighborhood of x_1. Fix $\gamma \in (0, g(y_1+) - g(y_1-))$ and observe that

$$\lim_{\varepsilon \to 0+} \xi^F(x_1 - 2\varepsilon) = g(y_1-), \qquad \lim_{\varepsilon \to 0+} \frac{h(x_1-\varepsilon) - h(x_1 - 2\varepsilon)}{h(x_1) - h(x_1 - 2\varepsilon)} = \frac{1}{2}.$$

We can thus find $\varepsilon > 0$ such that

$$\xi^F(x_1 - 2\varepsilon) > g(y_1-) - \gamma/2 \tag{4}$$

and

$$\frac{h(x_1-\varepsilon) - h(x_1 - 2\varepsilon)}{h(x_1) - h(x_1 - 2\varepsilon)} > \frac{\gamma}{g(y_1+) - g(y_1-) + \gamma}.$$

The measure $\zeta := \nu|_{(x_1-\varepsilon, x_1)}$ has mass $|\zeta| > 0$. Consider the probability measure

$$\nu' := \nu - \zeta + |\zeta|\big(\lambda\delta_{x_2} + (1-\lambda)\delta_{x_1-2\varepsilon}\big),$$

where $x_2 \in (x_1, \infty)$ is chosen to satisfy

$$\frac{h(x_1-\varepsilon) - h(x_1 - 2\varepsilon)}{h(x_2) - h(x_1 - 2\varepsilon)} > \frac{\gamma/2}{g(y_1+) - g(y_1) + \gamma/2} \tag{5}$$

and

$$\lambda := \frac{\frac{1}{|\zeta|}\int h d\zeta - h(x_1 - 2\varepsilon)}{h(x_2) - h(x_1 - 2\varepsilon)} \in \left(\frac{h_1(x_1-\varepsilon) - h(x_1 - 2\varepsilon)}{h(x_2) - h(x_1 - 2\varepsilon)}, 1\right). \tag{6}$$

It is easy to check that $\int h d(\nu' - \nu) = 0$; hence $\nu' \in \mathcal{F}$, by Lemma 1. To see that ν' is strictly better than ν, we use (4)–(6) and $F(x_2) > y_1$:

$$\begin{aligned}
\int \xi^F d(\nu' - \nu) &= \lambda \int (\xi^F(x_2) - \xi^F) d\zeta + (1-\lambda) \int (\xi^F(x_1 - 2\varepsilon) - \xi^F) d\zeta \\
&\geq \lambda(g(y_1+) - g(y_1-))|\zeta| + (1-\lambda)(\xi^F(x_1 - 2\varepsilon) - g(y_1-))|\zeta| \\
&> |\zeta|\left(\lambda[g(y_1+) - g(y_1-)] - (1-\lambda)\frac{\gamma}{2}\right) = |\zeta|\left(\lambda[g(y_1+) - g(y_1-) + \gamma/2] - \frac{\gamma}{2}\right) > 0.
\end{aligned}$$

(ii) *Case* $F(x_1-) < y_1 = F(x_1)$: In this case, $\{x \geq 0 : F(x) = y_1\}$ having zero length implies that F is nonconstant in any right neighborhood of x_1. Moreover, Lemma 2 implies that $g(y_1-) = g(y_1)$.
Fix $\gamma \in (0, (y_1 - F(x_1-))(g(y_1+) - g(y_1)))$ and $\varepsilon > 0$ such that $\xi^F(x_1 + \varepsilon) < g(y_1+) + \gamma$. We define $\zeta := \nu|_{[x_1, x_1+\varepsilon)}$ and

$$\nu' := \nu - \zeta + |\zeta|\delta_{x_2},$$

where $x_2 > x_1$ is to be determined. Because h is strictly increasing and $\nu(x_1, x_1+\varepsilon) > 0$, we see that $|\zeta|h(x_1) < \int h d\zeta < |\zeta|h(x_1+\varepsilon)$, and consequently, there exists $x_2 \in (x_1, x_1+\varepsilon)$ such that

$$\int h d(\nu' - \nu) = |\zeta|h(x_2) - \int h d\zeta = 0.$$

For this choice of x_2, we have $\nu' \in \mathcal{F}$ by Lemma 1. Moreover, ν' is strictly better than ν:

$$\begin{aligned}\int \xi^F d(\nu'-\nu) &= (y_1 - F(x_1-))[\xi^F(x_2) - \xi^F(x_1)] + \int (\xi^F(x_2) - \xi^F) d\nu|_{(x_1,x_1+\varepsilon)} \\ &\geq (y_1 - F(x_1-))(g(y_1+) - g(y_1)) + (g(y_1+) - \xi^F(x_1+\varepsilon))\nu(x_1, x_1+\varepsilon) \\ &\geq (y_1 - F(x_1-))(g(y_1+) - g(y_1)) - \gamma > 0\end{aligned}$$

by the choice of γ. □

3.2. Characterizing the Equilibrium

From now on, we shall work under the assumption that $\mu < \bar{\mu}_\infty$ and R is right continuous with $R(0) > R(1-) = R(1)$.

The general idea of the uniqueness argument is to analyze a minimization problem: in equilibrium, the opposing players act such as to minimize the value function of a given representative player, subject to the constraint that the opponents act symmetrically. This turns out to be substantially more involved than in the n-player case, as a result of the possible presence of atoms in the equilibrium distribution and the noninvertibility of the function g.

Let u^F be the value function of a representative player if the other players use $F \in \mathcal{F}$. Dynamic programming and optimal stopping theory yield

$$u^F(x_0) = \sup_{\tau<\infty} E[\xi^F(X_\tau)] = (\xi^F \circ h^{-1})^{\mathrm{conc}}(1) \leq (g \circ F \circ h^{-1})^{\mathrm{conc}}(1),$$

where h is the scale function (2) with normalization $h(x_0) = 1$, and "conc" denotes the concave envelope on $\mathbb{R}_+$. The last inequality is due to a possible breaking of ties (see (1)). Lemma 2 shows that the inequality must be an equality if F is a mean field equilibrium, even if ties do occur. On the other hand, if F^* is an equilibrium, we must have

$$\bar{R} = u^{F^*}(x_0) = \min_{F\in\mathcal{F}} u^F(x_0). \tag{7}$$

Indeed, given arbitrary $F \in \mathcal{F}$, a representative player can achieve $\bar{R}$ by also choosing F, and in equilibrium, this is the best possible performance, by symmetry. Combining the two arguments, any mean field equilibrium F^* must satisfy

$$(g \circ F^* \circ h^{-1})^{\mathrm{conc}}(1) = u^{F^*}(x_0) = \min_{F\in\mathcal{F}} u^F(x_0) \leq \min_{F\in\mathcal{F}} (g \circ F \circ h^{-1})^{\mathrm{conc}}(1).$$

That is,

$$F^* \in \arg\min_{F\in\mathcal{F}} \Phi(F)^{\mathrm{conc}}(1), \quad \text{where} \quad \Phi(F) := g \circ F \circ h^{-1}.$$

We also write $\Phi^{-1}(\phi) := g^{-1} \circ \phi \circ h$. We recall that g^{-1} denotes the right-continuous inverse of g; in particular, $g^{-1}(g(y)) \geq y$ and $g(g^{-1}(z)) \leq z$. Similarly, $\Phi^{-1}(\Phi(F)) \geq F$ and $\Phi(\Phi^{-1}(\phi)) \leq \phi$. Finally, we define

$$\begin{aligned}\bar{w}_F &:= \inf\{w \in [0, h(\infty)] : \Phi(F)(w) = g(1)\} \leq h(\infty), \\ \bar{y} &:= \inf\{y \in [0,1] : g(y) = g(1)\}, \\ \mathcal{F}' &:= \{F \in \mathcal{F} : \bar{w}_F > 1, \text{ and } \bar{w}_F < h(\infty) \text{ in the case } \bar{y} < 1\}.\end{aligned}$$

Lemma 5. *If F is a mean field equilibrium, then $F \in \mathcal{F}'$ and $\bar{w}_F = \inf\{w \in [0, h(\infty)] : F \circ h^{-1}(w) = 1\}$. In particular, $F \circ h^{-1}(\bar{w}_F) = 1$.*

Proof. We first show $F \in \mathcal{F}'$. Suppose $\bar{w}_F \leq 1$. Then

$$\Phi(F)^{\mathrm{conc}}(1) \geq \Phi(F)^{\mathrm{conc}}(\bar{w}_F) = \Phi(F)^{\mathrm{conc}}(\bar{w}_F+) = g(1).$$

Consider the distribution $G(x) = \lambda 1_{[0,x_0+\varepsilon)}(x) + 1_{[x_0+\varepsilon,\infty)}(x)$, where $\varepsilon > 0$ and $\lambda \in (0,\bar{y})$ are chosen so that $\int h dG = (1-\lambda)h(x_0+\varepsilon) = 1$. We have $\Phi(G) = g(\lambda)1_{[0,h(x_0+\varepsilon))} + g(1)1_{[h(x_0+\varepsilon),\infty)}$. The concave hull of this function is readily determined, and in view of $g(\lambda) < g(1)$, we arrive at $\Phi(G)^{\mathrm{conc}}(1) < g(1) \leq \Phi(F)^{\mathrm{conc}}(1)$, contradicting the optimality of F.

Suppose $\bar{y} < 1$ and $\bar{w}_F = h(\infty)$. Then for all $x < \infty$, we have $h(x) < h(\infty) = \bar{w}_F$, which implies $g(1) > \Phi(F)(h(x)) = g(F(x))$. But then $F(x) \leq \bar{y} < 1$ for all $x \in \mathbb{R}$, contradicting that F is the cdf of a probability measure on $\mathbb{R}$.

We next show $F\circ h^{-1}(\bar{w}_F)=1$. This is trivial if $\bar{w}_F=h(\infty)$, so we may assume that $\bar{w}_F<h(\infty)$. For any $w>\bar{w}_F$, we have $g(F\circ h^{-1}(w))=\Phi(F)(w)=g(1)$, which implies the following:

(i) If $\bar{y}=1$, then $F\circ h^{-1}(w)=1$. Because $w>\bar{w}_F$ is arbitrary, $F\circ h^{-1}(\bar{w}_F)=1$ by right continuity.

(ii) If $\bar{y}<1$ and $g(\bar{y})<g(1)$, then $F\circ h^{-1}(w)>\bar{y}$. In this case, $h^{-1}(w)\ge F^{-1}(\bar{y}+)$ for all $w>\bar{w}_F$, which further yields $h^{-1}(\bar{w}_F)\ge F^{-1}(\bar{y}+)$. By Lemma 3, F jumps from $\bar{y}$ to 1 at $F^{-1}(\bar{y}+)$. It follows that $F\circ h^{-1}(\bar{w}_F)\ge F(F^{-1}(\bar{y}+))=1$.

(iii) If $\bar{y}<1$ and $g(\bar{y})=g(1)$, then $F\circ h^{-1}(w)\ge\bar{y}$. In this case, we use Remark 4 to obtain $h^{-1}(w)\ge F^{-1}(\bar{y})=F^{-1}(\bar{y}+)$, and thus $h^{-1}(\bar{w}_F)\ge F^{-1}(\bar{y}+)$. We obtain the same conclusion as in (ii).

Finally, for $w<\bar{w}_F$, $\Phi(F)(w)<g(1)$ implies $F\circ h^{-1}(w)<1$. □

Lemma 6. *Let $F\in\mathcal{F}'$. Suppose there exists an increasing concave function $\phi\ge\Phi(F)$ on $[0,h(\infty)]$ satisfying $\phi(1)\le\Phi(F)^{\mathrm{conc}}(1)$ and*

$$\int_0^{h(\infty)}(1-g^{-1}\circ\phi(w))dw<1.$$

Then there exists $F'\in\mathcal{F}$ such that $\Phi(F')^{\mathrm{conc}}(1)<\Phi(F)^{\mathrm{conc}}(1)$, and consequently, F cannot be a mean field equilibrium.

Proof. Let ϕ be as stated. Note that $\bar{w}_F>1$ implies $\Phi(F)(1+)<g(1)$, which further yields $\Phi(F)^{\mathrm{conc}}(1)<g(1)=\Phi(F)(\bar{w}_F+)\le\phi(\bar{w}_F+)=\phi(\bar{w}_F)$. Let $\bar{w}_\phi:=\inf\{w\ge0:\phi(w)=g(1)\}$. Because $\phi(1)<g(1)$ and $\phi(\bar{w}_F+)=g(1)$, we know $1<\bar{w}_\phi\le\bar{w}_F$. Consider four cases.

(i) $\mu\le0$ *and* $\bar{y}=1$: In this case, $h(\infty)=\infty$ and g^{-1} is continuous at $g(1)$. Choose $\varepsilon\in(0,1)$ such that $\phi(\varepsilon)<\phi(1)$. Such an ε exists: as ϕ is increasing and concave, it must be strictly increasing before reaching $g(1)$. Let $\phi_\varepsilon(w):=\phi(\varepsilon w)$. Then ϕ_ε is concave on $\mathbb{R}_+$ and satisfies $\phi_\varepsilon(1)<\phi(1)$. Next, define $F_\lambda:=\Phi^{-1}(\lambda\phi+(1-\lambda)\phi_\varepsilon)$. One can check that F_λ is right continuous and satisfies $F_\lambda(\infty)=1$. We also have that for $\lambda\in[0,1)$, $\Phi(F_\lambda)^{\mathrm{conc}}(1)\le(\lambda\phi+(1-\lambda)\phi_\varepsilon)^{\mathrm{conc}}(1)=\lambda\phi(1)+(1-\lambda)\phi_\varepsilon(1)<\phi(1)\le\Phi(F)^{\mathrm{conc}}(1)$, showing that F_λ is strictly better than F. To reach the desired contradiction, it remains to show the feasibility of F_λ for λ sufficiently close to 1. We have

$$\int hdF_\lambda=\int_0^\infty(1-F_\lambda\circ h^{-1})(w)dw=\int_0^\infty(1-g^{-1}\circ(\lambda\phi+(1-\lambda)\phi_\varepsilon))(w)dw.$$

As g^{-1} is monotone, it has at most countably many points of discontinuity, and $\bar{y}=1$ implies that $g(1)$ is not one of them. For any $z<g(1)$, the set $\{w\ge0:\phi(w)=z\}$ has zero Lebesgue measure because ϕ is strictly increasing before reaching $g(1)$. It follows that as $\lambda\to1$, the integrand converges a.e. to $1-g^{-1}\circ\phi$. Using $0\le 1-g^{-1}\circ(\lambda\phi+(1-\lambda)\phi_\varepsilon)\le1-g^{-1}\circ\phi_\varepsilon\le1-g^{-1}\circ\Phi(F)(\varepsilon\cdot id)\le1-F\circ h^{-1}(\varepsilon\cdot id)$ and $\int_0^\infty(1-F\circ h^{-1}(\varepsilon w))dw=\varepsilon^{-1}\int hdF<\infty$, dominated convergence yields that

$$\lim_{\lambda\to1}\int hdF_\lambda=\int_0^\infty(1-g^{-1}\circ\phi)(w)dw<1.$$

By Lemma 1, this shows that F_λ is feasible for λ sufficiently close to 1.

(ii) $\mu\le0$ *and* $\bar{y}<1$: In this case, $1<\bar{w}_\phi\le\bar{w}_F<h(\infty)=\infty$. Choose $\varepsilon>0$ such that $\bar{w}_\phi+\varepsilon<h(\infty)$ and $\int_0^{\bar{w}_\phi}(1-g^{-1}\circ\phi)(w)dw+(1-\bar{y})\varepsilon<1$. Let ϕ'_ε denote the line connecting $(0,\phi(0))$ and $(\bar{w}_\phi+\varepsilon,g(1))$ and capped at level $g(1)$:

$$\phi'_\varepsilon(w)=\left(\phi(0)+\frac{g(1)-\phi(0)}{\bar{w}_\phi+\varepsilon}w\right)\wedge g(1).$$

Then ϕ'_ε is concave on $\mathbb{R}_+$ and satisfies $\phi'_\varepsilon(1)<\phi(1)$. As in the previous case, we define $F'_\lambda:=\Phi^{-1}(\lambda\phi+(1-\lambda)\phi'_\varepsilon)$. Then F'_λ is a cdf supported on $\mathbb{R}_+$, which satisfies $\Phi(F'_\lambda)^{\mathrm{conc}}(1)<\Phi(F)^{\mathrm{conc}}(1)$ for all $\lambda\in[0,1)$. To check the feasibility of F'_λ for λ close to 1, we write

$$\int hdF'_\lambda=\int_0^{\bar{w}_\phi}(1-g^{-1}(\lambda\phi(w)+(1-\lambda)\phi_\varepsilon(w)))dw+\int_{\bar{w}_\phi}^{\bar{w}_\phi+\varepsilon}(1-g^{-1}(\lambda g(1)+(1-\lambda)\phi_\varepsilon(w)))dw.$$

Using that ϕ is strictly increasing on $[0,\bar{w}_\phi]$, we obtain by bounded convergence that $\lim_{\lambda\to 1}\int hdF'_\lambda = \int_0^{\bar{w}_\phi}(1-g^{-1}\circ\phi)(w)dw + (1-\bar{y})\varepsilon < 1$.

(iii) $0<\mu<\bar{\mu}_\infty$ *and* $\bar{y}=1$: In this case, $\bar{w}_\phi \le h(\infty) < \infty$, and g^{-1} is continuous at $g(1)$. Let ℓ be the line segment connecting $(0,g(0))$ and $(h(\infty),g(1))$. We have

$$\int hd\Phi^{-1}(\ell) = \int_0^{h(\infty)}(1-\Phi^{-1}(\ell)\circ h^{-1}(w))dw = \int_0^{h(\infty)}(1-g^{-1}\circ\ell(w))dw$$
$$= \frac{h(\infty)}{g(1)-g(0)}\int_{g(0)}^{g(1)}(1-g^{-1}(y))dy = \frac{h(\infty)}{g(1)-g(0)}\left(\int_0^1 g(y)dy - g(0)\right) = \frac{h(\infty)}{h(\bar{x})} > 1.$$

Because $\phi(0)\ge g(0)=\ell(0)$ and $\phi(h(\infty)) = g(1) = \ell(h(\infty))$, by concavity, either $\phi>\ell$ on $(0,h(\infty))$ or $\phi=\ell$. The latter case is impossible, as $F\le\Phi^{-1}(\Phi(F))\le\Phi^{-1}(\phi)=\Phi^{-1}(\ell)$ would imply $F\notin\mathcal{F}$. Set $F''_\lambda := \Phi^{-1}(\lambda\phi+(1-\lambda)\ell)$. We again have $\Phi(F''_\lambda)^{\text{conc}}(1)\le\lambda\phi(1)+(1-\lambda)\ell(1)<\phi(1)\le\Phi(F)^{\text{conc}}(1)$ if $\lambda\in[0,1)$. Let

$$I(\lambda) := \int hdF''_\lambda = \int_0^{h(\infty)}(1-g^{-1}(\lambda\phi(w)+(1-\lambda)\ell(w)))dw.$$

Using the continuity of g^{-1} at $g(1)$, the strict monotonicity of $\lambda\phi+(1-\lambda)\ell$ before reaching $g(1)$, and the bounded convergence theorem, we deduce that $I(\cdot)$ is continuous on (0, 1), satisfying $I(1-) = \int_0^{h(\infty)}(1-g^{-1}\circ\phi(w))dw < 1$ and

$$I(0+) = \int_0^{h(\infty)}(1-g^{-1}\circ\ell)(w)dw = \int_0^{h(\infty)}(1-\Phi^{-1}(\ell)\circ h^{-1})(w)dw = \int hd\Phi^{-1}(\ell) > 1.$$

We may thus choose $\lambda_0\in(0,1)$ such that $I(\lambda_0)=1$. Then $F''_{\lambda_0}\in\mathcal{F}$, and the contradiction is complete.

(iv) $0<\mu<\bar{\mu}_\infty$ *and* $\bar{y}<1$: In this case, $\bar{w}_\phi<h(\infty)<\infty$. Let ϕ'_ε and F'_λ be constructed as in case (ii) with ε satisfying $\bar{w}_\phi+2\varepsilon<h(\infty)$. Define

$$G(x) := \begin{cases}\gamma F'_\lambda(x) & x<h^{-1}(\bar{w}_\phi+2\varepsilon),\\ 1 & x\ge h^{-1}(\bar{w}_\phi+2\varepsilon),\end{cases}$$

for some $\gamma\in(0,1)$ to be determined. We have $F'_\lambda(h^{-1}(\bar{w}_\phi+2\varepsilon)-)\ge F'_\lambda(h^{-1}(\bar{w}_\phi+\varepsilon)) = g^{-1}(\lambda\phi(\bar{w}_\phi+\varepsilon)+(1-\lambda)\phi'_\varepsilon(\bar{w}_\phi+\varepsilon))=1$ and

$$\int hdG = \gamma\int hdF'_\lambda + (1-\gamma)(\bar{w}_\phi+2\varepsilon).$$

In view of $\int hdF'_\lambda<1$ and $\bar{w}_\phi+2\varepsilon>1$, we can find $\gamma\in(0,1)$ such that $\int hdG=1$, and then G is feasible. We arrive at the desired contradiction after noting that $\Phi(G)^{\text{conc}}(1)\le\Phi(F'_\lambda)^{\text{conc}}(1)<\Phi(F)^{\text{conc}}(1)$. □

Lemma 7. *Let F be a mean field equilibrium. Define*

$$A = \bigcup_{y\notin J(g)} A_y,\quad \text{where}\quad A_y = \{w\in[0,h(\infty)]\cap\mathbb{R} : F\circ h^{-1}(w) = y\},$$

as well as $w_y = \inf A_y$. Then $w_y<\bar{w}_F$ for all $y\in J(g)$. Moreover, there exists a strictly increasing affine function $\ell_1\ge\Phi(F)$ satisfying $\ell_1(1)=\Phi(F)^{\text{conc}}(1)$ and

(i) *$\ell_1(\bar{w}_F)=g(1)$,*
(ii) *$\Phi(F)=\ell_1\wedge g(1)$ on A, and*
(iii) *$\Phi(F)(w_y)=\ell_1(w_y)$ for all $y\in J(g)$.*

Proof. Let F be an equilibrium; then $F\in\mathcal{F}'$ by Lemma 5. Define $\phi := \Phi(F)^{\text{conc}}$ and $\psi := \ell_1\wedge g(1)$, where ℓ_1 is an affine function passing through $(1,\phi(1))$ whose slope lies in the superdifferential of ϕ at $w=1$. We have $\ell_1\ge\psi\ge\phi\ge\Phi(F)$, $g^{-1}\circ\psi\ge g^{-1}\circ\Phi(F)\ge F\circ h^{-1}$ and $\psi(1)\le\ell_1(1)=\phi(1)=\Phi(F)^{\text{conc}}(1)$. By Lemma 5, $\Phi(F)(1+)<g(1)=\Phi(F)(\bar{w}_F+)\le\phi(\bar{w}_F)\le g(1)$, which implies that $\phi(1)<g(1)=\phi(\bar{w}_F)$. As ϕ is increasing and concave, it must be strictly increasing before reaching level $g(1)$. Consequently, ℓ_1 has positive slope. For any $y\in J(g)$, A_y has positive

length by Lemma 4. Because $F\circ h^{-1}=y<1$ on A_y, we must have $A_y\subseteq[0,\bar{w}_F)$ and $w_y<\bar{w}_F$. It remains to show properties (i)–(iii). Specifically, we show in what follows that if one of these properties does not hold, then $\int_0^{h(\infty)}(1-g^{-1}\circ\psi(w))dw<1$. Applying Lemma 6 with ψ being the increasing concave function, this contradicts that F is an equilibrium.

(i) Let $\bar{w}_1:=\ell_1^{-1}(g(1))$. Because $\ell_1\geq\Phi(F)$, we necessarily have $\bar{w}_1\leq\bar{w}_F$. Suppose $\bar{w}_1<\bar{w}_F$; then $F\circ h^{-1}<1$ in a right neighborhood of $\bar{w}_1$. Together with $g^{-1}\circ\psi\geq F\circ h^{-1}$, we obtain

$$\int_0^{h(\infty)}(1-g^{-1}\circ\psi(w))dw=\int_0^{\bar{w}_1}(1-g^{-1}\circ\psi(w))dw$$
$$\leq\int_0^{\bar{w}_1}(1-F\circ h^{-1}(w))dw<\int_0^{h(\infty)}(1-F\circ h^{-1}(w))dw=\int hdF\leq 1.$$

(ii) Suppose $\Phi(F)(w_0)<\psi(w_0)$ for some $w_0\in A$. As $F\circ h^{-1}(w_0)\notin J(g)$, we have $F\circ h^{-1}(w_0)\leq g^{-1}(\Phi(F)(w_0))<g^{-1}(\psi(w_0))$. By the right continuity of $F\circ h^{-1}$ and $g^{-1}\circ\psi$, it follows that $F\circ h^{-1}<g^{-1}\circ\psi$ in a right neighborhood of w_0. Thus

$$\int_0^{h(\infty)}(1-g^{-1}\circ\psi(w))dw<\int_0^{h(\infty)}(1-F\circ h^{-1}(w))dw=\int hdF\leq 1.$$

(iii) Let $y\in J(g)$. Suppose $\Phi(F)(w_y)<\ell_1(w_y)$. Let $w'_y:=\ell_1^{-1}(\Phi(F)(w_y))$. We have $w'_y<\ell_1^{-1}(\ell_1(w_y))=w_y<\bar{w}_F$ and $\ell_1(w'_y)=\Phi(F)(w_y)=g(y)<g(y+)\leq g(1)$. Define

$$F_1(x):=\begin{cases}F(x), & \text{if } x<h^{-1}(w'_y)\text{ or }x\geq h^{-1}(w_y),\\ F(h^{-1}(w_y))=y, & \text{if } h^{-1}(w'_y)\leq x<h^{-1}(w_y).\end{cases}$$

Clearly, $F_1\geq F$. For $x\in[h^{-1}(w'_y),h^{-1}(w_y))$, we have $h(x)\in[w'_y,w_y)$ and $F(x)=F\circ h^{-1}(h(x))<y=F_1(x)$ by the definition of w_y. It follows that $\int hdF_1<\int hdF\leq 1$. Now, observe that

$$\psi(w)\geq\Phi(F_1)(w)=\begin{cases}\Phi(F)(w) & \text{if } w<w'_y\text{ or }w\geq w_y,\\ g(y) & \text{if } w'_y\leq w<w_y.\end{cases}$$

This implies

$$\int_0^{h(\infty)}(1-g^{-1}\circ\psi(w))dw\leq\int_0^{h(\infty)}(1-g^{-1}\circ\Phi(F_1)(w))dw\leq\int_0^{h(\infty)}(1-F_1\circ h^{-1}(w))dw=\int hdF_1<1.\quad\square$$

Remark 6. When g is continuous, $A=[0,h(\infty)]\cap\mathbb{R}$, and Lemma 7 states that $\Phi(F)$ is affine before reaching level $g(1)$.

We can now complete the uniqueness argument.

Proof of Theorem 1—Uniqueness. Let F be any equilibrium. By Lemma 2, $\Phi(F)(0)=g(F(0))=g(0+)=g(0)$. Let ℓ_1 be the strictly increasing affine function given by Lemma 7. In particular, we have $\ell_1(1)=\Phi(F)^{\text{conc}}(1)$ and $\ell_1(\bar{w}_F)=g(1)$, and $\Phi(F)(w)=\ell_1(w)\wedge g(1)$ whenever $F\circ h^{-1}(w)\notin J(g)$ or $w=w_y=\inf\{w\in[0,h(\infty)]\cap\mathbb{R}:F\circ h^{-1}(w)=y\}<\bar{w}_F$ for some $y\in J(g)$.

We first find a formula for ℓ_1. Observe that either $F\circ h^{-1}(0)=F(0)\notin J(g)$ or $F(0)\in J(g)$ and $w_{F(0)}=0$. In both cases, $\ell_1(0)=\Phi(F)(0)=g(0)<g(1)$. We also have $\ell_1(1)=\Phi(F)^{\text{conc}}(1)=u^F(x_0)=\bar{R}$, by symmetry. This completely determines the shape of ℓ_1; namely,

$$\ell_1(w)=g(0)+(\bar{R}-g(0))w.$$

Next, recall A,A_y,w_y as defined in Lemma 7. We decompose $[0,h(\infty)]\cap\mathbb{R}$ into three disjoint parts: A_1, $A\backslash A_1$, and $\bigcup_{y\in J(g)}A_y$. Note that $A_1=[\bar{w}_F,h(\infty)]\cap\mathbb{R}=[\ell_1^{-1}(g(1)),h(\infty)]\cap\mathbb{R}$ by Lemma 5 and Lemma 7(i), and note that each A_y with $y\in J(g)$ has positive length by Lemma 4.

(i) On A_1, we have $F\circ h^{-1}\equiv 1$.

(ii) On $A\backslash A_1$, we have $g(F\circ h^{-1})=\Phi(F)=\ell_1\wedge g(1)=\ell_1$ by Lemma 7(ii). The strict monotonicity of ℓ_1 implies that $F\circ h^{-1}$ is strictly increasing on $A\backslash A_1$. Thus, A_y is a singleton for all $y\notin J(g)\bigcup\{1\}$. The relation $g(F\circ h^{-1})=\ell_1$ also implies $F\circ h^{-1}\leq g^{-1}\circ g(F\circ h^{-1})=g^{-1}(\ell_1)$. In view of Lemma 3, the inequality is, in fact, an equality. Indeed, any flat segment of g induces a gap in the range of F, which precisely excludes those points y for which $g^{-1}(g(y)) \neq y$, except possibly at the left endpoint of the flat segment—say, y_1. The exception only happens if F contains a flat segment at height y_1, which is equivalent to A_{y_1} having positive length. Thus, $F \circ h^{-1} = y_1$ is also ruled out on $A \backslash A_1$.

(iii) On each A_y with $y \in J(g)$, we use Lemma 7(iii) to obtain $g(y) = \Phi(F)(w_y) = \ell_1(w_y)$, which uniquely determines w_y.

In summary, we can decompose $[0, h(\infty)] \cap \mathbb{R}$ into (a) countably many intervals on which $F \circ h^{-1}$ is flat at some level $y \in J(g) \cup \{1\}$ and (b) the complementary set $A \backslash A_1$ on which $F \circ h^{-1} = g^{-1}(\ell_1)$. Each flat segment at level $y \in J(g)$ has the left endpoint $w_y = \ell_1^{-1}(g(y))$. To uniquely determine the right-continuous function $F \circ h^{-1}$, it only remains to specify, for each $y \in J(g)$, the right endpoint

$$\tilde{w}_y := \sup\{w \in [0, h(\infty)] \cap \mathbb{R} : F \circ h^{-1}(w) = y\} \leq \bar{w}_F$$

of the flat segment. To this end, let $y \in J(g)$ and $y' := \sup\{z \in [0,1] : g(z) = g(y+)\}$. We distinguish two cases:

Case 1: If $y' = y$, then g is nonconstant in any right neighborhood of y, which implies that $F \circ h^{-1}(\tilde{w}_y) = y$. (If $F \circ h^{-1}(\tilde{w}_y) > y$, then F would have an atom at $h^{-1}(\tilde{w}_y)$, and by Lemma 2, $g(F \circ h^{-1}(\tilde{w}_y)) = g(y+)$, contradicting the assumption that $y' = y$.) Let $w^{(m)} > \tilde{w}_y$ be a sequence such that $w^{(m)} \to \tilde{w}_y$, and let $y_m := F \circ h^{-1}(w^{(m)})$. By right continuity and the definition of $\tilde{w}_y$, we have $y_m \to y$ and $y_m > y$. For large m, we may assume $y_m < 1$. Observe that $w^{(m)} \geq w_{y_m} > \tilde{w}_y$, which implies $w_{y_m} \to \tilde{w}_y$. If $y_m \notin J(g)$, then $w^{(m)} = w_{y_m} \in A \backslash A_1$ and $g(y_m) = \ell_1(w^{(m)}) = \ell_1(w_{y_m})$ by (ii) in the proof of Theorem 1—Uniqueness. If $y_m \in J(g)$, then $g(y_m) = \ell_1(w_{y_n})$ by (iii) in the same proof. Combining the two cases and passing to the limit, we obtain $g(y+) = \ell_1(\tilde{w}_y)$.

Case 2: If $y' > y$, then by Lemma 3, F jumps from y to y' at $F^{-1}(y+) = h^{-1}(\tilde{w}_y)$. Hence, $F \circ h^{-1}(\tilde{w}_y) = y'$ and $\Phi(F)(\tilde{w}_y) = g(y') = g(y+)$. We have either $y' \notin J(g)$ or $y' \in J(g)$ with $\tilde{w}_y = w_{y'}$, and both lead to $\ell_1(\tilde{w}_y) = \Phi(F)(\tilde{w}_y) = g(y+)$.

In both cases, we have $\ell_1(\tilde{w}_y) = g(y+)$, which uniquely determines $\tilde{w}_y$.

Putting everything together and taking into account the right continuity of $F \circ h^{-1}$,

$$F \circ h^{-1}(w) = \begin{cases} y & \text{if } y \in J(g) \text{ and } w \in [\ell_1^{-1}(g(y)), \ell_1^{-1}(g(y+))), \\ 1 & \text{if } w \geq \ell_1^{-1}(g(1)), \\ g^{-1}(\ell_1(w)) & \text{otherwise.} \end{cases}$$

In summary, $F \circ h^{-1}(w) = g^{-1}(\ell_1(w) \wedge g(1))$, or $F(x) = F^*(x)$ after substituting $w = h(x)$. □

4. Optimal Reward Design

Consider a principal who may choose a normalized reward R (i.e., satisfying $R(0) = 0$ and $\bar{R} = 1$) and whose goal is to maximize the performance of the top $\alpha \in (0,1)$ fraction of players. More precisely, the aim is to maximize the lowest stopping position of all players in the ranks $[0, \alpha)$,

$$x_\alpha := F^{-1}((1-\alpha)+) = F_+^{-1}(1-\alpha),$$

where F is the equilibrium resulting from R and F_+^{-1} is the right-continuous inverse of F. See Remark 7 for the technical importance of using F_+^{-1} or, equivalently, of using the open interval $[0, \alpha)$ when defining the top ranks. Note that the constant $\bar{\mu}_\infty = \bar{\mu}_\infty(R)$ in Theorem 1 depends on R. For the following result, we assume $\mu \leq 0$ to ensure that $\mu < \bar{\mu}_\infty(R)$ holds for any reward R. Alternatively, one may relax this condition to $\mu < -[\sigma^2/(2x_0)]\log(1-\alpha)$ and restrict the principal to rewards R satisfying $\mu < \bar{\mu}_\infty(R)$.

Theorem 2. *Let $\alpha \in (0,1)$. Then*

$$R^*(r) = \frac{1}{\alpha} 1_{[0,\alpha)}(r)$$

is the unique normalized reward maximizing the performance x_α. The corresponding value is $x_\alpha^* = h^{-1}(1/\alpha)$, and the equilibrium distribution is

$$F^* = (1-\alpha) 1_{[0, x_\alpha^*)} + 1_{[x_\alpha^*, \infty)}.$$

Proof. Let R be an arbitrary normalized reward, and let $g(y) = R(1-y)$. By Theorem 1, the corresponding mean field equilibrium F is unique, and

$$F(x) = g^{-1}(h(x) \wedge g(1)) = 1 \wedge \inf\{y : g(y) > h(x) \wedge g(1)\}.$$

We have $F(x) > 1-\alpha$ if and only if $g(1-\alpha+\varepsilon) \le h(x) \wedge g(1)$ for some $\varepsilon = \varepsilon(x) > 0$; hence

$$\begin{aligned} F_+^{-1}(1-\alpha) &= \inf\{x \ge 0 : g(1-\alpha+\varepsilon) \le h(x) \wedge g(1) \text{ for some } \varepsilon > 0\} \\ &= h^{-1}(g((1-\alpha)+)). \end{aligned}$$

As h^{-1} is strictly increasing, maximizing this quantity is equivalent to maximizing $g((1-\alpha)+)$. Recalling that g is monotone and left continuous and that $\int_0^1 g(y)dy = 1$, the unique maximizer is given by $g^*(y) := R^*(1-y)$, and the corresponding maximum value is $h^{-1}(1/\alpha)$. By Theorem 1 (or Proposition 1), the corresponding equilibrium is F^*. □

Comparing with the results cited in the introduction (and recalled in more detail in Section 5.3), Theorem 2 gives a clear-cut answer to a question that remained partially open in the n-player setting. On the other hand, the result illustrates that the mean field analysis alone could easily lead to an oversimplified picture: the optimal design in the n-player game is *not* given by the cutoff reward at the target rank in most cases. See also Section 5 for further comparison of mean field and n-player games.

Remark 7. The principal's goal is to maximize $F_+^{-1}(1-\alpha)$ rather than the quantile $F^{-1}(1-\alpha)$. Indeed, the worst performance among the top α-fraction of players need not be the same as the best performance among the bottom $(1-\alpha)$-fraction. The equilibrium F^* of Theorem 2 has an atom of size α at x_α^* and an atom of size $1-\alpha$ at the origin. Thus, $(F^*)_+^{-1}(1-\alpha) = x_\alpha^*$, but $(F^*)^{-1}(1-\alpha) = 0$.

It is crucial to formulate the principal's problem in the form stated in the preceding: if instead we aim to maximize the best performance in the quantile $F^{-1}(1-\alpha)$, the optimization *fails to admit a solution.* To see this, note that for each $m \ge 1$, the cutoff reward $R^{(m)}(r) := (\alpha+1/m)^{-1} 1_{[0,\alpha+1/m)}(r)$ gives rise to the equilibrium $F^{(m)} = (1-\alpha-1/m)1_{[0,x_m)} + 1_{[x_m,\infty)}$, where $x_m = h^{-1}(1/(\alpha+1/m))$. Moreover, $(F^{(m)})^{-1}(1-\alpha) = x_m$ increases to $h^{-1}(1/\alpha)$ as $m \to \infty$. However, there exists no equilibrium distribution F achieving $F^{-1}(1-\alpha) = h^{-1}(1/\alpha)$. Indeed, by Theorem 2, such F would have to coincide with F^*, but $(F^*)^{-1}(1-\alpha) = 0 < h^{-1}(1/\alpha)$.

Remark 8. In analogy to the "price of anarchy," we can compare the principal's optimization over equilibria with a different problem where the planner can dictate the players' stopping strategy (regardless of equilibrium considerations). This problem can be stated as

$$\max_{F\in\mathcal{F}} F_+^{-1}(1-\alpha).$$

Using Lemma 1, we can check that the unique solution is $F = F^*$, the equilibrium distribution of Theorem 2. In particular, the "welfare" of the second-best principal who can only choose the reward function is equal to that of a planner who can dictate strategies. This fact extends to other objectives for the principal: from the explicit formula of the mean field equilibrium in Theorem 1, we see that any distribution $F \in \mathcal{F}$ with compact support can be attained in equilibrium under the normalized reward $R(r) = g(1-r)$ with

$$g^{-1}(y) := F(h^{-1}(y)).$$

As a result, finding an optimal reward is equivalent to finding an optimal target distribution. This explains why the first-best and the second-best solution coincide.

This consideration also shows a different avenue to Theorem 2: if one is only interested in this specific question rather than mean field equilibria for general reward functions, one can first argue that $\arg\max_{F\in\mathcal{F}} F_+^{-1}(1-\alpha) = F^*$ and then, as in Proposition 1, that F^* is the unique equilibrium for R^*.

5. Convergence to the Mean Field

To formulate the n-player game associated with our mean field contest, fix a decreasing, nonconstant reward vector, $(R_1, \ldots, R_n)$. Here, R_1 is interpreted as the reward for the best rank, whereas R_n is the worst. As in the mean field game, the players are ranked according to their level of stopping, and ties are split uniformly at random. The set $\mathcal{F}$ of feasible stopping distributions remains the same, and the definition of equilibrium is analogous. It is shown in Nutz and Zhang [42] that the n-player game admits a unique equilibrium $F_n^* \in \mathcal{F}$ as soon as the drift μ satisfies

$$\mu < \bar{\mu}_n := \frac{\sigma^2}{2x_0}\log\left(\frac{R_1 - R_n}{R_1 - \bar{R}_n}\right), \quad \text{where} \quad \bar{R}_n := \frac{1}{n}\sum_{k=1}^{n} R_k.$$

The equilibrium distribution has compact support $[0,\bar{x}_n]$ and cdf

$$F_n^*(x) = g_n^{-1}(u_n^*(x)), \quad \text{where} \quad \bar{x}_n = h^{-1}\left(\frac{R_1 - R_n}{\bar{R} - R_n}\right) \quad \text{and}$$
$$g_n(y) = \sum_{k=1}^{n} R_k \binom{n-1}{k-1} y^{n-k}(1-y)^{k-1}, \quad \text{and} \quad u_n^*(x) = \left[R_n + (\bar{R}_n - R_n)h(x)\right] \wedge R_1$$

is the equilibrium value function. In contrast to the mean field setting, F_n^* is always atomless. Moreover, g_n is strictly increasing and smooth—hence so is its (true) inverse g_n^{-1}.

5.1. Convergence of the *n*-Player Equilibrium

The next result shows that if the reward vector is induced by a reward function for the mean field game, the n-player equilibrium distributions F_n^* and value functions u_n^* converge to their mean field counterparts F^* and u^*, as described in Theorem 1.

Theorem 3. *Let $R:[0,1] \to \mathbb{R}_+$ be a reward function, $\mu < \bar{\mu}_\infty$, and define $R_k := R(k/n)$ for $k = 1,\dots,n$. Then as $n \to \infty$, using $(R_1,\dots,R_n)$ as the reward for the n-player game and R for the mean field game, the associated unique equilibrium distributions converge weakly, and the equilibrium value functions converge uniformly.*

Suppose, in addition, that $g(y) = R(1-y)$ is piecewise α-Hölder continuous and that its right-continuous inverse is piecewise β-Hölder continuous, where $\alpha,\beta \in (0,1]$.[4] *With W_1 denoting the 1-Wasserstein metric, we have the convergence rates*

$$\|u_n^* - u^*\|_\infty = O(n^{-\alpha}), \quad W_1(F_n^*, F^*) = O(n^{-\alpha'/2} + n^{-\alpha\beta}) \quad \forall \alpha' \in (0,\alpha] \cap (0,1).$$[5]

Remark 9. If we consider a generalized reward function R with $R(1-) < R(1)$, as discussed in Remark 2, the limit of the n-player equilibria selects a particular equilibrium among the infinitely many mean field equilibria—namely, the one detailed in Theorem 1. This follows from the fact that the proof of Theorem 3 does not use the condition $R(1-) = R(1)$.

Before proceeding with the proof, we state a formula that will be used in later arguments as well. Consider the empirical cdf of i.i.d. uniform random variables $\{U_i\}_{i=1,\dots,n-1}$ on $[0,1]$,

$$\hat{F}_{n-1}(y) = \frac{1}{n-1}\#\{i : U_i \le y\}.$$

Let $0 \le y_1 \le y_2 \le 1$. Among the $n-1$ random variables $\{U_i\}$, there are $I_{n-1} = (n-1)(1-\hat{F}_{n-1}(y_2))$ with values above y_2, $J_{n-1} = (n-1)\hat{F}_{n-1}(y_1)$ below y_1, and $K_{n-1} = (n-1)(\hat{F}_{n-1}(y_2) - \hat{F}_{n-1}(y_1))$ in between y_1 and y_2. Thus, we have the following formula for any function $\phi(i,j,k)$:

$$\begin{aligned}
&\sum_{\substack{i,j,k\ge 0\\ i+j+k=n-1}} \phi(i,j,k)\binom{n-1}{i,j,k}(1-y_2)^i y_1^j (y_2-y_1)^k \\
&= \sum_{\substack{i,j,k\ge 0\\ i+j+k=n-1}} \phi(i,j,k) P(I_{n-1}=i, J_{n-1}=j, K_{n-1}=k) \\
&= E\Big[\phi\Big((n-1)(1-\hat{F}_{n-1}(y_2)), (n-1)\hat{F}_{n-1}(y_1), (n-1)(\hat{F}_{n-1}(y_2)-\hat{F}_{n-1}(y_1))\Big)\Big].
\end{aligned} \tag{8}$$

Proof of Theorem 3. We have $R_n = R(1)$ and $R_1 = R(1/n) \to R(0)$ by the right continuity of R. Moreover, the Riemann sum $(1/n)\sum_{k=1}^n R_k \to \int_0^1 R(r)dr = \bar{R}$. It follows that $\bar{\mu}_n \to \bar{\mu}_\infty$ and $\bar{x}_n \to \bar{x}$, so that $\mu < \bar{\mu}_\infty$ ensures $\mu < \bar{\mu}_n$ for all n sufficiently large, and the equilibria are uniquely defined and compactly supported on $[0,\bar{x}_n] \subseteq [0,\bar{x}+1]$. The pointwise convergence of u_n^* to u^* is clear from their respective formulas. As these functions are increasing and u^* is continuous, the pointwise convergence is uniform on $[0,\bar{x}+1]$ (see, e.g., Resnick [44, proposition 2.1]). Outside the compact set $[0,\bar{x}+1]$, we have $\lim_n |u_n^* - u^*| = \lim_n |R_1 - R(0)| = 0$.

To show the weak convergence of the equilibrium distributions, we prove $F_n^*(x) \to F^*(x)$ whenever x is a point of continuity of F^*. We first argue that

$$g_n \to g \quad \text{at at every point of continuity of } g. \tag{9}$$

Taking $y_1 = y_2 = y$ and $\phi(i,j,k) = R_{i+1} = R((i+1)/n)$ in (8), we obtain

$$g_n(y) = E\left[R\left(\frac{(n-1)(1-\hat{F}_{n-1}(y))+1}{n}\right)\right] = E\left[g\left(\frac{n-1}{n}\hat{F}_{n-1}(y)\right)\right]. \tag{10}$$

By the strong law of large numbers, $\hat{F}_{n-1}(y) \to y$ a.s. for each y. If y is a point of continuity of g, it follows that $g((1-1/n)\hat{F}_{n-1}(y)) \to g(y)$ a.s., and the bounded convergence theorem yields $g_n(y) \to g(y)$ as claimed.

We have $F_n^*(x) = g_n^{-1}(z_n)$ for $z_n := u_n^*(x)$ and, similarly, $F^*(x) = g^{-1}(z)$ for $z := u^*(x)$. By the preceding, $z_n \to z$ for all x, and therefore we need to show that $g_n^{-1}(z_n) \to g^{-1}(z)$ whenever $z \in C$, where C is the set of continuity points of g^{-1}. Up to normalization, we may think of g_n, g as a cdf of weakly converging distributions. It is then known that the inverses g_n^{-1} converge to g^{-1} on the set where the left- and right-continuous inverses of g coincide and hence on C (see the proof of theorem 3.2.2 in Durrett [19]). We have $g_n^{-1}(z \pm \varepsilon) \to g^{-1}(z \pm \varepsilon)$ for $z \in C$ and $\varepsilon > 0$ with $z \pm \varepsilon \in C$. Using the fact that g_n^{-1}, g^{-1} are monotone, we deduce that $g_n^{-1}(z_n) \to g^{-1}(z)$ whenever $z_n \to z$ and $z \in C$, as desired.

Next, we derive the convergence rate under the additional piecewise Hölder condition. Let $0 = r_0 < r_1 < \cdots < r_m = 1$ be a finite partition of $[0,1]$ such that R is α-Hölder continuous on each interval (r_{i-1}, r_i). Write $\bar{R}^n := (1/n)\sum_{k=1}^n R_k$. We have $|R_1 - R(0)| = |R(1/n) - R(0)| = O(n^{-\alpha})$ and

$$|\bar{R}^n - \bar{R}| \le \int_0^1 |R(\lceil rn \rceil/n) - R(r)|dr \le \sum_{i=1}^m \int_{r_{i-1}+1/n}^{r_i - 1/n} \frac{C}{n^\alpha} dr + \frac{2m(R(0)-R(1))}{n} = O(n^{-\alpha}).$$

Throughout the proof, the constant C may vary from line to line. Using the above-mentioned convergence rates of R_1 and $\bar{R}^n$, as well as the Lipschitz continuity of h^{-1} and $(x,y) \longmapsto (x-R(1))/(y-R(1))$ when $y > (\bar{R}+R(1))/2 > R(1)$, one can easily show that

$$|\bar{x}_n - \bar{x}| = O(n^{-\alpha}). \tag{11}$$

Similarly, we use the convergence rates of R_1 and $\bar{R}^n$, and the Lipschitz continuity of $(x,y) \longmapsto [R(1) + (y - R(1))h] \wedge x$ uniformly in $h \in [0, h(\bar{x}+1)]$, to obtain

$$\|u_n^* - u^*\|_\infty = \max\left\{\sup_{0 \le x \le \bar{x}+1} |u_n^*(x) - u^*(x)|, |R_1 - R(0)|\right\} = O(n^{-\alpha}). \tag{12}$$

To show the convergence rate of the equilibrium distributions, note that

$$W_1(F_n, F) = \int_{\mathbb{R}} |F_n^*(x) - F^*(x)|dx = \int_0^{\bar{x}_n \vee \bar{x}} |F_n^*(x) - F^*(x)|dx \le I + J + |\bar{x}_n - \bar{x}|, \tag{13}$$

where

$$I := \int_0^{\bar{x}_n \wedge \bar{x}} |g_n^{-1}(u_n^*(x)) - g^{-1}(u_n^*(x))|dx, \quad J := \int_0^{\bar{x}_n \wedge \bar{x}} |g^{-1}(u_n^*(x)) - g^{-1}(u^*(x))|dx.$$

We deal with the two terms separately. Observe that u_n^* is invertible on $[0, \bar{x}_n]$. A change of variable leads to

$$I = \int_{R_n}^{u_n^*(\bar{x}_n \wedge \bar{x})} |g_n^{-1}(z) - g^{-1}(z)|(h^{-1})'\left(\frac{z - R_n}{\bar{R}^n - R_n}\right)\frac{dz}{\bar{R}^n - R_n}.$$

Set $g_n^{-1}(z) := 1$ for $z > R_1$. Because $h^{-1} \in C^1[0, h(\bar{x}+1)]$, we have for large n that

$$I \le C\int_{R(1)}^{R(0)} |g_n^{-1}(z) - g^{-1}(z)|dz = C\int_0^1 |g_n(y) - g(y)|dy.$$

Let $y_i := 1 - r_i$ so that g is α-Hölder continuous on each interval (y_{i-1}, y_i). Fix $\alpha' \in (0,\alpha] \cap (0,1)$ and $y \in I_{n,i} := (y_{i-1} + 2(n-1)^{-\alpha'/2}, y_i - 2(n-1)^{-\alpha'/2})$; it is easy to show that if $|\hat{F}_{n-1}(y) - y| \le (n-1)^{-\alpha'/2}$, then

$(1-1/n)\hat{F}_{n-1}(y)\in(y_{i-1},y_i)$ for n sufficiently large (independently of y). In this case, by (10), the α-Hölder continuity of g on (y_{i-1},y_i), Jensen's inequality, and the Dvoretzky–Kiefer–Wolfowitz inequality, we have that

$$\begin{aligned}
g_n(y)-g(y) &= E\left[g\left(\frac{n-1}{n}\hat{F}_{n-1}(y)\right)-g(y)\right]\\
&\le E\left[C\left(|\hat{F}_{n-1}(y)-y|+\frac{1}{n}\right)^{\alpha}\right]+CP\left(\sup_{y\in[0,1]}|\hat{F}_{n-1}(y)-y|>(n-1)^{-\alpha'/2}\right)\\
&\le C\left(E\left[\sup_{y\in[0,1]}|\hat{F}_{n-1}(y)-y|\right]+\frac{1}{n}\right)^{\alpha}+Ce^{-2(n-1)^{1-\alpha'}}\\
&\le C\left(\int_0^{\infty}2e^{-2(n-1)z^2}dz+\frac{1}{n}\right)^{\alpha}+Ce^{-2(n-1)^{1-\alpha'}}\\
&= O(n^{-\alpha/2})+O(e^{-2(n-1)^{1-\alpha'}})=O(n^{-\alpha/2}),\quad \text{uniformly in } y\in I_{n,i}.
\end{aligned}$$

Consequently,

$$\begin{aligned}
I &\le C\int_{\bigcup_i I_{n,i}}|g_n(y)-g(y)|dy+C\int_{(\bigcup_i I_{n,i})^c}|g_n(y)-g(y)|dy\\
&\le O(n^{-\alpha/2})+4C(n-1)^{-\alpha'/2}m(R(0)-R(1))=O(n^{-\alpha'/2}).
\end{aligned}\tag{14}$$

For the convergence of J, let $R(1)=z_0<z_1<\cdots<z_\ell=R(0)$ be a finite partition of $[R(1),R(0)]$ such that g^{-1} is β-Hölder continuous on each interval (z_{i-1},z_i). Note that u^* is invertible on $[0,\bar{x}]$. For any $z\in[R(0),R(1)]$,

$$|u_n^*((u^*)^{-1}(z))-z|=|u_n^*((u^*)^{-1}(z))-u^*((u^*)^{-1}(z))|\le\|u_n^*-u^*\|_\infty\le Cn^{-\alpha}$$

by (12). It follows that

$$\begin{aligned}
J &\le \int_{R(1)}^{R(0)}|g^{-1}(u_n^*\circ(u^*)^{-1}(z))-g^{-1}(z)|[(u^*)^{-1}]'(z)dz\\
&\le \sum_i\int_{z_{i-1}+Cn^{-\alpha}}^{z_i-Cn^{-\alpha}}C|u_n^*\circ(u^*)^{-1}(z)-z|^{\beta}dz+2Cn^{-\alpha}\ell\\
&\le Cn^{-\alpha\beta}(R(0)-R(1))+2Cn^{-\alpha}\ell=O(n^{-\alpha\beta}).
\end{aligned}\tag{15}$$

Combining (11), (13), (14) and (15), we obtain $W_1(F_n,F)=O(n^{-\alpha'/2}+n^{-\alpha\beta})$. □

Remark 10. A discussion related to Theorem 3 can be found in the work of Fang and Noe [21] on n-player capacity-constrained contests, which can be related to the present game via Skorokhod embedding. (Some results of the preprint by Fang and Noe [21] were later published as Fang et al. [22].) Namely, Fang and Noe [21, proposition 9] study the effect of scaling the n-player contest by multiplying the number of participants while dividing the reward at each rank. This basically corresponds to taking $\lim_n F_n^*$, if only for the particular case where R is a step function. An infinite player game is not considered so that the limiting distribution F^* cannot be recognized as a mean field object. Instead, the authors derive an involved algorithm (Fang and Noe [21, remark A-1]) to construct $F^*=\lim_n F_n^*$. The limit is a step function in this particular case, and the algorithm determines the $n-1$ jump locations and magnitudes. It seems that the simple representation (3), or the game-theoretic meaning of F^*, was not identified.

5.2. ε-Nash Equilibrium Property of the Mean Field Strategy

Recall that $\xi^F(x)$ denotes the payoff for stopping at x if all other players in the mean field game use F (see (1)). Analogously, we can define the expected payoff $\xi_n^F(x)$ in the n-player game. We say that $F^*\in\mathcal{F}$ is an *ε-Nash equilibrium* of the n-player game if

$$\int\xi_n^{F^*}dF^*\ge\int\xi_n^{F^*}dF-\varepsilon\quad\text{for all}\quad F\in\mathcal{F}.$$

That is, a player deviating unilaterally from F^* can improve her expected payoff by at most ε. Correspondingly, F^* is an $o(1)$-*Nash equilibrium* as $n \to \infty$ if for any $\varepsilon > 0$, the preceding holds for all large n, or equivalently,

$$\lim_{n\to\infty} \sup_{F\in\mathcal{F}} \left(\int \xi_n^{F^*} dF - \int \xi_n^{F^*} dF^* \right) = 0.$$

We can now state the main result of this subsection.

Theorem 4. *Let R be a reward function, $\mu < \bar{\mu}_\infty$, and let F^* be the associated mean field equilibrium. Define $R_k := R(k/n)$, $k = 1, \dots, n$ as reward for the n-player game. Then F^* is an $o(1)$-Nash equilibrium of the n-player game as $n \to \infty$ if and only if R is continuous.*

The positive result in Theorem 4 is consistent with a large body of literature (see the Introduction). That the continuity condition is sharp may be surprising. Indeed, we will show that if R has a jump and $\varepsilon > 0$ is small enough, then F^* is not an ε-Nash equilibrium for *all* large n. This is not related to atoms in the equilibrium but rather to the gap in the support of F^* caused by the jump $R(x) - R(x-)$ in reward and a stochastic knife-edge phenomenon. The idea of the proof is that a player can improve by suitably shifting some mass of the stopping distribution *into the gap.* A level of stopping inside the gap would imply the reward $R(x)$ in the mean field game, but in the n-player game, the result depends on the sample—the reward is approximately $R(x)$ in roughly half the samples, but the higher reward is $R(x-)$ in the other half. By shifting more mass from below the gap than from above (all while maintaining feasibility), the player can increase the payoff relative to F^*.

The proof of Theorem 4 occupies the remainder of this subsection. Throughout the proof, the rewards and F^* are defined as in Theorem 4. As a first step, we derive a convenient formula for $\xi_n^F(x)$. The probability that among players $2, \dots, n$ there are exactly i players stopping above x, j players below x, and k players at x is given by

$$\binom{n-1}{i,j,k} (1 - F(x))^i F(x-)^j (F(x) - F(x-))^k.$$

Such a configuration leads to an average payoff $(R_{i+1} + \cdots + R_{i+k+1})/(k+1)$ for player 1 as ties are broken randomly. It follows that

$$\xi_n^F(x) = \sum_{\substack{i,j,k\geq 0\\ i+j+k=n-1}} \frac{R_{i+1} + \cdots + R_{n-j}}{k+1} \binom{n-1}{i,j,k} (1 - F(x))^i F(x-)^j (F(x) - F(x-))^k.$$

This reduces to $g_n(F(x))$ if $F(x) = F(x-)$. Taking $\phi(i,j,k) = (R_{i+1} + \cdots + R_{i+k+1})/(k+1)$ in (8), we have the alternative representation:

$$\xi_n^F(x) = E\left[\frac{\sum_{\ell=(n-1)(1-\hat{F}_{n-1}(F(x)))+1}^{n-(n-1)\hat{F}_{n-1}(F(x-))} R_\ell}{(n-1)(\hat{F}_{n-1}(F(x)) - \hat{F}_{n-1}(F(x-))) + 1} \right]. \tag{16}$$

Lemma 8. *Let $F \in \mathcal{F}$ have an atom at x. Then $\lim_n \xi_n^F(x) = \xi^F(x)$.*

Proof. Let F have an atom at x. Write $y_1 = F(x-)$ and $y_2 = F(x)$. By (16),

$$\begin{aligned} \xi_n^F(x) &= E\left[\frac{\sum_{\ell=n-(n-1)\hat{F}_{n-1}(y_2)}^{n-(n-1)\hat{F}_{n-1}(y_1)} g\left(\frac{n-\ell}{n}\right)}{(n-1)(\hat{F}_{n-1}(y_2) - \hat{F}_{n-1}(y_1)) + 1} \right] \\ &= E\left[\frac{n}{(n-1)(\hat{F}_{n-1}(y_2) - \hat{F}_{n-1}(y_1)) + 1} \sum_{\ell=(n-1)\hat{F}_{n-1}(y_1)}^{(n-1)\hat{F}_{n-1}(y_2)} g\left(\frac{\ell}{n}\right) \frac{1}{n} \right]. \end{aligned}$$

Using the a.s. convergence of $\hat{F}_{n-1}(y)$ to y, we deduce that

$$\xi_n^F(x) \to \frac{1}{y_2 - y_1} \int_{y_1}^{y_2} g(y) dy = \xi^F(x). \quad \square$$

Lemma 9. *Let R be continuous. Then $\xi_n^{F^*}(\cdot)$ converges to $\xi^{F^*}(\cdot)$ uniformly.*

Proof. We first show that $\xi_n^{F^*}$ converges to ξ^{F^*} pointwise. The convergence at points of discontinuity of F^* holds by Lemma 8. At points of continuity, we have $\xi_n^{F^*} = g_n \circ F^*$ and $\xi^{F^*} = g \circ F^*$. The pointwise convergence then follows from (9) and the assumed continuity of g.

By Theorem 1, F^* has compact support $[0,\bar{x}]$. For $x>\bar{x}$, it is clear that $|\xi_n^{F^*}(x)-\xi^{F^*}(x)|=|g_n(1)-g(1)|=|R(1/n)-R(0)|\to 0$. To see that the convergence is also uniform on $[0,\bar{x}]$, we note that $\xi_n^{F^*}$ is an increasing function for each n. Moreover, the pointwise limit ξ^{F^*} is continuous: as g is continuous, we have $g(g^{-1}(z))=z$, and then $\xi^{F^*}(x)=g(F^*(x))=[R(1)+(\bar{R}-R(1))h(x)]\wedge R(0)$ is continuous as well (see Lemma 2 and Theorem 1). A standard argument for monotone functions then yields that the pointwise convergence is uniform. □

Lemma 10. *If g has a jump at $y\in(0,1)$, then* $\lim_n g_n(y)=[g(y)+g(y+)]/2$.

Proof. Let $\varepsilon>0$. Using (10), we have

$$\begin{aligned}
g_n(y)&=E\left[g\left(\frac{n-1}{n}\hat{F}_{n-1}(y)\right)\right]\\
&\geq g(y-\varepsilon)P\left(y-\varepsilon<\frac{n-1}{n}\hat{F}_{n-1}(y)\leq y\right)+g(y+)P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)>y\right)\\
&=g(y-\varepsilon)P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)>y-\varepsilon\right)+\big(g(y+)-g(y-\varepsilon)\big)P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)>y\right).
\end{aligned}$$

The strong law of large numbers implies $\hat{F}_{n-1}(y)\to y$ a.s., and hence

$$P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)>y-\varepsilon\right)\to 1.$$

By the central limit theorem, $\sqrt{n-1}(\hat{F}_{n-1}(y)-y)/\sqrt{y(1-y)}$ converges to $\mathcal{N}(0,1)$ in distribution. It follows that for any fixed $\gamma>0$ and $n\geq 1+y/[(1-y)\gamma^2]$,

$$\begin{aligned}
P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)>y\right)&=P\left(\frac{\sqrt{n-1}(\hat{F}_{n-1}(y)-y)}{\sqrt{y(1-y)}}>\frac{1}{\sqrt{n-1}}\sqrt{\frac{y}{1-y}}\right)\\
&\geq P\left(\frac{\sqrt{n-1}(\hat{F}_{n-1}(y)-y)}{\sqrt{y(1-y)}}>\gamma\right)\to 1-N(\gamma),
\end{aligned}$$

where $N(\cdot)$ is the standard normal cdf. Combining the two limits, we obtain

$$\liminf_n g_n(y)\geq g(y-\varepsilon)+(g(y+)-g(y-\varepsilon))(1-N(\gamma))=g(y-\varepsilon)N(\gamma)+g(y+)(1-N(\gamma)).$$

Similarly, we can show

$$\begin{aligned}
\limsup_n\, g_n(y)&\leq\big(g(y)-g(y+\varepsilon)\big)\liminf_n\, P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)\leq y\right)\\
&\quad+g(y+\varepsilon)\limsup_n\, P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)<y+\varepsilon\right)\\
&\quad+g(1)\limsup_n\, P\left(\frac{n-1}{n}\hat{F}_{n-1}(y)\geq y+\varepsilon\right)\\
&\leq\big(g(y)-g(y+\varepsilon)\big)N(\gamma)+g(y+\varepsilon).
\end{aligned}$$

Finally, we send $\gamma,\varepsilon\to 0$ and use the left continuity of g. □

Proof of Theorem 4. We prove the theorem in two parts.

- *Part* 1: *Sufficiency.* Let R be continuous, and let $\varepsilon>0$. Lemma 9 shows the existence of n_ε such that $\|\xi_n^{F^*}-\xi^{F^*}\|_\infty<\varepsilon/2$ whenever $n\geq n_\varepsilon$. Let $F\in\mathcal{F}$. For $n\geq n_\varepsilon$, noting that $\int\xi^{F^*}dF-\int\xi^{F^*}dF^*\leq 0$ by the equilibrium property of F^*,

$$\begin{aligned}
&\int\xi_n^{F^*}dF-\int\xi_n^{F^*}dF^*\\
&=\int\xi_n^{F^*}dF-\int\xi^{F^*}dF+\int\xi^{F^*}dF-\int\xi^{F^*}dF^*+\int\xi^{F^*}dF^*-\int\xi_n^{F^*}dF^*\\
&\leq\int\left|\xi_n^{F^*}-\xi^{F^*}\right|dF+\int\left|\xi_n^{F^*}-\xi^{F^*}\right|dF^*<\varepsilon.
\end{aligned}$$

This proves the $o(1)$-Nash property of F^*.

- *Part* 2: *Necessity*. Let $g(y) = R(1-y)$ have a jump at $y_0 \in (0,1)$. We show the stronger statement

$$\sup_{F \in \mathcal{F}} \liminf_n \left(\int \xi_n^{F^*} dF - \int \xi_n^{F^*} dF^* \right) > 0. \tag{17}$$

Let

$$a = h^{-1}\left(\frac{g(y_0) - R(1)}{\bar{R} - R(1)} \right) \quad \text{and} \quad b = h^{-1}\left(\frac{g(y_0+) - R(1)}{\bar{R} - R(1)} \right).$$

By Theorem 1, the associated mean field equilibrium F^* is flat on $[a,b)$ and $F^*(a-\eta) < F^*(a) = y_0 < F^*(b+\eta)$ for any $\eta > 0$. Suppose players $2, \ldots, n$ all use F^* with associated measure ν^*, and player 1 considers an alternative strategy of the form

$$\nu = \nu^* - \zeta + |\zeta|\delta_{a'}$$

for some $a' \in (a,b)$ and a subprobability $\zeta \le \nu^*$ with density $d\zeta/d\nu^* = f$. To ensure the feasibility of ν, we require $\int h d\nu = \int h d\nu^*$, which translates to

$$\int (h(a') - h) f d\nu^* = 0. \tag{18}$$

Our goal is to obtain a lower bound for the payoff difference

$$\int \xi_n^{F^*} d\nu - \int \xi_n^{F^*} d\nu^* = |\zeta| \xi_n^{F^*}(a') - \int \xi_n^{F^*} f d\nu^* \tag{19}$$

that is independent of n for n large. Because F^* is continuous at a', we have $\xi_n^{F^*}(a') = g_n(F^*(a')) = g_n(y_0)$. By Lemma 10,

$$\liminf_n \xi_n^{F^*}(a') = \liminf_n g_n(y_0) = \frac{g(y_0) + g(y_0+)}{2}. \tag{20}$$

Let $x \ge 0$. If F^* is continuous at x, we use (9) to get $\limsup_n \xi_n^{F^*}(x) = \limsup_n g_n(F^*(x)) \le g(F^*(x)+)$, whereas if F^* has a jump at x, we use Lemma 8 to get $\xi_n^{F^*}(x) \to \xi^{F^*}(x) = g(F^*(x))$.

Reverse Fatou's lemma then implies

$$\limsup_n \int \xi_n^{F^*} f d\nu^* \le \int g(F^*(\cdot)+) f d\nu^*|_{\mathbb{R}_+ \setminus \{a\}} + g(F^*(a)) f(a) \nu^*\{a\}. \tag{21}$$

Substituting (20) and (21) into (19), we obtain

$$\begin{aligned}
&\liminf_n \left(\int \xi_n^{F^*} d\nu - \int \xi_n^{F^*} d\nu^* \right) \\
&\ge \frac{g(y_0) + g(y_0+)}{2} \int f d\nu^* - \int g(F^*(\cdot)+) f d\nu^*|_{\mathbb{R}_+ \setminus \{a\}} - g(y_0) f(a) \nu^*\{a\} \\
&= \frac{g(y_0+) - g(y_0)}{2} f(a) \nu^*\{a\} + \int \left(\frac{g(y_0) + g(y_0+)}{2} - g(F^*(\cdot)+) \right) f d\nu^*|_{[0,a)} \\
&\quad - \int \left(g(F^*(\cdot)+) - \frac{g(y_0) + g(y_0+)}{2} \right) f d\nu^*|_{[b,\infty)}.
\end{aligned}$$

As $F^*(x) < y_0$ for $x < a$, and thus $g(F^*(x)+) \le g(y_0) = a$, we can further bound the aforementioned expression from below by

$$C_f := \frac{g(y_0+) - g(y_0)}{2} \int f d\nu^*|_{[0,a]} - \int \left(R(0) - \frac{g(y_0) + g(y_0+)}{2} \right) f d\nu^*|_{[b,\infty)}.$$

It remains to show that by choosing a suitable Radon–Nikodym derivative f, the lower bound C_f for the expected improvement can be made strictly positive. To this end, we pick

$$f(x) = 1_{(a-\eta,a]}(x) + \lambda 1_{[b,\infty)}(x)$$

for some constants $\lambda \in [0,1]$ and $\eta > 0$ to be determined. With this form of f, we always have $0 \le \zeta \le \nu^*$, and the feasibility condition (18) becomes

$$\lambda = \frac{\int (h(a') - h)d\nu^*|_{(a-\eta,a]}}{\int (h - h(a'))d\nu^*|_{[b,\infty)}} \in \left(0, \frac{h(a') - h((a-\eta)\vee 0)}{h(b) - h(a')} \cdot \frac{\nu^*(a-\eta,a]}{\nu^*[b,\infty)}\right].$$

We use this equality as the definition for λ. Then

$$\begin{aligned} C_f &= \frac{g(y_0+) - g(y_0)}{2}\nu^*(a-\eta,a] - \lambda\left(R(0) - \frac{g(y_0) + g(y_0+)}{2}\right)\nu^*[b,\infty) \\ &\geq \nu^*(a-\eta,a]\left(\frac{g(y_0+) - g(y_0)}{2} - \left(R(0) - \frac{g(y_0) + g(y_0+)}{2}\right)\frac{h(a') - h((a-\eta)\vee 0)}{h(b) - h(a')}\right), \end{aligned}$$

where the inequality is derived by replacing λ by its upper bound. Choose $a' - a$ and η sufficiently small so that

$$\frac{h(a') - h((a-\eta)\vee 0)}{h(b) - h(a')} < \min\left(\nu^*[b,\infty), \frac{g(y_0+) - g(y_0)}{2R(0) - g(y_0) - g(y_0+)}\right).$$

Then $\lambda \in (0,1)$ and $C_f > 0$. □

Remark 11. In Theorem 4, if R is α-Hölder continuous and has finitely many flat segments, the accuracy of the mean field approximation can be strengthened to $O(n^{-\alpha/2})$.

Indeed, by the sufficiency proof of Theorem 4, showing that F^* is an $O(n^{-\alpha/2})$-Nash equilibrium amounts to showing $\|\xi_n^{F^*} - \xi^{F^*}\|_\infty = O(n^{-\alpha/2})$. If x is a point of continuity of F^*, then $|\xi_n^{F^*}(x) - \xi^{F^*}(x)| = |g_n(F^*(x)) - g(F^*(x))| = O(n^{-\alpha/2})$ uniformly in x by (10) and the uniform $O(n^{-1/2})$-convergence of $E\left[\sup_{y\in[0,1]}|(1-1/n)\hat{F}_{n-1}(y) - y|\right]$. Whereas if F^* has an atom at x, we write $y_1 = F(x-)$, $y_2 = F(x)$ and $z_i = (1-1/n)\hat{F}_{n-1}(y_i)$. Again, using $E|z_i - y_i| = O(n^{-1/2})$ and the Hölder condition on g, we can strengthen Lemma 8 to

$$\begin{aligned} |\xi_n^{F^*}(x) - \xi^{F^*}(x)| &= \left|E\left[\frac{1}{z_2 - z_1}\sum_{\ell=nz_1}^{nz_2} g\left(\frac{\ell}{n}\right)\frac{1}{n}\right] - \frac{1}{y_2 - y_1}\int_{y_1}^{y_2} g(y)dy\right| \\ &= E\left|\frac{1}{z_2 - z_1}\int_{z_1}^{z_2+1/n} g\left(\frac{\lceil ny\rceil}{n}\right)dy - \frac{1}{y_2 - y_1}\int_{y_1}^{y_2} g(y)dy\right| \\ &\leq E\left[\left|\frac{1}{z_2 - z_1} - \frac{1}{y_2 - y_1}\right|R(0)(z_2 - z_1 + 1/n)\right] \\ &\quad + \frac{1}{y_2 - y_1}\int_{y_1}^{y_2}\left|g\left(\frac{\lceil ny\rceil}{n}\right) - g(y)\right|dy + E(|z_1 - y_1| + |z_2 - y_2 + 1/n|)R(0) \\ &= O(n^{-1/2}) + O(n^{-\alpha}) + O(n^{-1/2}) \leq O(n^{-\alpha/2}). \end{aligned}$$

Because R has finitely many flat segments, F^* has finitely many atoms by Lemma 2. As a result, the $O(n^{-\alpha/2})$-convergence is uniform in x.

Remark 12. When R is discontinuous, one can show that when all other players use the mean field equilibrium F^* in an n-player game, no alternative strategy F can generate an asymptotic gain exceeding half the maximum jump size of R. That is,

$$\limsup_n \int \xi_n^{F^*} dF - \int \xi_n^{F^*} dF^* \leq \frac{1}{2}\sup_k [g(y_k+) - g(y_k)] =: \varepsilon \quad \text{for all} \quad F \in \mathcal{F}, \tag{22}$$

where $\{y_k\} \subset (0,1)$ are the jump points of g. The proof uses

$$\int \xi_n^{F^*} dF - \int \xi_n^{F^*} dF^* \leq \int (\xi_n^{F^*} - \xi^{F^*})(dF - dF^*)$$

together with Lemmas 9 and 10. Note that the convergence in (22) is not necessarily uniform in F, so that (22) is weaker than the ε-Nash equilibrium property.

5.3. Convergence of the Optimal Reward Design

We have seen in Theorem 2 that the optimal design to maximize performance at a given target rank α is the cut-off reward at that same rank. As mentioned in the introduction, the best design in the prelimit is more complicated: for the n-player game with zero drift, the cutoff at a certain rank k_n^* is optimal for the expected performance at target rank k. A formula (recalled in the next paragraph) for k_n^* was found in Nutz and Zhang [42], and it is also noted that $k_n^* \geq k$, with $k_n^* > k$ unless k or k/n are small. For drift $\mu > 0$, a cutoff is again optimal, but the exact location of the cutoff is not known, whereas for $\mu < 0$, the optimal shape can look smoother than the sharp cutoff. In this section, we numerically compare the n-player game with the mean field limit for large n, focusing on $\mu = 0$ in order to have an exact result available for finite n.

We recall from Nutz and Zhang [42, proposition 3.11] that the optimal normalized reward for the expected kth rank performance in the n-player game with $\mu = 0$ is the cutoff at k_n^* (i.e., $R_i = 1/k_n^*$ for $i \leq k_n^*$ and $R_i = 0$ for $i > k_n^*$), where k_n^* is determined as

$$k_n^* = \max\left\{ j \geq k : \phi(k,j) \geq \frac{1}{j-1}\sum_{l=1}^{j-1}\phi(k,l)\right\}, \quad \phi(k,l) := \frac{(2n-k-l)!(k+l-2)!}{(n-l)!(l-1)!}.$$

The corresponding expected kth rank performance is

$$nx_0 \frac{n!}{(2n-1)!}\binom{n-1}{k-1}\frac{1}{k_n^*}\sum_{l=1}^{k_n^*}\phi(k,l). \tag{23}$$

If we scale k proportionally to n by fixing $k/n \approx \alpha \in (0,1)$, we can compare the optimal cutoff ratio k_n^*/n with the mean field optimal cutoff α. In the numerical example, we consider the median performance (i.e., $\alpha = 0.5$). A similar behavior can be observed for other choices of α.

Figure 1 shows that k_n^*/n converges to α as $n \to \infty$. The convergence is rather slow; for example, for $n = 1{,}024$, the optimal cutoff rank is still more than 9% larger than the mean-field optimum. This already suggests that using the mean field optimal design as a proxy for the n-player design may be problematic at least for moderate n.

Next, we consider the quality of the mean field proxy from the point of view of the principal: we fix the optimal design R^* from the mean field setting (Theorem 2) and compare the resulting expected performance in the n-player game with the performance (23) of the exact optimizer given by k_n^*. For comparison, we mention that the analogous question was considered in the Poissonian model of Nutz and Zhang [41] for the same performance functional of the principal, and there, the mean field proxy was shown to be $O(1/n)$-optimal for the n-player design problem.

Figure 2(a) shows not only that the performance of the proxy may be significantly inferior for finite n but, indeed, also that the performances diverge as $n \to \infty$, with the exact solution performing twice as well. The performance of the exact solution converges to the optimal performance in the mean field model as stated in Theorem 2, $x_\alpha^* = h^{-1}(1/\alpha) = 2x_0$, but the performance of the proxy does not.

Figure 1. (Color online) (a) Convergence of the optimal cutoff ratio k_n^*/n to α; (b) log-log plot of the difference $k_n^*/n - \alpha$, illustrating that k_n^*/n converges to α at a rate of approximately $O(n^{-r})$ for a fractional power r. Increments of n in all plots are chosen such as to avoid rounding effects related to the fact that k_n^* must be integer. On a finer scale for n, there are oscillations (see Nutz and Zhang [42, figure 3]) that, however, disappear in the large n limit.

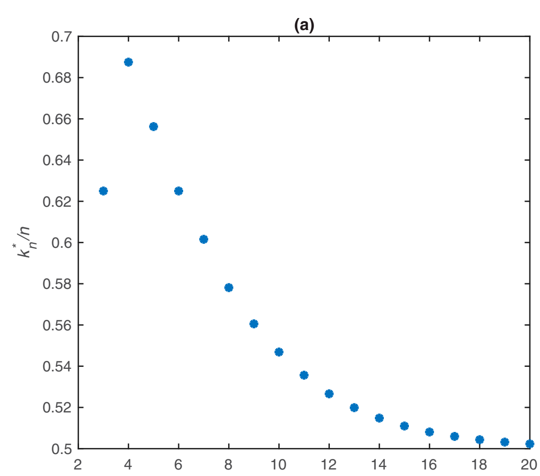

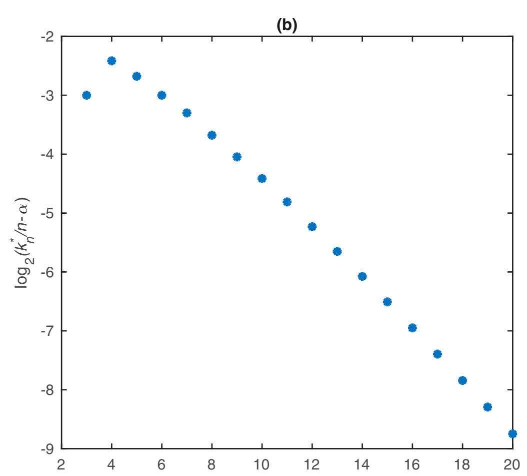

Figure 2. (Color online) (a) Performance of the mean field proxy diverges from the optimal design given by k_n^*; (b) median player's performance for all cutoff schemes when $n = 2^4, 2^6, 2^8, 2^{10}$. Red circles correspond to the mean field proxy (cutoff at rank αn or ratio α), and stars correspond to the exact n-player optimizer (cutoff at rank k_n^* or ratio k_n^*/n). As n increases, the blue and red points converge in the horizontal direction but nevertheless diverge in the vertical direction. Here, $\alpha = 0.5$ and $x_0 = 1$.

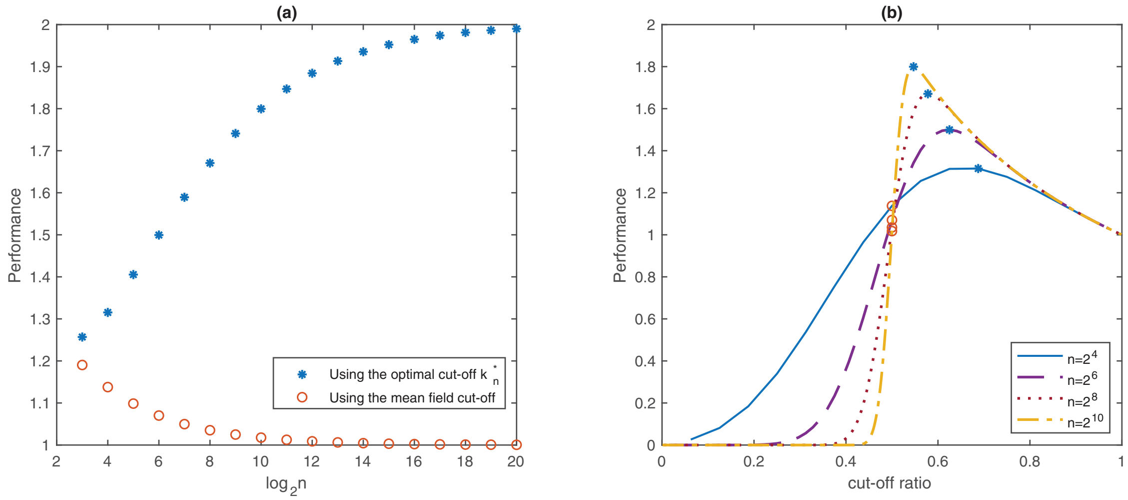

Figure 2(b) plots the same data points for some values of n, together with curves showing the performance of any cutoff strategy as a function of the cutoff location. For larger n, the curves are increasingly steep in a left neighborhood of the maximum: the vertical distance between the data points increases even though the horizontal distance decreases. In other words, the performance of R^* is increasingly inferior despite the cutoff location approximating the optimal location.

The reason lies in the lack of smoothness of the mean field game. Indeed, we know that the equilibrium distribution F_n^* induced by R^* in the n-player game converges weakly to the mean field equilibrium , which is a two-point distribution (Theorems 2 and 3). Although F_n^* is increasingly concentrated on the location of the limiting atoms at 0 and x_α^* for large n, the distribution is still smooth with connected support for finite n, so that the $(1-\alpha)$-quantile stretches far beyond x_α^*, causing the inferior performance.

We emphasize that the reason for the poor quality of the proxy observed here is very different from the knife-edge phenomenon leading to the negative result in Theorem 4 and quite possibly more relevant to applications.

Endnotes

[1] "Increase" and "decrease" are understood in the nonstrict sense in this paper.

[2] We use the same symbol for the distribution and its cdf when there is no danger of confusion. Note that if F has an atom at x, many players may share the same rank.

[3] See Obłój [43, section 9] for general background and a derivation. The extension to the present case with an absorbing boundary is immediate.

[4] A function $f : [a,b] \to \mathbb{R}$ is said to be piecewise α-Hölder continuous if $[a,b]$ is the union of finitely many intervals on which f is α-Hölder continuous.

[5] When R is globally Lipschitz continuous, the assertion also holds with $\alpha = \alpha' = 1$.

References

[1] Ankirchner S, Kazi-Tani N, Wendt J, Zhou C (2021) Large ranking games with non-observable diffusion control. Preprint, submitted February 11, https://hal.archives-ouvertes.fr/hal-03138716.
[2] Bardi M (2012) Explicit solutions of some linear-quadratic mean field games. *Networks Heterogeneous Media* 7(2):243–261.
[3] Bayraktar E, Cvitanić J, Zhang Y (2019) Large tournament games. *Ann. Appl. Probab.* 29(6):3695–3744.
[4] Bensoussan A, Frehse J, Yam SCP (2013) *Mean Field Games and Mean Field Type Control Theory* (Springer, New York).
[5] Campi L, Fischer M (2018) *N*-player games and mean-field games with absorption. *Ann. Appl. Probab.* 28(4):2188–2242.
[6] Cardaliaguet P, Delarue F, Lasry JM, Lions PL (2019) *The Master Equation and the Convergence Problem in Mean Field Games* (Princeton University Press, Princeton, NJ).
[7] Carmona R, Delarue F (2013) Probabilistic analysis of mean-field games. *SIAM J. Control Optim.* 51(4):2705–2734.
[8] Carmona R, Delarue F (2017) *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games* (Springer, Cham, Switzerland).
[9] Carmona R, Delarue F (2017) *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations* (Springer, Cham, Switzerland).

[10] Carmona R, Lacker D (2015) A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.* 25(3):1189–1231.
[11] Carmona R, Wang P (2021) Finite-state contract theory with a principal and a field of agents. *Management Sci.* 67(8):4725–4741.
[12] Carmona R, Delarue F, Lacker D (2017) Mean field games of timing and models for bank runs. *Appl. Math. Optim.* 76(1):217–260.
[13] Cecchin A, Fischer M (2020) Probabilistic approach to finite state mean field games. *Appl. Math. Optim.* 81(2):253–300.
[14] Cecchin A, Pra PD, Fischer M, Pelino G (2019) On the convergence problem in mean field games: A two state model without uniqueness. *SIAM J. Control Optim.* 57(4):2443–2466.
[15] Delarue F, Tchuendom RF (2020) Selection of equilibria in a linear quadratic mean-field game. *Stochastic Processes Appl.* 130(2):1000–1040.
[16] Demski JS, Sappington D (1984) Optimal incentive contracts with multiple agents. *J. Econom. Theory* 33(1):152–171.
[17] Djete M (2020) Mean field games of controls: On the convergence of Nash equilibria. Preprint, submitted June 19, https://doi.org/10.48550/arXiv.2006.12993.
[18] Djete M (2021) Large population games with interactions through controls and common noise: Convergence results and equivalence between open-loop and closed-loop controls. Preprint, submitted August 6, https://doi.org/10.48550/arXiv.2108.02992.
[19] Durrett R (2010) *Probability: Theory and Examples*, 4th ed. (Cambridge University Press, Cambridge, UK).
[20] Elie R, Mastrolia T, Possamaï D (2019) A tale of a principal and many, many agents. *Math. Oper. Res.* 44(2):440–467.
[21] Fang D, Noe T (2016) Skewing the odds: Taking risks for rank-based rewards. Preprint, submitted March 15, http://dx.doi.org/10.2139/ssrn.2747496.
[22] Fang D, Noe T, Strack P (2020) Turning up the heat: The discouraging effect of competition in contests. *J. Political Econom.* 128(5):1940–1975.
[23] Feng H, Hobson D (2015) Gambling in contests modelled with diffusions. *Decisions Econom. Finance* 38(1):21–37.
[24] Feng H, Hobson D (2016) Gambling in contests with random initial law. *Ann. Appl. Probab.* 26(1):186–215.
[25] Feng H, Hobson D (2016) Gambling in contests with regret. *Math. Finance* 26(3):674–695.
[26] Fischer M (2017) On the connection between symmetric N-player games and mean field games. *Ann. Appl. Probab.* 27(2):757–810.
[27] Green JR, Stokey NL (1983) A comparison of tournaments and contracts. *J. Political Econom.* 91(3):349–364.
[28] Hall WJ (1969) Embedding submartingales in Wiener processes with drift, with applications to sequential analysis. *J. Appl. Probab.* 6(3):612–632.
[29] Harris M, Kriebel CH, Raviv A (1982) Asymmetric information, incentives and intrafirm resource allocation. *Management Sci.* 28(6):604–620.
[30] Holmstrom B (1982) Moral hazard in teams. *Bell J. Econom.* 13(2):324–340.
[31] Huang M, Malhamé RP, Caines PE (2006) Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Comm. Inform. Systems* 6(3):221–251.
[32] Iseri M, Zhang J (2021) Set values for mean field games. Preprint, submitted September 2, https://doi.org/10.48550/arXiv.2107.01661.
[33] Lacker D (2016) A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields* 165(3–4):581–648.
[34] Lacker D (2020) On the convergence of closed-loop Nash equilibria to the mean field game limit. *Ann. Appl. Probab.* 30(4):1693–1761.
[35] Lacker D, Le Flem L (2021) Closed-loop convergence for mean field games with common noise. Preprint, submitted July 7, https://doi.org/10.48550/arXiv.2107.03273.
[36] Lasry JM, Lions PL (2006) Jeux à champ moyen. I. Le cas stationnaire. *Comptes Rendus Math. Acad. Sci. Paris* 343(9):619–625.
[37] Martin JS, Nutz M, Tan X (2020) Convergence to the mean field game limit: A case study. *Ann. Appl. Probab.* 30(1):259–286.
[38] Mookherjee D (1984) Optimal incentive schemes with many agents. *Rev. Econom. Stud.* 51(3):433–446.
[39] Nalebuff BJ, Stiglitz JE (1983) Prizes and incentives: Toward a general theory of compensation and competition. *Bell J. Econom.* 14(1):21–43.
[40] Nutz M (2018) A mean field game of optimal stopping. *SIAM J. Control Optim.* 56(2):1206–1221.
[41] Nutz M, Zhang Y (2019) A mean field competition. *Math. Oper. Res.* 44(4):1245–1263.
[42] Nutz M, Zhang Y (2022) Reward design in risk-taking contests. *SIAM J. Financial Math.* 13(1):129–146.
[43] Obłój J (2004) The Skorokhod embedding problem and its offspring. *Probab. Surveys* 1:321–390.
[44] Resnick SI (2007) *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling* (Springer, New York).
[45] Seel C (2015) Gambling in contests with heterogeneous loss constraints. *Econom. Lett.* 136(November):154–157.
[46] Seel C, Strack P (2013) Gambling in contests. *J. Econom. Theory* 148(5):2033–2048.
[47] Sun Y (2006) The exact law of large numbers via Fubini extension and characterization of insurable risks. *J. Econom. Theory* 126(1):31–69.
[48] Vojnović M (2016) *Contest Theory: Incentive Mechanisms and Ranking Methods* (Cambridge University Press, New York).