

Endogenous Formation of Limit Order Books: Dynamics Between Trades.*

Roman Gayduk and Sergey Nadtochiy[†]

Current version: June 19, 2017

Original version: May 26, 2016

Abstract

In this work, we present a continuous-time large-population game for modeling market microstructure between two consecutive trades. The proposed modeling framework is inspired by our previous work [23]. In this framework, the Limit Order Book (LOB) arises as an outcome of an equilibrium between multiple agents who have different beliefs about the future demand for the asset. The agents' beliefs may change according to the information they observe, triggering changes in their behavior. We present an example illustrating how the proposed models can be used to quantify the consequences of changes in relevant information signals. If these signals, themselves, depend on the LOB, then, our approach allows one to model the “indirect” market impact (as opposed to the “direct” impact that a market order makes on the LOB, by eliminating certain limit orders). On the mathematical side, we formulate the proposed modeling framework as a continuum-player control-stopping game. We manage to split the equilibrium problem into two parts. The first one is described by a two-dimensional system of Reflected Backward Stochastic Differential Equations (RBSDEs), whose solution components reflect against each other. The second one leads to an infinite-dimensional fixed-point problem for a discontinuous mapping. Both problems are non-standard, and we prove the existence of their solutions in the paper.

1 Introduction

In this paper, we continue the development of an equilibrium-based modeling framework for market microstructure, initiated in [23]. As in [23], we analyze the market microstructure in the context of an auction-style exchange (as most modern exchanges are), in which the participating agents can post limit or market orders. A crucial component of such a market is the *Limit Order Book (LOB)*, which contains all outstanding limit buy and sell orders (time and price prioritized), and whose shape and dynamics represent the liquidity of the market. We are interested in developing a modeling framework in which the shape of the LOB, and its dynamics, arise *endogenously* from the interactions between the agents. This is in contrast to many of the existing results on market microstructure, which assume that the shape and dynamics of the LOB are given exogenously. Among the many advantages of our approach is the possibility of modeling the reaction of the LOB to the changes in a relevant market indicator or in the rules of the exchange.¹

Herein, we extend the discrete time modeling framework proposed in [23] to continuous time, and restrict our analysis to the dynamics of the market between two consecutive trades. The latter simplifies the problem and is justified by the well known empirical fact that most changes in LOB are not due to trades. We manage to establish the existence, and obtain a numerically tractable representation, of an equilibrium in a general continuous time framework, in which the competing agents have different beliefs about the future demand for the asset.

*Partial support from the NSF grant DMS-1411824 is acknowledged by both authors.

[†]Address the correspondence to: Mathematics Department, University of Michigan, 530 Church St, Ann Arbor, MI 48104; sergeyn@umich.edu.

[‡]We would like to thank the anonymous referees whose constructive remarks helped us improve the paper.

¹We refer the reader to [23], whose introduction contains a more detailed explanation of the problems of market microstructure and a motivation for our study.

These beliefs determine the future distribution of the demand, given the (common) information observed thus far. The latter may, e.g., be generated by a relevant signal (or, market indicator). One can view such conditional distributions as the “models” that the market participants use to predict future demand, and which are based on the (commonly observed) relevant market indicators. Given the beliefs, the agents choose their optimal trading strategies (i.e. limit and market orders), aiming to maximize their expected profits, and reach an equilibrium. The modeling framework proposed herein can be used for predicting the reaction of a market to various changes in the relevant indicators. In particular, if the relevant market indicator depends on the LOB, our framework allows one to model the indirect market impact: i.e. how an initial change to the market may cause further changes to it, due to the information revealed by the initial change (as opposed to the direct impact, e.g., made by a market order eliminating a part of the LOB). An extreme example of such indirect impact is called “spoofing”, and it is an illegal activity aimed at manipulating the market. Our model can be used to quantify such indirect market impact, and it can be, ultimately, used to improve the optimal execution algorithms or to test the consequences of “spoofing” activity. We provide a simplistic example illustrating the potential applications of our model in Section 5, although an empirical investigation (including a more careful model specification, and its estimation), which is needed to make any specific conclusions about the actual market behavior, is left for future research.

On the mathematical side, the problem we analyze is the construction of an equilibrium in a control-stopping game with a continuum of players (cf. [3], [47], [7], for more on the general theory of continuum-player games). The main mathematical challenges stem from three sources: the complicated dependence between the individual payoffs and the controls of other players (which lacks the standard convexity and continuity properties), the presence of multiple participants (as compared to a two-player game) and the control-stopping nature of the game. Equilibria in the games with any number of players can often be constructed directly, by means of a system of Partial Differential Equations or a system of (Forward-) Backward Stochastic Differential Equations (BSDEs). However, in the case of multiple players, solving such systems numerically becomes very challenging. In such cases, the description of an equilibrium is, typically, limited to the proof of its existence, which, in turn, is obtained by an abstract fixed-point argument. However, even the latter method presents a challenge in the game considered herein. Namely, the complicated dependence structure between the players’ controls and individual payoffs, along with the control-stopping nature of the game, make it very challenging (or even impossible) to (a) to find a compact set of individual controls, which is sufficiently large to include any maximizer of the objective function, and (b) establish the continuity of the objective.² In order to overcome these challenges, we assume the existence of agents with “extremal beliefs” to split the problem into two parts: a control-stopping game with two players, and a pure control game (without stopping) with a continuum of players. Such a split simplifies our task dramatically, but both resulting problems remain challenging. The first one, concerned with the construction of an equilibrium in a two-player game, leads to a non-standard system of Reflected BSDEs (RBSDEs), whose components *reflect against each other*, and whose generator *lacks to desired regularity*. In Subsection 3.2, we prove the existence of a solution to this system, and, in Section 5, we show how it can be computed in a simple example. The second problem, concerned with the equilibrium in a continuum-player game (without stopping), is formulated as a fixed-point problem, and is solved in Subsection 4.1. This auxiliary game is complicated by the fact that it has a *discontinuous* objective function and does not possess the desired monotonicity properties. Nevertheless, an appropriate “mollification” technique is designed in Subsection 4.1 to construct a solution to the associated fixed-point problem, and, in turn, to describe an equilibrium in the original market microstructure game. One of the computational benefits of the solution method proposed herein is that the aforementioned fixed-point problem can be solved separately for each (t, ω) . In particular, it is not necessary to solve a forward-backward system at each step of the iteration, as it is, for example, done in a typical mean field game (see, e.g., [38], [8]). On the other hand, the local nature of the fixed-point problem causes additional measurability issues, in the proof of the existence result. All these issues are addressed in Subsection 4.1, and the main existence result is stated in Theorem 1, in Section 4.

The literature on market microstructure is vast. Most of the theoretical work is concerned with the problem of a single agent choosing an optimal trading strategy, consisting of limit and/or market orders. The trading environment (e.g. market impact) for this agent is either specified exogenously, or is determined by the agent herself, if she is the designated market-maker. The relevant publications include, among others, [2], [40], [15], [5], [43], [28], [12], [29], [48], [4], [45], [27], [10], and references therein. Nevertheless, none of these works attempt

²Alternatively, one can exploit the monotonicity properties of the objective, to apply a different type of fixed-point theorem. Nevertheless, such monotonicity is also lacking in the present setting.

to explain how the key market characteristics (e.g. the shape and dynamics of LOB) arise from the interaction between multiple market participants. Finally, several recent papers have applied an equilibrium-based approach to the problem of optimal execution (cf. [46], [31]). These papers describe an equilibrium between several agents solving an optimal execution problem, with the LOB (or, the market) against which these agents trade being specified exogenously, rather than being modeled as an output of the equilibrium. The endogenous formation of LOB in an auction-style exchange (i.e. without a designated market-maker) is investigated, e.g., in [42], [21], [25], [9], [37], [44], [18]. However, the models proposed in the aforementioned papers do not aim to represent the mechanics of an auction-style exchange with sufficient precision, which is needed to address the questions we investigate herein.

The paper is organized as follows. Section 2 describes the proposed continuum-player game and defines the associated equilibrium. Section 3 introduces an auxiliary two-player game. The latter is interesting in its own right, but its main purpose is to facilitate the construction of an equilibrium in the continuum-player game. The equilibria in the two-player game are described by a system of RBSDEs, whose solution components reflect against each other and whose generator does not satisfy the global Lipschitz and monotonicity properties. Proposition 1, in Subsection 3.2, provides the existence result for this system, which, to the best of our knowledge, has not been available before. Section 4 completes the construction of an equilibrium in the continuum-player game, stating the main result of the paper, Theorem 1. This section, in particular, describes the mollification technique for solving a fixed-point problem with discontinuity, which appears in the auxiliary continuum-player game. We believe that this method can be applied to other relevant fixed-point problems, with a similar type of discontinuity. Finally, in Section 5, we consider a numerical example, in which we compute the equilibrium strategies and show how our results can be used to study the indirect market impact (illustrated by the particular case of “spoofing”).

2 Modeling framework in continuous time

2.1 Preliminary constructions

We consider an auction-style exchange in which the trades may occur, and the limit orders may be posted, at any time $t \in [0, T]$. The market participants are split into two groups: the external investors, who are “impatient”, in the sense that they only submit market orders and need to execute immediately, and the strategic players, who can submit both market and limit orders, and who are willing to spend time doing so, in order to get a better execution price. In our model, we focus on the strategic players, who we refer to as *agents*, and we model the behavior of the external investors exogenously, via the *external demand*. The external demand for the asset is modeled using three components: the arrival times of the *potential external market orders*, the value of the *potential fundamental price* at these times, and the *elasticity* of the demand. In our previous investigation [23], we have considered a general family of discrete time games for an auction-style exchange, with the exogenous demand process given by a discretization of a (very general) continuous time demand process, over a chosen partition of $[0, T]$. One of the main conclusions of [23] can be, roughly, interpreted as follows: in order for a non-degenerate equilibrium³ to exist in a high-frequency limit (i.e. as the diameter of the partition vanishes), the agents have to be market-neutral – i.e. they should not expect the future fundamental price of the asset to increase or decrease. In other words, the results of [23] seem to imply that it is hopeless to search for an equilibrium in a continuous time game (i.e. with unlimited trading frequency) in which the agents have non-trivial trading signals about the direction of the future moves of the asset price. This may sound very discouraging, however, there is a subtle feature hidden in the setting considered in [23]. Namely, the assumptions of [23] imply that, in the limiting high-frequency regime, the (potential) external market orders arrive with an infinite frequency, while the beliefs of the agents (i.e. their trading signals) satisfy certain continuity properties. In other words, the agents’ signals are assumed to be persistent relative to the trades – they cannot change on the same time scale on which the market orders arrive. It turns out that this assumption is crucial, and, allowing the (potential) external market orders to arrive at a finite frequency, and making the agents’ beliefs be short-lived (i.e. only lasting until the next market order is executed), we can obtain a non-degenerate equilibrium in the continuous time (i.e. unlimited trading frequency) regime. Thus, herein, we model the arrival of

³Degeneracy of an equilibrium is defined formally in [23]. For the discussion presented herein, it suffices to know that degeneracy is an extremal state of the market, and the present work is concerned with the description of the typical (or, normal) states.

the (potential) external market orders via a (rather general) point process, and we assume that the game ends after the first trade occurs.

Let $(\Omega, \tilde{\mathbb{F}} = (\tilde{\mathcal{F}}_t)_{t \in [0, T]}, \mathbb{P})$ be a stochastic basis, satisfying the usual conditions, and supporting a (multidimensional) Brownian motion W and a Poisson random measure N . We assume that the compensator of N is finite on $[0, T] \times \mathbb{R}$ (i.e. N is the jump measure of a compound Poisson process) and that it is absolutely continuous w.r.t. Lebesgue measure in time and space. We denote by \mathbb{F}^W the usual augmented filtration generated by W . We assume that W and N are independent under \mathbb{P} . The arrival times of the potential external market orders and the values of the potential fundamental price at these times are described by a counting random measure M on $[0, T] \times (\mathbb{R} \setminus \{0\})$, defined as

$$M(A) = \int_0^T \int_{\mathbb{R}} \mathbf{1}_A(t, J_t(x)) N(dt, dx),$$

where $J : (t, x) \mapsto J_t(x)$ is a predictable random function (as defined in [30]). We assume that J is adapted to \mathbb{F}^W (in particular, it is independent of N). It is clear that the compensator of M is finite on $[0, T] \times \mathbb{R}$, it is absolutely continuous w.r.t. Lebesgue measure in time and space, and it is adapted to \mathbb{F}^W . Then, it can be represented as $\lambda_t f_t(x) dt dx$, with an \mathbb{R} -valued process $\lambda \geq 0$ and a random function $f : (t, x) \mapsto f_t(x) \geq 0$, progressively measurable and adapted to \mathbb{F}^W , and s.t. $\int_{\mathbb{R}} f_t(x) dx = 1$. Notice that, conditional on \mathcal{F}_T^W , M is a Poisson random measure with the compensator $\lambda_t f_t(x) dt dx$. The t -components of the atoms of M are the arrival times of the potential external market orders, and their x -components represent the values of the potential fundamental price at these times. A positive value of x corresponds to the arrival time of a potential external buy order, and a negative value corresponds to the arrival time of a potential external sell order. More precisely, we define the *fundamental price* process X (or, the *reservation price* process of external investors) as the jump process of M :

$$X_t = \int_{\mathbb{R}} x M(\{t\} \times dx). \quad (1)$$

Note that the X is the jump process of M , but it is *not* a cumulative jump process: it stays at $X_0 = 0$ at all times except the jump times (thus, X can also be interpreted as changes in the fundamental price). We choose $X_0 = 0$ to simplify the notation. In general, any $X_0 \in \mathbb{R}$ is possible, but the only effect it would have on the game is shifting all prices and values by X_0 . To develop a better understanding of the proposed framework, from the economic point of view, it may be useful to think of X_0 as the *last transaction price*, which occurred right before the current game started (although this is not important for the mathematical constructions). The process λ describes the intensity of arrival of the potential external market orders (both buy and sell). The function f_t is the probability density of the value of the potential fundamental price at time t . We refer to f as the density process of the jump sizes. When the jump size of the fundamental price (along with the demand elasticity, defined below) is not enough to trigger a trade, the jump remains “unregistered” by the agents, and the fundamental price returns to zero. The *elasticity* of the external demand for the asset is described by the progressively measurable random field $D : (t, p) \mapsto D_t(p)$, adapted to \mathbb{F}^W . We assume that, a.s., $D_t(\cdot)$ is a strictly decreasing continuous function taking value zero at zero. Then, the total external demand to buy and sell the asset at time t , at the price level p and at all more favorable prices, is equal to

$$D_t^+(p) = \max(0, D_t(p - X_t) \mathbf{1}_{\{X_t > 0\}}), \quad D_t^-(p) = -\min(0, D_t(p - X_t) \mathbf{1}_{\{X_t < 0\}}), \quad (2)$$

respectively.

At any time t , every agent (i.e. strategic player) is allowed to submit a *market order* or a *limit order*. The assumptions made further in the paper make it possible to submit a limit order at such a level that it may never get executed – this, effectively, allows the agents to wait (i.e. do nothing). We do not allow for any time-priority in the limit orders. Instead, we assume that the tick size is zero (the set of possible price levels is \mathbb{R}), and, hence an agent can achieve a priority by posting her order slightly above or below the competing ones (and arbitrarily close to them). The game stops at the terminal time T or at the time when the first trade occurs – whichever one is the earliest. The mechanics of order execution are explained in the next subsection. There is an infinite number of agents, and the inventory of an agent is measured in “shares per unit mass of agents” (see a discussion of this assumption in [23]). We assume that the agents are split into two groups: the ones whose initial inventory s is positive (the *long* agents, typically, indicated with a superscript “ a ”), and those whose initial inventory s is negative

(the *short* agents, indicated with a superscript “*b*”). We assume that the absolute size of each agent’s inventory is the same, $s \in \{-1, 1\}$, and that an agent with inventory s posts orders of size s . These assumptions are motivated by the results of our previous investigation [23], which demonstrate that, in equilibrium, the absolute value of agent’s inventory only scales the size of her orders proportionally, but does not change their type and location.⁴ We also assume that we are given a pair of measurable spaces of *beliefs*, \mathbb{A} and \mathbb{B} , and, for each $\alpha \in \mathbb{A} \cup \mathbb{B}$, there exists a *subjective probability measure* \mathbb{P}^α on $(\Omega, \tilde{\mathbb{F}})$, which is dominated by \mathbb{P} . An agent with beliefs α models the external demand under measure \mathbb{P}^α . The empirical distribution of the agents across beliefs is given by a pair of countably additive finite measures $\mu = (\mu^a, \mu^b)$, on \mathbb{A} and \mathbb{B} , respectively. Note that, because the game stops right after the first market order is executed, the empirical distribution μ remains constant throughout the game. We make the following assumption on the measures $\{\mathbb{P}^\alpha\}$.

Assumption 1. *Under every \mathbb{P}^α , W remains a Brownian motion, and the jump process of N is a process with conditionally independent increments w.r.t. \mathcal{F}_T^W (in the sense of [30]).*

The above assumption holds throughout the paper. It implies that, under every \mathbb{P}^α , X is a process with conditionally independent increments w.r.t. \mathcal{F}_T^W . Using this observation and the absolute continuity of \mathbb{P}^α w.r.t. \mathbb{P} , it is easy to deduce that, under every \mathbb{P}^α , the compensator of the jump measure of X , i.e. of the measure M , is given by

$$\lambda_t^\alpha f_t^\alpha(x) dt dx, \quad (3)$$

with some nonnegative \mathbb{F}^W -adapted λ and \mathbb{F}^W -progressively measurable f^α , s.t. $\int_{\mathbb{R}} f_t^\alpha(x) dx = 1$. The interpretation of λ^α and f^α is the same as the interpretation of λ and f , but under the measure \mathbb{P}^α . Note that we choose not to change the distribution of W under different measures \mathbb{P}^α for a technical reason – in order to avoid Z -dependence in the generator of the associated RBSDE system (43).

It is clear that Assumption 1 is satisfied if $Z_T^\alpha = d\mathbb{P}^\alpha/d\mathbb{P}$ is given by a stochastic exponential of a process that is an integral of \mathbb{F}^W -adapted random function w.r.t. compensated N . Namely,

$$dZ_t^\alpha = Z_{t-}^\alpha \int_{\mathbb{R}} \Gamma_t^\alpha(x) [N(dt, dx) - \lambda_t f_t(x) dt dx],$$

where $\Gamma^\alpha \geq -1$ is \mathbb{F}^W -progressively measurable. The compensator of N under \mathbb{P}^α is obtained by multiplying its compensator under \mathbb{P} by $1 + \Gamma^\alpha$, hence, Assumption 1 is clearly satisfied in this case (cf. [30]). In Section 5, we provide an example of a family of probability measures $\{\mathbb{P}^\alpha\}$ in the above form.

In the proposed setting, the compensator of X under \mathbb{P}^α , given by (3), represents the supply/demand signal used by the agents with beliefs α : in particular, it determines the arrival intensities of external buy and sell orders. Indeed, the value of X is determined uniquely by a path of W and a realization of the random measure N . As the compensator of N may be different under each \mathbb{P}^α , the resulting compensator of X may also vary, however, it always remains adapted to \mathbb{F}^W . Thus, the distribution of X under \mathbb{P}^α is uniquely determined by the choice of $(\lambda^\alpha, f^\alpha)$.⁵ As a result, the agents’ beliefs can be viewed as the “models” they use to map the observed information, given by W , into the predictive signal, given by (3).

2.2 The continuum-player game

Throughout the rest of this paper we, mostly, work with the filtration \mathbb{F}^W , hence, we denote $\mathbb{F} = \mathbb{F}^W$. The state of an agent is $(s, \alpha) \in (\{1\} \times \mathbb{A}) \cup (\{-1\} \times \mathbb{B}) =: \mathbb{S}$. Let us now discuss the controls of the agents and the order execution rules. First, we assume that α , representing the agent’s beliefs, does not change over time.⁶ Therefore, the state process of an agent represents only her inventory, which can only change once (because the game ends after the first trade). The control of every agent is given by a pair of processes $(p, v) = (p_t, v_t)_{t \in [0, T]}$,

⁴Note that the precise setting and the main questions of [23] are not the same as in the present paper. Nevertheless, the two modeling frameworks have many common features. In particular, in both cases, each agent is risk-neutral and infinitesimally small (hence, has no individual impact), which, ultimately, causes their equilibrium strategies to simply scale with the size of initial inventory.

⁵To have a complete model for the external demand, one also needs to know its elasticity D , but the latter is \mathbb{F} -adapted, hence, its distribution is the same under each \mathbb{P}^α .

⁶Note that the conditional distribution of the future demand can change dynamically, according to the new information revealed.

progressively measurable with respect to \mathbb{F} .⁷ The process p takes values in $\mathcal{P}(\mathbb{R})$, the space of probability measures on \mathbb{R} , equipped with the weak topology, while v takes values in \mathbb{R} . The second coordinate, v , determines the time at which the agent decides to submit a market order, and its formal definition is given below. The first coordinate, p_t , indicates the time- t distribution of the agent's limit orders across the price levels. For example, if p_t is a Dirac measure located at x , then, at time t , the agent posts all her limit orders at the price level x . The collection of all limit orders is described by the *Limit Order Book (LOB)*, which is a pair of process $\nu = (\nu_t^a, \nu_t^b)_{t \in [0, T]}$, with values in the finite sigma-additive measures on \mathbb{R} , adapted to \mathbb{F} . Herein, ν_t^a corresponds to the cumulative limit sell orders, and ν_t^b corresponds to the cumulative limit buy orders, posted at time t .⁸ The bid and ask prices at any time $t \in [0, T]$ are given by the random variables

$$p_t^b = Q^+(\nu_t^b), \quad p_t^a = Q^-(\nu_t^a),$$

respectively, where the functions Q^- and Q^+ act on sigma-additive measures κ on \mathbb{R} via

$$Q^+(\kappa) = \sup \text{supp}(\kappa), \quad Q^-(\kappa) = \inf \text{supp}(\kappa). \quad (4)$$

Notice that p_t^b and p_t^a are always well defined as extended random variables, but may take infinite values.

Assume that, at time t , an agent posts a limit sell order at the price level p' . If the demand to buy the asset at or below the price level p' , $D_t^+(p')$, exceeds the amount of all limit sell orders posted below p' at time t , i.e. $D_t^+(p') > \nu_t^a((-\infty, p'))$, then the limit sell order of the agent is executed. Analogous execution rules hold for the limit buy orders. Thus, if an agent follows the limit order strategy p , her limit order is (partially) executed by an external market order at the time

$$T^{p,a} = \inf\{t \in [0, T] : D_t^+(Q^-(p_t)) > \nu_t^a((-\infty, Q^-(p_t)))\},$$

$$T^{p,b} = \inf\{t \in [0, T] : D_t^-(Q^+(p_t)) > \nu_t^b((Q^+(p_t), \infty))\},$$

for the long and short agents, respectively. Let us clarify the meaning of the above formulas. Assume, for simplicity (and only for the sake of this example), that the demand elasticity curve, D , is deterministic. Note that $D_t^+ \equiv D_t^- \equiv 0$ unless X jumps at t . Thus, the above formulas say that a non-zero fraction of agent's limit orders is executed at time t , by an external order, if and only if X jumps at time t , and its jump is sufficiently large, so that the demand at the agent's "best limit order" is higher than the size of all limit orders with higher price priority. The latter, along with continuity of D , ensures that a non-zero fraction of agent's limit orders is executed at this time.

The value of v_t indicates the critical level of the bid or ask price (i.e. a threshold), at which the agent decides to submit a market order. We assume that the size of the agent's market order is equal to her inventory, and it is executed at the bid or ask price available at the time when the order is submitted. Thus, the agent will submit her own market order at the time

$$\tau^{v,a} = \inf\{t \in [0, T] : v_t \leq p_t^b\}, \quad \tau^{v,b} = \inf\{t \in [0, T] : v_t \geq p_t^a\},$$

for the long and short agents, respectively.⁹ The collection of all thresholds v is described by the pair of processes $\theta = (\theta_t^a, \theta_t^b)_{t \in [0, T]}$, with values in the finite sigma-additive measures on \mathbb{R} , adapted to \mathbb{F} .

Remark 1. *The above definitions of the execution times make use of the assumption that each agent is infinitesimally small, and, hence, her order is necessarily executed once the demand reaches it. They also use the following two implicit assumptions: each agent believes that her limit order will be executed first among all orders at the same price level, and her market order will be executed at the best price available. These assumptions and their connection to a finite-player game are discussed in [23].*

⁷It may seem natural to assume that the agents' filtration is enlarged by the information generated by the external trades – i.e. by the jumps of X that lead to a trade. Note that, since the game ends after the first trade, there may only be one such jump. Then, it is easy to see that the predictable filtration of the enlarged filtration, restricted to the time interval until the first trade, is \mathbb{F} itself. Naturally, we require the controls to be predictable.

⁸For convenience, we sometimes refer to ν_t as a "measure", rather than a "pair of measures".

⁹It is clear that, for every stopping time $\tau^{v,a/b}$ with respect to \mathbb{F} , there exists a process v_t , adapted to \mathbb{F} , such that $\tau^{v,a/b}$ has the above representation.

Recall that each agent is infinitesimal, hence, even if she executes a non-zero fraction of her inventory, this may not constitute a trade of non-zero size. We, therefore, define the first “significant” execution time as the first time when a non-zero mass of agents execute a non-zero fraction of their inventory (i.e. when a non-zero total inventory mass is traded). Consider the first significant execution times of external market orders:

$$T^a = \inf\{t \in [0, T] : D_t^+(p_t^a) > 0\}, \quad T^b = \inf\{t \in [0, T] : D_t^-(p_t^b) > 0\}, \quad (5)$$

Similarly, we define the first significant execution times of internal market orders:

$$\tau^a = \inf\{t \in [0, T] : \theta_t^a((-\infty, p_t^b]) > 0\}, \quad \tau^b = \inf\{t \in [0, T] : \theta_t^b([p_t^a, \infty)) > 0\}. \quad (6)$$

Finally, given (ν, X, D) , we define the *clearing prices*:

$$\begin{aligned} \tilde{p}_t^{c,a} &= \sup\{p < Q^+(\nu_t^a) : D_t^+(p) > \nu_t^a((-\infty, p))\}, \quad p_t^{c,a} = \tilde{p}_t^{c,a} \mathbf{1}_{\{\tilde{p}_t^{c,a} \geq p_t^a\}}, \\ \tilde{p}_t^{c,b} &= \inf\{p > Q^-(\nu_t^b) : D_t^-(p) > \nu_t^b((p, \infty))\}, \quad p_t^{c,b} = \tilde{p}_t^{c,b} \mathbf{1}_{\{\tilde{p}_t^{c,b} \leq p_t^b\}}. \end{aligned}$$

For a long agent with strategy (p, v) , the game ends at the time $T^{p,a} \wedge \tau^{v,a} \wedge T \wedge T^a \wedge T^b \wedge \tau^a \wedge \tau^b$ (and similarly for the short agents). If an agent has any inventory left at the end of the game, then it is marked to market.¹⁰ The precise rules for computing the payoff of a long agent, using strategy (p, v) , are described below.

If the game is terminated by an external market order: $T^{p,a} \wedge T^a \wedge T^b < T \wedge \tau^a \wedge \tau^b$ (note that equality is impossible, as the right hand side is predictable and the left hand side is totally inaccessible).

- If $T^{p,a} \wedge T^a < T^b$ (equality is impossible), then the payoff is

$$\int_{-\infty}^{\tilde{p}_t^{c,a}} z p_t(dz) + \int_{\tilde{p}_t^{c,a}}^{\infty} (p_t^{c,a} + p_t^b) p_t(dz), \quad \text{with } t = T^{p,a} \wedge T^a. \quad (7)$$

- If $T^b < T^{p,a} \wedge T^a$, then the payoff is $p_{T^b}^b + p_{T^b}^{c,b}$.

Notice that the remaining inventory of an agent is marked to the bid price shifted by the clearing price. This choice can be (heuristically) interpreted as follows. Assume that, after the trade, a new game starts, with the agents having the same distribution of inventory and the same beliefs about the distribution of future jumps of X (i.e. the same $\{(\lambda^\alpha, f^\alpha)\}$). Then, the only parameter that is different in the new game, as compared to the original one, is the value of X_0 , which, in the new game, becomes equal to the clearing price. As mentioned in the discussion following (1), the new value of X_0 will simply shift all prices and values in the new game by X_0 , hence, the bid price is shifted by the value of clearing price. Finally, it is easy to deduce (and will be shown later in the paper) that it is suboptimal for an agent to post a limit buy order at positive levels. Thus, if an external sell order is executed, the clearing price is non-positive, and, hence, the remaining inventory is marked to the current bid price shifted downwards (the opposite holds if an external buy order is executed).

If the game is terminated by an internal market order: $T \wedge \tau^a \wedge \tau^b < T^{p,a} \wedge T^a \wedge T^b$.


- If $\tau^b < \tau^a \wedge T$ then the payoff is $p_{\tau^b}^a$.
- If $\tau^a \wedge T \leq \tau^b$ then the payoff is $p_{\tau^a \wedge T}^b$.

To explain the above, assume, e.g., that an internal buy order occurs: i.e. $\tau^b < \tau^a \wedge T$. Note that the internal orders are different, because they are predictable. Hence, the long agents can act exactly at the time τ^b and “flock” their limit orders to the best ask price, p^a , to match the market orders from short agents (who initiated the internal buy order). On the other hand, if any of the agents (long or short) do not trade at p^a , they will mark their inventory to

¹⁰There is no canonical way to choose the marking-to-market rules in a setting where agents have no exogenously given valuation of the asset (and we insist on using such a setting, because we think of the agents as “pure speculators”). In particular, other marking rules are possible. Herein, we merely make a choice of marking rules which is economically meaningful.

the bid or ask price shifted by p^a , and, since $p^b \leq 0 \leq p^a$, it is easy to see that it is beneficial for all of them to trade at p^a .¹¹

The following diagram (containing a reference to equation (7)) describes the payoff of a *long* agent:



Similar rules apply to short agents. Formally, given (ν, θ, X, D) , the individual objective of an agent starting at the initial state $(1, \alpha)$ and using the control (p, v) is given by:

$$J^{(\nu, \theta), (p, v)}(1, \alpha) = \mathbb{E}^\alpha \left[\int_{\mathbb{R}} \left(z \mathbf{1}_{\{z \leq \hat{p}_{T^p, a}^{c, a}\}} + \left(p_{T^p, a}^b + p_{T^p, a}^{c, a} \right) \mathbf{1}_{\{z > \hat{p}_{T^p, a}^{c, a}\}} \right) p_{\hat{T}^p, a}(dz) \mathbf{1}_{\{\hat{T}^p, a < T^b \wedge \hat{\tau}^v, a \wedge \tau^b\}} \right. \\ \left. + \left(p_{T^b}^b + p_{T^b}^{c, b} \right) \mathbf{1}_{\{T^b < \hat{T}^p, a \wedge \hat{\tau}^v, a \wedge \tau^b\}} + \left(p_{\tau^b}^a \mathbf{1}_{\{\tau^b < \hat{\tau}^v, a\}} + p_{\hat{\tau}^v, a}^b \mathbf{1}_{\{\tau^b \geq \hat{\tau}^v, a\}} \right) \mathbf{1}_{\{\hat{T}^p, a \wedge T^b > \hat{\tau}^v, a \wedge \tau^b\}} \right] \quad (8)$$

where $\hat{T}^p, a = T \wedge T^{p, a} \wedge T^a$, $\hat{\tau}^v, a = T \wedge \tau^{v, a} \wedge \tau^a$, and we assume that $0 \cdot \infty = 0$. Similarly,

$$J^{(\nu, \theta), (p, v)}(-1, \alpha) = \mathbb{E}^\alpha \left[- \int_{\mathbb{R}} \left(z \mathbf{1}_{\{z \geq \hat{p}_{T^p, b}^{c, b}\}} + \left(p_{T^p, b}^a + p_{T^p, b}^{c, a} \right) \mathbf{1}_{\{z < \hat{p}_{T^p, b}^{c, b}\}} \right) p_{\hat{T}^p, b}(dz) \mathbf{1}_{\{\hat{T}^p, b < T^a \wedge \hat{\tau}^v, b \wedge \tau^a\}} \right. \\ \left. - \left(p_{T^a}^a + p_{T^a}^{c, a} \right) \mathbf{1}_{\{T^a < \hat{T}^p, b \wedge \hat{\tau}^v, b \wedge \tau^a\}} - \left(p_{\tau^a}^b \mathbf{1}_{\{\tau^a < \hat{\tau}^v, b\}} + p_{\hat{\tau}^v, b}^a \mathbf{1}_{\{\tau^a \geq \hat{\tau}^v, b\}} \right) \mathbf{1}_{\{\hat{T}^p, b \wedge T^a > \hat{\tau}^v, b \wedge \tau^a\}} \right] \quad (9)$$

where $\hat{T}^p, b = T \wedge T^{p, b} \wedge T^b$, $\hat{\tau}^v, b = T \wedge \tau^{v, b} \wedge \tau^b$. Every agent aims to maximize her objective. The above objectives may seem convoluted – this is because they are meant to provide a close approximation of the real-world execution rules and marking to market. In the next subsection, we establish a more transparent representation of the objectives.

In the following definitions, we assume that a stochastic basis, a Brownian motion W , a random measure M , a random field D , spaces \mathbb{A} and \mathbb{B} , an associated set of measures $\{\mathbb{P}^\alpha\}_{\alpha \in \mathbb{A} \cup \mathbb{B}}$, and the empirical distribution μ , are fixed and satisfy the assumptions made earlier in this section. (Nevertheless, it is shown in Subsection 2.3 that the input $(M, \{\mathbb{P}^\alpha\})$ can be replaced by the agents' signals $\{\lambda^\alpha, f^\alpha\}$.)

Definition 1. For a given market (ν, θ) and a state $(s, \alpha) \in \mathbb{S}$, a pair of \mathbb{F} -progressively measurable processes (p, v) is an **admissible control**, if the positive part of the expression inside the expectation in (8) (if $s = 1$) or (9) (if $s = -1$) has a finite expectation under \mathbb{P}^α .

Definition 2. For a given market (ν, θ) and state $(s, \alpha) \in \mathbb{S}$, we call an admissible control (p, v) **optimal** if

$$J^{(\nu, \theta), (p, v)}(s, \alpha) \geq J^{(\nu, \theta), (p', v')}(s, \alpha)$$

\mathbb{P} -a.s., for any admissible control (p', v') .

In the above, we make the standard assumption of games with a *continuum of players*: each agent is too small to affect the distribution of cumulative controls (described by ν) when she changes her control. Next, we define Nash equilibrium in the proposed game.

¹¹Of course, in practice, not all agents will act at the same time: only a fraction of them will submit the internal market orders at the end of the game, the others will move on to the next game, with updated X_0 . However, such “flocking” of agents at the end of the game (provided the game ends with an internal market order) is consistent with the empirical observation of “clustering trades”.

Definition 3. A given market (ν, θ) and a pair of \mathbb{F} -progressively measurable random fields $(p, v) : \Omega \times [0, T] \times \mathbb{S} \rightarrow \mathcal{P}(\mathbb{R}) \times \mathbb{R}$ form an **equilibrium**, if

1. for μ -a.e. $(s, \alpha) \in \mathbb{S}$, $(p(s, \alpha), v(s, \alpha))$ is an optimal control for (ν, θ) and (s, α) ,
2. and the following holds \mathbb{P} -a.s., for any $t < \bar{T} := T \wedge T^a \wedge T^b \wedge \tau^a \wedge \tau^b$ and any $x \in \mathbb{R}$:

$$\nu_t^a((-\infty, x]) = \int_{\mathbb{A}} p_t(1, \alpha; (-\infty, x]) \mu^a(d\alpha), \quad \nu_t^b((-\infty, x]) = \int_{\mathbb{B}} p_t(-1, \alpha; (-\infty, x]) \mu^b(d\alpha), \quad (10)$$

$$\theta_t^a((-\infty, x]) = \int_{\mathbb{A}} \mathbf{1}_{\{v_t(1, \alpha) \leq x\}} \mu^a(d\alpha), \quad \theta_t^b((-\infty, x]) = \int_{\mathbb{B}} \mathbf{1}_{\{v_t(-1, \alpha) \leq x\}} \mu^b(d\alpha). \quad (11)$$

Note that a trivial equilibrium, in which all agents stop immediately, is always possible. However, such equilibrium, clearly, is not sufficient for modeling purposes, and the existence of other, non-trivial, equilibria is far from obvious. In the remainder of this paper, we use an auxiliary two-player game (cf. Section 3) to identify a class of more realistic potential equilibria, in which the end time of the game is determined uniquely by the solution of an associated RBSDE system (cf. (44)), and we prove the existence of equilibrium in this class, in Theorem 1. Even though it is possible to construct models in which the resulting equilibrium is still trivial (i.e. the end time of the game is zero), this is not the case in general, as confirmed by the example in Section 5.

Remark 2. In the above definition, it is implicitly assumed that the empirical measure of the agents' states remains constant in time until the game is over for all players. This is, indeed, the case, if the equilibrium is such that, \mathbb{P} -a.s., for all $t < \bar{T}$, we have:

$$\mu \circ ((s, \alpha) \mapsto S_t(s, \alpha))^{-1} = \mu, \quad (12)$$

with

$$S_t(1, \alpha) = \mathbf{1}_{[0, T^{p(1, \alpha), a} \wedge \tau^{v(1, \alpha), a})}(t), \quad \text{and} \quad S_t(-1, \alpha) = -\mathbf{1}_{[0, T^{p(-1, \alpha), b} \wedge \tau^{v(-1, \alpha), b})}(t).$$

The condition (12) may fail if a non-zero mass of agents manages to execute their orders strictly before \bar{T} : i.e. if $T^{p(1, \alpha), a} \wedge \tau^{v(1, \alpha), a} < \bar{T}$ for a set of α with positive μ^a -measure, or $T^{p(-1, \alpha), b} \wedge \tau^{v(-1, \alpha), b} < \bar{T}$ for a set of α with positive μ^b -measure. The latter cannot occur due to external market orders, because they only arrive at a finite number of times and, before $T^a \wedge T^b \geq \bar{T}$, only a zero mass of agents can execute their limit orders against any such market order (cf. (5)). It is also true that, at any time t , before $\tau^a \wedge \tau^b \geq \bar{T}$, only a zero mass of agents can execute their internal market orders (cf. (6)). However, the set of such times t may be uncountable. Therefore, to ensure that μ remains constant and, hence, (12) holds, it suffices to consider only the equilibria satisfying, \mathbb{P} -a.s., for all t , except, possibly, a countable set:

$$v_t(1, \alpha) \geq v_t^a := Q^-(\theta_t^a), \quad v_t(-1, \alpha) \leq v_t^b := Q^+(\theta_t^b), \quad \forall \alpha \in \mathbb{A} \cup \mathbb{B}.$$

In the subsequent sections, we construct such an equilibrium.

2.3 Representation of the objective

In this section, we provide an equivalent representation of the objective of the agents, which makes it more tractable and more convenient for the analysis that follows. In addition, it shows that the main input parameters for the proposed equilibrium problem are the signals $\{(\lambda^\alpha, f^\alpha)\}_{\alpha \in \mathbb{A} \cup \mathbb{B}}$, forming the compensators of X under $\{\mathbb{P}^\alpha\}$, and the demand elasticity D (the latter is independent of α and, in many realistic models, can be deterministic). In particular, there is no need to keep track of the random measure N and the probability measures $\{\mathbb{P}^\alpha\}$ – they are only needed to show that the present setting fits within the standard framework for games with heterogeneous beliefs. The desired representation is derived following standard arguments, making use of the independence of the driving Poisson measure N and the Brownian motion W . First, we introduce new notation that will be used throughout the paper. For any $\alpha \in \mathbb{A} \cup \mathbb{B}$, $t \in [0, T]$, $p, x, y \in \mathbb{R}$ and $\kappa \in \mathcal{P}(\mathbb{R})$, we define the *instantaneous filling rates* for limit orders at levels x and y :

$$F_t^{+, \alpha}(x) = \int_{x \vee 0}^{\infty} f_t^\alpha(u) du, \quad F_t^{-, \alpha}(y) = \int_{-\infty}^{y \wedge 0} f_t^\alpha(u) du, \quad c_t^\alpha(x, y) = \lambda_t^\alpha (F_t^{-, \alpha}(y) + F_t^{+, \alpha}(x)). \quad (13)$$

Next, we define the *clearing price* as a function of the fundamental price x :

$$l_t^{c,a}(x) = \sup \{p < Q^+(\nu_t^a) : D_t(p - x) > \nu_t^a((-\infty, p))\}, \quad (14)$$

$$l_t^{c,b}(x) = \inf \{p > Q^-(\nu_t^b) : -D_t(p - x) > \nu_t^b((p, \infty))\}. \quad (15)$$

Notice that, if X has a positive jump at time t , then the clearing price at time t is given by $\tilde{p}_t^{c,a} = l_t^{c,a}(X_t)$. Similarly, if X has a negative jump at time t , then $\tilde{p}_t^{c,b} = l_t^{c,b}(X_t)$. Finally, we introduce the *instantaneous reward rates* from executed limit orders, distributed according to κ , with the bid and ask prices y and x :

$$h_t^{\alpha,a}(\kappa, x, y) = \lambda_t^\alpha \int_{(Q^-(\kappa) \wedge x) \vee 0}^\infty f_t^\alpha(u) \left[\int_{-\infty}^{l_t^{c,a}(u)} z \kappa(dz) + (y + l_t^{c,a}(u) \mathbf{1}_{\{l_t^{c,a}(u) \geq x\}}) \kappa((l_t^{c,a}(u), \infty)) \right] du \\ + \lambda_t^\alpha \int_{-\infty}^{y \wedge 0} f_t^\alpha(u) (y + l_t^{c,b}(u)) du, \quad (16)$$

$$h_t^{\alpha,b}(\kappa, x, y) = \lambda_t^\alpha \int_{-\infty}^{(Q^+(\kappa) \vee y) \wedge 0} f_t^\alpha(u) \left[\int_{l_t^{c,b}(u)}^\infty z \kappa(dz) + (x + l_t^{c,b}(u) \mathbf{1}_{\{l_t^{c,b}(u) \leq y\}}) \kappa((-\infty, l_t^{c,b}(u))) \right] du \\ + \lambda_t^\alpha \int_{x \vee 0}^\infty f_t^\alpha(u) (x + l_t^{c,a}(u)) du. \quad (17)$$

Using the above notation, we can obtain a simplified expression for the objective, given in the following lemma. Note that the expectation in this representation is taken under the reference measure, and the objective depends only on the cumulative actions (ν, θ) and on $(\{\lambda^\alpha, f^\alpha\}, D)$ (as the expressions in (13)–(17) depend only on $(\{\lambda^\alpha, f^\alpha\}, D)$).

Lemma 1. *Let Assumption 1 hold. Given a market (ν, θ) , for any $\alpha \in \mathbb{A} \cup \mathbb{B}$ and any admissible strategy (p, v) , we have:*

$$J^{(\nu, \theta), (p, v)}(1, \alpha) = \mathbb{E} \left[\int_0^{\hat{\tau}^{v,a} \wedge \tau^b} \exp \left(- \int_0^s c_u^\alpha (p_u^a \wedge Q^-(p_u), p_u^b) du \right) h_s^{\alpha,a}(p_s, p_s^a, p_s^b) ds \right] \quad (18)$$

$$+ \exp \left(- \int_0^{\hat{\tau}^{v,a} \wedge \tau^b} c_u^\alpha (p_u^a \wedge Q^-(p_u), p_u^b) du \right) (p_{\tau^b}^a \mathbf{1}_{\{\tau^b < \hat{\tau}^{v,a}\}} + p_{\hat{\tau}^{v,a}}^b \mathbf{1}_{\{\tau^b \geq \hat{\tau}^{v,a}\}}) \Big],$$

$$J^{(\nu, \theta), (p, v)}(-1, \alpha) = -\mathbb{E} \left[\int_0^{\hat{\tau}^{v,b} \wedge \tau^a} \exp \left(- \int_0^s c_u^\alpha (p_u^a, p_u^b \vee Q^+(p_u)) du \right) h_s^{\alpha,b}(p_s, p_s^a, p_s^b) ds \right] \quad (19)$$

$$+ \exp \left(- \int_0^{\hat{\tau}^{v,b} \wedge \tau^a} c_u^\alpha (p_u^a, p_u^b \vee Q^+(p_u)) du \right) (p_{\tau^a}^b \mathbf{1}_{\{\tau^a < \hat{\tau}^{v,b}\}} + p_{\hat{\tau}^{v,b}}^a \mathbf{1}_{\{\tau^a \geq \hat{\tau}^{v,b}\}}) \Big],$$

where $\hat{\tau}^{v,a} = T \wedge \tau^{v,a} \wedge \tau^a$, $\hat{\tau}^{v,b} = T \wedge \tau^{v,b} \wedge \tau^b$ and the expectations are taken under \mathbb{P} .

Proof: The proof follows easily by conditioning on W . Notice that, conditional on \mathcal{F}_T , M is a Poisson random measure, with the deterministic compensator $\lambda_t^\alpha f_t^\alpha(x) dt dx$, which is finite on $[0, T] \times \mathbb{R}$. Recall also that $D, \nu, \theta, p, v, p^a, p^b, \tau^{v,a}, \tau^{v,b}, \tau^a, \tau^b$, and all the random functions defined above the lemma, are adapted to \mathbb{F} . Conditional on \mathcal{F}_T , they become deterministic functions of time. Recall the fundamental price process, $X_t = \int_{\mathbb{R}} x M(\{t\} \times dx)$, and introduce

$$Y_t = X_t \left(\mathbf{1}_{\{X_t > (p_t^a \wedge Q^-(p_t)) \vee 0\}} + \mathbf{1}_{\{X_t < p_t^b \wedge 0\}} \right).$$

Notice that $\hat{T}^{p,a}$ is the time of the first positive jump of Y_t , and T^b is the time of its first negative jump. Notice also that, conditional on \mathcal{F}_T , the clearing price $\tilde{p}_t^{c,a}$ becomes a deterministic function of t and Y_t : $\tilde{p}_t^{c,a} = l_t^{c,a}(Y_t)$. Thus, conditional on \mathcal{F}_T , the expression inside the expectation in (8) becomes a function of the time and size of

the first jump of Y . Conditional on \mathcal{F}_T , X is the jump process of a Poisson random measure with the compensator $\lambda_t^\alpha f_t^\alpha(u) du dt$. It is also clear that, conditional on \mathcal{F}_T , Y is the jump process of a non-homogeneous compound Poisson process with intensity $c_t^\alpha (p_t^a \wedge Q^-(p_t), p_t^b)$, and with the distribution of jump sizes at time t given by

$$\frac{\lambda_t^\alpha f_t^\alpha(x)}{c_t^\alpha (p_t^a \wedge Q^-(p_t), p_t^b)} \left(\mathbf{1}_{\{x \leq p_t^b \wedge 0\}} + \mathbf{1}_{\{x \geq (p_t^a \wedge Q^-(p_t)) \vee 0\}} \right) dx.$$

A standard computation, then, yields (18). The equation (19) is derived similarly. The expectations in (18) and (19) are taken under \mathbb{P} , because the expressions inside the expectations are adapted to $\mathbb{F} = \mathbb{F}^W$, and W has the same distribution under \mathbb{P} and \mathbb{P}^α . ■

3 A two-player game

In this section, we consider an auxiliary non-zero-sum two-player control-stopping game. It is related to the continuum-player game, but the precise connection will be established in the subsequent sections. We refer the reader to [33], [34], and the references therein, for more on non-zero-sum two-player control-stopping games.¹² It is worth mentioning, however, that the present game does not fall within any of the classes considered before. A more detailed description of this class of games is carried out in our forthcoming work [22].

Assume that all the probabilistic constructions made in Subsection 2.1 are in place. Namely, we are given a stochastic basis, with a Brownian motion W , a Poisson measure N , a counting random measure M , a family of probability measures $\{\mathbb{P}^\alpha\}$, and with the demand elasticity process D , as described in Section 2. We assume that Assumption 1 holds. Assume, in addition, that $\mathbb{A} = \{\alpha^0\}$ and $\mathbb{B} = \{\beta^0\}$. Consider a two-player game, in which the first (long) player starts with the initial inventory 1 and has beliefs α^0 , and the second (short) player starts with the initial inventory -1 and has beliefs β^0 . The game proceeds according to the rules similar to those described in the previous section: each agent can post limit orders on the respective side of the book, or can terminate the game by submitting a market order. The execution of limit orders against the external market orders occurs in exactly the same way as described in the previous section. However, herein, at any given time, each agent is only allowed to post limit orders at a single location (i.e. the control p_t is a Dirac measure). In addition, the main difference between the present game and the one defined in the previous section is that, herein, each player has a non-zero mass and, hence, can affect the LOB. In fact, since there is only one player on each side of the book, the LOB is given by a combination of two Dirac measures: $\nu_t^a = \delta_{p_t^a}$, $\nu_t^b = \delta_{p_t^b}$, controlled by the locations of the players' limit orders: p^a for the long agent, and p^b for the short one. Clearly, p^a also coincides with the ask price, and p^b is the bid price. Note that each of these prices is now controlled by a single agent, which is not the case in the original game described in the previous section. The same is true for the stopping thresholds: θ^a and θ^b are given by Dirac measures, and the locations of these measures correspond to the thresholds v^a and v^b used by the long and short agents, respectively. In this new game (due to its simplicity), it turns out to be more convenient to work with the associated stopping times τ^a and τ^b . In fact, we will further constraint the agents' controls, so that $\tau^a = \tau^b =: \tau$ and $p_\tau^a = p_\tau^b = \bar{p}_\tau$. The meaning behind these constraints is clear: every agent assumes that the counterparty will execute a market order at exactly the same time as she does, and that these orders are executed at the same price. Taking into account the above considerations, we transform (8) into the objective of a long player:

$$\tilde{J}^{a,(p^b,\bar{p}),(p,\tau)} = \mathbb{E}^{\alpha^0} [p_{T^p,a} \mathbf{1}_{\{T^p,a < T^b \wedge \tau\}} + 2p_{T^b}^b \mathbf{1}_{\{T^b < T^p,a \wedge \tau\}} + \bar{p}_\tau \mathbf{1}_{\{T^p,a \wedge T^b > \tau\}}], \quad (20)$$

where p^b , \bar{p} and p are \mathbb{R} -valued \mathbb{F} -adapted processes, τ is a stopping time with values in $[0, T]$, and

$$T^b = \inf\{t \in [0, T] : X_t < p_t^b\}, \quad T^{p,a} = \inf\{t \in [0, T] : X_t > p_t\}, \quad X_t = M(\{t\} \times \mathbb{R}).$$

Similarly, for the short agents,

$$\tilde{J}^{b,(p^a,\bar{p}),(p,\tau)} = -\mathbb{E}^{\alpha^0} [p_{T^p,b} \mathbf{1}_{\{T^p,b < T^a \wedge \tau\}} + 2p_{T^a}^a \mathbf{1}_{\{T^a < T^p,b \wedge \tau\}} + \bar{p}_\tau \mathbf{1}_{\{T^p,b \wedge T^a > \tau\}}], \quad (21)$$

¹²See, e.g., [19], [36], [17], [6], and the references therein, for the related classical Dynkin games, which are zero-sum and stopping-only.

where p^a , \bar{p} and p are \mathbb{R} -valued \mathbb{F} -adapted processes, τ is a stopping time with values in $[0, T]$, and

$$T^a = \inf\{t \in [0, T] : X_t > p_t^a\}, \quad T^{p,b} = \inf\{t \in [0, T] : X_t < p_t\}.$$

Using Lemma 1, we deduce the following form of the objective functions

$$\begin{aligned} \tilde{J}^{a,(p^b,\bar{p}),(p,\tau)} &= \mathbb{E} \left[\int_0^\tau \exp \left(- \int_0^s c_u^{\alpha^0}(p_u, p_u^b) du \right) g_s^a(p_s, p_s^b) ds \right. \\ &\quad \left. + \exp \left(- \int_0^\tau c_u^{\alpha^0}(p_u, p_u^b) du \right) \bar{p}_\tau \right], \end{aligned} \quad (22)$$

where c^α is defined in (13) and

$$g_t^a(x, y) = \lambda_t^{\alpha^0} \left(2yF_t^{\alpha^0,-}(y) + xF_t^{\alpha^0,+}(x) \right). \quad (23)$$

Similarly,

$$\begin{aligned} \tilde{J}^{b,(p^a,\bar{p}),(p,\tau)} &= -\mathbb{E} \left[\int_0^\tau \exp \left(- \int_0^s c_u^{\beta^0}(p_u^a, p_u) du \right) g_s^b(p_s^a, p_s) ds \right. \\ &\quad \left. + \exp \left(- \int_0^\tau c_u^{\beta^0}(p_u^a, p_u) du \right) \bar{p}_\tau \right], \end{aligned} \quad (24)$$

where

$$g_t^b(x, y) = \lambda_t^{\beta^0} \left(yF_t^{\beta^0,-}(y) + 2xF_t^{\beta^0,+}(x) \right). \quad (25)$$

To ensure that the above expressions are well defined, and to analyze the equilibrium in a two-player game, we need to make the following assumptions.

Assumption 2. *There exists a constant $C' > 0$, s.t., \mathbb{P} -a.s., $|\lambda_t^\alpha|, |f_t^\alpha(x)| \leq C'$, for all $\alpha \in \mathbb{A} \cup \mathbb{B}$, $t \in [0, T]$ and $x \in \mathbb{R}$.*

We also assume that the possible price jump sizes are bounded.

Assumption 3. *There exists a constant $C_p > 0$, s.t., \mathbb{P} -a.s., $\text{supp}(f_t^\alpha) \subseteq [-C_p, C_p]$, for all $\alpha \in \mathbb{A} \cup \mathbb{B}$ and $t \in [0, T]$.*

Denote by \mathbb{S}^2 the set of continuous \mathbb{F} -adapted processes Y , such that $\sup_{0 \leq t \leq T} |Y_t| \in \mathbb{L}^2$. We say that the terminal execution price \bar{p} is admissible if $\bar{p} \in \mathbb{S}^2$. A control (p, τ) is admissible if p is \mathbb{F} -progressively measurable, satisfying, \mathbb{P} -a.s., $|p_t| \leq C_p$ for all $t \in [0, T]$, and τ is \mathbb{F} -stopping time. Next, we introduce the notions of optimality and equilibrium in the two-player game – they are analogous to Definitions 2–3.

Definition 4. *For a given admissible (p^b, \bar{p}) , we call an admissible control (p, τ) optimal for the long agent if*

$$\tilde{J}^{a,(p^b,\bar{p}),(p,\tau)} \geq \tilde{J}^{a,(p^b,\bar{p}),(p',\tau')},$$

for any admissible control (p', τ') . Similarly, for a given admissible (p^a, \bar{p}) , we call an admissible control (p, τ) optimal for the short agent if

$$\tilde{J}^{b,(p^a,\bar{p}),(p,\tau)} \geq \tilde{J}^{b,(p^a,\bar{p}),(p',\tau')},$$

for any admissible control (p', τ') .

Definition 5. *A combination $(p^a, p^b, \tau, \bar{p})$ is an equilibrium in the two-player game, if it is admissible and, given (p^b, \bar{p}) , the control (p^a, τ) is optimal for the long agent, while, given (p^a, \bar{p}) , the control (p^b, τ) is optimal for the short agent.*

In the next subsection, we characterize the equilibrium strategies via a system of Reflected Backward Stochastic Differential Equations (RBSDEs).

3.1 Characterizing the equilibria via a system of RBSDEs

The next assumptions are used to guarantee the uniqueness and regularity of the optimal control of an agent.

Assumption 4. \mathbb{P} -a.s., for any $\alpha \in \mathbb{A} \cup \mathbb{B}$ and $t \in [0, T]$, $f_t^\alpha(\cdot)$ is continuous in the interior of its support, $f_t^\alpha(0) = 0$, and $0 < F_t^{+, \alpha}(0) < 1$.

Assumption 5. \mathbb{P} -a.s., for any $\alpha \in \mathbb{A}$ and $t \in [0, T]$, $F_t^{+, \alpha}(\cdot)/f_t^\alpha(\cdot)$ is a decreasing function in the interior of $\text{supp}(f_t^\alpha) \cap \mathbb{R}_+$, vanishing at the right end of the interval. Similarly, \mathbb{P} -a.s., for any $\beta \in \mathbb{B}$ and $t \in [0, T]$, $F_t^{-, \beta}(\cdot)/f_t^\beta(\cdot)$ is an increasing function in the interior of $\text{supp}(f_t^\beta) \cap \mathbb{R}_-$, vanishing at the left end of the interval.

Remark 3. The monotonicity of $F_t^{+, \alpha}(\cdot)/f_t^\alpha(\cdot)$, for example, is implied by the log-concavity of the distribution of positive jumps (similarly, for the negative jumps). Instead of requiring that $F_t^{+, \alpha}(\cdot)/f_t^\alpha(\cdot)$ is decreasing, it suffices to assume that its growth rate is bounded from above by $1 - \varepsilon$, for a constant $\varepsilon > 0$ independent of (t, ω) .

To prove the existence of a solution to a system of RBSDEs characterizing the equilibria in a two-player game, we also need to assume that “the range of beliefs is relatively bounded”.

Assumption 6. There exists a constant $C > 0$, s.t., \mathbb{P} -a.s.:

$$\frac{1}{C} \leq \left| \frac{\lambda_t^{\alpha^0}}{\lambda_t^{\beta^0}} \right| \leq C, \quad \frac{1}{C} \leq \left| \frac{f_t^{\alpha^0}(x)}{f_t^{\beta^0}(x)} \right| \leq C, \quad \forall x \in \mathbb{R} \quad t \in [0, T].$$

First we analyze the individual optimization problem of an agent, taking the actions of the counterparty as given. Assume that we are given a process $\bar{p} \in \mathbb{S}^2$ and progressively measurable (p^a, p^b) , such that \mathbb{P} -a.s., $|p_t^a|, |p_t^b| \leq C_p, \forall t \in [0, T]$. Let us introduce the value functions of the agents:

$$\begin{aligned} V_t^a = \text{ess sup}_{\tau \in \mathcal{T}_t, p} \mathbb{E} \Big[& \int_t^\tau \exp \left(- \int_t^s c_u^{\alpha^0}(p_u, p_u^b) du \right) g_s^a(p_s, p_s^b) ds \\ & + \exp \left(- \int_t^\tau c_u^{\alpha^0}(p_u, p_u^b) du \right) \bar{p}_\tau \Big| \mathcal{F}_t \Big], \end{aligned} \quad (26)$$

$$\begin{aligned} V_t^b = \text{ess inf}_{\tau \in \mathcal{T}_t, p} \mathbb{E} \Big[& \int_t^\tau \exp \left(- \int_t^s c_u^{\beta^0}(p_u^a, p_u) du \right) g_s^b(p_s^a, p_s) ds \\ & + \exp \left(- \int_t^\tau c_u^{\beta^0}(p_u^a, p_u) du \right) \bar{p}_\tau \Big| \mathcal{F}_t \Big], \end{aligned} \quad (27)$$

where \mathcal{T}_t is the set of \mathbb{F} -stopping times with values in $[t, T]$, p is any \mathbb{F} -progressively measurable process, with $|p| \leq C_p$, and c^α, g^a and g^b are defined in (13), (23) and (25). In addition, we introduce the following random functions:

$$\begin{aligned} \mathcal{G}_t^{a,x}(y, z) &= -c_t^{\alpha^0}(x, z)y + g_t^a(x, z), \quad x, y \in \mathbb{R}, \\ \mathcal{G}_t^a(y, z) &= \sup_{x \in \mathbb{R}} \mathcal{G}_t^{a,x}(y, z) = -c_t^{\alpha^0}(P_t^a(y), z)y + g_t^a(P_t^a(y), z), \quad y \in \mathbb{R}, \end{aligned}$$

where P_t^a provides the optimal price location at the ask side, given in a feedback form:

$$P_t^a(y) = \inf_{p \in \mathbb{R}} \arg \max (p - y) F_t^{+, \alpha^0}(p), \quad y \in \mathbb{R}. \quad (28)$$

Similarly, for any admissible p^a , we define

$$P_t^b(y) = \sup_{p \in \mathbb{R}} \arg \max (y - p) F_t^{-, \beta^0}(p), \quad y \in \mathbb{R}, \quad (29)$$

$$\mathcal{G}_t^b(z, y) = -c_t^{\beta^0}(z, P_t^b(y))y + g_t^b(z, P_t^b(y)), \quad y \in \mathbb{R}.$$

The value of $P_t^a(y)$ can be described as the unique nonnegative solution p of

$$p - y = F_t^{+, \alpha^0}(p) / f_t^{\alpha^0}(p), \quad (30)$$

unless y is too large, in which case $P_t^a(y)$ is the upper boundary of the support of $f_t^{\alpha^0}$, or too small, in which case $P_t^a(y) = 0$. Similarly, $P_t^b(y)$ is the unique non-positive solution p of

$$y - p = F_t^{-, \beta^0}(p) / f_t^{\beta^0}(p), \quad (31)$$

or the lower boundary of the support of $f_t^{\beta^0}$, if y is too small, or zero, if y is too large.

Lemma 2. *Let Assumptions 1–5 hold. Then, the random functions P^a and P^b are progressively measurable and satisfy, \mathbb{P} -a.s., for all $t \in [0, T]$:*

$$0 \leq P_t^a(y) \leq C_p, \quad -C_p \leq P_t^b(y) \leq 0, \quad P_t^a(y) \geq y, \quad P_t^b(y) \leq y, \quad \forall y \in \mathbb{R},$$

and, in addition, $P_t^a(\cdot)$ and $P_t^b(\cdot)$ are non-decreasing and 1-Lipschitz.

Proof: The progressive measurability property and the above inequalities follow directly from Assumptions 2–4. The monotonicity and 1-Lipschitz property follow from Assumption 5 and the representations (30)–(31). ■

The above lemma, along with Assumptions 2–4, implies that, for any admissible (p, p^b, \bar{p}) , $\mathcal{G}_t^a(0, p_t^b)$ and $\mathcal{G}_t^{a, p_t^b}(0, p_t^b)$ are bounded processes, and that $\mathcal{G}_t^a(y, p_t^b)$ and $\mathcal{G}_t^{a, p_t^b}(y, p_t^b)$ are Lipschitz in y , uniformly over a.e. (t, ω) . This allows us to use Proposition 7.1 from [35], to show that, for any admissible (p, p^b, \bar{p}) , the process Y , which is a continuous modification of

$$Y_t := \hat{J}_t^{a, (p^b, \bar{p}), p} = \text{ess sup}_{\tau \in \mathcal{T}_t} \mathbb{E} \left[\int_t^\tau \exp \left(- \int_t^s c_u^{\alpha^0}(p_u, p_u^b) du \right) g_s^a(p_s, p_s^b) ds \right. \\ \left. + \exp \left(- \int_t^\tau c_u^{\alpha^0}(p_u, p_u^b) du \right) \bar{p}_\tau \middle| \mathcal{F}_t \right],$$

is the unique \mathbb{S}^2 solution of the affine RBSDE,

$$-dY_t = \mathcal{G}_t^{a, p_t^b}(Y_t, p_t^b) dt - Z_t dW_t + dK_t \quad 0 \leq t \leq T \quad (32)$$

$$Y_t \geq \bar{p}_t \quad 0 \leq t \leq T, \quad \int_0^T (Y_t - \bar{p}_t) dK_t = 0 \quad (33)$$

$$Y_T = \bar{p}_T, \quad (34)$$

where Z is a progressively measurable square-integrable (multidimensional) process, $K \in \mathbb{S}^2$ is increasing and satisfies $K_0 = 0$. Similarly, the existence results from [35] imply that

$$-dY_t = \mathcal{G}_t^a(Y_t, p_t^b) dt - Z_t dW_t + dK_t \quad 0 \leq t \leq T \quad (35)$$

$$Y_t \geq \bar{p}_t \quad 0 \leq t \leq T \quad \int_0^T (Y_t - \bar{p}_t) dK_t = 0 \quad (36)$$

$$Y_T = \bar{p}_T \quad (37)$$

has a unique solution (Y_t, Z_t, K_t) . Then, Theorem 7.2 in [35] implies that Y is a continuous modification of V^a , and that $p_t^a = P_t^a(Y_t)$ and $\tau^a = \inf\{s \geq 0: Y_s = \bar{p}_s\}$ form an optimal control for the long agent. Similarly, for a given admissible (p^a, \bar{p}) , there exists a unique solution (Y_t, Z_t, K_t) to

$$-dY_t = \mathcal{G}_t^b(p_t^a, Y_t) dt - Z_t dW_t - dK_t \quad 0 \leq t \leq T \quad (38)$$

$$Y_t \leq \bar{p}_t \quad 0 \leq t \leq T, \quad \int_0^T (\bar{p}_t - Y_t) dK_t = 0 \quad (39)$$

$$Y_T = \bar{p}_T, \quad (40)$$

Y is a continuous modification of V^b , and $p_t^b = P_t^b(Y_t)$ and $\tau^b = \inf\{s \geq 0: Y_s = \bar{p}_s\}$ form an optimal control for the short agent. It turns out that, because the optimal stopping time has to be the same for both agents in equilibrium, we can formulate a system of equations for V^a and V^b without \bar{p} . In order to state this result formally, we need to introduce the following random functions

$$\tilde{\mathcal{G}}_t^a(y, z) = \mathcal{G}_t^a(y, P_t^b(z)) = -c_t^{\alpha^0}(P_t^a(y), P_t^b(z))y + g_t^a(P_t^a(y), P_t^b(z)), \quad y, z \in \mathbb{R}, \quad (41)$$

$$\tilde{\mathcal{G}}_t^b(y, z) = \mathcal{G}_t^b(P_t^a(y), z) = -c_t^{\beta^0}(P_t^a(y), P_t^b(z))z + g_t^b(P_t^a(y), P_t^b(z)), \quad y, z \in \mathbb{R}, \quad (42)$$

where c^α , g^a and g^b are defined, respectively, in (13), (23) and (25), and P^a and P^b are given by (28) and (29).

Lemma 3. *Let Assumptions 1–6 hold. For any equilibrium $(p^a, p^b, \tau, \bar{p})$ in the two-player game (in the sense of Definition 5), the value functions of the agents, $V^a, V^b \in \mathbb{S}^2$, satisfy*

$$\begin{cases} -dV_t^a = \tilde{\mathcal{G}}_t^a(V_t^a, V_t^b)dt - Z_t^a dW_t + dK_t^a \\ -dV_t^b = \tilde{\mathcal{G}}_t^b(V_t^a, V_t^b)dt - Z_t^b dW_t - dK_t^b \\ V_t^a \geq V_t^b \quad \forall t \in [0, T], \quad \int_0^T (V_t^a - V_t^b)d(K_t^a + K_t^b) = 0 \\ V_T^a = V_T^b, \end{cases} \quad (43)$$

with some increasing processes $K^a, K^b \in \mathbb{S}^2$, starting at zero, and with progressively measurable square-integrable (Z^a, Z^b) . Moreover, $(\hat{p}^a, \hat{p}^b, \hat{\tau}, \bar{p})$ also form an equilibrium, with the same value functions, where: $\hat{p}_t^a = P_t^a(V_t^a)$, $\hat{p}_t^b = P_t^b(V_t^b)$ and $\hat{\tau} = \inf\{s \geq 0: V_s^a = V_s^b\}$. Conversely, given a solution to (43), we can define the optimal controls $(\hat{p}^a, \hat{p}^b, \hat{\tau})$ as above, and choose $\bar{p} = (1 - \eta)V^a + \eta V^b$, with any progressively measurable process η taking values in $(0, 1)$, to obtain an equilibrium $(\hat{p}^a, \hat{p}^b, \hat{\tau}, \bar{p})$.

Proof: Consider an equilibrium $(p^a, p^b, \tau, \bar{p})$. As discussed earlier, the standard results on BSDEs (cf. [35]) imply that (V^a, Z^a, K^a) solves (35)–(37), and (V^b, Z^b, K^b) solves (38)–(40) (both systems are considered with the same \bar{p}). It follows from the optimality of τ , via the standard theory, that $V_\tau^b = \bar{p}_\tau = V_\tau^a$. Consider the long agent. It is clear that the objective of the long agent cannot increase if we replace \bar{p} by V^b in its definition (cf. (22)). On the other hand, τ is optimal and $\bar{p}_\tau = V_\tau^b$, hence, the value function V^a remains the same if we replace \bar{p} by V^b in its definition (cf. (26)). Therefore, (V^a, Z^a, K^a) solves (35)–(37) with \bar{p} replaced by V^b . Similar argument applies to the short agent, and yields that (V^b, Z^b, K^b) solves (38)–(40) with \bar{p} replaced by V^a . Next, using the optimality of p^a and the comparison principle for the BSDE (32), we easily deduce that, for a.e. (t, ω) , p_t^a coincides with $\hat{p}_t^a = P_t^a(V_t^a)$ whenever $\lambda_t^{\alpha^0} > 0$ and $V_t^a < \sup \text{supp}(f_t^{\alpha^0})$. On the other hand, Assumption 6 implies that, if $\lambda_t^{\alpha^0} = 0$ or $V_t^a \geq \sup \text{supp}(f_t^{\alpha^0})$, then $\lambda_t^{\beta^0} = 0$ or $V_t^a \geq \sup \text{supp}(f_t^{\beta^0})$, and, in turn, $\mathcal{G}_t^b(p_t^a, V_t^b) = \mathcal{G}_t^b(\hat{p}_t^a, V_t^b)$. Thus, we conclude that V^b satisfies (38)–(40) with p^a replaced by \hat{p}^a . Similarly, we conclude that V^a satisfies (35)–(37) with p^b replaced by \hat{p}^b . Thus, (V^a, V^b) satisfy (43).

Next, consider a solution to (43). Choosing \bar{p} as shown in the statement of the lemma, we conclude that (V^a, Z^a, K^a) solves (35)–(37), with p^b replaced by \hat{p}^b . Then, the standard results (cf. [35]) imply that, given \hat{p}^b and \bar{p} , V^a is the value function of the long agent, and her optimal control is given by \hat{p}^a and

$$\inf\{s \geq 0: V_s^a \leq \bar{p}_s\} = \inf\{s \geq 0: V_s^a = V_s^b\} = \hat{\tau}.$$

Similar argument applies to the short agent, completing the proof. ■

3.2 Existence of a solution

In this subsection, we address the question of existence of a solution to the RBSDE (43). The main difficulty in analyzing (43) is the non-standard form of reflection: the components of the solution reflect against each other, as opposed to reflecting against a given boundary. Related equations have been analyzed in the literature on BSDEs arising in the problem of optimal switching: see, e.g., [13], [20], and the references therein. However, the exact form of reflection in (43) is different, and its generator does not possess the desired monotonicity properties,

making it impossible to prove the existence of a solution to (43) using the methods developed in optimal switching literature. Before we analyze the existence, it is convenient to consider the question of uniqueness. Note that there are two reflecting components of the solution, but only one minimality constraint, which indicates the potential lack of uniqueness of a solution to (43). The possibility of an arbitrary choice of η in Lemma 3 leads to the same conclusion. Indeed, a different choice of η produces a different \bar{p} , which results in a different pair of value functions (V^a, V^b) , which, nevertheless, have to solve the same system (43). This heuristic observation turns out to be correct and, in fact, allows us to construct a solution to (43). Consider a solution $(V^a, V^b, K^a, K^b, Z^a, Z^b)$ to (43). Introducing $K_t = K_t^a + K_t^b$, we notice that there must exist a process η , with values in $[0, 1]$, such that $dK_t^a = \eta_t dK_t$, $dK_t^b = (1 - \eta_t) dK_t$. Then, we introduce the new variables $(\tilde{Y}^1, \tilde{Y}^2)$, s.t. $\tilde{Y}_t^1 = V_t^a - V_t^b$ and $d\tilde{Y}_t^2 = (1 - \eta_t) dV_t^a + \eta_t dV_t^b$, to replace (V^a, V^b) . Assuming that the change of variables can be inverted, one obtains a system of RBSDEs for $(\tilde{Y}^1, \tilde{Y}^2)$, in which only the first component reflects against zero, and $\tilde{Y}_T^1 = 0$. Conversely, we can start by prescribing η and a terminal condition for \tilde{Y}^2 , solving the associated system of RBSDEs for $(\tilde{Y}^1, \tilde{Y}^2)$, and, then, recover (V^a, V^b) from $(\tilde{Y}^1, \tilde{Y}^2, \eta)$ via the above formulas. Naturally, the resulting (V^a, V^b) are expected to satisfy (43). This method seems to describe all solutions to (43), however, herein, we are only interested in constructing a particular one.¹³ Hence, we choose $\eta \equiv 1/2$ and $\tilde{Y}_T^2 = 0$, to obtain $Y^1 = \tilde{Y}^1 = V^a - V^b$ and $Y^2 = 2\tilde{Y}^2 = V^a + V^b$, which are expected to satisfy:

$$\begin{cases} -dY_t^1 = \mathcal{G}_t^1(Y_t^1, Y_t^2)dt - Z_t^1 dW_t + dK_t \\ Y_t^1 \geq 0, \quad \int_0^T Y_t^1 dK_t = 0, \quad Y_T^1 = 0 \\ -dY_t^2 = \mathcal{G}_t^2(Y_t^1, Y_t^2)dt - Z_t^2 dW_t, \quad Y_T^2 = 0 \end{cases} \quad (44)$$

where $Y^1, Y^2 \in \mathbb{S}^2$, the processes Z^1, Z^2 are progressively measurable and square-integrable, $K \in \mathbb{S}^2$ is increasing and satisfies $K_0 = 0$. In addition, we denote

$$\begin{aligned} \mathcal{G}_t^1(y^1, y^2) &= \tilde{\mathcal{G}}_t^a((y^1 + y^2)/2, (y^2 - y^1)/2) - \tilde{\mathcal{G}}_t^b((y^1 + y^2)/2, (y^2 - y^1)/2), \\ \mathcal{G}_t^2(y^1, y^2) &= \tilde{\mathcal{G}}_t^a((y^1 + y^2)/2, (y^2 - y^1)/2) + \tilde{\mathcal{G}}_t^b((y^1 + y^2)/2, (y^2 - y^1)/2) \end{aligned}$$

where $\tilde{\mathcal{G}}^a$ and $\tilde{\mathcal{G}}^b$ are defined in (41) and (42). The following lemma formalizes the connection between (44) and (43), and its proof follows easily by a direct verification.

Lemma 4. *Let (Y^1, Y^2, Z^1, Z^2, K) be a solution to (44). Then*

$$V^a = \frac{1}{2}Y^1 + \frac{1}{2}Y^2, \quad V^b = \frac{1}{2}Y^2 - \frac{1}{2}Y^1, \quad Z^a = \frac{1}{2}Z^1 + \frac{1}{2}Z^2, \quad Z^b = \frac{1}{2}Z^2 - \frac{1}{2}Z^1, \quad K^a = \frac{1}{2}K, \quad K^b = \frac{1}{2}K$$

form a solution to (43).

Note that the specific choice of $\eta \equiv 1/2$ corresponds to choosing an angle of reflection of the process (V^a, V^b) against the straight line “ $V^a = V^b$ ” in \mathbb{R}^2 . The specific angle chosen to obtain (44) implies orthogonal reflection against this line, and (44) arises after a simple rotation, which turns this line into a horizontal axis. The systems of RBSDEs with orthogonal reflection in a general convex domain have been analyzed in [24]. However, the latter results are not applicable in the present case, as the generator of (44) lacks the global Lipschitz property. Indeed, the generator can be written as

$$\mathcal{G}_t^1(y^1, y^2) = -c_t^1(y^1, y^2)y^1 + c_t^2(y^1, y^2)y^2 + g_t^1(y^1, y^2), \quad (45)$$

$$\mathcal{G}_t^2(y^1, y^2) = -c_t^2(y^1, y^2)y^1 - c_t^1(y^1, y^2)y^2 + g_t^2(y^1, y^2), \quad (46)$$

where

$$c_t^1(y^1, y^2) = \frac{1}{2}c_t^{\alpha^0}(P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)) + \frac{1}{2}c_t^{\beta^0}(P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2))$$

¹³It is an interesting topic for future research, to describe rigorously all solutions of (43).

$$\begin{aligned}
c_t^\alpha(y^1, y^2) &= \frac{1}{2} c_t^{\beta^0} (P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)) - \frac{1}{2} c_t^{\alpha^0} (P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)), \\
g_t^1(y^1, y^2) &= g_t^a (P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)) - g_t^b (P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)), \\
g_t^2(y^1, y^2) &= g_t^a (P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)) + g_t^b (P_t^a((y^1 + y^2)/2), P_t^b((y^2 - y^1)/2)),
\end{aligned}$$

with c^α , P^a , P^b , g^a and g^b defined in (13), (28), (29), (23) and (25). It is easy to see that every $c_t^i(\cdot, \cdot)$ and $g_t^i(\cdot, \cdot)$ is bounded and globally Lipschitz, uniformly over a.e. (t, ω) . However, due to the presence of the multipliers y^1 and y^2 , $g_t^i(\cdot, \cdot)$ is unbounded and does not possess the global Lipschitz property. In addition, the existence and uniqueness result established below (cf. Proposition 1) holds for any choice of constant η in $(0, 1)$, which, in turn, implies the “oblique” (i.e. non-orthogonal) reflection of (V^a, V^b) against the boundary, and brings the resulting system outside the scope of [24].

Recall that the existence result for BSDEs with linear growth, and without global Lipschitz property, has only been established in a one-dimensional case, whereas the present equation is multidimensional. Nevertheless, we can make use of the fact that the generator of (44) has the “correct” asymptotic behavior, to prove the existence of a solution. In particular, we exploit the fact that, due to the assumptions made earlier in this section, whenever $\|(Y_t^1, Y_t^2)\|$ becomes large, the generator $(\mathcal{G}_t^1, \mathcal{G}_t^2)$ pushes (Y_t^1, Y_t^2) in the direction in which the largest $|Y_t^i|$ decreases.

Proposition 1. *Let Assumptions 2–6 hold. Then, there exists a solution to (44), s.t. its components Y^1 and Y^2 are absolutely bounded by a constant. Such a solution is unique.*

Proof: Step 1: Existence for the fully capped system. For any constant $C > 0$, denote $\Psi_C(y) = (-C \vee y) \wedge C$. Clearly, this function is 1-Lipschitz in y and absolutely bounded by C . We fix arbitrary constants $\{C_i^j > 0\}$ and consider the fully capped system:

$$\begin{cases} -dY_t^1 = \left(-c_t^1(Y_t^1, Y_t^2) \Psi_{C_1^1}(Y_t^1) + c_t^2(Y_t^1, Y_t^2) \Psi_{C_1^2}(Y_t^2) + g_t^1(Y_t^1, Y_t^2) \right) dt - Z_t^1 dW_t + dK_t \\ -dY_t^2 = \left(-c_t^2(Y_t^1, Y_t^2) \Psi_{C_2^1}(Y_t^1) - c_t^1(Y_t^1, Y_t^2) \Psi_{C_2^2}(Y_t^2) + g_t^2(Y_t^1, Y_t^2) \right) dt - Z_t^2 dW_t \end{cases} \quad (47)$$

Here, and in some expressions that follow, we omit the terminal condition, barrier, and the minimality condition for K_t , as they remain unchanged throughout. Assumptions 2–6 imply that $c_t^1(y^1, y^2)$, $c_t^2(y^1, y^2)$, $g_t^1(y^1, y^2)$ and $g_t^2(y^1, y^2)$ are bounded and globally Lipschitz in (y^1, y^2) , uniformly over a.e. (t, ω) . Hence, the generator of (47) is globally Lipschitz in (y^1, y^2) (and independent of (Z^1, Z^2)), and the standard existence results for Lipschitz BSDEs (cf. for example, Theorem 2.2 in [49]) yield the existence (and uniqueness) of a solution to (47). Denote the Y -component of this solution (Y_t^{1c}, Y_t^{2c}) .

Step 2: Bounds on solution components via partial uncapping. We want to bound the components (Y_t^{1c}, Y_t^{2c}) , of the solution to the capped system, by using the control-stopping interpretation of the individual (R)BSDEs comprising our system. Consider the associated equation for Y^1 , with Y_t^{2c} being given:

$$\begin{cases} -dY_t^1 = \left(-c_t^1(Y_t^1, Y_t^{2c}) Y_t^1 + c_t^2(Y_t^1, Y_t^{2c}) \Psi_{C_1^2}(Y_t^{2c}) + g_t^1(Y_t^1, Y_t^{2c}) \right) dt - Z_t^1 dW_t + dK_t \\ Y_t^1 \geq 0, \quad \int_0^T Y_t^1 dK_t = 0, \quad Y_T^1 = 0 \end{cases} \quad (48)$$

Note that, as c_t^1, c_t^2, g^1 and $\Psi_{C_1^2}$ are bounded, this one-dimensional RBSDE has a continuous generator with linear growth in Y^1 , and, for example, by Theorem 4.1 in [49], it has a solution, which we denote Y_t^1 . Next, for Y^1 and Y^{2c} constructed above, we introduce the processes

$$\tilde{c}_t^1 = c_t^1(Y_t^1, Y_t^{2c}), \quad \tilde{c}_t^2 = c_t^2(Y_t^1, Y_t^{2c}), \quad \tilde{g}_t^1 = g_t^1(Y_t^1, Y_t^{2c}), \quad \tilde{g}_t^2 = g_t^2(Y_t^1, Y_t^{2c}),$$

and consider the one-dimensional RBSDE (for \tilde{Y}), obtained from (48) by pretending that the coefficients should depend on the solution itself:

$$\begin{cases} -d\tilde{Y}_t^1 = \left(-\tilde{c}_t^1 \tilde{Y}_t^1 + \tilde{c}_t^2 \Psi_{C_1^2}(\tilde{Y}_t^{2c}) + \tilde{g}_t^1 \right) dt - Z_t^1 dW_t + dK_t \\ \tilde{Y}_t^1 \geq 0, \quad \int_0^T \tilde{Y}_t^1 dK_t = 0, \quad \tilde{Y}_T^1 = 0 \end{cases} \quad (49)$$

Note that $\tilde{Y} = Y^1$ is the unique solution of this equation. On the other hand, the above RBSDE is affine in \tilde{Y} , and, for example, by Theorem 7.1 in [35], its unique solution admits the following interpretation, as the value function of an optimal stopping problem:

$$Y_t^1 = \sup_{\tau \in \mathcal{T}_t} \mathbb{E} \left[\int_t^\tau \exp \left(- \int_t^s \tilde{c}_u^1 du \right) \left(\tilde{c}_s^2 \Psi_{C_1^2}(Y_s^{2c}) + \tilde{g}_s^1 \right) ds \middle| \mathcal{F}_t \right]$$

We will use this representation to establish a bound on $|Y^1|$. First, note that, under our assumptions, there exist constants $C_0 > 0$ and $\lambda \in (0, 1)$, such that, for all t, y^1, y^2 , and a.e. ω , we have:

$$\begin{aligned} \left| \frac{g_t^i(y^1, y^2)}{c_t^1(y^1, y^2)} \right| &= 2 \left| \frac{g_t^a \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right) \pm g_t^b \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right)}{c_t^{\alpha^0} \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right) + c_t^{\beta^0} \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right)} \right| \leq C_0, \\ \left| \frac{c_t^2(y^1, y^2)}{c_t^1(y^1, y^2)} \right| &= \left| \frac{c_t^{\alpha^0} \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right) - c_t^{\beta^0} \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right)}{c_t^{\alpha^0} \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right) + c_t^{\beta^0} \left(P_t^a \left(\frac{y^1+y^2}{2} \right), P_t^b \left(\frac{y^2-y^1}{2} \right) \right)} \right| \leq \lambda < 1, \end{aligned}$$

with c^α, P^a, P^b, g^a and g^b defined in (13), (28), (29), (23) and (25). The first inequality holds with $C_0 = 5C_p$, and it follows from the boundedness of P^a, P^b and the jump sizes. The second one follows from Assumption 6. The above inequalities imply:

$$\left| \frac{\tilde{c}_t^2 \Psi_{C_1^2}(Y_t^{2c}) + \tilde{g}_t^1}{\tilde{c}_t^1} \right| \leq \lambda C_1^2 + C_0,$$

for all t and a.e. ω . The latter estimate, together with the following lemma, imply the desired upper bound:

$$|Y_t^1| \leq \lambda C_1^2 + C_0$$

for all t and a.e. ω .

Lemma 5. Consider any constant $C > 0$, any continuous function $S : [0, T] \rightarrow \mathbb{R}$, absolutely bounded by C , any nonnegative continuous function c on $[0, T]$, and any continuous function g on $[0, T]$, satisfying $|g| \leq C|c|$. For any $0 \leq t \leq \tau \leq T$, denote:

$$Y_{t,\tau} = \int_t^\tau \exp \left(- \int_t^s c(u) du \right) g(s) ds + \exp \left(- \int_t^\tau c(u) du \right) S(\tau).$$

Then

$$|Y_{t,\tau}| \leq C, \quad \forall 0 \leq t \leq \tau \leq T.$$

Proof: For any $0 \leq t \leq \tau \leq T$, we have

$$\begin{aligned} \left| \int_t^\tau \exp \left(- \int_t^s c(u) du \right) g(s) ds + \exp \left(- \int_t^\tau c(u) du \right) S(\tau) \right| \\ \leq - \int_t^\tau C d \left(\exp \left(- \int_t^s c(u) du \right) \right) + \exp \left(- \int_t^\tau c(u) du \right) C = C \end{aligned}$$

■

Thus, we have a solution Y^1 of (48) which satisfies $|Y_t^1| \leq \lambda C_1^2 + C_0$, \mathbb{P} -a.s., for all t . Then, for $C_1^1 \geq \lambda C_1^2 + C_0$, we have $\Psi_{C_1^1}(Y_t^1) = Y_t^1$, and, hence, Y^1 also solves

$$\begin{cases} -dY_t^1 = \left(-c_t^1(Y_t^1, Y_t^{2c}) \Psi_{C_1^1}(Y_t^1) + c_t^2 \Psi_{C_1^2}(Y_t^{2c}) + g_t^1(Y_t^1, Y_t^{2c}) \right) dt - Z_t^1 dW_t + dK_t \\ Y_t^1 \geq 0, \quad \int_0^T Y_t^1 dK_t = 0, \quad Y_T^1 = 0 \end{cases}$$

Note that the above RBSDE coincides with the Y^1 -equation in (47). This one-dimensional RBSDE has a globally Lipschitz generator and, thus, a unique solution. This implies that $Y^1 = Y^{1c}$, and we obtain the desired bound on Y^{1c} :

$$|Y_t^{1c}| \leq \lambda C_1^2 + C_0,$$

\mathbb{P} -a.s. for all t , provided $C_1^1 \geq \lambda C_1^2 + C_0$. Similarly, considering the Y^2 part of the capped system (47), with Y^{1c} fixed, we obtain

$$|Y_t^{2c}| \leq \lambda C_2^1 + C_0,$$

\mathbb{P} -a.s. for all t , provided $C_2^2 \geq \lambda C_2^1 + C_0$.

Step 3: Solution of the appropriately capped system solves the original system. To show that the solution (Y_t^{1c}, Y_t^{2c}) of (47) also solves the original system (44), we only need to show that, given the bounds on (Y^{1c}, Y^{2c}) , the capped system's generator coincides with the original generator, which translates into

$$\Psi_{C_1^1}(Y_t^{1c}) = Y_t^{1c}, \quad \Psi_{C_2^2}(Y_t^{2c}) = Y_t^{2c}, \quad \Psi_{C_1^2}(Y_t^{2c}) = Y_t^{2c}, \quad \Psi_{C_2^1}(Y_t^{1c}) = Y_t^{1c}.$$

The first two equalities are satisfied if

$$C_1^1 \geq \lambda C_1^2 + C_0, \quad C_2^2 \geq \lambda C_2^1 + C_0,$$

while the last two require

$$\lambda C_2^1 + C_0 \leq C_1^2, \quad \lambda C_1^1 + C_0 \leq C_2^2.$$

One can check these inequalities have a solution, as long as $\lambda < 1$. The “minimal” solution being

$$C_1^1 = C_2^1 = C_2^2 = C_1^2 = \frac{C_0}{1-\lambda}.$$

With the above choice of capping, the solution to (47) also solves (44), thus, showing the existence of a solution of (44). This solution is bounded by construction. The uniqueness of a bounded solution follows from the fact that, when (y^1, y^2) vary over a bounded set, the generator of (44) is Lipschitz, hence, the standard results yield uniqueness. ■ ■

Remark 4. *The above proof provides an existence result for any system (44), whose generator is given by (45)–(46), with arbitrary (bounded and Lipschitz) progressively measurable random functions $\{c^i, g^i\}$, as long as the following holds for a.e. (t, ω) and all $(y^1, y^2) \in \mathbb{R}^2$:*

$$\sum_{i=1}^2 |g_t^i(y^1, y^2)| \leq C_0 c_t^1(y^1, y^2), \quad |c_t^2(y^1, y^2)| \leq \lambda c_t^1(y^1, y^2),$$

with some constants $C_0 > 0$ and $\lambda \in (0, 1)$.

4 Equilibrium in the continuum-player game.

In this section we construct an equilibrium for the continuum-player game described in Section 2, in the sense of Definition 3. The main difficulty in constructing the equilibrium stems from the mixed control-stopping nature of the game (and, of course, the fact there are multiple participants). Therefore, we attempt to break the problem into two parts - isolating the “stopping” part of the game. In order to do this, it is convenient to make assumptions that guarantee the existence of the so-called “extremal” agents on each side of the book. These agents are called “extremal”, because their beliefs dominate the beliefs of the other agents on the same side of the book, in the sense explained below. We denote the extremal beliefs on the long side by α^0 , and, on the short side, by β^0 . In short, the agents with beliefs α^0 are the least bullish among the long ones, and the agents with beliefs β^0 are the least bearish among the short ones. The extremal agents, e.g., can be interpreted as market-makers, as they are closer to being market-neutral than any other agent on the same side of the book (recall that we, still, do not have any designated market makers in this game – market-neutrality is only one of the characteristics of a market maker).

Indeed, if one assumes that the long agents are bullish (which is natural, as, before the end of the game, the long agents choose to wait instead of submitting market orders), then, the agents with beliefs α^0 are the least bullish ones. In this section, we construct an equilibrium in which the time of the first internal market order and the bid and ask prices are determined by the extremal agents, while the rest of the shape of the LOB is due to the other agents' actions. The construction of an equilibrium, thus, splits into two parts. In the first part, the extremal agents find an equilibrium among themselves, using the results of the auxiliary two-player game, and determining the time of the first internal market order τ and the bid and ask prices p^a and p^b . In the second part, the other agents, taking (p^a, p^b, τ) as given, determine their optimal actions. Of course, we, ultimately, prove that the strategy of every agent is optimal in the overall market, consisting of both extremal and non-extremal agents. The resulting LOB ν has two atoms – at the bid and ask prices – comprised of the limit orders of the extremal and some of the non-extremal agents. The rest of the LOB contains limit orders of the non-extremal agents only.

In order to implement the above program, we assume that $\mathbb{A} = \{\alpha^0\} \cup \hat{\mathbb{A}}$ and $\mathbb{B} = \{\beta^0\} \cup \hat{\mathbb{B}}$. We assume that Assumptions 1–6 hold throughout this section. In addition, we make the following assumptions.

Assumption 7. For any $\alpha \in \hat{\mathbb{A}}$, $\beta \in \hat{\mathbb{B}}$ and a.e. (t, ω) , we have:

$$\begin{aligned} \lambda_t^\alpha F_t^{+, \alpha}(p) &\geq \lambda_t^{\alpha^0} F_t^{+, \alpha^0}(p), & \lambda_t^\beta F_t^{+, \beta}(p) &\leq \lambda_t^{\beta^0} F_t^{+, \beta^0}(p), & \forall p \geq 0, \\ \lambda_t^\alpha F_t^{-, \alpha}(p) &\leq \lambda_t^{\alpha^0} F_t^{-, \alpha^0}(p), & \lambda_t^\beta F_t^{-, \beta}(p) &\geq \lambda_t^{\beta^0} F_t^{-, \beta^0}(p), & \forall p \leq 0. \end{aligned}$$

Assumption 8. For any $\alpha \in \hat{\mathbb{A}}$, $\beta \in \hat{\mathbb{B}}$ and a.e. (t, ω) , we have:

$$\frac{F_t^{+, \alpha^0}(p)}{f_t^{\alpha^0}(p)} \leq \frac{F_t^{+, \alpha}(p)}{f_t^\alpha(p)}, \quad \frac{F_t^{-, \beta^0}(-p)}{f_t^{\beta^0}(-p)} \leq \frac{F_t^{-, \beta}(-p)}{f_t^\beta(-p)} \quad \forall p \geq 0.$$

Assumption 7 ensures that the distribution of the fundamental price at any time t , from an α -agent's perspective, dominates stochastically the respective distribution from the α^0 -agent's perspective. The opposite relation holds for the short agents. The first inequality in Assumption 8 ensures that $\log F_t^{+, \alpha^0}(\cdot)$ decays faster than $\log F_t^{+, \alpha}(\cdot)$, which is also consistent with the interpretation that α^0 -agents assign smaller probabilities to the large jumps of the fundamental price, and larger probabilities to the small jumps, as compared to the α -agents. Analogous interpretation holds for the second inequality in Assumption 8. Assumption 8 ensures that, in an empty LOB, the non-extremal agents would prefer to post their limit order further away from zero than the extremal ones do.

Lemma 6. Let Assumptions 1–8 hold. Fix any $\alpha \in \hat{\mathbb{A}}$ and $\beta \in \hat{\mathbb{B}}$. Then, for a.e. (t, ω) , the following holds for all $y \in \mathbb{R}$: $p \mapsto (p - y)F_t^{+, \alpha}(p)$ is non-decreasing in $p \in [y, P_t^a(y)]$, and $p \mapsto (y - p)F_t^{-, \beta}(p)$ is non-increasing in $p \in [P_t^b(y), y]$.

Proof: The statement follows easily by differentiating the target functions, recalling (30)–(31), and making use of Assumption 8. ■

We also need to make an assumption that limits the maximum possible demand size, as viewed by the extremal agents. Namely, the extremal agents believe that the external demand can never exceed the inventory held by these agents.

Assumption 9. For Leb $\otimes \mathbb{P}$ -a.e. (t, ω) , we have:

$$D_t \left(-Q^+ \left(f_t^{\alpha^0/\beta^0}(x) dx \right) \right) \leq \mu^a(\{\alpha^0\}), \quad -D_t \left(-Q^- \left(f_t^{\alpha^0/\beta^0}(x) dx \right) \right) \leq \mu^b(\{\beta^0\}),$$

where Q^+ and Q^- are defined in (4).

In order to construct an equilibrium, we need to impose certain topological conditions on the space of beliefs and on the mapping $\alpha \mapsto f^\alpha$.

Assumption 10. The spaces $\hat{\mathbb{A}}$ and $\hat{\mathbb{B}}$ are compact metric spaces, with the Borel sigma-algebras on them (i.e. μ^a and μ^b are measures with respect to the Borel sigma-algebras). In addition, for a.e. (t, ω) , the mapping $\alpha \mapsto f_t^\alpha$ is continuous as a mapping $\hat{\mathbb{A}} \rightarrow \mathbb{L}^1[0, C_p]$ and as a mapping $\hat{\mathbb{B}} \rightarrow \mathbb{L}^1[-C_p, 0]$.

Finally, we need to ensure that the demand size curve is “not too flat”.

Assumption 11. *There exists an increasing continuous (deterministic) function $\epsilon: [0, \infty) \rightarrow [0, \infty)$, s.t. $\epsilon(0) = 0$ and, for a.e. (t, ω) , $|D_t^{-1}(x) - D_t^{-1}(y)| \leq \epsilon(|x - y|)$, for all $x, y \in \mathbb{R}$.*

Now, we proceed to construct a special class of equilibria in the continuum-player game. As announced earlier, the equilibrium is constructed by, first, solving the auxiliary two-player game, as described in Section 3. In the two-player game, we assume that the two agents have beliefs α^0 and β^0 . Thus, we consider the unique bounded solution (Y^1, Y^2) to (44) and construct the associated (V^a, V^b) , which solve (43), according to Lemma 4. Then, Lemma 3 implies that (V^a, V^b) are the value functions of the two-player equilibrium $(\hat{p}^a, \hat{p}^b, \hat{\tau}, \bar{p})$, where

$$\hat{p}_t^a = P_t^a(V_t^a), \quad \hat{p}_t^b = P_t^b(V_t^b), \quad \hat{\tau} = \inf\{t \in [0, T] : V_t^a = V_t^b\}, \quad \bar{p}_t = \frac{1}{2}V_t^a + \frac{1}{2}V_t^b.$$

Let us introduce

$$p_t^a = \hat{p}_t^a \mathbf{1}_{\{t < \hat{\tau}\}} + \bar{p}_t \mathbf{1}_{\{t \geq \hat{\tau}\}}, \quad p_t^b = \hat{p}_t^b \mathbf{1}_{\{t < \hat{\tau}\}} + \bar{p}_t \mathbf{1}_{\{t \geq \hat{\tau}\}}. \quad (50)$$

Using these auxiliary quantities, we aim to construct an equilibrium for the continuum-player game, in which (ν, θ) satisfy the following two conditions. First,

$$\nu_t^a = \mu^a(\{\alpha^0\})\delta_{p_t^a} + \bar{\nu}_t^a, \quad \nu_t^b = \mu^b(\{\beta^0\})\delta_{p_t^b} + \bar{\nu}_t^b, \quad (51)$$

with progressively measurable $\bar{\nu}^a$ and $\bar{\nu}^b$ taking values in the space of sigma-additive measures on \mathbb{R} , such that, \mathbb{P} -a.s., for all $t \in [0, T]$, $\bar{\nu}_t^a$ is supported on $[p_t^a, C_p]$ and $\bar{\nu}_t^b$ is supported on $[-C_p, p_t^b]$.¹⁴ Second,

$$\theta_t^a = \mu^a(\mathbb{A})\delta_{V_t^a}, \quad \theta_t^b = \mu^b(\mathbb{B})\delta_{V_t^b}. \quad (52)$$

Note that, in such a market, we have

$$\tau^a = \tau^b = \hat{\tau}$$

The following theorem is the main result of this paper.

Theorem 1. *Let Assumptions 1–11 hold. Consider any solution (V^a, V^b) to (43) (whose existence is guaranteed by Proposition 1 and Lemma 4) and the associated (p^a, p^b) given by (50). Then, there exist progressively measurable measure-valued processes (ν, θ) and random fields $p, v : \Omega \times [0, T] \times \mathbb{S} \rightarrow \mathcal{P}(\mathbb{R}) \times \mathbb{R}$, which form an equilibrium, in the sense of Definition 3, and which satisfy (51)–(52) along with*

- $v_t(1, \alpha) = V_t^a, v_t(-1, \alpha) = V_t^b$, for all (t, ω, α) ,
- $p_t(1, \alpha^0) = p_t^a, p_t(-1, \beta^0) = p_t^b$, for all (t, ω) .

Remark 5. *Recall that there always exists a “trivial” equilibrium, in which all agents stop at time zero. However, such an equilibrium is unrealistic and does not appear to be useful from a modeling perspective. The main contribution of the above result is the existence of a potentially non-trivial equilibrium, in which the duration of the game, $\hat{\tau}$, is determined by a solution to (43), and there is no reason for it to be zero, in general. The latter is confirmed by the numerical experiments in Section 5.*

Remark 6. *Notice that, as announced in Remark 2, we have constructed an equilibrium, satisfying*

$$v_t(1, \alpha) = v_t^a = V_t^a, \quad v_t(-1, \alpha) = v_t^b = V_t^b, \quad \forall \alpha \in \mathbb{A} \cup \mathbb{B}, (t, \omega) \in [0, T] \times \Omega.$$

Therefore, in such an equilibrium, no agents execute market orders before the end of the game $\hat{\tau}$, and, hence, the empirical distribution μ remains constant and (12) holds.

The remainder of this section is devoted to the proof of Theorem 1. First, we show that, in a market (ν, θ) , satisfying (51)–(52), it is never (strictly) optimal for the agents to post limit sell orders below the ask price or to post limit buy orders above the bid price. In addition, it is never (strictly) optimal for the agents to submit a market order before $\hat{\tau}$. To achieve this, we need to compare the value functions of the agents to V^a and V^b , making use of Assumptions 7, 8.

¹⁴The components $\bar{\nu}^a$ and $\bar{\nu}^b$ are introduced for convenience, in order to indicate that $\nu_t^a(\{p_t^a\}) \geq \mu^a(\{\alpha^0\})$ and $\nu_t^b(\{p_t^b\}) \geq \mu^b(\{\beta^0\})$.

Lemma 7. *Let Assumptions 1–8 hold, and let (ν, θ) satisfy (51)–(52). Given any $\alpha \in \mathbb{A}$ and any admissible control (p, τ) , for a long agent with beliefs α , there exists an admissible control p' , s.t., \mathbb{P} -a.s., $\text{supp}(p'_t) \subset [p_t^a, \infty)$, for all $t \in [0, T]$, and $(p', \hat{\tau})$ does not decrease the objective value, i.e.*

$$J^{(\nu, \theta), (p, \tau)}(1, \alpha) \leq J^{(\nu, \theta), (p', \hat{\tau})}(1, \alpha).$$

Similarly, given any $\beta \in \mathbb{B}$ and any admissible control (p, τ) , for a short agent with beliefs β , there exists an admissible control p' , s.t., \mathbb{P} -a.s., $\text{supp}(p'_t) \subset (-\infty, p_t^b]$, for all $t \in [0, T]$, and $(p', \hat{\tau})$ does not decrease the objective value, i.e.

$$J^{(\nu, \theta), (p, \tau)}(-1, \beta) \leq J^{(\nu, \theta), (p', \hat{\tau})}(-1, \beta).$$

The proof of the above lemma is given in the appendix. This lemma has a straight-forward but useful corollary.

Corollary 1. *Let Assumptions 1–8 hold, and let (ν, θ) satisfy (51)–(52). Given any $\alpha \in \mathbb{A}$, let (p, τ) be an optimal strategy for the long agents with beliefs α , in the class of all admissible strategies satisfying: \mathbb{P} -a.s. $\text{supp}(p_t) \subset [p_t^a, \infty)$, for all $t \in [0, T]$, and $\tau = \hat{\tau}$. Then (p, τ) is optimal in the class of all admissible strategies, in the sense of Definition 2. Similarly, given any $\beta \in \mathbb{B}$, let (p, τ) be an optimal strategy for the short agents with beliefs β , in the class of all admissible strategies satisfying: \mathbb{P} -a.s. $\text{supp}(p_t) \subset (-\infty, p_t^b]$, for all $t \in [0, T]$, and $\tau = \hat{\tau}$. Then (p, τ) is optimal in the class of all admissible strategies, in the sense of Definition 2.*

Thus, no matter which limit order strategy p an agent is using, it is optimal for her to choose the following stopping threshold:

$$\hat{v}(s) = V^a \mathbf{1}_{\{s > 0\}} + V^b \mathbf{1}_{\{s < 0\}}.$$

This implies that, given a LOB ν in the form (51) and the stopping strategy \hat{v} as above, if an optimal limit order strategy $\hat{p}(s, \alpha)$ exists for any state (s, α) , then $(\hat{p}(s, \alpha), \hat{v})$ form an optimal control for the agents in state (s, α) , in the sense of Definition 2. Moreover, in such a case, θ , given by (52), satisfies the condition (11). Next, we need to construct a LOB ν , in the form (51), and the associated optimal limit order strategies for all agents, s.t. (10) is satisfied. In particular, the following lemma, whose proof is postponed to the appendix, shows that, for any ν in the form (51), the strategies (δ_{p^a}, V^a) and (δ_{p^b}, V^b) are optimal for the extremal agents.

Lemma 8. *Let Assumptions 1–9 hold, and let (ν, θ) satisfy (51)–(52). Then, given (ν, θ) , the strategy (δ_{p^a}, V^a) is optimal for a long agent with beliefs α^0 , and the strategy (δ_{p^b}, V^b) is optimal for a short agent with beliefs β^0 , in the sense of Definition 2.*

The remaining steps are carried out in the next subsection.

4.1 Equilibrium strategies of the non-extremal agents

In this subsection we construct the measure-valued processes (ν^a, ν^b) , in the form (51), and a progressively measurable random field $(\hat{p}_t(s, \alpha))$, such that the controls $(\hat{p}(1, \alpha), V^a)$ and $(\hat{p}(-1, \alpha), V^b)$ are optimal for the non-extremal agents with beliefs α , long and short, respectively (recall that the optimal strategies for the extremal agents are constructed in Lemma 8), and the fixed-point constraint (10) is satisfied. In view of Lemma 7, we can restrict the possible controls p to the those satisfying: $\text{supp}(p_t) \subset [p_t^a, \infty)$, for all $t \in [0, T]$. It is also obvious that we can restrict the support of p_t to be in $[-C_p, C_p]$. As the stopping strategy is fixed, for any $\alpha \in \hat{\mathbb{A}}$, the objective of a long player reduces to $\bar{J}_0^{\alpha, (p)}$, where

$$\bar{J}_t^{\alpha, (p)} = \mathbb{E} \left[\int_t^T \exp \left(- \int_t^s \bar{c}_u^\alpha(p_u^a, p_u^b) du \right) \bar{h}_s^{\alpha, a}(p_s, p_s^a, p_s^b) ds + \exp \left(- \int_t^T \bar{c}_u^\alpha(p_u^a, p_u^b) du \right) p_T^b | \mathcal{F}_t \right],$$

$$\bar{c}_t^\alpha(p_t^a, p_t^b) = c_t^\alpha(p_t^a, p_t^b) \mathbf{1}_{\{t \leq \hat{\tau}\}}, \quad \bar{h}_t^{\alpha, a}(p_t, p_t^a, p_t^b) = h_t^{\alpha, a}(p_t, p_t^a, p_t^b) \mathbf{1}_{\{t \leq \hat{\tau}\}},$$

and c^α and $h^{\alpha, a}$ defined in (13) and (16). Due to Assumptions 7 and 9, we have

$$l_t^{c, b}(x) = \inf \{ p > Q^-(\nu_t^b) : -D_t(p - x) > \nu_t^b((p, \infty)) \} = p_t^b \vee x, \quad \forall x \in \text{supp}(f_t^\alpha).$$

In addition, for any $z \geq p_t^a$,

$$\{u > 0 : l_t^{c,a}(u) \geq z\} = \{u > 0 : u \geq z - D_t^{-1}(\nu_t^a([p_t^a, z]))\},$$

and, hence, for any $B \geq p_t^a$,

$$\int_{p_t^a}^B f_t^\alpha(u)(l_t^{c,a}(u) - p_t^a)du = \int_0^{l_t^{c,a}(B) - p_t^a} \int_{u + p_t^a - D_t^{-1}(\nu_t^a([p_t^a, p_t^a + u]))}^{C_p} f_t^\alpha(y)dydu.$$

The above observations allow us to simplify the objective:

$$\begin{aligned} h_t^{\alpha,a}(p_t, p_t^a, p_t^b) &= \lambda_t^\alpha \int_{p_t^a}^\infty \left[(z - p_t^b) F_t^{+, \alpha}(z - D_t^{-1}(\nu_t^a([p_t^a, z]))) + p_t^b F_t^{+, \alpha}(p_t^b) \right. \\ &\quad \left. + \int_{p_t^a}^{z - D_t^{-1}(\nu_t^a([p_t^a, z]))} f_t^\alpha(u) l_t^{c,a}(u) du \right] p_t(dz) + 2\lambda_t^\alpha p_t^b F_t^{-, \alpha}(p_t^b) \\ &= \lambda_t^\alpha \int_{p_t^a}^{C_p} \left[(z - p_t^b) F_t^{+, \alpha}(z - D_t^{-1}(\nu_t^a([p_t^a, z]))) \right. \\ &\quad \left. + \int_0^{z - p_t^a} F_t^{+, \alpha}(u + p_t^a - D_t^{-1}(\nu_t^a([p_t^a, p_t^a + u]))) du \right] p_t(dz) + 2\lambda_t^\alpha p_t^b F_t^{-, \alpha}(p_t^b) + \lambda_t^\alpha p_t^b F_t^{+, \alpha}(p_t^a). \end{aligned}$$

Notice that the above objective does not depend on ν^b (for a given p^b), hence, we can separate the equilibrium problems of the long and short agents (this is only true for the non-extremal agents, of course). For simplicity, we only consider the problem of the long agents – the short agents can be treated similarly. Denote by κ_t and $\hat{\nu}_t^a$ the push-forward measures of p_t and ν_t^a , under the mapping $x \mapsto x - p_t^a$. Clearly, the measurability property is preserved by this transformation, hence, we can reformulate the equilibrium problem as a search for κ and $\hat{\nu}^a$, with the values in the space of measures with support in $[0, C_p]$. In the new variables, the objective takes a more convenient form. In particular, $h_t^{\alpha,a}(p_t, p_t^a, p_t^b) = \hat{h}_t^{\alpha,a}(\kappa_t, p_t^a, p_t^b)$, where

$$\begin{aligned} \hat{h}_t^{\alpha,a}(\kappa_t, p_t^a, p_t^b) &= \lambda_t^\alpha \int_0^{C_p} \left[(z + p_t^a - p_t^b) F_t^{+, \alpha}(z + p_t^a - D_t^{-1}(\hat{\nu}_t^a([0, z]))) \right. \\ &\quad \left. + \int_0^z F_t^{+, \alpha}(u + p_t^a - D_t^{-1}(\hat{\nu}_t^a([0, u]))) du \right] \kappa_t(dz) + 2\lambda_t^\alpha p_t^b F_t^{-, \alpha}(p_t^b) + \lambda_t^\alpha p_t^b F_t^{+, \alpha}(p_t^a). \end{aligned}$$

Note that $\bar{J}^{\alpha,(p)}$ solves a BSDE with the affine generator

$$\hat{\mathcal{G}}_t^\alpha(y) = \bar{c}_u^\alpha(p_u^a, p_u^b)y + \hat{h}_t^{\alpha,a}(\kappa_t, p_t^a, p_t^b).$$

In order to maximize $\bar{J}^{\alpha,(p)}$, it suffices to find a strategy κ which maximizes the above generator. The latter is, in turn, equivalent to maximizing $\hat{h}_t^{\alpha,a}(\cdot, p_t^a, p_t^b)$. Thus, we need to find a progressively measurable random field $(\kappa_t(\alpha))$, with values in $\mathcal{P}(\mathbb{R})$ (with the weak topology on it), s.t., for μ^a -a.e. $\alpha \in \hat{\mathbb{A}}$,

$$\kappa_t(\alpha) \in \operatorname{argmax}_{\kappa' \in \psi} \hat{h}_t^{\alpha,a}(\kappa', p_t^a, p_t^b) \quad (53)$$

holds for $dt \times \mathbb{P}$ -a.e. (t, ω) , where $\psi = \{p \in \mathcal{P}(\Pi) : \operatorname{supp}(p) \subseteq \Pi\}$ and $\Pi = [0, C_p]$. The standard BSDE results, then, imply that $\kappa(\alpha)$ is optimal for the agents in state $(1, \alpha)$, for μ^a -a.e. $\alpha \in \hat{\mathbb{A}}$. If, in addition, we ensure that the fixed-point constraint (10) is satisfied (and a similar construction holds for the short agents), we obtain an equilibrium in the continuum-player game, in the sense of Definition 3. Notice that we can rewrite

$$\hat{h}_t^{\alpha,a}(\kappa', p_t^a, p_t^b) = \lambda_t^\alpha \int_{\mathbb{R}} F_t(\alpha, p, \hat{\nu}_t^a) \kappa'(dp) + 2\lambda_t^\alpha p_t^b F_t^{-, \alpha}(p_t^b) + \lambda_t^\alpha p_t^b F_t^{+, \alpha}(p_t^a),$$

$$F_t(\alpha, p, \hat{\nu}_t^a) = (p + p_t^a - p_t^b) F_t^{+, \alpha} (p + p_t^a - D_t^{-1} (\hat{\nu}_t^a([0, p]))) + \int_0^p F_t^{+, \alpha} (u + p_t^a - D_t^{-1} (\hat{\nu}_t^a([0, u]))) du. \quad (54)$$

Assuming the extremal long agents post limit orders at p^a , the fixed-point constraint (10) (more precisely, the part of (10) that corresponds to the long agents) becomes:

$$\hat{\nu}_t^a([0, x]) = \mu^a(\{\alpha_0\}) + \int_{\hat{\mathbb{A}}} \kappa_t(\alpha; [0, x]) \mu^a(d\alpha), \quad \forall x \geq 0. \quad (55)$$

The above equations can be solved separately for different (t, ω) , hence, to this end, we fix (t, ω) and omit the t subscript whenever it causes no ambiguity. The statements that follow hold for a.e. (t, ω) . It turns out that it is more convenient to search for a measure

$$K(d\alpha, dx) = \kappa(\alpha; dx) \mu^a(d\alpha),$$

which is an element of $\mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$, the space of finite sigma-additive measures on $\hat{\mathbb{A}} \times \Pi$, with the first marginal μ^a . Transition from K to κ is accomplished via the usual disintegration. Thus, for a.e. (t, ω) , we need to find $(K, \nu) \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi) \times \mathcal{M}_{\mu^a(\mathbb{A})}(\Pi)$ solving the following system

$$\begin{cases} K \in \operatorname{argmax}_{K \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)} \int F(\alpha, p, \nu) K(d\alpha, dp), \\ \nu(dx) = \mu^a(\{\alpha_0\}) \delta_0(dx) + K(\hat{\mathbb{A}} \times dx), \end{cases} \quad (56)$$

where $\mathcal{M}_{\mu^a(\mathbb{A})}(\Pi)$ is the space of finite sigma-additive measures on Π , with the total mass $\mu^a(\mathbb{A}) = \mu^a(\{\alpha_0\}) + \mu^a(\hat{\mathbb{A}})$. The above system can be formulated as a fixed-point problem, in an obvious way. However, the main challenge in solving this problem stems from the fact that $F(\alpha, \cdot, \cdot)$ is not continuous: e.g. it may be discontinuous in p , if ν has atoms. Therefore, we replace F by its “mollified” version:

$$\hat{F}(\alpha, p, \nu) = \sup_{p' \in \Pi} F(\alpha, p', \nu) - |p' - p|.$$

The following lemma shows that we can replace F by \hat{F} in (56), and any solution to the new problem will solve the original one.

Lemma 9. *For any $\alpha \in \hat{\mathbb{A}}$ and $\nu \in \mathcal{M}_{\mu^a(\mathbb{A})}(\Pi)$, the function $p \mapsto \hat{F}(\alpha, p, \nu)$ is 1-Lipschitz in $p \in \Pi$, and*

$$\operatorname{argmax}_{p \in \Pi} \hat{F}(\alpha, p, \nu) = \operatorname{argmax}_{p \in \Pi} F(\alpha, p, \nu).$$

Proof: For convenience, we drop the dependence on (α, ν) . The first statement is clear from the definition. It is also clear that $\sup_{p \in \Pi} \hat{F}(p) = \sup_{p \in \Pi} F(p)$, and we denote this supremum by S . As \hat{F} is continuous in Π , it achieves its supremum, hence, it suffices to show that $F(p_0) = S$, for every p_0 such that $\hat{F}(p_0) = S$ (note that the opposite implication is obvious). Assume the contrary, then $F(p) \leq S - \varepsilon$, for some $\varepsilon > 0$ and all $p \in \Pi \cap (p_0 - \varepsilon, p_0 + \varepsilon)$ by the upper semi-continuity of F . Then, we obtain $\hat{F}(p_0) \leq S - \varepsilon$, which is a contradiction. To see that F is upper semi-continuous, notice that it is left-continuous, with only downward jumps, which follows directly from (54). ■

Summarizing the above discussion, to find a solution to (56), it suffices to find a fixed point of the following correspondence

$$\mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi) \ni K \mapsto \tilde{K}(\tilde{\nu}(K)),$$

where

$$\tilde{\nu}(K; dx) = \mu(\{\alpha_0\}) \delta_{p^a}(dx) + K(\hat{\mathbb{A}} \times dx) \in \mathcal{M}_{\mu^a(\mathbb{A})}(\Pi) \quad (57)$$

is single-valued, and

$$\tilde{K}(\nu) = \operatorname{argmax}_{K \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)} \int \hat{F}(\alpha, p, \nu) K(d\alpha, dp) \subset \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi). \quad (58)$$

Proposition 2. *Let Assumptions 10, 11 hold. Then, the correspondence $\mathcal{K}: K \mapsto \tilde{K}(\tilde{\nu}(K))$, defined by (57)–(58), has a fixed point.*

Proof: To prove the proposition, we use the Kakutani's theorem for correspondences (cf. Definition II.7.8.1 and Theorem II.7.8.6 in [26]). Note that $\mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$, equipped with the weak topology, is convex and compact (by Prokhorov's theorem). In addition, it can be viewed as a subspace of the dual of the space of continuous functions on $\hat{\mathbb{A}} \times \Pi$, which is semi-normed. Thus, in order to apply the Kakutani's theorem, it only remains to show that \mathcal{K} is upper hemi-continuous (uhc), with nonempty compact convex values. Notice also that $\tilde{K}(\nu)$ is convex by definition (as an argmax of a linear functional on a convex set), hence, \mathcal{K} is convex-valued, and we only need to show that it is uhc, with non-empty compact values. As $p \mapsto \tilde{\nu}(p)$ is a continuous function, and a composition of a continuous function and a uhc correspondence is a uhc correspondence, it suffices to verify that $\nu \mapsto \tilde{K}(\nu)$ is a uhc non-empty compact valued correspondence. To achieve this, we use the classical Berge's theorem (cf. [41], section E.3), which reduces to problem to the continuity of the function

$$(K, \nu) \mapsto \phi(K, \nu) = \int \hat{F}(\alpha, p, \nu) K(d\alpha, dp), \quad (59)$$

on $\mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi) \times \mathcal{M}_{\mu^a(\mathbb{A})}(\Pi)$, metrized via the Lévy-Prokhorov metric. In the remainder of the proof, we show that $\phi(K, \nu)$ is jointly continuous in (K, ν) . More precisely, $\phi(K, \nu)$ is continuous in K , and it is continuous in ν (with respect to Lévy-Prokhorov metric), uniformly over K .

First, we show that $\phi(K, \nu)$ is continuous in K . By the definition of weak topology, the desired continuity would follow from the joint continuity of $\hat{F}(\alpha, p, \nu)$ with respect to (α, p) . Due to Lemma 9, $\hat{F}(\alpha, p, \nu)$ is 1-Lipschitz in p (uniformly over $\alpha \in \hat{\mathbb{A}}$), hence, it suffices to check that $\hat{F}(\alpha, p, \nu)$ is continuous in α . The latter follows from the fact that $F(\alpha, p, \nu)$ is continuous in α , uniformly over $p \in \Pi$. Indeed, notice that, if, for some $\alpha' \in U(\alpha)$, we have $|F(\alpha', p, \nu) - F(\alpha, p, \nu)| \leq \varepsilon \forall p \in \Pi$, then

$$\hat{F}(\alpha', p, \nu) = F(\alpha', p', \nu) - |p' - p| \leq F(\alpha, p', \nu) - |p' - p| + \varepsilon \leq \hat{F}(\alpha, p, \nu) + \varepsilon,$$

which, together with the analogous symmetric inequality, shows that $|\hat{F}(\alpha', p, \nu) - \hat{F}(\alpha, p, \nu)| \leq \varepsilon$. The first equality in the above follows from the fact that F is upper semi-continuous in p (and bounded from above by $2C_p$), which is shown in the proof of Lemma 9, and, hence, the supremum in the definition of \hat{F} is achieved at some p' . To show that $F(\alpha, p, \nu)$ is continuous in α , uniformly over $p \in \Pi$, we recall (54), and the desired continuity follows directly from Assumption 10.

It remains to show that $\phi(K, \nu)$ is continuous in $\nu \in \mathcal{M}_{\mu^a(\mathbb{A})}(\Pi)$, uniformly over $K \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$. As every such K has a fixed finite total mass, due to the definition of ϕ , the desired continuity follows from the fact that $\hat{F}(\alpha, p, \nu)$ is continuous in ν , uniformly over $(\alpha, p) \in \hat{\mathbb{A}} \times \Pi$. To prove the latter, fix $\varepsilon > 0$, and let d_0 be Lévy-Prokhorov metric on $\mathcal{M}_{\mu^a(\mathbb{A})}(\Pi)$. Let us show that there exists an increasing continuous deterministic function $C_0: [0, \infty) \rightarrow [0, \infty)$, s.t. $C_0(0) = 0$ and

$$|\hat{F}(\alpha, p, \nu_1) - \hat{F}(\alpha, p, \nu_2)| \leq C_0(\varepsilon), \quad \forall p \in \Pi, \alpha \in \hat{\mathbb{A}}, d_0(\nu_1, \nu_2) \leq \varepsilon.$$

If we manage to show that there exists an increasing continuous deterministic function $B: [0, \infty) \rightarrow [0, \infty)$, s.t. $B(0) = 0$ and

$$F(\alpha, p, \nu_1) \leq F(\alpha, (p - \varepsilon) \vee 0, \nu_2) + B(\varepsilon), \quad (60)$$

then

$$\begin{aligned} \hat{F}(\alpha, p, \nu_1) &= F(\alpha, p', \nu_1) - |p' - p| \leq F(\alpha, (p' - \varepsilon) \vee 0, \nu_2) - |p' - p| + B(\varepsilon) \\ &\leq F(\alpha, (p' - \varepsilon) \vee 0, \nu_2) - |(p' - \varepsilon) \vee 0 - p| + B(\varepsilon) + \varepsilon \leq \hat{F}(\alpha, p, \nu_2) + B(\varepsilon) + \varepsilon. \end{aligned}$$

The latter, together with the analogous inequality in which ν_1 and ν_2 are switched, yields the desired uniform continuity of \hat{F} in ν . Thus, it is only left to prove (60). For any $p \in \Pi$, by the definition of the Lévy-Prokhorov metric, we have:

$$\nu_1([0, p]) \geq \nu_2([0, (p - \varepsilon) \vee 0]) - \varepsilon$$

and, hence, by Assumption 11,

$$-D^{-1}(\nu_1([0, p])) \geq -D^{-1}(\nu_2([0, (p - \varepsilon) \vee 0])) - \epsilon(\varepsilon).$$

Then, for any $p \in \Pi$,

$$p + p^a - D^{-1}(\nu_1([0, p])) \geq (p - \varepsilon) \vee 0 + p^a - D^{-1}(\nu_2([0, (p - \varepsilon) \vee 0])) - \epsilon(\varepsilon),$$

which implies

$$F^{+, \alpha}(p + p^a - D^{-1}(\nu_1^+(p))) \leq F^{+, \alpha}((p - \varepsilon) \vee 0 + p^a - D^{-1}(\nu_2^+((p - \varepsilon) \vee 0))) + M_f \epsilon(\varepsilon),$$

where we used the fact that f^α is bounded by some constant M_f . The above estimate, along with the boundedness of p^a , p^b and $F^{+, \alpha}$, yields the desired inequality (60) for the first term in (54). Integrating the above estimate, we obtain the analogous inequality for the last term in the right hand side of (54), thus, completing the proof. ■

Proposition 2 implies that, for a.e. (t, ω) , we can find $K_{t, \omega} \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$, s.t.

$$K_{t, \omega} \in \tilde{K}(\tilde{\nu}(K_{t, \omega})),$$

and, hence, $(K_{t, \omega}, \tilde{\nu}(K_{t, \omega}))$ satisfies (56). Next, we need to establish the measurability of $K_{t, \omega}$ with respect to (t, ω) . Namely, we need to show that there exists a progressively measurable mapping $(t, \omega) \mapsto K_{t, \omega} \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$, such that

$$K_{t, \omega} \in \operatorname{argmax}_{K' \in \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)} \phi_{t, \omega}(K', \tilde{\nu}(K_{t, \omega})), \quad (61)$$

for Leb $\otimes \mathbb{P}$ -a.e. (t, ω) , where ϕ and $\tilde{\nu}$ are defined in (59) and (57). We denote $S = [0, T] \times \Omega$, and let \mathcal{S} be the progressive sigma-algebra (defined w.r.t. the filtration \mathbb{F}) on S . We also denote $\mathbb{X} = \mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$ and introduce the correspondence $g_1: S \times \mathbb{X} \rightarrow \mathbb{X}$, given by

$$(t, \omega, K) \mapsto \operatorname{argmax}_{K' \in \mathbb{X}} \phi_{t, \omega}(K', \tilde{\nu}(K)).$$

Notice that \mathbb{X} is separable and metrizable, and consider the function $(t, \omega, K, K') \mapsto \phi_{t, \omega}(K', \tilde{\nu}(K))$, defined on $(S \times \mathbb{X}^2, \mathcal{S} \otimes \mathcal{B}(\mathbb{X}^2))$. Note that this function is continuous in K' (as shown in the proof of Proposition 2) and measurable in (t, ω, K) (as it is continuous in K and measurable in (t, ω) , as shown in the proof of Proposition 2), hence, it is a Carathéodory function. Then, the Measurable Maximum theorem (cf. Theorem 18.18 in [1]) implies that g_1 is a $(\mathcal{S} \otimes \mathcal{B}(\mathbb{X}))$ -measurable correspondence with nonempty and compact values. Consider another correspondence $g_2: S \rightarrow \mathbb{X}$, given by

$$(t, \omega) \mapsto \{K \in \mathbb{X} : K \in \operatorname{argmax}_{K' \in \mathbb{X}} \phi_{t, \omega}(K', \tilde{\nu}(K))\}.$$

Let us show how to measurably select from g_2 , for Leb $\otimes \mathbb{P}$ -a.e. (t, ω) . The standard measurable selection results (cf. Corollary 18.27 and Theorem 18.26 in [1]) imply that such a selection is possible if g_2 has $\mathcal{S} \otimes \mathcal{B}(\mathbb{X})$ -measurable graph and non-empty values. The latter follows from Proposition 2, and the former is guaranteed by the following lemma.

Lemma 10. *The correspondence g_2 has a $\mathcal{S} \otimes \mathcal{B}(\mathbb{X})$ -measurable graph.*

Proof: Denote this graph by Γ_{g_2} . Let $I_{\mathbb{X}}: \mathbb{X} \rightarrow \mathbb{X} \times \mathbb{X}$ be given by $I_{\mathbb{X}}(K) = (K, K)$. Then, $\Gamma_{g_2} = (\operatorname{id} \times I_{\mathbb{X}})^{-1}(\Gamma)$, where $\Gamma \subset S \times \mathbb{X} \times \mathbb{X}$ is given by

$$\begin{aligned} \Gamma = \{ & (t, \omega, K, K') | (t, \omega) \in S, K \in \mathbb{X}, K' \in \operatorname{argmax}_{K'' \in \mathbb{X}} \phi_{t, \omega}(K'', \tilde{\nu}(K)) \} \\ & \cap \{ (t, \omega, K, K) | (t, \omega) \in S, K \in \mathbb{X} \}. \end{aligned}$$

Clearly, $\operatorname{id} \times I_{\mathbb{X}}$ is a measurable map, and the set $\{(t, \omega, K, K) | (t, \omega) \in S, K \in \mathbb{X}\}$ is measurable. Therefore, we only need to check that

$$\{(t, \omega, K, K') | (t, \omega) \in S, K \in \mathbb{X}, K' \in \operatorname{argmax}_{K'' \in \mathbb{X}} \phi_{t, \omega}(K'', \tilde{\nu}(K))\}$$

is $\mathcal{S} \otimes \mathcal{B}(\mathbb{X}^2)$ -measurable. The latter set is precisely the graph of g_1 , and it is measurable as the correspondence g_1 is measurable (cf. Theorem 18.6 in [1]). ■

Thus, we conclude that there exists a progressively measurable K , with values in $\mathcal{M}_{\mu^a}(\hat{\mathbb{A}} \times \Pi)$, satisfying (61) for $\text{Leb} \otimes \mathbb{P}$ -a.e. (t, ω) . It only remains to construct κ from K , by disintegration. Let us introduce $A = S \times \hat{\mathbb{A}}$, equipped with the sigma-algebra $\mathcal{S} \otimes \mathcal{B}(\hat{\mathbb{A}})$, and the measure \mathbb{Q} on $A \times \Pi$, defined via $\mathbb{Q}(dt, d\omega, d\alpha, dp) = K_{t,\omega}(d\alpha, dp)dt\mathbb{P}(d\omega)$. Note that the marginal distribution of \mathbb{Q} on A is $\mu^a(d\alpha)dt\mathbb{P}(d\omega)$. Then, as the natural projection from $A \times \Pi$ to Π has a Borel range, Theorems 5.3 and 5.4 from [32] imply that there exists a kernel $\kappa: A \ni (t, \omega, \alpha) \mapsto \kappa_{t,\omega}(\alpha) \in \mathcal{P}(\Pi)$, which is a regular conditional distribution of the natural projection from $A \times \Pi$ to Π , given the natural projection from $A \times \Pi$ to $\hat{\mathbb{A}}$, under \mathbb{Q} . Namely, for every absolutely bounded measurable $f: A \times \Pi \rightarrow \mathbb{R}$, we have

$$\int_{A \times \Pi} f(t, \omega, \alpha, p) K_{t,\omega}(d\alpha, dp) dt \mathbb{P}(d\omega) = \int_{A \times \Pi} f(t, \omega, \alpha, p) \kappa_{t,\omega}(\alpha; dp) \mu^a(d\alpha) dt \mathbb{P}(d\omega). \quad (62)$$

The above property yields that $\hat{\nu}_{t,\omega}^a = \tilde{\nu}(K_{t,\omega})$ and $\kappa_{t,\omega}$ satisfy the fixed-point constraint (55). It only remains to show that κ satisfies (53), for $\text{Leb} \otimes \mathbb{P} \otimes \mu^a$ -a.e. (t, ω, α) . Assume that this is not the case, then, there exists a measurable set $B \subset [0, T] \times \Omega$, with positive measure, s.t. for any fixed $(t, \omega) \in B$, there exists a measurable set $C \subset \hat{\mathbb{A}}$, s.t. $\mu^a(C) > 0$ and, for all $\alpha \in C$,

$$\begin{aligned} \int_{\mathbb{R}} \hat{F}_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \kappa_{t,\omega}(\alpha; dp) &\leq \int_{\mathbb{R}} F_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \kappa_{t,\omega}(\alpha; dp) \\ &< \sup_{\kappa' \in \psi} \int_{\mathbb{R}} F_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \kappa'(\alpha; dp) = \sup_{\kappa' \in \psi} \int_{\mathbb{R}} \hat{F}_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \kappa'(\alpha; dp). \end{aligned}$$

The above inequality becomes non-strict for all $\alpha \in \hat{\mathbb{A}} \setminus C$. Then, for a fixed $(t, \omega) \in B$, we can choose a measurable $\tilde{\kappa}: \hat{\mathbb{A}} \rightarrow \mathcal{P}(\Pi)$ (in the same way as we chose a measurable K , except that, in this case, the measurability is required in the α -variable), s.t.

$$\sup_{\kappa' \in \psi} \int_{\mathbb{R}} \hat{F}_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \kappa'(\alpha; dp) = \int_{\mathbb{R}} \hat{F}_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \tilde{\kappa}(\alpha; dp), \quad \mu^a\text{-a.e. } \alpha \in \hat{\mathbb{A}}.$$

Thus, we obtain

$$\int_{\mathbb{R}} \hat{F}_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \kappa_{t,\omega}(\alpha; dp) < \int_{\mathbb{R}} \hat{F}_{t,\omega}(\alpha, p, \tilde{\nu}(K_{t,\omega})) \tilde{\kappa}(\alpha; dp),$$

for all $\alpha \in C$, and the non-strict inequality holds for all $\alpha \in \hat{\mathbb{A}}$. Integrating with respect to μ^a , and using (62) with $f(t, \omega, \alpha, p) = \hat{F}(t, \omega, \alpha, p, \tilde{\nu}(K_{t,\omega}))$, we obtain a contradiction with (61) on the set B (which has a positive measure). Thus, for μ^a -a.e. $\alpha \in \hat{\mathbb{A}}$, (53) holds for $\text{Leb} \otimes \mathbb{P}$ -a.e. (t, ω) . This means that, if we define $\hat{p}_t(\alpha)$ as the push-forward of $\kappa_t(\alpha)$, under the mapping $x \mapsto x + p_t^a$, the resulting strategy $\hat{p}(\alpha)$ maximizes the generator $\hat{\mathcal{G}}_t^\alpha(y)$, for any y and a.e. (t, ω) . Then, we define ν_t^a to be the push-forward of $\hat{\nu}_t^a$, under the mapping $x \mapsto x + p_t^a$, and use the standard BSDE results to conclude that, for μ^a -a.e. $\alpha \in \hat{\mathbb{A}}$,

$$J^{(\nu, \theta), (\hat{p}(\alpha), V^a)}(1, \alpha) = \bar{J}_0^{\alpha, (\hat{p}(\alpha))} \geq \bar{J}_0^{\alpha, (p')} = J^{(\nu, \theta), (p', V^a)}(1, \alpha)$$

holds for all admissible strategies p' , which means that $\hat{p}(\alpha)$ is optimal for the long agents with beliefs α . With such a choice of ν^a and \hat{p} , the fixed-point condition on ν^a , given in (10), is satisfied, as it is equivalent to (55) (assuming the extremal long agents post limit orders at p^a , which is optimal for them). This, along with Corollary 1, implies that $(\hat{p}(\alpha), V^a)$ is an optimal strategy for the long agents with beliefs $\alpha \in \hat{\mathbb{A}}$. The short agents are treated similarly. Thus, we complete the proof of Theorem 1.

5 Example

In this section, we consider the simplest concrete example of our model and show how it can be used. Consider a stochastic basis $(\Omega, \tilde{\mathbb{F}} = (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$, with a Poisson random measure N , whose compensator is $\lambda_t f_t(x) dx dt$,

as described in Subsection 2.1. We assume that $J_t(x) = x$ (i.e. $M \equiv N$), so that N is the jump measure of the (potential) fundamental price process X . We also assume that $T = 20$, $\lambda_t \equiv 1$ and f_t is the density of a uniform distribution on $[-C_0, C_0]$, where the constant C_0 is chosen to be sufficiently large, so that this interval contains the supports of all f^α described below. We take $\mathbb{A} = \{\alpha_0\} \cup \hat{\mathbb{A}}$, $\mathbb{B} = \{\beta_0\} \cup \hat{\mathbb{B}}$, where

$$\hat{\mathbb{A}} = \left\{ \frac{i}{K} \mid 0 \leq i < K \right\}, \quad \hat{\mathbb{B}} = \left\{ -\frac{i}{K} \mid 0 \leq i < K \right\}$$

are the uniform partitions of unit intervals, and $K = 500$ is used for most of the computations herein. The restrictions of μ^a (resp. μ^b) on $\hat{\mathbb{A}}$ (resp. $\hat{\mathbb{B}}$) assign a mass of $1/K$ to every point of the corresponding discrete space. Note that this implies $\mu^a(\hat{\mathbb{A}}) = \mu^b(\hat{\mathbb{B}}) = 1$. We also define $\mu^a(\{\alpha^0\}) = \mu^b(\{\beta^0\}) = 0.1$.

Next, we consider a collection of positive numbers $\{\lambda^{+, \alpha}, \lambda^{-, \alpha}, C^{+, \alpha}, C^{-, \alpha}\}_{\alpha \in \mathbb{A} \cup \mathbb{B}}$, and define

$$f^\alpha(x) = \frac{\lambda^{+, \alpha}}{(\lambda^{+, \alpha} + \lambda^{-, \alpha})C^{+, \alpha}} \mathbf{1}_{[0, C^{+, \alpha}]}(x) + \frac{\lambda^{-, \alpha}}{(\lambda^{+, \alpha} + \lambda^{-, \alpha})C^{-, \alpha}} \mathbf{1}_{[-C^{-, \alpha}, 0]}(x), \quad \lambda^\alpha = \lambda^{+, \alpha} + \lambda^{-, \alpha}.$$

Herein, we use $C^{+, \alpha_0} = C^{-, \alpha_0} = C^{+, \beta_0} = C^{-, \beta_0} = 0.5$ and

$$C^{+, \alpha} = a + b\alpha, \quad C^{-, \alpha} = C^{-, \alpha_0}, \quad \forall \alpha \in \hat{\mathbb{A}}, \quad C^{-, \beta} = a - b\beta, \quad C^{+, \beta} = C^{+, \beta_0}, \quad \forall \beta \in \hat{\mathbb{B}},$$

with $a = 0.5$ and $b = 10$. Finally, for any $\alpha \in \mathbb{A} \cup \mathbb{B}$, we introduce

$$\Gamma^\alpha(x) = \frac{\lambda^\alpha f^\alpha(x)}{\lambda f(x)} - 1, \quad dZ_t^\alpha = Z_{t-}^\alpha \int_{\mathbb{R}} \Gamma^\alpha(x) [N(dt, dx) - \lambda f(x) dt dx],$$

and define $\mathbb{P}^\alpha \ll \mathbb{P}$ by its Radon-Nikodym density Z_T^α . One can easily check, using the general results in [30] (or in [16], for the deterministic case, used herein) that, under such \mathbb{P}^α , N is a Poisson random measure with the compensator $\lambda^\alpha f^\alpha(x) dx dt$.

We assume that the demand elasticity is deterministic, constant in time, and linear in price:

$$D_t(p) = -kp,$$

with the elasticity parameter $k = 0.2$. With the above choice of $(C^{\pm, \alpha_0}, C^{\pm, \beta_0}, \mu^a(\{\alpha^0\}), \mu^b(\{\beta^0\}), k)$, it is easy to see that Assumption 9 is satisfied. Notice that the choice of $\lambda^{\pm, \alpha}$, for $\alpha \in \hat{\mathbb{A}} \cup \hat{\mathbb{B}}$, does not affect the equilibrium, as long as Assumptions 7 and 8 are satisfied. This is, clearly, the case if we choose $\lambda^{\pm, \alpha} = \lambda^{\pm, \alpha_0}$ and $\lambda^{\pm, \beta} = \lambda^{\pm, \beta_0}$, for $\alpha \in \mathbb{A}$ and $\beta \in \mathbb{B}$. Herein, we consider several different sets of values for $(\lambda^{\pm, \alpha_0}, \lambda^{\pm, \beta_0})$.

Let us construct an equilibrium in this example. Notice that, in the present case, the Brownian motion W does not affect the jump intensities and, in turn, the agents' objectives, hence, the RBSDE system (44) becomes a system of reflected ODEs. We can solve it easily, using a simple Euler scheme, then, recover the value functions (V^a, V^b) , as shown in Lemma 4, and construct the bid and ask prices, (p^a, p^b) , in the feedback form, as shown in Lemma 3. We implement this strategy with the parameters chosen above, and with $\lambda^{+, \alpha_0} = 2.5$, $\lambda^{-, \alpha_0} = 1$, $\lambda^{+, \beta_0} = 1$, $\lambda^{-, \beta_0} = 2.5$ (so that the extremal ask agents are bullish whereas the extremal bid agents are bearish). The results are shown in the left part of Figure 1. Using the same parameters, we consider the book beyond the best bid and ask prices. In order to construct it, we solve the fixed-point problem (56) numerically. The latter is achieved by limiting the set of possible price levels for the limit orders to a finite set (i.e. to a partition of a large interval), which reduces (56) to a finite-dimensional fixed-point problem. In addition, we allow each agent to post a limit order at a single price level only, which further simplifies the problem.¹⁵ Thus, we find a solution by the standard recursive iteration, maximizing, at each step, the objective over a finite set. The resulting optimal limit order strategies of the agents (at time zero) are plotted in the right part of Figure 1, as a function of the agents' beliefs $\alpha \in \hat{\mathbb{A}} \cup \hat{\mathbb{B}}$. Notice that the optimal limit order strategy $p(\cdot)$ is piece-wise constant. It is worth mentioning

¹⁵Note that this restriction does not compromise the optimality of the agents' actions, provided a fixed point can be found. Indeed, it is a well known phenomenon that, in a continuum-player game, an equilibrium with pure controls also provides an equilibrium for a setting with distributed controls. This is, in fact, one of the advantages of the continuum-player games. We consider distributed controls only to prove that the equilibrium does exist, which is much harder (if at all possible) to show for a setting with pure controls.

that this discreteness seems to be inherent in the model and not just an artifact of the discretization of prices or beliefs that we chose herein, as the results do not change when we increase the number of possible beliefs (K) and price levels.

Finally, we demonstrate how the proposed framework can be used to model the *indirect market impact*, which appears when an initial change to the LOB creates “feedback loop” and causes further changes. Note that the initial change may be triggered by a trade (which is the case in the classical models of optimal execution) or by a new limit order. An extreme example of the latter is the so-called “spoofing” – i.e. posting a large limit order with the goal to make the price of the asset move in the opposite direction.¹⁶ To the best of our knowledge, to date, there exists no model capable of explaining how exactly this activity causes the LOB (and, in particular, the price) to change. To model this process, we modify the present example by assuming that $(\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0})$ are, in fact, functions of a relevant market indicator, which we denote by I :

$$\lambda^{+, \alpha^0} = 2.3 \exp(I s), \quad \lambda^{-, \alpha^0} = \exp(-I s), \quad \lambda^{+, \beta^0} = \exp(I s), \quad \lambda^{-, \beta^0} = 2.3 \exp(-I s), \quad (63)$$

where $s = 2.6$ is the sensitivity. We further assume that I is the so-called *market imbalance*: the ratio of the size of all limit orders at the best bid over the size of all limit orders at the best ask, less one. It is a well known empirical fact (cf. [14], [11], [39]) that such an indicator has a predictive power for the direction of the next price move. Note that I is a function of the LOB, which, in turn, is an outcome of an equilibrium, in which I is the input. Strictly speaking, our results do not guarantee the existence of an equilibrium with this additional fixed-point constraint. In fact, an equilibrium with “feedback beliefs”, given by (63), can be viewed as a fixed-point of the following mapping:

$$(\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0}) \mapsto \nu \mapsto I \mapsto (\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0}), \quad (64)$$

where $(\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0}) \mapsto \nu$ maps the numbers $(\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0})$ into an equilibrium LOB ν , as it is done in the first part of this section, and $I \mapsto (\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0})$ is given by (63). Herein, we do not prove a general existence result for the aforementioned fixed point, but we can compute it numerically by applying the associated mapping iteratively (assuming the iterations do converge). In particular, the top right part of Figure 2 shows an example of LOB arising in equilibrium with feedback beliefs, given by (63), with $I_0 = .0984456$.

Our next goal is to show how the market may move from one equilibrium to another, once the LOB is perturbed (assuming the feedback beliefs (63)). It is worth mentioning that there is no canonical way to describe how agents achieve an equilibrium. Nevertheless, we propose a specific algorithm, based on the iterations of (64), with the following rationale behind it. For any parameters $\bar{\lambda} = (\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0})$ (given as functions of time), the agents know their equilibrium strategies: $(p(\bar{\lambda}), v(\bar{\lambda}))$, which can be computed as shown in the first part of this example (and whose existence follows from the main result of this paper). If the LOB is perturbed, I changes, and, in turn, the parameters change from $\bar{\lambda}$ to $\bar{\lambda}'$, via (63). Then, the agents change their strategies to $(p(\bar{\lambda}'), v(\bar{\lambda}'))$, which form an equilibrium with respect to the new set of parameters $\bar{\lambda}'$. In the new equilibrium, the LOB, and, hence, the imbalance I , may change, causing further change to the parameters, and so on, until the agents reach a set of parameters that coincides with the previous one (or, almost coincides, from a numerical point of view). We believe that this algorithm for moving to a new equilibrium makes economic sense, although, of course, it is not the only possible choice. Mathematically, it corresponds to iterating the mapping (64). To illustrate this approach, we add an extra limit buy order of size 0.05, located at the best bid price, to the previously obtained equilibrium LOB – as shown in the bottom right part of Figure 2. This implies a change to the imbalance I and, in turn, to the agents’ parameters $(\lambda^{\pm, \alpha^0}, \lambda^{\pm, \beta^0})$, via (63). Hence, the agents adjust their controls to reach a new equilibrium, then, re-calculate the parameters with the new imbalance, and so on. Figure 3 shows what happens to the LOB and to the functions (V^a, V^b) in the first five iterations. We can see that the initial change in imbalance makes the agents more bullish about the asset, and they tend to move their limit orders higher. In particular, the size of the best bid queue increases, while the size of the best ask queue decreases, further increasing the market imbalance. The left part of Figure 3 also shows that, starting from step three, the value functions V^a and V^b coincide at time zero, which means that the agents, in fact, choose to submit an internal market order, terminating the game. The latter constitutes an equilibrium with feedback beliefs (63).¹⁷ This experiment, in particular, shows why the predictive

¹⁶We stress that intentional spoofing is an illegal activity.

¹⁷It is important to notice that an equilibrium with feedback beliefs (63), typically, is not unique, but the proposed algorithm leads to a specific one. The resulting equilibrium is degenerate, in the sense that the game ends immediately, but, of course, there exist other equilibria.

power of market imbalance is a “self-fulfilling prophecy”: the fact that the agents base their beliefs about the type of the next market order on the market imbalance, itself, implies that a sufficient change in market imbalance will, indeed, trigger a market order of the right type.

Of course, the analysis provided in the second part of this section is merely an example, which is meant to illustrate a potential application of our theoretical results. Namely, our main results show that a single iteration of the mapping (64) is well defined. Nevertheless, a rigorous analysis of the resulting iterative scheme, including its convergence to a fixed point, is missing. In general, it would also be interesting to consider other indicators: e.g. choosing the size and direction of the last trade as the relevant indicator, would allow one to model the indirect impact of a market order on the LOB (in addition to the obvious, direct, impact, resulting from the immediate execution of limit orders). In our future research, we plan to find appropriate model specifications which would allow us to conduct a more thorough analysis of the indirect market impact, within the proposed setting, and to test the predictions of our model against the market data.

6 Appendix

Proof of Lemma 7. We consider a long agent with beliefs α and introduce

$$\begin{aligned} \bar{J}_t^{\alpha, (p, \tau)} &= \mathbb{E} \left[\int_t^\tau \exp \left(- \int_t^s \bar{c}_u^\alpha (p_u^a \wedge Q^-(p_u), p_u^b) du \right) \bar{h}_s^{\alpha, a} (p_s, p_s^a, p_s^b) ds \right. \\ &\quad \left. + \exp \left(- \int_t^\tau \bar{c}_u^\alpha (p_u^a \wedge Q^-(p_u), p_u^b) du \right) p_{\tau \wedge \hat{\tau}}^b | \mathcal{F}_t \right], \end{aligned}$$

where

$$\bar{c}_t^\alpha(x, y) = c_t^\alpha(x, y) \mathbf{1}_{\{t \leq \hat{\tau}\}}, \quad \bar{h}_t^{\alpha, a}(\kappa, x, y) = h_t^{\alpha, a}(\kappa, x, y) \mathbf{1}_{\{t \leq \hat{\tau}\}}, \quad x, y \in \mathbb{R} \quad \kappa \in \mathcal{P}(\mathbb{R}),$$

with c^α and $h^{\alpha, a}$ defined in (13) and (16). Next, for any $t \in [0, T]$, any $\alpha \in \mathbb{A}$, and any admissible p , we introduce

$$Y_t^{\alpha, p} = \text{ess sup}_{\tau \in \mathcal{T}_t} \bar{J}_t^{\alpha, (p, \tau)}, \quad (65)$$

The standard results on RBSDEs imply that $Y^{\alpha, p}$ is the unique \mathbb{S}^2 solution of the affine RBSDE,

$$-dY_t^{\alpha, p} = \bar{\mathcal{G}}_t^{\alpha, p}(Y_t^{\alpha, p}) dt - Z_t dW_t + dK_t \quad 0 \leq t \leq T \quad (66)$$

$$Y_t^{\alpha, p} \geq p_{t \wedge \hat{\tau}}^b \quad 0 \leq t \leq T, \quad \int_0^T (Y_t^{\alpha, p} - p_{t \wedge \hat{\tau}}^b) dK_t = 0 \quad (67)$$

$$Y_T^{\alpha, p} = p_{\hat{\tau}}^b, \quad (68)$$

where

$$\bar{\mathcal{G}}_t^{\alpha, p}(y) = -\bar{c}_t^\alpha(p_t^a \wedge Q^-(p_t), p_t^b) y + \bar{h}_t^{\alpha, a}(p_t, p_t^a, p_t^b) = [-c_t^\alpha(p_t^a \wedge Q^-(p_t), p_t^b) y + h_t^{\alpha, a}(p_t, p_t^a, p_t^b)] \mathbf{1}_{\{t < \hat{\tau}\}},$$

with c^α and $h^{\alpha, a}$ defined in (13) and (16). Recall that V^a satisfies (43), with the generator

$$\mathcal{G}_t^a(y, p_t^b) = 2\lambda_t^{\alpha^0} p_t^b F_t^{\alpha^0, -}(p_t^b) - \lambda_t^{\alpha^0} F_t^{\alpha^0, -}(p_t^b) y + \lambda_t^{\alpha^0} P_t^a(y) F_t^{\alpha^0, +}(P_t^a(y)) - \lambda_t^{\alpha^0} F_t^{\alpha^0, +}(P_t^a(y)) y.$$

It is easy to deduce that

$$\bar{\mathcal{G}}_t^{\alpha^0, p^a}(V_t^a) = \mathcal{G}_t^a(V_t^a, p_t^b) \mathbf{1}_{\{t < \hat{\tau}\}}.$$

Hence, $(V_{t \wedge \hat{\tau}}^a)$ satisfies the same RBSDE as $(Y_t^{\alpha^0, p^a})$. From the comparison principle, we conclude that $Y_t^{\alpha^0, p^a} = V_{t \wedge \hat{\tau}}^a$. On the other hand, for any $\alpha \in \mathbb{A}$, let us choose $p_t = \delta_{p_t^a}$, to obtain:

$$\bar{\mathcal{G}}_t^{\alpha, p^a}(Y_t^{\alpha^0, p^a}) = \bar{\mathcal{G}}_t^{\alpha, p^a}(V_t^a) = \left[\lambda_t^\alpha p_t^b F_t^{\alpha, -}(p_t^b) + \lambda_t^\alpha \int_{-\infty}^{p_t^b} f_t^\alpha(u) l_t^{c, b}(u) du - \lambda_t^\alpha F_t^{\alpha, -}(p_t^b) V_t^a \right]$$

$$\begin{aligned}
& + \lambda_t^\alpha p_t^a F_t^{+, \alpha}(p_t^a) - \lambda_t^\alpha F_t^{+, \alpha}(p_t^a) V_t^a \mathbf{1}_{\{t < \hat{\tau}\}} \\
& \geq \left[\lambda_t^\alpha F_t^{-, \alpha}(p_t^b)(p_t^b - V_t^a) + \lambda_t^\alpha F_t^{+, \alpha}(p_t^a)(p_t^a - V_t^a) + \lambda_t^{\alpha^0} p_t^b F_t^{-, \alpha^0}(p_t^b) \right] \mathbf{1}_{\{t < \hat{\tau}\}},
\end{aligned}$$

where $l^{c,b}$ is defined in (15), and the last inequality is based on the Assumptions 7, 9, and on the monotonicity of $l_t^{c,b}(\cdot)$, which imply

$$\lambda_t^\alpha \int_{-\infty}^{p_t^b} f_t^\alpha(u) l_t^{c,b}(u) du \geq \lambda_t^{\alpha^0} \int_{-\infty}^{p_t^b} f_t^{\alpha^0}(u) l_t^{c,b}(u) du = \lambda_t^{\alpha^0} p_t^b F_t^{-, \alpha^0}(p_t^b).$$

Notice that, by construction, $p_t^b \leq V_t^b \leq V_t^a \leq p_t^a$. Then, Assumption 7 implies

$$\lambda_t^\alpha F_t^{-, \alpha}(p_t^b)(2p_t^b - V_t^a) \geq \lambda_t^{\alpha^0} F_t^{-, \alpha^0}(p_t^b)(2p_t^b - V_t^a), \quad \lambda_t^\alpha F_t^{+, \alpha}(p_t^a)(p_t^a - V_t^a) \geq \lambda_t^{\alpha^0} F_t^{+, \alpha^0}(p_t^a)(p_t^a - V_t^a).$$

Thus, we obtain:

$$\bar{\mathcal{G}}_t^{\alpha^0, p^a} \left(Y_t^{\alpha^0, p^a} \right) \leq \bar{\mathcal{G}}_t^{\alpha, p^a} \left(Y_t^{\alpha^0, p^a} \right).$$

Using the comparison principle for RBSDEs, we conclude that $Y_t^{\alpha, p^a} \geq Y_t^{\alpha^0, p^a} = V_{t \wedge \hat{\tau}}^a$. Consider an arbitrary strategy (p, τ) . By switching between p^a and p , we can construct a new strategy p' , such that $Y_t^{\alpha, p'} \geq V_{t \wedge \hat{\tau}}^a \vee Y_t^{\alpha, p}$, for all t . More precisely, we define

$$\bar{\mathcal{G}}_t^{\alpha, p'}(y) = \bar{\mathcal{G}}_t^{\alpha, p^a}(y) \vee \bar{\mathcal{G}}_t^{\alpha, p}(y),$$

and solve the RBSDE (66)–(68). By the standard argument, the Y -component of the solution is $Y^{\alpha, p'}$, where p'_t is defined to be equal to $\delta_{p_t^a}$ if the maximum in the above equation is achieved at $\bar{\mathcal{G}}_t^{\alpha, p^a}(Y_t^{\alpha, p'})$, and it is equal to p_t otherwise. The comparison principle implies that $Y_t^{\alpha, p'} \geq Y_t^{\alpha, p^a} \vee Y_t^{\alpha, p} \geq V_{t \wedge \hat{\tau}}^a \vee Y_t^{\alpha, p}$. Then, the standard results on RBSDEs imply that the optimal stopping time associated with $Y^{\alpha, p'}$ is

$$\inf\{t \in [0, T] : Y_t^{\alpha, p'} \leq p_{t \wedge \hat{\tau}}^b\} = \inf\{t \in [0, T] : V_t^a \leq p_{t \wedge \hat{\tau}}^b\} = \hat{\tau}.$$

Thus,

$$J_0^{(\nu, \theta), (p, \tau)}(1, \alpha) = \bar{J}_0^{\alpha, (p, \tau)} \leq Y_0^{\alpha, p} \leq Y_0^{\alpha, p'} = \bar{J}_0^{\alpha, (p', \hat{\tau})} = J_0^{(\nu, \theta), (p', \hat{\tau})}(1, \alpha).$$

Next, we show that the control p can be chosen so that, \mathbb{P} -a.s., for all t , $\text{supp}(p_t) \subset [p_t^a, \infty)$. Consider any control p . By switching, if necessary, between p^a and p , we can ensure that $Y_t^{\alpha, p} \geq V_{t \wedge \hat{\tau}}^a$. Then, for $t < \hat{\tau}$, the generator of $Y^{\alpha, p}$ is given by

$$\begin{aligned}
\bar{\mathcal{G}}_t^{\alpha, p}(y) &= -c_t^\alpha(p_t^a \wedge Q^-(p_t), p_t^b) y + h_t^{\alpha, a}(p_t, p_t^a, p_t^b) \\
&= -\lambda_t^\alpha F_t^{+, \alpha}(p_t^a \wedge Q^-(p_t)) y + \lambda_t^\alpha p_t^b F_t^{+, \alpha}(Q^-(p_t) \wedge p_t^a) p_t((p_t^a, \infty)) \\
&+ \lambda_t^\alpha \int_{(Q^-(p_t) \wedge p_t^a) \vee 0}^{p_t^a} f_t^\alpha(u) \int_{-\infty}^{p_t^a} [(z \wedge u + (p_t^b - u) \mathbf{1}_{\{z > u\}})] p_t(dz) du + \lambda_t^\alpha F_t^{+, \alpha}(p_t^a) \int_{-\infty}^{p_t^a} z p_t(dz) \\
&- \lambda_t^\alpha F_t^{-, \alpha}(p_t^b) y + 2\lambda_t^\alpha p_t^b F_t^{-, \alpha}(p_t^b) - \lambda_t^\alpha p_t^b F_t^{+, \alpha}(p_t^a) p_t((p_t^a, \infty)) \\
&+ \lambda_t^\alpha \int_{p_t^a}^\infty f_t^\alpha(u) \int_{p_t^a}^\infty [(z \wedge l_t^{c, a}(u) + p_t^b \mathbf{1}_{\{z > l_t^{c, a}(u)\}})] p_t(dz) du
\end{aligned}$$

Let us estimate the first four terms in the right hand side of the above (i.e. the ones that depend on $p_t(dx)$ restricted to $x < p_t^a$):

$$\begin{aligned}
& -\lambda_t^\alpha F_t^{+, \alpha}(p_t^a \wedge Q^-(p_t)) y + \lambda_t^\alpha p_t^b F_t^{+, \alpha}(Q^-(p_t) \wedge p_t^a) p_t((p_t^a, \infty)) \\
& + \lambda_t^\alpha \int_{(Q^-(p_t) \wedge p_t^a) \vee 0}^{p_t^a} f_t^\alpha(u) \int_{-\infty}^{p_t^a} [(z \wedge u + (p_t^b - u) \mathbf{1}_{\{z > u\}})] p_t(dz) du + \lambda_t^\alpha F_t^{+, \alpha}(p_t^a) \int_{-\infty}^{p_t^a} z p_t(dz) \\
& \leq \lambda_t^\alpha \sup_{x \leq p_t^a} \left[(-y + p_t^b) F_t^{+, \alpha}(x) + p_t((-\infty, p_t^a]) \sup_{z \in [x, p_t^a]} [(z - p_t^b) F_t^{+, \alpha}(z)] \right].
\end{aligned}$$

Notice that, for $t < \hat{\tau}$ and $y = Y_t^{\alpha,p}$, we have $p_t^b - y \leq 0$, and, hence,

$$\begin{aligned} & \sup_{x \leq p_t^a} \left[(-y + p_t^b) F_t^{+, \alpha}(x) + p_t((-\infty, p_t^a]) \sup_{z \in [x, p_t^a]} \left[(z - p_t^b) F_t^{+, \alpha}(z) \right] \right] \\ &= \sup_{z \leq p_t^a} \left[(z - y) F_t^{+, \alpha}(z) + p_t^b p_t((p_t^a, \infty)) F_t^{+, \alpha}(z) \right]. \end{aligned}$$

Due to Lemma 6, the function $z \mapsto (z - y) F_t^{+, \alpha}(z)$ is nondecreasing in $z \leq P_t^a(y)$. As $p_t^b \leq 0$, the function $z \mapsto p_t^b p_t((p_t^a, \infty)) F_t^{+, \alpha}(z)$ is also nondecreasing, and, hence, the above supremum is attained at $z = p_t^a$, provided $P_t^a(y) \geq p_t^a$. The latter does hold for $t < \hat{\tau}$ and $y = Y_t^{\alpha,p}$, as $P_t^a(\cdot)$ is non-decreasing, $p_t^a = P_t^a(V_t^a)$ and $Y_t^{\alpha,p} \geq V_t^a$. Thus, the generator $\bar{\mathcal{G}}_t^{\alpha,p}(Y_t^{\alpha,p})$ does not decrease if we replace p by

$$p'_t(dx) = p_t(dx) \mathbf{1}_{[p_t^a, \infty)} + p_t((-\infty, p_t^a)) \delta_{p_t^a}(dx).$$

In other words,

$$\bar{\mathcal{G}}_t^{\alpha,p}(Y_t^{\alpha,p}) \leq \bar{\mathcal{G}}_t^{\alpha,p'}(Y_t^{\alpha,p}).$$

The comparison principle, then, yields $Y_t^{\alpha,p} \leq Y_t^{\alpha,p'}$. Moreover, the optimal stopping strategy associated with $Y^{\alpha,p'}$ is $\hat{\tau}$. Repeating the argument used earlier in this proof, we conclude that any strategy $(p, \hat{\tau})$ can be modified to $(p', \hat{\tau})$, satisfying the properties stated in the lemma, so that the objective value does not decrease. The case of short agents is treated similarly. ■

Proof of Lemma 8 Consider a long agent with beliefs α^0 . In view of Corollary 1, it suffices to show the optimality in the class of strategies $(p, \hat{\tau})$, with $\text{supp}(p_t) \subset [p_t^a, \infty)$. Notice that Assumption 9 implies:

$$l_t^{c,a}(x) = x \wedge p_t^a, \quad \forall x \in \text{supp}\left(f_t^{\alpha^0}\right).$$

Using the above observation, we recall the constructions from the proof of Lemma 7, to obtain, for any strategy p and all $t < \hat{\tau}$:

$$\begin{aligned} \bar{\mathcal{G}}_t^{\alpha^0,p}(y) &= -\lambda_t^{\alpha^0} F_t^{+, \alpha^0}(p_t^a) y - \lambda_t^{\alpha^0} F_t^{-, \alpha^0}(p_t^b) y + 2\lambda_t^{\alpha^0} p_t^b F_t^{-, \alpha^0}(p_t^b) \\ &\quad + \lambda_t^{\alpha^0} F_t^{+, \alpha^0}(p_t^a) (p_t^a p_t(\{p_t^a\}) + (p_t^a + p_t^b) p_t((p_t^a, \infty))). \end{aligned}$$

As $p_t^b \leq 0$, the above expression is maximized at $p_t = \delta_{p_t^a}$. Using the comparison principle for the RBSDE satisfied by $Y^{\alpha^0,p}$, we conclude that $p = \delta_{p_t^a}$ produces the largest $Y^{\alpha^0,p}$ and, hence, the largest objective value for the long agents with beliefs α^0 . The case of short agents is treated similarly. ■

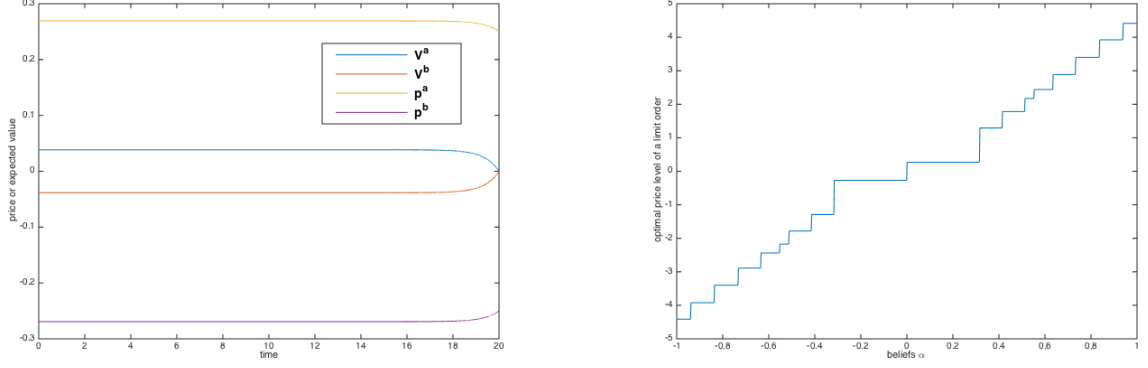


Figure 1: On the left: value functions (V^b, V^a) (red and blue), and the bid and ask prices (p^b, p^a) (purple and orange), as functions of time. On the right: the optimal price level of a limit order, as a function of the beliefs $\alpha \in \hat{\mathbb{A}} \cup \hat{\mathbb{B}}$. Parameters: $\lambda^{+, \alpha_0} = 2.5$, $\lambda^{-, \alpha_0} = 1$, $\lambda^{+, \beta_0} = 1$, $\lambda^{-, \beta_0} = 2.5$.




Figure 2: Left: LOB at time zero, with $\lambda^{+, \alpha_0} = 2.5$, $\lambda^{-, \alpha_0} = 1$, $\lambda^{+, \beta_0} = 1$, $\lambda^{-, \beta_0} = 2.5$. Right: equilibrium LOB at time zero, with the parameters depending on the market imbalance I (top), and the same LOB, with an additional (yellow) limit order (bottom).

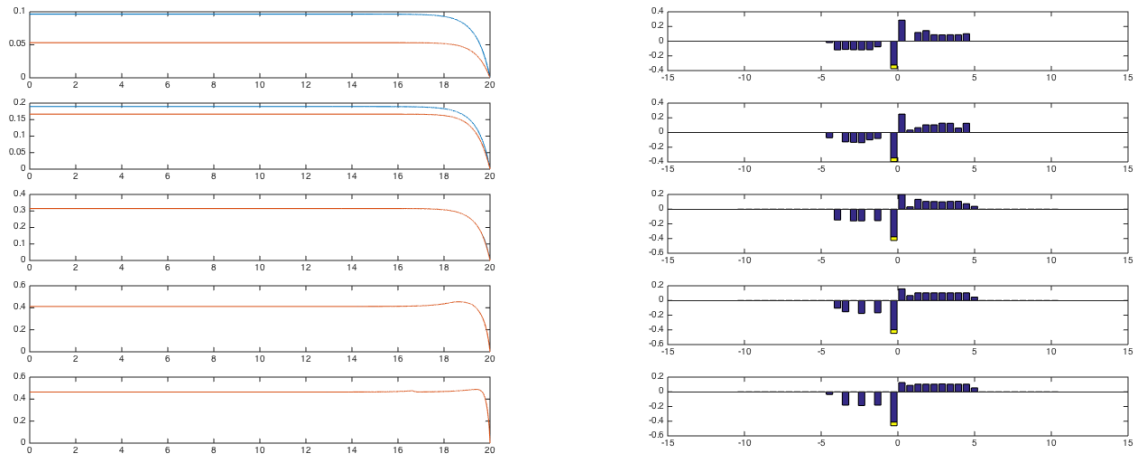


Figure 3: On the left: value functions (V^b, V^a) (red and blue), as functions of time. On the right: LOB at each step of the convergence to a new equilibrium.

References

- [1] C. Aliprantis and K. Border. *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media, 2006.
- [2] R. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10:1–18, 2003.
- [3] R. Aumann. Markets with a continuum of traders. *Econometrica*, 32:39–50, 1964.
- [4] M. Avellaneda and S. Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.
- [5] E. Bayraktar and M. Ludkovski. Optimal trade execution in illiquid markets. *Mathematical Finance*, 21(4):681–701, 2011.
- [6] T. R. Bielecki, S. Crépey, M. Jeanblanc, and M. Rutkowski. Arbitrage pricing of defaultable game options with applications to convertible bonds. *Quantitative Finance*, 8(8):795–810, 2008.
- [7] G. Carmona. *Existence and stability of Nash equilibrium*. World Scientific Publishing Co. Pte. Ltd, 2013.
- [8] R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SICON*, 51(4):27052734, 2013.
- [9] R. Carmona and K. Webster. A belief-driven order book model. <http://www.princeton.edu/~rcarmona/download/fe/CW2.pdf>, working paper.
- [10] A. Cartea, R.-F. Donnelly, and S. Jaimungal. Algorithmic trading with model uncertainty. *To appear in SIAM Journal on Financial Mathematics*.
- [11] A. Cartea, R.-F. Donnelly, and S. Jaimungal. Enhancing trading strategies with order book signals. *SSRN*.
- [12] A. Cartea and A. Jaimungal. Optimal execution with limit and market orders. *Quantitative Finance*, 15(8):1279–1291, 2015.

- [13] J. Chassagneux, R. Elie, and I. Kharroubi. A note on existence and uniqueness for solutions of multidimensional reflected BSDEs. *Electronic Communications in Probability*, 16:120–128, 2011.
- [14] R. Cont, A. Kukanov, and S. Stoikov. The price impact of Order Book events. *Journal of Financial Econometrics*, 12(1):47–88, 2014.
- [15] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563, 2010.
- [16] R. Cont and P. Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, 2004.
- [17] B. J. Cvitanic and I. Karatzas. Backward stochastic differential equations with reflection and dynkin games. *Annals of Probability*, 24(4):2024–2056, 1996.
- [18] S. Du and H. Zhu. Welfare and optimal trading frequency in dynamic double auctions. *NBER Working Paper No. 20588*, 2014.
- [19] E. B. Dynkin. Game variant of a problem on optimal stopping. *Soviet Math. Dokl.*, 10:270–274, 1969.
- [20] R. Elie and I. Kharroubi. BSDE representations for optimal switching problems with controlled volatility. *Stochastics and Dynamics*, 14(3), 2014.
- [21] T. Foucault. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets*, 2:99–134, 1999.
- [22] R. Gayduk and S. Nadtochiy. Control-stopping games for market microstructure and beyond. *Working paper*, 2017.
- [23] R. Gayduk and S. Nadtochiy. Liquidity effects of trading frequency. *Accepted in Mathematical Finance*, 2017.
- [24] A. Gegout-Petit and E. Pardoux. Equations différentielles stochastiques rétrogrades réfléchies dans un convexe. *Stochastics and Stochastics Reports*, 57:111–128, 1996.
- [25] R. L. Goettler, C. A. Parlour, and U. Rajan. Equilibrium in a dynamic limit order market. *The Journal of Finance*, 60(5):2149–2192, 2005.
- [26] A. Granas and J. Dugundji. *Fixed point theory*. Springer Science & Business Media, 2013.
- [27] O. Guéant. Optimal market making. *arXiv:1605.01862*.
- [28] O. Guéant and C.-A. Lehalle. General intensity shapes in optimal liquidation. *Mathematical Finance*, 3:457–495, 2015.
- [29] F. Guilbaud and H. Pham. Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1):79–94, 2013.
- [30] J. Jacod and A. Shiryaev. *Limit theorems for stochastic processes*. Springer-Verlag Berlin Heidelberg, 2003.
- [31] S. Jaimungal and M. Nourian. Mean-field game strategies for a major-minor agent optimal execution problem. <http://ssrn.com/abstract=2578733>, working paper.
- [32] O. Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- [33] I. Karatzas, Q. Li, and R. J. Elliott. BSDE approach to non-zero-sum stochastic differential games of control and stopping. In *Stochastic Processes, Finance and Control: A Festschrift in Honor of Robert J. Elliott*, pages 105–153. World Scientific, 2011.

- [34] I. Karatzas and W. Suderth. Stochastic games of control and stopping for a linear diffusion. In *Random Walk, Sequential Analysis and Related Topics: A Festschrift in Honor of Y. S. Chow*, pages 100–117. World Scientific, 2006.
- [35] N. E. Karoui, C. Kapoudjian, E. Pardoux, S. Peng, and M. C. Quenez. Reflected solutions of backward SDE's, and related obstacle problems for PDE's. *Annals of Probability*, 25(2):702–737, 1997.
- [36] Y. Kifer. Optimal stopping in games with continuous time. *Theory of Probability and Applications*, 16:545–550, 1971.
- [37] A. Lachapelle, J.-M. Lasry, C.-A. Lehalle, and P.-L. Lions. Efficiency of the price formation process in presence of high frequency participants: a mean field game analysis. *arXiv:1305.6323*, preprint, 2013.
- [38] J.-M. Lasry and P.-L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2:229–260, 2007.
- [39] A. Lipton, U. Pesavento, and M. G. Sotiropoulos. Trade arrival dynamics and quote imbalance in a limit order book. *arXiv:1312.0514*.
- [40] C. Lorenz and A. Schied. Drift dependence of optimal trade execution strategies under transient price impact. *Finance and Stochastics*, 17:743–770, 2013.
- [41] E. Ok. *Real analysis with economic applications*. Princeton University Press, 2007.
- [42] C. A. Parlour. Price dynamics in limit order markets. *The Review of Financial Studies*, 11(4):789–816, 1998.
- [43] S. Predoiu, G. Shaikhet, and S. Shreve. Optimal execution of a general one-sided limit-order book. *SIAM Journal on Financial Mathematics*, 2:183–212, 2011.
- [44] I. Rosu. A dynamic model of the limit order book. *Review of Financial Studies*, 22:4601–4641, 2009.
- [45] M. Saglam and S. Stoikov. Option market making under inventory risk. *Review of Derivatives Research*, 12(1):55–79, 2009.
- [46] A. Schied and T. Zhang. A state-constrained differential game arising in optimal portfolio liquidation. *Mathematical Finance*, to appear.
- [47] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.
- [48] S. Stoikov and R. Waeber. Optimal asset liquidation using limit order book information. *doi:10.2139/ssrn.2113827*, working paper, 2012.
- [49] W. Zhen and X. Hua. Multi-dimensional reflected backward stochastic differential equations and the comparison theorem. *Acta Mathematica Scientia*, 30(5):1819–1836, 2010.