ORIGINAL ARTICLE

MATHEMATICAL FINANCE An International Journal of Mathematics, Statistics and Financial Economics WILEY

Mean field and *n*-agent games for optimal investment under relative performance criteria

Daniel Lacker[1] | Thaleia Zariphopoulou[2,3]

[1]Industrial Engineering & Operations Research, Columbia University, New York

[2]Departments of Mathematics and IROM, The University of Texas at Austin, Austin, Texas

[3]The Mathematical Institute, University of Oxford, Oxford, UK

Correspondence
Daniel Lacker, Industrial Engineering & Operations Research, Columbia University, 306 Mudd, 500 West 120 St, New York City, NY 10027.
Email: dl3133@columbia.edu

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 investment horizon in log-normal markets. We construct explicit constant equilibrium strategies for both the finite population games and the corresponding mean field games, which we show are unique in the class of constant 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 underinvest, 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 (MFGs), 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 *idiosyncratic 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. We choose to work with one-dimensional stocks S^i for mere simplicity, but our analysis would adapt with purely notational changes to cover the case where each S^i is a vector of stocks available to agent i. In this multidimensional setting, the "single stock" case would model the realistic situation of a large number of agents trading in the same vector of stocks.

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.

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 ith 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 Espinosa and Touzi (2015).

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 ith agent's relative risk tolerance and relative competition weight. The geometric mean is used here largely for its tractability, but it also admits a natural interpretation: Quantities of the form $e^{r_1}, \ldots, e^{r_n}$ have geometric mean $\exp(\frac{1}{n}\sum_{i=1}^{n} r_i)$, which indicates that the geometric mean of wealths is simply the exponential of the arithmetic mean of returns. In this sense, the agents are using *returns* rather than *absolute wealth* in measuring relative performance.

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 ith 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 ith stock, respectively. The values $\pi_t^{i,*}$ may be negative, indicating that the agent shorts the stock.

Competition among fund managers is well documented in investment practice for both mutual and hedge funds; see, for example, Agarwal, Daniel, and Naik (2003); Basak and Makarov (2015); Brown, Goetzmann, and Park (2001); Chevalier and Ellison (1997); Ding, Getmansky, Liang, and Wermers (2008); Gallaher, Kaniel, and Starks (2006); Kempf and Ruenzi (2008); Li and Tiwari (2006); Sirri and Tufano (1998). 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 Basak and Makarov (2015) 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 Brennan (1975); Coval and Moskowitz (1999); Merton (1987). Other representative works include Boyle, Garlappi, Uppal, and Wang (2012); Kacperczyk, Sialm, and Zheng (2005); Liu (2014); Mitton and Vorkink (2007); Van Nieuwerburgh and Veldkamp (2009, 2010); Uppal and Wang (2003). 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 constants (i.e., chosen at time zero). 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 (but fixed) choice of competitors' constant 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 unique 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 constant Nash equilibrium in each model turns out to be the sum of two components. The first is the traditional Merton portfolio (see Merton, 1971), 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 portfolios (see, e.g., the so-called "nirvana" cases in Kim & Omberg, 1996).

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 $i = 1, \dots, n$. We show that these limiting strategies can be derived intrinsically, as equilibria of suitable 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-agent 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-agent games. For instance, Bielagk, Lionnet, and Dos Reis (2017), Espinosa and Touzi (2015), and Frei and Dos Reis (2011) study n-agent models similar to our CARA utility model but notably including equilibrium pricing and portfolio constraints, leading to difficult n-dimensional quadratic backward stochastic differential equations (BSDE) systems. An 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 $(\pi^*, \overline{X})$ 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-agent 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 $(\pi^*, \overline{X})$, 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. While its use in MFG theory appears to be new, this notion of geometric mean of a measure is essentially a special case of the well-studied concept of *generalized mean* (see Hardy, Littlewood, & Pólya, 1952, chapter III).

As in the finite population CARA and CRRA cases, we focus on the tractable class of equilibria in which the strategy π^* is 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 that agree with the limiting expressions from the n-agent games. In each model, the solution technique is analogous to the n-agent 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-agent counterparts and exhibit the same qualitative behavior and two-component structure discussed above.

MFGs, first introduced in Lasry and Lions (2007) and Huang, Malhamé, and Caines (2006), have by now found numerous applications in economics and finance, notably including models of income inequality (Gabaix, Lasry, Lions, & Moll, 2016), economic growth (Lucas Jr & Moll, 2014), limit order book formation (Gayduk & Nadtochiy, 2016), systemic risk (Carmona, Fouque, & Sun, 2015), optimal

execution (Huang, Jaimungal, & Nourian, 2017; Cardaliaguet & Lehalle, 2016), and oligopoly market models (Chan & Sircar, 2015), to name but a few. The closest works to ours are the static model of Guéant, Lasry, and Lions (2011, Section 6), which is a competitive variant of the Markowitz model, and the stochastic growth model of Huang and Nguyen (2016), 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 Bensoussan, Sung, Yam, and Yung (2016), Carmona, Delarue, and Lachapelle (2013), and Carmona et al. (2015), such examples are scarce, especially in the presence of common noise. The only other examples we know of are those in Guéant et al. (2011, Sections 5 and 7) as well as the more recent (Sun, 2016), 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* (cf. Carmona & Delarue, 2014). Moreover, we wish to emphasize the manner in which we incorporate different *types* of agents, by randomizing $\zeta = (\xi, \delta, \theta, \mu, \nu, \sigma)$ as described above. Several previous works on MFGs (e.g., Huang et al., 2006) incorporated finitely many types 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. While randomizing types is a standard technique in static games with a continuum of agents, the idea has scarcely appeared in the (dynamic) MFG literature; to the best of our knowledge, it has appeared only in Cardaliaguet and Lehalle (2016).

The paper is organized as follows. In Section 2, we present the exponential model and study both the *n*-agent 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 equilibrium and mean field equilibrium (MFE), 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 to the performances 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 agent 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 ith 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 $B, W^1, \dots, W^n$ 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

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

for some $\mu, \sigma > 0$ independent of i. Notably, the single stock case was studied in Espinosa and Touzi (2015) and Frei and Dos Reis (2011) 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 represents the (discounted by the bond) amount invested in the ith stock. The ith agent's wealth $(X_t^i)_{t \in [0,T]}$ then solves

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

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 ith 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

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

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

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

which highlights how the competition weight θ_i determines the ith 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 2.1. 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]

Remark 2.2. 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 $(\pi^{1,*}, \ldots, \pi^{n,*}) \in \mathbb{R}^n$. Note also that the definition of a constant Nash equilibrium still requires that the optimality condition (8) holds for *every* choice of alternative strategy, not just constant ones.

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

Theorem 2.3. *Assume that for all $i = 1, \dots, n$ we have $\delta_i > 0$, $\theta_i \in [0, 1]$, $\mu_i > 0$, $\sigma_i \geq 0$, $v_i \geq 0$, and $\sigma_i + v_i > 0$. Define the constants*

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

There are two cases:

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

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

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 2.4 (**Single stock**). *Assume that for all $i = 1, \dots, n$ we have $\mu_i = \mu > 0$, $\sigma_i = \sigma > 0$, and $v_i = 0$. Define the constants*

$$\bar{\delta} := \frac{1}{n} \sum_{k=1}^{n} \delta_k \quad and \quad \bar{\theta} := \frac{1}{n} \sum_{k=1}^{n} \theta_k.$$

There are two cases:

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

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

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

Proof. Apply Theorem 2.3, taking note of the simplifications $\varphi_n = \bar{\delta}\mu/\sigma$ and $\psi_n = \bar{\theta}$. □

Remark 2.5. 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 ith 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

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

We prefer our original formulation mainly because it results in simpler formulas for the equilibrium strategies in Theorem 2.3 and Corollary 2.4. 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.

Remark 2.6. Even in the absence of competition, there are well-known technical issues with exponential preferences in expected utility optimization, in that, the wealth may become arbitrarily negative. This has been studied and partially addressed, essentially by carefully redefining the class of admissible controls. In particular, one can define admissible strategies such that wealth processes are supermartingales under all martingale measures with finite entropy. One must then solve the dual problem, but some technical issues may still remain; see Delbaen et al. (2002) and Schachermayer (2001). On the other hand, more recent work has identified financially meaningful admissibility classes, which often (but not always) contain the desired optimizer; see Biagini and Černý (2011) and Biagini and Sîrbu (2014).

Proof of Theorem 2.3. 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 ith agent solves the optimization problem

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

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

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

$$dY_t = \widehat{\mu\alpha} dt + \widehat{\sigma\alpha} dB_t + \frac{1}{n} \sum_{k \neq i} \nu_k \alpha_k dW_t^k, \qquad Y_0 = \frac{1}{n} \sum_{k \neq i} x_0^k,$$

and where we have abbreviated

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

In the sequel, we will also use the abbreviation

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

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

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

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

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

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

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

with $f(T) = 1$ and

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

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

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

The maximum in (12) is achieved at

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

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

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

We have thus constructed a smooth solution of the HJB equation and calculated its associated feedback policy, which is constant and thus admissible. Using the explicit form (14) and the admissibility of this candidate control, we can establish a verification theorem following well-known arguments in stochastic optimization (Fleming & Soner, 2006; Pham, 2009; Touzi, 2012).

Alternatively, we note 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 $\gamma_i := \frac{1}{\delta_i}(1 - \frac{\theta_i}{n})$. Similar problems have been studied in Henderson

(2002) and Musiela and Zariphopoulou (2004), to which we refer the reader for more detailed arguments for the specific optimization problem at hand.

Therefore, for a candidate portfolio vector $(\alpha_1, \ldots, \alpha_n)$ to be a constant 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

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

which implies that

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

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 $\sigma_i + v_i > 0$ by assumption, this implies $\psi_n = 0$, which is a contradiction.

Remark 2.7. 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 $\widehat{v\alpha}$ are now known. However, we omit this tedious calculation.

It remains an open problem to determine if there exist Nash equilibria that are not constant, e.g., equilibria of the feedback form $\pi^{i,*} = \pi^{i,*}(t, x_1, \ldots, x_n)$, depending on time and/or the agents' wealths. Actually, an adaptation our argument could show that the constant Nash equilibrium we derived is, in fact, unique among the broader class of equilibria involving wealth-independent but potentially time-dependent strategies $\pi = \pi(t)$. The only delicate point is the solvability of the HJB equation (12), which would force us to consider time-dependence with appropriate smoothness and growth conditions.

In addition, focusing on constant Nash equilibria can be justified as follows. It is well known that, for log-normal models, exponential utilities lead to constant strategies. This holds not only for the plain investment problem but also, in the absence of competition, for indifference-type investment problems. The individual optimization problems we encountered in (11) are directly analogous to these standard

indifference-type problems. Wealth independence is well documented in this context, even in general semimartingale models. This is precisely the reason exponential utilities are so popular in the areas of asset price equilibrium and indifference valuation.

2.2 | The MFG

In this section, we study the limit as $n \to \infty$ of the n-agent game analyzed in the previous section.

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

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

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 2.3 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 (cf. (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 $\pi^{i,*}$ (cf. (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

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

The 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 2.9 will formalize this intuition.

2.2.1 | Formulating the MFG

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 both 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-agent 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 2.8. 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-agent 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 2.3, 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

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

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

Definition 2.9. Let $\pi^* \in \mathcal{A}_{\mathrm{MF}}$ be an admissible strategy, and consider the $\mathcal{F}_T^B$-measurable random variable $\overline{X} := \mathbb{E}[X_T^* \mid \mathcal{F}_T^B]$, where $(X_t^*)_{t \in [0,T]}$ is the wealth process in (19) corresponding to the

strategy π^*. We say that π^* is a MFE if π^* is optimal for the optimization problem (20) corresponding to this choice of $\overline{X}$.

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

Typically, an 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 $\mathbb{E}[X_T^* \mid \mathcal{F}_T^B]$. If we have a fixed point in the sense that the *consistency condition*, $\mathbb{E}[X_T^* \mid \mathcal{F}_T^B] = \overline{X}$, holds, then π^* is an 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 independent and identically distributed (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 $\mathbb{E}[X_T^* \mid \mathcal{F}_T^B]$. 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 Carmona, Delarue, and Lacker (2016) for a detailed discussion of MFGs with common noise, and to Cardaliaguet, Delarue, Lasry, and Lions (2015) and Lacker (2016) for general results on limits of n-agent 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 (Sun, 2006).

2.2.2 | An alternative formulation of the MFG

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, note that for a fixed $\mathcal{F}_T^B$-measurable random variable $\overline{X}$ we have

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

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

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

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 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 an MFE in the proof of Theorem 2.10, 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 MFG

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-agent equilibrium strategies in the sense of (18).

Theorem 2.10. *Assume that, a.s., $\delta > 0$, $\theta \in [0, 1]$, $\mu > 0$, $\sigma \geq 0$, $v \geq 0$, and $\sigma + v > 0$. Define the constants*

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

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*

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

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*-agent game, presented in Corollary 2.4.

Corollary 2.11 (Single stock). *Suppose (μ, v, σ) are deterministic, with $v = 0$ and $\mu, \sigma > 0$. Define the constants*

$$\bar{\delta} := \mathbb{E}[\delta] \quad and \quad \bar{\theta} := \mathbb{E}[\theta].$$

There are two cases:

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

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

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

Proof of Theorem 2.10. 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 $\mathbb{E}[\alpha^2] < \infty$.

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

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

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 $\mathbb{E}[\upsilon(\xi - \theta\overline{\xi}, 0)]$, where $\upsilon(x, t)$ is the unique (smooth, strictly concave, and strictly increasing in x) solution of the HJB equation

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

with terminal condition $\upsilon(x, T) = -e^{-x/\delta}$. 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

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

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

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

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

$$\rho := \frac{\left(\mu + \frac{\theta}{\delta}\sigma\overline{\sigma\alpha}\right)^2}{2(\sigma^2 + v^2)} - \frac{\theta}{\delta}\overline{\mu\alpha} - \frac{1}{2}\left(\frac{\theta}{\delta}\overline{\sigma\alpha}\right)^2. \tag{26}$$

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

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

In fact, $\pi^* = \pi^*(x, t)$ is $\mathcal{F}_0^{MF}$-measurable and does not depend on (x, t). The optimality of π^* for the problem (24) follows.

Recalling Definition 2.9, we see that for the candidate control α to be a constant MFE, we need $\alpha = \pi^*$. In light of (27), π^* is a constant MFE if it solves the equation

$$\pi^* = \delta\frac{\mu}{\sigma^2 + v^2} + \theta\frac{\sigma}{\sigma^2 + v^2}\overline{\sigma\pi^*}. \tag{28}$$

Multiply both sides by σ and average to find that $\overline{\sigma\pi^*}$ must satisfy

$$\overline{\sigma\pi^*} = \mathbb{E}\left[\delta\frac{\mu\sigma}{\sigma^2 + v^2}\right] + \mathbb{E}\left[\theta\frac{\sigma^2}{\sigma^2 + v^2}\right]\overline{\sigma\pi^*} = \varphi + \psi\overline{\sigma\pi^*}. \tag{29}$$

We then have the following cases:

(i) If $\psi < 1$, the above yields $\overline{\sigma\pi^*} = \varphi/(1 - \psi)$, and using Equation (28) we prove part (1).

(ii) If $\psi = 1$ but $\varphi \neq 0$, then Equation (29) has no solution, and, as a result, 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 would imply $\sigma = 0 < v$ a.s., which would in turn yield $\psi = 0$, a contradiction. This completes the proof of part (2).

Remark 2.12. Note that the proof above yields a tractable formula for the equilibrium value function of the representative agent. Because 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

$$v(\xi - \theta\overline{\xi}, 0) = -\exp\left(-\frac{1}{\delta}(\xi - \theta\overline{\xi}) - \rho T\right). \tag{30}$$

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

$$\rho = \frac{1}{2(\sigma^2 + v^2)}\left(\mu + \frac{\theta}{\delta}\frac{\varphi}{1 - \psi}\sigma\right)^2 - \frac{\theta}{\delta}\left(\widetilde{\psi} + \frac{\widetilde{\varphi}\varphi}{1 - \psi}\right), \tag{31}$$

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

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

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

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

Remark 2.13. Equation (30) essentially provides the solution of the so-called *master equation*; see Bensoussan, Frehse, and Yam (2015) or Carmona and Delarue (2014) for an introduction to the master equation in MFG theory. Indeed, the master equation is a partial differential equation (PDE) providing a function $U = U(x, m, t)$, where $(x, t) \in \mathbb{R} \times [0, T]$ and m is a probability measure on $\mathbb{R}$. The value $U(x, m, t)$ is naturally interpreted as the value at time t for a representative agent starting with wealth $X_t = x$, when the distribution of other agents' wealths is m. In our case, this value is nothing but

$$U(x, m, t) = v(x - \theta\bar{m}, t) = -\exp\left(-\frac{1}{\delta}(x - \theta\bar{m}) - \rho(T - t)\right),$$

where $\bar{m}$ is the mean of the measure m. We do not attempt to make this any more precise, as the concept of a master equation has not yet been settled for models with different types of agents.

2.3 | Discussion of the equilibrium

We focus most of the discussion on the MFE of Theorem 2.10 and Corollary 2.11, as the n-agent equilibria of Theorem 2.3 and Corollary 2.4 have essentially the same structure.

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 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, so 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 quantity can be viewed 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 the quality of the other stocks increases, as measured by their Sharpe ratio (higher μ/σ on average).
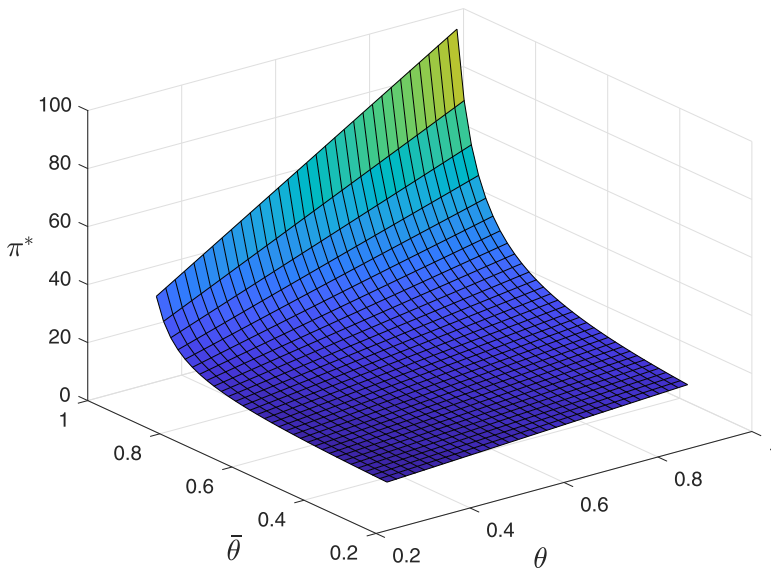
FIGURE 1 Single stock case (Corollary 2.11): π^* versus θ and $\bar{\theta}$, with $\delta = 5$, $\bar{\delta} = 6$, and $\mu = \sigma = 1$ [Color figure can be viewed at wileyonlinelibrary.com]

Some of the effects of competition are more transparent in the single stock case of Corollary 2.11. The resulting MFE π^* clearly resembles the Merton portfolio but with *effective risk tolerance parameter*

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

We always have $\delta_{\text{eff}} > \delta$ if $\theta > 0$, and the difference $\delta_{\text{eff}} - \delta$ increases with θ, with $\bar{\delta}$, and with $\bar{\theta}$. 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 $\bar{\delta}$ and $\bar{\theta}$ 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. At the extreme ends, as both θ and $\bar{\theta}$ approach 1, π^* blows up very quickly; that is, a highly competitive agent in a population of highly competitive agents invests significantly in the risky asset. This is illustrated in Figure 1.

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

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

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

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

Finally, we comment on the effect of population size on the equilibrium in the n-agent game given in Theorem 2.3. The only real difference compared to the mean field setting is the rescaling of v_k^2 by $(1 - \theta_k/n)$ wherever it appears in Theorem 2.3, with no rescaling present in the single stock case of Corollary 2.4. It is not yet clear how to properly interpret this rescaling, and it is worth noting that the change of variables discussed in Remark 2.5 does not significantly change the situation. Interpreting the average wealth $\frac{1}{n-1}\sum_{j\neq i} X_T^j$ as a liability term in an indifference valuation problem, as we have mentioned before, seems promising, but we do not pursue this any further here.

3 | CRRA RISK PREFERENCES

In this section, we focus on power and logarithmic (CRRA) 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 Basak and Makarov (2015) and more recently in Anthropelos, Geng, and Zariphopoulou (2017) under forward relative performance criteria.

3.1 | The n-agent game

We consider an n-agent 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

$$dX_t^i = \pi_t^i X_t^i \left(\mu_i dt + v_i dW_t^i + \sigma_i dB_t \right), \tag{32}$$

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 ith 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,

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

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

$$J_i(\pi^1, \ldots, \pi^n) := \mathbb{E}\left[U\left(X_T^i \overline{X}_T^{-\theta_i}; \delta_i\right)\right], \quad \text{where} \quad \overline{X}_T = \left(\prod_{k=1}^n X_T^k\right)^{1/n}. \tag{33}$$

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 mean 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

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

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 2.1, but with the new objective function defined in (33) above. We find a unique *constant* equilibrium in the following theorem, which we subsequently specialize to the single stock case.

Theorem 3.1. *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*

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

and

$$\psi_n := \frac{1}{n}\sum_{k=1}^n \theta_k(\delta_k - 1)\frac{\sigma_k^2}{\sigma_k^2 + \nu_k^2(1 + (\delta_k - 1)\theta_k/n)}. \tag{35}$$

There exists a unique constant equilibrium, given by

$$\pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2(1 + (\delta_i - 1)\theta_i/n)} - \theta_i(\delta_i - 1)\frac{\sigma_i}{\sigma_i^2 + \nu_i^2(1 + (\delta_i - 1)\theta_i/n)}\frac{\varphi_n}{1 + \psi_n}. \tag{36}$$

Moreover, we have the identity

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

Corollary 3.2 (Single stock). *Assume that for all $i = 1, \ldots, n$ we have $\mu_i = \mu > 0$, $\sigma_i = \sigma > 0$, and $\nu_i = 0$. Define the constants*

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

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

Remark 3.3. As in Remark 2.5 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 (33) 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

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

However, this is only valid if $(1 - 1/\delta'_i)(1 + \frac{1}{n-1}\theta'_i) < 1$, which ensures that $\delta_i > 0$. This certainly holds for sufficiently large n. We favor our original parameterization because of the relative simplicity of the formulas in Theorem 3.1 and Corollary 3.2, and because there is no difference in the $n \to \infty$ limit.

Proof of Theorem 3.1. The proof is similar to that of Theorem 2.3, 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 (32) with constant weights α_k and $X_0^k = x_0^k$.

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

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

In turn,

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

where we abbreviate

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

$$\widehat{\Sigma\alpha^2} := \frac{1}{n} \sum_{k \neq i} \Sigma\alpha_k^2, \quad \text{and} \quad \widehat{(v\alpha)^2} := \frac{1}{n} \sum_{k \neq i} v_k^2 \alpha_k^2.$$

Thus, the process Y_t solves

$$\frac{dY_t}{Y_t} = \eta dt + \frac{1}{n} \sum_{k \neq i} v_k \alpha_k dW_t^k + \widehat{\sigma\alpha} dB_t, \quad Y_0 = \left(\prod_{k \neq i} x_0^k \right)^{1/n}, \tag{37}$$

with

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

The ith agent then solves the optimization problem

$$\sup_{\pi^i \in \mathcal{A}} \mathbb{E} \left[U \left((X_T^i)^{1-\theta_i/n} Y_T^{-\theta_i}; \delta_i \right) \right], \tag{38}$$

where

$$dX_t^i = \pi_t^i X_t^i \left(\mu_i dt + v_i dW_t^i + \sigma_i dB_t \right), \quad X_0^i = x_0^i,$$

with $(Y_t)_{t \in [0,T]}$ solving (37). We then obtain that the value (38) is equal to $v(X_0^i, Y_0, 0)$, where $v(x, y, t)$ solves the HJB equation

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

$$+ \frac{1}{2} \left(\widehat{\sigma\alpha}^2 + \frac{1}{n} \widehat{(v\alpha)^2} \right) y^2 v_{yy} + \eta y v_y = 0, \tag{39}$$

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

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

Applying the first-order conditions, the maximum in (39) is attained by

$$\pi^{i,*}(x, y, t) = -\frac{\mu_i x v_x(x, y, t) + \sigma_i \widehat{\sigma\alpha} x y v_{xy}(x, y, t)}{\left(\sigma_i^2 + v_i^2 \right) x^2 v_{xx}(x, y, t)}. \tag{40}$$

In turn, Equation (39) reduces to

$$v_t - \frac{1}{2} \frac{(\mu_i x v_x + \sigma_i \widehat{\sigma\alpha} x y v_{xy})^2}{\left(\sigma_i^2 + v_i^2 \right) x^2 v_{xx}} + \frac{1}{2} \left(\widehat{\sigma\alpha}^2 + \frac{1}{n} \widehat{(v\alpha)^2} \right) y^2 v_{yy} + \eta y v_y = 0. \tag{41}$$

Working as in the proof of Theorem 2.3, we deduce that the above HJB equation has a unique smooth solution (in an appropriate class of time-separable and space-homogeneous solutions), and the optimal feedback control in (40) reduces to

$$\pi^{i,*} = \frac{\delta_i \mu_i - \sigma_i \widehat{\sigma\alpha} \theta_i \left(\delta_i - 1 \right)}{(\sigma_i^2 + v_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))}. \tag{42}$$

We prove this in two cases:

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

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

reduces Equation (41) to $(1 - 1/\delta_i)^{-1} f'(t) + \rho f(t) = 0$, with $f(T) = 1$, where

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

$$+ \frac{1}{2}\left(\widehat{\sigma\alpha}^2 + \frac{1}{n}\widehat{(v\alpha)^2}\right)\theta_i(1 + \theta_i(1 - 1/\delta_i)).$$

We easily deduce that the solution of (41) 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 (40) yields (42).

(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 (41) to $f'(t) + \rho = 0$, with $f(T) = 0$ and

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

In turn, the solution of (41) is given by

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

and (40) reduces to $\pi^{i,*} = \mu_i/(\sigma_i^2 + v_i^2)$, which is consistent with (42) for $\delta_i = 1$.

We conclude the proof as follows. 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 (42) 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

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

Solving for α_i yields,

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

$$= \frac{\mu_i - \sigma_i \overline{\sigma\alpha}\theta_i(1 - 1/\delta_i)}{\left(\sigma_i^2 + v_i^2\right)(1 - (1 - \theta_i/n)(1 - 1/\delta_i)) - \sigma_i^2(\theta_i/n)(1 - 1/\delta_i)}$$

$$= \frac{\mu_i - \sigma_i \overline{\sigma\alpha}\theta_i(1 - 1/\delta_i)}{\sigma_i^2/\delta_i + v_i^2(1 - (1 - \theta_i/n)(1 - 1/\delta_i))} = \frac{\mu_i \delta_i - \sigma_i \overline{\sigma\alpha}\theta_i(\delta_i - 1)}{\sigma_i^2 + v_i^2(1 - \theta_i/n + \delta_i\theta_i/n)}. \tag{43}$$

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

$$\overline{\sigma\alpha} = \varphi_n - \psi_n \overline{\sigma\alpha}, \tag{44}$$

where φ_n, ψ_n are as in (34) and (35). Because $1 + \psi_n > 0$, Equation (44) holds if and only if $\overline{\sigma\alpha} = \varphi_n/(1 + \psi_n)$. We then deduce from (43) that the equilibrium strategy $\alpha_i = \pi^{i,*}$ is given by (36).

Remark 3.4. Note that Equation (43) above has a *unique* solution for *all* parameter values. In contrast, the analogous equation (29) in the exponential case has no solutions for certain parameter values, which is why there were two cases in Theorem 2.3.

It is worth highlighting that, for the CRRA case, we assume that relative performance concerns appear multiplicatively, and not additively. There are two reasons for this. First, as discussed in the Introduction, this is natural in modeling preferences, which depend on relative *return* as opposed to relative *wealth*; see also Basak and Makarov (2015) for a discussion. The second reason is mathematical tractability. We have already seen that using the geometric mean leads to explicit solutions.

To formulate an analogous problem using an arithmetic mean, we may consider the following two possibilities. First, we may modify the optimization criterion to be of the form

$$U\left(\frac{X_T^i}{\frac{1}{n}\sum_{i=1}^n X_T^i}; \delta_i\right).$$

The challenge here is that the ratio appearing in the first argument cannot be expressed as the solution of a one-dimensional stochastic differential equation (SDE). The proofs of both Theorems 3.1 and 3.6 exploited the fact that the geometric mean of geometric Brownian motions remains a geometric Brownian motion, whereas the arithmetic mean enjoys no such properties.

Alternatively, we may use an optimization criterion of the form $U(X_T^i - \frac{1}{n}\sum_{i=1}^n X_T^i; \delta_i)$, which, however, runs into more serious problems because $U(x, \delta_i)$ is well defined and finite only for $x > 0$ (or $x \geq 0$ if $\delta_i > 1$). Hence, this criterion would enforce the hard constraint $X_T^i > \frac{1}{n}\sum_{i=1}^n X_T^i$ a.s., which raises the two natural questions of how this constraint propagates to previous times and whether this results in a meaningful class of solutions.

In short, using an arithmetic mean criterion would give rise to interdependent state and control constraints, which will likely render the problem intractable, and, at worst, could lead to trivial or meaningless solutions.

3.2 | The MFG

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

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

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

$$\mathcal{Z}^p := (0, \infty) \times (0, \infty) \times [0, 1] \times (0, \infty) \times [0, \infty) \times [0, \infty) \tag{45}$$

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 3.1 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 $\pi^{i,*}$ (cf. (36)) should converge to

$$\lim_{n \to \infty} \pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2} - \theta_i(\delta_i - 1) \frac{\sigma_i}{\sigma_i^2 + \nu_i^2} \frac{\varphi}{1 - \psi}, \tag{46}$$

where

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

As in the exponential case, we will demonstrate that this limiting strategy is indeed the equilibrium of an MFG, 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

$$dX_t = \pi_t X_t(\mu dt + \nu dW_t + \sigma dB_t), \quad X_0 = \xi, \tag{47}$$

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

$$\sup_{\pi \in \mathcal{A}_{\mathrm{MF}}} \mathbb{E}\left[U(X_T \overline{X}^{-\theta}; \delta)\right], \tag{48}$$

where $(X_t)_{t \in [0,T]}$ is given by (47).

The definition of an MFE is analogous to Definition 2.9. 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

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

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 3.5. Let $\pi^* \in \mathcal{A}_{\mathrm{MF}}$ be an admissible strategy, and consider the $\mathcal{F}_T^B$-measurable random variable $\overline{X} := \exp \mathbb{E}[\log X_T^* \mid \mathcal{F}_T^B]$, where $(X_t^*)_{t \in [0,T]}$ is the wealth process in (47) corresponding to the strategy π^*. We say that π^* is an MFE if π^* is optimal for the optimization problem (48) 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 an MFE.

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

Theorem 3.6. *Assume that, a.s., $\delta > 0$, $\theta \in [0,1]$, $\mu > 0$, $\sigma \geq 0$, $v \geq 0$, and $\sigma + v > 0$. Define the constants*

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

where we assume that both expectations exist and are finite.
There exists a unique constant MFE, given by

$$\pi^* = \delta \frac{\mu}{\sigma^2 + v^2} - \theta(\delta - 1) \frac{\sigma}{\sigma^2 + v^2} \frac{\varphi}{1 + \psi}. \tag{49}$$

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-agent game, presented in Corollary 3.2:

Corollary 3.7 (Single stock). *Suppose (μ, v, σ) are deterministic, with $v = 0$ and $\mu, \sigma > 0$. Define the constants*

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

There exists a unique constant MFE, given by

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

Proof of Theorem 3.6. As in the exponential case, we first reduce the optimal control problem (48) 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 (47) with an admissible constant strategy α. That is, α is an $\mathcal{F}_0^{MF}$-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

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

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

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

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

$$d\check{X}_t^\alpha = \left(\overline{\mu\alpha} - \frac{1}{2}\overline{\Sigma\alpha^2}\right)dt + \overline{\sigma\alpha}dB_t, \tag{50}$$

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

$$dY_t = de^{\check{X}_t^\alpha} = Y_t\left(\eta dt + \overline{\sigma\alpha}dB_t\right), \quad Y_0 = \overline{\xi}, \tag{51}$$

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

To solve the stochastic optimization problem (48), we equivalently solve

$$\sup_{\pi \in \mathcal{A}_{MF}} \mathbb{E}\left[U(X_T Y_T^{-\theta}; \delta)\right] \tag{52}$$

with

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

and $(Y_t)_{t \in [0,T]}$ solving (51). Then, as in the discussion of Section 2.2.2, the value of (52) 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

$$v_t + \max_{\pi \in \mathbb{R}}\left(\frac{1}{2}\Sigma\pi^2 x^2 v_{xx} + \pi\left(\mu x v_x + \sigma\overline{\sigma\alpha}xy v_{xy}\right)\right) + \frac{1}{2}\overline{\sigma\alpha}^2 y^2 v_{yy} + \eta y v_y = 0, \tag{53}$$

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

Applying the first-order conditions, the maximum in (53) is attained by

$$\pi^*(x, y, t) = -\frac{\mu x v_x(x, y, t) + \sigma \overline{\sigma \alpha} x y v_{xy}(x, y, t)}{\Sigma x^2 v_{xx}(x, y, t)}. \tag{54}$$

In turn, Equation (53) reduces to

$$v_t - \frac{(\mu x v_x + \sigma \overline{\sigma \alpha} x y v_{xy})^2}{2\Sigma x^2 v_{xx}} + \frac{1}{2} \overline{\sigma \alpha}^2 y^2 v_{yy} + \eta y v_y = 0. \tag{55}$$

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

$$\pi^*(x, y, t) = \Sigma^{-1} \left(\mu \delta - \theta(\delta - 1) \sigma \overline{\sigma \alpha} \right). \tag{56}$$

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

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

We easily deduce that the solution of (55) is

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

and that (54) yields (56).

(ii) Suppose $\delta = 1$. It is easily checked that the solution v of (55) 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, (54) becomes $\pi^*(x, y, t) = \mu/\Sigma$, which is consistent with (56) for $\delta = 1$.

Recalling Definition 3.5, we see that for the candidate control α to be a constant MFE, we need $\alpha = \pi^*$. In light of (3.5), π^* is a constant MFE if it solves the equation

$$\pi^* = \Sigma^{-1} \left(\mu \delta - \theta(\delta - 1) \sigma \overline{\sigma \pi^*} \right). \tag{57}$$

Multiply both sides by σ and average to find that $\overline{\sigma \pi^*}$ must satisfy

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

We then deduce that $\overline{\sigma \alpha} = \varphi/(1 + \psi)$, and plugging this into (56) we obtain (49).

3.3 | Discussion of the equilibrium

Some of the structural properties of the equilibrium are similar to those observed in the CARA model in Section 2.3. We again focus the discussion here on the mean field case of Theorem 3.6 and Corollary 3.7, as the n-agent equilibria of Theorem 3.1 and Corollary 3.2 have essentially the same structure. The only difference is the rescaling of v_k^2 by $(1 + (\delta_k - 1)\theta_k/n)$ wherever it appears in Theorem 3.1. We first discuss the general case of Theorem 3.6, before concentrating on the single stock case of Corollary 3.7.

3.3.1 | The general case

The MFE π^* of Theorem 3.6 can be written as the sum of two components, $\pi^* = \pi_{(1)}^* + \pi_{(2)}^*$, where $\pi_{(1)}^* = \delta\mu/(\sigma^2 + v^2)$ is the classical Merton portfolio, and

$$\pi_{(2)}^* := -\theta(\delta - 1)\frac{\sigma}{\sigma^2 + v^2}\frac{\varphi}{1 + \psi}.$$

The second component $\pi_{(2)}^*$ isolates the linear effect of the competition parameter θ. Notably, $\pi_{(2)}^*$ vanishes when $\theta = 0$.

Interestingly, the effect of competition is quite different in the CRRA model than in the CARA model, in the sense that competition now 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 $(1 - \delta)$, 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 is also easily deduced from the original problem formulation.

In fact, a highly risk-tolerant and competitive agent may choose to short the stock. That is, if $\delta > 1$ and θ is close to 1, π^* may be negative. This typically occurs when δ is much higher than their population averages, or, in other words, when the representative agent is very risk tolerant and competitive relative to the other agents.

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 (47) using the strategy π^*). The geometric average wealth of the population at time $t \in [0, T]$ is $Y_t := \log \mathbb{E}[\exp(X_t^*)|\mathcal{F}_T^B]$, and, as we saw in the proof of Theorem 3.6, it satisfies

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

Alternatively, the ratio $\varphi/(1 + \psi)$ can be interpreted directly in terms of the type distribution. Define $R = \sigma^2/(\sigma^2 + v^2)$, 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 the quality of the other stocks increases, 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 $(1 - \delta)$. These complexities are more easily unraveled in the single stock case.

3.3.2 | Single stock case

From the results of Corollary 3.7, we may write the equilibrium portfolio in the single stock case as

$$\pi^* = ((1 - k\theta)\delta + k\theta)\frac{\mu}{\sigma^2}, \tag{58}$$

where

$$k := \frac{\bar{\delta}}{1 + \overline{\theta(\delta - 1)}}. \tag{59}$$

The equilibrium π^* can thus be written as a Merton portfolio, $\pi^* = \delta_{\text{eff}}\mu/\sigma^2$, with *effective risk tolerance parameter*

$$\delta_{\text{eff}} := (1 - k\theta)\delta + k\theta = \delta - \frac{\theta(\delta - 1)\bar{\delta}}{1 + \overline{\theta(\delta - 1)}}. \tag{60}$$

This representation 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 $\overline{\theta(\delta - 1)} = \bar{\theta}(\bar{\delta} - 1)$. If $\bar{\delta} > 1$, then $|\delta_{\text{eff}} - \delta|$ is decreasing in $\bar{\theta}$. That is, if the average risk tolerance is high, then, as the population becomes more competitive (i.e., $\bar{\theta}$ increases), the representative agent behaves less competitively in the sense that δ_{eff} moves closer to δ. On the other hand, if $\bar{\delta} < 1$, then $|\delta_{\text{eff}} - \delta|$ is increasing in $\bar{\theta}$. 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 $\bar{\delta} = 1$, then $\bar{\theta}$ and θ play no role whatsoever.

More interesting is the joint effect of $(\delta, \bar{\theta})$ on π^*, when the other parameters are fixed. Still assuming θ and δ are uncorrelated, notice that the value of k can range between 1 and $\bar{\delta}$, as $\bar{\theta}$ varies between 0 and 1. Hence, if $\theta\bar{\delta} > 1$, then there is a critical value, $\bar{\theta}_{\text{crit}} := (\theta\bar{\delta} - 1)/(\bar{\delta} - 1)$, at which the effect of δ on π^* changes sign. When the population is highly competitive (i.e., $\bar{\theta} > \bar{\theta}_{\text{crit}}$), the investment π^* in the risky asset increases with the risk tolerance δ, as one might expect. On the other hand, when the population is less competitive (i.e., $\bar{\theta} < \bar{\theta}_{\text{crit}}$), π^* is decreasing in δ. This effect is illustrated in Figure 2.

A similar transition appears in the joint effect of (δ, θ) on π^*, when the other parameters are fixed. When $k \leq 1$ (which is equivalent to $\bar{\delta} \leq 1$ if we assume θ and δ are uncorrelated), then the risky investment π^* is increasing in the risk tolerance δ, for any value of θ. On the other hand, if $k < 1$, then π^* is increasing in δ if and only if $\theta < 1/k$. This situation is illustrated in Figure 3. Note that these effects are more pronounced if δ and θ are positively correlated and less pronounced if they are negatively correlated.

There are two ways to explain the counterintuitive phenomenon described in the previous two paragraphs, in which π^* may be decreasing in δ for certain (fixed) values of the other parameters. This regime happens precisely when $1 < k\theta$. Assuming again that θ and δ are uncorrelated, the latter inequality is equivalent to

$$\theta > \frac{1}{\delta}(1 - \bar{\theta}) + \bar{\theta}.$$
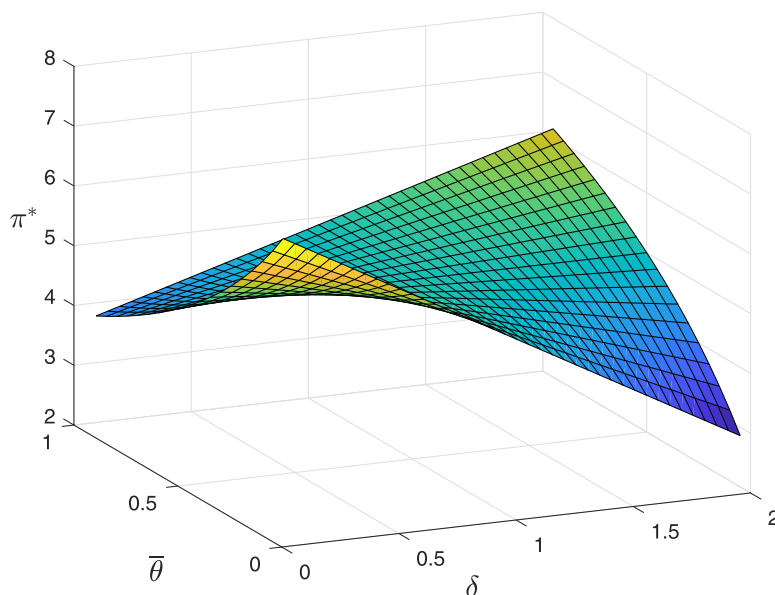
FIGURE 2 Single stock case (Corollary 3.7): π^* versus δ and $\overline{\theta}$, with $\theta = 3/4$, $\overline{\delta} = 2$, $\mu = 5$, and $\sigma = 1$; here, θ and δ are uncorrelated, and $\overline{\theta}_{\text{crit}} = 1/2$ [Color figure can be viewed at wileyonlinelibrary.com]

Assume for the sake of argument that $\overline{\delta}$ is extremely large. In the limit $\overline{\delta} \to \infty$, we see that $1 < k\theta$ if and only if $\theta > \overline{\theta}$, and the expression for π^* becomes

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

There are two immediate observations:

 (i) If $\theta < \overline{\theta}$ and δ is large, then π^* is positive and large. That is, less competitive agents go long.

(ii) If $\theta > \overline{\theta}$ and δ is large, then π^* is very negative. That is, competitive agents go short.

The first case, $\theta < \overline{\theta}$, is natural: A less competitive agent behaves more like a Merton investor, with π^* increasing in δ. On the other hand, we may explain the second regime, $\theta > \overline{\theta}$, as follows. Still assuming that the average risk tolerance is very large, we know from (i) that most other agents will take large long positions in the stock. If our representative agent were to go long, he could not afford to accept enough risk to go *as long* as the other investors. Then, if the stock price goes up, the representative agent may achieve a high absolute performance but would suffer in terms of relative performance, as the other agents who invested even more in the stock would earn even higher returns. Hence, a natural strategy is to short the stock, to focus on gains from outperforming competitors when the stock price drops, rather than focusing on absolute performance. This is still reasonable if the representative agent is himself quite risk tolerant, willing to accept risk in the opposite direction as the more Merton-like investors. Note that these effects are less pronounced but still present outside of the asymptotic regime $\overline{\delta} \to \infty$.

There is an alternative explanation in the spirit of (Basak & Makarov, 2015, pp. 13–14). A risk-averse agent typically seeks to minimize volatility by investing less in the stock than a more risk-tolerant agent.
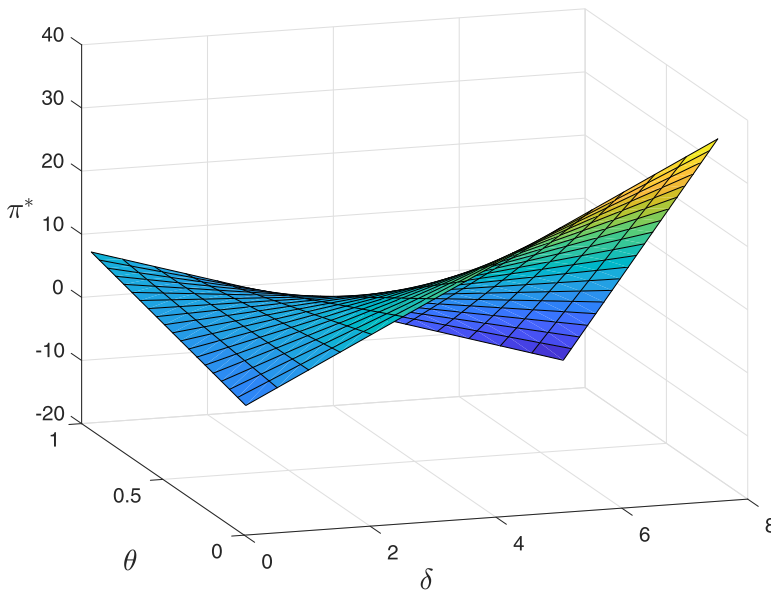
FIGURE 3 Single stock case (Corollary 3.7): π^* versus δ and θ, with $\bar{\theta} = 1/5$, $\bar{\delta} = 2$, and $\mu = 5$, $\sigma = 1$; here δ and θ are uncorrelated, and $k = 5/3$ [Color figure can be viewed at wileyonlinelibrary.com]

However, relative performance concerns provide an additional source of volatility. A risk-averse agent may then invest heavily in the stock in an attempt to mitigate losses from being outperformed.

3.3.3 | Some additional special cases

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, which means that the agents are not at all competitive. On the other hand, if $v = 0$ 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 optimal portfolio becomes

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

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

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

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 investment horizon and either CARA or CRRA risk preferences, and they trade individual stocks with log-normal dynamics driven by both common and idiosyncratic noises. They face competition in that their individual utility criteria depend on both their individual wealth as well as the wealth of the others. We have explicitly constructed the associated constant Nash equilibrium and MFE.

Our study points to several directions for future research. A first direction is to further analyze the finite population problems by using concepts from indifference valuation. Indeed, as we mentioned in the proof of Theorem 2.3, we may identify the effect of competition as a liability and, in turn, solve an indifference valuation problem. Similarly, for the CRRA case, one may relate the competition to a multiplicative liability factor. There is a fundamental difference, of course, between the classical indifference pricing problems and the ones herein; namely, 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.

Herein, the fund managers are only concerned with maximization of utility coming from terminal wealth (both absolute and relative to other agents), but one could also incorporate intermediate consumption. There are two natural ways to do this, discussed only for the CRRA model for concreteness. First, one might add a utility-of-consumption term to the optimization criterion (33) and modify the individual wealth process to

$$dX_t^i = \pi_t^i X_t^i \left(\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t \right) - C_t^i dt.$$

While the calculations may be tedious, we expect that this problem is tractable. A more interesting approach would build on the first by incorporating *relative consumption standards*, modeled as constraints in the form of lower bounds on the rate of consumption, which could themselves depend on the consumption of other agents. This setting would reflect the more realistic situation in which individual consumption standards are affected by the behavior of other agents.

An important assumption of the model herein is that each agent has full information of the individual preference and market parameters of each other agent. This is also the main modeling ingredient in Basak and Makarov (2015) and is partially defended by the fact that fund managers post their returns publicly, and from these announcements certain information can in turn be inferred by their competitors. While this is undoubtedly a considerable modeling limitation, our results give new solutions to existing problems, especially for MFGs with nonquadratic criteria. Furthermore, this assumption of common knowledge may be relaxed if one introduces ambiguity around the publicly posted competitors' parameters. This ambiguity may be, for example, modeled through an error margin depending on individual views. This would give rise to a class of interesting MFGs with filtering.

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 Basak and Makarov (2015), and in Anthropelos et al. (2017) 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 Anthropelos et al. (2017), and would lead naturally to a new class of MFGs. It would also allow for further extending the concept of benchmarking under forward criteria, introduced in Musiela and Zariphopoulou (2009).

ACKNOWLEDGMENTS

ENDNOTES

[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.

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

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

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

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

REFERENCES

Agarwal, V., Daniel, N. D., & Naik, N. Y. (2003). *Flows, performance, and managerial incentives in hedge funds*, EFA 2003 Annual Conference Paper No. 501.

Anthropelos, M., Geng, T., & Zariphopoulou, T. (2017). Competitive investment strategies under forward performance criteria, in preparation.

Basak, S., & Makarov, D. (2015). *Competition among portfolio managers and asset specialization*, Paris December 2014 Finance Meeting EUROFIDAI-AFFI Paper.

Bensoussan, A., Frehse, J., & Yam, S. C. P. (2015). The master equation in mean field theory. *Journal de Mathématiques Pures et Appliquées*, *103*, 1441–1474.

Bensoussan, A., Sung, K. C. J., Yam, S. C. P., & Yung, S. P. (2016). Linear-quadratic mean field games. *Journal of Optimization Theory and Applications*, *169*, 496–529.

Biagini, S., & Černỳ, A. (2011). Admissible strategies in semimartingale portfolio selection. *SIAM Journal on Control and Optimization*, *49*, 42–72.

Biagini, S., & Sîrbu, M. (2014). A note on admissibility when the credit line is infinite. *Stochastics An International Journal of Probability and Stochastic Processes*, *84*, 157–169.

Bielagk, J., Lionnet, A., & Dos Reis, G. (2017). Equilibrium pricing under relative performance concerns. *SIAM Journal on Financial Mathematics*, *8*, 435–482.

Boyle, P., Garlappi, L., Uppal, R., & Wang, T. (2012). Keynes meets Markowitz: The trade-off between familiarity and diversification. *Management Science*, *58*, 253–272.

Brennan, M. J. (1975). 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*, 483–496.

Brown, S. J., Goetzmann, W. N., & Park, J. (2001). Careers and survival: Competition and risk in the hedge fund and CTA industry. *The Journal of Finance*, *56*, 1869–1886.

Cardaliaguet, P., Delarue, F., Lasry, J.-M., & Lions, P.-L. (2015). The master equation and the convergence problem in mean field games, arXiv preprint arXiv:1509.02505.

Cardaliaguet, P., & Lehalle, C.-A. (2016). Mean field game of controls and an application to trade crowding, arXiv preprint arXiv:1610.09904 (2016).

Carmona, R., & Delarue, F. (2014). The master equation for large population equilibriums. *Stochastic Analysis and Applications*, *2014*, 77–128.

Carmona, R., Delarue, F., & Lachapelle, A. (2013). Control of McKean–Vlasov dynamics versus mean field games. *Mathematics and Financial Economics*, *7*, 131–166.

Carmona, R., Delarue, F., & Lacker, D. (2016). Mean field games with common noise. *The Annals of Probability*, *44*, 3740–3803.

Carmona, R., Fouque, J.-P., & Sun, L.-H. (2015). Mean field games and systemic risk. *Communications in Mathematical Sciences*, *13*, 911–933.

Chan, P., & Sircar, R. (2015). Bertrand and Cournot mean field games. *Applied Mathematics & Optimization*, *71*, 533–569.

Chevalier, J., & Ellison, G. (1997). Risk taking by mutual funds as a response to incentives. *Journal of Political Economy*, *105*, 1167–1200.

Coval, J. D., & Moskowitz, T. J. (1999). Home bias at home: Local equity preference in domestic portfolios. *The Journal of Finance*, *54*, 2045–2073.

Delbaen, F., Grandits, P., Rheinländer, T., Samperi, D., Schweizer, M., & Stricker, C. (2002). Exponential hedging and entropic penalties. *Mathematical Finance*, *12*, 99–123.

Ding, B., Getmansky, M., Liang, B., & Wermers, R. (2008). Investor flows and share restrictions in the hedge fund industry. Available at SSRN: https://ssrn.com/abstract=2692598

Espinosa, G.-E., & Touzi, N. (2015). Optimal investment under relative performance concerns. *Mathematical Finance*, *25*, 221–257.

Fleming, W., & Soner, H. M. (2006). *Controlled Markov processes and viscosity solutions* (Vol. 25). New York: Springer Science & Business Media.

Frei, C., & Dos Reis, G. (2011). A financial market with interacting investors: Does an equilibrium exist? *Mathematics and Financial Economics*, *4*, 161–182.

Gabaix, X., Lasry, J.-M., Lions, P.-L., & Moll, B. (2016). The dynamics of inequality. *Econometrica*, *84*, 2071–2111.

Gallaher, S., Kaniel, R., & Starks, L. T. (2006). Madison Avenue meets Wall Street: Mutual fund families, competition and advertising. Available at SSRN: https://ssrn.com/abstract=879775

Gayduk, R., & Nadtochiy, S. (2016). Endogenous formation of limit order books: dynamics between trades, arXiv preprint arXiv:1605.09720.

Guéant, O., Lasry, J.-M., & Lions, P.-L. (2011). *Mean field games and applications, Paris-Princeton lectures on mathematical finance 2010*. New York: Springer.

Hardy, G. H., Littlewood, J. E., & Pólya, G. (1952). *Inequalities*. Cambridge, UK: Cambridge University Press.

Henderson, V. (2002). Valuation of claims on nontraded assets using utility maximization. *Mathematical Finance*, *12*, 351–373.

Huang, M., Malhamé, R. P., & Caines, P. E. (2006). Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, *6*, 221–252.

Huang, M., & Nguyen, S. L. (2016). Mean field games for stochastic growth with relative utility. *Applied Mathematics & Optimization*, *74*, 643–668.

Huang, X., Jaimungal, S., & Nourian, M. (2017). Mean-field game strategies for a major-minor agent optimal execution problem. Available at SSRN: https://ssrn.com/abstract=2578733

Kacperczyk, M., Sialm, C., & Zheng, L. (2005). On the industry concentration of actively managed equity mutual funds. *The Journal of Finance*, *60*, 1983–2011.

Kempf, A., & Ruenzi, S. (2008). Tournaments in mutual-fund families. *Review of Financial Studies*, *21*, 1013–1036.

Kim, T. S., & Omberg, E. (1996). Dynamic nonmyopic portfolio behavior. *Review of Financial Studies*, *9*, 141–161.

Lacker, D. (2016). A general characterization of the mean field limit for stochastic differential games. *Probability Theory and Related Fields*, *165*, 581–648.

Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, *2*, 229–260.

Li, W., & Tiwari, A. (2006). On the consequences of mutual fund tournaments. Available at SSRN: https://ssrn.com/abstract=687123

Liu, H. (2014). Solvency constraint, underdiversification, and idiosyncratic risks. *Journal of Financial and Quantitative Analysis*, *49*, 409–430.

Lucas Jr, R. E., & Moll, B. (2014). Knowledge growth and the allocation of time. *Journal of Political Economy*, *122*, 1–51.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, *3*, 373–413.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, *42*, 483–510.

Mitton, T., & Vorkink, K. (2007). Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies*, *20*, 1255–1288.

Musiela, M., & Zariphopoulou, T. (2004). An example of indifference prices under exponential preferences. *Finance and Stochastics*, *8*, 229–239.

Musiela, M., & Zariphopoulou, T. (2009). Portfolio choice under dynamic investment performance criteria. *Quantitative Finance*, *9*, 161–170.

Pham, H. (2009). *Continuous-time stochastic control and optimization with financial applications* (Vol. 61). New York: Springer Science & Business Media.

Schachermayer, W. (2001). Optimal investment in incomplete markets when wealth may become negative. *Annals of Applied Probability*, *11*, 694–734.

Sirri, E. R., & Tufano, P. (1998). Costly search and mutual fund flows. *The Journal of Finance*, *53*, 1589–1622.

Sun, L.-H. (2016). Systemic risk and interbank lending, arXiv preprint arXiv:1611.06672.

Sun, Y. (2006). The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, *126*, 31–69.

Touzi, N. (2012). *Optimal stochastic control, stochastic target problems, and backward SDE* (Vol. 29). New York: Springer Science & Business Media.

Uppal, R., & Wang, T. (2003). Model misspecification and underdiversification. *The Journal of Finance*, *58*, 2465–2486.

Van Nieuwerburgh, S., & Veldkamp, L. (2009). Information immobility and the home bias puzzle. *The Journal of Finance*, *64*, 1187–1215.

Van Nieuwerburgh, S., & Veldkamp, L. (2010). Information acquisition and under-diversification. *The Review of Economic Studies*, *77*, no. 2, 779–805.