



Incomplete Stochastic Equilibria with Exponential Utilities Close to Pareto Optimality

Constantinos Kardaras, Hao Xing, and Gordan Žitković

Abstract We study existence and uniqueness of continuous-time stochastic Radner equilibria in an incomplete markets model. An assumption of “smallness” type—imposed through the new notion of “closeness to Pareto optimality”—is shown to be sufficient for existence and uniqueness. Central role in our analysis is played by a fully-coupled nonlinear system of quadratic BSDEs.

Introduction

The equilibrium problem

The focus of the present paper is the problem of existence and uniqueness of a competitive (Radner) equilibrium in an incomplete continuous-time stochastic model of a financial market. A discrete version of our model was introduced by Radner in [26] as an extension of the classical Arrow-Debreu framework, with the goal of understanding how asset prices in financial (or any other) markets are formed, under minimal assumption on the ingredients or the underlying market structure. One of those assumptions is often market completeness; more precisely, it is usually postulated that the range of various types of transactions the markets allow is such that the wealth distribution among agents, after all the trading is done, is Pareto optimal, i.e., that no further redistribution of wealth can make one agent better off

Constantinos Kardaras
Department of Statistics, London School of Economics and Political Science, e-mail: k.kardaras@lse.ac.uk

Hao Xing
Department of Finance, Questrom School of Business, Boston University, e-mail: haoxing@bu.edu

Gordan Žitković
Department of Mathematics, University of Texas at Austin, e-mail: gordanz@math.utexas.edu

without hurting somebody else. Real markets are not complete; in fact, as it turns out, the precise way in which completeness fails matters greatly for the output and should be understood as an a-priori constraint. Indeed, it is instructive to ask the following questions: Why are markets incomplete in the first place? Would rational economic agents not continue introducing new assets into the market, as long as it is still useful? The answer is that they, indeed, would, were it not for exogenously-imposed constraints out there, no markets exist for most contingencies; those markets that do exist are heavily regulated, transactions costs are imposed, short selling is sometimes prohibited, liquidity effects render replication impossible, etc. Instead of delving into the modeling issues regarding various types of *completeness constraints*, we point the reader to [31] where a longer discussion of such issues can be found.

The “fast-and-slow” model

The particular setting we subscribe to here is one of the simplest from the financial point of view. It, nevertheless, exhibits many of the interesting features found in more general incomplete structures and admits a straightforward continuous-time formulation. It corresponds essentially to the so-called “fast-and-slow” completeness constraint, introduced in [31].

One of the ways in which the “fast-and-slow” completeness constraint can be envisioned is by allowing for different speeds at which information of two different kinds is incorporated and processed. The discrete-time version of the model is described in detail in [25, p. 213], where it goes under the heading of “short-lived” asset models. Therein, at each node in the event tree, the agents have access to a number of short-lived assets, i.e., assets whose life-span ends in one unit of time, at which time all the dividends are distributed. The prices of such assets are determined in the equilibrium, but their number is typically not sufficient to guarantee local (and therefore global) completeness of the market. In our, continuous time model, the underlying filtration is generated by two independent Brownian motions (B and W). Positioned the “node” (ω, t) , we think of dB_t and dW_t as two independent symmetric random variables, realized at time $t + dt$, with values $\pm\sqrt{dt}$. Allowing the agents to insure each other only with respect to the risks contained in dB , we denote the (equilibrium) price of such an “asset” by $-\lambda_t dt$. As already hinted to above, one possible economic rationale behind this type of constraint is obtained by thinking of dB as the readily-available (fast) information, while dW models slower information which will be incorporated into the process λ_t indirectly, and only at later dates. For simplicity, we also fix the spot interest rate to 0, allowing agents to transfer wealth from t to $t + dt$ costlessly and profitlessly. While, strictly speaking, this feature puts us in the partial-equilibrium framework, this fact will not play a role in our analysis, chiefly because our agents draw their utility only from the terminal wealth (which is converted to the consumption good at that point).

For mathematical convenience, and to be able to access the available continuous-time results, we concatenate all short-lived assets with payoffs dB_t and prices $-\lambda_t dt$

into a single asset $B_t^\lambda = B_t + \int_0^t \lambda_u du$. It should not be thought of as an asset that carries a dividend at time T , but only as a single-object representation of the family of all infinitesimal, short-lived assets.

As a context for the "fast-and-slow" constraint, we consider a finite number I of agents; we assume that all of their utility functions are of exponential type, but allow for idiosyncratic risk-aversion parameters and non-traded random endowments. The exponential nature of the agents' utilities is absolutely crucial for all of our results as it induces a "backward" structure to our problem, which, while still very difficult to analyze, allows us to make a significant step forward.

The representative-agent approach, and its failure in incomplete markets

The classical and nearly ubiquitous approach to existence of equilibria in complete markets is using the so-called representative-agent approach. Here, the agents' endowments are first aggregated and then split in a Pareto-optimal way. Along the way, a pricing measure is produced, and then, a-posteriori, a market is constructed whose unique martingale measure is precisely that particular pricing measure. As long as no completeness constraints are imposed, this approach works extremely well, pretty much independently of the shape of the agents' utility functions (see, e.g., [14, 13, 18, 19, 20, 9, 1, 30] for a sample of continuous-time literature). A convenient exposition of some of these and many other results, together with a thorough classical literature overview can be found in the Notes section of Chapter 4. of [21]).

The incomplete case requires a completely different approach and what were once minute details, now become salient features. The failure of representative-agent methods under incompleteness are directly related to the inability of the market to achieve Pareto optimality by wealth redistribution. Indeed, when not every transaction can be implemented through the market, one cannot reduce the search for the equilibrium to a finite-dimensional "manifold" of Pareto-optimal allocations. Even more dramatically, the whole nature of what is considered a solution to the equilibrium problem changes. In the complete case, one simply needs to identify a market-clearing valuation measure. In the present "fast-and-slow" formulation, the very family of all replicable claims (in addition to the valuation measure) has to be determined. This significantly impacts the "dimensionality" of the problem and calls for a different toolbox.

Our probabilistic-analytic approach

The direction of the present paper is partially similar to that of [31], where a much simpler model of the "fast-and-slow" type is introduced and considered. Here, however, the setting is different and somewhat closer to [29] and [8]. The fast component

is modeled by an independent Brownian motion, instead of the one-jump process. Also, unlike in any of the above papers, pure PDE techniques are largely replaced or supplemented by probabilistic ones, and much stronger results are obtained.

Doing away with the Markovian assumption, we allow for a collection of unbounded random variables, satisfying suitable integrability assumptions, to act as random endowments and characterize the equilibrium as a (functional of a) solution to a nonlinear system of quadratic Backward Stochastic Differential Equations (BSDE). Unlike single quadratic BSDE, whose theory is by now quite complete (see e.g., [23, 5, 6, 12, 15, 3] for a sample), the systems of quadratic BSDEs are much less understood. The main difficulty is that the comparison theorem may fail to hold for BSDE systems (see [17]). Moreover, Frei and dos Reis (see [16]) constructed a quadratic BSDE system which has bounded terminal condition but admits no solution. The strongest general-purpose result seems to be the one of Tevzadze (see [28]), which guarantees existence under an “ \mathbb{L}^∞ -smallness” condition placed on the terminal conditions.

Like in [28], but unlike in [31] or [8], our general result imposes no regularity conditions on the agents’ random endowments. Unlike [28], however, our smallness conditions come in several different forms. First, we show existence and uniqueness when the random-endowment allocation among agents is close to a Pareto optimal one. In contrast to [28], we allow here for unbounded terminal conditions (random endowments), and measure their size using an “entropic” BMO-type norm strictly weaker than the \mathbb{L}^∞ -norm. In addition, the equilibrium established is unique in a global sense (as in [24], where a different quadratic BSDE system is studied).

Another interesting feature of our general result is that it is largely independent of the number of agents. This leads to the following observation: the equilibrium exists as soon as “sufficiently many sufficiently homogeneous” (under an appropriate notion of homogeneity) agents share a given total endowment, which is not assumed to be small. This is precisely the natural context of a number of competitive equilibrium models with a large number of small agents, none of whom has a dominating sway over the price.

Another parameter our general result is independent of is the time horizon T . Indirectly, this leads to our second existence and uniqueness result which holds when the time horizon is sufficiently small, but the random endowments are not limited in size. Under the additional assumption of Malliavin differentiability, a lower bound on how small the horizon has to be to guarantee existence and uniqueness turns out to be inversely proportional to the size of the (Malliavin) derivatives of random endowments. This extends [8, Theorem 3.1] to a non-Markovian setting. Interestingly, both the \mathbb{L}^∞ -smallness of the random endowments and the smallness of the time-horizon are implied by the small-entropic-BMO-norm condition mentioned above, and the existence theorems under these conditions can be seen as special cases of our general result.

Some notational conventions

As we will be dealing with various classes of vector-valued random variables and stochastic processes, we try to introduce sufficiently compact notation to make reading more palatable.

A time horizon $T > 0$ is fixed throughout. An equality sign between random variables signals almost-sure equality, while one between two processes signifies Lebesgue-almost everywhere, almost sure equality; any two processes that are equal in this sense will be identified; this, in particular, applied to indistinguishable càdlàg processes. Given a filtered probability space $(\Omega, \mathcal{F}_T, \mathbb{F} = \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$ satisfying the usual conditions, \mathcal{T} denotes the set of all $[0, T]$ -valued \mathbb{F} -stopping times, and \mathcal{P}^2 denotes the set of all predictable processes $\{\mu_t\}_{t \in [0, T]}$ such that $\int_0^T \mu_t^2 dt < \infty$, a.s. The integral $\int_0^\cdot \mu_u dB_u$ of $\mu \in \mathcal{P}^2$ with respect to an \mathbb{F} -Brownian motion B is alternatively denoted by $\mu \cdot B$, while the stochastic (Doléans-Dade) exponential retains the standard notation $\mathcal{E}(\cdot)$. The \mathbb{L}^p -spaces, $p \in [1, \infty]$ are all defined with respect to $(\Omega, \mathcal{F}_T, \mathbb{P})$ and \mathbb{L}^0 denotes the set of (\mathbb{P} -equivalence classes) of finite-valued random variables on this space. For a continuous adapted process $\{Y_t\}_{t \in [0, T]}$, we set

$$\|Y\|_{\mathcal{S}^\infty} = \|\sup_{t \in [0, T]} |Y_t|\|_{\mathbb{L}^\infty},$$

and denote the space of all such Y with $\|Y\|_{\mathcal{S}^\infty} < \infty$ by \mathcal{S}^∞ . For $p \geq 1$, the space of all $\mu \in \mathcal{P}^2$ with $\|\mu\|_{H^p}^p = \mathbb{E} \left[\int_0^T |\mu_u|^p du \right] < \infty$ is denoted by H^p , an alias for the Lebesgue space \mathbb{L}^p on the product $[0, T] \times \Omega$.

Given a probability measure $\hat{\mathbb{P}}$ and a $\hat{\mathbb{P}}$ -martingale M , we define its BMO-norm by

$$\|M\|_{\text{BMO}(\hat{\mathbb{P}})}^2 = \sup_{\tau \in \mathcal{T}} \left\| \mathbb{E}_\tau^{\hat{\mathbb{P}}} [\langle M \rangle_T - \langle M \rangle_\tau] \right\|_{\mathbb{L}^\infty},$$

where $\mathbb{E}_\tau^{\hat{\mathbb{P}}}[\cdot]$ denotes the conditional expectation $\mathbb{E}^{\hat{\mathbb{P}}}[\cdot | \mathcal{F}_\tau]$ with respect to \mathcal{F}_τ , computed under $\hat{\mathbb{P}}$. The set of all $\hat{\mathbb{P}}$ -martingales M with finite $\|M\|_{\text{BMO}(\hat{\mathbb{P}})}$ is denoted by $\text{BMO}(\hat{\mathbb{P}})$, or, simply, BMO, when $\hat{\mathbb{P}} = \mathbb{P}$. When applied to random variables, $X \in \text{BMO}(\hat{\mathbb{P}})$ means that $X = M_T$, for some $M \in \text{BMO}(\hat{\mathbb{P}})$. In the same vein, we define (for some, and then any, $(\hat{\mathbb{P}}, \mathbb{F})$ -Brownian motion B)

$$\text{bmo}(\hat{\mathbb{P}}) = \{\mu \in \mathcal{P}^2 : \mu \cdot B \in \text{BMO}(\hat{\mathbb{P}})\},$$

with the norm $\|\mu\|_{\text{bmo}(\hat{\mathbb{P}})} = \|\mu \cdot B\|_{\text{BMO}(\hat{\mathbb{P}})}$. The same convention as above is used: the dependence on $\hat{\mathbb{P}}$ is suppressed when $\hat{\mathbb{P}} = \mathbb{P}$.

Many of our objects will take values in \mathbb{R}^I , for some fixed $I \in \mathbb{N}$. Those are typically denoted by bold letters such as $\mathbf{E}, \mathbf{G}, \boldsymbol{\mu}, \boldsymbol{\nu}, \boldsymbol{\alpha}$, etc. If specific components are needed, they will be given a superscript - e.g., $\mathbf{G} = (G^i)_i$. Unquantified variables i, j always range over $\{1, 2, \dots, I\}$. The topology of \mathbb{R}^k is induced by the Euclidean norm $|\cdot|_2$, defined by $|\mathbf{x}|_2 = \sqrt{\sum_k |x^k|^2}$ for $\mathbf{x} \in \mathbb{R}^k$. All standard operations and

relations (including the absolute value $|\cdot|$ and order \leq) between \mathbb{R}^k -valued variables are considered componentwise.

1 The Equilibrium Problem and its BSDE Reformulation

We work on a filtered probability space $(\Omega, \mathcal{F}_T, \mathbb{F} = \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$, where \mathbb{F} is the standard augmentation of the filtration generated by a two-dimensional standard Brownian motion $\{(B_t, W_t)\}_{t \in [0, T]}$. The augmented natural filtrations \mathbb{F}^B and \mathbb{F}^W of the two Brownian motions B and W will also be considered below.

1.1 The financial market, its agents, and equilibria

Our model of a financial market features one liquidly traded **risky asset**, whose value, denoted in terms of a prespecified numéraire which we normalize to 1, is given by

$$dB_t^\lambda = \lambda_t dt + dB_t, \quad t \in [0, T], \quad (1)$$

for some $\lambda \in \mathcal{P}^2$. Given that it will play a role of a “free parameter” in our analysis, the volatility in (1) is normalized to 1; this way, λ can simultaneously be interpreted as the **market price of risk**. The reader should consult the section ‘The “fast-and-slow” model’ in the introduction for the proper economic interpretation of this asset as a concatenation of a continuum of infinitesimally-short-lived securities.

We assume there is a finite number $I \in \mathbb{N}$ of **economic agents**, all of whom trade the risky asset as well as the aforementioned riskless, numéraire, asset of constant value 1. The preference structure of each agent is modeled in the von Neumann-Morgenstern framework via the following two elements:

- i) an exponential **utility function** with **risk tolerance coefficient** $\delta^i > 0$:

$$U^i(x) = -\exp(-x/\delta^i), \quad x \in \mathbb{R}, \text{ and}$$

- ii) a **random endowment** $E^i \in \mathbb{L}^0(\mathcal{F}_T)$.

The pair (E, δ) , where $E = (E^i)_i$, $\delta = (\delta^i)_i$, of endowments and risk-tolerance coefficients fully characterizes the behavior of the agents in the model; we call it the **population characteristics**— E is the **initial allocation** and δ the **risk profile**. In general, any \mathbb{R}^I -valued random vector will be referred to as an **allocation**.

Each agent maximizes the expected utility of trading and random endowment:

$$\mathbb{E} \left[U^i(\pi \cdot B_T^\lambda + E^i) \right] \rightarrow \max. \quad (2)$$

Here $\{\pi_t\}_{t \in [0, T]}$ is a one-dimensional process which represents the number of shares of the asset kept by the agent at time t . As usual, this strategy is financed by investing

in or borrowing from the interestless numéraire asset, as needed. To describe the admissible strategies of the agent, we follow the convention in [11]:

For $\lambda \in \mathcal{P}^2$, we denote by \mathcal{M}_a^λ the set of absolutely continuous local martingale measures for B^λ , i.e., all probability measures $\mathbb{Q} \ll \mathbb{P}$ such that $\mathbb{E}^\mathbb{Q}[h(B_\tau^\lambda - B_\sigma^\lambda)] = 0$ for all pairs of stopping times $\sigma \leq \tau \leq T$ and for all bounded \mathcal{F}_σ -measurable random variables h . For a probability measure $\mathbb{Q} \ll \mathbb{P}$, let $H(\mathbb{Q}|\mathbb{P})$ be the relative entropy of \mathbb{Q} with respect to \mathbb{P} , i.e., $H(\mathbb{Q}|\mathbb{P}) = \mathbb{E} \left[\frac{d\mathbb{Q}}{d\mathbb{P}} \log \frac{d\mathbb{Q}}{d\mathbb{P}} \right] \geq 0$. For $\lambda \in \mathcal{P}^2$ such that $\mathcal{M}^\lambda \neq \emptyset$, where

$$\mathcal{M}^\lambda = \{ \mathbb{Q} \in \mathcal{M}_a^\lambda \mid H(\mathbb{Q}|\mathbb{P}) < \infty \},$$

a strategy π is said to be λ -**admissible** if $\pi \in \mathcal{A}^\lambda$, where

$$\mathcal{A}^\lambda = \{ \pi \in \mathcal{P}^2 \mid \pi \cdot B^\lambda \text{ is a } \mathbb{Q}\text{-martingale for all } \mathbb{Q} \in \mathcal{M}^\lambda \}.$$

We note that the set \mathcal{A}^λ corresponds - up to finiteness of the utility - exactly to the set Θ_2 in [11]. This admissible class contains, in particular, all $\pi \in \mathcal{P}^2$ such that $\pi \cdot B^\lambda$ is bounded (uniformly in t and ω).

Definition 1 (Equilibrium)

Given a population with characteristics (E, δ) , a process $\lambda \in \mathcal{P}^2$ with $\mathcal{M}^\lambda \neq \emptyset$ is called an **equilibrium (market price of risk)** if there exists an I -tuple $(\pi^i)_i$ such that

- i) each π^i is an *optimal strategy* for the agent i under λ , i.e.

$$\pi^i \in \operatorname{argmax}_{\pi \in \mathcal{A}^\lambda} \mathbb{E} \left[U^i(\pi \cdot B_T^\lambda + E^i) \right],$$

- ii) the *market clears*, i.e., $\sum_i \pi^i = 0$.

The set of all equilibria is denoted by $\Lambda_\delta(E, \mathbb{P})$, or simply, $\Lambda_\delta(E)$, when the the probability \mathbb{P} is clear from the context.

Remark 1 The assumptions on the agents’ random endowments that we introduce below and the proof techniques we employ make it clear that bmo is a natural space to search for equilibria in. There is, however, no compelling economic argument to include bmo into the *definition* of an equilibrium, so we do not. It turns out, nevertheless, that whenever an equilibrium λ is mentioned in the rest of the paper it will be in the bmo context, and we will assume automatically that any equilibrium market price of risk belongs to bmo. In particular, all uniqueness statements we make will be with respect to bmo as the ambient space.

1.2 A simple risk-aware reparametrization

It turns out that a simple reparametrization in our “ingredient space” leads to substantial notational simplification. It also sheds some light on the economic meaning

of various objects. The main idea is to think of the risk-tolerance coefficients as numéraires, as they naturally carry the same currency units as wealth. When expressed in risk-tolerance units, the random endowments and strategies become unitless and we introduce the following notation

$$G = \frac{1}{\delta} E, \text{ i.e., } G^i = \frac{1}{\delta^i} E^i, \quad \text{and} \quad \rho = \frac{1}{\delta} \pi, \text{ i.e., } \rho^i = \frac{1}{\delta^i} \pi^i. \quad (3)$$

Since \mathcal{A}^λ is invariant under this reparametrization, the equilibrium conditions become

$$\rho^i \in \operatorname{argmax}_{\rho \in \mathcal{A}_i^\lambda} \mathbb{E} \left[U(\rho \cdot B_T^\lambda + G^i) \right] \quad \text{and} \quad \sum_i \alpha^i \rho^i = 0, \quad (4)$$

where $U(x) = -\exp(-x)$, and $\alpha^i = \delta^i / (\sum_j \delta^j) \in (0, 1)$ - with $\sum_i \alpha^i = 1$ - are the **(relative) weights** of the agents. The set of all equilibria with risk-denominated random endowments $G = (G^i)_i$ and relative weights $\alpha = (\alpha^i)_i$ is denoted by $\Lambda_\alpha(G, \mathbb{P})$ (this notation overload should not cause any confusion in the sequel).

Since the market-clearing condition in (4) now involves the relative weights α^i as “conversion rates”, it is useful to introduce the **aggregation operator** $A : \mathbb{R}^I \rightarrow \mathbb{R}$ by

$$A[x] = \sum_i \alpha^i x^i, \quad \text{for } x \in \mathbb{R}^I, \quad (5)$$

so that the market-clearing condition now simply reads $A[\rho] = 0$, pointwise.

1.3 A solution of the single-agent utility-maximization problem

Before we focus on the questions of existence and uniqueness of an equilibrium, we start with the single agent’s optimization problem. Here we suppress the index i and first introduce an assumptions on the risk-denominated random endowment:

$$G \text{ is bounded from above and } G \in \text{EBMO}, \quad (6)$$

where EBMO denotes the set of all $G \in \mathbb{L}^0$ for which there exists (necessarily unique) processes m^G and n^G in bmo, as well a constant X_0^G , such that $G = X_T^G$, where

$$X_t^G = X_0^G + \int_0^t m_u^G dB_u + \int_0^t n_u^G dW_u + \frac{1}{2} \int_0^t ((m_u^G)^2 + (n_u^G)^2) du. \quad (7)$$

The supermartingale X^G admits the following representation

$$X_t^G = -\log \mathbb{E}_t[\exp(-G)], \text{ so that } U(X_t^G) = \mathbb{E}_t[U(G)] \text{ for } t \in [0, T], \quad (8)$$

and can be interpreted as the certainty-equivalent process (without access to the market) of G , expressed in the units of risk tolerance.

Remark 2

1. When G is bounded from above, as we require it to be in (6), a sufficient condition for $G \in \text{EBMO}$ is $e^{-G} \in \text{BMO}$. This follows directly from the boundedness of the (exponential) martingale $e^{-X_t^G}$ away from zero.
2. The condition (6) amounts to the membership $M^G \in \text{BMO}$, where $M^G = m^G \cdot B + n^G \cdot W$. Then $-M^G \in \text{BMO}$ and, by Theorem 3.1, p. 54 in [22], $\mathcal{E}(-M^G)$ satisfies the reverse Hölder inequality (R_p) with some $p > 1$. Therefore, for $\varepsilon < p - 1$, we have

$$\begin{aligned} \mathbb{E}[e^{-(1+\varepsilon)G}] &= \mathbb{E}[e^{-(1+\varepsilon)(X_0^G + M_T^G + \frac{1}{2}\langle M^G \rangle_T)}] \\ &= e^{-(1+\varepsilon)X_0^G} \mathbb{E}\left[\left(\mathcal{E}(-M^G)_T\right)^{1+\varepsilon}\right] < \infty. \end{aligned}$$

On the other hand, by (1) above, we clearly have $\mathbb{L}^\infty \subseteq \text{EBMO}$, so

$$G \in \mathbb{L}^\infty \Rightarrow G \in \text{EBMO} \Rightarrow \mathbb{E}[e^{-(1+\varepsilon)G}] < \infty \text{ for some } \varepsilon > 0.$$

In particular our condition (6), while implied by the boundedness of G , itself implies the conditions $G^+ = \max\{G, 0\} \in \mathbb{L}^\infty$, $e^{-G} \in \cup_{p>1} \mathbb{L}^p$, imposed in [11].

We recall in Proposition 1 some results about the nature of the optimal solution to the utility-maximization problem (2) from [11]; the proof is given in Section 3 below.

Proposition 1 (Single agent’s optimization problem: existence and duality)

Suppose that $\lambda \in \text{bmo}$ and that G satisfies (6). Then both primal and dual problems have finite values and the following statements hold:

1. *There exists a unique $\rho^{\lambda,G} \in \mathcal{A}^\lambda$ such that*

$$\rho^{\lambda,G} \in \operatorname{argmax}_{\rho \in \mathcal{A}^\lambda} \mathbb{E}\left[U(\rho \cdot B_T^\lambda + G)\right].$$

2. *There exists a unique $\mathbb{Q}^{\lambda,G} \in \mathcal{M}^\lambda$ such that*

$$\mathbb{Q}^{\lambda,G} \in \operatorname{argmin}_{\mathbb{Q} \in \mathcal{M}^\lambda} (H(\mathbb{Q}|\mathbb{P}) + \mathbb{E}^\mathbb{Q}[G]).$$

3. *There exists a constant $c^{\lambda,G}$ such that*

$$c^{\lambda,G} + \rho^{\lambda,G} \cdot B_T^\lambda + G = -\log(Z_T^{\lambda,G}), \text{ where } Z_T^{\lambda,G} = \frac{d\mathbb{Q}^{\lambda,G}}{d\mathbb{P}}. \tag{9}$$

The process $\rho^{\lambda,G}$ and the probability measure $\mathbb{Q}^{\lambda,G}$ are called the **primal** and the **dual optimizers**, respectively. While they were first obtained by convex-duality methods, they also admit a BSDE representation (see, e.g., [27]), where a major role is played by (the risk-denominated version) of the so-called **certainty-equivalent process**:

$$Y_t^{\lambda,G} = U^{-1}\left(\mathbb{E}_t\left[U(\rho^{\lambda,G} \cdot B_T^\lambda - \rho_t^{\lambda,G} \cdot B_t^\lambda + G)\right]\right), \quad t \in [0, T]. \tag{10}$$

The optimality of $\rho^{\lambda, G}$ implies that

$$U(Y_t^{\lambda, G}) = \operatorname{esssup}_{\rho \in \mathcal{A}^I} \mathbb{E}_t \left[U(\rho \cdot B_T^\lambda - \rho \cdot B_t^\lambda + G) \right], \quad t \in [0, T]. \quad (11)$$

Hence $Y_t^{\lambda, G}$ can be interpreted as the risk-denominated certainty equivalent of the agent i , when he/she trades optimally from t onwards, starting from no wealth. Finally, with

$$Z_t^{\lambda, G} = \mathbb{E}_t \left[\frac{dQ_t^{\lambda, G}}{dP} \right] = \mathcal{E}(-\lambda \cdot B - v^{\lambda, G} \cdot W)_t, \quad t \in [0, T] \text{ for some } v^{\lambda, G} \in \mathcal{P}^2, \quad (12)$$

we have the following BSDE characterization for single agent's optimization problem.

Lemma 1 (Single agent's optimization problem: a BSDE characterization)

For $\lambda \in bmo$ and G satisfying (6), let $Y^{\lambda, G}$ be as in (10), let $\mu^{\lambda, G} = \lambda - \rho^{\lambda, G}$ and let $v^{\lambda, G}$ be defined by (12). Then the triplet $(Y^{\lambda, G}, \mu^{\lambda, G}, v^{\lambda, G})$ is the unique solution to the BSDE

$$dY_t = \mu_t dB_t + v_t dW_t + \left(\frac{1}{2} v_t^2 - \frac{1}{2} \lambda_t^2 + \lambda_t \mu_t \right) dt, \quad Y_T = G, \quad (13)$$

in the class where $(\mu, v) \in bmo$. Such a unique solution also satisfies $Y^{\lambda, G} - X^G \in \mathcal{S}^\infty$.

Given the results of Propositions 1 and 1 above, we fix the notation $Y^{\lambda, G}$, $\mu^{\lambda, G}$, $v^{\lambda, G}$, $Q^{\lambda, G}$, $Z^{\lambda, G}$ and $\rho^{\lambda, G}$ for λ and G . We also introduce the vectorized versions $\mathbf{Y}^{\lambda, G}$, $\boldsymbol{\mu}^{\lambda, G}$, $\mathbf{v}^{\lambda, G}$, $\mathbf{Q}^{\lambda, G}$, and $\mathbf{Z}^{\lambda, G}$, so that, e.g., $\boldsymbol{\mu}^{\lambda, G} = (\mu^{\lambda, G^i})_i$ and $\mathbf{G} = (G^i)_i$.

1.4 A BSDE characterization of equilibria

The BSDE-based description in Lemma 1 of the solution of a single agent's optimization problem is the main ingredient in the following characterization, whose proof is given in Subsection 3.3 below. We use the risk-aware parametrization introduced in Subsection 1.2, and remind the reader that $\Lambda_\alpha(\mathbf{G})$ denotes the set of all equilibria in bmo when $\mathbf{G} = (G^i)_i$ are the agents' risk-denominated random endowments and $\alpha = (\alpha^i)_i$ are the relative weights.

Theorem 1 (BSDE characterization of equilibria)

For a process $\lambda \in bmo$, and an allocation \mathbf{G} which satisfies (6) componentwise, the following are equivalent:

1. $\lambda \in \Lambda_\alpha(\mathbf{G})$, i.e., λ is an equilibrium for the population (\mathbf{G}, α) .
2. $\lambda = A[\boldsymbol{\mu}]$ for some solution $(\mathbf{Y}, \boldsymbol{\mu}, \mathbf{v})$ of the BSDE system:

$$dY_t = \boldsymbol{\mu}_t dB_t + \mathbf{v}_t dW_t + \left(\frac{1}{2} \mathbf{v}_t^2 - \frac{1}{2} A[\boldsymbol{\mu}_t]^2 + A[\boldsymbol{\mu}_t] \boldsymbol{\mu}_t \right) dt, \quad Y_T = \mathbf{G}, \quad (14)$$

with $(\boldsymbol{\mu}, \mathbf{v}) \in bmo^I$.

Remark 3

1. Spelled out “in coordinates”, the system (14) becomes

$$\begin{cases} dY_t^i = \mu_t^i dB_t + \nu_t^i dW_t + \left(\frac{1}{2}(\nu_t^i)^2 - \frac{1}{2}(\sum_j \alpha^j \mu_t^j)^2 + (\sum_j \alpha^j \mu_t^j) \mu_t^i \right) dt, \\ Y_T^i = G^i, \quad i \in \{1, 2, \dots, I\}, \end{cases} \tag{14}$$

and the market-clearing condition $\lambda = A[\mu_t]$ reads $\lambda = \sum_j \alpha^j \mu^j$.

2. While quite meaningless from the competitive point of view, in the case $I = 1$ of the above characterization still admits a meaningful interpretation. The notion of an equilibrium here corresponds to the choice of λ under which an agent, with risk-denominated random endowment $G \in \text{EBMO}$ would choose not to invest in the market at all. The system (14) reduces to a single equation

$$dY_t = \mu_t dB_t + \nu_t dW_t + \left(\frac{1}{2} \mu_t^2 + \frac{1}{2} \nu_t^2 \right) dt, \quad Y_T = G,$$

which admits a unique solution, namely $Y = X^G$, so that $\lambda = m^G$ is the unique equilibrium. This case also singles out the space EBMO as the natural environment for the random endowments G^i in this context.

2 Main Results

We first present our main result, then discuss its implications on models with short time horizons or a large population of agents. All proofs are postponed until Section 3.

2.1 Equilibria close to Pareto optimality

Whenever equilibrium is discussed, Pareto optimality is a key concept. Passing to the more-convenient risk-aware notation, we remind the reader the following definition, where, as usual, $A[x] = \sum_i \alpha^i x^i$:

Definition 2 For $\xi \in \mathbb{L}^0(\mathcal{F}_T)$, an allocation ξ is called ξ -feasible if $A[\xi] \leq \xi$. An allocation ξ is said to be **Pareto optimal** if there is no $A[\xi]$ -feasible allocation $\tilde{\xi}$, such that $\mathbb{E}[U(\tilde{\xi}^i)] \geq \mathbb{E}[U(\xi^i)]$ for all i , and $\mathbb{E}[U(\tilde{\xi}^i)] > \mathbb{E}[U(\xi^i)]$ for some i .

In our setting, Pareto optimal allocations admit a very simple characterization; this is a direct consequence of the classical result [4] of Borch so we omit the proof.

Lemma 2 *A (sufficiently integrable) allocation ξ is Pareto optimal if and only if its components agree up to a constant, i.e., if there exist $\xi^c \in \mathbb{L}^0(\mathcal{F}_T)$ and constants $(c^i)_i$ such that $\xi^i = \xi^c + c^i$ for all i .*

Next, we introduce a concept which plays a central role in our main result. Given a population with the (risk-denominated) initial allocation \mathbf{G} whose components satisfy (6), let $(m^i, n^i) \in \text{bmo}$ be an alias for the pair (m^{G^i}, n^{G^i}) defined in (7). We define **distance to Pareto optimality** $H(\mathbf{G})$ of \mathbf{G} by

$$H(\mathbf{G}) = \inf_{\xi^c} \max_i \|(m^i - m^c, n^i - n^c)\|_{\text{bmo}(\mathbb{P}^c)},$$

where the infimum is taken over the set of $\xi^c \in \text{EBMO}$, with $(m^c, n^c) = (m^{\xi^c}, n^{\xi^c})$ as in (7), and the probability measure \mathbb{P}^c is given by

$$d\mathbb{P}^c / d\mathbb{P} = \mathcal{E}(-m^c \cdot B - n^c \cdot W)_T = \exp(-\xi^c) / \mathbb{E}[\exp(-\xi^c)]. \tag{15}$$

Remark 4

1. Suppose that $H(\mathbf{G}) = 0$ and that the infimum is attained. Then $(m^i, n^i) = (m^c, n^c)$, for all i , implying that all components of \mathbf{G} coincide with ξ^c up to some additive constants, making \mathbf{G} Pareto optimal. On the other hand, since each agent has exponential utility, shifting all components of \mathbf{G} by the same amount ξ^c is equivalent to a measure change from \mathbb{P} to \mathbb{P}^c . Therefore, $\lambda \in \Lambda_\alpha(\mathbf{G}, \mathbb{P})$ if and only if $\lambda - m^c \in \Lambda_\alpha(\mathbf{G} - \xi^c, \mathbb{P}^c)$, i.e., translation in endowments does not affect the wellposedness of the equilibrium. As a consequence, to show $\Lambda_\alpha(\mathbf{G}, \mathbb{P}) \neq \emptyset$, it suffices to prove $\Lambda_\alpha(\mathbf{G} - \xi^c, \mathbb{P}^c) \neq \emptyset$ for some ξ^c , which is the strategy we follow below.
2. Our “distance to Pareto optimality” is conceptually similar to the “coefficient of resource utilization” of Debreu (see [10]), well known in economics. There, however, seems to be no simple and direct mathematical connection between the two.

In our first main result below, we assume that \mathbf{G} is sufficiently close to *some* Pareto optimal allocation, i.e., that $H(\mathbf{G}) \leq \epsilon^*$, for some sufficiently small ϵ^* :

Theorem 2 (Existence and uniqueness close to Pareto optimality)

Let (6) hold for all components in \mathbf{G} . There exists a sufficiently small constant ϵ^ , independent of the number of agents I , such that if*

$$H(\mathbf{G}) \leq \epsilon^*, \tag{16}$$

Then there exists a unique equilibrium $\lambda \in \text{bmo}$. Moreover, the triplet $(\mathbf{Y}^{\lambda, \mathbf{G}}, \boldsymbol{\mu}^{\lambda, \mathbf{G}}, \boldsymbol{\nu}^{\lambda, \mathbf{G}})$, defined in Lemma 1, is the unique solution to (14) with $(\boldsymbol{\mu}^{\lambda, \mathbf{G}}, \boldsymbol{\nu}^{\lambda, \mathbf{G}}) \in \text{bmo}^I$.

Remark 5 A similar global uniqueness has been obtained in [24, Theorem 4.1] for a different quadratic BSDE system arising from a price impact model.

The proof of Theorem 2 will be presented in Section 2.1. For the time being, let us discuss two important cases in which (16) holds:

- First, given $\xi^c \in \text{EBMO}$ and $1 \leq i \leq I$, let X^{G^i} and X^{ξ^c} be defined by (7) with terminal conditions G^i and ξ^c , respectively. A simple calculation shows that

$$d(X_t^{G^i} - X_t^{\xi^c}) = (m_t^i - m_t^c) dB_t^c + (n_t^i - n_t^c) dW_t^c + \frac{1}{2} \left((m_t^i - m_t^c)^2 + (n_t^i - n_t^c)^2 \right) dt,$$

with the terminal condition $G^i - \xi^c$, for a two-dimensional \mathbb{P}^c -Brownian motion (B^c, W^c) , where \mathbb{P}^c is given by (15). If, furthermore, $G^i - \xi^c \in \mathbb{L}^\infty$, it follows that

$$\begin{aligned} \|(m^i - m^c, n^i - n^c)\|_{\text{bmo}(\mathbb{P}^c)}^2 &= 2 \sup_\tau \|\mathbb{E}_\tau^{\mathbb{P}^c} [X_T^{G^i} - \xi^c] - (X_\tau^{G^i} - \xi^c)\|_{\mathbb{L}^\infty} \\ &\leq 4 \|G^i - \xi^c\|_{\mathbb{L}^\infty}. \end{aligned}$$

Therefore, assumption (16) holds, if

$$\inf_{\xi^c} \max_i \|G^i - \xi^c\|_{\mathbb{L}^\infty} \leq \frac{(\epsilon^*)^2}{4}. \tag{17}$$

- The second case in which (16) can be verified is in the case of a "large" number of agents. Indeed, an interesting feature of (17) is its lack of dependence on I , leading to the existence of equilibria in an economically meaningful asymptotic regime. Given a **total endowment** $E_\Sigma \in \mathbb{L}^\infty$ to be shared among I agents, i.e., $\sum_i E^i = E_\Sigma$, one can ask the following question: how many and what kind of agents need to share this total endowment so that they can form a financial market in which an equilibrium exists? The answer turns out to be “sufficiently many sufficiently homogeneous agents”. In order show that, we first make precise what we mean by sufficiently homogeneous. For the population characteristics $\mathbf{E} = (E^i)_i$ and $\boldsymbol{\delta} = (\delta^i)_i$, with $\mathbf{E} \in (\mathbb{L}^\infty)^I$, we define the **endowment heterogeneity index** $\chi^E(\mathbf{E}) \in [0, 1]$ by

$$\chi^E(\mathbf{E}) = \max_{i,j} \frac{\|E^i - E^j\|_{\mathbb{L}^\infty}}{\|E^i\|_{\mathbb{L}^\infty} + \|E^j\|_{\mathbb{L}^\infty}}.$$

We think of a population of agents as “sufficiently homogeneous” if $\chi^E(\mathbf{E}) \leq \chi_0^E$ for some, given, critical index χ_0^E . With this in mind, we have the following corollary of Theorem 2:

Corollary 1 (Existence of equilibria for sufficiently many sufficiently homogeneous agents)

Given a critical endowment homogeneity index $\chi_0^E \in [0, \frac{1}{2})$, a critical risk tolerance $\delta_0 > 0$, as well as the total endowment $E_\Sigma \in \mathbb{L}^\infty$, there exists $I_0 = I_0(\|E_\Sigma\|_{\mathbb{L}^\infty}, \chi_0^E, \delta_0) \in \mathbb{N}$, so that any population $(\mathbf{E}, \boldsymbol{\delta}) = (E^i, \delta^i)_i$ satisfying

$$I \geq I_0, \quad \sum_i E^i = E_\Sigma, \quad \chi^E(\mathbf{E}) \leq \chi_0^E, \quad \text{and} \quad \min_i \delta^i \geq \delta_0,$$

admits an equilibrium.

Condition (17) can be thought of as a smallness-in-size assumption placed on the random endowments, possibly after translation. It turns out that it can be “traded” for a smallness-in-time condition which we now describe. We start by briefly recalling the notion of Malliavin differentiation on the Wiener space. Let Φ be the set of random variables ζ of the form $\zeta = \varphi(I(h^1), \dots, I(h^k))$, where $\varphi \in C_b^\infty(\mathbb{R}^k, \mathbb{R})$ (smooth

functions with bounded derivatives of all orders) for some k , $h^j = (h^{j,b}, h^{j,w}) \in \mathbb{L}^2([0, T]; \mathbb{R}^2)$ and $\mathcal{I}(h^j) = h^{j,b} \cdot B_T + h^{j,w} \cdot W_T$, for each $j = 1, \dots, k$. If $\zeta \in \Phi$, we define its **Malliavin derivative** as the 2-dimensional process

$$D_\theta \zeta = \sum_{j=1}^k \frac{\partial \varphi}{\partial x_j}(\mathcal{I}(h^1), \dots, \mathcal{I}(h^k)) h_\theta^j, \quad \theta \in [0, T].$$

We denote by $D_\theta^b \zeta$ and $D_\theta^w \zeta$ the two components of $D_\theta \zeta$ and for $\zeta \in \Phi$, $p \geq 1$, define the norm

$$\|\zeta\|_{1,p} = \left[\mathbb{E} \left[|\zeta|^p + \left(\int_0^T |D_\theta \zeta|^2 d\theta \right)^{p/2} \right] \right]^{1/p}.$$

For $p \in [1, \infty)$, the Banach space $\mathbb{D}^{1,p}$ is the closure of Φ under $\|\cdot\|_{1,p}$. For $p = \infty$, we define $\mathbb{D}^{1,\infty}$ as the set of all those $G \in \mathbb{D}^{1,1}$ with $D^b G, D^w G \in \mathcal{S}^\infty$.

Corollary 2 (Existence of equilibria on sufficiently small time horizons)

Suppose that (6) holds for all components of \mathbf{G} and that there exists $\xi^c \in \text{EBMO}$ such that $G^i - \xi^c \in \mathbb{D}^{1,\infty}$ for all i . Then a unique equilibrium exists as soon as

$$T < T^* = \frac{(\epsilon^*)^2}{\max_i \left(\|D^b(G^i - \xi^c)\|_{\mathcal{S}^\infty}^2 + \|D^w(G^i - \xi^c)\|_{\mathcal{S}^\infty}^2 \right)}. \tag{18}$$

Remark 6 In a Markovian setting where $\mathbf{G} = \mathbf{g}(B_T, W_T)$, for some functions $\mathbf{g} = (g^i)_i$, we only need to assume there exists some $g^c \in \mathbb{L}^\infty$ such that $\partial_b(g^i - g^c), \partial_w(g^i - g^c) \in \mathbb{L}^\infty$, for any i , where $\partial_b(g^i - g^c)$ and $\partial_w(g^i - g^c)$ are weak derivatives of $g^i - g^c$. A similar “smallness in time” result has been proven in [8, Theorem 3.1] (and in [30] in a simpler model) in a Markovian setting. Corollary 2 extends the result of [8] to a non-Markovian setting.

3 Proofs

3.1 Proof of Proposition 1

For $\lambda \in \text{bmo}$, we record that $\mathcal{M}^\lambda \neq \emptyset$. Indeed, thanks to the bmo property of λ , the process $Z^\lambda = \mathcal{E}(-\lambda \cdot B)$ is a martingale and satisfies the reverse Hölder inequality R_p for some $p > 1$ (see [22, Theorem 3.1]). That, in turn, implies the reverse Hölder inequality $R \log R$, and, so, the probability \mathbb{Q}^λ defined via $d\mathbb{Q}^\lambda/d\mathbb{P} = Z_T^\lambda$ satisfies $H(\mathbb{Q}^\lambda|\mathbb{P}) < \infty$, and, consequently $\mathbb{Q}^\lambda \in \mathcal{M}^\lambda$.

The statements of Proposition 1 will follow from [11, Theorem 2.2], once we verify that Z^λ satisfies the reverse Hölder inequality $R \log R$ under $\overline{\mathbb{P}}$ as well, where $d\overline{\mathbb{P}}/d\mathbb{P} = e^{-G}/\mathbb{E}[e^{-G}]$. For that, we note that $e^{-G}/\mathbb{E}[e^{-G}] = \mathcal{E}(-m^G \cdot B - n^G \cdot W)_T$, where (m^G, n^G) is as in (7). Given $\lambda \in \text{bmo}$, the bmo property of (m^G, n^G) and [22,

Theorem 3.6] imply that $\lambda - m^G \in \text{bmo}(\bar{\mathbb{P}})$, and, so, $Z^\lambda = \mathcal{E}(-(\lambda - m) \cdot \bar{B})_T$, where $\bar{B} = \int_0^\cdot m_u du + B$ is a $\bar{\mathbb{P}}$ -martingale. It remains to use the same argument as in the previous paragraph to show that Z^λ indeed satisfies the reverse Hölder inequality $R \log R$ under $\bar{\mathbb{P}}$.

3.2 Proof of Lemma 1

Let $(m, n) = (m^G, n^G)$ from (7); more generally, we suppress the superscripts λ and G throughout to increase legibility. A combination of (9) and (10) yields that

$$Y = -c - \rho \cdot B^\lambda - \log Z,$$

and a simple calculation confirms that (Y, μ, ν) satisfies (13). Next, we show $Y - X \in \mathcal{S}^\infty$. We start by defining the probability measure $\bar{\mathbb{P}}$ via $d\bar{\mathbb{P}}/d\mathbb{P} = \mathcal{E}(-m \cdot B - n \cdot W)_T$ so that under $\bar{\mathbb{P}}$, $D = Y - X$ is the certainty-equivalent process corresponding to the zero endowment. By (11), we have $D \geq 0$ as well as

$$dD_t = (\mu_t - m_t) d\bar{B} + (\nu_t - n_t) d\bar{W} + \left(\frac{1}{2}(\nu_t - n_t)^2 - \frac{1}{2}(\lambda_t - m_t)^2 + (\lambda_t - m_t)(\mu_t - m_t) \right) dt, \text{ with } D_T = 0, \quad (19)$$

where $\bar{B} = B + \int_0^\cdot m_u du$ and $\bar{W} = W + \int_0^\cdot n_u du$ are $\bar{\mathbb{P}}$ -Brownian motions. Using the notation \mathbb{Q}^λ , as well as the argument of Proof of Proposition 1 above, we can deduce that $\mathbb{Q}^\lambda \in \mathcal{M}^{\lambda-m}$ (where \mathbb{P} in the definition of $\mathcal{M}^{\lambda-m}$ is replaced by $\bar{\mathbb{P}}$). We claim that

$$D_\tau \leq H_\tau(\mathbb{Q}^\lambda | \bar{\mathbb{P}}), \quad \text{for any } \tau \in \mathcal{T}. \quad (20)$$

Proposition 1, applied under $\bar{\mathbb{P}}$ and with zero random endowment produces the dual optimizer $\mathbb{Q}^{\lambda,G}$, with $\bar{\mathbb{P}}$ -density $Z^{\lambda-m,G}$. If we project both sides of the equality $\bar{c}^{\lambda,G} + \rho^{\lambda,G} \cdot B_T^\lambda = -\log(Z_T^{\lambda-m,G})$ under $\mathbb{Q}^{\lambda,G}$ onto \mathcal{F}_τ we obtain

$$D_\tau = H_\tau(\mathbb{Q}^{\lambda,G} | \bar{\mathbb{P}}).$$

No integrability issues arise here since $H(\mathbb{Q}^{\lambda,G} | \bar{\mathbb{P}}) < \infty$ and $\rho^{\lambda,G} \cdot B^\lambda$ is a $\mathbb{Q}^{\lambda,G}$ -martingale (by part (iii) of Proposition 1). The required inequality (20) follows from the optimality of $\mathbb{Q}^{\lambda,G}$ in part (ii) of Proposition 1.

The right-hand side of (20) can be written as

$$H_\tau(\mathbb{Q}^\lambda | \bar{\mathbb{P}}) = \mathbb{E}_\tau^{\mathbb{Q}^\lambda} \left[\frac{1}{2} \int_\tau^T (\lambda_t - m_t)^2 dt - \int_\tau^T (\lambda_t - m_t) dB_t^\lambda \right] \leq \frac{1}{2} \|\lambda - m\|_{\text{bmo}(\mathbb{Q}^\lambda)}^2.$$

Given that both λ and m belong to bmo we have $\lambda - m \in \text{bmo}(\mathbb{Q}^\lambda)$ by [22, Theorem 3.6]. Therefore, we can combine (20) and the fact that $D \geq 0$ to conclude that $D \in \mathcal{S}^\infty$. Consequently, it suffices to apply the standard bmo -estimate for quadratic BSDEs

(see Lemma 9) to (19), to obtain $(\mu - m, \nu - n) \in \text{bmo}(\overline{\mathbb{P}})$. Since $(m, n) \in \text{bmo}$, another application of [22, Theorem 3.6] confirms that $(\mu, \nu) \in \text{bmo}$.

Lastly, we show that there can be at most one solution to (13) with $(\mu, \nu) \in \text{bmo}$. Let (Y, μ, ν) and $(\tilde{Y}, \tilde{\mu}, \tilde{\nu})$ be two solutions with $(\mu, \nu), (\tilde{\mu}, \tilde{\nu}) \in \text{bmo}$. For $\delta Y = \tilde{Y} - Y$, we have

$$d(\delta Y)_t = \delta \mu_t dB_t^\lambda + \delta \nu_t dW_t^{\bar{\nu}}, \quad \delta Y_T = 0.$$

Here $\delta \mu = \tilde{\mu} - \mu$, $\delta \nu = \tilde{\nu} - \nu$, $\bar{\nu} = \frac{1}{2}(\nu + \tilde{\nu})$, and $W^{\bar{\nu}} = W + \int_0^\cdot \bar{\nu}_t dt$ is a $\mathbb{Q}^{\lambda, \bar{\nu}}$ -Brownian motion, where $\mathbb{Q}^{\lambda, \bar{\nu}}$ is defined via $d\mathbb{Q}^{\lambda, \bar{\nu}}/d\mathbb{P} = \mathcal{E}(-\lambda \cdot B - \bar{\nu} \cdot W)_T$. By [22, Theorem 3.6], both $\delta \mu \cdot B^\lambda$ and $\delta \nu \cdot W^{\bar{\nu}}$ are $\text{BMO}(\mathbb{Q}^{\lambda, \bar{\nu}})$ -martingales. Hence $\delta Y_T = 0$ implies that $\delta Y = 0$ and, consequently, $\delta \mu = \delta \nu = 0$.

3.3 Proof of Theorem 1

(1) \Rightarrow (2). Given an equilibrium $\lambda \in \Lambda_\alpha(\mathbf{G})$ and $i \in \{1, 2, \dots, I\}$, let ρ^{λ, G^i} be the primal optimizer of agent i , and let (Y^i, μ^i, ν^i) be defined as in Lemma 1 where (13) has the terminal condition $Y_T^i = G^i$. Since λ is an equilibrium, $\sum_i \alpha^i \rho^{\lambda, G^i} = 0$, and so $\lambda = \lambda - \sum_i \alpha^i \rho^{\lambda, G^i} = \sum_i \alpha^i \mu^i$, for $\mu^i = \lambda - \rho^{\lambda, G^i}$, implying that $(Y, \mu, \nu) = (Y^i, \mu^i, \nu^i)_i$ solves the system (14). The property $(\mu, \nu) \in \text{bmo}^I$ follows from Lemma 1.

(2) \Rightarrow (1). Given a solution (Y, μ, ν) of (14), we set $\lambda = \sum_i \alpha^i \mu^i$. This way, individual equations in (14) turn into BSDEs of the form (13). If we set $\rho^{\lambda, i} = \lambda - \mu^i$ the market clearing condition $\sum_i \alpha_i \rho^{\lambda, i} = 0$ holds. Since $(\mu^i, \nu^i) \in \text{bmo}$ the uniqueness part of Lemma 1 implies that λ, ρ^i maximizes single-agents' utilities.

3.4 Proof of Theorem 2

In order to prove Theorem 2, we start with a refinement of the classical result on uniform equivalence of bmo spaces (see Theorem 3.6, p. 62 in [22]), based on a result of Chinkvinidze and Mania (see [7]).

Lemma 3 *Let $\sigma \in \text{bmo}$ be such that $\|\sigma\|_{\text{bmo}} =: \sqrt{2}R$ for some $R < 1$. If $\hat{\mathbb{P}} \sim \mathbb{P}$ is such that $d\hat{\mathbb{P}} = \mathcal{E}(\sigma \cdot \tilde{B})_T d\mathbb{P}$, for some \mathbb{F} -Brownian motion \tilde{B} , then, for all $\zeta \in \text{bmo}$, we have*

$$(1 + R)^{-1} \|\zeta\|_{\text{bmo}} \leq \|\zeta\|_{\text{bmo}(\hat{\mathbb{P}})} \leq (1 - R)^{-1} \|\zeta\|_{\text{bmo}}. \tag{21}$$

Proof Since $M = \sigma \cdot \tilde{B}$ is a BMO-martingale, Theorem 3.6. in [22] states that the spaces bmo and $\text{bmo}(\hat{\mathbb{P}})$ coincide and that the norms $\|\cdot\|_{\text{bmo}}$ and $\|\cdot\|_{\text{bmo}(\hat{\mathbb{P}})}$ are uniformly equivalent. This norm equivalence is refined in [7]; Theorem 2 there implies that

$$(1 + R)^{-1} \|\zeta\|_{\text{bmo}} \leq \|\zeta\|_{\text{bmo}(\hat{\mathbb{P}})} \leq (1 + \hat{R}) \|\zeta\|_{\text{bmo}}, \text{ where } \hat{R} = \sqrt{\frac{1}{2} \|\sigma\|_{\text{bmo}(\hat{\mathbb{P}})}^2}. \tag{22}$$

Clearly, only the second inequality in (21) needs to be discussed; it is obtained by substituting $\zeta = \sigma$ into the second inequality in (22):

$$\sqrt{2}\hat{R} = \|\sigma\|_{\text{bmo}(\hat{\mathbb{P}})} = (1 + \hat{R})\|\sigma\|_{\text{bmo}} \leq \sqrt{2}(1 + \hat{R})R, \text{ so that } (1 + \hat{R}) \leq (1 - R)^{-1}.$$

Coming back to Theorem 2, suppose that (16) is satisfied. Then there exists $\xi^c \in \text{EBMO}$ such that

$$\max_i \|(m^i - m^c, n^i - n^c)\|_{\text{bmo}(\mathbb{P}^c)} \leq \epsilon^*. \tag{23}$$

To simplify notation, we introduce $\mathbf{m} = (m^i)_i$ and $\mathbf{n} = (n^i)_i$. A calculation shows that (component-by-component)

$$\begin{aligned} d(Y_t - \xi_t^c) &= (\boldsymbol{\mu}_t - m_t^c) dB_t^c + (\mathbf{v}_t - n_t^c) dW_t^c \\ &\quad + \left(\frac{1}{2}(\mathbf{v}_t - n_t^c)^2 - \frac{1}{2}(\lambda_t - m_t^c)^2 + (\lambda_t - m_t^c)(\boldsymbol{\mu}_t - m_t^c) \right) dt, \\ \mathbf{Y}^T - \xi_T^c &= \mathbf{G} - \xi^c, \end{aligned}$$

where $\lambda = A[\boldsymbol{\mu}]$, $\xi_t^c = -\log(\mathbb{E}_t[\exp(-\xi^c)])$, and B^c, W^c are \mathbb{P}^c -Brownian motions. This is exactly the type of system covered in (14). Therefore, to ease notation, we treat, throughout this section, \mathbb{P} as \mathbb{P}^c , B as B^c , W as W^c , and $\mathbf{G}, \lambda, \boldsymbol{\mu}, \mathbf{v}$ as their shifted versions, i.e., eg. \mathbf{G} as $\mathbf{G} - \xi^c$, λ as $\lambda - m^c$, etc. As a result, (23) translates to

$$\max_i \|(m^i, n^i)\|_{\text{bmo}} \leq \epsilon^*. \tag{24}$$

We proceed by setting up a framework for the Banach fixed-point theorem. First observe that since $(m^i, n^i) \in \text{bmo}$ for all i , then bmo is a natural space in which the fixed-point theorem can be applied. Given $\lambda \in \text{bmo}$ and $\mathbf{G} = (G^i)_i$, let $\mathbf{Y}^\lambda = (Y^{\lambda, G^i})_i$ and $\mathbf{X} = (X^{G^i})_i$, denote the agents' certainty-equivalent processes with and without assess the market, respectively; we also set $(\boldsymbol{\mu}^{\lambda, \mathbf{G}}, \mathbf{v}^{\lambda, \mathbf{G}}) = (\boldsymbol{\mu}^{\lambda, G^i}, \mathbf{v}^{\lambda, G^i})_i$, where $(\boldsymbol{\mu}^{\lambda, G^i}, \mathbf{v}^{\lambda, G^i})_i$ is defined in Lemma 1. This allows us to define (a simple transformation of) the **excess-demand map**

$$F : \lambda \mapsto A[\boldsymbol{\mu}^{\lambda, \mathbf{G}}],$$

where the aggregation operator $A[\cdot]$ is defined in (5). The significance of this map lies in the simple fact that λ is an equilibrium if and only if $F(\lambda) = \lambda$, i.e., if λ is a fixed point of F .

Before proceeding to studying properties of F , we first record the following a-priori estimate on λ in equilibrium.

Lemma 4 *If $\lambda \in \text{bmo}$ is an equilibrium, then*

$$\|\lambda\|_{\text{bmo}} \leq \max_i \|(m^i, n^i)\|_{\text{bmo}}.$$

Proof Aggregating all single equations in (14) and (7), we obtain

$$dA[Y_t^\lambda - X_t] = (\lambda_t - A[m_t])dB_t + A[v_t^\lambda - n_t]dW_t + \frac{1}{2}(\lambda_t^2 + A[(v_t^\lambda)^2])dt - \frac{1}{2}A[m_t^2 + n_t^2]dt.$$

Let $(\sigma_n)_n$ be a reducing sequence for local martingale part above. For any $\tau \in \mathcal{T}$, integrating the previous dynamics from $\tau \wedge \sigma_n$ to σ_n and projecting onto \mathcal{F}_τ yields

$$\begin{aligned} \mathbb{E}_\tau [A[Y_{\sigma_n}^\lambda - X_{\sigma_n}]] - A[Y_{\tau \wedge \sigma_n}^\lambda - X_{\tau \wedge \sigma_n}] &= \\ &= \frac{1}{2}\mathbb{E}_\tau \left[\int_{\tau \wedge \sigma_n}^{\sigma_n} (\lambda_t^2 + A[(v_t^\lambda)^2])dt \right] - \frac{1}{2}\mathbb{E}_\tau \left[\int_{\tau \wedge \sigma_n}^{\sigma_n} A[m_t^2 + n_t^2]dt \right]. \end{aligned} \tag{25}$$

Sending $n \rightarrow \infty$, since $Y^\lambda - X \geq 0$ (component-by-component) and is also bounded (see Lemma 1) and $A[X_T] = A[G] = A[Y_T^\lambda]$, we obtain

$$\begin{aligned} \|\lambda\|_{\text{bmo}}^2 &\leq \|\lambda^2 + A[(v^\lambda)^2]\|_{\text{bmo}} \leq \|A[m^2 + n^2]\|_{\text{bmo}} \\ &\leq A[\|(m, n)\|_{\text{bmo}}^2] \leq \max_i \|(m^i, n^i)\|_{\text{bmo}}^2. \end{aligned}$$

For the third inequality, note that $\mathbb{E}_\tau[\int_\tau^T A[m_t^2 + n_t^2]dt] \leq A[\|(m, n)\|_{\text{bmo}}^2]$ holds for all stopping times τ . □

For arbitrary $\lambda \in \text{bmo}$, the following estimate gives an explicit upper bound on the (nonnegative) difference $D^{\lambda,i} = Y^{\lambda,i} - X^i$.

Lemma 5 *Suppose that $\|\lambda\|_{\text{bmo}} < \sqrt{2}$. Then,*

$$0 \leq \sqrt{D^{\lambda,i}} \leq \frac{\|\lambda\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}}}{\sqrt{2} - \|\lambda\|_{\text{bmo}}}, \quad \text{for all } i.$$

Proof Let \mathbb{Q}^λ be the probability such that $d\mathbb{Q}^\lambda = Z_T^\lambda d\mathbb{P}$, where $Z^\lambda = \mathcal{E}(-\lambda \cdot B)$. Since $\mathbb{Q}^\lambda \in \mathcal{M}^\lambda$, then the argument that leads to (20) also implies that

$$Y_\tau^{\lambda,i} \leq H_\tau(\mathbb{Q}^\lambda | \mathbb{P}) + \mathbb{E}_\tau^{\mathbb{Q}^\lambda} [G^i], \quad \text{for any } \tau \in \mathcal{T}. \tag{26}$$

On the right-hand side of (26),

$$H_\tau(\mathbb{Q}^\lambda | \mathbb{P}) = \mathbb{E}_\tau^{\mathbb{Q}^\lambda} \left[\frac{1}{2} \int_\tau^T \lambda_u^2 du - \int_\tau^T \lambda_u dB_u^\lambda \right] \leq \frac{1}{2} \|\lambda\|_{\text{bmo}(\mathbb{Q}^\lambda)}^2.$$

Since $\|\lambda\|_{\text{bmo}(\mathbb{Q}^\lambda)} \leq \sqrt{2}\|\lambda\|_{\text{bmo}} / (\sqrt{2} - \|\lambda\|_{\text{bmo}})$, as follows from Lemma 3, we obtain

$$H_\tau(\mathbb{Q}^\lambda | \mathbb{P}) \leq \frac{\|\lambda\|_{\text{bmo}}^2}{(\sqrt{2} - \|\lambda\|_{\text{bmo}})^2}.$$

Furthermore, recalling that $X_T^i = G^i$ and $dX_t^i = m_t^i dB_t + n_t^i dW_t + \frac{1}{2}((m_t^i)^2 + (n_t^i)^2)dt$, we note that

$$\mathbb{E}_\tau^{\mathbb{Q}^\lambda} [G^i] = \mathbb{E}_\tau [(Z_T^\lambda / Z_\tau^\lambda) G^i] = \mathbb{E}_\tau [(Z_T^\lambda / Z_\tau^\lambda) X_T^i].$$

Given that Z^λ is a BMO-martingale and $\|(m^i, n^i)\|_{\text{bmo}} < \infty$, the integration-by-parts formula implies that

$$\begin{aligned} \mathbb{E}_\tau[(Z_T^\lambda/Z_\tau^\lambda)X_T^i] &= \\ &= X_\tau^i - \mathbb{E}_\tau \left[\int_\tau^T (Z_u^\lambda/Z_\tau^\lambda)\lambda_u m_u^i du \right] + \frac{1}{2}\mathbb{E}_\tau \left[\int_\tau^T (Z_u^\lambda/Z_\tau^\lambda) \left((m_u^i)^2 + (n_u^i)^2 \right) du \right] \\ &= X_\tau^i - \mathbb{E}_\tau^{\mathbb{Q}^\lambda} \left[\int_\tau^T \lambda_u m_u^i du \right] + \frac{1}{2}\mathbb{E}_\tau^{\mathbb{Q}^\lambda} \left[\int_\tau^T \left((m_u^i)^2 + (n_u^i)^2 \right) du \right]. \end{aligned}$$

A use of Holder’s inequality then gives

$$\begin{aligned} \mathbb{E}_\tau^{\mathbb{Q}^\lambda} [G^i] - X_\tau^i &\leq \|\lambda\|_{\text{bmo}(\mathbb{Q}^\lambda)}\|m^i\|_{\text{bmo}(\mathbb{Q}^\lambda)} + \frac{1}{2}\|(m^i, n^i)\|_{\text{bmo}(\mathbb{Q}^\lambda)}^2 \\ &\leq \frac{2\|\lambda\|_{\text{bmo}}\|(m^i, n^i)\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}}^2}{(\sqrt{2} - \|\lambda\|_{\text{bmo}})^2}, \end{aligned}$$

where, again, the last inequality follows from Lemma 3. A Combination of the above estimates shows that

$$D_\tau^{\lambda,i} = Y_\tau^{\lambda,i} - X_\tau^i \leq \left(\frac{\|\lambda\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}}}{\sqrt{2} - \|\lambda\|_{\text{bmo}}} \right)^2,$$

which completes the proof. □

Lemma 6 Suppose that $\lambda \in \text{bmo}$ satisfies

$$\|\lambda\|_{\text{bmo}} < \frac{\sqrt{2} - \|(m^i, n^i)\|_{\text{bmo}}}{2}.$$

Then, it holds that

$$\begin{aligned} \|(\mu^{\lambda,i}, \nu^{\lambda,i})\|_{\text{bmo}} &\leq \\ &\leq \frac{(\sqrt{2} + \|(m^i, n^i)\|_{\text{bmo}})\|(m^i, n^i)\|_{\text{bmo}} + \|\lambda\|_{\text{bmo}}(\|\lambda\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}})}{\sqrt{2} - 2\|\lambda\|_{\text{bmo}} - \|(m^i, n^i)\|_{\text{bmo}}}. \end{aligned}$$

In particular, the previous is also a bound for both $\|\mu^{\lambda,i}\|_{\text{bmo}}$ and $\|\nu^{\lambda,i}\|_{\text{bmo}}$.

Proof Set $Y = Y^\lambda$, $\mu = \mu^\lambda$ and $\nu = \nu^\lambda$ to increase legibility, and define

$$f^i = \frac{\|\lambda\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}}}{\sqrt{2} - \|\lambda\|_{\text{bmo}}},$$

and $D = Y - X$. Note that $D_T^i = 0$ and $0 \leq D^i \leq (f^i)^2$ from Lemma 5. Since

$$dD_t^i = (\mu_t^i - m_t^i)dB_t + (\nu_t^i - n_t^i)dW_t + \frac{1}{2} \left((\nu_t^i)^2 - \lambda_t^2 + 2\mu_t^i \lambda_t - (m_t^i)^2 - (n_t^i)^2 \right) dt,$$

an application of Itô's lemma gives

$$\begin{aligned} d(D_t^i)^2 &= 2D_t^i(\mu_t^i - m_t^i)dB_t + 2D_t^i(v_t^i - n_t^i)dW_t \\ &\quad + D_t^i\left((v_t^i)^2 - \lambda_t^2 + 2\mu_t^i\lambda_t - (m_t^i)^2 - (n_t^i)^2\right)dt \\ &\quad + \left((\mu_t^i - m_t^i)^2 + (v_t^i - n_t^i)^2\right)dt. \end{aligned}$$

Next, we take a reducing sequence $(\sigma_n)_n$ for the local martingales on the right-hand side above, as well as an arbitrary $\tau \in \mathcal{T}$. If we integrate the above dynamics between $\sigma_n \wedge \tau$ and σ_n , and use the facts that $(v^i)^2 \geq 0$, $\lambda^2 - 2\mu^i\lambda \leq (\mu^i - \lambda)^2$, and $D^i \geq 0$, we obtain

$$\begin{aligned} (D_{\sigma_n}^i)^2 &\geq (D_{\sigma_n}^i) - (D_{\tau \wedge \sigma_n}^i)^2 \geq 2 \int_{\tau \wedge \sigma_n}^{\sigma_n} D_t^i(\mu_t^i - m_t^i)dB_t + 2 \int_{\tau \wedge \sigma_n}^{\sigma_n} D_t^i(v_t^i - n_t^i)dW_t \\ &\quad - \int_{\tau \wedge \sigma_n}^{\sigma_n} D_t^i\left((\mu_t^i - \lambda_t)^2 + (m_t^i)^2 + (n_t^i)^2\right)dt \\ &\quad + \int_{\tau \wedge \sigma_n}^{\sigma_n} \left((\mu_t^i - m_t^i)^2 + (v_t^i - n_t^i)^2\right)dt. \end{aligned}$$

Given that $D^i \leq (f^i)^2$, a projection of both sides above on \mathcal{F}_τ yields

$$\begin{aligned} &\mathbb{E}_\tau \left[\int_{\tau \wedge \sigma_n}^{\sigma_n} \left((\mu_t^i - m_t^i)^2 + (v_t^i - n_t^i)^2\right) dt \right] \\ &\leq \mathbb{E}_\tau [D_{\sigma_n}^i] + (f^i)^2 \mathbb{E}_\tau \left[\int_{\tau \wedge \sigma_n}^{\sigma_n} \left((\mu^i - \lambda)^2 + (m_t^i)^2 + (n_t^i)^2\right) dt \right]. \end{aligned}$$

Sending $n \rightarrow \infty$ first on the right-hand side then the left, helped by the facts that D^i is bounded and $D_T^i = 0$, implies that

$$\|(\mu^i, v^i) - (m^i, n^i)\|_{\text{bmo}}^2 \leq (f^i)^2 \left(\|\mu^i - \lambda\|_{\text{bmo}}^2 + \|(m^i, n^i)\|_{\text{bmo}}^2 \right).$$

Taking square roots on both sides, and using the elementary inequality $\sqrt{x^2 + y^2} \leq |x| + |y|$ for any x, y , and the fact that $\|\mu^i - \lambda\|_{\text{bmo}} \leq \|(\mu^i, v^i)\|_{\text{bmo}} + \|\lambda\|_{\text{bmo}}$, we obtain

$$\|(\mu^i, v^i) - (m^i, n^i)\|_{\text{bmo}} \leq f^i \left(\|\lambda\|_{\text{bmo}} + \|(\mu^i, v^i)\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}} \right).$$

Finally, since $\|(\mu^i, v^i)\|_{\text{bmo}} \leq \|(\mu^i, v^i) - (m^i, n^i)\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}}$, it follows that

$$(1 - f^i)\|(\mu^i, v^i)\|_{\text{bmo}} \leq \|(m^i, n^i)\|_{\text{bmo}} + f^i \left(\|\lambda\|_{\text{bmo}} + \|(m^i, n^i)\|_{\text{bmo}} \right),$$

from which the result follows after simple algebra. \square

Define $\mathcal{B}(r) = \{\lambda \in \text{bmo} : \|\lambda\|_{\text{bmo}} \leq r\}$. The following result shows that the excess-demand map F maps $\mathcal{B}(r)$ into itself for an appropriate choice of r .

Lemma 7 *There exists a sufficiently small ϵ^* independent of the number of the agents I , such that whenever $\max_i \|(m^i, n^i)\|_{bmo} \leq \sqrt{2}\epsilon$ for $\epsilon \leq \epsilon^*$, F maps $\mathcal{B}(2\epsilon)$ into itself.*

Proof Suppose that $\max_i \|(m^i, n^i)\|_{bmo} \leq \sqrt{2}\epsilon$ for some $\epsilon \in (0, 1)$ determined later. Let us consider $\lambda \in \mathcal{B}(\sqrt{2}\epsilon a)$, where $a \in [1, 1/\epsilon)$ will also be determined later. Our goal is to choose a sufficiently small ϵ such that $A[\mu^\lambda] \in \mathcal{B}(\sqrt{2}\epsilon a)$ for some $a \in [1, 1/\epsilon)$, whenever λ is chosen from the same ball. If this task is successful, given $a \geq 1$, Lemma 4 implies that all possible equilibria are already in the same ball. Hence the local uniqueness immediately implies global uniqueness in bmo.

For $\lambda \in \mathcal{B}(\sqrt{2}\epsilon a)$, Lemma 5 gives

$$0 \leq \sqrt{D^{\lambda, i}} \leq \frac{\epsilon(1+a)}{1-a\epsilon} =: \phi(\epsilon, a).$$

Note that ϕ is an increasing function of both arguments. For Lemma 6 we need $\phi < 1$. Therefore, only $\epsilon \in (0, 1)$ such that $\phi(\epsilon, 1) < 1$ can be used, i.e., $\epsilon \in (0, 1/3)$. Taking $\epsilon \in (0, 1/3)$ and $a \in [1, 1/\epsilon)$, in order to have $\phi(\epsilon, a) < 1$, it is necessary and sufficient that

$$a < \frac{1-\epsilon}{2\epsilon} =: \bar{a}(\epsilon).$$

Note that \bar{a} is decreasing in ϵ with $\bar{a}(0+) = \infty$ and $\bar{a}(1/3) = 1$, and that $\bar{a}(\epsilon) < 1/\epsilon$ holds for all $\epsilon \in (0, 1/3)$.

Now, in order to have $\|\mu^{\lambda, i}\|_{bmo} \leq \sqrt{2}\epsilon a$, by Lemma 6 we need to ensure that

$$\frac{2(1+\epsilon)\epsilon + 2a\epsilon^2(1+a)}{\sqrt{2}(1-2a\epsilon-\epsilon)} \leq a\sqrt{2}\epsilon,$$

or, equivalently, that

$$q(a, \epsilon) := 3\epsilon a^2 - (1-2\epsilon)a + (1+\epsilon) \leq 0.$$

Fix $a > 1$, say $a = \sqrt{2}$, there exists a sufficiently small ϵ^* such that $q(\sqrt{2}, \epsilon) \leq 0$ for any $\epsilon \leq \epsilon^*$. Note that the choice of ϵ^* is independent of the number of the agent I . For such choice of ϵ , we have $\|\mu^{\lambda, i}\|_{bmo} \leq 2\epsilon$ for all i . As a weighted sum of individual component, $\|F[\lambda]\|_{bmo} \leq A[\|\mu^\lambda\|_{bmo}]$, hence $F[\lambda] \in \mathcal{B}(2\epsilon)$ as well. \square

Finally we check that F is a contraction on $\mathcal{B}(2\epsilon)$ for sufficiently small ϵ .

Lemma 8 *There exists a sufficiently small ϵ^* independent of the number of the agents I , such that whenever $\max_i \|(m^i, n^i)\|_{bmo} \leq \sqrt{2}\epsilon$ for $\epsilon \leq \epsilon^*$, F is a contraction on $\mathcal{B}(2\epsilon)$.*

Proof We drop the superscript i to increase legibility. Set $\delta Y = Y^\lambda - Y^{\bar{\lambda}}$, and note that $\|\delta Y\|_{S^\infty} < \infty$ from Lemma 5 and $\delta Y_T = 0$. Set $(\mu, \nu) = (\mu^\lambda, \nu^\lambda)$ and $(\bar{\mu}, \bar{\nu}) = (\mu^{\bar{\lambda}}, \nu^{\bar{\lambda}})$. Denote $\bar{\lambda} = (\lambda + \bar{\lambda})/2$, $\bar{\mu} = (\mu + \bar{\mu})/2$, and $\bar{\nu} = (\nu + \bar{\nu})/2$. Calculation using (13) gives

$$\begin{aligned} d\delta Y_t &= (\mu_t - \bar{\mu}_t)dB_t + (\nu_t - \bar{\nu}_t)dW_t + \frac{1}{2} \left(\nu_t^2 - \bar{\nu}_t^2 + \bar{\lambda}_t^2 - \lambda_t^2 + 2\mu_t \lambda_t - 2\bar{\mu}_t \bar{\lambda}_t \right) dt \\ &= (\mu_t - \bar{\mu}_t)dB_t^{\bar{\lambda}} + (\nu_t - \bar{\nu}_t)dW_t^{\bar{\nu}} - (\lambda_t - \bar{\lambda}_t)(\bar{\lambda}_t - \bar{\mu}_t)dt, \end{aligned}$$

where $B^{\bar{\lambda}} = B + \int_0^\cdot \lambda_t dt$, $W^{\bar{\nu}} = W + \int_0^\cdot \bar{\nu}_t dt$ are Brownian motions under $\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}$, and $\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}$ is defined via $d\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}/d\mathbb{P} = \mathcal{E}(-\bar{\lambda} \cdot B - \bar{\nu} \cdot W)_T$. For an arbitrary $\tau \in \mathcal{T}$, integrating the previous dynamics on $[\tau, T]$, taking conditional expectation $\mathbb{E}_\tau^{\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}}$ on both sides, (both local martingales are $\text{BMO}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})$ -martingales, due to $\mu, \tilde{\mu}, \nu, \tilde{\nu} \in \text{bmo}$ from Lemma 6 and [22, Theorem 3.6]), and finally using $\delta Y_T = 0$, we obtain

$$|\delta Y_\tau| \leq \mathbb{E}_\tau^{\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}} \left[\int_\tau^T |\lambda_t - \tilde{\lambda}_t| |\bar{\lambda}_t - \bar{\mu}_t| dt \right] \leq \|\bar{\lambda} - \bar{\mu}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})} \|\lambda - \tilde{\lambda}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})}.$$

This implies that

$$\|\delta Y\|_{\mathcal{S}^\infty} \leq \|\bar{\lambda} - \bar{\mu}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})} \|\lambda - \tilde{\lambda}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})}. \tag{27}$$

To establish the Lipschitz continuity of F , we use Itô’s formula to get

$$\begin{aligned} d(\delta Y_t)^2 &= 2\delta Y_t(\mu_t - \tilde{\mu}_t)dB_t^{\bar{\lambda}} + 2\delta Y_t(\nu_t - \tilde{\nu}_t)dW_t^{\bar{\nu}} - 2\delta Y_t(\lambda_t - \tilde{\lambda}_t)(\bar{\lambda}_t - \bar{\mu}_t)dt \\ &\quad + \left((\mu_t - \tilde{\mu}_t)^2 + (\nu_t - \tilde{\nu}_t)^2 \right) dt. \end{aligned}$$

For an arbitrary $\tau \in \mathcal{T}$, an integration of the above dynamics between τ and T , and using (27) and $\delta Y_T = 0$, yields that

$$\begin{aligned} \mathbb{E}_\tau^{\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}} \left[\int_\tau^T \left((\mu_t - \tilde{\mu}_t)^2 + (\nu_t - \tilde{\nu}_t)^2 \right) dt \right] &\leq \\ &\leq 2\|\delta Y\|_{\mathcal{S}^\infty} \mathbb{E}_\tau^{\mathbb{Q}^{\bar{\lambda}, \bar{\nu}}} \left[\int_\tau^T (\lambda_t - \tilde{\lambda}_t)(\bar{\lambda}_t - \bar{\mu}_t) dt \right] \\ &\leq 2\|\bar{\lambda} - \bar{\mu}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})}^2 \|\lambda - \tilde{\lambda}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})}^2, \end{aligned}$$

which, in turn, implies that

$$\|(\tilde{\mu}, \tilde{\nu}) - (\mu, \nu)\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})} \leq \sqrt{2} \|\bar{\lambda} - \bar{\mu}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})} \|\lambda - \tilde{\lambda}\|_{\text{bmo}(\mathbb{Q}^{\bar{\lambda}, \bar{\nu}})}.$$

Note that Lemma 6 and the estimates in Lemma 7 also imply that $\|\bar{\nu}\|_{\text{bmo}} \leq 2\epsilon$, where 2ϵ is taken from Lemma 7. Therefore, $\|(\bar{\lambda}, \bar{\nu})\|_{\text{bmo}} \leq 4\epsilon$ and, similarly, $\|\bar{\lambda} - \bar{\mu}\|_{\text{bmo}} \leq 4\epsilon$. Therefore, it follows from Lemma 3 that

$$\begin{aligned} \|(\tilde{\mu}, \tilde{\nu}) - (\mu, \nu)\|_{\text{bmo}} &\leq \sqrt{2} \frac{1 + 2\sqrt{2}\epsilon}{(1 - 2\sqrt{2}\epsilon)^2} \|\bar{\lambda} - \bar{\mu}\|_{\text{bmo}} \|\lambda - \tilde{\lambda}\|_{\text{bmo}} \\ &\leq \frac{1 + 2\sqrt{2}\epsilon}{(1 - 2\sqrt{2}\epsilon)^2} 8\epsilon \|\lambda - \tilde{\lambda}\|_{\text{bmo}}. \end{aligned}$$

Choosing sufficiently small ϵ so that $\frac{1 + 2\sqrt{2}\epsilon}{(1 - 2\sqrt{2}\epsilon)^2} 8\epsilon < 1$, the proof is complete after aggregating all components. \square

Proof (of Theorem 2) We have shown in the sequence of lemmas above that, when (24) holds, the excess-demand map F is a contraction on $\mathcal{B}(2\epsilon)$ and that $(\mu^\lambda, \nu^\lambda) \in \text{bmo}^I$. The Banach fixed point theorem implies that F has a unique fixed point λ with $\|\lambda\|_{\text{bmo}} \leq 2\epsilon$. Therefore the system (14) admits a solution (Y, μ, ν) with $(\mu, \nu) \in \text{bmo}^I$. Hence λ is an equilibrium by Theorem 1. For the uniqueness of equilibrium, Lemma 4 implies that any equilibrium λ satisfies $\|\lambda\|_{\text{bmo}} \leq \max_i \|(m^i, n^i)\|_{\text{bmo}} \leq \sqrt{2}\epsilon$. However, we have already shown that there can be only one equilibrium λ in $\mathcal{B}(2\epsilon)$. Therefore we immediately have global uniqueness of equilibrium. Given the unique λ , by Lemma 1, (Y, μ, ν) is the unique solution to (14) with $(\mu, \nu) \in \text{bmo}^I$. \square

3.5 Proof of Corollary 1

Summing both sides of $\|E^i - E^j\|_{\mathbb{L}^\infty} \leq \chi_0^E (\|E^i\|_{\mathbb{L}^\infty} + \|E^j\|_{\mathbb{L}^\infty})$ over j , we obtain

$$\begin{aligned} I\|E^i\|_{\mathbb{L}^\infty} - \|E_\Sigma\|_{\mathbb{L}^\infty} &\leq \|IE^i - \sum_j E^j\|_{\mathbb{L}^\infty} \leq \sum_j \|E^i - E^j\|_{\mathbb{L}^\infty} \\ &\leq \chi_0^E I\|E^i\|_{\mathbb{L}^\infty} + \chi_0^E \sum_j \|E^j\|_{\mathbb{L}^\infty}, \end{aligned}$$

which implies

$$(1 - \chi_0^E)\|E^i\|_{\mathbb{L}^\infty} \leq \frac{1}{I}\|E_\Sigma\|_{\mathbb{L}^\infty} + \chi_0^E \frac{1}{I} \sum_j \|E^j\|_{\mathbb{L}^\infty}.$$

Summing both sides of the previous inequality over i yields

$$\sum_i \|E^i\|_{\mathbb{L}^\infty} \leq \frac{1}{1-2\chi_0^E} \|E_\Sigma\|_{\mathbb{L}^\infty}.$$

The previous two inequalities combined then imply

$$\|E^i\|_{\mathbb{L}^\infty} \leq \frac{1}{1-2\chi_0^E} \frac{1}{I} \|E_\Sigma\|_{\mathbb{L}^\infty}, \quad \text{for all } i.$$

Therefore

$$\max_i \frac{\|E^i\|_{\mathbb{L}^\infty}}{\delta^i} \leq \frac{1}{1-2\chi_0^E} \frac{1}{I\delta_0} \|E_\Sigma\|_{\mathbb{L}^\infty},$$

where the right-hand side is strictly less than $(\epsilon^*)^2/4$ for sufficiently large I . Hence (17) is satisfied when $I \geq I_0$, for some I_0 , and the existence of equilibrium follows from Theorem 2.

3.6 Proof of Corollary 2

Throughout the proof, we treat G as $G - \xi^c$ and suppress the superscript i when we work with each component.

Recalling (6) and Remark 2, we have $\mathbb{E}[G^2] < \infty$, which combined with the assumption $D^b G, D^w G \in \mathcal{S}^\infty$ implies $G \in \mathbb{D}^{1,2}$. Let $G = \mathbb{E}[G] + M_T$, where $M_T = \bar{m} \cdot B_T + \bar{n} \cdot W_T$ for some (\bar{m}, \bar{n}) . Clark-Ocone formula implies that $\mathbb{E}_\theta[D_\theta G] = (\bar{m}_\theta, \bar{n}_\theta)$, for any $\theta \leq T$, hence $(\bar{m}, \bar{n}) \in \mathcal{S}^\infty$ as well. As a result, there exists a constant C such that $\langle M \rangle_T \leq CT$, implying that G has at most Gaussian tail by Bernstein inequality (see Equation (4.i) in [2]), hence $\mathbb{E}[\exp(-2G)] < \infty$. Now combining the previous inequality with $D^b G, D^w G \in \mathcal{S}^\infty$, we obtain $\exp(-G) \in \mathbb{D}^{1,2}$, consequently, $V_t = \mathbb{E}_t[\exp(-G)] \in \mathbb{D}^{1,2}$ and

$$D_\theta^k V_t = -\mathbb{E}_t[e^{-G} D_\theta^k G] \quad \text{for all } \theta \leq t \leq T \text{ and } k = b \text{ or } w.$$

Applying Clark-Ocone formula to V_t yields

$$V_t = \mathbb{E}[V_t] + \int_0^t \mathbb{E}_\theta[D_\theta^b V_t] dB_\theta + \int_0^t \mathbb{E}_\theta[D_\theta^w V_t] dW_\theta.$$

On the other hand, $dV_\theta = -V_\theta m_\theta dB_\theta - V_\theta n_\theta dW_\theta$. Therefore $\mathbb{E}_\theta[D_\theta^b V_t] = -V_\theta m_\theta$ and $\mathbb{E}_\theta[D_\theta^w V_t] = -V_\theta n_\theta$, for $\theta \leq t$. Hence,

$$m_\theta = -\frac{\mathbb{E}_\theta[D_\theta^b V_t]}{V_\theta} = \frac{\mathbb{E}_\theta[e^{-G} D_\theta^b G]}{\mathbb{E}_\theta[e^{-G}]} \leq \|D^b G\|_{\mathcal{S}^\infty},$$

which implies $\|m\|_{\mathcal{S}^\infty} \leq \|D^w G\|_{\mathcal{S}^\infty}$, and similarly, $\|n\|_{\mathcal{S}^\infty} \leq \|D^w G\|_{\mathcal{S}^\infty}$.

The statement now follows from Theorem 2 since, for $T < T^*$, where T^* is given in Corollary 2, we have

$$\begin{aligned} \max_i \|(m^i, n^i)\|_{\text{bmo}}^2 &< T^* \max_i (\|m^i\|_{\mathcal{S}^\infty}^2 + \|n^i\|_{\mathcal{S}^\infty}^2) \\ &\leq T^* \max_i (\|D^b G^i\|_{\mathcal{S}^\infty}^2 + \|D^w G^i\|_{\mathcal{S}^\infty}^2) \leq (\epsilon^*)^2. \end{aligned}$$

3.7 An a-priori bmo-estimate

Lemma 9 (An a-priori bmo-estimate for a single BSDE)

Given $\lambda \in \mathcal{P}^2$, let (Y, μ, ν) be a solution of the BSDE

$$dY_t = \mu_t dB_t + \nu_t dW_t + \left(\frac{1}{2}\nu_t^2 - \frac{1}{2}\lambda_t^2 + \mu_t \lambda_t\right) dt, \quad Y_T = \xi.$$

If $Y \in \mathcal{S}^\infty$, then $(\mu, \nu) \in \text{bmo}$.

Proof For $\beta > 1$ and two stopping times $\tau \leq \sigma \in \mathcal{T}$, Itô's formula yields

$$\begin{aligned}
 e^{-\beta Y_\sigma} &\geq e^{-\beta Y_\sigma} - e^{-\beta Y_\tau} = -\beta \int_\tau^\sigma e^{-\beta Y_u} (\mu_u dB_u + \nu_u dW_u) \\
 &\quad - \beta \int_\tau^\sigma e^{-\beta Y_u} \left(\frac{1}{2} \nu_u^2 - \frac{1}{2} \lambda_u^2 + \lambda_u \mu_u \right) du + \frac{1}{2} \beta^2 \int_\tau^\sigma e^{-\beta Y_u} (\mu_u^2 + \nu_u^2) du \\
 &\geq -\beta \int_\tau^\sigma e^{-\beta Y_u} (\mu_u dB_u + \nu_u dW_u) + \frac{1}{2} (\beta^2 - \beta) \int_\tau^\sigma e^{-\beta Y_u} (\mu_u^2 + \nu_u^2) du,
 \end{aligned}$$

where we used the elementary fact that $a^2 - b^2 + 2bc \leq a^2 + c^2$, for all a, b, c . We pick a reducing sequence $\{\sigma_n\}_{n \in \mathbb{N}}$ for the stochastic integral above, project onto \mathcal{F}_τ , and then let $n \rightarrow \infty$ to get

$$\begin{aligned}
 e^{\beta \|Y\|_{S^\infty}} &\geq \frac{1}{2} (\beta^2 - \beta) \mathbb{E}_\tau \left[\int_\tau^T e^{\beta Y_u} (\mu_u^2 + \nu_u^2) du \right] \\
 &\geq \frac{1}{2} (\beta^2 - \beta) e^{-\beta \|Y\|_{S^\infty}} \mathbb{E}_\tau \left[\int_\tau^T (\mu_u^2 + \nu_u^2) dt \right].
 \end{aligned}$$

This implies

$$\mathbb{E}_\tau \left[\int_\tau^T (\mu_u^2 + \nu_u^2) du \right] \leq \frac{2}{\beta^2 - \beta} e^{2\beta \|Y\|_{S^\infty}}.$$

Since the above inequality holds for arbitrary $\tau \in \mathcal{T}$, the statement follows. □

References

1. Anderson, R.M., Raimondo, R.C.: Equilibrium in continuous-time financial markets: Endogenously dynamically complete markets. *Econometrica* **76**(4), 841–907 (2008)
2. Barlow, M., Jacka, S., Yor, M.: Inequality for a pair of processes stopped at a random time. *Proc. London Math. Soc.* **52**(3), 142–172 (1986)
3. Barrieu, P., El Karoui, N.: Monotone stability of quadratic semimartingales with applications to unbounded general quadratic BSDEs. *Ann. Probab.* **41**, 1831–2853 (2013)
4. Borch, K.: The safety loading of reinsurance premiums. *Skand. Aktuarietidskr.* **1960**, 163–184 (1961) (1960)
5. Briand, P., Hu, Y.: BSDE with quadratic growth and unbounded terminal value. *Probab. Theory Relat. Fields* **136**(4), 604–618 (2006)
6. Briand, P., Hu, Y.: Quadratic BSDEs with convex generators and unbounded terminal conditions. *Probab. Theory Relat. Fields* **141**(3-4), 543–567 (2008)
7. Chikvinidze, B., Mania, M.: New proofs of some results on bounded mean oscillation martingales using Backward stochastic differential equations. *J. Theor. Probab.* **27**, 1213–1228 (2014)
8. Choi, J.H., Larsen, K.: Taylor approximation of incomplete Radner equilibrium models (2014). To appear in *Finance Stoch.*
9. Dana, R.A., Pontier, M.: On existence of an Arrow-Radner equilibrium in the case of complete markets. A remark. *Math. Oper. Res.* **17**(1), 148–163 (1992)
10. Debreu, G.: The coefficient of resource utilization. *Econometrica* **19**(3), 273–292 (1951)
11. Delbaen, F., Grandits, P., Rheinländer, T., Samperi, D., Schweizer, M., Stricker, C.: Exponential hedging and entropic penalties. *Math. Finance* **12**(2), 99–123 (2002)
12. Delbaen, F., Hu, Y., Bao, X.: Backward SDEs with superquadratic growth. *Probab. Theory Relat. Fields* **150**, 145–192 (2011)

13. Duffie, D.: Stochastic equilibria: existence, spanning number, and the “no expected financial gain from trade” hypothesis. *Econometrica* **54**(5), 1161–1183 (1986)
14. Duffie, D., Huang, C.F.: Implementing Arrow-Debreu equilibria by continuous trading of few long-lived securities. *Econometrica* **53**(6), 1337–1356 (1985)
15. Elie, R., Briand, P.: A simple constructive approach to quadratic BSDEs with or without delay. *Stoch. Process. Appl.* **123**(8), 2921–2939 (2013)
16. Frei, C., dos Reis, G.: A financial market with interacting investors: does an equilibrium exist? *Math. Financ. Econ.* **4**(3), 161–182 (2011)
17. Hu, Y., Peng, S.: On the comparison theorem for multidimensional BSDEs. *C. R. Math. Acad. Sci. Paris* **343**(2), 135–140 (2006)
18. Karatzas, I., Lakner, P., Lehoczky, J.P., Shreve, S.E.: Equilibrium in a simplified dynamic, stochastic economy with heterogeneous agents. In: *Stochastic analysis*, pp. 245–272. Academic Press, Boston, MA (1991)
19. Karatzas, I., Lehoczky, J.P., Shreve, S.E.: Existence and uniqueness of multi-agent equilibrium in a stochastic, dynamic consumption/investment model. *Math. Oper. Res.* **15**(1), 80–128 (1990)
20. Karatzas, I., Lehoczky, J.P., Shreve, S.E.: Equilibrium models with singular asset prices. *Math. Finance* **1**, 11–29 (1991)
21. Karatzas, I., Shreve, S.E.: *Methods of mathematical finance, Applications of Mathematics (New York)*, vol. 39. Springer-Verlag, New York (1998)
22. Kazamaki, N.: Continuous exponential martingales and BMO, *Lecture Notes in Mathematics*, vol. 1579. Springer-Verlag, Berlin (1994)
23. Kobylanski, M.: Backward stochastic differential equations and partial differential equations with quadratic growth. *Ann. Probab.* **28**(2), 558–602 (2000)
24. Kramkov, D., Pulido, S.: A system of quadratic BSDEs arising in a price impact model (2014). To appear in *Ann. Appl. Probab.*
25. Magill, M., Quinzii, M.: *Theory of Incomplete Markets*, Volume 1. MIT Press, Cambridge and London (1996)
26. Radner, R.: Equilibrium under uncertainty. *Econometrica* **36**(1), 31–58 (1982)
27. Rouge, R., El Karoui, N.: Pricing via utility maximization and entropy. *Math. Finance* **10**(2), 259–276 (2000). INFORMS Applied Probability Conference (Ulm, 1999)
28. Tevzadze, R.: Solvability of backward stochastic differential equations with quadratic growth. *Stochastic Process. Appl.* **118**(3), 503–515 (2008)
29. Zhao, Y.: *Stochastic equilibria in a general class of incomplete brownian market environments*. Ph.D. thesis, The University of Texas at Austin (2012)
30. Žitković, G.: Financial equilibria in the semimartingale setting: complete markets and markets with withdrawal constraints. *Finance Stoch.* **10**(1), 99–119 (2006)
31. Žitković, G.: An example of a stochastic equilibrium with incomplete markets. *Finance Stoch.* **16**(2), 177–206 (2012)