MEAN FIELD AND N-AGENT GAMES FOR OPTIMAL INVESTMENT UNDER RELATIVE PERFORMANCE CRITERIA

DANIEL LACKER* AND THALEIA ZARIPHOPOULOU†

Abstract. We analyze a family of portfolio management problems under relative performance criteria, for fund managers having CARA or CRRA utilities and trading in a common time horizon in log-normal markets. We construct explicit time-independent equilibrium strategies for both the finite population games and the corresponding mean field games, which we show are unique in the class of time-independent equilibria. In the CARA case, competition drives agents to invest more in the risky asset than they would otherwise, while in the CRRA case competitive agents may over- or under-invest, depending on their levels of risk tolerance.

1. Introduction

This paper is a contribution to both the theory of finite population and mean field games and to optimal portfolio management under competition and relative performance criteria. For the former, we construct explicit solutions for both n-player and mean field games, providing a new family of tractable solutions. For the latter, we formulate a new class of competition and relative performance optimal investment problems for agents having exponential (CARA) and power (CRRA) utilities, for both a finite number and a continuum of agents.

The finite-population case consists of n *fund managers* (or *agents*) trading between a common riskless bond and an individual stock. The price of each stock is modeled as a log-normal process driven by two independent Brownian motions. The first Brownian motion is the same for all prices, representing a common market noise, while the second is idiosyncratic, specific to each individual stock. Precisely, the i^{th} fund specializes in stock i whose price $(S_t^i)_{t\geq 0}$ is given by

$$\frac{dS_t^i}{S_t^i} = \mu_i dt + \nu_i dW_t^i + \sigma_i dB_t, \tag{1}$$

with constant market parameters $\mu_i > 0$, $\sigma_i \geq 0$, and $\nu_i \geq 0$, with $\sigma_i + \nu_i > 0$. The (one-dimensional) standard Brownian motions $B, W^1, W^2, \ldots, W^n$ are independent. When $\sigma_i > 0$, the process B induces a correlation between the stocks, and thus we call B the *common noise* and W^i an *indiosyncratic noise*.

Our setup covers the important special case in which all funds trade in the same stock; that is, $\mu_i = \mu$, $\nu_i = 0$, and $\sigma_i = \sigma$ for all $i = 1, \ldots, n$, for some $\mu, \sigma > 0$ independent of i. In this setting, all stocks are identical and the model is to be interpreted as n agents investing in a single common stock. These n agents differ in their risk preferences but otherwise face the same market opportunities.

All fund managers share a common time horizon, $T > 0$, and aim to maximize their expected utility at T. The utility functions $U_1, \ldots, U_n$ are agent-specific functions of both terminal wealth, X_T^i, and a "competition component," $\overline{X}_T$, which depends on the terminal wealths of all agents. We study two representative cases, related to the popular exponential and power utilities.

*Division of Applied Mathematics, Brown University. Supported by the National Science Foundation under Award No. DMS 1502980; daniel_lacker@brown.edu.

†Dpts. of Mathematics and IROM, The University of Texas at Austin and the Oxford-Man Institute, University of Oxford; zariphop@math.utexas.edu.

For the exponential case, we assume that competition affects the wealth additively and is modeled through the arithmetic average wealth of all agents,

$$U_i\left(X_T^i, \overline{X}_T\right) = -e^{-\frac{1}{\delta_i}\left(X_T^i - \theta_i \overline{X}_T\right)}, \quad \text{where} \quad \overline{X}_T = \frac{1}{n}\sum_{k=1}^{n} X_T^k. \tag{2}$$

The parameters $\delta_i > 0$ and $\theta_i \in [0,1]$ represent the i^{th} agent's absolute risk tolerance and absolute competition weight, with small (resp. high) values of θ_i denoting low (resp. high) relative performance concern. This model is similar to and largely inspired by that of Espinosa and Touzi [18].

For the power case, the competition affects the wealth multiplicatively and is modeled through the geometric average wealth of all agents,

$$U_i\left(X_T^i, \overline{X}_T\right) = \frac{1}{1-1/\delta_i}\left(X_T^i \overline{X}_T^{-\theta_i}\right)^{1-1/\delta_i}, \quad \text{where} \quad \overline{X}_T = \left(\prod_{k=1}^{n} X_T^k\right)^{1/n}. \tag{3}$$

Now, the parameters $\delta_i > 0$ and $\theta_i \in [0,1]$ represent the i^{th} agent's relative risk tolerance and relative competition weight.

The aim is to identify Nash equilibria, namely, to find investment strategies $(\pi_t^{1,*}, \ldots, \pi_t^{n,*})_{t\in[0,T]}$ such that $\pi_t^{i,*}$ is the optimal stock allocation exercised by the i^{th} agent in response to the strategy choices of all other competitors, for $i = 1, \ldots, n$. As is usually the case for exponential and power risk preferences, $\pi_t^{i,*}$ is taken to be the absolute wealth and the fraction of wealth invested in the i^{th} stock, respectively.

Competition among fund managers is well documented in investment practice for both mutual and hedge funds; see, for example, [1, 2, 7, 14, 16, 21, 30, 34, 42]. As it is argued in these works, competition can stem, for example, from career advancement motives, seeking higher money inflows from their clients, preferential compensation contracts. In most of these works, only the case of two managers has been considered and in discrete time (or two period) models, with variations of criteria involving risk neutrality, relative performance with respect to an absolute benchmark or a critical threshold, or constraints on the managers' risk aversion parameters. More recently, the authors in [2] proposed a continuous time log-normal model for two fund managers with power utilities.

Asset specialization is also well documented in the finance literature, starting with Brennan [6, 15, 37]. Other representative works include [5, 29, 35, 38, 40, 41, 45]. As it is argued in these works, a variety of factors prompt managers to specialize in individual stocks or asset classes, such as familiarity, learning cost reduction, ambiguity aversion, solvency requirements, trading costs and constraints, liquidation risks, and informational frictions.

For tractability, we search only for Nash equilibria in which the investment strategies are time-independent. This restriction is quite natural, given the log-normality of the stock prices, the scaling properties of the CARA and CRRA utilities, and the form of the associated competition components. To construct such an equilibrium, we first solve each single agent's optimization problem given an arbitrary choice of competitors' strategies. Incorporating the competition component $\overline{X}$ as an additional uncontrolled state process leads to a single Hamilton-Jacobi-Bellman (HJB) equation, which we show has a separable smooth solution. Together with the first order conditions, this yields the candidate policies in a closed-form. We then construct the equilibrium through a set of compatibility conditions, which also provide criteria for existence and uniqueness. As an intermediate step, we use arguments from indifference valuation to obtain verification results for these smooth solutions. Specifically, we interpret each HJB equation as the one solved by the writer of an individual liability determined by the competition component.

The unique time-independent Nash equilibrium in each model turns out to be the sum of two components. The first is the traditional Merton-type portfolio (c.f. [36]), which is optimal for the individual expected utility problem without any relative performance concerns. The second component depends on the individual competition parameter and on other quantities involving the risk tolerance and competition parameters of all agents as well as the market parameters of all stocks. Naturally, this second component disappears when there is no competition.

In the exponential model, it turns out that competition always results in higher investment in the risky asset. This is not, however, the case for the power model, mainly because the sign of the second component might not be always fixed. This sign depends on the value of the relative risk tolerance, particularly whether it is larger or smaller than one; this is to be expected given well known properties of CRRA utilities and their optimal portolios (see, for example, the so called "nirvana" cases in [31]).

In the noteworthy special case of a single stock, common to all agents, the equilibrium strategies are simpler. For both the exponential and the power cases, the Nash equilibrium is of Merton type but with a modified risk tolerance, which depends linearly on the individual risk tolerance and competition parameters, with the coefficients of this linear function depending on the population averages of these parameters.

The expressions for the equilibrium strategies simplify when the number of agents n tends to infinity. The limiting expressions depend solely on the limit of the empirical distribution of the *type vectors* $\zeta_i = (x_0^i, \delta_i, \theta_i, \mu_i, \nu_i, \sigma_i)$, for $i = 1, \ldots, n$. We show that these limiting strategies can be derived intrinsically, as equilibria of suitable mean field games (MFGs). Intuitively, the finite set of agents becomes a continuum, with each individual trading between the common bond and her individual stock while also competing with the rest of the (infinite) population through a relative performance functional affecting her expected terminal utility.

Although explicit solutions are available for our n-player games, the MFG framework is worth introducing in this context in part because it extends naturally to more complex models, such as those involving portfolio constraints or general utility functions. In such models, we expect the MFG framework to be more tractable than the n-player games. For instance, [18, 19] study n-player models similar to our CARA utility model but notably including portfolio constraints, and both papers run into a difficult n-dimensional quadratic BSDE system. A MFG formulation would likely be more tractable, at least reducing the dimensionality of the problem, though we do not tackle such an analysis in this paper.

The MFG is defined in terms of a representative agent who is assigned a *random* type vector $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ at time zero, determining her initial wealth ξ, preference parameters (δ, θ), and market parameters (μ, ν, σ). The randomness of the type vector encodes the distribution of the (continuum of) agents' types.

For the exponential case, the MFG problem is to find a pair $\left(\pi^*, \overline{X}\right)$ with the following properties. The investment strategy π^* optimizes, in analogy to (2),

$$\sup_{\pi} \mathbb{E}\left[-e^{-\frac{1}{\delta}\left(X_T - \overline{X}\right)}\right], \tag{4}$$

where X_T is the wealth of the representative agent and $\overline{X}$ the average wealth of the continuum of agents. Furthermore, at this optimum, the consistency condition $\overline{X} = \mathbb{E}[X_T^* | \mathcal{F}_T^B]$ must hold, where $(\mathcal{F}_t^B)_{t\in[0,T]}$ is the filtration generated by the common noise B, and X^* is the optimal wealth determined by π^*.

For the power case, the aggregate wealth $\overline{X}$ must be consistent with its n-player form in (3). With this in mind, note that the geometric mean of a positive random variable Y can be written as $\exp \mathbb{E}[\log Y]$, whether or not the distribution of Y is discrete. This points to the MFG problem of finding a pair $\left(\pi^*, \overline{X}\right)$, such that π^* optimizes

$$\sup_{\pi} \mathbb{E}\left[\frac{1}{1-1/\delta}\left(X_T \overline{X}^{-\theta}\right)^{1-1/\delta}\right], \tag{5}$$

and, furthermore, the consistency condition $\overline{X} = \exp \mathbb{E}[\log X_T^* | \mathcal{F}_T^B]$ holds.

As in the finite population CARA and CRRA cases, we focus on the tractable class of MFE in which the strategy π^* is time-independent. While constant in time, such strategies are still *random*, measurable with respect to the (time-zero-measurable) random type vector. We solve the MFG problems directly, constructing equilibria which agree with the limiting expressions from the n-player games. In each model, the solution technique is analogous to the n-player setting in that we treat the aggregate wealth term as an uncontrolled state process, find a smooth separable solution of a single HJB equation, and then enforce the consistency condition. The resulting MFG strategies take similar but notably simpler forms than their n-player counterparts and exhibit the same qualitative behavior and two-component structure discussed above.

Mean field games, first introduced in [33] and [25], have by now found numerous applications in economics and finance, notably including models of income inequality [20], economic growth [28], limit order book formation [22], systemic risk [12], optimal execution [27], and oligopoly market models [13], to name but a few. The closest works to ours are the static model of [24, Section 6], which is a competitive variant of the Markowitz model, and the stochastic growth model of Huang and Nguyen [26] which has some mathematical features in common with our power utility model. That said, our work appears to be the first application of MFG theory to portfolio optimization.

Our results add two new examples of *explicitly* solvable MFG models. Beyond the linear quadratic models of [4, 10, 12], such examples are scarce, especially in the presence of common noise. The only other examples we know of are those in [24, Sections 5 and 7] as well as the more recent [43], which is linear-quadratic aside from a square root diffusion term. In fact, our models permit an explicit solution of the so-called *master equation* (c.f. [9]). Moreover, the manner in which we incorporate different *types* of agents, by randomizing $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ as described above, seems to be novel; several works on MFGs have incorporated finitely many types [25] by tracking a vector of mean field interactions, one for each type, but our approach has the advantage of seamlessly incorporating (uncountably) infinitely many types.

The paper is organized as follows. In Section 2, we present the exponential model and study both the n-player game and the MFG. In Section 3, we present the analogous results for power and logarithmic utilities. For both classes, we provide qualitative comments on the Nash and mean field equilibria, in Sections 2.3 and 3.3, respectively. We conclude in Section 4 with a discussion of open questions and future research directions.

2. CARA risk preferences

We consider *fund managers* (henceforth, *agents*) with exponential risk preferences with constant individual (absolute) risk tolerances. Agents are also concerned with how their performance is measured in relation the one of their competitors. This is modeled as an additive penalty term depending on the average wealth, and weighted by an investor specific comparison parameter.

We begin our analysis with the exponential class because of its additive scaling properties, which allow for substantial tractability. Furthermore, the exponential class provides a direct connection with indifference valuation, used in solving the underlying HJB equation.

2.1. The n-agent game.

We introduce a game of n agents who trade in a common investment horizon $[0, T]$. Each of these agents trades between an "individual" stock and a riskless bond. The latter is common to all agents, serves as the numeraire and offers zero interest rate.

Stock prices are taken to be log-normal, as described in the introduction, each driven by two independent Brownian motions. Precisely, the price $(S_t^i)_{t\in[0,T]}$ of the stock i traded by the i^{th} agent solves (1), with given market parameters $\mu_i > 0$, $\sigma_i \geq 0$, and $\nu_i \geq 0$, with $\sigma_i + \nu_i > 0$. The independent Brownian motions $B, W^1, \ldots, W^n$ are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which we endow with the natural filtration $(\mathcal{F}_t)_{t\in[0,T]}$ generated by these $n+1$ Brownian motions. Recall that the *single stock case* is when

$$(\mu_i, \sigma_i) = (\mu, \sigma), \quad \text{and } \nu_i = 0, \quad \text{for } i = 1, \ldots, n,$$

for some $\mu, \sigma > 0$ independent of i. Notably, the single stock case was studied in [18] and [19] in greater generality, incorporating portfolio constraints and more general stock price dynamics.

Each agent $i = 1, \ldots, n$ trades using a self-financing strategy, $(\pi_t^i)_{t\in[0,T]}$, which represent the (discounted by the bond) amount invested in the i^{th} stock. The i^{th} agent's wealth $(X_t^i)_{t\in[0,T]}$ then solves

$$dX_t^i = \pi_t^i(\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t), \tag{6}$$

with $X_0^i = x_0^i \in \mathbb{R}$. A portfolio strategy is deemed admissible if it belongs to the set $\mathcal{A}$, which consists of self-financing $\mathbb{F}$-progressively measurable real-valued processes $(\pi_t)_{t\in[0,T]}$ satisfying $\mathbb{E}\int_0^T |\pi_t|^2 dt < \infty$.

The i^{th} agent's utility is a function $U_i : \mathbb{R}^2 \to \mathbb{R}$ of *both* her individual wealth, x, and the average wealth of all agents, m. It is of the form

$$U_i(x, m) := -\exp\left(-\frac{1}{\delta_i}(x - \theta_i m)\right).$$

We will refer to the constants $\delta_i > 0$ and $\theta_i \in [0, 1]$ as the *personal risk tolerance* and *competition weight* parameters, respectively.[1] If agents $i = 1, \ldots, n$ choose admissible strategies $\pi^1, \ldots, \pi^n$, the payoff for agent i is given by

$$J_i(\pi^1, \ldots, \pi^n) := \mathbb{E}\left[-\exp\left(-\frac{1}{\delta_i}\left(X_T^i - \theta_i \overline{X}_T\right)\right)\right], \quad \text{with} \quad \overline{X}_T = \frac{1}{n}\sum_{k=1}^n X_T^k, \tag{7}$$

where the dynamics of $(X_t^i)_{t\in[0,T]}$ are as in (6). Alternatively, we may express the above as

$$J_i(\pi^1, \ldots, \pi^n) = \mathbb{E}\left[-\exp\left(-\frac{1}{\delta_i}\left((1-\theta_i)X_T^i + \theta_i(X_T^i - \overline{X}_T)\right)\right)\right],$$

which highlights how the competition weight θ_i determines an agent's risk preference for *absolute wealth* versus *relative wealth*. An agent with large θ_i (close to one) is thus more concerned with relative wealth than absolute wealth.

These interdependent optimization problems are resolved competitively, applying the concept of Nash equilibrium in the above investment setting.

Definition 1. *A vector $(\pi^{1,*}, \ldots, \pi^{n,*})$ of admissible strategies is a (Nash) equilibrium if, for all $\pi^i \in \mathcal{A}$ and $i = 1, \ldots, n$,*

$$J_i(\pi^{1,*}, \ldots, \pi^{i,*}, \ldots, \pi^{n,*}) \geq J_i(\pi^{1,*}, \ldots, \pi^{i-1,*}, \pi^i, \pi^{i+1,*}, \ldots, \pi^{n,*}). \tag{8}$$

A constant (Nash) equilibrium is one in which, for each i, $\pi^{i,*}$ is constant in time, i.e., $\pi_t^{i,*} = \pi_0^{i,*}$ for all $t \in [0, T]$.[2]

[1] Note that (δ_i, θ_i), $i = 1, \ldots, n$, are unitless, because all wealth processes are discounted by the riskless bond.

[2] Our notion of Nash equilibrium is more accurately known as an *open-loop* Nash equilibrium. A popular alternative is *closed-loop* Nash equilibrium, in which agents choose strategies in terms of feedback functions as opposed to stochastic processes. However, for *constant* strategies, the open-loop and closed-loop concepts coincide. That is, a constant (open-loop) Nash equilibrium is also a closed-loop Nash equilibrium, and vice versa.

Remark 2. Because the filtration $\mathbb{F}$ is Brownian, it holds for any admissible strategy $\pi \in \mathcal{A}$ that π_0 is nonrandom. With this in mind, a constant Nash equilibrium will be identified with a vector $(\pi^{1,*}, \dots, \pi^{n,*}) \in \mathbb{R}^n$.

Our first main finding provides conditions for existence and uniqueness of a constant Nash equilibrium and also constructs it explicitly.

Theorem 3. *Assume that for all $i = 1, \dots, n$ we have $\delta_i > 0$, $\theta_i \in [0,1]$, $\mu_i > 0$, $\sigma_i \geq 0$, $\nu_i \geq 0$, and $\sigma_i + \nu_i > 0$. Define the constants*

$$\varphi_n := \frac{1}{n} \sum_{k=1}^{n} \delta_k \frac{\mu_k \sigma_k}{\sigma_k^2 + \nu_k^2(1 - \theta_k/n)} \qquad \text{and} \qquad \psi_n := \frac{1}{n} \sum_{k=1}^{n} \theta_k \frac{\sigma_k^2}{\sigma_k^2 + \nu_k^2(1 - \theta_k/n)}. \tag{9}$$

There are two cases:

(i) If $\psi_n < 1$, there exists a unique constant equilibrium, given by

$$\pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2(1 - \theta_i/n)} + \theta_i \frac{\sigma_i}{\sigma_i^2 + \nu_i^2(1 - \theta_i/n)} \frac{\varphi_n}{1 - \psi_n}. \tag{10}$$

Moreover, we have the identity

$$\frac{1}{n} \sum_{k=1}^{n} \sigma_k \pi^{k,*} = \frac{\varphi_n}{1 - \psi_n}.$$

(ii) If $\psi_n = 1$, there is no constant equilibrium.

An important corollary covers the special case of a single stock.

Corollary 4 (Single stock)**.** *Assume that for all $i = 1, \dots, n$ we have $\mu_i = \mu > 0$, $\sigma_i = \sigma > 0$, and $\nu_i = 0$, with μ and σ independent of i. Define the constants*

$$\overline{\delta} := \frac{1}{n} \sum_{k=1}^{n} \delta_k \qquad \text{and} \qquad \overline{\theta} := \frac{1}{n} \sum_{k=1}^{n} \theta_k.$$

There are two cases:

(i) If $\overline{\theta} < 1$, there exists a unique constant equilibrium, given by

$$\pi^{i,*} = \left(\delta_i + \theta_i \frac{\overline{\delta}}{1 - \overline{\theta}} \right) \frac{\mu}{\sigma^2}.$$

(ii) If $\overline{\theta} = 1$, there is no constant equilibrium.

Proof. Apply Theorem 3, taking note of the simplifications $\varphi_n = \overline{\delta}\mu/\sigma$ and $\psi_n = \overline{\theta}$. $\square$

Remark 5. For a given agent i, it is arguably more natural to replace the average wealth $\overline{X}_T$ in the payoff functional J_i defined in (7) with the average over all *other* agents, *not including herself*, i.e., $\overline{X}_T^{(-i)} = \frac{1}{n-1} \sum_{k \neq i} X_T^k$. Fortunately, there is a one-to-one mapping between the two formulations, so there is no need to solve both separately. Indeed, suppose the i^{th} agent's payoff is

$$\mathbb{E}\left[-\exp\left(-\frac{1}{\delta_i'} \left(X_T^i - \theta_i' \overline{X}_T^{(-i)} \right) \right) \right],$$

for some parameters $\theta_i' \in [0,1]$ and $\delta_i' > 0$. By matching coefficients it is straightforward to show that

$$\frac{1}{\delta_i'} \left(X_T^i - \theta_i' \overline{X}_T^{(-i)} \right) = \frac{1}{\delta_i} \left(X_T^i - \theta_i \overline{X}_T \right),$$

when $\theta_i \in [0,1]$ and $\delta_i > 0$ are defined by

$$\delta_i = \frac{\delta_i'}{1+\frac{1}{n-1}\theta_i'} \qquad \text{and} \qquad \theta_i = \frac{\theta_i'}{\frac{n-1}{n}+\frac{1}{n}\theta_i'}.$$

We prefer our original formulation mainly because it results in simpler formulas for the equilibrium strategies in Theorem 3 and Corollary 4. Moreover, for large n, this choice has negligible effect on the strategy $\pi^{i,*}$, as the differences $|\delta_i - \delta_i'|$ and $|\theta_i - \theta_i'|$ vanish.

Proof of Theorem 3. Let i be fixed. Assume that all other agents, $k \neq i$, follow constant investment strategies, denoted by $\alpha_k \in \mathbb{R}$. Let $(X_t^k)_{t\in[0,T]}$ be the associated wealth processes,

$$X_t^k = x_0^k + \alpha_k\left(\mu_k t + \nu_k W_t^k + \sigma_k B_t\right),$$

and also define

$$Y_t := \frac{1}{n}\sum_{k\neq i} X_t^k.$$

Then, the i^{th} agent solves the optimization problem

$$\sup_{\pi\in\mathcal{A}} \mathbb{E}\left[-\exp\left(-\frac{1}{\delta_i}\left(\left(1-\frac{\theta_i}{n}\right)X_T^i - \theta_i Y_T\right)\right)\right], \tag{11}$$

where $(X_t^i)_{t\in[0,T]}$ and $(Y_t)_{t\in[0,T]}$ have dynamics (c.f. (6))

$$\begin{aligned} dX_t^i &= \pi_t(\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t), & X_0^i &= x_0^i,\\ dY_t &= \widehat{\mu\alpha}dt + \widehat{\sigma\alpha}dB_t + \frac{1}{n}\sum_{k\neq i}\nu_k\alpha_k dW_t^k, & Y_0 &= \frac{1}{n}\sum_{k\neq i} x_0^k, \end{aligned}$$

and where we have abbreviated

$$\widehat{\mu\alpha} := \frac{1}{n}\sum_{k\neq i}\mu_k\alpha_k \qquad \text{and} \qquad \widehat{\sigma\alpha} := \frac{1}{n}\sum_{k\neq i}\sigma_k\alpha_k.$$

In the sequel, we will also use the abbreviation

$$\widehat{(\nu\alpha)^2} := \frac{1}{n}\sum_{k\neq i}\nu_k^2\alpha_k^2.$$

One then sees that the stochastic optimization problem (11) can be alternatively viewed as the one solved by an agent who is the "writer" of a liability $G(Y_T) := \frac{\theta_i}{1-\theta_i/n}Y_T$, having exponential preferences with risk aversion $\gamma_i := \frac{1}{\delta_i}\left(1-\frac{\theta_i}{n}\right)$. This problem was explicitly solved in [39], but for the reader's convenience we include some of the details here.

The value of the supremum in (11) is equal to $v(X_0^i, Y_0, 0)$, where $v(x,y,t)$ solves the Hamilton-Jacobi-Bellman (HJB) equation

$$\begin{aligned} v_t + \max_{\pi\in\mathbb{R}}&\left(\frac{1}{2}(\sigma_i^2+\nu_i^2)\pi^2 v_{xx} + \pi\left(\mu_i v_x + \sigma_i\widehat{\sigma\alpha}v_{xy}\right)\right) \\ &+\frac{1}{2}\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right)v_{yy} + \widehat{\mu\alpha}v_y = 0, \end{aligned} \tag{12}$$

for $(x,y,t) \in \mathbb{R}\times\mathbb{R}\times[0,T]$, with terminal condition

$$v(x,y,T) = -e^{-\gamma_i(x-G(y))} = -\exp\left(-\frac{1}{\delta_i}\left(\left(1-\frac{\theta_i}{n}\right)x - \theta_i y\right)\right).$$

Applying the first order conditions, equation (12) reduces to

$$v_t - \frac{1}{2}\frac{(\mu_i v_x + \sigma_i \widehat{\sigma\alpha} v_{xy})^2}{(\sigma_i^2+\nu_i^2)v_{xx}} + \frac{1}{2}\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right) v_{yy} + \widehat{\mu\alpha} v_y = 0.$$

Making the ansatz $v(x,y,t) = -\exp\left(-\frac{1}{\delta_i}((1-\frac{\theta_i}{n})x - \theta_i y)\right) f(t)$ yields, for $t \in [0,T]$,

$$f'(t) - \rho f(t) = 0.$$

with $f(T) = 1$ and

$$\rho := \frac{(\mu_i + \theta_i \delta_i^{-1}\sigma_i \widehat{\sigma\alpha})^2}{2(\sigma_i^2+\nu_i^2)} - \theta_i \widehat{\mu\alpha} - \frac{1}{2}\theta_i^2\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right). \tag{13}$$

Therefore, $f(t) = e^{-\rho(T-t)}$ and, in turn,

$$v\,(x,y,t) = -\exp\left(-\frac{1}{\delta_i}\left(\left(1-\frac{\theta_i}{n}\right)x - \theta_i y\right) - \rho(T-t)\right). \tag{14}$$

We note that the above HJB equation has a unique smooth solution given by v above, and thus the solution of the n-agent game is also unique. Indeed, uniqueness is established as follows. First, one can show that the HJB equation (12) has a unique viscosity solution in the class of functions that are strictly concave and strictly increasing in x (see, for example, [17]). On the other hand, v is by construction smooth as well as strictly concave and strictly increasing in x, and thus it coincides with this unique weak solution.

The i^{th} agent's optimal feedback control achieves the maximum in (12) and is thus given by

$$\pi^{i,*}(x,y,t) := -\frac{\mu_i v_x(x,y,t) + \sigma_i \widehat{\sigma\alpha} v_{xy}(x,y,t)}{(\sigma_i^2+\nu_i^2)v_{xx}(x,y,t)}.$$

Direct calculations yield that $\pi^{i,*}$ is constant,

$$\pi^{i,*} = \frac{\delta_i^{-1}(1-\theta_i/n)(\mu_i + \theta_i\delta_i^{-1}\sigma_i\widehat{\sigma\alpha})}{(\sigma_i^2+\nu_i^2)\delta_i^{-2}(1-\theta_i/n)^2} = \frac{\delta_i\mu_i + \theta_i\sigma_i\widehat{\sigma\alpha}}{(\sigma_i^2+\nu_i^2)(1-\theta_i/n)}.$$

Therefore, for a candidate portfolio vector $(\alpha_1,\ldots,\alpha_n)$ to be a Nash equilibrium, we need $\pi^{i,*} = \alpha_i$, for $i = 1,\ldots,n$. Let

$$\overline{\sigma\alpha} := \frac{1}{n}\sum_{k=1}^{n}\sigma_k\alpha_k = \widehat{\sigma\alpha} + \frac{1}{n}\sigma_i\alpha_i.$$

Then, we must have

$$\alpha_i = \pi^{i,*} = \frac{\delta_i\mu_i + \theta_i\sigma_i\overline{\sigma\alpha}}{(\sigma_i^2+\nu_i^2)(1-\theta_i/n)} - \frac{\theta_i\sigma_i^2}{n(\sigma_i^2+\nu_i^2)(1-\theta_i/n)}\alpha_i,$$

which implies that

$$\begin{aligned}
\alpha_i &= \frac{\delta_i\mu_i + \theta_i\sigma_i\overline{\sigma\alpha}}{(\sigma_i^2+\nu_i^2)(1-\theta_i/n)}\left(1 + \frac{\theta_i\sigma_i^2}{n(\sigma_i^2+\nu_i^2)(1-\theta_i/n)}\right)^{-1} \\
&= \frac{\delta_i\mu_i + \theta_i\sigma_i\overline{\sigma\alpha}}{(\sigma_i^2+\nu_i^2)(1-\theta_i/n) + \sigma_i^2\theta_i/n} = \frac{\delta_i\mu_i + \theta_i\sigma_i\overline{\sigma\alpha}}{\sigma_i^2+\nu_i^2(1-\theta_i/n)}.
\end{aligned} \tag{15}$$

Multiplying both sides by σ_i and then averaging over $i = 1,\ldots,n$, gives

$$\overline{\sigma\alpha} = \varphi_n + \psi_n\overline{\sigma\alpha}, \tag{16}$$

with φ_n, ψ_n as in (9). For equality (15) to hold, equality (16) must hold as well. There are three cases:

(i) If $\psi_n < 1$, then (16) yields $\overline{\sigma\alpha} = \varphi_n/(1-\psi_n)$, and the equilibrium control is well defined and given by (10).

(ii) If $\psi_n = 1$ and $\varphi_n > 0$, then equation (16) has no solution and thus no constant equilibria exist.

(iii) The remaining case is $\psi_n = 1$ and $\varphi_n = 0$, in which case equation (16) has infinitely many solutions. This, however, cannot occur. Indeed, because $\delta_i, \mu_i > 0$ for all i, we can only have $\varphi_n = 0$ if $\sigma_i = 0$ for all i. Recalling that $\sigma_i + \nu_i > 0$ by assumption, this implies $\psi_n = 0$, which is a contradiction.

□

Remark 6. One can also compute the equilibrium value function $v(x, y, t)$ of agent i, by explicitly computing ρ defined in (13), as the quantities $\widehat{\mu\alpha}$, $\widehat{\sigma\alpha}$, and $\widehat{\nu\alpha}$ are now known. However, we omit this tedious calculation.

2.2. **The mean field game.** In this section we study the limit as $n \to \infty$ of the n-player game analyzed in the previous section.

We start with an informal argument, to build intuition and motivate the upcoming definition. For the n-player game, we define for each agent $i = 1, \ldots, n$ the *type vector*

$$\zeta_i := (x_0^i, \delta_i, \theta_i, \mu_i, \nu_i, \sigma_i).$$

These type vectors induce an empirical measure, called the *type distribution*, which is the probability measure on the *type space*

$$\mathcal{Z}^e := \mathbb{R} \times (0,\infty) \times [0,1] \times (0,\infty) \times [0,\infty) \times [0,\infty), \tag{17}$$

given by

$$m_n(A) = \frac{1}{n}\sum_{i=1}^n 1_A(\zeta_i), \quad \text{for Borel sets } A \subset \mathcal{Z}^e.$$

We then see that for each agent i the equilibrium strategy $\pi^{i,*}$ computed in Theorem 3 depends only on her own type vector ζ_i and the distribution m_n of all type vectors. Indeed, the constants φ_n and ψ_n (c.f. (9)) are obtained simply by integrating appropriate functions under m_n.

Assume now that as the number of agents becomes large, $n \to \infty$, the above empirical measure m_n has a weak limit m, in the sense that $\int_{\mathcal{Z}^e} f\,dm_n \to \int_{\mathcal{Z}^e} f\,dm$ for every bounded continuous function f on $\mathcal{Z}^e$. For example, this holds almost surely if the ζ_i's are i.i.d. samples from m. Let $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ denote a random variable with this limiting distribution m. Then, we should expect the optimal strategy $\pi^{i,*}$ (c.f. (10)) to converge to

$$\lim_{n\to\infty} \pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2} + \theta_i \frac{\sigma_i}{\sigma_i^2 + \nu_i^2}\frac{\varphi}{1-\psi}, \tag{18}$$

where

$$\varphi := \lim_{n\to\infty} \varphi_n = \mathbb{E}\left[\delta \frac{\mu\sigma}{\sigma^2 + \nu^2}\right] \qquad \text{and} \qquad \psi := \lim_{n\to\infty} \psi_n = \mathbb{E}\left[\theta \frac{\sigma^2}{\sigma^2+\nu^2}\right].$$

The *mean field game* (MFG) defined next allows us to derive the limiting strategy (18) as the outcome of a self-contained equilibrium problem, which intuitively represents a game with a continuum of agents with type distribution m. Rather than directly modeling a continuum of agents, we follow the MFG paradigm of modeling a single representative agent, who we view as randomly selected from the population. The probability measure m represents the distribution of type parameters among the continuum of agents; equivalently, the representative agent's type vector is a random variable with law m. Heuristically, each agent in the continuum trades in a single stock driven by two Brownian motions, one of which is unique to this agent and one of which is common to all agents. The equilibrium concept introduced in Definition 8 will formalize this intuition.

2.2.1. *Formulating the mean field game.* To formulate the MFG, we now assume that the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ supports yet another independent (one-dimensional) Brownian motion, W, as well as a random variable

$$\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma),$$

independent of W and B, and with values in the space $\mathcal{Z}^e$ defined in (17). This random variable ζ is called the *type vector*, and its distribution is called the *type distribution*.

Let $\mathbb{F}^{\mathrm{MF}} = (\mathcal{F}_t^{\mathrm{MF}})_{t\in[0,T]}$ denote the smallest filtration satisfying the usual assumptions for which ζ is $\mathcal{F}_0^{\mathrm{MF}}$-measurable and W and B are adapted. Let also $\mathbb{F}^B = (\mathcal{F}_t^B)_{t\in[0,T]}$ denote the natural filtration generated by the Brownian motion B.[3]

The *representative agent's* wealth process solves

$$dX_t = \pi_t(\mu dt + \nu dW_t + \sigma dB_t), \quad X_0 = \xi, \tag{19}$$

where the portfolio strategy must belong to the admissible set $\mathcal{A}_{\mathrm{MF}}$ of self-financing $\mathbb{F}^{\mathrm{MF}}$-progressively measurable real-valued processes $(\pi_t)_{t\in[0,T]}$ satisfying $\mathbb{E}\int_0^T |\pi_t|^2 dt < \infty$. The random variable ξ is the initial wealth of the representative agent, whereas (μ, ν, σ) are the market parameters. In the sequel, the parameters δ and θ will affect the risk preferences of the representative agent.

In this mean field setup, the *single stock case* refers to the case where (μ, ν, σ) is nonrandom, with $\nu = 0$, $\mu > 0$, and $\sigma > 0$. In the context of the limiting argument above, this corresponds to the n-player game in which $\mu_i = \mu$, $\nu_i = \nu = 0$, and $\sigma_i = \sigma$ for all i. Note that each agent among the continuum may still have different preference parameters, captured by the fact that δ and θ are random.

Remark 7. There are two distinct sources of randomness in this model. One source comes from the Brownian motions W and B which drive the stock price processes over time. The second source is static and comes from the random variable ζ, which describes the distribution of type vectors (which includes initial wealth, individual preference parameters, and market parameters) among a large (in fact, continuous) population. One can then think of a continuum of agents, each of whom is assigned an i.i.d. type vector at time zero, and the agents interact *after* these assignments are made.

To see how to formulate the representative agent's optimization problem, let us first recall how the Nash equilibrium in the n-player game was constructed. We first solved the optimization problem (12) faced by each single agent i, in which the strategies of the other agents $k \neq i$ were treated as fixed. However, instead of fixing the strategies of the other agents, we could have just fixed the mean terminal wealth $\frac{1}{n}\sum_{k\neq i} X_T^k$, as this is effectively the only source of interaction between the agents. This was precisely the idea behind the proof of Theorem 3, and this guides the upcoming formulation of the MFG.

To this end, suppose that $\overline{X}$ is a given random variable, representing the average wealth of the continuum of agents. The representative agent has no influence on $\overline{X}$, as but one agent amid a continuum. The objective of the representative agent is thus to maximize the expected payoff

$$\sup_{\pi\in\mathcal{A}_{\mathrm{MF}}} \mathbb{E}\left[-\exp\left(-\frac{1}{\delta}\left(X_T - \theta\overline{X}\right)\right)\right], \tag{20}$$

[3]One might as well formulate the MFG on a different probability space from the n-player game, but we prefer to avoid introducing additional notation.

where $(X_t)_{t\in[0,T]}$ is given by (19). We are now ready to introduce the main definition of this section.

Definition 8. *Let $\pi^* \in \mathcal{A}_{\mathrm{MF}}$ be an admissible strategy, and consider the $\mathcal{F}_T^B$-measurable random variable $\overline{X} := \mathbb{E}[X_T^* \,|\, \mathcal{F}_T^B]$, where $(X_t^*)_{t\in[0,T]}$ is the wealth process in* (19) *corresponding to the strategy π^*. We say that π^* is a mean field equilibrium (MFE) if π^* is optimal for the optimization problem* (20) *corresponding to this choice of $\overline{X}$.*

A constant MFE is an $\mathcal{F}_0^{\mathrm{MF}}$-measurable random variable π^* such that, if $\pi_t := \pi^*$ for all $t \in [0,T]$, then $(\pi_t)_{t\in[0,T]}$ is a MFE.

Typically, a MFE is computed as a fixed point. One starts with a generic $\mathcal{F}_T^B$-measurable random variable $\overline{X}$, solves (20) for an optimal π^*, and then computes $\mathbb{E}[X_T^* \,|\, \mathcal{F}_T^B]$. If we have a fixed point in the sense that the *consistency condition*, $\mathbb{E}[X_T^* \,|\, \mathcal{F}_T^B] = \overline{X}$, holds, then π^* is a MFE. Intuitively, every agent in the continuum faces an independent noise W, an independent type vector ζ, and the same common noise B. Therefore, conditionally on B, all agents face i.i.d. copies of the same optimization problem. Heuristically, the law of large numbers suggests that the average terminal wealth of the whole population should be $\mathbb{E}[X_T^* \,|\, \mathcal{F}_T^B]$. This consistency condition illustrates the distinct roles played by the two Brownian motions W and B faced by the representative agent.

Perhaps more intuitively clear is the case where $\sigma = 0$ a.s., so there is no common noise term. In this case, the consistency condition could be replaced with $\overline{X} = \mathbb{E}[X_T^*]$, owing to the fact that each agent in the continuum faces an i.i.d. copy of the same optimization problem.

We refer the reader to [11] for a detailed discussion of mean field games with common noise, and to [32, 8] for general results on limits of n-player games. Alternatively, the so-called "exact law of large numbers" provides another way to formalize this idea of averaging over a continuum of (conditionally) independent agents [44].

2.2.2. *An alternative formulation of the mean field game.* It is worth emphasizing that the optimization problem (20) treats the type vector ζ as a genuine source of randomness, in addition to the stochasticity coming from the Brownian motions. However, an alternative interpretation is given below which will also help in solving the MFG.

As our starting point, for a fixed $\mathcal{F}_T^B$-measurable random variable $\overline{X}$, a dynamic programming argument shows that

$$\sup_{\pi\in\mathcal{A}_{\mathrm{MF}}} \mathbb{E}\left[-e^{-\frac{1}{\delta}\left(X_T-\theta\overline{X}\right)}\right] = \mathbb{E}\left[u(\zeta)\right], \tag{21}$$

where $u(\cdot)$ is a value function defined for (deterministic) elements $\zeta_0 = (x_0, \delta_0, \theta_0, \mu_0, \nu_0, \sigma_0)$ of the type space $\mathcal{Z}^e$ by

$$u(\zeta_0) := \sup_{\pi} \mathbb{E}\left[-\exp\left(-\frac{1}{\delta_0}\left(\widetilde{X}_T^{\zeta_0,\pi} - \theta_0\overline{X}\right)\right)\right], \tag{22}$$

with

$$d\widetilde{X}_t^{\zeta_0,\pi} = \pi_t\left(\mu_0 dt + \nu_0 dW_t + \sigma_0 dB_t\right), \ \widetilde{X}_0^{\zeta_0,\pi} = x_0,$$

and where the supremum is over square-integrable processes which are progressively measurable with respect to the filtration generated by the Brownian motions W and B (noting that the random variable ζ is absent from this filtration). For a deterministic type vector $\zeta_0 \in \mathcal{Z}^e$, the quantity $u(\zeta_0)$ can then be interpreted as the value of the optimization problem (22) *faced by an agent of type* ζ_0. On the other hand, the original optimization problem on the left-hand side of (21) gives the optimal expected value faced by an agent *before the random assignment of types* at time 0.

This new interpretation will be used somewhat implicitly to compute a MFE in the proof of Theorem 9 below, in the following manner. We may write $u(\zeta_0) = v_{\zeta_0}(x_0, 0)$ as the time-zero value of the solution of a HJB equation, $v_{\zeta_0}(x,t)$, for $(x,t) \in \mathbb{R} \times [0,T]$.[4] The optimal value on the left-hand side of (21) is then the expectation of these time-zero values, $\mathbb{E}[v_\zeta(\xi, 0)]$, when the *random* type vector ζ is used.

Similarly, the optimal strategy $\pi^{\zeta_0,*}$ in (22) depends on the fixed value of ζ_0. The optimal strategy for the left-hand side of (21) is then obtained by plugging in the random type vector, yielding $\pi^{\zeta,*} \in \mathcal{A}_{\mathrm{MF}}$. This justifies the interpretation of the strategy $\pi^{\zeta_0,*}$ as the strategy chosen by an agent with type vector ζ_0.

2.2.3. *Solving the mean field game.* Next, we present the second main result, in which we construct a constant MFE and also provide conditions for its uniqueness. The result also confirms that the MFG formulation is indeed appropriate, as the MFE we obtain agrees with the limit of the n-player equilibrium strategies in the sense of (18).

Theorem 9. *Assume that, a.s., $\delta > 0$, $\theta \in [0,1]$, $\mu > 0$, $\sigma \geq 0$, $\nu \geq 0$, and $\sigma + \nu > 0$. Define the constants*

$$\varphi := \mathbb{E}\left[\delta \frac{\mu\sigma}{\sigma^2 + \nu^2}\right] \qquad \text{and} \qquad \psi := \mathbb{E}\left[\theta \frac{\sigma^2}{\sigma^2 + \nu^2}\right],$$

where we assume that both expectations exist and are finite.

There are two cases:

(1) If $\psi < 1$, there exists a unique constant MFE, given by

$$\pi^* = \delta \frac{\mu}{\sigma^2 + \nu^2} + \theta \frac{\sigma}{\sigma^2 + \nu^2} \frac{\varphi}{1 - \psi}. \tag{23}$$

Moreover, we have the identity

$$\mathbb{E}[\sigma \pi^*] = \frac{\varphi}{1 - \psi}.$$

(2) If $\psi = 1$, there is no constant MFE.

Next, we highlight the single stock case, noting that the form of the solution is essentially the same as in the n-player game, presented in Corollary 4.

Corollary 10 (Single stock). *Suppose (μ, ν, σ) are deterministic, with $\nu = 0$ and $\mu, \sigma > 0$. Define the constants*

$$\overline{\delta} := \mathbb{E}[\delta] \qquad \text{and} \qquad \overline{\theta} := \mathbb{E}[\theta].$$

There are two cases:

(1) If $\overline{\theta} < 1$, there exists a unique constant MFE, given by

$$\pi^* = \left(\delta + \theta \frac{\overline{\delta}}{1 - \overline{\theta}}\right) \frac{\mu}{\sigma^2}.$$

(2) If $\overline{\theta} = 1$, there is no constant MFE.

Proof of Theorem 9. The first step in constructing a constant MFE is to solve the stochastic optimization problem in (20), for a given choice of $\overline{X}$. First, observe that it suffices to restrict our attention to random variables $\overline{X}$ of the form $\overline{X} = \mathbb{E}[X_T^\alpha | \mathcal{F}_T^B]$, where X^α solves (19) for some admissible strategy $\alpha \in \mathcal{A}_{\mathrm{MF}}$. However, because we are searching only for constant MFE, we may fix a constant strategy, i.e., an $\mathcal{F}_0^{\mathrm{MF}}$-measurable random variable α with $\mathbb{E}[\alpha^2] < \infty$.

[4]There is some redundancy in this notation, as x_0 is already part of the vector ζ_0.

It is convenient to define, for $t \in [0,T]$,

$$\overline{X}_t := \mathbb{E}[X_t^\alpha | \mathcal{F}_T^B].$$

Noting that $\overline{X}_T = \overline{X}$, the key idea is then to identify the dynamics of the process $(\overline{X}_t)_{t\in[0,T]}$ and incorporate it into the state process of the control problem (20). Because $(\xi,\mu,\sigma,\nu,\alpha)$, W, and B are independent, we must have

$$\overline{X}_t = \overline{\xi} + \overline{\mu\alpha} t + \overline{\sigma\alpha} B_t,$$

where we use the notation $\overline{M} = \mathbb{E}[M]$ for an integrable random variable M.

In turn, for $\pi \in \mathcal{A}_{\mathrm{MF}}$, we define, for $t \in [0,T]$,

$$Z_t^\pi := X_t^\pi - \theta \overline{X}_t,$$

with $(X_t^\pi)_{t\in[0,T]}$ solving (19). Then,

$$dZ_t^\pi = (\mu\pi_t - \theta\overline{\mu\alpha})dt + \nu\pi_t dW_t + (\sigma\pi_t - \theta\overline{\sigma\alpha})dB_t,$$

with $Z_0^\pi = \xi - \theta\overline{\xi}$.

We have thus absorbed $\overline{X}$ as part of the controlled state process. As a result, instead of solving the original control problem (20) we can equivalently solve the Merton-type problem,

$$\sup_{\pi\in\mathcal{A}_{\mathrm{MF}}} \mathbb{E}\left[-\exp\left(-\frac{1}{\delta} Z_T^\pi\right)\right]. \tag{24}$$

As in the discussion in Section 2.2.2, the above supremum equals $\mathbb{E}[v(\xi - \theta\overline{\xi}, 0)]$, where $v(x,t)$ is the unique (smooth, strictly concave and strictly increasing in x) solution of the HJB equation

$$v_t + \max_\pi \left(\frac{1}{2}\left(\nu^2\pi^2 + (\sigma\pi - \theta\overline{\sigma\alpha})^2\right) v_{xx} + (\mu\pi - \theta\overline{\mu\alpha})v_x\right) = 0, \tag{25}$$

with terminal condition $v(x,T) = -e^{-x/\delta}$. We stress that this HJB equation is random, in the sense that it depends on the $\mathcal{F}_0^{\mathrm{MF}}$-measurable type parameters $(\delta,\theta,\mu,\nu,\sigma)$.

Equation (25) simplifies to

$$v_t - \frac{1}{2}\frac{\left(\mu v_x - \theta\sigma\overline{\sigma\alpha} v_{xx}\right)^2}{(\sigma^2+\nu^2)v_{xx}} - \theta\overline{\mu\alpha} v_x = 0.$$

Making the ansatz $v(x,t) = -e^{-x/\delta} f(t)$, the above reduces to

$$f'(t) - \rho f(t) = 0,$$

with $f(T) = 1$, and with ρ given by the $\mathcal{F}_0^{\mathrm{MF}}$-measurable random variable

$$\rho := -\frac{\theta}{\delta}\overline{\mu\alpha} + \frac{\left(\mu + \frac{\theta}{\delta}\sigma\overline{\sigma\alpha}\right)^2}{2(\sigma^2+\nu^2)}. \tag{26}$$

Thus, $f(t) = e^{-\rho(T-t)}$, and $v(x,t) = -e^{-x/\delta} f(t)$. Furthermore, the optimal feedback control achieving the maximum in (25) is given by

$$\pi^*(x,t) = -\frac{\mu v_x(x,t) - \theta\sigma\overline{\sigma\alpha} v_{xx}(x,t)}{(\sigma^2+\nu^2)v_{xx}(x,t)} = \delta\frac{\mu}{\sigma^2+\nu^2} + \theta\frac{\sigma}{\sigma^2+\nu^2}\overline{\sigma\alpha}, \tag{27}$$

In fact, $\pi^* = \pi^*(x,t)$ is $\mathcal{F}_0^{\mathrm{MF}}$-measurable and does not depend on (x,t).

Next, we construct the MFE. To this end, observe that the original constant strategy α is a MFE if and only if

$$\mathbb{E}[X_T^\alpha | \mathcal{F}_T^B] = \mathbb{E}[X_T^{\pi^*} | \mathcal{F}_T^B], \ a.s.,$$

or, equivalently,

$$\overline{\xi} + \overline{\mu\alpha} T + \overline{\sigma\alpha} B_T = \overline{\xi} + \overline{\mu\pi^*} T + \overline{\sigma\pi^*} B_T, \ a.s.$$

Taking expectations, we conclude that α is a constant MFE if and only if

$$\overline{\mu\alpha} = \overline{\mu\pi^*} \qquad \text{and} \qquad \overline{\sigma\alpha} = \overline{\sigma\pi^*}.$$

However, we have from (27) that $\overline{\sigma\alpha} = \overline{\sigma\pi^*}$ if and only if

$$\overline{\sigma\alpha} = \mathbb{E}\left[\delta\frac{\mu\sigma}{\sigma^2+\nu^2}\right] + \mathbb{E}\left[\theta\frac{\sigma^2}{\sigma^2+\nu^2}\right]\overline{\sigma\alpha} = \varphi + \psi\overline{\sigma\alpha}. \tag{28}$$

We then have the following cases:

(i) If $\psi < 1$, the above yields $\overline{\sigma\alpha} = \varphi/(1-\psi)$, and using equation (27) we prove part (1).

(ii) If $\psi = 1$ but $\varphi \neq 0$, then equation (28) has no solutions, and there can be no constant MFE.

(iii) The remaining case is $\psi = 1$ and $\varphi = 0$. However, this cannot happen. Indeed, if this were the case, then $\varphi = 0$ and the restrictions on the parameters imply $\sigma = 0 < \nu$ a.s., which would yield $\psi = 0$, which is a contradiction. This completes the proof of part (2).

□

Remark 11. Note that the proof above yields a tractable formula for the equilibrium value function of the representative agent. Since the controlled process $(Z_t^\pi)_{t\in[0,T]}$ starts from $Z_0^\pi = \xi - \theta\overline{\xi}$, the time-zero value to the representative agent (also called $u(\zeta)$ in Section 2.2.2) is given by

$$v(\xi - \theta\overline{\xi}, 0) = -\exp\left(-\frac{1}{\delta}(\xi - \theta\overline{\xi}) - \rho T\right). \tag{29}$$

It is now straightforward to explicitly compute ρ in (26), using the values of $\overline{\mu\alpha}$ and $\overline{\sigma\alpha}$. Indeed,

$$\rho = \frac{1}{2(\sigma^2+\nu^2)}\left(\mu + \frac{\theta}{\delta}\frac{\varphi}{1-\psi}\sigma\right)^2 - \frac{\theta}{\delta}\left(\widetilde{\psi} + \frac{\widetilde{\varphi}\varphi}{1-\psi}\right), \tag{30}$$

where the constants $\widetilde{\varphi}$ and $\widetilde{\psi}$ are defined by

$$\widetilde{\psi} = \mathbb{E}\left[\delta\frac{\mu^2}{\sigma^2+\nu^2}\right] \qquad \text{and} \qquad \widetilde{\varphi} = \mathbb{E}\left[\theta\frac{\mu\sigma}{\sigma^2+\nu^2}\right].$$

Notably, in the single stock case, this simplifies further to

$$\rho = \left(1 + \left(\frac{\overline{\delta}\theta}{\delta(1-\overline{\theta})}\right)^2\right)\frac{\mu^2}{2\sigma^2}.$$

In fact, the equation (29) essentially provides the solution of the so-called *master equation*; see [9] or [3] for an introduction to the master equation in MFG theory.

2.3. **Discussion of the equilibrium.** We focus the discussion on the mean field equilibria of Theorem 9 and Corollary 10, as the n-player equilibria of Theorem 3 and Corollary 4 have essentially the same structure. The only difference is the rescaling of ν_k^2 by $(1-\theta_k/n)$ wherever it appears in Theorem 3.

Recall first that the MFE π^* is $\mathcal{F}_0^{\mathrm{MF}}$-measurable or, equivalently, ζ-measurable where $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ is the type vector. The randomness of ζ captures the distribution of type vectors among the population, while a single realization of ζ can be interpreted as the type vector of a single representative agent. Hence, we interpret the investment strategy π^* as the equilibrium strategy *adopted by those agents with type vector* ζ.

The equilibrium portfolio π^* consists of two components. The first, $\delta\mu/(\sigma^2+\nu^2)$, is the classical Merton-type portfolio in the absence of relative performance concerns. The second component is always nonnegative, vanishing only in the absence of competition, i.e., when $\theta = 0$.

It increases linearly with the competition weight θ, and we find that competition always increases the allocation in the risky asset.

The representative agent's strategy π^* is influenced by the other agents only through the quantity $\varphi/(1-\psi) = \mathbb{E}[\sigma\pi^*]$. This can be seen as the volatility of aggregate wealth. Indeed, let X^* denote the wealth process corresponding to π^* (i.e., the solution of (19)). The average wealth of the population at time $t \in [0,T]$ is $Y_t := \mathbb{E}[X_t^*|\mathcal{F}_T^B]$. A straightforward computation using the independence of ζ, W, and B yields

$$Y_t = \mathbb{E}[\xi] + \mathbb{E}[\mu\pi^*]t + \mathbb{E}[\sigma\pi^*]B_t.$$

Alternatively, we may interpret the ratio $\varphi/(1-\psi)$ in terms of the type distribution. Define $R = \sigma^2/(\sigma^2+\nu^2)$, which is the fraction of the representative agent's stock's variance driven by the common noise B. Then $\varphi = \mathbb{E}[R\delta\mu/\sigma]$ is computed by multiplying each agent's Sharpe ratio by her risk tolerance parameter and the weight R, then averaging over all agents. Similarly, $\psi = \mathbb{E}[R\theta]$ is the average competition parameter, weighted by R. Several important factors will lead to an increase in $\varphi/(1-\psi)$, and thus an increase in the investment π^* in the risky asset. Namely, π^* increases as other agents become more risk tolerant (higher δ on average), as other agents become more competitive (higher θ on average), or as we increase the quality of the other stocks as measured by their Sharpe ratio (higher $\delta\mu/\sigma$ on average).

Some of the effects of competition are clearer in the single stock case of Corollary 10. The resulting MFE π^* clearly resembles the Merton-type portfolio but with *effective risk tolerance parameter*

$$\delta_{\text{eff}} := \delta + \theta\frac{\overline{\delta}}{1-\overline{\theta}}.$$

We always have $\delta_{\text{eff}} > \delta$ if $\theta > 0$, and the difference $\delta_{\text{eff}} - \delta$ increases with θ, with $\overline{\delta}$, and with $\overline{\theta}$. That is, the representative agent invests more in the risky asset if she is more competitive, if other agents tend to be more risk tolerant, or if other agents tend to be more competitive. In the latter cases, when $\overline{\delta}$ and $\overline{\theta}$ increase, we can interpret the increase in π^* as an effort, on the part of the representative agent, to "keep up" with a population more willing to take risk.

A few other special cases are worth discussing. If $\sigma = 0$ a.s., there is no common noise. In this case, $\varphi = \psi = 0$, and in turn the MFE is equal to the Merton portfolio. All agents act independently uncompetitively, not taking into account the performance of their competitors.

On the other hand, if $\nu = 0$ a.s., there is no independent noise, and we have the simplifications $\psi = \mathbb{E}[\theta]$ and $\varphi = \mathbb{E}[\delta\mu/\sigma]$. If $\mathbb{E}[\theta] < 1$, then we have

$$\pi^* = \delta\frac{\mu}{\sigma^2} + \frac{\theta}{\sigma\left(1-\mathbb{E}\left[\theta\right]\right)}\mathbb{E}\left[\delta\frac{\mu}{\sigma}\right].$$

If $\nu = 0$ a.s. and also $\mathbb{E}[\theta] = 1$, then $\theta = 1$ a.s. and $\psi = 1$. In this case, every agent is concerned exclusively with relative and not absolute performance, and there is no equilibrium.

Another degenerate case is when all agents have the same type vector, i.e., when ζ is deterministic. Then, the MFE is common for all agents and (assuming $\theta < 1$) reduces to

$$\pi^* = \frac{\delta\mu}{\left(1-\theta\right)\sigma^2+\nu^2}.$$

3. CRRA risk preferences

In this section we focus on power and logarithmic (CARA) utilities. Given the homogeneity properties of the power risk preferences, we choose to measure relative performance using a multiplicative and not additive factor. Such cases were analyzed for a two-agent setting in [2] and more recently in [23] under forward relative performance criteria.

3.1. **The n-agent game.** We consider an n-player game analogous to that of Section 2.1, but where each agent has a CRRA utility. We work on the same filtered probability space of Section 2.1, and we assume that the n stocks have the same dynamics as in (1).

The n agents trade in a common investment horizon. As is common in power utility models, the strategy π_t^i is taken to be the *fraction* (as opposed to the amount) of wealth that agent i invests in the stock S^i at time t. Her discounted wealth is then given by

$$dX_t^i = \pi_t^i X_t^i \left(\mu_i dt + \nu_i dW_t + \sigma_i dB_t\right), \tag{31}$$

with initial endowment $X_0^i = x_0^i$. The class of admissible strategies is as before the set $\mathcal{A}$ of self-financing $\mathbb{F}$-progressively measurable processes $(\pi_t)_{t\in[0,T]}$ satisfying $\mathbb{E}\int_0^T |\pi_t|^2 dt < \infty$.

The i^{th} agent's utility is a function $U_i : \mathbb{R}_+^2 \to \mathbb{R}$ of both her individual wealth, x, and the geometric average wealth of all agents, m. Specifically,

$$U_i(x, m) := U(xm^{-\theta_i}; \delta_i),$$

where $U(x;\delta)$ is defined for $x > 0$ and $\delta > 0$ by

$$U(x;\delta) := \begin{cases} \left(1 - \frac{1}{\delta}\right)^{-1} x^{1-\frac{1}{\delta}}, & \text{for } \delta \neq 1, \\ \log x & \text{for } \delta = 1. \end{cases}$$

The constant parameters $\delta_i > 0$ and $\theta_i \in [0,1]$ are the personal relative risk tolerance and competition weight parameters, respectively. [5] If agents $i = 1, \ldots, n$ choose admissible strategies $\pi^1, \ldots, \pi^n$, the payoff for agent i is given by

$$J_i(\pi^1, \ldots, \pi^n) := \mathbb{E}\left[U\left(X_T^i \overline{X}_T^{-\theta_i}; \delta_i\right)\right], \quad \text{where} \quad \overline{X}_T = \left(\prod_{k=1}^n X_T^k\right)^{1/n}. \tag{32}$$

Notice that here, unlike in the exponential utility model, agents measure relative wealth using the *geometric mean*, rather than the arithmetic mean. Working with the geometric mean instead of the arithmetic means renders the problem tractable, as it allows us to exploit the homogeneity of the utility function.

The above expected utility may be rewritten more illustratively as

$$J_i(\pi^1, \ldots, \pi^n) = \mathbb{E}\left[U\left((X_T^i)^{1-\theta_i}(R_T^i)^{\theta_i}; \delta_i\right)\right],$$

where $R_T^i = X_T^i / \overline{X}_T$ is the *relative return* for agent i. This clarifies the role of the competition weight θ_i as governing the trade-off between absolute and relative wealth to agent i, as in the exponential utility model. As before, an agent with a higher value of θ_i is more concerned with relative wealth than with absolute wealth.

The notion of (Nash) equilibrium is defined exactly as in Definition 1, but with the new objective function defined in (32) above. We find a unique *constant* equilibrium in the following theorem, which we subsequently specialize to the single stock case.

Theorem 12. *Assume that for all $i = 1, \ldots, n$ we have $x_0^i > 0$, $\delta_i > 0$, $\theta_i \in [0,1]$, $\mu_i > 0$, $\sigma_i \geq 0$, $\nu_i \geq 0$, and $\sigma_i + \nu_i > 0$. Define the constants*

$$\varphi_n := \frac{1}{n}\sum_{k=1}^n \delta_k \frac{\mu_k \sigma_k}{\sigma_k^2 + \nu_k^2(1 + (\delta_k - 1)\theta_k/n)} \tag{33}$$

[5]For CRRA utilities, it is more common to use the *relative risk aversion* parameter $\gamma_i = 1/\delta_i$, but our choice of parametrization ensures that the *relative risk tolerance* is precisely

$$\delta_i = -\frac{U_x(x;\delta_i)}{xU_{xx}(x;\delta_i)}.$$

and

$$\psi_n := \frac{1}{n}\sum_{k=1}^n \theta_k(\delta_k - 1)\frac{\sigma_k^2}{\sigma_k^2 + \nu_k^2(1+(\delta_k-1)\theta_k/n)}. \tag{34}$$

There exists a unique constant equilibrium, given by

$$\pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2(1+(\delta_i-1)\theta_i/n)} - \theta_i(\delta_i - 1)\frac{\sigma_i}{\sigma_i^2 + \nu_i^2(1+(\delta_i-1)\theta_i/n)}\frac{\varphi_n}{1+\psi_n}. \tag{35}$$

Moreover, we have the identity

$$\frac{1}{n}\sum_{k=1}^n \sigma_k \pi^{k,*} = \frac{\varphi_n}{1+\psi_n}.$$

Corollary 13 (Single stock)**.** *Assume that for all $i = 1, \dots, n$ we have $\mu_i = \mu > 0$, $\sigma_i = \sigma > 0$, and $\nu_i = 0$, with μ and σ independent of i. Define the constants*

$$\overline{\delta} := \frac{1}{n}\sum_{k=1}^n \delta_k \qquad \text{and} \qquad \overline{\theta(\delta-1)} := \frac{1}{n}\sum_{k=1}^n \theta_k(\delta_k - 1).$$

There exists unique constant equilibrium, given by

$$\pi^{i,*} = \left(\delta_i - \frac{\theta_i(\delta_i - 1)\overline{\delta}}{1+\overline{\theta(\delta-1)}}\right)\frac{\mu}{\sigma^2}.$$

Proof. Apply Theorem 12, taking note of the simplifications $\varphi_n = \overline{\delta}\mu/\sigma$ and $\psi_n = \overline{\theta(\delta-1)}$. □

Remark 14. As in Remark 5 in the exponential utility model, one might modify our payoff structure so that agent i excludes herself from the geometric mean $\overline{X}_T$. That is, one might replace the payoff functional J_i defined in (32) by

$$\mathbb{E}\left[U\left(X_T^i(\overline{X}_T^{(-i)})^{-\theta_i'}; \delta_i'\right)\right], \quad \text{where} \quad \overline{X}_T^{(-i)} = \left(\prod_{k\neq i} X_T^k\right)^{1/(n-1)},$$

for some parameters $\theta_i' \in [0,1]$ and $\delta_i' > 0$. By modifying the preference parameters, we may view this payoff as a special case of ours. Indeed, by matching coefficients it is straightforward to show that

$$U\left(X_T^i(\overline{X}_T^{(-i)})^{-\theta_i'}; \delta_i'\right) = c_i U\left(X_T^i \overline{X}_T^{-\theta_i}; \delta_i\right),$$

for some constant $c_i > 0$ (which does not influence the optimal strategies), when $\theta_i \in [0,1]$ and $\delta_i > 0$ are defined by

$$\delta_i = \frac{\delta_i'}{\delta_i' - (\delta_i' - 1)\left(1 + \frac{1}{n-1}\theta_i'\right)} \qquad \text{and} \qquad \theta_i = \frac{\theta_i'}{\frac{n-1}{n} + \frac{1}{n}\theta_i'}.$$

However, this is only valid if $(1 - 1/\delta_i')\left(1 + \frac{1}{n-1}\theta_i'\right) < 1$, which ensures that $\delta_i > 0$. This certainly holds for sufficiently large n. We favor our original parametrization because of the relative simplicity of the formulas in Theorem 12 and Corollary 13, and because there is no difference in the $n \to \infty$ limit.

Proof of Theorem 12. The proof is similar to that of Theorem 3, so we only highlight the main steps. Fix an agent i and constant strategies $\alpha_k \in \mathbb{R}$, for $k \neq i$. Define

$$Y_t := \left(\prod_{k\neq i} X_t^k\right)^{1/n},$$

where X_t^k solves (31) with constant weights α_k and $X_0^k = x_0^k$.

Setting $\Sigma_k := \sigma_k^2 + \nu_k^2$, we deduce that

$$d\left(\log X_t^k\right) = \left(\mu_k\alpha_k - \frac{1}{2}\Sigma_k\alpha_k^2\right)dt + \nu_k\alpha_k dW_t^k + \sigma_k\alpha_k dB_t.$$

In turn,

$$d\left(\log Y_t\right) = \frac{1}{n}\sum_{k\neq i} d\log X_t^k = \left(\widehat{\mu\alpha} - \frac{1}{2}\widehat{\Sigma\alpha^2}\right)dt + \frac{1}{n}\sum_{k\neq i}\nu_k\alpha_k dW_t^k + \widehat{\sigma\alpha}dB_t,$$

where we abbreviate

$$\widehat{\mu\alpha} := \frac{1}{n}\sum_{k\neq i}\mu_k\alpha_k, \qquad \widehat{\sigma\alpha} := \frac{1}{n}\sum_{k\neq i}\sigma_k\alpha_k,$$

$$\widehat{\Sigma\alpha^2} := \frac{1}{n}\sum_{k\neq i}\Sigma\alpha_k^2 \qquad \text{and} \qquad \widehat{(\nu\alpha)^2} := \frac{1}{n}\sum_{k\neq i}\nu_k^2\alpha_k^2.$$

Thus, the process Y_t solves

$$\frac{dY_t}{Y_t} = \eta dt + \frac{1}{n}\sum_{k\neq i}\nu_k\alpha_k dW_t^k + \widehat{\sigma\alpha}dB_t, \quad Y_0 = \left(\prod_{k\neq i} x_0^k\right)^{1/n}, \tag{36}$$

with

$$\eta := \widehat{\mu\alpha} - \frac{1}{2}\left(\widehat{\Sigma\alpha^2} - \widehat{\sigma\alpha}^2 - \frac{1}{n}\widehat{(\nu\alpha)^2}\right).$$

The i^{th} agent then solves the optimization problem

$$\sup_{\pi\in\mathcal{A}} \mathbb{E}\left[U\left((X_T^i)^{1-\theta_i/n}Y_T^{-\theta_i};\delta_i\right)\right], \tag{37}$$

where

$$dX_t^i = X_t^i\pi_t(\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t),\ X_0^i = x_0^i,$$

and where $(Y_t)_{t\in[0,T]}$ solves (36). We then obtain that the value (37) is equal to $v(X_0^i, Y_0, 0)$, where $v(x,y,t)$ solves the HJB equation

$$v_t + \max_{\pi\in\mathbb{R}}\left(\frac{1}{2}(\sigma_i^2+\nu_i^2)\pi^2x^2v_{xx} + \pi\left(\mu_i x v_x + \sigma_i\widehat{\sigma\alpha}xyv_{xy}\right)\right) \tag{38}$$

$$+\frac{1}{2}\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right)y^2v_{yy} + \eta y v_y = 0,$$

for $(x,y,t)\in\mathbb{R}_+\times\mathbb{R}_+\times[0,T]$, with terminal condition

$$v(x,y,T) = U(x^{1-\theta_i/n}y^{-\theta_i};\delta_i).$$

Applying the first order conditions, the maximum in (38) is attained by

$$\pi^{i,*}(x,y,t) = -\frac{\mu_i x v_x(x,y,t) + \sigma_i\widehat{\sigma\alpha}xyv_{xy}(x,y,t)}{(\sigma_i^2+\nu_i^2)x^2v_{xx}(x,y,t)}. \tag{39}$$

In turn, equation (38) reduces to

$$v_t - \frac{1}{2}\frac{(\mu_i x v_x + \sigma_i\widehat{\sigma\alpha}xyv_{xy})^2}{(\sigma_i^2+\nu_i^2)x^2v_{xx}} + \frac{1}{2}\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right)y^2v_{yy} + \eta y v_y = 0. \tag{40}$$

Next, we claim that the above HJB equation has a unique smooth solution (in an appropriate class), and the optimal feedback control in (39) reduces to

$$\pi^{i,*} = \frac{\delta_i \mu_i - \sigma_i \widehat{\sigma\alpha}\theta_i(\delta_i - 1)}{(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))}. \tag{41}$$

We prove this in two cases:

(i) Suppose $\delta_i \neq 1$. Making the ansatz

$$v(x, y, t) = U(x^{1-\theta_i/n} y^{-\theta_i}; \delta_i) f(t) = (1 - 1/\delta_i)^{-1} (x^{(1-\theta_i/n)} y^{-\theta_i})^{1-1/\delta_i} f(t)$$

reduces equation (40) to $(1 - 1/\delta_i)^{-1} f'(t) + \rho f(t) = 0$, with $f(T) = 1$, where

$$\rho := \frac{(\mu_i(1 - \theta_i/n) - \sigma_i \widehat{\sigma\alpha}\theta_i(1 - \theta_i/n)(1 - 1/\delta_i))^2}{2(\sigma_i^2 + \nu_i^2)(1 - \theta_i/n)(1 - (1 - \theta_i/n)(1 - 1/\delta_i))} - \eta\theta_i$$
$$+ \frac{1}{2}\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right)\theta_i(1 + \theta_i(1 - 1/\delta_i)).$$

We easily deduce that the solution of (40) is

$$v(x, y, t) = (1 - 1/\delta_i)^{-1} (x^{(1-\theta_i/n)} y^{-\theta_i})^{1-1/\delta_i} e^{\rho(1-1/\delta_i)(T-t)},$$

and that (39) yields (41).

(ii) Suppose $\delta_i = 1$. Making the ansatz

$$v(x, y, t) = U(x^{1-\theta_i/n} y^{-\theta_i}; \delta_i) + f(t) = \left(1 - \frac{\theta_i}{n}\right)\log x - \theta_i \log y + f(t)$$

reduces equation (40) to $f'(t) + \rho f(t) = 0$, with $f(T) = 1$ and

$$\rho := \frac{\mu_i^2(1 - \theta_i/n)}{2(\sigma_i^2 + \nu_i^2)} - \theta_i\eta + \frac{1}{2}\theta_i\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(\nu\alpha)^2}\right).$$

In turn, the solution of (40) is given by

$$v(x, y, t) = \left(1 - \frac{\theta_i}{n}\right)\log x - \theta_i \log y + \rho(T - t),$$

and (39) reduces to $\pi^{i,*} = \mu_i/(\sigma_i^2 + \nu_i^2)$, which is consistent with (41) for $\delta_i = 1$.

With (41) now established, we conclude the proof. For $(\alpha_1, \ldots, \alpha_n)$ to be a constant equilibrium, we must have $\pi^{i,*} = \alpha_i$, for each $i = 1, \ldots, n$. Using (41) and abbreviating

$$\overline{\sigma\alpha} := \frac{1}{n}\sum_{k=1}^{n} \sigma_k\alpha_k = \widehat{\sigma\alpha} + \frac{1}{n}\sigma_i\alpha_i,$$

we deduce that we must have

$$\alpha_i = \frac{\mu_i - \sigma_i\overline{\sigma\alpha}\theta_i(1 - 1/\delta_i) + \sigma_i^2\alpha_i(\theta_i/n)(1 - 1/\delta_i)}{(\sigma_i^2 + \nu_i^2)(1 - (1 - \theta_i/n)(1 - 1/\delta_i))}.$$

Solving for α_i yields,

$$\begin{aligned}
\alpha_i &= \frac{\mu_i - \sigma_i\overline{\sigma\alpha}\theta_i(1 - 1/\delta_i)}{(\sigma_i^2 + \nu_i^2)(1 - (1 - \theta_i/n)(1 - 1/\delta_i))}\left(1 - \frac{\sigma_i^2(\theta_i/n)(1 - 1/\delta_i)}{(\sigma_i^2 + \nu_i^2)(1 - (1 - \theta_i/n)(1 - 1/\delta_i))}\right)^{-1} \\
&= \frac{\mu_i - \sigma_i\overline{\sigma\alpha}\theta_i(1 - 1/\delta_i)}{(\sigma_i^2 + \nu_i^2)(1 - (1 - \theta_i/n)(1 - 1/\delta_i)) - \sigma_i^2(\theta_i/n)(1 - 1/\delta_i)} \\
&= \frac{\mu_i - \sigma_i\overline{\sigma\alpha}\theta_i(1 - 1/\delta_i)}{\sigma_i^2/\delta_i + \nu_i^2(1 - (1 - \theta_i/n)(1 - 1/\delta_i))} = \frac{\mu_i\delta_i - \sigma_i\overline{\sigma\alpha}\theta_i(\delta_i - 1)}{\sigma_i^2 + \nu_i^2(1 - \theta_i/n + \delta_i\theta_i/n)}.
\end{aligned} \tag{42}$$

Multiplying both sides by σ_i and averaging over $i = 1, \dots, n$ give

$$\overline{\sigma\alpha} = \varphi_n - \psi_n \overline{\sigma\alpha}, \tag{43}$$

where φ_n, ψ_n are as in (33) and (34). Because $1 + \psi_n > 0$, equation (43) holds if and only if $\overline{\sigma\alpha} = \varphi_n/(1 + \psi_n)$. We then deduce from (42) that the equilibrium strategy $\alpha_i = \pi^{i,*}$ is given by (35). □

Remark 15. Note that equation (42) above has a *unique* solution for *all* parameter values. In contrast, the analogous equation (28) in the exponential case has no solutions for certain parameter values, which is why there were two cases in Theorem 3.

3.2. **The mean field game.** This section studies the limit as $n \to \infty$ of the n-player game analyzed in the previous section, analogously to the treatment of the exponential case in Section 2.2.

We proceed with some informal arguments. Recall that the *type vector* of agent i is

$$\zeta_i := (x_0^i, \delta_i, \theta_i, \mu_i, \nu_i, \sigma_i).$$

As before, the type vectors induce an empirical measure, which is the probability measure on

$$\mathcal{Z}^p := (0, \infty) \times (0, \infty) \times [0, 1] \times (0, \infty) \times [0, \infty) \times [0, \infty) \tag{44}$$

given by

$$m_n(A) = \frac{1}{n} \sum_{i=1}^{n} 1_A(\zeta_i), \quad \text{for Borel sets } A \subset \mathcal{Z}^p.$$

Similarly to the exponential case, for a given agent i, the equilibrium strategy $\pi^{i,*}$ computed in Theorem 12 depends only on her own type vector ζ_i and the distribution m_n of type vectors, and this enables the passage to the limit.

Assume now that m_n has a weak limit m, in the sense that $\int_{\mathcal{Z}^p} f\, dm_n \to \int_{\mathcal{Z}^p} f\, dm$ for every bounded continuous function f on $\mathcal{Z}^p$. Let $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ denote a random variable with distribution m. Then, the optimal strategy $\pi^{i,*}$ (c.f. (35)) should converge to

$$\lim_{n\to\infty} \pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2} - \theta_i(\delta_i - 1) \frac{\sigma_i}{\sigma_i^2 + \nu_i^2} \frac{\varphi}{1 - \psi}, \tag{45}$$

where

$$\varphi := \lim_{n\uparrow\infty} \varphi_n = \mathbb{E}\left[\delta \frac{\mu\sigma}{\sigma^2 + \nu^2}\right] \qquad \text{and} \qquad \psi := \lim_{n\uparrow\infty} \psi_n = \mathbb{E}\left[\theta(\delta - 1) \frac{\sigma^2}{\sigma^2 + \nu^2}\right].$$

As in the exponential case, we will demonstrate that this limiting strategy is indeed the equilibrium of a mean field game, which we formulate analogously to Section 2.2.1.

Recall that W and B are independent Brownian motions and that the random variable $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ is independent of W and B. For the power case, the type vector ζ now takes values in the space $\mathcal{Z}^p$. Furthermore, the filtration $\mathbb{F}^{\mathrm{MF}}$ is the smallest one satisfying the usual assumptions for which ζ is $\mathcal{F}_0^{\mathrm{MF}}$-measurable and W and B are adapted. Finally, recall that $\mathbb{F}^B = (\mathcal{F}_t^B)_{t\in[0,T]}$ denotes the natural filtration generated by the Brownian motion B.

The *representative agent's* wealth process solves

$$dX_t = \pi_t X_t(\mu dt + \nu dW_t + \sigma dB_t), \ X_0 = \xi, \tag{46}$$

where the investment weight π belongs to the admissible set $\mathcal{A}_{\mathrm{MF}}$ of $\mathbb{F}^{\mathrm{MF}}$-progressively measurable real-valued processes satisfying $\mathbb{E}\int_0^T |\pi_t|^2 dt < \infty$. Notice that, for all admissible π, the wealth process $(X_t)_{t\in[0,T]}$ is strictly positive, as $\xi > 0$ a.s.

We denote by $\overline{X}$ an $\mathcal{F}_0^{\mathrm{MF}}$-measurable random variable representing the geometric mean wealth among the continuum of agents. Then, the objective of the representative agent is to maximize the expected payoff

$$\sup_{\pi \in \mathcal{A}_{\mathrm{MF}}} \mathbb{E}\left[U(X_T \overline{X}^{-\theta}; \delta)\right], \tag{47}$$

where $(X_t)_{t\in[0,T]}$ is given by (46).

The definition of a mean field equilibrium is analogous to Definition 8. However, one needs to extend the notion of geometric mean appropriately to a continuum of agents. The geometric mean of a measure m on $(0,\infty)$ is most naturally defined as

$$\exp\left(\int_{(0,\infty)} \log y \; dm(y)\right),$$

at least when $\log y$ is m-integrable. Indeed, when m is the empirical measure of n points $(y_1, \ldots, y_n)$, this reduces to the usual definition $(y_1 y_2 \cdots y_n)^{1/n}$.

Definition 16. *Let $\pi^* \in \mathcal{A}_{\mathrm{MF}}$ be an admissible strategy, and consider the $\mathcal{F}_T^B$-measurable random variable $\overline{X} := \exp \mathbb{E}[\log X_T^* \,|\, \mathcal{F}_T^B]$, where $(X_t^*)_{t\in[0,T]}$ is the wealth process in (46) corresponding to the strategy π^*. We say that π^* is a mean field equilibrium (MFE) if π^* is optimal for the optimization problem (47) corresponding to this choice of $\overline{X}$.*

A constant MFE is an $\mathcal{F}_0^{\mathrm{MF}}$-measurable random variable π^* such that, if $\pi_t := \pi^*$ for all $t \in [0,T]$, then $(\pi_t)_{t\in[0,T]}$ is a MFE.

The following theorem characterizes the constant MFE, recovering the limiting expressions derived above from the n-player equilibria.

Theorem 17. *Assume that, a.s., $\delta > 0$, $\theta \in [0,1]$, $\mu > 0$, $\sigma \geq 0$, $\nu \geq 0$, and $\sigma + \nu > 0$. Define the constants*

$$\varphi := \mathbb{E}\left[\delta \frac{\mu\sigma}{\sigma^2 + \nu^2}\right] \qquad \text{and} \qquad \psi := \mathbb{E}\left[\theta(\delta - 1)\frac{\sigma^2}{\sigma^2+\nu^2}\right],$$

where we assume that both expectations exist and are finite.

There exists a unique constant MFE, given by

$$\pi^* = \delta \frac{\mu}{\sigma^2 + \nu^2} - \theta(\delta - 1)\frac{\sigma}{\sigma^2+\nu^2}\frac{\varphi}{1+\psi}. \tag{48}$$

Moreover, we have the identity

$$\mathbb{E}[\sigma\pi^*] = \frac{\varphi}{1+\psi}.$$

In the single stock case, the form of the solution is essentially the same as in the n-player game, presented in Corollary 13:

Corollary 18 (Single stock)**.** *Suppose (μ, ν, σ) are deterministic, with $\nu = 0$ and $\mu, \sigma > 0$. Define the constants*

$$\overline{\delta} := \mathbb{E}[\delta] \qquad \text{and} \qquad \overline{\theta(\delta-1)} := \mathbb{E}[\theta(\delta-1)].$$

There exists a unique constant MFE, given by

$$\pi^* = \left(\delta - \frac{\theta(\delta-1)\overline{\delta}}{1 + \overline{\theta(\delta-1)}}\right)\frac{\mu}{\sigma^2}.$$

Proof of Theorem 17. As in the exponential case, we first reduce the optimal control problem (47) to a low-dimensional Markovian one. To this end, it suffices to restrict our attention to random variables $\overline{X}$ of the form

$$\overline{X} = \exp \mathbb{E}[\log X_T^\alpha | \mathcal{F}_T^B],$$

where X^α is the wealth process of (46) with an admissible constant strategy α. That is, α is an $\mathcal{F}_0^{\mathrm{MF}}$-measurable random variable satisfying $\mathbb{E}[\alpha^2] < \infty$.

Define

$$Y_t := \exp \mathbb{E}[\log X_t^\alpha | \mathcal{F}_T^B].$$

Note that $Y_t = \exp \mathbb{E}[\log X_t^\alpha | \mathcal{F}_t^B]$, for $t \in [0, T)$, because $(B_s - B_t : s \in [t, T])$ and X_t^α are independent. In analogy to the exponential case, we identify the dynamics of Y and, in turn, treat it, as an additional (uncontrolled) state process.

To this end, first use Itô's formula to get

$$d\left(\log X_t^\alpha\right) = \left(\mu\alpha - \frac{1}{2}(\sigma^2 + \nu^2)\alpha^2\right) dt + \nu\alpha dW_t + \sigma\alpha dB_t.$$

Define $\check{X}_t^\alpha := \mathbb{E}[\log X_t^\alpha | \mathcal{F}_T^B]$, and note as with Y_t above that $\check{X}_t^\alpha = \mathbb{E}[\log X_t^\alpha | \mathcal{F}_t^B]$, for $t \in [0, T)$. Setting

$$\Sigma := \sigma^2 + \nu^2,$$

and noting that $(\xi, \mu, \sigma, \nu, \alpha)$, W and B are independent, we compute

$$d\check{X}_t^\alpha = \left(\overline{\mu\alpha} - \frac{1}{2}\overline{\Sigma\alpha^2}\right) dt + \overline{\sigma\alpha} dB_t, \tag{49}$$

where, again, we use the notation $\overline{M} = \mathbb{E}[M]$ for a generic integrable random variable M. In turn,

$$dY_t = de^{\check{X}_t^\alpha} = Y_t\left(\eta dt + \overline{\sigma\alpha} dB_t\right), \quad Y_0 = \overline{\xi}, \tag{50}$$

where $\eta := \overline{\mu\alpha} - \frac{1}{2}(\overline{\Sigma\alpha^2} - \overline{\sigma\alpha}^2)$.

To solve the stochastic optimization problem (47), we equivalently solve

$$\sup_{\pi \in \mathcal{A}_{\mathrm{MF}}} \mathbb{E}\left[U(X_T Y_T^{-\theta}; \delta)\right] \tag{51}$$

with

$$dX_t = \pi_t X_t(\mu dt + \nu dW_t + \sigma dB_t)$$

and with $(Y_t)_{t\in[0,T]}$ solving (50). Then, as in the discussion of Section 2.2.2, the value of (51) is equal to $\mathbb{E}[v(\xi, \overline{\xi}, 0)]$, where $v = v(x, y, t)$ is the unique smooth (strictly concave and strictly increasing in x) solution of the HJB equation

$$v_t + \max_{\pi \in \mathbb{R}}\left(\frac{1}{2}\Sigma\pi^2 x^2 v_{xx} + \pi\left(\mu x v_x + \sigma\overline{\sigma\alpha} x y v_{xy}\right)\right) + \frac{1}{2}\overline{\sigma\alpha}^2 y^2 v_{yy} + \eta y v_y = 0, \tag{52}$$

with terminal condition $v(x, y, T) = U(xy^{-\theta}; \delta)$. Notice that this HJB equation is random, because of its dependence on the $\mathcal{F}_0^{\mathrm{MF}}$-measurable type parameters.

Applying the first order conditions, the maximum in (52) is attained by

$$\pi^*(x, y, t) = -\frac{\mu x v_x(x, y, t) + \sigma\overline{\sigma\alpha} x y v_{xy}(x, y, t)}{\Sigma x^2 v_{xx}(x, y, t)}. \tag{53}$$

In turn, equation (52) reduces to

$$v_t - \frac{(\mu x v_x + \sigma\overline{\sigma\alpha} x y v_{xy})^2}{2\Sigma x^2 v_{xx}} + \frac{1}{2}\overline{\sigma\alpha}^2 y^2 v_{yy} + \eta y v_y = 0. \tag{54}$$

Next, we claim that, for all (x, y, t),

$$\pi^*(x, y, t) = \Sigma^{-1} (\mu\delta - \theta(\delta - 1)\sigma\overline{\sigma\alpha}) . \tag{55}$$

We prove this in two cases:

(i) Suppose $\delta \neq 1$. Making the ansatz

$$v(x, y, t) = U(xy^{-\theta}; \delta)f(t) = (1 - 1/\delta)^{-1} x^{1-1/\delta} y^{-\theta(1-1/\delta)} f(t),$$

reduces equation (54) to $f'(t) + \rho f(t) = 0$, with $f(T) = 1$, where

$$\rho := \frac{(\mu\delta - \theta(\delta - 1)\sigma\overline{\sigma\alpha})^2}{2\Sigma(\delta - 1)} - \eta\theta(1 - 1/\delta)^{-1} + \frac{1}{2}\overline{\sigma\alpha}^2\theta(\theta + (1 - 1/\delta)^{-1}).$$

We easily deduce that the solution of (54) is

$$v(x, y, t) = (1 - 1/\delta)^{-1} x^{1-1/\delta} y^{-\theta(1-1/\delta)} \exp(\rho(T - t)),$$

and that (53) yields (55).

(ii) Suppose $\delta = 1$. It is easily checked that the solution v of (54) is given by

$$v(x, y, t) = \log x - \theta \log y + \rho(T - t),$$

with

$$\rho := \frac{\mu^2}{2\Sigma} - \eta\theta + \frac{1}{2}\theta\overline{\sigma\alpha}^2.$$

In this case, (53) becomes $\pi^*(x, y, t) = \mu/\Sigma$, which is consistent with (55) for $\delta = 1$.

With (55). now established, we conclude the proof. Note that the original constant strategy α is a MFE if and only if

$$\exp \mathbb{E}[\log X_T^{\pi^*} | \mathcal{F}_T^B] = \exp \mathbb{E}[\log X_T^{\alpha} | \mathcal{F}_T^B], \ a.s.$$

Recalling (49), this is equivalent to

$$\overline{\xi} + \left(\overline{\mu\pi^*} - \frac{1}{2}\overline{\Sigma|\pi^*|^2}\right) T + \overline{\sigma\pi^*} B_T = \overline{\xi} + \left(\overline{\mu\alpha} - \frac{1}{2}\overline{\Sigma\alpha^2}\right) T + \overline{\sigma\alpha} B_T, \ a.s.$$

Taking expectations yields that the above is equivalent to

$$\overline{\mu\pi^*} - \frac{1}{2}\overline{\Sigma|\pi^*|^2} = \overline{\mu\alpha} - \frac{1}{2}\overline{\Sigma\alpha^2} \qquad \text{and} \qquad \overline{\sigma\pi^*} = \overline{\sigma\alpha}.$$

Multiplying both sides of (55) by σ and taking expectations, we find that $\overline{\sigma\pi^*} = \overline{\sigma\alpha}$ if and only if

$$\overline{\sigma\alpha} = \mathbb{E}\left[\delta\frac{\mu\sigma}{\Sigma}\right] - \mathbb{E}\left[\theta(\delta - 1)\frac{\sigma^2}{\Sigma}\right]\overline{\sigma\alpha} = \varphi - \psi\overline{\sigma\alpha}.$$

We then deduce that $\overline{\sigma\alpha} = \varphi/(1 + \psi)$, and plugging this into (55) we obtain (48). $\square$

3.3. **Discussion of the equilibrium.** Some of the structure of the equilibrium is similar to what we observed in the CARA model, discussed in Section 2.3. We again focus the discussion here on the mean field case of Theorem 17 and Corollary 18, as the n-player equilibria of Theorem 12 and Corollary 13 have essentially the same structure. The only difference is the rescaling of ν_k^2 by $(1 + (\delta_k - 1)\theta_k/n)$ wherever it appears in Theorem 12.

As in the CARA model, the MFE π^* in the CRRA model is the sum of two components. The first, $\delta\mu/(\sigma^2 + \nu^2)$, is the classical Merton portfolio. The second,

$$\pi^*_{(2)} := -\theta(\delta - 1)\frac{\sigma}{\sigma^2 + \nu^2}\frac{\varphi}{1 + \psi},$$

displays a variety of behaviors. Again, this term is linear in θ and vanishes when $\theta = 0$, i.e., when the representative agent is not competitive.

Interestingly, however, the effect of competition is quite different compared to CARA model, in the sense that now competition leads some agents to invest *less* in the risky asset than they would in the absence of competition. Indeed, the sign of $\pi^*_{(2)}$ is the same as that of $1-\delta$, at least assuming $\theta > 0$ and $\sigma > 0$. Thus agents with $\delta < 1$ invest *more* as θ increases, whereas agents with $\delta > 1$ invest *less*. In particular, we have $\pi^*_{(2)} = 0$ when $\delta = 1$. That is, agents with log utility are not competitive, which can be deduced easily from the original problem formulation.

In fact, the effect of competition is so strong that agents may even go short. That is, π^* may be negative. This typically occurs when δ and θ are much higher than their population averages, or, in other words, when the representative agent is very risk tolerant and competitive relative to other agents. More precise criteria are easiest to see in the single stock case, detailed in equation (56) below.

The representative agent's strategy π^* is influenced by the other agents only through the quantity $\varphi/(1+\psi) = \mathbb{E}[\sigma\pi^*]$, and as in Section 2.3 we can view this quantity as the volatility of the aggregate wealth. Indeed, let X^* denote the wealth process corresponding to π^* (i.e., the solution of (46)). The geometric average wealth of the population at time $t \in [0,T]$ is $Y_t := \log \mathbb{E}[\exp(X^*_t)|\mathcal{F}^B_T]$. We saw in the proof of Theorem 17 that $(Y_t)_{t\in[0,T]}$ satisfies

$$dY_t = Y_t\left(\eta dt + \mathbb{E}[\sigma\pi^*]dB_t\right).$$

Alternatively, the ratio $\varphi/(1+\psi)$ can be interpreted directly in terms of the type distribution. Define $R = \sigma^2/(\sigma^2+\nu^2)$, and note that $\varphi = \mathbb{E}[R\delta\mu/\sigma]$ and $\psi = \mathbb{E}[R\theta(\delta-1)]$. Notice that the assumptions on the parameter ranges ensure that $1+\psi > 0$. As before, the numerator φ increases as we increase the quality of the other stocks, as measured by their Sharpe ratio. However, the ratio $\varphi/(1+\psi)$ may not increase as the population becomes more risk tolerant (i.e., as δ increases on average), as both the numerator and denominator increase in this case.

The dependence of $\varphi/(1+\psi)$ and thus of π^* on the type distribution is rather complex. The distribution of competition weights θ appears only through ψ, and its effect is mediated by the risk tolerance δ. Loosely speaking, the population average of θ can have a positive or negative effect on $\pi^*_{(2)}$ depending on the "typical" sign of $(1-\delta)$.

In the single stock case of Corollary 18, the MFE π^* clearly resembles the Merton-type portfolio but with *effective risk tolerance parameter*

$$\delta_{\mathrm{eff}} := \delta - \frac{\theta(\delta-1)\overline{\delta}}{1+\overline{\theta(\delta-1)}}. \tag{56}$$

This simplifies some of the complex dependencies of π^* on the type distribution mentioned in the previous paragraph. For instance, suppose θ and δ are uncorrelated, so that $\overline{\theta(\delta-1)} = \overline{\theta}(\overline{\delta}-1)$. If $\overline{\delta} > 1$, then $|\delta_{\mathrm{eff}} - \delta|$ is decreasing in $\overline{\theta}$. That is, if the average risk tolerance is high, then, as the population becomes *more* competitive (i.e. $\overline{\theta}$ increases), the representative agent behaves less competitively in the sense that δ_{eff} moves closer to δ. On the other hand, if $\overline{\delta} < 1$, then $|\delta_{\mathrm{eff}} - \delta|$ is increasing in $\overline{\theta}$. That is, if the average risk tolerance is low, then, as the population becomes *more* competitive, the representative agent behaves more competitively in the sense that δ_{eff} moves away from δ. Again, if $\overline{\delta} = 1$, then $\overline{\theta}$ plays no role whatsoever.

A few other special cases are worth discussing. If $\sigma = 0$ a.s., there is no common noise. In this case, $\varphi = \psi = 0$, and in turn the MFE is equal to the Merton portfolio, meaning the agents are not competitive. On the other hand, if $\nu = 0$ a.s., there is no independent noise. In this case, $\varphi = \mathbb{E}[\delta\mu/\sigma]$ and $\psi = \mathbb{E}[\theta(\delta-1)]$, and the portfolio becomes

$$\pi^* = \delta\frac{\mu}{\sigma^2} - \frac{\theta(\delta-1)}{\sigma(1+\mathbb{E}[\theta(\delta-1)])}\mathbb{E}\left[\delta\frac{\mu}{\sigma}\right].$$

Lastly, if all agents have the same type vector (i.e., ζ is deterministic), then π^* is deterministic and equal to

$$\frac{\delta\mu}{(1+\theta(\delta-1))\sigma^2+\nu^2}.$$

4. Conclusions and extensions

We have considered optimal portfolio management problems under relative performance concerns for both finite and infinite populations. The agents have a common time horizon and either CARA or CRRA risk preference and trade individual stocks with log-normal dynamics driven by both a common and an idiosyncratic noise. The agents face competition in that their utility criterion depends both on their individual wealth as well as the wealth of the others. We have explicitly constructed the associated Nash and mean field equilibria.

Some natural mathematical questions remain open. For instance, while the equilibria we construct are unique among the class of *constant* equilibria, it is not clear if there can exist alternative non-constant equilibria. In a similar spirit, one might try to generalize our results to market parameters (μ, ν, σ) with time- or state-dependence, but this would significantly complicated the analysis.

Our study points to several directions for future research. A first direction is to further analyze the finite-population problems by using elements of indifference valuation. Indeed, as we mentioned in the proof of Theorem 3, we may identify the effect of competition as a liability and in turn solve an indifference valuation problem. There is a fundamental difference, of course, between the classical indifference pricing problem and the one herein; for us, the liability is essentially endogenous, as it depends on the actions of the agents. Nevertheless, employing indifference valuation arguments is expected to yield a clearer financial interpretation of the equilibrium strategies by relating them to indifference hedging strategies. It will also permit an analysis of sensitivity effects of varying agents' population size using arguments from the so-called relative indifference valuation. Such questions are left for future work.

In a different direction, a natural generalization of our model would allow agents to invest in any of the stocks, not just the individual stock assigned to them. Such a case has been recently analyzed in [2], and in [23] under forward performance criteria. Important questions arise on the effects of competition to asset specialization. While such generalization might be intractable for the finite population setting, a mean field formulation may provide a more tractable framework for studying the interactive role of competition and asset familiarity, specialization and competition.

Finally, one may extend the current model to dynamically evolving markets and rolling horizons. Such generalization may be analyzed under forward performance criteria, extending the results of [23], and would lead naturally to a new class of mean field games. It would also allow for further extending the concept of "forward benchmarking," introduced in [39].

References

1. V. Agarwal, N.D. Daniel, and N.Y. Naik, *Flows, performance, and managerial incentives in hedge funds*, EFA 2003 Annual Conference Paper No. 501, 2003.
2. S. Basak and D. Makarov, *Competition among portfolio managers and asset specialization*, Paris December 2014 Finance Meeting EUROFIDAI-AFFI Paper, 2015.
3. A. Bensoussan, J. Frehse, and S.C.P. Yam, *The master equation in mean field theory*, Journal de Mathématiques Pures et Appliquées **103** (2015), no. 6, 1441–1474.

4. A. Bensoussan, K.C.J. Sung, S.C.P. Yam, and S.P. Yung, *Linear-quadratic mean field games*, Journal of Optimization Theory and Applications **169** (2016), no. 2, 496–529.
5. P. Boyle, L. Garlappi, R. Uppal, and T. Wang, *Keynes meets Markowitz: The trade-off between familiarity and diversification*, Management Science **58** (2012), no. 2, 253–272.
6. M.J. Brennan, *The optimal number of securities in a risky asset portfolio when there are fixed costs of transacting: Theory and some empirical results*, Journal of Financial and Quantitative analysis **10** (1975), no. 03, 483–496.
7. S.J. Brown, W.N. Goetzmann, and J. Park, *Careers and survival: Competition and risk in the hedge fund and CTA industry*, The Journal of Finance **56** (2001), no. 5, 1869–1886.
8. P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions, *The master equation and the convergence problem in mean field games*, arXiv preprint arXiv:1509.02505 (2015).
9. R. Carmona and F. Delarue, *The master equation for large population equilibriums*, Stochastic Analysis and Applications 2014, Springer, 2014, pp. 77–128.
10. R. Carmona, F. Delarue, and A. Lachapelle, *Control of McKean–Vlasov dynamics versus mean field games*, Mathematics and Financial Economics **7** (2013), no. 2, 131–166.
11. R. Carmona, F. Delarue, and D. Lacker, *Mean field games with common noise*, The Annals of Probability **44** (2016), no. 6, 3740–3803.
12. R. Carmona, J.-P. Fouque, and L.-H. Sun, *Mean field games and systemic risk*, Communications in Mathematical Sciences **13** (2015), no. 4, 911–933.
13. P. Chan and R. Sircar, *Bertrand and Cournot mean field games*, Applied Mathematics & Optimization **71** (2015), no. 3, 533–569.
14. J. Chevalier and G. Ellison, *Risk taking by mutual funds as a response to incentives*, Journal of Political Economy **105** (1997), no. 6, 1167–1200.
15. J.D. Coval and T.J. Moskowitz, *Home bias at home: Local equity preference in domestic portfolios*, The Journal of Finance **54** (1999), no. 6, 2045–2073.
16. B. Ding, M. Getmansky, B. Liang, and R. Wermers, *Investor flows and share restrictions in the hedge fund industry*, Working paper (2008).
17. D. Duffie, W. Fleming, H.M. Soner, and T. Zariphopoulou, *Hedging in incomplete markets with HARA utility*, Journal of Economic Dynamics and Control **21** (1997), no. 4, 753–782.
18. G.-E. Espinosa and N. Touzi, *Optimal investment under relative performance concerns*, Mathematical Finance **25** (2015), no. 2, 221–257.
19. C. Frei and G. Dos Reis, *A financial market with interacting investors: does an equilibrium exist?*, Mathematics and Financial Economics **4** (2011), no. 3, 161–182.
20. X. Gabaix, J.-M. Lasry, P.-L. Lions, and B. Moll, *The dynamics of inequality*, Econometrica **84** (2016), no. 6, 2071–2111.
21. S. Gallaher, R. Kaniel, and L.T. Starks, *Madison avenue meets wall street: Mutual fund families, competition and advertising*, Working paper, 2006.
22. R. Gayduk and S. Nadtochiy, *Endogenous formation of limit order books: dynamics between trades*, arXiv preprint arXiv:1605.09720 (2016).
23. T. Geng and T. Zariphopoulou, *Passive and competitive strategies under relative forward performance criteria*, Available at SSRN 2870040 (2017).
24. O. Guéant, J.-M. Lasry, and P.-L. Lions, *Mean field games and applications*, Paris-Princeton lectures on mathematical finance 2010, Springer, 2011, pp. 205–266.
25. M. Huang, R.P. Malhamé, and P.E. Caines, *Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle*, Communications in Information & Systems **6** (2006), no. 3, 221–252.
26. M. Huang and S.L. Nguyen, *Mean field games for stochastic growth with relative utility*, Applied Mathematics & Optimization **74** (2016), no. 3, 643–668.
27. X. Huang, S. Jaimungal, and M. Nourian, *Mean-field game strategies for a major-minor agent optimal execution problem*, Available at SSRN 2578733 (2017).
28. R.E. Lucas Jr and B. Moll, *Knowledge growth and the allocation of time*, Journal of Political Economy **122** (2014), no. 1, 1–51.
29. M. Kacperczyk, C. Sialm, and L. Zheng, *On the industry concentration of actively managed equity mutual funds*, The Journal of Finance **60** (2005), no. 4, 1983–2011.
30. A. Kempf and S. Ruenzi, *Tournaments in mutual-fund families*, Review of Financial Studies **21** (2008), no. 2, 1013–1036.

31. T.S. Kim and E. Omberg, *Dynamic nonmyopic portfolio behavior*, Review of financial studies **9** (1996), no. 1, 141–161.
32. D. Lacker, *A general characterization of the mean field limit for stochastic differential games*, Probability Theory and Related Fields **165** (2016), no. 3, 581–648.
33. J.-M. Lasry and P.-L. Lions, *Mean field games*, Japanese Journal of Mathematics **2** (2007), no. 1, 229–260.
34. W. Li and A. Tiwari, *On the consequences of mutual fund tournaments*, Working paper (2006).
35. H. Liu, *Solvency constraint, underdiversification, and idiosyncratic risks*, Journal of Financial and quantitative Analysis **49** (2014), no. 02, 409–430.
36. R.C. Merton, *Optimum consumption and portfolio rules in a continuous-time model*, Journal of Economic Theory **3** (1971), no. 4, 373–413.
37. ______, *A simple model of capital market equilibrium with incomplete information*, The Journal of Finance **42** (1987), no. 3, 483–510.
38. T. Mitton and K. Vorkink, *Equilibrium underdiversification and the preference for skewness*, Review of Financial studies **20** (2007), no. 4, 1255–1288.
39. M. Musiela and T. Zariphopoulou, *An example of indifference prices under exponential preferences*, Finance and Stochastics **8** (2004), no. 2, 229–239.
40. S. Van Nieuwerburgh and L. Veldkamp, *Information immobility and the home bias puzzle*, The Journal of Finance **64** (2009), no. 3, 1187–1215.
41. ______, *Information acquisition and under-diversification*, The Review of Economic Studies **77** (2010), no. 2, 779–805.
42. E.R. Sirri and P. Tufano, *Costly search and mutual fund flows*, The journal of finance **53** (1998), no. 5, 1589–1622.
43. L.-H. Sun, *Systemic risk and interbank lending*, arXiv preprint arXiv:1611.06672 (2016).
44. Y. Sun, *The exact law of large numbers via Fubini extension and characterization of insurable risks*, Journal of Economic Theory **126** (2006), no. 1, 31–69.
45. R. Uppal and T. Wang, *Model misspecification and underdiversification*, The Journal of Finance **58** (2003), no. 6, 2465–2486.