



Optimal bubble riding with price-dependent entry: a mean field game of controls with common noise

Ludovic Tangpi²  · Shichun Wang¹

Received: 28 June 2023 / Accepted: 9 January 2024 / Published online: 27 March 2024

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2024

Abstract

In this paper we further extend the optimal bubble riding model proposed in Tangpi and Wang (Optimal bubble riding: a mean field game with varying entry times, 2022) by allowing for price-dependent entry times. Agents are characterized by their individual *entry threshold* that represents their belief in the strength of the bubble. Conversely, the growth dynamics of the bubble is fueled by the influx of players. Price-dependent entry naturally leads to a mean field game of controls with common noise and random entry time, for which we provide an existence result. The equilibrium is obtained by first solving discretized versions of the game in the weak formulation and then examining the measurability property in the limit. In this paper, the common noise comes from two sources: the price of the asset which all agents trade, and also the the exogenous bubble burst time, which we also discretize and incorporate into the model via progressive enlargement of filtration.

Keywords Asset Bubbles · Mean field games with common noise · Control interaction · Random times

1 Introduction

Financial bubbles have become a topic of growing concern in the recent past. The classical view of Beaver [10] that investors only follow “information content” is clearly not applicable to bubbles. Empirical evidence demonstrates the substantial stock price premium following vacuous company announcements of broad intentions to enter the cryptocurrency market [3]

The authors wish to thank Camilo Hernández for fruitful discussions. L.T. is partially supported by the NSF Grant DMS-2005832, the NSF CAREER award DMS-2143861 and the AMS Claytor-Gilmer fellowship. S.W. is partially supported by the NSF grant DMS-2005832.

✉ Ludovic Tangpi
ludovic.tangpi@princeton.edu

Shichun Wang
shichun.wang@princeton.edu

¹ Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA

² Operations Research and Financial Engineering, Princeton University, Bendheim Center for Finance, Princeton, NJ 08544, USA

and investors' overreaction to description of Blockchain activities in firms' 8-K disclosures [22]. Journeying back another thirty years, a similar "gold rush" occurs during the dot-com bubble. Instead of avoiding the overpriced technology segment, sophisticated investors such as hedge funds invest heavily in the bubble while acknowledging an inevitable burst in the future [33, 41]. The herding behavior is more prevalent now due to the increasing democratization of investing, as evidenced by the dramatic surge of retail traders during the "meme-stock" frenzy [24]. Recent advancements in generative artificial intelligence (AI) unleash a frenzy both on Wall Street and among retail investors, pushing up stock prices of big technology companies. As NVIDIA's market capitalization marches across the trillion dollar line, many begin to suspect overvaluation in the semiconductor market. However, despite the unprecedented interest rate hikes and the recent turmoil in the cryptocurrency market, enthusiasm towards a potential AI bubble has not dampened. The intricate interplay between the intention to leverage rapid growth and the apprehension towards a future price adjustment provides the motivation for our equilibrium-based model proposed in [62], which we further investigate in this paper.

Substantial empirical evidence points to the inaccuracy of viewing bubbles merely as "irrational exuberance". A wealth of literature exists on the topic of rationality behind bubbles. The famous "greater fool" model formulated by [5], as well as a more recent adaptation [54], pinpoints the driving factor behind bubble riding as the perception that others will acquire the overpriced asset in the future. Abreu and Brunnermeier [1] explore the idea of "information asymmetry" from another perspective by giving investors different entry times and various priors on the bubble formation time (see also [29, 30] for extensions of this model). Sotes-Paladino and Zapatero [61] use a dynamic trading model to show that sophisticated, risk-averse money managers can invest in overvalued non-benchmark asset due to the presence of convex incentives. The authors in [7, 58] argue that a "chain of middlemen" could also spur the escalation of the asset price.

Despite extensive debates surrounding how a bubble is formed, or even defined, it is commonly agreed that the influx of investors and capital is what sustains the overvaluation. In other words, it is natural to model these events in a large population setting. This is the motivation behind the use of a mean field game (MFG) in our previous paper [62], which should be understood as the infinite population limit of symmetric stochastic differential games [16, 26, 46, 52]. First introduced by Lasry and Lions [49–51] and also by Huang et al. [35, 36], MFGs provide tractable solutions compared to large but finite population games. We refer the readers to the monographs of Carmona and Delarue [17, 18] for the probabilistic approach and to the notes of Cardaliaguet [15] for the partial differential equation (PDE) approach to MFGs.

Our previous study [62] introduced a class of MFGs with varying entry times. Players begin to take advantage of the rising price trajectory at different times during the "awareness window", a period viewed by Abreu and Brunnermeier [1] as a measure of heterogeneity among traders. The inflow of traders in turn fuels the price dynamics, whose drift is a function of the number of players currently in the game. We modeled the burst of the bubble as the minimum of exogenous and endogenous burst times. While an exogenous crash occurs due to events independent of trading, an endogenous crash happens when the average inventory of the players in the game falls below a threshold. We also included price impact as a second source of interactions among the agents through the controls, leading to an *extended* MFG. We proved existence of MFG equilibrium using the method initiated in [19]. Leveraging established methods on filtration enlargement, we were able to decompose the optimal strategy into before-and-after-burst segments, each part being progressively measurable with respect to the original filtration. Numerically, our model discovers that the equilibrium strat-

egy attempts to delay the burst time and therefore sustain the growth if the bubble is large enough.

The aim of this paper is to remove two major limitations of the model in [62]. The first is that the entry times could not depend on the price of the traded asset. They were instead modeled as independent samples from the same pre-determined distribution. However, an intuitive entry criterion for bubble riders is the first time the price crosses a certain threshold, which we use to characterize player influx in the present work. The second improvement is that we allow for an unbounded drift term in the price dynamics. In particular, since players' entry affects the price, the drift term will depend on the running maximum process of the price itself. We also provide a well-posedness result for this path-dependent dynamics.

As natural as these changes may seem, they require the model to incorporate a “common noise” which is famously challenging because the law of the population has to react to the realization of the noise. Just as in the case of a classical MFG, there are two approaches to deal with common noise. The analytic approach either reformulates the problem into a coupled system of stochastic PDEs or a deterministic, but infinite dimensional, PDE called *the master equation* (see Cardaliaguet et al. [16] for a careful presentation). Given a sufficiently smooth solution of the master equation, one can usually obtain strong results on the MFG equilibrium such as uniqueness or even regularity. However, almost all well-posedness results of the master equation require the Lasry-Lions monotonicity condition [51], or the “displacement monotonicity” condition [2, 32, 37] (see also the “anti-monotonicity” condition [56]). On the other hand, the probabilistic approach introduced by Carmona et al. [20] avoids making this assumption by a compactness argument. Since the monotonicity condition is too strong for our model, we take the latter route for constructing MFG equilibrium. A notable drawback to this compactness approach is that the controls might only be measurable with respect to a larger filtration. A well-known *immersion* property is enforced to ensure fairness in observing that additional information. Immersion is a crucial property in the theory of filtration enlargement [44, 65], stochastic control [43, 45], the theory of conditional McKean-Vlasov SDEs [48] and of course MFGs [20]. For an extensive discussion and generalization on both methods of tackling common noise, see [18]. Other recent extensions on related topics include MFGs with finite state space [9], restoring uniqueness of equilibrium [25], incorporating absorption [14], convergence from N -player games [26, 47], and MFGs with interactions through controls [28]. To our knowledge, there are not general existence results on the equilibrium of extended MFGs with common noise, which will be our main contribution in this work (Theorem 2.8). It is worth noting that although we provide a more intuitive model by incorporating common noise compared to our previous version in [62], the result is certainly weaker and less explicit, especially for numerical analysis.

The paper is structured as follows. In Sect. 2, we recall the features of the N -player model for bubble riding from [62] and also introduce the new mechanism for price-dependent entry. Then we formulate the limit mean field game with common noise and state the assumptions as well as the main existence result. The proof is broken down into two parts. In Sect. 3 we first show an intermediate step with a weaker notion of admissibility. Then we strengthen the result in Sect. 4 by reducing the filtration for the controls.

2 Model setup

2.1 The N-player game

2.1.1 Price-dependent entry

Suppose the price trajectory P_t starts at $P_0 > 0$ at time $t = 0$ when the bubble starts. Each player i is characterized by $p^i \stackrel{\text{i.i.d.}}{\sim} \nu_p$ on $[P_0, \infty)$, a price threshold for the player to deem the bubble worth riding. The player enters the game at

$$t^i := \inf\{t \geq 0 : P_t \geq p^i\} \wedge (T + 1).$$

The value $T + 1$ is chosen arbitrarily, but strictly greater than T , in order to bound t^i if the price never reaches the threshold. Note that in contrast to [62], the entry times are still random even if we fix the individual information p^i due to their dependence on the common noise in P . If the price trajectory is càdlàg and jumps are non-positive, then $P_{t^i} = p^i$ on the event that $t^i \leq T$. We assume that there are initial players with thresholds P_0 who are already in the game at $t = 0$. That is, $\nu_p(\{P_0\}) > 0$. Let $\vec{p} = (p^1, \dots, p^N)$ be the vector of thresholds. Then the number of players in the game by time t is

$$N_{in}(t, P; \vec{p}) = \sum_{i=1}^N \mathbb{1}_{\{t^i \leq t\}} = \sum_{i=1}^N \mathbb{1}_{\{\max_{s \leq t} P_s \geq p^i\}} = N F_p^{N, \vec{p}}(\max_{s \leq t} P_s) \quad (1)$$

where $F_p^{N, \vec{p}}$ is the empirical CDF of the thresholds (p^1, \dots, p^N) .

2.1.2 Price dynamics in the bubble phase

Let P^+ denote the price process in the bubble phase before the burst. The growth of the bubble should depend on the inflow of players, $N_{in}(t, P; \vec{p})$, which by (1) is a function of the maximum process of the bubble price itself. This motivates the following price dynamics before burst

$$dP_t^+ = b(t, \max_{s \leq t} P_s^+, P_t^+) dt + \sigma_0 dB_t, \quad P_0^+ = P_0, \quad (2)$$

where b is called the bubble trend function. Because the price grows with entry, b depends on the thresholds \vec{p} and should be monotone increasing in its second argument. We present the generalized form of the two examples for b given in [62] with price-dependent entry.

Example 2.1 (Exponential Bubble) Abreu and Brunnermeier [1] assumed a fully deterministic model with exponential price trajectory. In our setting, this translates to

$$dP_t^+ = \ell_t P_t^+ dt + \sigma_0 dB_t, \quad \ell_t = \ell F_p^{N, \vec{p}}(\max_{s \leq t} P_s^+),$$

where $\ell > 0$ stands for the peak growth rate of the bubble. If we assume that everyone enters at $t = 0$, that is having $p^i = P_0$ for all i , we obtain the model used in [1] with rate ℓ .

Example 2.2 (LPPL Bubble) The Johansen-Ledoit-Sornette (JLS) model proposed by Johansen et al. [39] uses an assumption on the hazard rate h_t of the burst time and arrives at a mean trajectory following the log-periodic power law (LPPL). While we model the burst time

very differently, we can match the shape of the process by taking $dP_t^+ = h_t P_t^+ dt + \sigma_0 dB_t$ with

$$h_t = A(t_c - t)^{\ell_t - 1} + C(t_c - t)^{\ell_t - 1} \cos(\omega \ln(t_c - t) - \phi),$$

with parameters A, C, ω, ϕ and critical time t_c set to T . In particular, the vanilla JLS model uses $\ell_t = \ell \in (0, 1)$ to measure the power law acceleration of prices, which we generalize by incorporating the impact from the players' entry, namely

$$\ell_t = \ell F_p^{N, \vec{p}}(\max_{s \leq t} P_s^+), \quad \ell \in (0, 1).$$

Similar to the previous example, if all players enter at $t = 0$, the model reduces to a standard LPPL. See [40, 59, 60] for more detailed analysis of the JLS model.

It is not obvious that the path-dependent SDE (2) is strongly solvable, since b is not necessarily bounded or Lipschitz, as in the case of empirical CDF. We defer the well-posedness result to the Appendix (Proposition A.1), where monotonicity of the drift is the key condition that ensures existence and uniqueness of a strong solution. The risk of the bubble bursting is not reflected in (2) since this is the dynamics given that the bubble is still alive. We take a constant diffusion as the time horizon is relatively short-term by nature of a bubble.

2.1.3 Inventory and trading rate

Each player i has initial endowment $K_0^i \stackrel{\text{i.i.d.}}{\sim} \nu_K$ on \mathbb{R} and goes “all in” when they enters the bubble ride. For simplicity, assume that there is no transaction cost when joining a bubble ride. That is, each player joins the game with K_0^i/p^i shares of the bubble asset. We also allow for negative values of K_0^i , which represent an initial short (attack) position on the asset. Note that allowing price-dependent entry fixes a shortcoming of the original model proposed in [62] where only the initial inventory is assumed to be i.i.d., which implies that the players who enter later will have more initial wealth due to the higher asset price at entry.

Suppose that there is a common horizon $T > 0$. By choosing her trading rate $\alpha^i = (\alpha_t^i)_{t^i \leq t \leq T}$ after entry, the player can control her inventory trajectory by

$$dX_t^i = \alpha_t^i dt + \sigma dW_t^i, \quad X_t^i = 0 \text{ on } t < t^i, \quad X_{t^i}^i = K_0^i/p^i.$$

where $\sigma > 0$ is fixed and W^i, \dots, W^N are independent 1-dimensional Brownian motions corresponding to random streams of demand [19, 53]. A positive α_t corresponds to buying and a negative α_t corresponds to selling. We require $\alpha_t^i = 0$ on $t < t^i$ before entry for each $i \in \{1, \dots, N\}$, and α_t^i takes values in a compact interval $A \subset \mathbb{R}$.

2.1.4 Burst time and post-burst price dynamics

Following [1] and [62], we allow the bubble to burst for both *exogenous* and *endogenous* reasons. An exogenous burst will be modeled as a non-anticipative random time, more specifically a totally inaccessible stopping time τ that is independent from the market information $(B, \{W^i\}_{i=1, \dots, N})$. On the other hand, an endogenous burst occurs when the inventory of the active players (i.e. those who have entered) becomes too low to sustain the frenzy of the bubble. Define the empirical measure of the inventory μ_t^N and the average inventory $\bar{\mu}_t$ as

$$\mu_t^N := \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}, \quad \bar{\mu}_t^N := \frac{\mathbb{1}_{\{N_{in}(t, P; \vec{p}) \neq 0\}}}{N_{in}(t, P; \vec{p})} \sum_{i=1}^N X_t^i = \frac{\mathbb{1}_{\{N_{in}(t, P; \vec{p}) \neq 0\}}}{F_p^{N, \vec{p}}(\max_{s \leq t} P_s)} \int_{\mathbb{R}} x \mu_t^N(dx). \quad (3)$$

Note that our assumption on the existence of initial players allows us to drop the indicator for all $t \in [0, T]$ when $N \rightarrow \infty$. For a given inventory threshold function $\zeta : [0, T] \rightarrow \mathbb{R}_+$, define the endogenous burst as

$$\bar{\tau}^N(\mu^N) := \inf \left\{ t > \min_{i \in \{1, 2, \dots, N\}} t^i : \inf_{s \in [0, t]} \bar{\mu}_s^N \leq \zeta_t \right\} \wedge T.$$

The true burst time is defined as the first of the two events:

$$\tau^*(\mu^N) := \tau \wedge \bar{\tau}^N(\mu^N).$$

At burst time, the price drops by a fraction β_t of the bubble component γ_t , defined by

$$\gamma_t := \int_0^t b(s, \max_{u \leq s} P_u^+, P_s^+) ds. \quad (4)$$

The function $\beta : [0, T] \rightarrow \mathbb{R}_+$ is referred to as the “size” of the bubble [1], or the “loss amplitude” in the JLS model [60].

The setup becomes an optimal execution problem after the crash. Trades convey information which has a long-term impact to the price dynamics. When the bubble is present, the frenzy of the bubble growth dominates the impact from selling. However, after the crash, the asset price is governed by price impact within the short horizon. We use the game-theoretic extension of the model by Almgren and Chriss [6], where the aggregate trading rate determines the instantaneous price impact.

Let $\rho : A \rightarrow \mathbb{R}$ be a concave (hence also continuous) function for the instantaneous impact. See [13, 55, 57] for reasons of the concavity of price impact. Define the empirical measures of controls θ^N and the permanent price impact term $\langle \rho, \theta_t^N \rangle_{F_p^{N, \bar{p}}}$ as

$$\theta_t^N := \frac{1}{N} \sum_{i=1}^N \delta_{\alpha_t^i}, \quad \langle \rho, \theta_t^N \rangle_{F_p^{N, \bar{p}}} := \frac{\mathbb{1}_{\{N_{in}(t, P; \bar{p}) \neq 0\}}}{N_{in}(t, P; \bar{p})} \sum_{i=1}^N \rho(\alpha_t^i) = \frac{\mathbb{1}_{\{N_{in}(t, P; \bar{p}) \neq 0\}}}{F_p^{N, \bar{p}}(\max_{s \leq t} P_s)} \int_A \rho(a) \theta_t^N(da). \quad (5)$$

This is the second source of interaction among players currently in the game. Since by definition α^i and X^i are both 0 before entry, there is a factor $1/F_p^N$ in both (3) and (5) before the integral. After burst, the bubble trend is no longer present, so the price P^- after burst follows on $[\tau^*, T]$

$$dP_t^- = \langle \rho, \theta_t^N \rangle_{F_p^{N, \bar{p}}} dt + \sigma_0 dB_t, \quad P_{\tau^*}^- = P_{\tau^*}^+ - \beta_{\tau^*} \gamma_{\tau^*}. \quad (6)$$

Define $D_t^* = \mathbb{1}_{\{t \geq \tau^*\}}$. Using chain rule on $P_t = P_t^+(1 - D_t^*) + P_t^- D_t^*$ gives us the price dynamics

$$\begin{aligned} dP_t &= (1 - D_t^*) dP_t^+ - P_t^+ dD_t^* + D_t^* dP_t^- + P_t^- dD_t^* \\ &= \mathbb{1}_{\{t < \tau^*\}} dP_t^+ + \mathbb{1}_{\{t \geq \tau^*\}} dP_t^- - \gamma_{\tau^*} \beta_{\tau^*} dD_t^*. \end{aligned} \quad (7)$$

2.1.5 Objective and equilibrium

Player i ’s cash process is modeled by

$$dK_t^i = -(\alpha_t^i P_t + \kappa(\alpha_t^i)) dt, \quad K_0^i \sim \nu_k.$$

where $\kappa(\cdot)$ is a continuous, strictly convex function satisfying $\kappa(0) = 0$ that measures the temporary price impact that affects only the individual trader and not the price itself. The

well-known linear temporary impact [6] corresponds to κ being quadratic. See also [19, Section 2.1] for choosing κ as the antiderivative of ρ . Note that the cash process remains at the initial endowment until the player enters the game, since α^i is kept at 0. Under the usual self-financing condition, the pre-burst wealth V^i of this player follows

$$\begin{aligned} dV_t^i &= dK_t^i + X_t^i dP_t + P_t dX_t^i \\ &= \left(-\kappa(\alpha_t^i) + X_t^i b(t, \max_{s \leq t} P_s^+, P_t^+) \mathbb{1}_{\{t < \tau^*\}} + X_t^i \langle \rho, \theta_t^N \rangle_{F_p^{N, \tilde{p}}} \mathbb{1}_{\{t \geq \tau^*\}} \right) dt \\ &\quad - X_t^i \beta_t \gamma_t dD_t^* + \sigma_0 X_t^i dB_t + \sigma P_t dW_t^i. \end{aligned}$$

The players are allowed to continue trading until T , even if the burst has already happened. But by definition of riding a bubble, the players do not believe in the fundamental value of the asset. Therefore, we impose a quadratic terminal inventory penalty $c(X_T^i)^2$ with $c > 0$ to encourage selling. For a fixed $\phi > 0$, we also impose a quadratic running inventory cost $\phi(X_t^i)^2$ which Cartea et al. [21] refer to as *ambiguity aversion*. Adding these costs to the negative of increase in wealth, we have the total cost of player i that she wants to minimize:

$$\begin{aligned} J^{N,i}(\alpha, \tilde{p}) &:= \mathbb{E} \left[-(V_T^i - V_{t^i}^i) + \int_{t^i}^T \phi(X_t^i)^2 dt + c(X_T^i)^2 \right] = \mathbb{E} \left[c(X_T^i)^2 + X_{\tau^*}^i \beta_{\tau^*} \gamma_{\tau^*} \right] \\ &\quad + \mathbb{E} \left[\int_{t^i}^T \left(\kappa(\alpha_t^i) + \phi(X_t^i)^2 - X_t^i b(t, \max_{s \leq t} P_s^+, P_t^+) \mathbb{1}_{\{t < \tau^*\}} \right) \right. \\ &\quad \left. + X_t^i \langle \rho, \theta_t^N \rangle_{F_p^{N, \tilde{p}}} \mathbb{1}_{\{t \geq \tau^*\}} dt \right] \end{aligned} \quad (8)$$

for given vectors of strategies $\alpha = (\alpha^1, \dots, \alpha^N)$ and entry thresholds $\tilde{p} = (p^1, \dots, p^N)$. The interaction among the players appears both in the price impact term through the average trading speed and also the burst time through the average inventory. We refer the readers to [62] for more details on the model.

It is well-known that finite-player games of this type quickly becomes intractable as N increases. Since the phenomenon of bubble riding fits naturally in the large-population setting, we shift our focus directly to the mean field limit of the game described above.

2.2 Mean field game setup

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space that supports independent (K_0, W, B, τ) , whose law under \mathbb{P} is $\nu_K \times \mathcal{W} \times \mathcal{W} \times \nu_\tau$ where \mathcal{W} is the (one dimensional) Wiener measure. Let $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ be a \mathbb{P} -completed filtration defined on this probability space such that W, B are (\mathbb{F}, \mathbb{P}) Brownian motions, the initial wealth K_0 is \mathcal{F}_0 -measurable, and exogenous burst time τ is not an \mathbb{F} -stopping time. Let \mathbb{G} be the smallest filtration, on which τ is a stopping time, that contains \mathbb{F} . We will see in later sections that by construction, τ will in fact be a \mathbb{G} -totally inaccessible stopping time under mild assumptions. K_0 and \mathcal{P} are the initial endowment and entry threshold, respectively. The initial inventory at entry is K_0/\mathcal{P} .

Let $\mathcal{B}(E)$ denote the Borel subsets of a Polish space (E, d) , and let $\mathcal{P}(E)$ denote the set of all probability measures on $\mathcal{B}(E)$. Unless specified otherwise, $\mathcal{P}(E)$ is equipped with the topology of weak convergence of measures, and $\mathcal{P}(E)$ is also a Polish space. Denote the Wasserstein space (or order 1) by $\mathcal{P}_1(E)$, that is

$$\mathcal{P}_1(E) := \left\{ \mu \in \mathcal{P}(E) : \int_E d(x_0, x) \mu(dx) < \infty \right\}.$$

where the choice of $x_0 \in E$ is arbitrary. Equip $\mathcal{P}_1(E)$ with the 1-Wasserstein distance

$$\mathcal{W}_1(\mu, \mu') := \inf_{\pi \in \Pi(\mu \times \mu')} \int_{E \times E} d(x, y) d\pi(x, y).$$

Let $\mathcal{X} := C([0, T], \mathbb{R})$ denote the space of all continuous functions from $[0, T]$ to \mathbb{R} equipped with the sup–norm. Let $D([0, T], \mathbb{R})$ denote the space of all càdlàg functions from $[0, T]$ to \mathbb{R} . For a fixed $t^* \leq T$, define

$$\mathcal{X}^{t^*} := \left\{ \mathbf{x}^{t^*} \in D([0, T], \mathbb{R}) : \mathbf{x}_t^{t^*} = 0 \text{ on } [0, t^*), \text{ continuous on } [t^*, T] \right\} \text{ and } \mathcal{X}^* := \bigcup_{t^* \in [0, T]} \mathcal{X}^{t^*}.$$

For each $\mathbf{x}^{t^*} \in \mathcal{X}^*$, we require t^* to be the largest value such that $\mathbf{x}^{t^*} \in \mathcal{X}^{t^*}$ to avoid redundancies. Suppose $t_1, t_2 \in [0, T]$. Let $\mathbf{x}^{t_1}, \mathbf{y}^{t_2} \in \mathcal{X}^*$. Notice that the standard sup distance $d(\mathbf{x}^{t_1}, \mathbf{y}^{t_2}) = \|\mathbf{x}^{t_1} - \mathbf{y}^{t_2}\|_\infty$ is no longer suitable on \mathcal{X}^* because it does not allow two processes to be close unless $t_1 = t_2$. Therefore, for each $\mathbf{x}^{t^*} \in \mathcal{X}^*$, we can define its continuous counterpart $\bar{\mathbf{x}}^{t^*} \in C([0, T], \mathbb{R})$ as

$$\bar{\mathbf{x}}_t^{t^*} := \begin{cases} \mathbf{x}_{t^*}^{t^*} & t \in [0, t^*) \\ \mathbf{x}_t^{t^*} & t \geq t^*. \end{cases}$$

We then define a metric on $\mathcal{X}^* \subset D([0, T], \mathbb{R})$ to be

$$d_{\mathcal{X}^*}(\mathbf{x}^{t_1}, \mathbf{y}^{t_2}) := \|\bar{\mathbf{x}}^{t_1} - \bar{\mathbf{y}}^{t_2}\|_\infty + |t_1 - t_2|, \quad \text{and} \quad \|\mathbf{x}^{t^*}\|_{\mathcal{X}^*} := \|\bar{\mathbf{x}}^{t^*}\|_\infty + t^*. \quad (9)$$

Define $\Theta^A := \mathcal{P}([0, T] \times A)$ with time marginal being Lebesgue, and define

$$\mathcal{M} := \left\{ \mu \in \mathcal{P}(\mathcal{X}^*) : \int_{\mathcal{X}^*} \|\mathbf{x}\|_{\mathcal{X}^*} \mu(d\mathbf{x}) < \infty \right\} = \mathcal{P}_1(\mathcal{X}^*). \quad (10)$$

Equip \mathcal{M} and Θ^A with the topology of 1-Wasserstein convergence and weak convergence, respectively. Each $\theta \in \Theta^A$ can uniquely disintegrate into $\theta(dt, da) = \theta_t(da)dt$ with some measurable map $t \mapsto \theta_t \in \mathcal{P}(A)$. Let F_p be the CDF of the price threshold distribution v_p . Each $\mu \in \mathcal{M}$ can also be viewed as a $\mathcal{P}(\mathbb{R})$ -valued process $t \mapsto \mu_t$, where $\mu_t = \mu \circ \mathcal{E}_t^{-1}$ with $\mathcal{E}_t : \mathcal{X}^* \rightarrow \mathbb{R}$ being the time coordinate mapping. Let μ and θ be the law of X and α . Then the average inventory among the players in the game and their price impact are the natural limit of (3) and (5), namely:

$$\bar{\mu}_t := \frac{1}{F_p(\max_{s \leq t} P_s)} \langle x, \mu_t \rangle, \quad \langle \rho, \theta_t \rangle_{F_p} := \frac{1}{F_p(\max_{s \leq t} P_s)} \langle \rho, \theta_t \rangle.$$

2.2.1 Burst time

While the exogenous burst time is the same as the N -player game, the endogenous burst now depends on $\bar{\mu}$. To ensure that there are already players in the game at time $t = 0$, we assume that $v_p(\{P_0\}) > 0$ and define

$$\bar{\tau}(\mu) := \inf \left\{ t \in [0, T] : \inf_{s \in [0, t]} \bar{\mu}_s \leq \zeta_t \right\} \wedge T,$$

where $\zeta : [0, T] \rightarrow \mathbb{R}^+$ is deterministic, continuous and strictly increasing with $\zeta_0 \in (0, \mathbb{E}[K_0]/P_0)$. The upper bound is set so that the bubble at least survives the initial players. These conditions guarantee enough regularity of $\bar{\tau}$ for the equilibrium to exist (see [62],

Sections 2.1.5 and 6.1]). The actual burst time is again $\tau^*(\mu) = \bar{\tau}(\mu) \wedge \tau$. Throughout the paper, we work under the following assumptions.

Assumption 1 (E)

- (E1) A is a compact interval that includes 0.
- (E2) v_τ is absolutely continuous with respect to Lebesgue measure on \mathbb{R}^+ satisfying $v_\tau(\tau > T) > 0$. Moreover, its deterministic, non-negative intensity process $k : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is bounded by some $C_k > 0$ on $[0, T]$.
- (E3) K_0 has all moments. $\rho : A \rightarrow \mathbb{R}$ is locally bounded.

For a càdlàg process Y , define $M_t^Y = \sup_{0 \leq s \leq t} Y_s$. Observe from Examples 2.1 and 2.2 that the bubble trend function b depends on the running maximum of the price P naturally through the CDF of entry thresholds. That is, the dynamics of P depends on $F_p(M_t^P)$ where F_p is not necessarily Lipschitz continuous. Therefore, the SDE defining the price process may not be well-posed. We will show in the appendix that appropriate growth and monotonicity conditions on b below, which we make as standing assumptions on the bubble, allow to obtain well-posedness.

Assumption 2 (B)

- (B1) The bubble function $b : [0, T] \times [0, 1] \times \mathbb{R} \rightarrow \mathbb{R}$ is non-negative and satisfies the assumptions in Proposition A.1.
- (B2) The bubble size $\beta : [0, T] \times \Omega \rightarrow \mathbb{R}$ is a positive, continuous, bounded \mathbb{F} -progressively measurable process.

Using Proposition A.1, the price dynamics (7) is well defined for a fixed $(\mu, \theta) \in \mathcal{M} \times \Theta^A$, namely

$$P_t = \mathbb{1}_{\{t < \tau^*\}} P_t^+ + \mathbb{1}_{\{t \geq \tau^*\}} P_t^-,$$

where the pre-burst price P^+ follows (2) and the post-burst price P^- follows

$$\begin{aligned} P_t^- &= P_{\tau^*}^+ - \beta_{\tau^*} \gamma_{\tau^*} + \int_{\tau^*}^t \langle \rho, \theta_s \rangle_{F_p} ds + \sigma_0 (B_t - B_{\tau^*}) \\ &= P_0 + \gamma_{\tau^*} + \sigma_0 B_{\tau^*} - \beta_{\tau^*} \gamma_{\tau^*} + \int_{\tau^*}^t \langle \rho, \theta_s \rangle_{F_p} ds + \sigma_0 (B_t - B_{\tau^*}) \\ &= P_0 + \int_{\tau^*}^t \langle \rho, \theta_s \rangle_{F_p} ds + \sigma_0 B_t + (1 - \beta_{\tau^*}) \gamma_{\tau^*}. \end{aligned}$$

The bubble component γ is defined by (4). Before crash, the bubble component dominates the drift of the price dynamics, whereas the impact term takes over after the the crash. Note that the price has exactly one jump at τ^* , and the jump size $-\beta_{\tau^*} \gamma_{\tau^*}$ is always negative.

2.2.2 Entry time

Since the underlying asset starts at a known value P_0 , the entry threshold should be at least this amount, that is, v_p is a distribution on $[P_0, \infty)$. Consider the product probability space $(\Omega \times [P_0, \infty), \mathcal{F} \otimes \mathcal{B}([P_0, \infty)), \mathbb{P} \otimes v_p)$. We naturally extend (K_0, W, B, τ) on the product space. The representative player enters the game at a random \mathbb{G} -stopping time $\mathcal{T}(p^*)$, $p^* \in [P_0, \infty)$, where $\mathcal{T} : \Omega \times [P_0, \infty) \rightarrow [0, T] \cup \{T + 1\}$ is given by

$$\mathcal{T}(p^*) := \inf\{t \in [0, T] : P_t \geq p^*\} \wedge (T + 1), \quad \mathbb{P} - a.s. \quad (11)$$

The value $T + 1$ is again arbitrarily chosen. Then for every $p^* \in [P_0, \infty)$, $\mathcal{T}(p^*)$ is a bounded stopping time with respect to the filtration generated by B .

Lemma 2.3 *Consider $\mathcal{T}: \Omega \times [P_0, \infty) \rightarrow [0, T] \cup \{T + 1\}$ in equation (11). For \mathbb{P} -almost every ω , the function $\mathcal{T}(\omega, \cdot)$ is strictly increasing until and if it reaches $T + 1$. Moreover, it is left continuous with right limit on $[P_0, \infty)$, and \mathcal{T} is jointly measurable.*

Proof Monotonicity is obvious. To see that it is strict, we can first write $\mathcal{T}(p^*) = \inf\{t \geq 0 : M_t^P \geq p^*\} \wedge (T + 1)$. Note that the price process P is \mathbb{P} -almost surely continuous except at τ^* where there is a non-positive jump. Therefore, M^P is a monotone increasing, \mathbb{P} -almost surely continuous process, which implies strict monotonicity of \mathcal{T} .

For each ω such that M^P is continuous, $\mathcal{T}(p^*) = T + 1$ for $p^* > M_T^P$. For $p^* \leq M_T^P$, we have $\mathcal{T}(p^*) \in [0, T]$ and $M_{\mathcal{T}(p^*)} = p^*$. Take an increasing sequence $p_n \uparrow p^*$. Then $\mathcal{T}(p_n)$ is also an increasing sequence which converges to some $t \leq \mathcal{T}(p^*)$ as $n \rightarrow \infty$. If $t < \mathcal{T}(p^*)$, then we can find $t' \in (t, \mathcal{T}(p^*))$ such that $p_n \leq M_{t'}^P \leq p^*$ for all n by monotonicity of M^P . p_n converging to p^* implies $M_{t'}^P = p^*$ which contradicts the definition of $\mathcal{T}(p^*)$. The existence of right limit follows a similar argument. Joint measurability follows from Lemma A.2. \square

2.2.3 Admissibility of controls

By continuity of P , given $\mathcal{T} = t^*$ we can also recover the price threshold by $p^* = P_{t^*}$. However, it is still useful to define admissibility of controls in two separate ways: one in terms of entry times, and the other in terms of entry thresholds. To simplify notation, we denote by $\mathcal{PM}(\mathbb{G})$ (resp. \mathbb{F}) the σ -algebra generated by the \mathbb{G} (resp. \mathbb{F})-progressively measurable subsets of $\Omega \times [0, T]$.

Definition 2.4 Define the following sets for admissible controls:

- For $t^* \in [0, T]$, let $\mathbb{A}(t^*)$ denote the set of square integrable, $\mathcal{PM}(\mathbb{G})$ -measurable processes $\alpha: \Omega \times [0, T] \rightarrow A$ such that $\alpha_t = 0$ for $t \in [0, t^*)$. We also set $\mathbb{A}(T + 1)$ to be the singleton of the constant $\mathbf{0}$ process.
- A *time-admissible control* α is a process $\Omega \times [0, T] \times [0, T] \ni (\omega, t, t^*) \mapsto \alpha_t^{t^*}(\omega) \in A$ that is $\mathcal{PM}(\mathbb{G}) \otimes \mathcal{B}([0, T])$ -measurable such that for almost all t^* , $\alpha^{t^*} \in \mathbb{A}(t^*)$. Let \mathbb{A} denote all such strategies.
- A *price-admissible control* α is a process $\Omega \times [0, T] \times [P_0, \infty) \ni (\omega, t, p^*) \mapsto \alpha_t(\omega, p^*) \in A$ that is $\mathcal{PM}(\mathbb{G}) \otimes \mathcal{B}([P_0, \infty))$ -measurable such that for v_p -almost all p^* , $\alpha_t(\cdot, p^*) = 0$ a.s. on the random interval $[0, \mathcal{T}(p^*) \wedge T]$. Let \mathbb{A}^* denote all such strategies.

By Lemma 2.3, any time-admissible control $\alpha \in \mathbb{A}$ induces a price-admissible control by $\alpha_t(\cdot, p^*) = \alpha_t^{\mathcal{T}(\cdot, p^*)}(\cdot)$. Let \mathcal{P} denote the identity map for the last coordinate on $[P_0, \infty)$, which represents the entry threshold with law v_p . Then for each $\alpha \in \mathbb{A}^*$, the corresponding state process satisfies

$$X_t^{\mathcal{T}, \alpha} = \mathbb{1}_{\{t \geq \mathcal{T}\}} K_0 / \mathcal{P} + \int_0^t \alpha_s ds + \sigma (W_{t \vee \mathcal{T}} - W_T), \quad t \in [0, T]. \quad (12)$$

2.2.4 Objective and equilibrium

Following the same derivation from the N-player game, using (8) we can define the running cost function $f : [0, T] \times \mathbb{R} \times \mathbb{R} \times [0, T] \times \mathbb{R} \times A \rightarrow \mathbb{R}$:

$$f(t, x, b, \eta, \varrho, a) = \kappa(a) + \phi x^2 - x(\mathbf{b} \mathbf{1}_{\{t < \eta\}} + \varrho \mathbf{1}_{\{t \geq \eta\}}) \quad (13)$$

and the terminal cost function $g : \Omega \times \mathbb{R} \times \mathbb{R} \times [0, T] \rightarrow \mathbb{R}$:

$$g(x_1, x_2, \eta) = cx_1^2 + \beta_\eta \gamma_\eta x_2. \quad (14)$$

Allowing $C > 0$ to vary in each step, by Assumptions (2), (1) and Proposition A.1 we have

$$\begin{aligned} \mathbb{E} \left[\sup_{\eta \in [0, T], \alpha \in \mathbb{A}^*} |g(X_T^{T, \alpha}, X_\eta^{T, \alpha}, \eta)|^2 \right] &\leq C \mathbb{E} \left[1 + \left(\frac{K_0}{p^*} \right)^4 + \sigma^4 W_T^4 + \sup_{\eta \in [0, T]} \gamma_\eta^2 + \sup_{\eta \in [0, T]} X_\eta^2 \right] \\ &\leq C \mathbb{E} \left[1 + \left(\int_0^T b(t, M_t^P, P_t) dt \right)^2 \right] \\ &\leq C \mathbb{E} \left[1 + \int_0^T |M_t^P|^2 dt \right] \\ &\leq C \mathbb{E} \left[1 + T \sup_{t \in [0, T]} P_t^2 \right] < \infty. \end{aligned}$$

Given a price process P , define the process $b_t^P := b(t, M_t^P, P_t)$. For a fixed $\theta = (\theta_t)_{t \in [0, T]}$ and $\mu = (\mu_t)_{t \in [0, T]}$, the objective which the representative player minimizes over \mathbb{A}^* is:

$$J^{\mu, \theta}(\alpha) = \mathbb{E} \left[g(X_T^{T, \alpha}, X_{\tau^*(\mu)}^{T, \alpha}, \tau^*(\mu)) + \int_{T \wedge T}^T f(s, X_s^{T, \alpha}, b_s^P, \tau^*(\mu), \langle \rho, \theta_s \rangle, \alpha_s) ds \right].$$

If the player does not enter by time T , her total cost is 0. This is also true by construction, see Remark 2.5 (C4) below.

Remark 2.5 We make a note that the following set of properties of the cost structure will be utilized in the proof.

(C1) The running cost function $f : [0, T] \times \mathbb{R} \times \mathbb{R} \times [0, T] \times \mathbb{R} \times A \rightarrow \mathbb{R}$ is (jointly) Borel measurable and can be decomposed as

$$f(t, x, b, \eta, \varrho, a) = f_a(t, x, a) + f_b(t, x, b) \mathbf{1}_{\{0 \leq t < \eta\}} + f_c(t, x, \varrho) \mathbf{1}_{\{t \geq \eta\}}.$$

For each t , $f_a(t, \cdot, \cdot)$, $f_b(t, \cdot, \cdot)$ and $f_c(t, \cdot, \cdot)$ are continuous. In addition, there exists $\ell_f > 0$ such that for all $(t, b, \varrho, x) \in [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}$:

$$|f_a(t, x, a)| + |f_b(t, x, b)| + |f_c(t, x, \varrho)| \leq \ell_f (1 + |x|^2 + |b|^2).$$

(C2) The terminal cost function $g : \Omega \times \mathbb{R} \times \mathbb{R} \times [0, T] \rightarrow \mathbb{R}$ is almost surely continuous in (x_1, x_2, η) . In addition, there exists $C > 0$ such that $\mathbb{E} \left[\sup_{\eta \in [0, T], \alpha \in \mathbb{A}^*} |g(X_T^{T, \alpha}, X_\eta^{T, \alpha}, \eta)|^2 \leq C \right]$.

(C3) f is strictly convex in (a, x) ; g is convex in x_1 and x_2 .

(C4) $f(t, 0, b, \eta, \varrho, 0) = 0$ for any $(t, b, \eta, \varrho) \in [0, T] \times \mathbb{R} \times [0, T] \times \mathbb{R}$. $g(0, 0, \eta) = 0$ for all $\eta \in [0, T]$.

Although we will focus on the specific case of the model with cost functions (13) and (14), most of our results remain true for arbitrary costs satisfying (C1)–(C4).

2.2.5 Identical threshold case

A special case is where everyone has the same threshold $p^* = P_0$ and thus enters all at the beginning. Then the bubble function b does not depend on M^p . Suppose further that b also does not depend on P . Then the game reduces to a fixed entry time case in [62, Proposition A.10].

2.2.6 Common noise and admissible setup

Unlike idiosyncratic noise, the presence of common noise does not vanish even when the number of players goes to infinity. As a consequence, we need to consider “random versions” of (μ, θ) , which we denote as (μ, ϑ) , to represent the *conditional* probability measures given the common noise. Specifically, the probability setup should also support random variable $\mathfrak{P} := (\mu, \vartheta) : \Omega \rightarrow \mathcal{M} \times \Theta^A$. Therefore, for $\alpha \in \mathbb{A}^*$, the objective a representative agent minimizes is

$$J^{\mu, \vartheta}(\alpha) = \mathbb{E} \left[g(X_T^{\mathcal{T}, \alpha}, X_{\tau^*}^{\mathcal{T}, \alpha}, \tau^*(\mu)) + \int_{T \wedge T}^T f(s, X_s^{\mathcal{T}, \alpha}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, \alpha_s) ds \right], \quad (15)$$

where $X^{\mathcal{T}, \alpha}$ follows (12).

In our setup, there are two sources of common noise to the players: a Brownian motion B from the price process P and a jump process $D_t = \mathbb{1}_{\{\tau \leq t\}}$ for the exogenous burst. For any stochastic process Z and random variable ξ , define their natural filtration $\mathbb{F}^{Z, \xi} := (\mathcal{F}_t^{Z, \xi})_{t \in [0, T]}$ where $\mathcal{F}_t^{Z, \xi}$ is the \mathbb{P} -completion of $\sigma((\xi, Z_s)_{s \in [0, t]})$. Intuitively, (μ, ϑ) are conditional laws given (B, D) , so if we view \mathfrak{P} as a $\mathcal{P}(\mathbb{R}) \times \mathcal{P}(A)$ -valued process, it should be $\mathbb{F}^{B, D}$ -adapted. The natural filtration to work with is the completion of $\mathbb{F}^{K_0, \mathcal{P}, W, B, D}$. An equilibrium of this type is called a *strong* solution, which is known to be very hard to obtain (see e.g. the monograph [18]). Instead, we look for a weak equilibrium in the sense of [20, 45] where we only require (μ, ϑ) to be the conditional law of state and control processes given both the common noise (B, D) and the law process \mathfrak{P} itself.

We collect all the components from this section in the next definition in a more general setting where we do not assume that the underlying probability space has a product structure.

Definition 2.6 An *admissible probability setup* is a filtered probability space $(\Omega, \mathcal{F}, \mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}, \mathbb{P})$ satisfying the usual conditions that supports the following mutually independent random elements:

1. A two-dimensional Brownian motion (W, B) .
2. \mathcal{G}_0 -measurable initial data $\mathcal{I} := (K_0, \mathcal{P}) \in \mathbb{R} \times [P_0, \infty)$ with law $\nu_K \otimes \nu_P$.
3. A \mathbb{G} -stopping time τ with law ν_τ , from which we can define the jump process $D_t := \mathbb{1}_{\{\tau \leq t\}}$.

If an admissible probability setup additionally supports $\mathfrak{P} = (\mu, \vartheta)$ taking values in $\mathcal{M} \times \Theta^A$, we can then define $\tau^*(\mu)$, the price process P and random entry time

$$\mathcal{T} := \inf\{t \in [0, T] : P_t - \mathcal{P} \geq 0\} \wedge (T + 1).$$

Observe that \mathcal{T} may not be defined for *every* threshold value $p^* \in [P_0, \infty)$ that \mathcal{P} takes, making this setup slightly weaker. Similarly, we will also weaken the notion of price-admissibility

and let \mathbb{A}^* denote the set of processes $\Omega \times [0, T] \ni (\omega, t) \mapsto \alpha_t(\omega) \in A$ that is \mathbb{G} -progressive measurable such that \mathbb{P} -almost surely, $\alpha_t \mathbb{1}_{\{t \in [0, T]\}} = 0$. In fact, Lemma 2.3 ensures that our p^* -by- p^* construction is also a particular case under this new definition.

It is worth noting that if τ is also independent from \mathfrak{P} , then τ will be an $\mathbb{F}^{\mathcal{I}, B, W, D, \mathfrak{P}}$ -totally inaccessible stopping time (see Remark 3.5). In particular, if \mathbb{G} is just $\mathbb{F}^{\mathcal{I}, B, W, D, \mathfrak{P}}$, this would be a desired feature for the exogenous burst time because the admissible controls can only react to it once τ occurs but cannot anticipate it.

Definition 2.7 A weak MFG equilibrium with strong control is an admissible probability setup $(\Omega, \mathcal{F}, \mathbb{G}, \mathbb{P})$ that supports a \mathcal{G}_T -measurable random variable $\mathfrak{P} = (\hat{\mu}, \hat{\vartheta}) : \Omega \rightarrow \mathcal{M} \times \Theta^A$, paired with optimal control $\hat{\alpha} \in \mathbb{A}^*$ and corresponding state process $X^{\mathcal{I}, \hat{\alpha}}$ satisfying (12) such that

- (1) The filtration $\mathbb{G} = \mathbb{F}^{\mathcal{I}, B, W, D, \mathfrak{P}}$.
- (2) $\hat{\alpha}$ minimizes over \mathbb{A}^* the objective $J^{\hat{\mu}, \hat{\vartheta}}$ defined in (15).
- (3) $\hat{\mu}$ is a version of the conditional law of $X^{\mathcal{I}, \hat{\alpha}}$ given (B, D, \mathfrak{P}) under \mathbb{P} . That is,

$$\hat{\mu}_t(\cdot) = \mathbb{P}\left(X_t^{\mathcal{I}, \hat{\alpha}} \in \cdot \mid \mathcal{F}_t^{B, D, \mathfrak{P}}\right) \text{ for almost all } t \in [0, T].$$

- (4) $\hat{\vartheta}$ is a version of the conditional law of $\hat{\alpha}$ given (B, D, \mathfrak{P}) under \mathbb{P} . That is,

$$\hat{\vartheta}_t(\cdot) = \mathbb{P}\left(\hat{\alpha}_t \in \cdot \mid \mathcal{F}_t^{B, D, \mathfrak{P}}\right) \text{ for almost all } t \in [0, T].$$

Theorem 2.8 Under Assumptions (2) and (1), there exists a weak MFG equilibrium with strong control.

3 Existence of MFG solutions with weak control

3.1 Weak controls

The term “strong control” in Definition 2.7 refers to the fact that $\hat{\alpha}$ is an A valued process that is $\mathbb{F}^{\mathcal{I}, B, W, D, \mathfrak{P}}$ -progressive. We shall prove Theorem 2.8 by following the chain of arguments presented in [20]. Specifically, we use a fixed point and compactness argument by discretizing the common noise (B, D) and then taking weak limit to obtain an equilibrium. To ensure that the limit exists, we first work with *relaxed controls* in a larger filtration.

3.1.1 Relaxed controls

Since the space of uniformly bounded functions is not compact, a standard workaround when analyzing extended MFGs, especially in the presence of common noise, is to consider *relaxed controls*. A relaxed control is a randomized strategy taking values in Γ where

$$\Gamma := \{\gamma \in \mathcal{P}([0, T] \times A) \text{ with time marginal being the Lebesgue measure}\}.$$

Any $\gamma \in \Gamma$ can be characterized, with dt -a.s. uniqueness, by the form $\gamma(dt, da) = \gamma_t(da)dt$ where $t \mapsto \gamma_t \in \mathcal{P}(A)$ is a Borel measurable mapping. Therefore, we can view each $\gamma \in \Gamma$ as a $\mathcal{P}(A)$ -valued process. For a given admissible probability setup, the set of admissible relaxed controls is defined as

$$\bar{\Gamma} := \{\gamma \in \Gamma \text{ that is } \mathbb{G}\text{-progressive such that } \mathbb{P} - \text{almost surely } \gamma_t \mathbb{1}_{\{t \in [0, T]\}} = \delta_0 \ \forall t \in [0, T]\}.$$

A *strict* control refers to the case where γ_t is \mathbb{P} -almost surely a Dirac measure almost everywhere. The state process corresponding to a relaxed control $\gamma \in \Pi^r$ is

$$X_t^{\mathcal{T}, \gamma} = \mathbb{1}_{\{t \geq T\}} K_0 / \mathcal{P} + \int_0^t \int_A a \gamma(ds, da) + \sigma (W_{t \wedge \mathcal{T}} - W_T), \quad t \in [0, T]. \quad (16)$$

Define Θ as the subset of $\mathcal{P}([0, T] \times \mathcal{P}(A))$ whose first projection is Lebesgue measure dt on $[0, T]$. Any $\theta \in \Theta$ can be characterized, with dt a.s. uniqueness, by $\{\theta_t \in \mathcal{P}(\mathcal{P}(A))\}_{t \in [0, T]}$ such that $\theta(dt, dq) = \theta_t(dq)dt$. We naturally extend any bounded measurable function $F : \mathcal{P}(A) \rightarrow \mathbb{R}$ to $\underline{F} : \mathcal{P}(\mathcal{P}(A)) \rightarrow \mathbb{R}$ by

$$\underline{F}(\theta) := \int_{\mathcal{P}(A)} \theta(dq) F(q).$$

In particular, $\underline{F}(\delta_q) = F(q)$ for $q \in \mathcal{P}(A)$. Recall from Remark 2.5 (C1) that we have separability between a and q in the cost f . Therefore, when evaluating f (or rather its extension) on an element of $\mathcal{P}(\mathcal{P}(A))$, we can drop the underline from the notation to avoid further confusion. In particular, for a bounded measurable function $\rho : A \rightarrow \mathbb{R}$, sometimes we slightly abuse the notation by using $\langle \rho, \theta_t \rangle$ to mean $\left\langle \rho, \int_{\mathcal{P}(A)} \theta_t(dq) \right\rangle$ if θ is in Θ instead of Θ^A . Endow Θ with the stable topology, which is the weakest topology making the map $\theta \rightarrow \int \phi d\theta$ continuous, for each bounded measurable function $\phi : [0, T] \times \mathcal{P}(A) \rightarrow \mathbb{R}$ that is continuous in the measure variable for each t . Since A is convex, compact and metrizable, so is Θ . See [38] for details.

The version of objective function (15) for relaxed controls is

$$J^{\mu, \vartheta}(\gamma) = \mathbb{E} \left[g(X_T^{\mathcal{T}, \gamma}, X_{t^*}^{\mathcal{T}, \gamma}, \tau^*(\mu)) + \int_{T \wedge T}^T \int_A f(s, X_s^{\mathcal{T}, \gamma}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, a) \gamma(da, ds) \right]. \quad (17)$$

Notice from (17) and (16) that A -valued controls are naturally embedded in the space of relaxed controls in the form of strict controls.

3.1.2 Immersion property and lifted environment

In this section, we will also weaken the first requirement in Definition 2.7 and work with a filtration \mathbb{G} that is potentially larger than $\mathbb{F}^{\mathcal{T}, B, W, D, \mathcal{P}}$. Allowing more information into the system immediately requires extra care to ensure fairness in observing that additional information. A widely-used procedure is to check that $\mathbb{F}^{\mathcal{T}, B, W, D, \mathcal{P}}$ is *immersed* in \mathbb{G} . This notion of fairness is also called the (H)-hypothesis, natural extension, or compatibility.

Definition 3.1 A filtration \mathbb{H} is said to be *immersed* in another filtration \mathbb{F} defined on the same probability space if $\mathbb{H} \subset \mathbb{F}$ and every square integrable \mathbb{F} -martingale is a square integrable \mathbb{H} -martingale. An \mathbb{F} -adapted càdlàg process $\mu = (\mu_t)_{t \geq 0}$ with values in a Polish space is *compatible* with \mathbb{F} if its natural filtration \mathbb{F}^μ is immersed in \mathbb{F} .

The following proposition is a useful characterization of this property and explains how compatibility weakens the adaptedness to a conditional independence requirement, which is mainly a property of laws. See e.g. [18, Proposition 1.3] for a proof.

Proposition 3.2 *On probability space $(\Omega, \mathcal{F}, \mathbb{P})$, consider two filtrations $\mathbb{H} = (\mathcal{H}_t)_{t \in [0, T]} \subset \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$. The following statements are equivalent.*

1. \mathcal{H} is immersed in \mathbb{F} .
2. \mathcal{H}_T is conditionally independent from \mathcal{F}_t given \mathcal{H}_t for every $t \in [0, T]$.
3. For any $\zeta \in \mathbb{L}^1(\mathcal{F}_t)$, $\mathbb{E}[\zeta | \mathcal{H}_T] = \mathbb{E}[\zeta | \mathcal{H}_t]$.

Specifically, the third statement in Proposition 3.2 allows us to eventually recover a strong control in Sect. 4 from a larger filtration. To ensure that we carry enough information in the smaller filtration for the immersion property to eventually hold, we will work with a *lifted environment* [18]. That is, instead of $\mathfrak{P} = (\mu, \vartheta)$, we require the admissible probability setup to support a random *joint* probability measure \mathfrak{M} that represents the conditional law of $(X^T, W, \gamma, \mathcal{I})$ given common noise. Let \mathfrak{M}^x and \mathfrak{M}^y denote its first and third marginals, which serves the same purpose of (μ, ϑ) in the objective (17). Now we have all the ingredients to define a solution with weak control.

Definition 3.3 A weak MFG equilibrium with weak control is an admissible probability setup $(\Omega, \mathcal{F}, \mathbb{G}, \mathbb{P})$ that supports a \mathcal{G}_T -measurable random variable $\mathfrak{M} : \Omega \rightarrow \mathcal{P}(\mathcal{X}^* \times \mathcal{X} \times \Gamma \times \mathbb{R}^2)$, paired with optimal *relaxed* control $\hat{\gamma} \in \mathbb{I}^\Gamma$ and corresponding state process $X^{T, \hat{\gamma}}$ satisfying (16) such that

1. The filtration $\mathbb{F}^{T, B, W, D, \mathfrak{M}}$ is immersed in \mathbb{G} .
2. $\hat{\gamma}$ minimizes over \mathbb{I}^Γ the relaxed objective $J^{\mathfrak{M}^x, \mathfrak{M}^y}$ defined in (17).
3. \mathfrak{M}^x is a version of the conditional law of $X^{T, \hat{\gamma}}$ given (B, D, \mathfrak{M}) under \mathbb{P} .
4. \mathfrak{M}^y is a version of the conditional law of $\hat{\gamma}$ given (B, D, \mathfrak{M}) under \mathbb{P} .

It is worth noting that both definitions are weak in the probabilistic sense, where the probability space is part of the solution. They are also both weak in the sense of control theory, where the equilibrium strategy is not necessarily measurable with respect to the Brownian motions, but potentially depends on additional randomness.

The usual fixed point argument using compactness no longer applies to these conditional probability measures as their domain becomes too large. To combat the infinite dimensionality issue, Carmona et al. [20] discretizes time and space to reduce the common noise to a finite dimension process and then pass to the limit. We adapt the discretization scheme from [18], also used in [14]. In this section, our goal is to prove the following intermediate result.

Theorem 3.4 Under Assumptions (2) and (1), there exists a weak MFG equilibrium with weak control.

3.2 Weak formulation and enlargement of filtration

Since the probability space is part of the solution, it is convenient to work on the canonical space with the product structure in Sect. 2.2.2. We will also work under the weak formulation as in [19]. Define

$$\Omega_1 := \mathbb{R} \times \mathcal{X}, \quad \Omega_0 = \mathcal{X} \times \mathbb{R}_+, \quad \Omega := \Omega_1 \times \Omega_0, \quad \Omega_c := \Omega \times [P_0, \infty)$$

and let $(K_0, W, B, \tau, \mathcal{P})$ be the corresponding identity maps on $\Omega_c = \mathbb{R} \times \mathcal{X} \times \mathcal{X} \times \mathbb{R}_+ \times [P_0, \infty)$. Let \mathcal{F} be a σ -algebra carrying the above random variables. Define the corresponding probability measures

$$\mathbb{Q}_1 := \nu_K \otimes \mathcal{W}, \quad \mathbb{Q}_0 := \mathcal{W} \otimes \nu_\tau, \quad \mathbb{Q} := \mathbb{Q}_1 \otimes \mathbb{Q}_0, \quad \mathbb{P} := \mathbb{Q} \otimes \nu_p.$$

Define entry time T in a p^* -by- p^* way on Ω_c as (11). Lemma 2.3 ensures that we have an admissible probability setup. Let X^T denote the uncontrolled state variable on the product

space $[0, T] \times \Omega_c$:

$$X_t^T := K_0/\mathcal{P} + \sigma(W_t - W_T) \text{ for } t \geq T \text{ and } X_t^T := 0 \text{ for } t \in [0, T]. \quad (18)$$

Given $\alpha \in \mathbb{A}^*$ define

$$\frac{d\mathbb{P}^\alpha}{d\mathbb{P}} := \mathcal{E}\left(\int_0^\cdot \sigma^{-1}\alpha_s dW_s\right)_T, \quad W_t^\alpha := W_t - \int_0^t \sigma^{-1}\alpha_s ds. \quad (19)$$

By Girsanov's theorem, and square integrability of α , W^α is a Brownian motion under \mathbb{P}^α and X^T follows the state SDE (12) under \mathbb{P}^α . Given (μ, ϑ) , the cost under the weak formulation is

$$J_{\text{weak}}^{\mu, \vartheta}(\alpha) := \mathbb{E}^{\mathbb{P}^\alpha} \left[g(X_T^T, X_{\tau^*(\mu)}^T, \tau^*(\mu)) + \int_{T \wedge T}^T f(s, X_s^T, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, \alpha_s) ds \right] \quad (20)$$

If we fix a price threshold p^* , then $\mathbb{Q}^{\alpha(p^*)}$ is defined on Ω by $\frac{d\mathbb{Q}^{\alpha(p^*)}}{d\mathbb{Q}}$ in a similar way as \mathbb{P}^α .

3.2.1 Progressive enlargement of filtration

We now recall some facts regarding filtration enlargement. Let $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ be a filtration supporting (W, B, \mathcal{I}) and is independent from τ . Let $(\mathcal{D}_t)_{t \in [0, T]}$ denote the natural filtration of the exogenous burst time jump process $(D_t)_{t \in [0, T]}$. Define the progressively enlarged filtration

$$\mathcal{G}_t := \mathcal{F}_t \vee \mathcal{D}_t, \quad \mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}.$$

Note that \mathbb{G} is the smallest filtration which contains \mathbb{F} and such that τ is a \mathbb{G} -stopping time. Since τ is independent from \mathbb{F} , Proposition 3.2 implies that \mathbb{F} is immersed in \mathbb{G} . In particular, W and B remain (\mathbb{G}, \mathbb{P}) -Wiener processes.

For any \mathbb{G} -predictable process h , Assumption (E2) implies there exists a unique \mathbb{F} -predictable process f such that $h_t \mathbb{1}_{\{t \leq \tau\}} = f_t \mathbb{1}_{\{t \leq \tau\}}$ (see [27, Page 186 (a)]). Since D is a \mathbb{G} -submartingale, by Doob–Meyer decomposition we can find a unique, \mathbb{F} -predictable, increasing compensator process K with $K_0 = 0$ and such that:

$$M_t^\tau = D_t - \int_0^t (1 - D_{s-}) dK_s \quad (21)$$

is a (\mathbb{P}, \mathbb{G}) -martingale. Under Assumption (E2), we have $dK_t = k_t dt$.

Remark 3.5 *The random time τ is a \mathbb{G} -inaccessible stopping time if either of the two following conditions is satisfied (see e.g. [8, 12]):*

1. *Every \mathbb{F} -martingale is a.s. continuous;*
2. *τ avoids all \mathbb{F} -stopping times. That is, $\mathbb{P}(\tau = L) = 0$ for any \mathbb{F} -stopping time L .*

For example, if \mathbb{F} is just the \mathbb{P} -completed Brownian filtration with the initial enlargement of K and \mathcal{P} , by martingale representation theorem (1) would be satisfied. Under the non-atomic condition in Assumption (E2), if there is independence between τ and \mathbb{F} , then (2) holds.

3.3 Proof of Theorem 3.4

In this section, we prove the existence of equilibrium with weak controls using backward stochastic differential equations (BSDE).

3.3.1 BSDEs with random entry times

We begin by introducing a few notation of spaces and norms. For a filtration \mathbb{H} and probability measure $\tilde{\mathbb{Q}}$ on Ω , define the following spaces of processes on $[s, t] \subseteq [0, T]$:

- Let $\mathcal{S}_{\mathbb{H}, \tilde{\mathbb{Q}}}^2[s, t]$ denote the space of \mathbb{R} -valued \mathbb{H} -progressively measurable, càdlàg processes Y on $\Omega \times [s, t]$ satisfying

$$\|Y\|_{\mathcal{S}_{\mathbb{H}, \tilde{\mathbb{Q}}}^2[s, t]} := \mathbb{E}^{\tilde{\mathbb{Q}}} \left[\sup_{u \in [s, t]} |Y_u|^2 \right]^{\frac{1}{2}} < \infty.$$

- Let $\mathcal{H}_{\mathbb{H}, \tilde{\mathbb{Q}}}^2[s, t]$ denote \mathbb{R} -valued \mathbb{H} -predictable processes Z on $\Omega \times [s, t]$ satisfying

$$\|Z\|_{\mathcal{H}_{\mathbb{H}, \tilde{\mathbb{Q}}}^2[s, t]} := \mathbb{E}^{\tilde{\mathbb{Q}}} \left[\int_s^t |Z_u|^2 du \right]^{\frac{1}{2}} < \infty.$$

- Let $\mathcal{H}_{\mathbb{H}, \tilde{\mathbb{Q}}, D}^2[s, t]$ denote \mathbb{R} -valued \mathbb{H} -predictable processes U on $\Omega \times [s, t]$ satisfying

$$\|U\|_{\mathcal{H}_{\mathbb{H}, \tilde{\mathbb{Q}}, D}^2[s, t]} := \mathbb{E}^{\tilde{\mathbb{Q}}} \left[\int_s^t |U_u|^2 dD_u \right]^{\frac{1}{2}} < \infty.$$

We drop $\tilde{\mathbb{Q}}$ from notation when \mathbb{Q} is the probability measure. Respectively, for a probability measure $\tilde{\mathbb{Q}}_c$ on Ω_c , define $\mathcal{S}_{\mathbb{H}, c, \tilde{\mathbb{Q}}_c}^2[s, t]$, $\mathcal{H}_{\mathbb{H}, c, \tilde{\mathbb{Q}}_c}^2[s, t]$ and $\mathcal{H}_{\mathbb{H}, c, \tilde{\mathbb{Q}}_c, D}^2[s, t]$ in the same way for processes on $\Omega_c \times [s, t]$. In particular, when $\tilde{\mathbb{Q}}_c = \mathbb{P}$, $\mathcal{H}_{\mathbb{H}, c}^2[s, t]$ denotes \mathbb{R} -valued \mathbb{H} -predictable processes Z on $\Omega_c \times [s, t]$ satisfying

$$\|Z\|_{\mathcal{H}_{\mathbb{H}, c}^2[s, t]} := \mathbb{E}^{\mathbb{P}} \left[\int_s^t |Z_u|^2 du \right]^{\frac{1}{2}} = \mathbb{E}^{\mathbb{Q}} \left[\int_{[P_0, \infty)} \int_s^t |Z_u(p^*)|^2 du \nu_p(dp^*) \right]^{\frac{1}{2}} < \infty.$$

We drop $[s, t]$ from notation when considering the whole interval $[0, T]$.

Since we take the weak formulation to MFGs, we can rewrite the objective function (20) using the solution to a BSDE. Define the Hamiltonian by

$$H(t, x, \mathbf{b}, \boldsymbol{\eta}, \varrho, z, a) = f(t, x, \mathbf{b}, \boldsymbol{\eta}, \varrho, a) + \sigma^{-1} a z.$$

By Remark 2.5 and Assumption (E1), for each $(t, x, m, p, \boldsymbol{\eta}, \varrho, z)$, there exists a unique element in A that minimizes $H(t, x, m, \boldsymbol{\eta}, \varrho, z, \cdot)$. For our model, the minimizer is a function of z only, which we denote as $\hat{a}(z)$. Let h denote the minimized Hamiltonian, that is

$$\begin{aligned} h(t, x, \mathbf{b}, \boldsymbol{\eta}, \varrho, z) := H(t, x, \mathbf{b}, \boldsymbol{\eta}, \varrho, z, \hat{a}(z)) &= \kappa(\hat{a}(z)) + \phi x^2 - x(\mathbf{b} \mathbb{1}_{\{t < \eta\}} + \varrho \mathbb{1}_{\{t \geq \eta\}}) \\ &+ \sigma^{-1} \hat{a}(z) z. \end{aligned} \tag{22}$$

Remark 3.6 We point out some properties of \hat{a} and h that will be utilized later.

- (S1) For a general f and g satisfying the properties in Remark 2.5, \hat{a} is a jointly measurable function of (t, x, z) and continuous in z by Berge's maximum theorem. In our case, the unique minimizer $\hat{a}(\cdot)$ only depends on z .
- (S2) The minimized Hamiltonian h is Lipschitz in z , and it is jointly continuous in (x, z, ϱ) for fixed $(t, m, p, \boldsymbol{\eta})$.

Recall the definition of M^τ in (21). For a given p^* , consider a generic type of BSDEs on the enlarged filtration \mathbb{G} solved on $[\mathcal{T}(p^*), T]$

$$\begin{aligned} Y_t &= g(X_T^{\mathcal{T}(p^*)}, X_{\tau^*}^{\mathcal{T}(p^*)}, \tau^*(\mu)) + \int_t^T h(s, X_s^{\mathcal{T}(p^*)}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, Z_s) ds \\ &\quad - \int_t^T Z_s dW_s - \int_t^T \mathfrak{Z}_s dB_s - \int_t^T U_s dM_s^\tau - \int_t^T dM_s, \quad t \in [\mathcal{T}(p^*), T]. \end{aligned} \quad (23)$$

where M is a martingale orthogonal to (W, B, M^τ) . A solution to the BSDE (23) is a process $(Y, Z, \mathfrak{Z}, U, M) \in \mathcal{S}_{\mathbb{G}}^2[\mathcal{T}(p^*), T] \times \mathcal{G}_{\mathbb{G}}^2[\mathcal{T}(p^*), T] \times \mathcal{H}_{\mathbb{G}}^2[\mathcal{T}(p^*), T] \times \mathcal{H}_{\mathbb{G}, D}^2[\mathcal{T}(p^*), T] \times \mathcal{S}_{\mathbb{G}}^2[\mathcal{T}(p^*), T]$ on the probability space $(\Omega, \mathcal{F}, \mathbb{G}, \mathbb{P})$. If the pre-enlarged filtration \mathbb{F} is generated by the Brownian motions, then $M \equiv 0$. Note that the BSDE above is solved on a random interval even after conditioning on a p^* . The following proposition addresses the solvability of this BSDE. To differentiate the two types of admissibility, we denote a time-admissible control in \mathbb{A} by α and price-admissible control in \mathbb{A}^* by α^τ .

Proposition 3.7 *Suppose that $\mathbb{G} = \mathbb{F}^{\mathcal{T}, W, B, D}$ and fix a \mathbb{G} -progressive $\mathfrak{P} = (\mu, \vartheta)$. Given p^* , for each $t^* \in [0, T]$, there exists a unique solution $(Y^{t^*}, Z^{t^*}, \mathfrak{Z}^{t^*}, U^{t^*})$ to the following BSDE*

$$\begin{aligned} Y_t &= g(X_T^{t^*}, X_{\tau^*}^{t^*}, \tau^*(\mu)) + \int_t^T h(s, X_s^{t^*}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, Z_s) ds \\ &\quad - \int_t^T Z_s dW_s - \int_t^T \mathfrak{Z}_s dB_s - \int_t^T U_s dM_s^\tau, \quad t \in [t^*, T]. \end{aligned} \quad (24)$$

where X^{t^*} follows

$$X_t^{t^*} = \mathbb{1}_{\{t \geq t^*\}} (K_0/p^* + \sigma(W_t - W_{t^*})).$$

If we define the process $\hat{\alpha}_t^{t^*} = \mathbb{1}_{\{t \geq t^*\}} \hat{\alpha}(Z_t^{t^*}) \in \mathbb{A}^{t^*}$ for each t^* , then $\hat{\alpha}$ is time admissible and induces a price-admissible control $\hat{\alpha}^{\mathcal{T}(\cdot)} \in \mathbb{A}^*$. Moreover, $\hat{\alpha}^\tau$ minimizes (20) over \mathbb{A}^* .

Proof Using Assumption (E2) and (21), we can rewrite (24) as

$$\begin{aligned} Y_t &= g(X_T^{t^*}, X_{\tau^*}^{t^*}, \tau^*(\mu)) + \int_t^T h(s, X_s^{t^*}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, Z_s) + U_s k_s \mathbb{1}_{\{0 \leq s < \tau\}} ds \\ &\quad - \int_t^T Z_s dW_s - \int_t^T \mathfrak{Z}_s dB_s - \int_t^T U_s dD_s, \quad t \in [t^*, T]. \end{aligned}$$

Well-posedness follows from [31, Theorem 53.1]. We need to show that $\hat{\alpha}$ is jointly measurable when composing the t^* -by- t^* solutions. We first show that $t^* \mapsto \hat{\alpha}^{t^*}$ is \mathbb{P} -almost surely left-continuous in $\mathcal{H}_{\mathbb{G}}^2$.

Suppose we have a sequence $t_n^* \uparrow t^* \in [0, T]$, and let $\alpha_n^{t^*}$ and α^{t^*} be the corresponding control processes. Then we have

$$\begin{aligned} \|\alpha_n^{t^*} - \alpha^{t^*}\|_{\mathcal{H}_{\mathbb{G}}^2} &= \mathbb{E} \left[\int_0^T |\alpha_n^{t^*} - \alpha^{t^*}|^2 dt \right] \\ &= \mathbb{E} \left[\int_{t_n^*}^{t^*} |\alpha_n^{t^*}|^2 dt \right] + \mathbb{E} \left[\int_{t^*}^T |\alpha_n^{t^*} - \alpha^{t^*}|^2 dt \right] \\ &= \mathbb{E} \left[\int_{t_n^*}^{t^*} |\alpha_n^{t^*}|^2 dt \right] + \mathbb{E} \left[\int_{t^*}^T \left| \hat{\alpha}(Z_n^{t^*}) - \hat{\alpha}(Z_t^{t^*}) \right|^2 dt \right]. \end{aligned}$$

The first term goes to 0 by dominated convergence theorem since A is assumed to be bounded. To show the convergence of the second term, by continuity of \hat{a} it suffices to show $Z^{t_n^*} \xrightarrow{n \rightarrow \infty} Z^{t^*}$ in $\mathcal{H}_{\mathbb{Q}}^2[t^*, T]$. By the stability of BSDE solutions (e.g. [31, Proposition 54.2]), we have

$$\begin{aligned} \|Z^{t_n^*} - Z^{t^*}\|_{\mathcal{H}_{\mathbb{Q}}^2[t^*, T]}^2 &= \mathbb{E} \left[\int_{t^*}^T \left| Z_t^{t_n^*} - Z_t^{t^*} \right|^2 dt \right] \\ &\leq C \mathbb{E} \left[\left| g(X_T^{t_n^*}, X_{\tau^*}^{t_n^*}, \tau^*) - g(X_T^{t^*}, X_{\tau^*}^{t^*}, \tau^*) \right|^2 \right] \\ &\quad + C \mathbb{E} \left[\int_{t^*}^T \left| h(s, X_s^{t_n^*}, b_s^P, \tau^*, \langle \rho, \vartheta_s \rangle, Z_s^{t^*}) \right| \right. \\ &\quad \left. - h(s, X_s^{t^*}, b_s^P, \tau^*, \langle \rho, \vartheta_s \rangle, Z_s^{t^*})^2 ds \right]. \end{aligned}$$

It is easy to check that for any t , we have $X_t^{t_n^*} \xrightarrow{n \rightarrow \infty} X_t^{t^*}$, and by (S2) left-continuity is proved. The second condition in Lemma A.2 is satisfied, and invoking the lemma yields time-admissibility.

Therefore, for each $p^* \in [P_0, \infty)$, we can solve (23) with random entry time on $[\mathcal{T}(p^*), T]$ and obtain $\hat{\alpha}^{\mathcal{T}(p^*)}$. Note that altering p^* only affects the initial inventory K_0/p^* and entry time $\mathcal{T}(p^*)$. Since $P_0 > 0$, and also by Lemma 2.3, the function $p^* \mapsto \hat{\alpha}_t^{\mathcal{T}(p^*)}(\omega)$ is left-continuous for $\mathbb{Q} \times \lambda$ -almost every (ω, t) . Therefore, Lemma A.2 again implies joint measurability and therefore price-admissibility of $\hat{\alpha}^{\mathcal{T}(\cdot)}$.

Observe that (20) can be rewritten as $J_{weak}^{\mu, \vartheta}(\alpha^{\mathcal{T}}) = \mathbb{E}^{\nu_p}[J_{weak}^{\mu, \vartheta}(\alpha^{\mathcal{T}}|p^*)]$ for each price-admissible $\alpha^{\mathcal{T}}$, where $J_{weak}^{\mu, \vartheta}(\cdot|p^*)$ is the conditional objective given $\mathcal{P} = p^*$. That is

$$J_{weak}^{\mu, \vartheta}(\alpha^{\mathcal{T}}|p^*) := \mathbb{E}^{\hat{\alpha}^{\mathcal{T}(p^*)}} \left[g(X_T^{\mathcal{T}(p^*)}, X_{\tau^*}^{\mathcal{T}(p^*)}, \tau^*(\mu)) + \int_{\mathcal{T}(p^*) \wedge T}^T f(s, X_s^{\mathcal{T}(p^*)}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, \alpha_s^{\mathcal{T}(p^*)}) ds \right].$$

Then optimality of $\hat{\alpha}^{\mathcal{T}}$ follows if we show conditional optimality of $\hat{\alpha}^{\mathcal{T}(p^*)}$ for each p^* , which we fix from this point on. Take any price-admissible strategy $\beta^{\mathcal{T}}$. We can uniquely solve the following BSDE

$$\begin{aligned} Y_t^{\beta, p^*} &= g(X_T^{\mathcal{T}(p^*)}, X_{\tau^*}^{\mathcal{T}(p^*)}, \tau^*(\mu)) + \int_t^T H(s, X_s^{\mathcal{T}(p^*)}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, Z_s^{\beta, p^*}, \beta_s^{\mathcal{T}(p^*)}) ds \\ &\quad - \int_t^T Z_s^{\beta, p^*} dW_s - \int_t^T \mathfrak{Z}_s^{\beta, p^*} dB_s - \int_t^T U_s^{\beta, p^*} dM_s^{\tau}, \quad t \in [0, T]. \end{aligned} \tag{25}$$

We can also solve (25) on $[0, T]$ with $\hat{\alpha}$ as input. Then by (22) and uniqueness, the solution coincides with the solution of (23) on the interval $[\mathcal{T}(p^*), T]$. Comparison principle of (25) ([62, Proposition 4.3]) implies $Y_0^{\hat{\alpha}^{\mathcal{T}}, p^*} \leq Y_0^{\beta^{\mathcal{T}}, p^*}$ \mathbb{Q} -almost surely. Optional stopping theorem and Remark 2.5 (C4) imply

$$\begin{aligned} J_{weak}^{\mu, \vartheta}(\hat{\alpha}^{\mathcal{T}}|p^*) &= \mathbb{E}^{\mathbb{Q}^{\hat{\alpha}^{\mathcal{T}(p^*)}}} \left[Y_0^{\hat{\alpha}, p^*} \right] = \mathbb{E}^{\mathbb{Q}} \left[Y_0^{\hat{\alpha}, p^*} \right] \leq \mathbb{E}^{\mathbb{Q}} \left[Y_0^{\beta, p^*} \right] = \mathbb{E}^{\mathbb{Q}^{\beta^{\mathcal{T}(p^*)}}} \left[Y_0^{\beta, p^*} \right] \\ &= J_{weak}^{\mu, \vartheta}(\beta^{\mathcal{T}}|p^*). \end{aligned}$$

□

For the remainder of this section up until Remark 3.16, we take $\mathbb{G} = \mathbb{R}^{\mathcal{I}, W, B, D}$, so Proposition 3.7 applies. This result implies that for a given (μ, ϑ) we can find $\hat{\alpha}^{\mathcal{T}}$ by solving the problem p^* -by- p^* on Ω , and it is well defined for every p^* , not just ν_p almost every p^* . On the other hand, by uniqueness of the optimizer, we can also obtain $\hat{\alpha}^{\mathcal{T}}$ by solving the BSDE directly on the product space Ω_c .

Corollary 3.8 *Given $(\mu, \vartheta) : \Omega_0 \rightarrow (\mathcal{M}, \Theta)$ that is $\mathbb{F}^{B, D, \mathcal{T}}$ -progressive, there exists a unique solution $(Y, Z, \mathfrak{Z}, U) \in \mathcal{S}_{\mathbb{G}, c}^2 \times \mathcal{G}_{\mathbb{G}, c}^2 \times \mathcal{H}_{\mathbb{G}, c}^2 \times \mathcal{H}_{\mathbb{G}, c, D}^2$ to the following BSDE on the product space in the \mathbb{G} filtration*

$$\begin{aligned} Y_t &= g(X_t^{\mathcal{T}}, X_{\tau^*(\mu)}^{\mathcal{T}}, \tau^*(\mu)) + \int_t^T h(s, X_s^{\mathcal{T}}, b_s^P, \tau^*(\mu), \langle \rho, \vartheta_s \rangle, Z_s) ds \\ &\quad - \int_t^T Z_s dW_s - \int_t^T \mathfrak{Z}_s dB_s - \int_t^T U_s dM_s^{\tau}, \quad t \in [\mathcal{T} \wedge T, T]. \end{aligned} \quad (26)$$

Additionally, the process defined by $\hat{\alpha}_t := \mathbb{1}_{\{t \geq T\}} \hat{\alpha}(Z_t) \in \mathbb{A}^*$ is $\mathbb{P} \otimes dt$ almost surely identical to the one constructed in Proposition 3.7 and minimizes (20).

For any price-admissible $\alpha \in \mathbb{A}^*$, let \tilde{X}^{α} denote the solution of (12) defined on $\Omega_c \times [0, T]$. Then by Girsanov's theorem

$$\mathbb{P} \circ (\tilde{X}^{\alpha}, \mathcal{I}, \mathcal{T}, W, B)^{-1} = \mathbb{P}^{\alpha} \circ (X^{\mathcal{T}}, \mathcal{I}, \mathcal{T}, W^{\alpha}, B)^{-1}. \quad (27)$$

Under Assumption (E1), for any $p > 0$ we have

$$\sup_{\alpha \in \mathbb{A}^*} \mathbb{E}^{\mathbb{P}^{\alpha}} \left[\sup_{p^* \in [P_0, \infty)} \|X^{\mathcal{T}(p^*)}\|_{\mathcal{X}^*}^p \right] = \sup_{\alpha \in \mathbb{A}^*} \mathbb{E} \left[\sup_{p^* \in [P_0, \infty)} \|\tilde{X}^{\alpha, \mathcal{T}(p^*)}\|_{\mathcal{X}^*}^p \right] < \infty. \quad (28)$$

3.3.2 Fixed point from discretization

Instead of conditional measure flows given common noise (B, D) , we look at a piecewise constant approximation process. Suppose for $N \in \mathbb{N}$, a partition $\{0 = t_0 < t_1 < \dots < t_{N-1} < t_N = T\}$ on $[0, T]$ and a finite set (some grid on space) $\Lambda_N \subset \mathbb{R}$ are given. Define the Λ_N -valued finite process V^N on \mathcal{X} by

$$V_t^N(\beta) := \sum_{i=1}^N v_{i-1}(\beta) \mathbb{1}_{\{t \in [t_{i-1}, t_i)\}} + v_{N-1}(\beta) \mathbb{1}_{\{t=T\}}, \quad (29)$$

where each $v_i : \mathcal{X} \rightarrow \Lambda_N$ is a $\mathcal{F}_{t_i}^B$ measurable random variable taking values in the finite set, for $i = 0, \dots, N-1$. These knots are meant to approximate the Brownian common noise $B_{t_i}(\beta)$. We also have an additional source of common noise: the external burst time τ , which requires us to discretize the jump process D .

Let \mathcal{X}_D denote the space of processes on $[0, T]$ of the form $D_t(\eta) = \mathbb{1}_{\{\eta \leq t\}}$ for some $\eta \in [0, T]$. Equip \mathcal{X}_D with the natural metric $d(D(\eta), D(\eta')) = |\eta - \eta'|$. For $N \in \mathbb{N}$, define the \mathcal{X}_D valued process on \mathbb{R}_+ by:

$$D_t^N(\eta) := \sum_{i=1}^N \mathbb{1}_{\{\eta \leq t_{i-1}\}} \mathbb{1}_{\{t \in [t_{i-1}, t_i)\}} + \mathbb{1}_{\{\eta \leq t_{N-1}\}} \mathbb{1}_{\{t=T\}}. \quad (30)$$

Let $\mathcal{V}_N := \{A_1, \dots, A_{|\mathcal{V}_N|}\}$ denote the (finite) σ -algebra generated by (V^N, D^N) , and choose v_i 's such that $\mathbb{P}(A_k) > 0$ for every k . We now define the input domain for conditional

laws of the state and control. At this stage, we can work with A -valued controls. However, in anticipation of taking the limit in the space of relaxed controls, we switch to strict controls now. For $\alpha \in \mathbb{A}^*$, call $\gamma(\alpha) \in \Gamma$ its corresponding strict control. Define

$$\mathcal{M}_N := \left\{ \left(\mathcal{L}^{\alpha,1}(X_t^T), \dots, \mathcal{L}^{\alpha,|\mathcal{V}_N|}(X_t^T) \right) : \text{for some } \alpha \in \mathbb{A}^* \right\} \subset \mathcal{P}_1(\mathcal{X}^*)^{|\mathcal{V}_N|} \quad (31)$$

$$\Theta_N := \left\{ \left(\mathcal{L}^{\alpha,1}(\gamma(\alpha)), \dots, \mathcal{L}^{\alpha,|\mathcal{V}_N|}(\gamma(\alpha)) \right) : \text{for some } \alpha \in \mathbb{A}^* \right\} \subset \Theta^{|\mathcal{V}_N|} \quad (32)$$

where for each $\alpha \in \mathbb{A}^*$, $k \in 1, 2, \dots, |\mathcal{V}_N|$, $\mathcal{L}^{\alpha,k}$ denotes the conditional law under \mathbb{P}^α given A_k . Equip each coordinate of \mathcal{M}_N with the Wasserstein metric. Denote by $\vec{m} = (m_1, \dots, m_{|\mathcal{V}_N|})$ an element in \mathcal{M}_N .

Lemma 3.9 *The process $t \mapsto \mathbb{E}^{\mathbb{P}^\alpha}[X_t^T | A_k]$ is \mathbb{Q}_0 -almost surely continuous for each k for all $\alpha \in \mathbb{A}^*$. Consequentially, the mapping $\mathcal{M}_N \ni \vec{m} \mapsto \bar{\tau}(\mu^N)$ is \mathbb{Q}_0 -almost surely continuous on the closure of \mathcal{M}_N , where for $\omega_0 = (\beta, \eta) \in \Omega_0 = \mathcal{X} \times \mathbb{R}_+$, $\mu^N((\beta, \eta)) := \sum_{k=1}^{|\mathcal{V}_N|} m_k \mathbb{1}_{\{(\beta, D(\eta)) \in A_k\}}$.*

Proof Take a sequence $t_n \xrightarrow{n \rightarrow \infty} t_\infty$ in $[0, T]$. Note that the event that $X_{t_n}^T$ does not converge to $X_{t_\infty}^T$ is, up-to a \mathbb{P}^α -null set, contained in $\{\mathcal{T} = t_\infty\}$. Then by dominated convergence theorem, it suffices to show that for all t , $\mathbb{P}^\alpha(\{\mathcal{T} = t_\infty\} | A_k) = 0$ for each A_k , which is implied by $\mathbb{P}^\alpha(\mathcal{T} = t_\infty) = \mathbb{P}(\mathcal{T} = t_\infty) = 0$. This follows from price dynamics (2) and the fact that the only jump of P is negative by Assumptions (2). Then continuity of mean processes for each k implies continuity of the endogenous burst mapping (see the proof of [62, Theorem 6.1]). Taking closure in Wasserstein space preserves the continuity of the mean processes. \square

Lemma 3.10 *The set \mathcal{M}_N is Wasserstein pre-compact and convex in $\mathcal{P}(\mathcal{X}^*)^{|\mathcal{V}_N|}$, and Θ_N is Wasserstein compact and convex in $\Theta^{|\mathcal{V}_N|}$.*

Proof The statement on Θ_N is immediate given that A is compact and convex. For convexity of \mathcal{M}_N , take $\alpha^1, \alpha^2 \in \mathbb{A}^*$. It suffices to show that for all $\lambda \in [0, 1]$, there is $\alpha \in \mathbb{A}^*$ such that

$$\frac{d\mathbb{P}^\alpha}{d\mathbb{P}} = \lambda \frac{d\mathbb{P}^{\alpha^1}}{d\mathbb{P}} + (1 - \lambda) \frac{d\mathbb{P}^{\alpha^2}}{d\mathbb{P}}.$$

See [62, Lemma 4.7] for a proof. As for Wasserstein pre-compactness, it suffices to show this for each marginal $k \in \{1, \dots, |\mathcal{V}_N|\}$. One can easily adapt the proof of [62, Lemma 5.9], but with conditional expectations given A_k instead of regular expectations. Indeed, using the fact that $\mathbb{E}[\xi | A_k] \leq \mathbb{E}[\xi] / \mathbb{P}(A_k)$ for any integrable non-negative random variable ξ , the constant of the bounds in the proof of [62, Lemma 5.9] is multiplied to an additional term $1 / \mathbb{P}(A_k)$, which is finite for all k by assumption. We shall skip the details to reduce redundancy. \square

Suppose we take any $\vec{m}^N = (m_1^N, \dots, m_{|\mathcal{V}_N|}^N) \in \mathcal{M}_N$ and $\vec{a}^N = (a_1^N, \dots, a_{|\mathcal{V}_N|}^N) \in \Theta_N$. Define the inputs to the optimization problem: for each $(\beta, \eta) = \omega_0 \in \Omega_0 = \mathcal{X} \times \mathbb{R}_+$,

$$\mu^N(\omega_0) := \sum_{k=1}^{|\mathcal{V}_N|} m_k^N \mathbb{1}_{\{(\beta, D(\eta)) \in A_k\}} \in \mathcal{P}_1(\mathcal{X}), \quad \vartheta^N(\omega_0) := \sum_{k=1}^{|\mathcal{V}_N|} a_k^N \mathbb{1}_{\{(\beta, D(\eta)) \in A_k\}} \in \Theta. \quad (33)$$

Again (μ^N, ϑ^N) can be viewed as a $\mathbb{F}^{B,D}$ -measurable random process taking values (μ_t^N, ϑ_t^N) in $\mathcal{P}(\mathbb{R}) \times \mathcal{P}(\mathcal{P}(A))$. By Proposition 3.7 we obtain strict, optimal control $\hat{\alpha}^{T,N} \in \mathbb{A}^*$ with $\hat{\alpha}_t^{T,N} = \hat{\alpha}(Z_t^{T,N})$ along with the probability measure $\mathbb{P}^N := \mathbb{P}^{\hat{\alpha}^{T,N}}$ on Ω_c , and $W_t^N :=$

$W_t - \int_0^t \sigma^{-1} \hat{\alpha}_s^{\mathcal{T},N} ds$ is a Brownian motion under \mathbb{P}^N . Denote by γ^N the optimal control in the strict relaxed form $\gamma(\hat{\alpha}^{\mathcal{T},N}) = (\delta_{\hat{\alpha}_t^{\mathcal{T},N}} dt)_{t \in [0, T]}$. Define output conditional measures $(\vec{\mathfrak{m}}^{N,out}, \vec{\mathfrak{a}}^{N,out}) \in \mathcal{M}_N \times \Theta_N$:

$$\mathfrak{m}_k^{N,out}(\cdot) := \frac{\mathbb{P}^N(A_k \cap \{X^T \in \cdot\})}{\mathbb{P}^N(A_k)}, \quad \mathfrak{a}_k^{N,out}(\cdot) := \frac{\mathbb{P}^N(A_k \cap \{\gamma^N \in \cdot\})}{\mathbb{P}^N(A_k)}. \quad (34)$$

We have now defined our fixed point mapping:

$$\Phi^N : \mathcal{M}_N \times \Theta_N \ni (\vec{\mathfrak{m}}^N, \vec{\mathfrak{a}}^N) \mapsto (\vec{\mathfrak{m}}^{N,out}, \vec{\mathfrak{a}}^{N,out}) \in \mathcal{M}_N \times \Theta_N.$$

Lemma 3.11 *For each $N \in \mathbb{N}$, the mapping Φ^N is continuous.*

Proof We shall fix N and drop the notation to avoid confusion with the proof steps for sequential continuity. Recall $\hat{\alpha}^{\mathcal{T}}$ is obtained by solving the optimization problem p^* -by- p^* . Let $\mathbb{P}^{\hat{\alpha}^{\mathcal{T}(p^*)}}$ denote the conditional probability measure on Ω from Girsanov transform for a given $p^* \in [P_0, \infty)$. Take a sequence of vectors $(\vec{\mathfrak{m}}^n, \vec{\mathfrak{a}}^n) \xrightarrow{n \rightarrow \infty} (\vec{\mathfrak{m}}^\infty, \vec{\mathfrak{a}}^\infty) \in \mathcal{M}_N \times \Theta_N$ in the sense that each coordinate converges in Wasserstein distance, and define correspondingly (μ^n, ϑ^n) for $n \in \mathbb{N} \cup \{\infty\}$. Then \mathbb{Q}_0 almost surely, μ^n converges to μ^∞ in \mathcal{M} and ϑ^n converges to ϑ^∞ in Θ . For $n \in \mathbb{N} \cup \{\infty\}$ define the discretized conditional probability measures μ^n and ϑ^n as in (33), and let $Z^{\mathcal{T}(p^*),n}$ be part of the unique solution to the BSDE (23) with $M \equiv 0$ and input (μ^n, ϑ^n) . Let $\hat{\alpha}_t^{\mathcal{T}(p^*),n} = \hat{\alpha}(Z_t^{\mathcal{T}(p^*),n})$ be the optimal control given $\mathcal{P} = p^*$.

We first show that $\hat{\alpha}^{\mathcal{T}(p^*),n} \xrightarrow{n \rightarrow \infty} \hat{\alpha}^{\mathcal{T}(p^*),\infty}$ in $\mathcal{H}_{\mathbb{G}}^2$ for every $p^* \in [P_0, \infty)$. Recall that optimal controls are continuous in $Z^{\mathcal{T}(p^*)}$, so it suffices to show $Z^{\mathcal{T}(p^*),n} \xrightarrow{n \rightarrow \infty} Z^{\mathcal{T}(p^*),\infty}$ in $\mathcal{H}_{\mathbb{G}}^2$. Using the stability property of BSDE solutions [31, Proposition 54.2], convergence of $Z^{\mathcal{T}(p^*),n}$ is immediately implied if we show

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[|\Delta_n g|^2 + \int_{\mathcal{T}(p^*)}^T |\Delta_n h_s|^2 ds \right] = 0,$$

where $\Delta_n g := g \left(X_T^{\mathcal{T}(p^*)}, X_{\tau^*(\mu^n)}^{\mathcal{T}(p^*)}, \tau^*(\mu^n) \right) - g \left(X_T^{\mathcal{T}(p^*)}, X_{\tau^*(\mu^\infty)}^{\mathcal{T}(p^*)}, \tau^*(\mu^\infty) \right)$ and

$$\begin{aligned} \Delta_n h_s &:= h(s, X_s^{\mathcal{T}(p^*)}, b_s^P, \tau^*(\mu^n), \langle \rho, \vartheta_s^n \rangle, Z_s^{\mathcal{T}(p^*),\infty}) - h(s, X_s^{\mathcal{T}(p^*)}, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, \\ &\quad Z_s^{\mathcal{T}(p^*),\infty}). \end{aligned}$$

By continuity of g we know that $\Delta_n g$ converges to 0 in probability if both $\tau^*(\mu^n) \xrightarrow{n \rightarrow \infty} \tau^*(\mu^\infty)$ and $X_{\tau^*(\mu^n)} \xrightarrow{n \rightarrow \infty} X_{\tau^*(\mu^\infty)}$ in probability. Lemma 3.9 gives us convergence in burst time. Also, the inventory is continuous everywhere except at entry $\mathcal{T}(p^*)$. Therefore, $X_{\tau^*(\mu^n)}$ does not converge in probability to $X_{\tau^*(\mu^\infty)}$ only if $\mathcal{T}(p^*) = \tau^*(\mu^\infty)$, which also has probability 0. Therefore,

$$\mathbb{P}(|\Delta_n g|^2 \xrightarrow{n \rightarrow \infty} 0) \geq 1 - \mathbb{P}(\{\mathcal{T}(p^*) = \tau^*(\mu^\infty)\}) = 1.$$

By dominated convergence theorem, we have $\mathbb{E}[|\Delta_n g|^2] \xrightarrow{n \rightarrow \infty} 0$.

Now let I_n denote the random interval $[\tau^*(\mu^n) \wedge \tau^*(\mu^\infty), \tau^*(\mu^n) \vee \tau^*(\mu^\infty)]$. Observe that

$$|\mathbb{1}_{\{0 \leq t < \tau^*(\mu^n)\}} - \mathbb{1}_{\{0 \leq t < \tau^*(\mu^\infty)\}}|^2 = |1 - \mathbb{1}_{\{t \geq \tau^*(\mu^n)\}} - 1 + \mathbb{1}_{\{t \geq \tau^*(\mu^\infty)\}}|^2 = \mathbb{1}_{\{t \in I_n\}}.$$

Define $I_n^{p^*} := I_n \cap [\mathcal{T}(p^*), T]$. Remark 2.5 (C1) implies

$$\begin{aligned} \mathbb{E} \left[\int_{\mathcal{T}(p^*)}^T |\Delta_n h_s|^2 ds \right] &\leq \mathbb{E} \left[\int_{\mathcal{T}(p^*)}^T \left| f_b(s, X_s^{\mathcal{T}(p^*)}, b_s^P) \mathbb{1}_{\{0 \leq t < \tau^*(\mu^n)\}} - f_b(s, X_s^{\mathcal{T}(p^*)}, b_s^P) \mathbb{1}_{\{0 \leq t < \tau^*(\mu^\infty)\}} \right|^2 ds \right] \\ &\quad + \mathbb{E} \left[\int_{\mathcal{T}(p^*)}^T \left| f_c(s, X_s^{\mathcal{T}(p^*)}, \langle \rho, \vartheta_s^n \rangle) \mathbb{1}_{\{s \geq \tau^*(\mu^n)\}} - f_c(s, X_s^{\mathcal{T}(p^*)}, \langle \rho, \vartheta_s^\infty \rangle) \mathbb{1}_{\{s \geq \tau^*(\mu^\infty)\}} \right|^2 ds \right] \\ &\leq \mathbb{E} \left[\int_{I_n^{p^*}} |f_b(s, X_s^{\mathcal{T}(p^*)}, b_s^P)|^2 ds \right] + 2\mathbb{E} \left[\int_{I_n^{p^*}} |f_c(s, X_s^{\mathcal{T}(p^*)}, \langle \rho, \vartheta_s^\infty \rangle)|^2 ds \right] \\ &\quad + 2\mathbb{E} \left[\int_{\mathcal{T}(p^*)}^T \left| f_c(s, X_s^{\mathcal{T}(p^*)}, \langle \rho, \vartheta_s^n \rangle) - f_c(s, X_s^{\mathcal{T}(p^*)}, \langle \rho, \vartheta_s^\infty \rangle) \right|^2 \mathbb{1}_{\{s \geq \tau^*(\mu^n)\}} ds \right]. \end{aligned}$$

Lemma 3.9 implies $\mathbb{1}_{\{t \in I_n^{p^*}\}} \xrightarrow{n \rightarrow \infty} \mathbb{1}_{\{t = \tau^*(\mu^\infty)\}}$ almost surely. Under Assumption (E1) and Remark 2.5 (C1), all terms converge to 0 by dominated convergence theorem and Fubini's theorem. Therefore, for all $p^* \in [P_0, \infty)$, $\hat{\alpha}^{\mathcal{T}(p^*)}, n \xrightarrow{n \rightarrow \infty} \hat{\alpha}^{\mathcal{T}(p^*)}, \infty$ in $\mathcal{H}_{\mathbb{G}}^2$. Since A is assumed to be bounded, this implies that $\hat{\alpha}^n \xrightarrow{n \rightarrow \infty} \hat{\alpha}^\infty$ in the $\mathcal{H}_{\mathbb{G}, c}^2$ sense on the product space Ω_c .

For each $n \in \mathbb{N} \cup \{\infty\}$, denote by $\mathbb{P}^n := \mathbb{P}^{\hat{\alpha}^n}$ the probability measure on Ω_c from Girsanov transformation. By construction, $\mathbb{P}^n \ll \mathbb{P}$ for all n and

$$\frac{d\mathbb{P}^n}{d\mathbb{P}^\infty} = \mathcal{E} \left(\int_0^\cdot \sigma^{-1} (\hat{\alpha}_t^n - \hat{\alpha}_t^\infty) dW_t \right)_T.$$

Therefore, by boundedness of A we can calculate the relative entropy

$$\mathcal{H}(\mathbb{P}^\infty | \mathbb{P}^n) = -\mathbb{E}^{\mathbb{P}^\infty} \left[\log \frac{d\mathbb{P}^n}{d\mathbb{P}^\infty} \right] = \frac{1}{2} \mathbb{E}^{\mathbb{P}^\infty} \left[\int_0^T \sigma^{-2} |\hat{\alpha}_t^n - \hat{\alpha}_t^\infty|^2 dt \right] \xrightarrow{n \rightarrow \infty} 0,$$

Pinsker's inequality implies that \mathbb{P}^n converges to \mathbb{P}^{∞, p^*} in total variation. By triangular inequality and the convergence in $\mathcal{H}_{\mathbb{G}, c}^2$ of controls, we have

$$\mathbb{P}^n \circ (\hat{\alpha}_t^n)^{-1} \xrightarrow{n \rightarrow \infty} \mathbb{P}^\infty \circ (\hat{\alpha}_t^\infty)^{-1} \text{ in } dt\text{-measure.}$$

Bounded convergence theorem yields $\mathbb{P}^n \circ (\gamma(\hat{\alpha}^n))^{-1} \xrightarrow{n \rightarrow \infty} \mathbb{P}^\infty \circ (\gamma(\hat{\alpha}^\infty))^{-1}$ in the stable topology. Since $\mathbb{P}(A_k) > 0$ for each k , this implies convergence of $\mathfrak{a}_k^{n, out} \xrightarrow{n \rightarrow \infty} \mathfrak{a}_k^{\infty, out}$ as well. Boundedness of A ensures Wasserstein convergence as well. $\mathbb{P}^n \xrightarrow{n \rightarrow \infty} \mathbb{P}^\infty$ in total variation also implies weak convergence of $\mathfrak{m}_k^{n, out} \xrightarrow{n \rightarrow \infty} \mathfrak{m}_k^{\infty, out}$ for each k . To show Wasserstein convergence, it suffices to show uniform integrability (see e.g. [63, Theorem 6.9]).

$$\begin{aligned} &\lim_{R \rightarrow \infty} \limsup_{n \rightarrow \infty} \sup_{k \in \{1, \dots, |\mathcal{V}_N|\}} \int_{\{\|\mathfrak{x}\|_{\mathcal{X}^*} > R\}} \|\mathfrak{x}\|_{\mathcal{X}^*} \mathfrak{m}_k^{n, out}(d\mathfrak{x}) \\ &\leq \lim_{R \rightarrow \infty} \sup_{\alpha \in \mathbb{A}^*} \sup_{k \in \{1, \dots, |\mathcal{V}_N|\}} \frac{1}{\mathbb{P}(A_k)} \mathbb{E}^{\mathbb{P}^\alpha} \left[\|X^T\|_{\mathcal{X}^*} \mathbb{1}_{\{\|X^T\|_{\mathcal{X}^*} > R\}} \right] = 0 \end{aligned}$$

which follows from (28). Therefore, since $|\mathcal{V}_N|$ is finite:

$$\Phi^N(\vec{\mathfrak{m}}^n, \vec{\mathfrak{a}}^n) \xrightarrow{n \rightarrow \infty} \Phi^N(\vec{\mathfrak{m}}^\infty, \vec{\mathfrak{a}}^\infty) \text{ in } \mathcal{M}_N \times \Theta_N,$$

and continuity of Φ^N holds for all N . \square

Proposition 3.12 *The mapping Φ^N admits a fixed point $(\vec{\mathfrak{m}}^N, \vec{\mathfrak{a}}^N) \in \mathcal{M}_N \times \Theta_N$ for all $N \in \mathbb{N}$.*

Proof Let $\overline{\mathcal{M}}_N$ denote the closure of \mathcal{M}_N , which by Lemma 3.10 is convex and compact. Note that the input (μ^N, ϑ^N) to the BSDE is still well-defined for $\tilde{\mathbf{m}}^N \in \overline{\mathcal{M}}_N$, except that ϑ^N might not be the law of a strict control anymore. Therefore, we can define Φ^N on the larger domain $\overline{\mathcal{M}}_N \times \Theta_N$. Moreover, the Wasserstein closure preserves the continuity of the mean process. This implies that $\tau^*(\cdot)$ is still continuous on $\overline{\mathcal{M}}_N$ (see the proof of [62, Theorem 6.1]), so continuity of Φ^N still holds. Applying Brouwer's fixed point theorem (e.g. [4, Corollary 17.56]) yields a fixed point $(\tilde{\mathbf{m}}^N, \tilde{\mathbf{a}}^N)$ of Φ^N . However, the range of Φ^N is still strictly in $\mathcal{M}_N \times \Theta_N$ since allowing μ to take values in $\overline{\mathcal{M}}_N$ does not affect how we construct the optimal control and its corresponding state process. Then $(\tilde{\mathbf{m}}^N, \tilde{\mathbf{a}}^N) \in \mathcal{M}_N \times \Theta_N$ is also a fixed point of Φ had we not enlarged the domain. \square

For each $N \in \mathbb{N}$, let $\alpha^N \in \mathbb{A}^*$ be the equilibrium strategy from a fixed point $(\tilde{\mathbf{m}}^N, \tilde{\mathbf{a}}^N)$ from Proposition 3.12 and $\mathbb{P}^N = \mathbb{P}^{\alpha^N}$. Define the corresponding random laws (μ^N, ϑ^N) as in (33). Then by construction, α^N minimizes $J_{\text{weak}}^{\mu^N, \vartheta^N}$ defined in (20) over \mathbb{A}^* .

Remark 3.13 For each $N \in \mathbb{N}$, we work with the same filtration $\mathbb{F}^{\mathcal{T}, W, B, D}$. In fact, by the argument in [62, Proposition A.10], $\hat{\alpha}^N$ also minimizes (20) had we allowed for a larger filtration $\widetilde{\mathbb{F}} \supseteq \mathbb{F}^{\mathcal{T}, W, B, D}$. We provide a proof to keep the paper self-contained.

Proof We fix $N \in \mathbb{N}$ and recall that the uncontrolled state process is defined by

$$X_t^{\mathcal{T}} := K_0 / \mathcal{P} + \sigma(W_t - W_T) \text{ for } t \geq \mathcal{T} \text{ and } X_t^{\mathcal{T}} := 0 \text{ for } t \in [0, \mathcal{T}).$$

Recall from Remark 3.6 that $\hat{a}(\cdot)$ is a continuous function that minimizes the Hamiltonian. For each $\beta \in \mathbb{A}^*$ defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P}, \widetilde{\mathbb{F}})$, by [31, Theorem 53.1] there exists a unique solution $(Y^\beta, Z^\beta, \mathfrak{Z}^\beta, U^\beta, N^\beta) \in \mathcal{S}_{\widetilde{\mathbb{F}}}^2 \times \mathcal{G}_{\widetilde{\mathbb{F}}}^2 \times \mathcal{H}_{\widetilde{\mathbb{F}}}^2 \times \mathcal{H}_{\widetilde{\mathbb{F}}, D}^2 \times \mathcal{S}_{\widetilde{\mathbb{F}}}^2$ to the following BSDE

$$\begin{aligned} Y_t &= g(X_T^{\mathcal{T}}, X_{\tau^*(\mu^N)}^{\mathcal{T}}, \tau^*(\mu^N)) + \int_t^T H(s, X_s^{\mathcal{T}}, b_s^P, \tau^*(\mu^N), \langle \rho, \vartheta_s^N \rangle, Z_s, \beta_s) ds \\ &\quad - \int_t^T Z_s dW_s - \int_t^T \mathfrak{Z}_s dB_s - \int_t^T U_s dM_s^{\tau} - \int_t^T dN_s, \quad t \in [0, T]. \end{aligned} \quad (35)$$

Remark 3.6 (S2) also implies well-posedness for the following BSDE solved on $(\Omega, \mathcal{F}, \mathbb{P}, \widetilde{\mathbb{F}})$:

$$\begin{aligned} Y_t &= g(X_T^{\mathcal{T}}, X_{\tau^*(\mu^N)}^{\mathcal{T}}, \tau^*(\mu^N)) + \int_t^T \mathbb{1}_{\{s \geq \mathcal{T}\}} h(s, X_s^{\mathcal{T}}, b_s^P, \tau^*(\mu^N), \langle \rho, \vartheta_s^N \rangle, Z_s) ds \\ &\quad - \int_t^T Z_s dW_s - \int_t^T \mathfrak{Z}_s dB_s - \int_t^T U_s dM_s^{\tau} - \int_t^T dN_s, \quad t \in [0, T]. \end{aligned} \quad (36)$$

whose unique solution we denote by $(\hat{Y}, \hat{Z}, \hat{\mathfrak{Z}}, \hat{U}, \hat{N})$. However, since (μ^N, ϑ^N) are $\mathbb{F}^{B, D}$ measurable, uniqueness of the solution implies that \hat{N} is $\mathbb{P}^N \otimes dt$ almost surely zero, and $(\hat{Y}, \hat{Z}, \hat{\mathfrak{Z}}, \hat{U})$ coincides with the solution of (26) on $[\mathcal{T}, T]$. Recall from Corollary 3.8 and the construction of the fixed point mapping Φ^N that $\mathbb{P}^N \otimes dt$ almost surely, we must also have $\alpha_t^N = \hat{a}(\hat{Z}_t) \mathbb{1}_{\{t \geq \mathcal{T}\}}$.

Recall from Remark 2.5 (C4) that for any admissible control, the Hamiltonian is 0 before entry. Then the generator of (35) is $\mathbb{P} \otimes dt$ -almost surely greater than the generator of (36), and they are equal when we take $\beta = \alpha^N$. If a comparison principle for the general BSDE (35) holds, then Remark 3.13 follows from the argument in the proof of Proposition 3.7. To ease some notation, for $t \in [0, T]$ we denote

$$\Delta H_t(\beta, \alpha^N) := H(t, X_t^{\mathcal{T}}, b_t^P, \tau^*(\mu^N), \langle \rho, \vartheta_t^N \rangle, Z_t^\beta, \beta_t) - H(t, X_t^{\mathcal{T}}, b_t^P, \tau^*(\mu^N), \langle \rho, \vartheta_t^N \rangle, Z_t^{\alpha^N}, \alpha_t^N).$$

Due to the presence of compensated martingale M^τ and orthogonal martingale N , additional conditions are required for comparison principle to hold. In light of [23], a sufficient condition is the existence of an equivalent measure $\tilde{\mathbb{P}}$ to \mathbb{P} such that

$$S_t := - \int_t^T \Delta H_s(\beta, \alpha^N) ds + \int_t^T (Z_s^\beta - Z_s^{\alpha^N}) dW_s + \int_t^T (\mathfrak{Z}_s^\beta - \mathfrak{Z}_s^{\alpha^N}) dB_s \\ + \int_t^T (U_s^\beta - U_s^{\alpha^N}) dM_s^\tau + \int_t^T dN_s^\beta - \int_t^T dN_s^{\alpha^N}, \quad t \in [0, T] \quad (37)$$

is a martingale under $\tilde{\mathbb{P}}$. This probability can be constructed by defining

$$\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} := \mathcal{E} \left(\int_0^{\cdot} \frac{\Delta H_t(\beta, \alpha^N)}{Z_t^\beta - Z_t^{\alpha^N}} \mathbb{1}_{\{Z_t^\beta - Z_t^{\alpha^N} \neq 0\}} dW_t \right)_T.$$

Since H is Lipschitz in z , we can apply Girsanov's theorem and conclude by [23, Theorem 2] that comparison principle holds. Denote again by \mathbb{P}^β the probability measure corresponding to β , which is defined as in (19). Recall that $\mathbb{P}^N, \mathbb{P}^\beta$ and \mathbb{P} all agree at time $t = 0$ and the Hamiltonian is 0 for any admissible strategy before T . Using comparison principle and applying optional stopping theorem yield

$$J_{\text{weak}}^{\mu^N, \vartheta^N}(\alpha^N) = \mathbb{E}^{\mathbb{P}^N} \left[Y_T^{\alpha^N} \right] = \mathbb{E}^{\mathbb{P}^N} \left[Y_0^{\alpha^N} \right] = \mathbb{E}^{\mathbb{P}} \left[Y_0^{\alpha^N} \right] \leq \mathbb{E}^{\mathbb{P}^\beta} \left[Y_0^\beta \right] = \mathbb{E}^{\mathbb{P}^\beta} \left[Y_T^\beta \right] = J_{\text{weak}}^{\mu^N, \vartheta^N}(\beta).$$

The inequality holds for all \mathbb{F} -admissible β . \square

Remark 3.14 Since the state variable is linearly controlled, for each $\gamma \in \mathbb{I}^*$ we can define an A -valued control $\alpha_t = \int_A a \gamma_t(da)$ in \mathbb{A}^* that corresponds to the same state process. Moreover, Remark 2.5 (C3) and Jensen's inequality imply that this control α is at least as good as γ in terms of minimizing the objective value. See [34, Theorem A.9] and [43, Theorem 4.11] for more rigorous arguments on existence of strict controls achieving the same value and regarding the weak formulation of the problem. In other words, given inputs (μ, ϑ) , the optimal A -valued control is also optimal over relaxed controls for the relaxed objective (17).

3.3.3 Solutions as weak limit points

Before taking $N \rightarrow \infty$, we shall “lift” the marginal laws of \mathcal{X}^T and γ to a joint law in order to carry more information along the way. Let $\gamma^N = \gamma(\alpha^N)$ be the strict control corresponding to the equilibrium strategy α^N from the fixed point mapping. Recall that $\mathcal{I} = (K_0, \mathcal{P})$, taking values in $\mathbb{R} \times [p^*, \infty)$ denotes the initial wealth and entry threshold. Define on Ω_0 the lifted random law $\mathfrak{M}^N \in \mathcal{P}(\mathcal{X}^* \times \mathcal{X} \times \Gamma \times \mathbb{R}^2)$ by

$$\mathfrak{M}^N(\omega_0) := \sum_{k=1}^{|\mathcal{V}_N|} \mathbb{1}_{\{(\beta, \eta) \in A_k\}} \mathcal{L}^{\alpha^N, k} \left((X^T, W^N, \gamma^N, \mathcal{I}) \right), \quad \omega_0 = (\beta, \eta) \in \Omega_0. \quad (38)$$

Recall the notation from Definition 3.3 that $\mathfrak{M}^{N,x}$ and $\mathfrak{M}^{N,\gamma}$ denote the first and third marginals of \mathfrak{M}^N , which are the conditional law of (X^T, γ^N) under (V^N, D^N) .

Lemma 3.15 The sequence $\mathbb{P}^N \circ (B, \tau, \mathfrak{M}^N, X^T, W^N, \gamma^N, \mathcal{I})^{-1}$ is tight.

Proof Assumption (E1) implies that

$$\mathbb{E}^{\mathbb{P}^N} \left[\sup_{t \in [0, T]} |X_t^T| \right] \leq C$$

for some $C > 0$ that does not depend on N . Let s be a \mathbb{G} stopping time and $\delta > 0$. Then

$$\mathbb{E}^{\mathbb{P}^N} \left[|X_{s+\delta}^T - X_s^T| \right] \leq \mathbb{E}^{\mathbb{P}^N} \left[\int_s^{s+\delta} |\hat{\alpha}_t^N| dt + \sigma |W_{\rho+\delta}^N - W_\rho^N| \right] \leq C\delta^{1/2}$$

for a possibly different C . These two conditions are sufficient for Aldous' criterion for tightness of $\mathbb{P}^N \circ (X^T)^{-1}$. For each N , (W^N, B) remain independent Brownian motions under \mathbb{P}^N , and (τ, \mathcal{I}) also retain the same law throughout by independence. Compactness of A implies the tightness of $\mathbb{P}^N \circ (\gamma^N)^{-1}$.

Now tightness of $\mathbb{P}^N \circ (\mathfrak{M}^N)^{-1}$ in the weak topology follows from the tightness of $\mathbb{P}^N \circ (X^T, W^N, \gamma^N, \mathcal{I})^{-1}$ (see the proof of [18, Lemma 3.16]). As suggested by Lemma 3.9, we will need to equip $\mathcal{P}(\mathcal{X}^* \times \mathcal{X} \times \Gamma \times \mathbb{R}^2)$ with the Wasserstein metric in order to guarantee continuity of $\bar{\tau}$, where the product space is equipped with the l_1 metric. By [17, Corollary 5.6] and boundedness of A , the proof of [18, Lemma 3.16] implies that it suffices to show uniform square-integrability of $\|X^T\|_{\mathcal{X}^*} + \|W^N\|_\infty$. Since W^N is a Brownian motion under \mathbb{P}^N , we only need to show that

$$\lim_{R \rightarrow \infty} \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\|X^T\|_{\mathcal{X}^*} \mathbb{1}_{\{\|X^T\|_{\mathcal{X}^*} \geq R\}} \right] < \infty,$$

which is implied by (28). Finally, tightness of the marginals implies that of the joint law. \square

Remark 3.16 Let $\bar{\mathbb{P}} := \mathbb{P}^\infty \circ (B, \tau, \mathfrak{M}^\infty, X^T, W^\infty, \gamma^\infty, \mathcal{I})^{-1}$ be a limit point in Lemma 3.15. Since we work with a weak limit, we only care about the law, not the processes themselves. Therefore, without loss of generality, we can take $\bar{\Omega}_0 := \mathbb{R}_+ \times \mathcal{X} \times \mathcal{P}_1(\mathcal{X}^* \times \mathcal{X} \times \Gamma \times \mathbb{R}^2)$, $\bar{\Omega}_1 := \mathcal{X}^* \times \mathcal{X} \times \Gamma \times \mathbb{R}^2$. Then let $(\tau, B, \mathfrak{M}^\infty)$, $(X, W^\infty, \gamma^\infty, \mathcal{I})$ be the respective canonical processes on $\bar{\Omega} := \bar{\Omega}_0 \times \bar{\Omega}_1$, and $\bar{\mathbb{P}}$ a probability measure on $\bar{\Omega}$. For each N , define also on $\bar{\Omega}$ the law $\bar{\mathbb{P}}^N := \mathbb{P}^N \circ (B, \tau, \mathfrak{M}^N, X^T, W^N, \gamma^N, \mathcal{I})^{-1}$. We can obviously drop the ∞ from the notation (or even N , since we can always work on the canonical space), but we keep it to avoid confusion.

Define the jump process D , price process P and entry time \mathcal{T} , thanks to the strong solvability from Proposition A.1. Take \mathbb{G} to be the $\bar{\mathbb{P}}$ -completed natural filtration generated by $(B, D, \mathfrak{M}^\infty, X^T, W^\infty, \gamma^\infty, \mathcal{I})$, which is again the progressively enlarged version of $\mathbb{F}^{B, \mathfrak{M}^\infty, X, W^\infty, \gamma^\infty, \mathcal{I}}$ by $D_t = \mathbb{1}_{\{\tau \leq t\}}$, so τ is a \mathbb{G} -inaccessible stopping time. Let $\mathcal{F} = \mathcal{G}_T \vee \sigma(\tau)$. Weak convergence of γ^N to γ^∞ implies that X^T with entry time \mathcal{T} satisfies the relaxed state SDE (16) on $(\bar{\Omega}, \mathcal{F}, \bar{\mathbb{P}}, \mathbb{G})$.

Note that the existence of fixed points of the discretized problem and tightness result both hold for arbitrary choice of discretized process V^N . Now we need to specify the time and space grids to ensure V^N approximates B well enough. We will prove Theorem 3.4 by verifying that \mathfrak{M}^∞ defined on $(\bar{\Omega}, \mathcal{F}, \bar{\mathbb{P}}, \mathbb{G})$ satisfies the three required conditions in Definition 3.3, namely consistency, compatibility and optimality.

3.3.4 Consistency

Under a slight abuse of notation of re-indexing N , we take the same processes used in [14, 18] where time is discretized to the dyadic mesh and space projected to a more refined grid. Specifically, for a fixed $N \in \mathbb{N}$, let $t_i = \frac{iT}{2^N}$ for $i = 0, \dots, 2^N$. Set $v_0 = 0$ and $v_i = v_{i-1} + \Pi^{(N)}(B_{t_i} - B_{t_{i-1}})$, where the projection map $\Pi^{(N)} : \mathbb{R} \rightarrow \mathbb{R}$ is defined as

$$\Pi^{(N)}(x) = \begin{cases} 4^{-N} \lfloor 4^N x \rfloor & |x| \leq 4^N \\ 4^N \text{sign}(x) & |x| > 4^N \end{cases}.$$

Then on the event $E_N := \{\|B\|_\infty \leq 4^N - 1\}$ the process V^N satisfies

$$\left| V_{t_i}^N - B_{t_i} \right| \leq \frac{1}{2^N}, \quad \forall N \in \mathbb{N}$$

and so we always have

$$\|V^N - B\|_\infty \leq \frac{1}{2^N} + \sup_{s, t \in [0, T]: |s-t| \leq 2^{-N}} |B_s - B_t|.$$

The right hand side converges to 0 in \mathbb{P} probability. Since B remains a Brownian motion under each \mathbb{P}^N , $\mathbb{P}^N(E_N)$ converges to 1. Therefore, for all $\varepsilon > 0$ we have

$$\lim_{N \rightarrow \infty} \mathbb{P}^N \left(\|V^N - B\|_\infty \leq \varepsilon \right) = 1. \quad (39)$$

It is also obvious from the definition of D^N that

$$\lim_{N \rightarrow \infty} \mathbb{P}^N \left(d(D^N, D) \leq \varepsilon \right) = 1. \quad (40)$$

With this choice of V^N , we have the following consistency property in the limit.

Lemma 3.17 *For all bounded, uniformly continuous functions $h^0 : \mathcal{X} \times \mathcal{P}_1(\bar{\Omega}_1) \times \mathcal{X}_D \rightarrow \mathbb{R}$ and $h^1 : \mathcal{X}^* \times \mathcal{X} \times \Gamma \times \mathbb{R}^2 \rightarrow \mathbb{R}$, we have*

$$\bar{E} \left[h^0(B, \mathfrak{M}^\infty, D) h^1(X^T, W^\infty, \gamma^\infty, \mathcal{I}) \right] = \bar{E} \left[h^0(B, \mathfrak{M}^\infty, D) \int_{\bar{\Omega}_1} h^1(x, w, q, \iota) d\mathfrak{M}^\infty(x, w, q, \iota) \right].$$

Proof This equality holds at the discretization level by (38), namely

$$\begin{aligned} \mathbb{E}^N [h^0(V^N, \mathfrak{M}^N, D^N) h^1(X^T, W^N, \gamma^N, \mathcal{I})] \\ = \mathbb{E}^N \left[h^0(V^N, \mathfrak{M}^N, D^N) \int_{\bar{\Omega}_1} h^1(x, w, q, \iota) d\mathfrak{M}^N(x, w, q, \iota) \right]. \end{aligned}$$

It is obvious that for all $\varepsilon > 0$:

By (39) and (40), uniform continuity of the function h^0 allows us to swap (V^N, D^N) with (B, D) on both sides of the equality above without changing the limits, if they exist. Boundedness of h_0, h_1 and weak convergence enable us to take $N \rightarrow \infty$ and retain equality in the limit. \square

Lemma 3.17 says that \mathfrak{M}^∞ is a version of conditional distribution of $(X^T, W^\infty, \gamma^\infty, \mathcal{I})$ given $(B, \mathfrak{M}^\infty, D)$, which implies the consistency of marginals $\mathfrak{M}^{\infty, x}$ and $\mathfrak{M}^{\infty, \gamma}$. We want to carry the conditional joint laws because of the compatibility condition.

3.3.5 Optimality

For each $N \in \mathbb{N} \cup \{\infty\}$, define $\mathfrak{P}^N := (\mu^N, \vartheta^N) = (\mathfrak{M}^{N, x}, \mathfrak{M}^{N, \gamma})$ to be the first and third marginals of \mathfrak{M}^N . Define the state process corresponding to relaxed control any $\gamma \in \Gamma$ as in (16) but in the environment \mathbb{P}^N using W^N , namely

$$X^{N, \gamma} := \mathbb{1}_{\{t \geq T\}} K_0 / \mathcal{P} + \int_0^t \int_A a \gamma(ds, da) + \sigma (W_{t \vee T}^N - W_T^N), \quad t \in [0, T].$$

Then in particular, $X^{N,\gamma^N} = X^T$. Recalling (17) the objective function for relaxed controls, we now define for each N the objective under the environment \mathbb{P}^N :

$$J^N(\gamma) := \mathbb{E}^{\bar{P}^N} \left[g(X_T^{N,\gamma}, X_{\tau^*(\mu^N)}^{N,\gamma}, \tau^*(\mu^N)) + \int_{T \wedge T}^T \int_A f(s, X_s^{N,\gamma}, b_s^P, \tau^*(\mu^N), \langle \rho, \vartheta_s^N \rangle, a) \gamma(da, ds) \right].$$

Lemma 3.18 $\lim_{N \rightarrow \infty} J^N(\gamma^N) = J^\infty(\gamma^\infty)$.

Proof Since the convergence in Lemma 3.15 is weak, we need to uniformly approximate f, g by bounded functions. For $k \in \mathbb{N}$ and $x \in \mathbb{R}$, denote by \underline{x}_k the projection of x on $[-k, k]$. Define $f^k : [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times [0, T] \times \mathbb{R} \times A \rightarrow \mathbb{R}$ and $g : \Omega \times \mathbb{R} \times \mathbb{R} \times [0, T] \rightarrow \mathbb{R}$ by:

$$f^k(t, x, b, \eta, \varrho, a) = \kappa(a) + \phi \underline{x}_k^2 - \underline{x}_k (\underline{b}_k \mathbb{1}_{\{t < \eta\}} + \varrho \mathbb{1}_{\{t \geq \eta\}}), \quad (41)$$

$$g^k(x, y, \eta) = c \underline{x}_k^2 + \beta_\eta \underline{\eta}_k \underline{y}_k. \quad (42)$$

Recall that γ here is the bubble component defined in (4). Since we only care about γ at the burst time, we can equivalently take $\gamma_t = \int_0^t b_s^P ds$. Since the price impact functions κ and ϱ are continuous, compactness of A implies that for each $k \in \mathbb{N}$, there exists some $C_k > 0$ such that

$$\left| g^k(X_T^T, X_{\tau^*(\mu^N)}^T, \tau^*(\mu^N)) \right| + \left| \int_{T \wedge T}^T \int_A f^k(s, X_s^T, b_s^P, \tau^*(\mu^N), \langle \rho, \vartheta_s^N \rangle, a) \gamma(da, ds) \right| \leq C_k.$$

For $N \in \mathbb{N} \cup \{\infty\}$ and $k \in \mathbb{N}$, define the approximated objective $J^{N,k}$ on \mathbb{P}^N by

$$J^{N,k}(\gamma) := \mathbb{E}^{\mathbb{P}^N} \left[g^k(X_T^T, X_{\tau^*(\mu^N)}^T, \tau^*(\mu^N)) + \int_{T \wedge T}^T \int_A f^k(s, X_s^T, b_s^P, \tau^*(\mu^N), \langle \rho, \vartheta_s^N \rangle, a) \gamma(da, ds) \right].$$

Then weak convergence implies that $\lim_{N \rightarrow \infty} |J^{N,k}(\gamma^N) - J^{\infty,k}(\gamma^\infty)| = 0$. To shorten the notation, let $\tau^N := \tau^*(\mu^N)$. Using (28), we have

$$\begin{aligned} \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\sup_{t \in [0, T]} \left| X_t^T - \underline{X}_{t,k}^T \right|^2 \right] &= \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\sup_{t \in [0, T]} \left| X_t^T - k \right|^2 \mathbb{1}_{\{\|X_t^T\| > k\}} \right] \\ &\leq \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\|X^T\|_{\mathcal{A}^*}^2 \mathbb{1}_{\{\|X^T\|_{\mathcal{A}^*} > k\}} \right] \xrightarrow{k \rightarrow \infty} 0. \end{aligned}$$

Recall that P has the same law under \mathbb{P}^N for each N . Then similarly, Assumption (2) and Proposition A.1 together imply

$$\begin{aligned} \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\sup_{t \in [0, T]} \left| b_t^P - \underline{b}_{t,k}^P \right|^2 \right] &= \mathbb{E}^{\mathbb{P}^1} \left[\sup_{t \in [0, T]} \left| b_t^P - \underline{b}_{t,k}^P \right|^2 \right] \leq \mathbb{E}^{\mathbb{P}^1} \left[\|b^P\|_\infty \mathbb{1}_{\{\|b^P\|_\infty > k\}} \right] \xrightarrow{k \rightarrow \infty} 0. \\ \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\sup_{t \in [0, T]} \left| \gamma_t - \underline{\gamma}_{t,k} \right|^2 \right] &= \mathbb{E}^{\mathbb{P}^1} \left[\sup_{t \in [0, T]} \left| \gamma_t - \underline{\gamma}_{t,k} \right|^2 \right] \leq \mathbb{E}^{\mathbb{P}^1} \left[\|\gamma\|_\infty^2 \mathbb{1}_{\{\|\gamma\|_\infty > k\}} \right] \\ &\leq \mathbb{E}^{\mathbb{P}^1} \left[T \|b^P\|_\infty \mathbb{1}_{\{T \|b^P\|_\infty > k\}} \right] \xrightarrow{k \rightarrow \infty} 0. \end{aligned}$$

These uniform integrability properties, along with Assumption (2) and the separability condition in Remark 2.5 (C1), imply that there exists $C > 0$ such that

$$\begin{aligned} \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[|g - g^k|(X_T^{\mathcal{T}}, X_{\tau^N}^{\mathcal{T}}, \tau^N) \right] &\leq C \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\left| (X_T^{\mathcal{T}})^2 - \underline{(X_T^{\mathcal{T}})^2}_k \right| \right] \\ &+ C \sqrt{\sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} [|\gamma_{\tau^N} - \underline{\gamma_{\tau^N}}_k|^2]} + C \sqrt{\sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} [|X_{\tau^N}^{\mathcal{T}} - \underline{X_{\tau^N}^{\mathcal{T}}}_k|^2]} \xrightarrow{k \rightarrow \infty} 0, \end{aligned} \quad (43)$$

and also

$$\begin{aligned} \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\int_{T \wedge T}^T \int_A |f - f^k|(s, X_s^{\mathcal{T}}, b_s^P, \tau^N, \langle \rho, \vartheta_s^N \rangle, a) \gamma^N(da, ds) \right] \\ \leq C \int_0^T \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\left| (X_s^{\mathcal{T}})^2 - \underline{(X_s^{\mathcal{T}})^2}_k \right| \right] ds + C \int_0^T \sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} \left[\left| X_s^{\mathcal{T}} - \underline{X_s^{\mathcal{T}}}_k \right|^2 \right] ds \quad (44) \\ + C \int_0^T \sqrt{\sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} [|b_s^P - \underline{b_s^P}_k|^2]} + \sqrt{\sup_{N \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^N} [|X_s^{\mathcal{T}} - \underline{X_s^{\mathcal{T}}}_k|^2]} ds \xrightarrow{k \rightarrow \infty} 0. \end{aligned}$$

Therefore, for any fixed $k \in \mathbb{N}$,

$$\begin{aligned} \left| J^N(\gamma^N) - J^\infty(\gamma^\infty) \right| &\leq \left| J^N(\gamma^N) - J^{N,k}(\gamma^N) \right| + \left| J^{N,k}(\gamma^N) - J^{\infty,k}(\gamma^\infty) \right| \\ &\quad + \left| J^{\infty,k}(\gamma^\infty) - J^\infty(\gamma^\infty) \right|. \end{aligned}$$

Taking limit $N \rightarrow \infty$ gives

$$\lim_{N \rightarrow \infty} \left| J^N(\gamma^N) - J^\infty(\gamma^\infty) \right| \leq \sup_{N \in \mathbb{N}} \left| J^N(\gamma^N) - J^{N,k}(\gamma^N) \right| + \left| J^{\infty,k}(\gamma^\infty) - J^\infty(\gamma^\infty) \right|.$$

Taking $k \rightarrow \infty$ on the right hand side and using (43) and (44) give the result. \square

Let $\beta \in \mathbb{I}^\Gamma$ be another \mathbb{G} -admissible relaxed strategy. Following the proof of Lemma 3.18, we also have $J^N(\beta) \xrightarrow{N \rightarrow \infty} J(\beta)$. Remarks 3.13 and 3.14 together imply that $J^N(\gamma^N) \leq J^N(\beta)$ for each N . Taking $N \rightarrow \infty$ on both sides we have $J^\infty(\gamma^\infty) \leq J^\infty(\beta)$ for all $\beta \in \mathbb{I}^\Gamma$, so optimality is proved.

Now recall that for an A -valued control $\alpha \in \mathbb{A}^*$, we denote by $\gamma(\alpha)$ its corresponding strict control in the space of relaxed controls, where each time marginal is the Dirac measure at α_t . Using the optimality lemma above, we can in fact show that γ^∞ must be a strict control.

Lemma 3.19 *There is a version of γ^∞ that is $\mathbb{F}^{B, X^{\mathcal{T}}, W^{\infty}, \mathcal{I}, D}$ -progressively measurable that is a strict control taking the form $\gamma^\infty = \gamma(\hat{\alpha}^\infty)$ for some $\hat{\alpha}^\infty \in \mathbb{A}^*$.*

Proof Define $\alpha_t^\infty := \int_A a \gamma_t^\infty(da)$ for $t \in [0, T]$ and $\tilde{\gamma} := \gamma(\alpha^\infty)$. Then $\alpha^\infty \in \mathbb{A}^*$ and $\tilde{\gamma} \in \mathbb{I}^\Gamma$ is a strict control. It is obvious that $\tilde{\gamma}$ and γ^∞ both give rise to the same state process $X^{\mathcal{T}}$ according to (16). Using strict convexity of f in a and Jensen's inequality, we have

$$\begin{aligned} J^\infty(\tilde{\gamma}) &= \bar{\mathbb{E}} \left[g(X_T^{\mathcal{T}}, X_{\tau^*(\mu^\infty)}^{\mathcal{T}}, \tau^*(\mu^\infty)) + \int_{T \wedge T}^T \int_A f(s, X_s^{\mathcal{T}}, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, a) \tilde{\gamma}(da, ds) \right] \\ &= \bar{\mathbb{E}} \left[g(X_T^{\mathcal{T}}, X_{\tau^*(\mu^\infty)}^{\mathcal{T}}, \tau^*(\mu^\infty)) + \int_{T \wedge T}^T f(s, X_s^{\mathcal{T}}, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, \alpha_s^\infty) ds \right] \\ &\leq \bar{\mathbb{E}} \left[g(X_T^{\mathcal{T}}, X_{\tau^*(\mu^\infty)}^{\mathcal{T}}, \tau^*(\mu^\infty)) + \int_{T \wedge T}^T \int_A f(s, X_s^{\mathcal{T}}, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, a) \gamma^\infty(da, ds) \right] \\ &= J^\infty(\gamma^\infty). \end{aligned}$$

The inequality is strict (which contradicts with optimality of γ^∞) unless $\gamma^\infty = \gamma(\alpha^\infty)$. Lebesgue differentiation theorem allows us to define $\hat{\alpha}^\infty \in \mathbb{A}^*$ by

$$\hat{\alpha}_t^\infty = \begin{cases} \lim_{n \rightarrow \infty} n \int_{(t-1/n)^+}^t \alpha_s^\infty ds & \text{if the limit exists} \\ 0 & \text{otherwise.} \end{cases}$$

Then $\bar{\mathbb{P}} \otimes dt$ almost surely, $\hat{\alpha}_t^\infty = \alpha_t^\infty$. Note that $\hat{\alpha}^\infty$ shares the same measurability with $\int_0^\cdot \alpha_s^\infty ds$, which by (12) is $\mathbb{F}^{X^T, W^\infty, B, \mathcal{I}, D}$ measurable. \square

A consequence of the lemma above is that we can drop either γ^∞ or X^T from the definition of \mathbb{G} and simply consider $\mathbb{G} = \mathbb{F}^{B, \mathfrak{M}^\infty, X^T, W^\infty, \mathcal{I}, D} = \mathbb{F}^{B, \mathfrak{M}^\infty, \gamma^\infty, W^\infty, \mathcal{I}, D}$. Moreover, both optimality and consistency still hold for $\gamma^\infty = \gamma(\hat{\alpha}^\infty)$. In fact, this is the case for every limit point of the sequence in Lemma 3.15.

3.3.6 Compatibility

Following Definition 3.3, we need to show that $\bar{\mathbb{F}} := \mathbb{F}^{\mathcal{I}, \mathfrak{M}^\infty, W^\infty, B, D}$ is immersed in $\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}$ defined above. We need to keep in mind that while \mathcal{I} is \mathcal{G}_0 -measurable, it is not $\mathcal{F}_0^{\mathcal{X}^T}$ -measurable due to random entry, which is why we need to treat \mathcal{I} separately.

Lemma 3.20 *The filtration \mathbb{G} is compatible with $(\mathcal{I}, \mathfrak{M}^\infty, W^\infty, B, D)$.*

Proof By Proposition 3.2, it suffices to show that for all $t \in [0, T]$, $\mathcal{F}_t^{\mathcal{X}^T}$ is conditionally independent from $\bar{\mathcal{F}}_T = \mathcal{F}_T^{\mathcal{I}, \mathfrak{M}^\infty, W^\infty, B, D}$ given $\bar{\mathcal{F}}_t = \mathcal{F}_t^{\mathcal{I}, \mathfrak{M}^\infty, W^\infty, B, D}$. We follow the proof of [20, Lemma 3.7].

Lemma 3.17 implies that W^∞ is a $\bar{\mathbb{P}}$ -Brownian motion independent from $(B, \mathfrak{M}^\infty, D)$. Fix $t \in [0, T]$. Consider three bounded functions ϕ_t^m, ϕ_{t+}^w and ϕ_t^1 where $\phi_t^m : \mathcal{P}_1(\bar{\Omega}_1) \rightarrow \mathbb{R}$ is $\mathcal{F}_t^{\mathfrak{M}^\infty}$ measurable, $\phi_{t+}^w : \mathcal{X} \rightarrow \mathbb{R}$ is $\sigma(W_s - W_t : s \in [t, T])$ measurable, and $\phi_t^1 : \bar{\Omega}_1 \rightarrow \mathbb{R}$ is $\mathcal{F}_t^{\mathcal{I}, X^T, W^\infty}$ measurable. By Lemma 3.17 and property of Brownian motion we have

$$\begin{aligned} \bar{E} \left[\phi_t^m(\mathfrak{M}^\infty) \int_{\bar{\Omega}_1} \phi_t^1 d\mathfrak{M}^\infty \right] \bar{E} [\phi_{t+}^w(W^\infty)] &= \bar{E} \left[\phi_t^m(\mathfrak{M}^\infty) \phi_t^1(X^T, W^\infty, \gamma^\infty, \mathcal{I}) \phi_{t+}^w(W^\infty) \right] \\ &= \bar{E} \left[\phi_t^m(\mathfrak{M}^\infty) \int_{\bar{\Omega}_1} \phi_{t+}^w(w) \phi_t^1(x, w, q, \iota) d\mathfrak{M}^\infty(x, w, q, \iota) \right]. \end{aligned}$$

Since this holds for all ϕ_t^m , $\bar{\mathbb{P}}$ almost surely we have

$$\bar{E} [\phi_{t+}^w(W^\infty)] \mathbb{E}^{\mathfrak{M}^\infty} [\phi_t^1(X^T, \gamma^\infty, \mathcal{I}, W^\infty)] = \mathbb{E}^{\mathfrak{M}^\infty} [\phi_{t+}^w(W^\infty) \phi_t^1(X^T, \gamma^\infty, \mathcal{I}, W^\infty)] \quad (45)$$

where by $\mathbb{E}^{\mathfrak{M}^\infty} [\phi(X^T, \gamma^\infty, \mathcal{I}, W^\infty)]$ we mean the integral $\int_{\bar{\Omega}_1} \phi d\mathfrak{M}^\infty$ for $\phi : \bar{\Omega}_1 \rightarrow \mathbb{R}$. Note that this expectation is $\bar{\mathcal{F}}_t$ -measurable if ϕ is $\mathcal{F}_t^{\mathcal{I}, X^T, W^\infty}$ -measurable.

Additionally, consider bounded functions $\phi^l, \phi_t^x, \varphi_t, \varphi_T$ where $\phi^l : \mathbb{R}^2 \rightarrow \mathbb{R}$ is Borel measurable, $\phi_t^x : \mathcal{X}^* \rightarrow \mathbb{R}$ is $\mathcal{F}_t^{\mathcal{X}^T}$ measurable, φ_t and φ_T are functions from $\mathcal{X} \times \mathcal{P}_1(\bar{\Omega}_1) \times \mathcal{X}_D$

to \mathbb{R} that are $\mathcal{F}_t^{B, \mathfrak{M}^\infty, D}$ and $\mathcal{F}_T^{B, \mathfrak{M}^\infty, D}$ measurable, respectively. Using (45) and Lemma 3.17, we have

$$\begin{aligned}
& \bar{E}[\phi_t^x(X^T)\varphi_T(B, \mathfrak{M}^\infty, D)\phi^i(\mathcal{I})\phi_{t+}^w(W^\infty)\phi_t^w(W^\infty)\varphi_t(B, \mathfrak{M}^\infty, D)] \\
&= \bar{E}\left[\mathbb{E}^{\mathfrak{M}^\infty}[\phi_t^x(X^T)\phi^i(\mathcal{I})\phi_{t+}^w(W^\infty)\phi_t^w(W^\infty)](\varphi_T \cdot \varphi_t)(B, \mathfrak{M}^\infty, D)\right] \\
&= \bar{E}\left[\mathbb{E}^{\mathfrak{M}^\infty}[\phi_t^x(X^T)\phi^i(\mathcal{I})\phi_t^w(W^\infty)](\varphi_T \cdot \varphi_t)(B, \mathfrak{M}^\infty, D)\right]\bar{E}[\phi_{t+}^w(W^\infty)] \\
&= \bar{E}\left[\bar{\mathbb{E}}\left[\phi_t^x(X^T)\phi^i(\mathcal{I})\phi_t^w(W^\infty) \mid \bar{\mathcal{F}}_t\right](\varphi_T \cdot \varphi_t)(B, \mathfrak{M}^\infty, D)\right]\bar{E}[\phi_{t+}^w(W^\infty)] \\
&= \bar{E}\left[\bar{\mathbb{E}}\left[\phi_t^x(X^T)\phi^i(\mathcal{I})\phi_t^w(W^\infty) \mid \bar{\mathcal{F}}_t\right]\varphi_t(B, \mathfrak{M}^\infty, D)\bar{E}\left[\varphi_T(B, \mathfrak{M}^\infty, D) \mid \bar{\mathcal{F}}_t\right]\right]\bar{E}[\phi_{t+}^w(W^\infty)] \\
&= \bar{E}\left[\bar{E}[\phi_t^x(X^T) \mid \bar{\mathcal{F}}_t]\bar{E}\left[\varphi_T(B, \mathfrak{M}^\infty, D) \mid \bar{\mathcal{F}}_t\right]\varphi_t(B, \mathfrak{M}^\infty, D)\phi^i(\mathcal{I})\phi_{t+}^w(W^\infty)\phi_t^w(W^\infty)\right], \tag{46}
\end{aligned}$$

where the last equality follows from the independence of $\phi_{t+}^w(W^\infty)$ and $\bar{\mathcal{F}}_t$. Since ϕ^i and ϕ_t^w are arbitrary, we can replace them with bounded $\phi^i \cdot \varphi^i$ and $\phi_t^w \cdot \varphi_t^w$, each with the same corresponding measurability requirements. Then by definition of conditional expectation we have

$$\begin{aligned}
& \bar{E}[\phi_t^x(X^T)\varphi_T(B, \mathfrak{M}^\infty, D)\phi^i(\mathcal{I})\phi_{t+}^w(W^\infty)\phi_t^w(W^\infty) \mid \bar{\mathcal{F}}_t] \\
&= \bar{E}\left[\phi_t^x(X^T) \mid \bar{\mathcal{F}}_t\right]\bar{E}\left[\varphi_T(B, \mathfrak{M}^\infty, D)\phi^i(\mathcal{I})\phi_{t+}^w(W^\infty)\phi_t^w(W^\infty) \mid \bar{\mathcal{F}}_t\right].
\end{aligned}$$

We conclude by noting that $\mathcal{F}_T^{W^\infty}$ is generated by $\phi_{t+}^w(W^\infty)\phi_t^w(W^\infty)$ with arbitrary ϕ_{t+}^w and ϕ_t^w . \square

We have then finished the proof of Theorem 3.4.

4 Strong control and Proof of Theorem 2.8

Recall from Lemma 3.19 that the weak control found in the previous section is in fact a strict control $\gamma(\alpha^\infty)$, and α^∞ is $\mathbb{F}^{B, X^T, W^\infty, \mathfrak{M}^\infty, \mathcal{I}, D}$ progressive. In order to obtain an equilibrium with strong control, we will show that α^∞ is $\mathbb{F}^{B, W^\infty, \mathfrak{P}^\infty, \mathcal{I}, D}$ measurable after bringing the lifted environment \mathfrak{M}^∞ back to the “original” environment $\mathfrak{P}^\infty = (\mu^\infty, \vartheta^\infty) = (\mathfrak{M}^{\infty, x}, \mathfrak{M}^{\infty, y})$. We now begin the proof of Theorem 2.8.

Proof of Theorem 2.8 We split the proof into three steps. First we show that the original environment \mathfrak{P}^∞ carries enough information for the equilibrium. Then we adopt the proof of [19, Proposition 4.4] to recover an A -valued control. Lastly, we end the proof by arguing that τ remains a totally inaccessible stopping time.

Step 1: Back to original environment

The reason for lifting the environment is solely for the proof of the compatibility lemma 3.20, in particular the first and third equality in (46). Recall from (38) that we took \mathfrak{M}^N to be the joint conditional law of $(X^T, W^N, \gamma^N, \mathcal{I})$ given (B, D) under \mathbb{P}^N . We did this to ease the notation in the consistency and compatibility lemmas. Notice that we did not need the full fledged joint law in deriving (46), but only the product of the marginals. This implies

that for fixed $N \in \mathbb{N}$, we could alternatively define for each $(\beta, \eta) \in \mathcal{X} \times \mathbb{R}_+$:

$$\tilde{\mathfrak{M}}^N(\beta, \eta) := \sum_{k=1}^{|\mathcal{V}_N|} \mathbb{1}_{\{(\beta, D(\eta)) \in A_k\}} \mathcal{L}^{\alpha^N, k}(X^T) \otimes \mathcal{L}^{\alpha^N, k}(W^N) \otimes \mathcal{L}^{\alpha^N, k}(\gamma^N) \otimes \mathcal{L}^{\alpha^N, k}(\mathcal{I}).$$

This version still carries the necessary inputs \mathfrak{P}^N to the BSDE (26) as its first and third marginals, and tightness of $\mathbb{P}^N \circ (\tilde{\mathfrak{M}}^N)^{-1}$ follows immediately from that of $\mathbb{P}^N \circ (\mathfrak{M}^N)^{-1}$ in Lemma 3.15. Then we take a limit point $\bar{\mathbb{P}} := \mathbb{P}^\infty \circ (B, \tau, \tilde{\mathfrak{M}}^\infty, X^T, W^\infty, \gamma^\infty, \mathcal{I})^{-1}$ and follow the same argument in Remark 3.16 to work on the canonical space. In particular, $\tilde{\mathfrak{M}}^\infty$ is the canonical process on $\mathcal{P}_1(\bar{\Omega}_1)$. Following the argument in Lemma 3.17, the fixed point property for each $N \in \mathbb{N}$ now leads to a weaker consistency in the limit. Namely, for all bounded, uniformly continuous, \mathbb{R} -valued functions $h^0, h_x^1, h_w^1, h_\gamma^1, h_\iota^1$ with respective domains $\mathcal{X} \times \mathcal{P}_1(\bar{\Omega}_1) \times \mathcal{X}_D, \mathcal{X}^*, \mathcal{X}, \Gamma, \mathbb{R}^2$, we have

$$\begin{aligned} \bar{E} & \left[h^0(B, \tilde{\mathfrak{M}}^\infty, D) \int_{\mathcal{X}^*} h_x^1(x) d\tilde{\mathfrak{M}}^\infty(x) \int_{\mathcal{X}} h_w^1(w) d\tilde{\mathfrak{M}}^\infty(w) \int_{\Gamma} h_\gamma^1(q) d\tilde{\mathfrak{M}}^\infty(q) \int_{\mathbb{R}^2} h_\iota^1(\iota) d\tilde{\mathfrak{M}}^\infty(\iota) \right] \\ & = \bar{E}[h^0(B, \tilde{\mathfrak{M}}^\infty, D) h_x^1(X^T) h_w^1(W^\infty) h_\gamma^1(\gamma^\infty) h_\iota^1(\mathcal{I})] \\ & = \bar{E} \left[h^0(B, \tilde{\mathfrak{M}}^\infty, D) \int_{\bar{\Omega}_1} h_x^1(x) h_w^1(w) h_\gamma^1(q) h_\iota^1(\iota) d\tilde{\mathfrak{M}}^\infty(x, w, q, \iota) \right], \end{aligned} \quad (47)$$

which results from taking $N \rightarrow \infty$ of the following equalities by construction

$$\begin{aligned} \bar{E} & \left[h^0(V^N, \tilde{\mathfrak{M}}^N, D^N) \int_{\mathcal{X}^*} h_x^1(x) d\tilde{\mathfrak{M}}^N(x) \int_{\mathcal{X}} h_w^1(w) d\tilde{\mathfrak{M}}^N(w) \int_{\Gamma} h_\gamma^1(q) d\tilde{\mathfrak{M}}^N(q) \int_{\mathbb{R}^2} h_\iota^1(\iota) d\tilde{\mathfrak{M}}^N(\iota) \right] \\ & = \bar{E}[h^0(V^N, \tilde{\mathfrak{M}}^N, D^N) h_x^1(X^T) h_w^1(W^N) h_\gamma^1(\gamma^N) h_\iota^1(\mathcal{I})] \\ & = \bar{E} \left[h^0(V^N, \tilde{\mathfrak{M}}^N, D^N) \int_{\bar{\Omega}_1} h_x^1(x) h_w^1(w) h_\gamma^1(q) h_\iota^1(\iota) d\tilde{\mathfrak{M}}^N(x, w, q, \iota) \right]. \end{aligned}$$

This is also sufficient for the consistency requirement in Definition 3.3. Similarly, with $\tilde{\mathfrak{M}}^\infty$ the equality (45) holds only for ϕ_t^1 taking the form of a product, separable in each coordinate. This weaker property, however, is sufficient for (46) and hence the compatibility requirement. Since the optimality property only depends on the marginals and thus is not influenced, we can replace \mathfrak{M}^∞ with $\tilde{\mathfrak{M}}^\infty$ in the final filtration $\mathbb{G} = \mathbb{F}^{B, X^T, W^\infty, \tilde{\mathfrak{M}}^\infty, \mathcal{I}, D}$.

Note from (47) that $\bar{\mathbb{P}}$ almost surely, $\tilde{\mathfrak{M}}^\infty$ is a product measure of its four marginals by uniqueness of measures on the product space. More importantly, its second and fourth are almost surely the Wiener measure and $(\nu_K \otimes \nu_p)$, respectively, since for each $N \in \mathbb{N}$, (W^N, \mathcal{I}, B, D) are mutually independent under \mathbb{P}^N . Being complete, the filtration $\mathbb{F}^{\mathfrak{P}^\infty}$ in the original environment coincides with $\mathbb{F}^{\tilde{\mathfrak{M}}^\infty}$ from the lifted environment. Therefore, we can equivalently take $\mathbb{G} = \mathbb{F}^{B, X^T, W^\infty, \mathfrak{P}^\infty, \mathcal{I}, D}$, and the compatibility condition reads that $\bar{\mathbb{F}} := \mathbb{F}^{B, W^\infty, \mathfrak{P}^\infty, \mathcal{I}, D}$ is immersed in \mathbb{G} .

Step 2: Strong control via optional projection

To further strengthen the measurability property of γ from \mathbb{G} to $\bar{\mathbb{F}}$, we follow the proof of [20, Proposition 4.4]. Recall that the state equation (12) with W^∞ as the Brownian motion is satisfied by X^T and $\hat{\alpha}^\infty$. By optional projection we can find $\bar{\mathbb{F}}$ -optional processes \bar{X}^T and $\bar{\alpha}$ such that for any finite $\bar{\mathbb{F}}$ -stopping time ρ :

$$\bar{X}_\rho^T := \bar{E}[X_\rho^T | \bar{\mathcal{F}}_\rho], \quad \bar{\alpha}_\rho := \bar{E}[\hat{\alpha}_\rho^\infty | \bar{\mathcal{F}}_\rho], \quad a.s. \quad (48)$$

Since $\bar{\mathbb{F}}$ is immersed in \mathbb{G} , Proposition 3.2 implies that for each $0 \leq s \leq t \leq T$,

$$\bar{X}_s^T = \bar{E}[X_s^T | \bar{\mathcal{F}}_t], \quad \bar{\alpha}_s = \bar{E}[\hat{\alpha}_s^\infty | \bar{\mathcal{F}}_t], \quad a.s. \quad (49)$$

Using Fubini's theorem for conditional expectation along with (49) on (12), we can replace \bar{X}^T by a modification such that $\bar{\mathbb{P}}$ almost surely

$$\bar{X}_t^T = \mathbb{1}_{\{t \geq T\}} K_0 / \mathcal{P} + \int_0^t \bar{\alpha}_s ds + \sigma(W_{t \vee T}^\infty - W_T^\infty), \quad t \in [0, T].$$

Notice that given $\mathfrak{P}^\infty = (\tilde{\mathfrak{M}}^{\infty,x}, \tilde{\mathfrak{M}}^{\infty,y}) \equiv (\mu^\infty, \vartheta^\infty)$, the bubble burst time $\tau^*(\mu^\infty)$ is a $\bar{\mathbb{F}}$ -stopping time, and the bubble component $\gamma_{\tau^*(\mu^\infty)}$ is $\bar{\mathcal{F}}_{\tau^*(\mu^\infty)}$ measurable. Recall also from Remark 2.5 (C1) that the running cost f depends on $\tau^*(\mu^\infty)$ only through D , which is $\bar{\mathbb{F}}$ -adapted. Then by conditional Jensen's inequality, Remark 2.5 (C3) and (48),

$$\begin{aligned} J^{\mu^\infty, \vartheta^\infty}(\hat{\alpha}^\infty) &= \bar{E} \left[\int_{T \wedge T}^T f(s, X_s^T, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, \hat{\alpha}_s) ds + c(X_T^T)^2 + \beta_{\tau^*(\mu^\infty)} \gamma_{\tau^*(\mu^\infty)} X_{\tau^*(\mu^\infty)}^T \right] \\ &= \bar{E} \left[\int_{T \wedge T}^T \bar{E} \left[f(s, X_s^T, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, \hat{\alpha}_s) | \bar{\mathcal{F}}_s \right] ds \right] \\ &\quad + c \bar{E} \left[\bar{E} \left[|X_T^T|^2 | \bar{\mathcal{F}}_T \right] \right] + \bar{E} \left[\beta_{\tau^*(\mu^\infty)} \gamma_{\tau^*(\mu^\infty)} \bar{E} \left[X_{\tau^*(\mu^\infty)}^T | \bar{\mathcal{F}}_{\tau^*(\mu^\infty)} \right] \right] \\ &\geq \bar{E} \left[\int_{T \wedge T}^T f(s, \bar{X}_s^T, b_s^P, \tau^*(\mu^\infty), \langle \rho, \vartheta_s^\infty \rangle, \bar{\alpha}_s) ds + c(\bar{X}_T^T)^2 + \beta_{\tau^*(\mu^\infty)} \gamma_{\tau^*(\mu^\infty)} \bar{X}_{\tau^*(\mu^\infty)}^T \right] \\ &= J^{\mu^\infty, \vartheta^\infty}(\bar{\alpha}). \end{aligned}$$

By strict convexity of f in (x, a) , the inequality is strict unless $\hat{\alpha}^\infty$ and X^T are both already $\bar{\mathbb{F}}$ adapted. Strict inequality would lead to a contradiction to optimality of $\hat{\alpha}^\infty$ among \mathbb{G} -progressive controls, since $\bar{\alpha}$ is $\bar{\mathbb{F}}$ -optional, hence also \mathbb{G} -progressive.

Step 3: Exogenous burst time as totally inaccessible stopping time

The section above implies that we can take $\mathbb{G} = \bar{\mathbb{F}} = \bar{\mathbb{F}}^{\mathcal{I}, B, W^\infty, D, \mathfrak{P}^\infty}$ to begin with. This concludes the proof for the existence statement of Theorem 2.8. It remains to show that the exogenous burst time τ is a $\bar{\mathbb{F}}^{\mathcal{I}, B, W^\infty, D, \mathfrak{P}^\infty}$ -totally inaccessible stopping time. In light of Remark 3.5 and Assumption (E2), it suffices to remark that τ is independent from $(\mathcal{I}, B, W^\infty, \mathfrak{P}^\infty)$, which follows from the independence between τ and $(\mathcal{I}, B, W^N, \mathfrak{P}^N)$ for each N . Therefore, we have found an equilibrium with strong control as defined in Definition 2.7, concluding the proof of Theorem 2.8. \square

5 Concluding remarks

In this paper we proposed a more realistic extension of the bubble riding game introduced in [62]. In contrast to [62] where agents were assumed to enter the game at independent and identically distributed times on an awareness window $[0, \eta]$, here we allow players to enter the game when the price trajectory of the bubble asset reaches a given threshold. We also allow the initial inventory to depend on the initial (cash) investment and the price level at time of entry. Due to these improvements on the model, the resulting MFG in the $N \rightarrow \infty$ limit is one with common noise in addition to non-standard features such as random entry times, interaction through the controls and possible jump of the state processes. Because the coefficients of the game do not satisfy the usual monotonicity conditions assumed in common noise MFG theory, we have to settle for existence of equilibria in a suitable weak form (see Definition 2.7). In short, the weaker, more realistic model assumptions made in the present

paper result in weak, abstract equilibrium strategies whereas the stronger model assumptions made in [62] result in stronger equilibrium strategies that can be numerically simulated thus providing interesting economical insights.

Appendix A Two auxiliary results

For a càdlàg process Y , denote by $M_t^Y = \sup_{0 \leq s \leq t} Y_s$ its running maximum. Recall from (1) and the price dynamics of the N-player game that the bubble trend function b naturally depends on $F_p(M_t^P)$, which is not Lipschitz in M^P . In general, the dynamics of asset price in the bubble phase is not well-posed. However, as the bubble is fueled by players' entry, b should be increasing in $F_p(M_t^P)$, hence also increasing in M_t^P at each time $t \in [0, T]$ since F_p is a CDF. This monotonicity property of the path-dependent SDE (2) restores unique solvability.

Proposition A.1 *The following path-dependent SDE*

$$X_t = x + \int_0^t \tilde{b}(s, M_s^X, X_s) ds + \sigma_0 B_t \quad (50)$$

has a unique strong solution satisfying $\mathbb{E}[\|X\|_\infty^2] < \infty$ if for each fixed $t \in [0, T]$:

(1) There exists $C > 0$ such that for all $\mathbf{x} \in C([0, T]; \mathbb{R})$:

$$|\tilde{b}(t, M_t^x, \mathbf{x}_t)| \leq C (1 + M_t^{|x|}).$$

(2) $\tilde{b}(t, \cdot, \cdot)$ is increasing (not necessarily strictly) in each argument.

Proof We adapt the proof of [11, Theorem 4.1]. The first condition guarantees a weak solution satisfying the integrability condition that is unique in law (see [42, Proposition 5.3.6 and Remark 5.3.8]). By the well-known result of Yamada and Watanabe [64], we only need to show pathwise uniqueness. Suppose X and Y are two solutions on the same probability space with respect to the same Brownian motion B . Observing that $X - Y$ is absolutely continuous, by Tanaka's formula we get

$$\begin{aligned} X_t \vee Y_t &= Y_t + (X_t - Y_t)^+ = Y_t + \int_0^t \mathbb{1}_{\{X_s > Y_s\}} d(X_s - Y_s) \\ &= x + \sigma_0 B_t + \int_0^t \mathbb{1}_{\{X_s > Y_s\}} \tilde{b}(s, M_s^X, X_s) ds + \int_0^t \mathbb{1}_{\{X_s \leq Y_s\}} \tilde{b}(s, M_s^Y, Y_s) ds. \quad (51) \\ Y_t \vee X_t &= x + \sigma_0 B_t + \int_0^t \mathbb{1}_{\{Y_s > X_s\}} \tilde{b}(s, M_s^Y, Y_s) ds + \int_0^t \mathbb{1}_{\{Y_s \leq X_s\}} \tilde{b}(s, M_s^X, X_s) ds. \end{aligned}$$

We can equate the above expressions for all t , implying that for almost every t we have

$$\mathbb{1}_{\{X_t = Y_t\}} (\tilde{b}(t, M_t^Y, Y_t) - \tilde{b}(t, M_t^X, X_t)) = 0. \quad (52)$$

We now show that if $X_s > Y_s$, then $M_s^X \geq M_s^Y$. Define

$$s_0 := \sup\{u \in [0, s] : X_u = Y_u\}.$$

The case is trivial if $s_0 = 0$.

On the event $\{s_0 > 0\}$, continuity of X and Y implies that $X_t > Y_t$ for all $t \in (s_0, s)$. Suppose $M_s^X < M_s^Y$, then there must exist $s^* \in [0, s_0)$ where $Y_{s^*} = M_s^Y > M_s^X \geq X_{s^*}$. Then define

$$s_1 := \inf\{u \in [s^*, s_0] : X_u = Y_u\}.$$

By continuity again, $Y_t > X_t$ for all $t \in [s^*, s_1]$. By definition of s^* , we must also have $M_t^Y > M_t^X$ for all $t \in [s^*, s_1]$. Monotonicity of \tilde{b} leads to a contradiction

$$0 > X_{s^*} - Y_{s^*} = \int_{s^*}^{s_1} \tilde{b}(t, M_t^Y, Y_t) - \tilde{b}(t, M_t^X, X_t) dt \geq 0.$$

Therefore, $M_s^X \geq M_s^Y$ and in particular, $M_s^{X \vee Y} = M_s^X$. We can then rewrite (51) as

$$\begin{aligned} X_t \vee Y_t &= x + \sigma_0 B_t + \int_0^t \tilde{b}(s, M_s^{X \vee Y}, X_s \vee Y_s) ds \\ &\quad + \int_0^t \mathbb{1}_{\{X_s=Y_s\}} \left(\tilde{b}(s, M_s^Y, Y_s) - \tilde{b}(s, M_s^{X \vee Y}, X_s \vee Y_s) \right) ds \\ &= x + \sigma_0 B_t + \int_0^t \tilde{b}(s, M_s^{X \vee Y}, X_s \vee Y_s) ds \\ &\quad + \int_0^t \mathbb{1}_{\{\{X_s=Y_s\} \cap \{M_s^X > M_s^Y\}\}} \left(\tilde{b}(s, M_s^Y, Y_s) - \tilde{b}(s, M_s^{X \vee Y}, X_s \vee Y_s) \right) ds. \end{aligned}$$

where the last line vanishes by (52). Therefore, $X \vee Y$ also satisfies (50). Similarly, one can show $X \wedge Y$ is also a solution. Then by uniqueness of law, we have $\mathbb{E}[|X - Y|] = \mathbb{E}[X \vee Y - X \wedge Y] = 0$ which leads to pathwise-uniqueness and completes the proof. The integrability property easily follows from Grönwall's inequality. \square

The following measure theoretic result is probably well known. We give a proof since we could not find a directly citable reference.

Lemma A.2 *Let (S, Σ, μ) be a complete measurable space. A function $f : S \times \mathbb{R} \rightarrow \mathbb{R}$ is jointly measurable if for all $x \in \mathbb{R}$:*

- (1) *$f(\cdot, x)$ is measurable.*
- (2) *$f(\cdot, x_n)$ converges to $f(\cdot, x)$ in μ -measure for any increasing sequence $x_n \uparrow x$.*

Proof First let $E \subseteq \mathbb{R}$ be any closed set and let $X = \{x_m\}_{m \geq 1}$ be a countable, dense subset of \mathbb{R} . For $\varepsilon > 0$, denote by $\mathcal{O}_\varepsilon(E)$ the open set $\{x \in \mathbb{R} : \inf_{e \in E} |x - e| < \varepsilon\}$. We claim that for μ -almost every $s \in S$ and any $x \in \mathbb{R}$, $f(s, x) \in E$ if and only if for each $n \in \mathbb{N}$, there is $x_m \in X \cap (x - \frac{1}{n}, x]$ such that $f(s, x_m) \in \mathcal{O}_{\frac{1}{n}}(E)$. Note that we can always approximate any $x \in \mathbb{R}$ by an increasing sequence $\{x_{m_k}\}_{k \geq 1}$ with elements in X such that the functions $f(\cdot, x_{m_k})$ converge μ -almost everywhere to $f(\cdot, x)$. The claim follows almost immediately. Denoting by f^{-1} the preimage of f , joint measurability is proved by writing

$$f^{-1}(E) = \bigcap_{n=1}^{\infty} \bigcup_{m=1}^{\infty} \left\{ s \in S : f(s, x_m) \in \mathcal{O}_{\frac{1}{n}}(E) \right\} \times \left[x_m, x_m + \frac{1}{n} \right).$$

\square

References

1. Abreu, D., Brunnermeier, M.K.: Bubbles and crashes. *Econometrica* **71**(1), 173–204 (2003)
2. Ahuja, Saran: Wellposedness of mean field games with common noise under a weak monotonicity condition. *SIAM J. Control Optim.* **54**(1), 30–48 (2016)
3. Akyildirim, E., Corbet, S., Cumming, D., Lucey, B., Sensoy, A.: Riding the wave of crypto-exuberance: the potential misusage of corporate blockchain announcements. *Technol. Forecast. Soc. Change* **159**, 120191 (2020)
4. Aliprantis, C.D., Border, K.C.: *Infinite Dimensional Analysis: A Hitchhiker's Guide*, 3rd edn. Springer, Cham (2006)
5. Allen, F., Morris, S., Postlewaite, A.: Finite bubbles with short sale constraints and asymmetric information. *J. Econ. Theory* **61**(2), 206–229 (1993)
6. Almgren, R., Chriss, N.: Optimal execution of portfolio transactions. *J. Risk* **3**, 5–39 (2001)
7. Awaya, Y., Iwasaki, K., Watanabe, M.: Rational bubbles and middlemen. *Theor. Econ.* **17**(4), 1559–1587 (2022)
8. Azéma, J., Jeulin, T., Knight, F.B. Marc Y.: Le théorème d'arrêt en une fin d'ensemble prévisible. In: *Séminaire de Probabilités XXVII*, pp. 133–158. Springer, Berlin (1993)
9. Bayraktar, E., Cecchin, A., Cohen, A., Delarue, F.: Finite state mean field games with wright–fisher common noise. *Journal de Mathématiques Pures et Appliquées* **147**, 98–162 (2021)
10. Beaver, W.H.: The information content of annual earnings announcements. *J. Account. Res.* **6**, 67–92 (1968)
11. Belfadli, R., Hamadéne, S., Ouknine, Y.: On one-dimensional stochastic differential equations involving the maximum process. *Stoch. Dyn.* **9**(02), 277–292 (2009)
12. Blanchet-Scalliet, C., Jeanblanc, M.: Hazard rate for credit risk and hedging defaultable contingent claims. *Finance Stoch.* **8**(1), 145–159 (2004)
13. Bouchaud, J.-P., Farmer, J.D., Lillo, F.: Chapter 2: How markets slowly digest changes in supply and demand? In: Thorsten, H., Schenk-Hoppé, K.R. (eds.) *Handbook of Financial Markets: Dynamics and Evolution, Handbooks in Finance*, pp. 57–160. North-Holland (2009)
14. Burzoni, M., Campi, L.: Mean field games with absorption and common noise with a model of bank run. *Stoch. Process. Appl.* **164**, 206–241 (2023)
15. Cardaliaguet, P.: Notes on mean field games. In: Technical Report, Université de Paris - Dauphine (2012)
16. Cardaliaguet, P., Delarue, F., Lasry, J.-M., Lions, P.-L.: *The Master Equation and the Convergence Problem in Mean Field Games*: (AMS-201). Princeton University Press, (2019). ISBN 9780691193717
17. Carmona, R., Delarue, F.: *Probabilistic Theory of Mean Field Games with Applications*, vol. I. Springer, 1st edition (2018a)
18. Carmona, R., Delarue, F.: *Probabilistic Theory of Mean Field Games with Applications*, vol. II. Springer, 1st edition (2018b)
19. Carmona, R., Lacker, D.: A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.* **25**(3), 1189–1231 (2015)
20. Carmona, R., Delarue, F., Lacker, D.: Mean field games with common noise. *Ann. Probab.* **44**(6), 3740–3803 (2016)
21. Cartea, A., Donnelly, R.F., Jaimungal, S.: Algorithmic trading with model uncertainty. *SIAM J. Financ. Math.* **8**(1), 635–671 (2017)
22. Cheng, Stephanie F., De Franco, Gus, Jiang, Haibo, Lin, Pengkai: Riding the blockchain mania: public firms' speculative 8-k disclosures. *Manag. Sci.* **65**(12), 5901–5913 (2019)
23. Cohen, S.N., Elliott, R.J., Pearce, C.E.M.: A general comparison theorem for backward stochastic differential equations. *Adv. Appl. Probab.* **42**(3), 878–898 (2010)
24. Credit Suisse Corporate Insights Group: *The Investor Landscape: Four Evolving Themes and Their Implications*. Technical Report, Credit Suisse, NY (2022)
25. Delarue, F.: Restoring uniqueness to mean-field games by randomizing the equilibria. *Stoch. Partial Differ. Equ. Anal. Comput.* **7**(4), 598–678 (2019)
26. Delarue, F., Lacker, D., Ramanan, K.: From the master equation to mean field game limit theory: large deviations and concentration of measure. *Ann. Probab.* **48**, 211–263 (2020)
27. Dellacherie, C., Maisonneuve, B., Paul-André M.: *Probabilités et potentiel, chapitres XVII–XXIV, Processus de Markov (fin). Compléments de calcul stochastique*, vol. 5. Hermann (1992)
28. Djete, M.F.: Mean field games of controls: on the convergence of Nash equilibria. *Ann. Appl. Probab.* **33**(4), 2824–2862 (2023)
29. Doblas-Madrid, A.: A robust model of bubbles with multidimensional uncertainty. *Econometrica* **80**(5), 1845–1893 (2012)
30. Doblas-Madrid, A.: A finite model of riding bubbles. *J. Math. Econ.* **65**, 154–162 (2016)

31. Pardoux, É.: Generalized discontinuous backward stochastic differential equations. In: El Karoui, N., Mazliak, L. (eds.) *Backward Stochastic Differential Equations*, Chapter IV, pp. 207–219. Addison Wesley Longman Limited (1997)
32. Gangbo, W., Mészáros, A.R., Mou, C., Zhang, J.: Mean field games master equations with nonseparable hamiltonians and displacement monotonicity. *Ann. Probab.* **50**(6), 2178–2217 (2022)
33. Griffin, J.M., Harris, J.H., Shu, T., Topaloglu, S.: Who drove and burst the tech bubble? *J. Finance* **66**(4), 1251–1290 (2011)
34. Haussmann, U.G., Lepeltier, J.P.: On the existence of optimal controls. *SIAM J. Control Optim.* **28**(4), 851–852 (1990)
35. Huang, M., Malhamé, R.P., Caines, P.E.: Large population stochastic dynamic games: closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.* **6**(3), 221–252 (2006)
36. Huang, M., Caines, P.E., Malhamé, R.P.: The Nash certainty equivalence principle and McKean–Vlasov systems: an invariance principle and entry adaptation. In: 2007 46th IEEE Conference on Decision and Control, pp. 121–126 (2007)
37. Jackson, J., Tangpi, L.: Quantitative convergence for displacement monotone mean field games with controlled volatility (2023)
38. Jacod, J., Mémin, J.: Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité. *Séminaire de probabilités de Strasbourg, Tome 15*, 529–546 (1981)
39. Johansen, A., Sornette, D., Ledoit, O.: Predicting financial crashes using discrete scale invariance. *J. Risk* **1**(4), 5–32 (1999)
40. Johansen, A., Ledoit, O., Sornette, D.: Crashes at critical points. *Int. J. Theor. Appl. Finance* **3**, 219–255 (2000)
41. Brunnermeier, M.K., Nagel, S.: Hedge funds and the technology bubble. *J. Finance* **59**(5), 2013–2040 (2004)
42. Karatzas, I., Shreve, S.E.: *Brownian Motion and Stochastic Calculus*. Graduate Texts in Mathematics, 2nd edn. Springer, New York (1991)
43. El Karoui, N., Du, N.H., Jeanblanc-Picqué, M.: Compactification methods in the control of degenerate diffusions: existence of an optimal control. *Stochastics* **20**(3), 169–219 (1987)
44. Kharroubi, I., Lim, T.: Progressive enlargement of filtrations and Backward SDEs with jumps. *J. Theor. Probab.* **27**, 683–724 (2014)
45. Kurtz, T.: Weak and strong solutions of general stochastic models. *Electron. Commun. Probab.* **19**, 1–16 (2014)
46. Lacker, D.: A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Relat. Fields* **165**(3), 581–648 (2016)
47. Lacker, D., Le Flem, L.: Closed-loop convergence for mean field games with common noise. *Ann. Appl. Probab.* **33**(4), 2681–2733 (2023)
48. Lacker, D., Shkolnikov, M., Zhang, J.: Superposition and mimicking theorems for conditional McKean Vlasov equations. *J. Eur. Math. Soc.* (2020)
49. Lasry, J.-M., Lions, P.-L.: Jeux à champ moyen. I-le cas stationnaire. *Comptes Rendus Mathématique* **343**(9), 619–625 (2006)
50. Lasry, J.-M., Lions, P.-L.: Jeux à champ moyen. II-horizon fini et contrôle optimal. *Comptes Rendus Mathématique* **343**(10), 679–684 (2006)
51. Lasry, J.-M., Lions, P.-L.: Mean field games. *Jpn. J. Math.* **2**(1), 229–260 (2007)
52. Laurière, M., Tangpi, L.: Convergence of large population games to mean field games with interaction through the controls. *SIAM J. Math. Anal.* **54**(3), 3535–3574 (2022)
53. Leal, L.S.: Topics in High-Frequency Optimal Execution and Microstructure of Product Repricings. Ph.D. thesis, Princeton University (2022)
54. Liu, F., Conlon, J.R.: The simplest rational greater-fool bubble model. *J. Econ. Theory* **175**, 38–57 (2018)
55. Tóth, B., Mastromatteo, I., Bouchaud, J.P.: Anomalous impact in reaction-diffusion financial models. *Phys. Rev. Lett.* **113**, 268701 (2014)
56. Mou, C., Zhang, J.: Mean field game master equations with anti-monotonicity conditions (2022)
57. Pohl, M., Ristig, A., Schachermayer, W., Tangpi, L.: The amazing power of dimensional analysis: quantifying market impact. *Market Microstruct. Liqu.* **3**(3–4), 1850004 (2018)
58. Rubinstein, A., Wolinsky, A.: Middlemen. *Q. J. Econ.* **102**(3), 581–593 (1987)
59. Sornette, D.: Critical market crashes. *Phys. Rep.* **378**(1), 1–98 (2003)
60. Sornette, D., Woodard, R., Yan, W., Zhou, W.-X.: Clarifications to questions and criticisms on the Johansen–Ledoit–Sornette financial bubble model. *Phys. A Stat. Mech. Appl.* **392**(19), 4417–4428 (2013)
61. Sotes-Paladino, J., Zapatero, F.: Riding the bubble with convex incentives. *Rev. Financ. Stud.* **32**(4), 1416–1456 (2018)

62. Tangpi, L., Wang, S.: Optimal Bubble Riding: A Mean Field Game with Varying Entry Times (2022). arXiv preprint:<https://arxiv.org/abs/2209.04001>
63. Villani, C.: Optimal Transport. Springer, Berlin (2009)
64. Yamada, T., Watanabe, S.: On the uniqueness of solutions of stochastic differential equations. *J. Math. Kyoto Univ.* **11**(1), 155–167 (1971)
65. Yor, M.: Application d'un lemme de Jeulin au grossissement de la filtration brownienne. *Séminaire de probabilités de Strasbourg* **14**, 189–199 (1980)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.