Mean Field Games of Control and Cryptocurrency Mining

Nicolás García*† Ronnie Sircar* H. Mete Soner*‡

May 2025, revised June 2026

Abstract

This paper studies Mean Field Games (MFGs) in which agent dynamics are given by jump processes of controlled intensity, with mean-field interaction via the controls and affecting the jump intensities. We establish the existence of MFG equilibria in a general discrete-time setting, and prove a limit theorem as the time discretization goes to zero, establishing equilibria in the continuous-time setting for a class of MFGs of intensity control. This motivates numerical schemes that involve directly solving discrete-time games as opposed to coupled Hamilton-Jacobi-Bellman and Kolmogorov equations. As an example of the general theory, we consider cryptocurrency mining competition, modeled as an MFG both in continuous and discrete time, and illustrate the effectiveness of the discrete-time algorithm to solve it.

1 Introduction

The study of Mean Field Games (MFGs) dates back to the foundational papers [21] and [26], and we refer to [12] for a comprehensive review of the literature, most of which has focused on continuous-time models. There are comparatively few works treating discrete-time setups, of which we highlight [17], which considers a finite state and horizon setup where agents control their transition probabilities, [10] which allows for a Polish state space incorporating mean-field interactions only through the costs, and a discrete-time MFG with countable state space studied in [1]. Additionally, [30], [29], and [31] consider linear dynamics in a variety of discrete-time settings, and [32] analyzes the infinite-horizon discounted-cost problem with a Polish state space and MFG interaction via the states, among others.

Motivated by the cryptocurrency mining MFG model in [28], we study a class of problems set both in discrete-time and continuous-time in which the controlled dynamics follow a jump Markov process and the mean-field interaction is via the controls. The latter property complicates the continuous-time analysis due to the loss of regularity in the control measure flow when compared to the state measure flow in the more typical setup involving mean-field interactions of state. MFGs with jump process dynamics in the continuous-time and finite state setting are considered in [5, 16], where agents' controls are their transition probabilities and they interact via the empirical (joint) state measures. The techniques involved are primarily based on ODE and PDE methods related to the associated master equations. We also mention [3] and [4] which study mean field game models for queuing systems under heavy-traffic scaling, including numerical approximation and

*Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ, 08540, USA.
†Acknowledges support from NSERC CGSD.
‡Partially supported by the National Science Foundation grant DMS 2406762.

convergence from large strategic-server systems. A probabilistic treatment of MFGs with jumps is given in [6], where relaxed controls are used in a weak formulation to provide general existence results for MFGs with jump-diffusion dynamics where both jump size and intensity are controlled, and where the mean-field interaction is limited to the states. In our analysis, we also utilize relaxed control techniques as developed in [25] for diffusion MFGs of state and [13] for diffusion MFGs of control. Our results complement [8, 34], and directly apply to the cryptocurrency MFG model from [28] for which we provide theoretical existence guarantees.

We begin our analysis by considering discrete-time finite-horizon MFGs in Section 2, in which we let the state, noise, and control spaces be arbitrary Polish, and consider general transition dynamics. Following the construction from [32] which in turn was inspired by [22], we characterize the MFG equilibria as fixed points of a set-valued operator and establish fixed-point existence by way of Kakutani's theorem. In comparison to [32], we allow for mean-field interaction via the controls, and prove existence in the finite-time horizon setting under a weaker growth assumption on the transition dynamics.

In Section 3, we consider a concrete continuous-time MFG with state dynamics given by a jump process and prove that it arises as the limit of analogous discrete-time models. We assume that the drift and the intensity coefficients do not depend on the state process, and the MFG interaction is via the controls affecting only the intensity but not the drift of the agent's jump processes. We consider this concrete setup in order to avoid routine but technical details which can be found in the classical approximations literature for stochastic optimal control including [14, 23, 24].

In Section 4, we apply our results from Sections 2 and 3 to establish MFG existence for the cryptocurrency mining MFG model of [28] as well as its discrete-time analogue. We compute the discrete-time MFG using damped fixed-point iterations, and reproduce the qualitative equilibrium behavior which was established in [28] using a finite difference scheme for the associated coupled PDEs. We remark on the uniqueness of equilibrium in Section 4.4.

Notation and Terminology: For an integer d, let $\mathcal{S}^d$ denote the d-dimensional simplex, $[d] := \{1, 2, \cdots, d\}$, and define $\Delta_i : \mathbb{R}^d \to \mathbb{R}^d$ by $\Delta_i x := (x_1 - x_i, x_2 - x_i, \cdots, x_d - x_i)$. Given a Polish space $\mathcal{X}$, let $\mathcal{B}(\mathcal{X})$ denote the Borel σ-algebra on $\mathcal{X}$ and let $\mathcal{P}(\mathcal{X})$ denote the set of probability measures on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$. Also, $\mathcal{M}(\mathcal{X})$ will denote the set of non-negative finite measures on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$. We endow these spaces with the weak topology, i.e., $\mu_n \to \mu$ if $\lim_{n\to\infty} \int_{\mathcal{X}} f d\mu_n = \int_{\mathcal{X}} f d\mu$ for any continuous and bounded $f : \mathcal{X} \to \mathbb{R}$. To emphasize convergence in the weak sense, we write $\mu_n \Rightarrow \mu$, or if the objects are probability measures, $\mu_n \xrightarrow{\mathcal{L}, n\to\infty} \mu$. Let $\mathbb{R}_{\geq 0}$ denote the non-negative reals. For a random variable X, $\mathcal{L}(X)$ denotes its distribution. For $\mu \in \mathcal{P}(\mathcal{X})$, $\delta_{\{a\}}$ denotes the Dirac measure located at $a \in \mathcal{X}$. If $\mathcal{X} \subseteq \mathbb{R}$ we write $\overline{\mu} := \int_{\mathcal{X}} x\mu(dx)$ for the mean of the distribution. We say that a sequence of maps $g_n : \mathcal{X} \to \mathbb{R}$ converge continuously to $g : \mathcal{X} \to \mathbb{R}$ if whenever $x_n \xrightarrow{n\to\infty} x$, we have that $g_n(x_n) \xrightarrow{n\to\infty} g(x)$.

2 Mean Field Games in Discrete Time

This section establishes the existence of equilibria for discrete-time, finite-horizon MFGs where the mean-field interaction, which affects the dynamics and costs, is through both the controls and states. We motivate the problem by first introducing the N-player game.

2.1 N-Player Game Formulation

We begin with the following N-player setup, where:

- The agents' state processes take values in a Polish space $\mathcal{X}$ and evolve in discrete time steps $k = 0, 1, 2, \ldots, T$. We denote by x_k^i the state of agent $i \in [N] := \{1, 2, \ldots, N\}$ at time k.
- The initial states x_0^i are drawn i.i.d. from a measure $\mu_0 \in \mathcal{P}(\mathcal{X})$.
- The action space $\mathcal{U}$ is assumed to be Polish, $a_k^i \in \mathcal{U}$ denotes the action of agent $i \in [N]$ at time k, and $e_k := (e_k^c, e_k^s) \in \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X})$ denotes the empirical distribution of the agents' controls and states at time k.
- The transition dynamics of each agent are given by a Markov transition kernel

  $$\rho : \mathcal{X} \times \mathcal{U} \times \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}) \to \mathcal{P}(\mathcal{X}) \quad \text{so that} \quad x_{k+1}^i \sim \rho(\cdot | x_k^i, a_k^i, e_k^c, e_k^s)$$

  for each $k = 0, 1, 2, \ldots, T-1$.

A control policy for player i consists of a sequence $(\pi_k^i)_{k=0}^T$ of $\mathcal{P}(\mathcal{U})$-valued random variables adapted to the filtration

$$\mathcal{F}_0^i := \sigma(x_0^i, e_0^s), \quad \mathcal{F}_k^i := \sigma\Big(\mathcal{F}_{k-1}^i \cup \sigma(x_k^i, a_{k-1}^i, e_{k-1}^c, e_k^s)\Big), \quad \text{for all } k = 1, 2, \ldots, T.$$

Conditioned on $\mathcal{F}_k^i$, the action a_k^i of agent i is drawn randomly (and independent of any other random quantity) from the distribution π_k^i (i.e. $a_k^i \sim \pi_k^i$).

The last step in the specification of the model is to define the optimality criterion, which is in terms of the one-step running and terminal cost functions

$$c : \mathcal{X} \times \mathcal{U} \times \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}) \to [0, \infty), \qquad \varphi : \mathcal{X} \times \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}) \to [0, \infty).$$

Fixing an N-tuple of control policies $\pi^{(N)} := (\pi^{N,1}, \pi^{N,2}, \ldots, \pi^{N,N})$, the i-th agent incurs a cost

$$J^i(\pi^{(N)}) := J^i(\pi^{N,i}, \pi^{N,-i}) := \mathbb{E}^{\pi^{(N)}}\Big[\sum_{k=0}^{T-1} c(x_k^i, a_k^i, e_k) + \varphi(x_T^i, e_T)\Big],$$

where the superscript $\pi^{(N)}$ on the expectation denotes that the control actions of the agents are determined according to their respective policies from $\pi^{(N)}$, and where $\pi^{N,-i}$ denotes the policies of all agents except for the i-th one. Agents wish to select their respective policies to minimize costs, and a solution to the N-player game consists of an equilibrium, which is a joint policy $\tilde{\pi}^{(N)}$ such that

$$J^i(\tilde{\pi}^{(N)}) = \inf_{\pi} J^i(\pi, \tilde{\pi}^{N,-i}),$$

for every $i = 1, \ldots, N$. This completes the N-player game setup.

2.2 Mean Field Game Formulation

We proceed with the representative agent problem which characterizes the corresponding MFG. We denote the state and control of the representative agent at time k by $x_k \in \mathcal{X}$ and $a_k \in \mathcal{U}$, respectively. In this case, the representative agent has identical dynamics as in the N-player game, but the sequences of empirical measures $(e^s_k)_{k=0}^T$ and $(e^c_k)_{k=0}^T$ are taken to be deterministic. Denoting $\delta_k = (e^c_k, e^s_k) \in \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X})$, the dynamics follow

$$x_{k+1} \sim \rho(\cdot|x_k, a_k, \delta_k), \quad a_l \sim \pi_l(\cdot|x_l), \quad k = 0, 1, \ldots, T-1, \quad l = 0, 1, \ldots, T \tag{2.1}$$

where, because $\delta \coloneqq (\delta_k)_{k=0}^T$ is now a parameter, control policies become random measures adapted to the filtration of the state process. We let Π denote the set of such policies, but will in fact search for an optimal policy within the following smaller set of Markov control policies.

The transition kernel depends only on the marginal state and control laws. Joint state-action laws are used later in the fixed-point construction so that randomized Markov controls can be recovered by disintegration.

Definition 2.1. *A control policy is called Markov if it is a sequence $(\pi_k)_{k=0}^T$ where each π_k is measurable w.r.t. $\sigma(x_k)$. Under such a policy, the measure $\pi_k(\cdot|x_k)$ used to generate the control at a given time k depends only on the state at that time. We denote by $\mathbb{M}$ the set of all such policies.*

It is well known that in a Markov Decision Problem (MDP) setting, the restriction to Markov policies does not result in a larger value function (see for example [32, Proposition 3.2]). Given a fixed sequence of probability measures $\delta \coloneqq (\delta_k)_{k=0}^T \subseteq \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X})$, the representative agent's control problem consists of determining a Markov policy π^* such that

$$J(\pi^*, \delta) = \inf_{\pi \in \mathbb{M}} J(\pi, \delta) \quad \text{for} \quad J(\pi, \delta) \coloneqq E^\pi\Big[\sum_{k=0}^{T-1} c(x_k, a_k, \delta_k) + \varphi(x_T, \delta_T)\Big],$$

with dynamics given by (2.1).

For convenience in defining an MFG equilibrium, we introduce

$$\Phi : (\mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}))^{T+1} \to 2^\Pi \quad \text{by} \quad \Phi(\delta) \coloneqq \{\pi^* \in \Pi : \pi^* \text{ minimizes } J(\cdot, \delta)\}.$$

We additionally define a map $\Lambda : \Pi \to (\mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}))^{T+1}$ by constructing $\Lambda(\pi) \coloneqq (\delta^c_k, \delta^s_k)_{k=0}^T$ iteratively using the initial state law μ_0 as follows: First, set $\delta^s_0 = \mu_0$. For all other $k \geq 0$, define

$$\delta^c_k(\cdot) = \int_{\mathcal{X}} \mathcal{P}^\pi_k(\cdot|x)\delta^s_k(dx), \quad \delta^s_{k+1}(\cdot) = \int_{\mathcal{X}} \int_{\mathcal{U}} \rho(\cdot|x, a, \delta_k)\mathcal{P}^\pi_k(da|x)\delta^s_k(dx),$$

where $\mathcal{P}^\pi_k(\cdot|x)$ denotes the conditional law of a_k given the event $\{x_k = x\}$ under the fixed flow of measures δ and control policy π which specify the dynamics. The above equations are analogous to the Kolmogorov PDE in continuous-time and $\delta = \Lambda(\pi)$ represents the sequence of distributions over the control and state space in the infinite-player limit when all agents use policy π and are initially distributed on the state space according to μ_0. The maps Λ and Φ allow us to define the equilibrium compactly.

Definition 2.2. A pair $(\pi, \delta) \in \mathbb{M} \times (\mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}))^{T+1}$ is called *an MFG equilibrium* if and only if

$$\pi \in \Phi(\delta) \qquad \textit{and} \qquad \delta = \Lambda(\pi).$$

As discussed immediately following Definition 2.1, the restriction to Markov policies has no impact on the value function, thus an MFG equilibrium over policies in $\mathbb{M}$ is also an MFG equilibrium over policies in Π. Our approach to establish an existence theorem is to construct a set-valued operator whose fixed points correspond to MFG equilibria and use Kakutani's fixed-point theorem to guarantee the existence of a fixed point. We closely follow the structure from the proof of [32, Theorem 3.3]. In the next two subsections, we construct spaces for the set-valued operator and specify its construction, establish that its fixed points correspond to equilibria, and prove the existence of fixed points. We now proceed with the required assumptions.

2.3 Model Assumptions

This section contains the assumptions required for the existence of an MFG equilibrium. Fix a continuous moment $w : \mathcal{X} \to [1, \infty)$ on the state space, which is a map for which there exists a sequence of compact sets $(H_n)_{n=1}^{\infty} \subseteq \mathcal{X}$ which are increasing (in the sense that $H_n \subseteq H_{n+1}$ for every n), such that $\lim_{n\to\infty} \inf_{x\in\mathcal{X}\setminus H_n} w(x) = \infty$, and which satisfies $w(\cdot) \geq 1 + d_{\mathcal{X}}(\cdot, x_0)^p$ for some $x_0 \in \mathcal{X}$ and some $p \geq 1$, where $d_{\mathcal{X}}$ denotes a metric on $\mathcal{X}$ compatible with its topology.

To treat two cases simultaneously, let $v := \mathbb{1}$ be the function of $\mathcal{X}$ identically equal to one when both c and φ are assumed bounded, and $v := w$ otherwise. Slightly abusing notation by using the same name for each, we define the v-norm on maps $g : \mathcal{X} \to \mathbb{R}$ and on measures $\mu \in \mathcal{P}(\mathcal{X})$ by

$$\|g\|_v := \sup_{x\in\mathcal{X}} \frac{|g(x)|}{v(x)} \quad \text{and} \quad \|\mu\|_v := \int_{\mathcal{X}} v(x)\mu(dx).$$

Moreover, we let $B_v(\mathcal{X})$ denote the space of all real-valued measurable functions on $\mathcal{X}$ with finite v-norm and let $C_v(\mathcal{X}) \subseteq B_v(\mathcal{X})$ denote the subset of continuous functions. Both of these are Banach spaces: If $v = \mathbb{1}$, this is an elementary result, and the $v = w$ case follows from almost identical arguments as the $v = \mathbb{1}$ case. Finally, we define

$$\mathcal{P}_v(\mathcal{X}) := \{\mu \in \mathcal{P}(\mathcal{X}) : \|\mu\|_v < \infty\} = \{\mu \in \mathcal{P}(\mathcal{X}) : \int_{\mathcal{X}} v(x)\mu(dx) < \infty\}. \tag{2.2}$$

The following are the main assumptions required for the equilibrium existence theorem.

Assumption 2.3.

(i) The maps c and φ are continuous.

(ii) $\mathcal{U}$ is compact and $\mathcal{X}$ is locally compact (every point has a compact neighborhood).

(iii) There exists a constant $\alpha \geq 1$ such that

$$\sup_{(a,\delta^c,\delta^s)\in\mathcal{U}\times\mathcal{P}(\mathcal{U})\times\mathcal{P}(\mathcal{X})} \int_{\mathcal{X}} w(y)\rho(dy|x, a, \delta^c, \delta^s) \leq \alpha w(x) \quad \textit{for all } x \in \mathcal{X}.$$

(iv) The stochastic kernel ρ is weakly continuous in the sense that if $(x_n, a_n, \delta^c_n, \delta^s_n) \xrightarrow{n\to\infty} (x, a, \delta^c, \delta^s)$, then $\rho(\cdot|(x_n, a_n, \delta^c_n, \delta^s_n)) \xrightarrow{n\to\infty} \rho(\cdot|(x, a, \delta^c, \delta^s))$ under the topology of weak convergence. In addition, $\int_{\mathcal{X}} w(y)\rho(dy|x, a, \delta^c, \delta^s)$ is continuous in the variables x, a, δ^c, and δ^s.

(v) The initial law μ_0 satisfies $M := \int_{\mathcal{X}} v(x)\mu_0(dx) < \infty$.

(vi) There exists $R > 0$ satisfying

$$\sup_{(a,\delta)\in\mathcal{U}\times\mathcal{P}(\mathcal{U})\times\mathcal{P}(\mathcal{X})} c(x, a, \delta) \leq R\ v(x), \ \textit{and} \sup_{\delta\in\mathcal{P}(\mathcal{U})\times\mathcal{P}(\mathcal{X})} \varphi(x, \delta) \leq R\ v(x), \quad \forall x \in \mathcal{X}.$$

Before proceeding, we note that Assumption 2.3.(iii) and Assumption 2.3.(vi) ensure uniform control of moments and costs, which is essential for keeping the sequence of measures tight and bounded in the v-norm. In particular, Assumption 2.3.(iii) propagates moment bounds through the dynamics, while Assumption 2.3(vi) guarantees uniform bounds on costs and value functions via the Bellman operator. These conditions are therefore key to the compactness and norm estimates required in the fixed-point argument. Finally, we note that all the results in this section can be generalized to the time-inhomogeneous setting where the transition kernels depend on time if we assume that (iii) and (iv) hold for any choice of time index. For simplicity of setup and notation, we do not pursue this extension here.

2.4 Existence of MFG Equilibria

This section contains a proof of the following existence theorem.

Theorem 2.4. *Under Assumption 2.3, there exists an MFG equilibrium in the sense of Definition 2.2.*

The proof can be found at the end of the current section. Before proceeding with it, we introduce some required notions and construct a set-valued map whose fixed points correspond to MFG equilibria.

Definition 2.5. *(Value Function) Consider the non-homogeneous Markov Decision Process whose k-th step transition kernel is given by $\rho(\cdot|x, a, \delta_k)$, for a fixed flow of measures $\delta := (\delta_k)_{k=0}^T$ and costs as defined above. We let $V_k^\delta : \mathcal{X} \to \mathbb{R}$ denote the value function at time $k = 0, 1, \dots, T$. In other words, we define $V_T^\delta := \varphi$ and iteratively define*

$$V_k^\delta(x) := \min_{a\in\mathcal{U}} \Big(c(x, a, \delta_k) + \int_{\mathcal{X}} V_{k+1}^\delta(y)\rho(dy|x, a, \delta_k) \Big)$$

for $k = 0, 1, 2, \cdots, T-1$. Furthermore, let $V^\delta := (V_k^\delta)_{k=0}^T$.

Recall that $M = \int_{\mathcal{X}} v(x)\mu_0(dx)$ and the definition of $\mathcal{P}_v(\mathcal{X})$ in (2.2). We now define the spaces

$$\mathcal{P}_v^k(\mathcal{X}) := \{\mu \in \mathcal{P}_v(\mathcal{X}) : \int_{\mathcal{X}} w(x)\mu(dx) \le \alpha^k M\},$$

$$\mathcal{P}_v^k(\mathcal{X} \times \mathcal{U}) := \{\mu \in \mathcal{P}(\mathcal{X} \times \mathcal{U}) : \mu_1 \in \mathcal{P}_v^k(\mathcal{X})\},$$

where in the second set, μ_1 denotes the state marginal of μ (i.e. $\mu_1(\cdot) = \mu(\cdot \times \mathcal{U})$). Taking R and α as in Assumption 2.3, define for $k = 0, 1, \dots, T$

$$L_k = R \sum_{l=0}^{T-k} \alpha^l \quad \text{and observe that} \quad L_k = R + \alpha L_{k+1} \quad \text{for every } k = 0, 1, 2, \cdots, T-1.$$

Next, define for every $k = 0, 1, \dots, T$

$$C_v^k(\mathcal{X}) := \{u \in C_v(\mathcal{X}) : \|u\|_v \le L_k\}, \quad C := \prod_{k=0}^T C_v^k(\mathcal{X}), \quad \Xi := \prod_{k=0}^T \mathcal{P}_v^k(\mathcal{X} \times \mathcal{U}).$$

The following lemma establishes regularity properties of the value function that will be required later.

Lemma 2.6. *For any $\nu \in \Xi$ we have that $V^\nu \in C$.*

Proof. First, observe that by definition $V_T^\nu = \varphi \in C_v^T$. Inductively, we have

$$\begin{aligned}
\frac{1}{v(x)} V_k^\nu(x) &= \min_{a\in\mathcal{U}} \Big[\frac{1}{v(x)} c(x, a, \nu_k) + \int_{\mathcal{X}} \frac{1}{v(x)} V_{k+1}^\nu(y)\rho(dy|x, a, \nu_k) \Big] \\
&\le \min_{a\in\mathcal{U}} \Big[R + \frac{1}{v(x)} \int_{\mathcal{X}} L_{k+1} v(y)\rho(dy|x, a, \nu_k) \Big] \\
&\le R + \alpha L_{k+1} = L_k.
\end{aligned}$$

By assumption, φ is continuous and thus inductively, and using [7, Proposition 7.32], which establishes that the Bellman operator (see the following definition) preserves continuity, continuity of all the value functions follows. □

Definition 2.7. *(Bellman operator) For a given $\nu \in \Xi$, the Bellman operator acting on maps $u : \mathcal{X} \to \mathbb{R}$ is defined by setting*

$$T_k^\nu u(x) \coloneqq \min_{a\in\mathcal{U}} \Big[c(x,a,\nu_k) + \int_{\mathcal{X}} u(y)\rho(dy|x,a,\nu_k)\Big] \quad \text{for } k = 0,1,2,\dots,T,$$

and defining the operator $T^\nu : C \to C$ by

$$(T^\nu u)_k \coloneqq \begin{cases} T_k^\nu u_{k+1} & \text{for } k = 0,1,\cdots,T-1 \\ \varphi & \text{for } k = T. \end{cases}$$

By definition, the value function is a fixed point of this operator. The following result establishes that in fact T^ν maps C into itself, and will be required when applying Kakutani's theorem.

Lemma 2.8. *Let $\nu \in \Xi$ be arbitrary. Then for all $k = 0,1,\dots,T-1$ the operator T_k^ν maps $C_v^{k+1}(\mathcal{X})$ into $C_v^k(\mathcal{X})$.*

Proof. Let $u \in C_v^{k+1}(\mathcal{X})$. Since the Bellman operator preserves continuity (see [7, Proposition 7.32]) we immediately have that $T_k^\nu u$ is continuous. Because

$$\begin{aligned}\left\|T_k^\nu u\right\|_v = \sup_{x\in\mathcal{X}} \frac{|(T_k^\nu u)(x)|}{v(x)} &= \sup_{x\in\mathcal{X}} \frac{\left|\min_{a\in\mathcal{U}}\left[c(x,a,\nu_k) + \int_{\mathcal{X}} u(y)\rho(dy|x,a,\nu_k)\right]\right|}{v(x)} \\ &\le \sup_{(x,a)\in\mathcal{X}\times\mathcal{U}} \frac{c(x,a,\nu_k) + \int_{\mathcal{X}} |u(y)|\rho(dy|x,a,\nu_k)}{v(x)} \\ &\le \sup_{(x,a)\in\mathcal{X}\times\mathcal{U}} \frac{c(x,a,\nu_k) + L_{k+1}\int_{\mathcal{X}} v(y)\rho(dy|x,a,\nu_k)}{v(x)} \\ &\le \sup_{(x,a)\in\mathcal{X}\times\mathcal{U}} \frac{Rv(x) + \alpha L_{k+1}v(x)}{v(x)} = L_k,\end{aligned}$$

it follows that $T_k^\nu u \in C_v^k(\mathcal{X})$, as claimed. □

We are ready to define a set-valued operator whose fixed points will correspond to MFG equilibria. Recall that a pair $(\pi,\delta) \in \Pi \times (\mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{X}))^{T+1}$ is an MFG equilibrium if and only if we have both $\pi \in \Phi(\delta)$ (optimality) and $\delta \in \Lambda(\pi)$ (consistency). Let $\Gamma : \Xi \to 2^{\mathcal{P}(\mathcal{X}\times\mathcal{U})^{T+1}}$ denote a set-valued operator defined by $\Gamma(\nu) = C(\nu) \cap B(\nu)$ where

$$\begin{aligned} C(\nu) &\coloneqq \big\{\nu' \in \mathcal{P}(\mathcal{X}\times\mathcal{U})^{T+1} : \\ &\qquad \nu'_{k+1,1}(\cdot) = \int_{\mathcal{X}\times\mathcal{U}} \rho(\cdot|x,a,\nu_k)\nu_k(dx,da) \text{ for every } k \ge 0, \quad \nu'_{0,1} = \mu_0\big\}, \\ B(\nu) &\coloneqq \Big\{\nu' \in \mathcal{P}(\mathcal{X}\times\mathcal{U})^{T+1} : \\ &\qquad \nu'_k\Big(\Big\{(x,a) : c(x,a,\nu_k) + \int_{\mathcal{X}} V_{k+1}^\nu(y)\rho(dy|x,a,\nu_k) = T_k^\nu V_{k+1}^\nu(x)\Big\}\Big) = 1 \text{ for every } k \ge 0\Big\}. \end{aligned}$$

Lemma 2.9. *Given a fixed point ν of Γ, one can obtain a MFG equilibrium via disintegration.*

Proof. Let $\nu \in \Xi$ be a fixed point of Γ. We then define a Markov randomized control policy by disintegrating the measure ν_k into its marginal on $\mathcal{X}$ and the conditional measures, which we define to be π_k. In other words, we have

$$\nu_k(dx, da) = \pi_k(da|x)\nu_k^s(dx) \quad \text{for every } k = 0, 1, 2, \cdots, T.$$

Writing $\pi \coloneqq (\pi_k)_{k=0}^T$, it follows that (ν, π) constitutes an MFG equilibrium for our setup, with optimality and consistency following from the definitions of C and B above. In particular, optimality follows because the random control is supported on the set of maximizer of the Bellman operator (see [20, Theorem 17.1] for a proof of this result). ☐

We are almost ready to complete the MFG equilibrium existence proof by establishing existence of a fixed point of Γ. Before doing this, however, we prove the following lemma.

Lemma 2.10. *For any compact set $K \subseteq \mathcal{X}$ and any time step $k = 0, 1, \ldots, T$ we have*

$$\lim_{n\to\infty} \sup_{x\in K} |V_k^{\nu^{(n)}}(x) - V_k^\nu(x)| = 0.$$

Proof. We establish the result by backwards induction on the time step. The result is trivially true for $k = T$ since $V_T^{\nu^{(n)}} = V_T^\nu = \varphi$. Suppose now that the result holds for all time steps $k+1$ or larger. Then we have

$$\begin{aligned}\sup_{x\in K} |V_k^{\nu^{(n)}}(x) - V_k(x)| \leq &\sup_{x\in K, a\in\mathcal{U}} |c(x, a, \nu_k^{(n)}) - c(x, a, \nu_k)| \\ &+ \sup_{x\in K, a\in\mathcal{U}} \Big| \int_{\mathcal{X}} V_{k+1}^{\nu^{(n)}}(y)\rho(dy|x, a, \nu_k^{(n)}) - \int_{\mathcal{X}} V_{k+1}(y)\rho(dy|x, a, \nu_k)\Big|.\end{aligned}$$

Recall that continuous convergence is equivalent to uniform convergence on compact sets when the limit function is continuous. Using the joint continuity of c in its three arguments, it follows that $c(\cdot, \cdot, \nu_k^{(n)})$ converges continuously to $c(\cdot, \cdot, \nu_k)$. By inductive hypothesis and using [33, Theorem 3.3] with $y \mapsto L_{k+1}v(y)$ the dominating function, we also have continuous convergence of the maps $G_n(x, a) \coloneqq \int_{\mathcal{X}} V_{k+1}^{\nu^{(n)}}(y)\rho(dy|x, a, \nu_k^{(n)})$ to $G(x, a) \coloneqq \int_{\mathcal{X}} V_{k+1}^\nu(y)\rho(dy|x, a, \nu_k)$ and hence uniform convergence on compacts. As such, both terms above converge to zero and the result follows. ☐

Proof. To show that Γ admits a fixed point, we use [2, 17.55 Corollary (Kakutani–Fan–Glicksberg)]. We follow the arguments from the proof of [32, Theorem 3.3] (in turn adapted from the arguments of [22]), but due to the finite horizon and less restrictive growth assumptions, we do not have that the Bellman operator is a contraction. Nevertheless, the direct argument in Lemma 2.10 allows us to conclude uniform convergence of the value function over compact sets, which is required in the proof. We must check that Γ maps a non-empty compact convex subset of a locally convex Hausdorff topological vector-space into itself, and when restricted to this subset, its graph is closed, and it has non-empty convex values. Of course, our candidate subset is Ξ, and we take the ambient space to be the $T+1$-fold tuple of finite, signed measures on $\mathcal{X} \times \mathcal{U}$. First, we note that the set $\Gamma(\nu)$ can be shown to be non-empty via a measurable selector argument as in [22, Theorem 1], which in turn makes use of the arguments [19, p. 54]. Next, observe that per [18, Proposition E.8] the set of measures $\mathcal{P}_v^k(\mathcal{X})$ is tight. It can easily be shown that this set is closed in the weak topology, hence using Prokhorov's theorem we conclude that it is compact. As $\mathcal{U}$ is compact, we conclude that $\mathcal{P}_v^k(\mathcal{X}\times\mathcal{U})$ and thus Ξ are also compact, the latter in the product weak topology. Clearly, Ξ is convex, thus satisfies the conditions required to use [2, 17.55 Corollary (Kakutani–Fan–Glicksberg)].

We proceed by showing that $\Gamma(\nu)\subseteq\Xi$ for any $\nu\in\Xi$. Indeed, it suffices to show that for an arbitrary $\nu\in\Xi$ we have that $C(\nu)\in\Xi$. As such, let $\nu'\in C(\nu)$. We need to check that $\nu'_k\in\mathcal{P}_v^k(\mathcal{X}\times\mathcal{U})$ for every $k\geq 0$. In other words, we need to check that $\nu'_{k,1}\in\mathcal{P}_v^k(\mathcal{X})$ for all $k\geq 0$. By definition, we know that $\nu'_{0,1}=\mu_0$ thus we have that

$$\int_{\mathcal{X}} w(x)\nu'_{0,1}(dx)=\int_{\mathcal{X}} w(x)\mu_0(dx)=\alpha^0 M,$$

and so by definition, $\nu'_{0,1}\in\mathcal{P}_v^0(\mathcal{X})$. For any other k, simply observe that by definition of $C(\nu)$, we have

$$\begin{aligned}\int_{\mathcal{X}} w(y)\nu'_{k+1,1}(dy)&=\int_{\mathcal{X}\times\mathcal{U}}\int_{\mathcal{X}} w(y)\rho(dy|x,a,\nu_k)\nu_k(dx,da)\\&\leq\int_{\mathcal{X}\times\mathcal{U}}\alpha w(x)\nu_k(dx,da)=\int_{\mathcal{X}} w(x)\alpha\nu_{k,1}(dx)\leq\alpha^{k+1}M,\end{aligned}$$

where the penultimate inequality follows from Assumption 2.3.(iii) and the last inequality by the fact that $\nu\in\Xi$. Convexity of $\Gamma(\nu)$ for $\nu\in\Xi$ is immediate, due to convexity of each of $C(\nu)$ and $B(\nu)$.

It remains to show closedness of the graph of Γ. We follow the arguments from [32, Proposition 3.9] which are in turn based on the proof of [22, Theorem 1]. Let $\Xi\times\Xi\ni(\nu^{(n)},\zeta^{(n)})\xrightarrow{n\to\infty}(\nu,\zeta)$ where $\zeta^{(n)}\in\Gamma(\nu^{(n)})$ for each $n\in\mathbb{N}$. We are done if we can show that $\zeta\in\Gamma(\nu)$. By assumption, we have

$$\zeta^{(n)}_{k+1,1}(\cdot)=\int_{\mathcal{X}\times\mathcal{U}}\rho(\cdot|x,a,\nu_k^{(n)})\nu_k^{(n)}(dx,da),$$

thus $\zeta\in C(\nu)$ follows if we can pass to the limit. Indeed, we note that because Ξ is closed, $\zeta^{(n)}\xrightarrow{n\to\infty}\zeta$ implies $\zeta^{(n)}_{k+1}\xrightarrow{n\to\infty}\zeta_{k+1}$ which implies $\zeta\in\Xi$. To pass to the limit in the right hand side of the previous display, let $g:\mathcal{X}\to\mathbb{R}$ be continuous and bounded. Then, by the regularity assumption Assumption 2.3.(iv) on ρ, we have that $(x,a)\mapsto\int_{\mathcal{X}} g(y)\rho(dy|x,a,\nu_k^{(n)})$ converges continuously to $(x,a)\mapsto\int_{\mathcal{X}} g(y)\rho(dy|x,a,\nu_k)$. Using the constant $\|g\|_\infty$ as the dominating function, we apply [33, Theorem 3.3] to conclude that

$$\int_{\mathcal{X}\times\mathcal{U}}\int_{\mathcal{X}} g(y)\rho(dy|x,a,\nu_k^{(n)})\nu_k^{(n)}(dx,da)\xrightarrow{n\to\infty}\int_{\mathcal{X}\times\mathcal{U}}\int_{\mathcal{X}} g(y)\rho(dy|x,a,\nu_k)\nu_k(dx,da)$$

hence since g was arbitrary we conclude that

$$\int_{\mathcal{X}\times\mathcal{U}}\rho(\cdot|x,a,\nu_k^{(n)})\nu_k^{(n)}(dx,da)\xrightarrow{w,n\to\infty}\int_{\mathcal{X}\times\mathcal{U}}\rho(\cdot|x,a,\nu_k)\nu_k(dx,da)$$

as desired. To establish $\zeta\in B(\nu)$, we define the maps

$$F_k^{(n)}(x,a)\coloneqq c(x,a,\nu_k^{(n)})+\int_{\mathcal{X}} V_{k+1}^{\nu^{(n)}}(y)\rho(dy|x,a,\nu_k^{(n)})$$

for $n\in\mathbb{N}$ and time-step k. We define F similarly using ν instead of $\nu^{(n)}$. By assumption, we have for a given time step k that

$$1=\zeta^{(n)}\Big\{(x,a):F_k^{(n)}(x,a)=V_k^{\nu^{(n)}}(x)\Big\}$$

for every $n \in \mathbb{N}$. We now pass this optimal-action support condition to the limit. Lemma 2.10, together with the continuity assumptions on c and ρ, gives local uniform convergence of $F_k^{(n)}$ to F_k and of $V_k^{\nu^{(n)}}$ to V_k^{ν}. Fix $\varepsilon > 0$. By tightness of $\{\zeta_k^{(n)}\}_{n\geq 1}$ and of ζ_k, choose a compact set $K \subset X \times U$ and a relatively compact open neighborhood G of K, so that $\overline{G}$ is compact, with $\zeta_k(K) > 1 - \varepsilon$ and, for all large n, $\zeta_k^{(n)}(G) > 1 - \varepsilon$. On the compact set $\overline{G}$, the above local uniform convergence implies that, for every $\delta > 0$ and all sufficiently large n, the set on which $F_k^{(n)}(x, a) = V_k^{\nu^{(n)}}(x)$ and which lies in G is contained in

$$A_{k,\delta} \coloneqq \{(x, a) : |F_k(x, a) - V_k^{\nu}(x)| \leq \delta\}.$$

Since $A_{k,\delta}$ is closed, the Portmanteau theorem [9, Theorem 2.1] gives

$$\zeta_k(A_{k,\delta}) \geq \limsup_{n\to\infty} \zeta_k^{(n)}(A_{k,\delta}) \geq 1 - \varepsilon.$$

Letting $\varepsilon \downarrow 0$ and then $\delta \downarrow 0$, and using continuity of $F_k - V_k^{\nu}$, yields

$$\zeta_k\{(x, a) : F_k(x, a) = V_k^{\nu}(x)\} = 1.$$

Thus $\zeta \in B(\nu)$. Combining this with the preceding argument, which gives $\zeta \in C(\nu)$, we have $\zeta \in C(\nu) \cap B(\nu) = \Gamma(\nu)$. Hence the graph of Γ is closed. □

3 Continuous-Time MFG Equilibria as Discrete-Time Limits

We now turn to the problem of establishing convergence of discrete-time MFG equilibria, as studied in Section 2, to continuous-time MFG equilibria for models involving jump process dynamics of controlled intensity, with interaction via the control mean, and with a finite time-horizon. Our motivation to study this problem is two-fold: first, no general existence result for such dynamics (to the best of our knowledge) exists in the literature, yet such dynamics have been used in concrete scenarios such as the cryptocurrency model in [28], with equilibria conjectured to exist from the convergence of numerical schemes. Second, the discrete-time to continuous-time convergence result provides rigorous justification for the solving of a discrete-time MFG problem via fixed-point iterations as an approximate solution to the continuous-time MFG, and we illustrate the effectiveness of this scheme by numerically solving the discrete-time version of the cryptocurrency mining mean field game studied in [28].

3.1 MFG of Controlled Intensity

Although we prove our convergence result for the relatively simple dynamics in Definition 3.2 below, the result can likely be established for much more general jump-diffusion MFG dynamics. The dynamics in Definition 3.2 are motivated by the cryptocurrency model discussed in the following sections. Our method uses the convergence methods for stochastic optimal control problems considered in detail in [24] and involves weak formulations, relaxed controls, and compactness arguments.

Definition 3.1. *(Relaxed Control) Let $\mathcal{U}$ denote a compact subset of $\mathbb{R}$. A relaxed control on $\mathcal{U}$ is a random measure $m \in \mathcal{M}(\mathcal{U} \times [0, T])$ such that, almost surely, $m(\mathcal{U} \times [0, t]) = t$ for every $t \in [0, T]$. The relaxed control is admissible w.r.t. a filtration $(\mathcal{F}_t)_{t\in[0,T]}$ iff $t \mapsto m_t(B)$ is progressively measurable w.r.t. $(\mathcal{F}_t)_{t\in[0,T]}$ for any fixed $B \in \mathcal{B}(\mathcal{U})$, where $m_t(da)dt = m(da, dt)$.*

We make use of relaxed controls because, if $(m^{(n)})_{n=1}^{\infty}$ denotes a sequence of relaxed controls on $\mathcal{U}$ (each possibly defined on its own probability space), then their process laws admit a weak limit. This follows because the subset of $\mathcal{M}(\mathcal{U} \times [0,T])$ whose second marginal is Lebesgue, endowed with the topology of weak convergence of measures, is a compact Polish space. Thus $(\mathcal{L}(m^{(n)}))_{n=1}^{\infty} \subseteq \mathcal{P}(\mathcal{M}(\mathcal{U} \times [0,T]))$ is a tight family and, by Prokhorov's theorem, admits a weakly convergent subsequence. In contrast, it would be difficult to extract a limit from the laws of arbitrary progressively measurable control processes on $[0,T]$ taking values in $\mathcal{U}$.

Definition 3.2. *(Agent Dynamics) Fix an intensity function $\lambda : \mathcal{U}^2 \to \mathbb{R}_{\geq 0}$, and constants $c > 0$ and $r > 0$. We say that state and jump processes X and N along with a $\mathcal{P}(\mathcal{U} \times \mathcal{U})$-valued process $\pi = (\pi_t)_{t \in [0,T]}$, all defined on a common filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} := (\mathcal{F}_t)_{t \in [0,T]}, P)$, satisfy the dynamics considered herein iff*

- *π is progressively measurable with respect to $\mathbb{F}$ and its second marginal $\pi_t(\mathcal{U} \times \cdot)$ is deterministic. We will write*

$$m_t(\cdot) := \pi_t(\cdot \times \mathcal{U}) \quad \text{and} \quad \eta_t(\cdot) := \pi_t(\mathcal{U} \times \cdot).$$

- *The processes satisfy*

$$X_t = X_0 - \int_0^t \int_{\mathcal{U}} ca m_s(da) ds + rN_t = X_0 - \int_0^t \int_{\mathcal{U}} ca \pi_s(da \times \mathcal{U}) ds + rN_t, \tag{3.1}$$

where N is an $\mathbb{F}$-adapted unit-jump process with stochastic intensity given by

$$\lambda_t^{\pi} := \int_{\mathcal{U}} \int_{\mathcal{U}} \lambda(\alpha, h) \pi_t(d\alpha, dh).$$

We will interpret $(m_t)_{t \in [0,T]}$ as a relaxed control process and $(\eta_t)_{t=0}^{T}$ as a relaxation of the average control of the agent population.

Remark 3.3. *(Intensity Control Representation) Unlike the more classical controlled diffusion setup where a Brownian filtration is fixed* a priori*, the jump process in Definition 3.2 cannot be fixed before the control process is specified when working with intensity control models. It is therefore convenient to work under a weak formulation (as done in [23] when establishing limit theorems for stochastic control problems), where the underlying probability space is allowed to vary with the control. We refer the reader to [11, Chapter VII.2] for more details on the formulation of jump-intensity control problems.*

By taking $\mathcal{U} = [0, L]$ for some $L > 0$ and taking $\lambda(a, h) = a/(a + Mh)$ for some large constant $M > 0$ representing the number of players, we recover a relaxed version of the dynamics considered in [28], where $(\eta_t)_{t \in [0,T]}$ is interpreted as a relaxation of the background hash-rate of the agent population (and not the distribution of controls themselves), and X denotes the representative agent wealth process. In the model from [28], a representative miner in a proof-of-work cryptocurrency network (such as Bitcoin) selects a hashing rate $a \geq 0$ to affect the likelihood of success in block discovery. Block discoveries arrive according to a Poisson process, and each successful discovery generates a reward $r > 0$. However, hashing comes at a marginal cost (of electricity) $c > 0$, hence the linear drift in (3.1). The miner's goal is to maximize expected utility φ of wealth X at time T.

The following characterization [11, T9 Theorem] (due to Watanabe) of a stochastic intensity Poisson process will be used to characterize limits of discrete-time processes.

Definition 3.4. *Let $(N_t)_{t\in[0,T]}$ be a non-explosive (see [11, Eq 1.2]) jump process with unit jumps, adapted to a filtration $\mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}$, and let $\lambda := (\lambda_t)_{t\in[0,T]}$ be a $\mathbb{F}$- progressively measurable process. Then N is a jump process with stochastic intensity λ iff*

$$N_t - \int_0^t \lambda_s \, ds, \quad t \in [0,T]$$

is a martingale.

We now give a precise definition of the weak MFG equilibria that we consider in this section.

Definition 3.5. *A relaxed MFG equilibrium of controlled jump intensity for a given intensity function $\lambda : \mathcal{U}^2 \to \mathbb{R}_{\geq 0}$, terminal reward function φ, and constants $c > 0$ and $r > 0$, is a tuple $\mathcal{T} := (\Omega, \mathcal{F}, \mathbb{F}, P, \pi, N, m, \eta, X)$ satisfying Definition 3.2, and such that:*

- *the control is optimal in the sense that for any other tuple $(\Omega', \mathcal{F}', \mathbb{F}', P', \pi', N', m', \eta', X')$ satisfying Definition 3.2 and such that $\eta_t = \eta'_t$ for Lebesgue almost every $t \in [0,T]$, we have that $E^P[\varphi(X_T)] \geq E^{P'}[\varphi(X'_T)]$;*
- *and consistency holds, which is the requirement that for Lebesgue almost every $t \in [0,T]$, we have*
$$E\big[\int_{\mathcal{U}} a\pi_t(da|h)\big] = \int_\Omega \int_{\mathcal{U}} a\pi_t(\omega)(da|h)P(d\omega) = h$$
for η_t-almost every $h \in \mathcal{U}$, where $\pi_t(\omega)(da|h)\eta_t(dh) = \pi_t(\omega)(da, dh)$ (we use slightly different disintegration notation here due to the presence of time subscripts).

If η_t happens to be a Dirac probability measure for Lebesgue almost every $t \in [0,T]$, then we say that the MFG equilibrium is sharp. If for almost every $\omega \in \Omega$, the measure $\pi_t(\omega)(\cdot|h)$ is a Dirac for Lebesgue almost every $t \in [0,T]$ and η_t-almost every h, we say that the MFG equilibrium is of sharp controls.

Here h is the relaxed population-average control coordinate, while a is the representative agent's relaxed control conditional on h. Thus η_t is not the law of individual controls but a relaxation of the average population control. If η_t is non-Dirac, it represents chattering of the population average among several values.

In the event where η_t is a Dirac mass at $h(t) \in \mathcal{U}$, then $h(t)$ represents the average population control and we recover $E\big[\int a m_t(da)\big] = h(t)$, which further reduces to $E[a_t] = h(t)$ in the case where $m_t = \delta_{\{a_t\}}$ is a sharp control. Note also that integrating both sides of (3.5) we obtain

$$\overline{\eta}_t = \int_{\mathcal{U}} h\eta_t(dh) = \int_{\mathcal{U}} E\big[\int_{\mathcal{U}} a\pi_t(da|h)\big]\eta_t(dh) = E\Big[\int a\pi_t(da \times \mathcal{U})\Big] = E[\overline{m}_t],$$

as expected.

3.2 Time-Discretized MFG & Convergence Theorem

We aim to establish an MFG existence result for the above setup as a limit of discrete-time MFG equilibria. We use the following natural (due to the Poisson limit theorem) discrete-time approximations, indexed by $n \in \mathbb{N}$.

Definition 3.6. *(Discretization Scheme) Fix a bounded non-negative map $\lambda : \mathcal{U}^2 \to \mathbb{R}_{\geq 0}$. For $n \in \mathbb{N}$ large enough so that $\|\lambda\|_\infty / 2^n \leq 1$, the discrete-time state process $(x_k^{(n)})_{k=0}^{2^n T}$ follows*

$$x_{k+1}^{(n)} \sim \alpha_k^{(n)} \delta_{\{x_k^{(n)} - (c a_k^{(n)}/2^n) + r\}}(\cdot) + \big(1 - \alpha_k^{(n)}\big) \delta_{\{x_k^{(n)} - (c a_k^{(n)}/2^n)\}}(\cdot),$$

where, conditional on $(x_k^{(n)}, a_k^{(n)})$, the random variable $x_{k+1}^{(n)}$ is independent of all other random quantities up to time k. We allow for two schemes by taking either:

$$\textit{Scheme 1:} \quad \alpha_k^{(n)} = \frac{1}{2^n} \int_{\mathcal{U}} \lambda(a_k^{(n)}, h) \eta_k^{(n)}(dh), \quad \textit{Scheme 2:} \quad \alpha_k^{(n)} = \frac{1}{2^n} \lambda(a_k^{(n)}, \overline{\eta}_k^{(n)}).$$

It is assumed that $x_0^{(n)} \sim \mu_0$. The optimality criterion is the maximization of $E[\varphi(x_{2^n T}^{(n)})]$.

In either scheme, $(\eta_k^{(n)})_{k=0}^{2^n T} \subseteq \mathcal{P}(\mathcal{U})$, which we shall refer to as the parametrizing sequence, is a fixed but arbitrary deterministic sequence, and $a_k^{(n)}$ is the $\mathcal{U}$-valued control at time k. We will later apply Theorem 2.4 to the transition kernel defined by the discretization schemes, thus we will have $a_k^{(n)} \sim \pi_k^{(n)}$ where $(\pi_k^{(n)})_{k=0}^{2^n T}$ is a sequence of random probability measures on $\mathcal{U}$, adapted to the state process filtration. More generally, we say that:

Definition 3.7. *(Discrete-time Admissibility) The state/control pair $(x^{(n)}, a^{(n)})$ is admissible for a given scheme if it is defined on some common probability space equipped with a filtration $(\mathcal{F}_k)_{k=0}^{2^n T}$ to which the processes are adapted, and, conditional on $\mathcal{F}_k$, the law of $x_{k+1}^{(n)}$ is given by the transition kernel (for the given scheme) from Definition 3.6, for every $k = 0, 1, \ldots, 2^n T - 1$. Of course, we also impose that $x_0^{(n)} \sim \mu_0$.*

Remark 3.8. *(Relaxed Discrete-Time Controls) When relating optimality of discrete-time models to continuous-time ones, it will be useful to allow for measure-valued controls in Definition 3.6. For a $\mathcal{P}(\mathcal{U})$-valued relaxed control process $m_k^{(n)}$, the dynamics will follow*

$$x_{k+1}^{(n)} \sim \alpha_k^{(n)} \delta_{\{x_k^{(n)} - c 2^{-n} \int a m_k^{(n)}(da) + r\}}(\cdot) + \big(1 - \alpha_k^{(n)}\big) \delta_{\{x_k^{(n)} - c 2^{-n} \int a m_k^{(n)}(da)\}}(\cdot),$$

with

$$\alpha_k^{(n)} = \frac{1}{2^n} \int_{\mathcal{U}} \int_{\mathcal{U}} \lambda(a, h) m_k^{(n)}(da) \eta_k^{(n)}(dh) \quad \textit{or} \quad \alpha_k^{(n)} = \frac{1}{2^n} \int_{\mathcal{U}} \lambda(a, \overline{\eta}_k^{(n)}) m_k^{(n)}(da)$$

in Scheme 1 and Scheme 2, respectively. Note that in this case, randomized and relaxed controls are fundamentally distinct.

The convergence theorem will hold under different subsets of the following assumptions.

Assumption 3.9.

1. *$\mathcal{U}$ is a closed interval, and the intensity function λ is continuous in both of its variables, and hence uniformly continuous, and bounded given that its domain is compact.*
2. *μ_0 is compactly supported.*
3. *φ is non-decreasing, continuous and bounded. Also, $\lambda(\cdot, h)$ is concave, and Lipschitz uniformly over the choice of $h \in \mathcal{U}$.*
4. *φ is strictly increasing and $\lambda(\cdot, h)$ is strictly concave for fixed $h \in \mathcal{U}$.*

Before stating the convergence theorem, we show that there is no gain (in terms of the optimality criteria) of enlarging from ordinary to relaxed controls in the discrete-time setup.

Lemma 3.10. *Suppose that 1, 2, and 3 from Assumption 3.9 hold. Then, under either Scheme 1 or Scheme 2, relaxed controls do not improve the value function.*

Proof. By induction, it suffices to establish the claim for the penultimate time-step $k = 2^nT - 1$. Take a fixed state value $x \in \mathbb{R}$ and let $m_k^{(n)}$ denote the measure-valued control at time step k. Taking some $m \in \mathcal{P}(\mathcal{U})$ and recalling that φ is assumed non-decreasing, thus

$$
\begin{aligned}
&E[\varphi(x_{k+1}^{(n)})|x_k^{(n)} = x, m_k^{(n)} = m] \\
&= \Big[\frac{1}{2^n}\int_{\mathcal{U}}\int_{\mathcal{U}} \lambda(a,h)m(da)\eta_k^{(n)}(dh)\Big]\Big[\varphi(x + r - 2^{-n}c\overline{m}) - \varphi(x - 2^{-n}c\overline{m})\Big] + \varphi(x - 2^{-n}c\overline{m}) \\
&\leq \Big[\frac{1}{2^n}\int_{\mathcal{U}} \lambda(\overline{m},h)\eta_k^{(n)}(dh)\Big]\Big[\varphi(x + r - 2^{-n}c\overline{m}) - \varphi(x - 2^{-n}c\overline{m})\Big] + \varphi(x - 2^{-n}c\overline{m}) \\
&= E[\varphi(x_{k+1}^{(n)})|x_k^{(n)} = x, m_k^{(n)} = \delta_{\{\overline{m}\}}],
\end{aligned}
$$

thus given a measure-valued control action $m \in \mathcal{P}(\mathcal{U})$, one never decreases the expected terminal utility by replacing it with the ordinary control $\overline{m} \in \mathcal{U}$. Inductively, it is straightforward to see that the value functions at arbitrary time steps are always non-decreasing, thus the above argument holds for any $k \in \{0, 1, \dots, 2^nT - 1\}$. □

We now state and prove the main result from this section.

Theorem 3.11. *Suppose that 1, 2, and 3 from Assumption 3.9 hold. Then there exists an MFG equilibrium in the sense of Definition 3.5 for the continuous-time game of controlled intensity. If additionally 4 holds, then the equilibrium is of sharp control.*

The proof is organized in four steps. In step 1, we interpolate the discrete-time controlled Bernoulli chains (Definition 3.6) to obtain continuous-time processes, establish that the sequence of joint state/control laws of the interpolations admits a weak limit, and extract a limiting flow of control measure. In step 2, we establish that the limiting law corresponds to a controlled jump process with the dynamics specified in Definition 3.5. In step 3, consistency with respect to the flow of measure obtained from step 2 is established. In step 4, we establish optimality. We then conclude with a brief note regarding the last assertion of the theorem.

Step 1 (Discrete-Time Approximations): For simplicity and without loss of generality, let the finite time horizon T of the continuous-time model be an integer. Consider the MFG arising from the discretization introduced in Definition 3.6 under Scheme 2, where we use $y^{(n)}$ to denote the state process, and where we always assume n is sufficiently large so that the discretizations are well-defined. Applying Theorem 2.4, we obtain an MFG equilibrium $((a_k^{(n)})_{k=0}^{T2^n}, (\zeta_k^{(n)})_{k=0}^{T2^n})$, where the ζ's are elements of $\mathcal{P}(\mathcal{U})$ satisfying $\mathcal{L}(a_k^{(n)}) = \zeta_k^{(n)}$ for every k, and where $(a_k^{(n)})_{k=0}^{T2^n}$ is the optimal control process for the fixed sequence $(\zeta_k^{(n)})_{k=0}^{T2^n}$ with respect to the maximization of $E[\varphi(y_{2^nT}^{(n)})]$.

We now define new discretizations under Definition 3.6 Scheme 1 with state process denoted by $x^{(n)}$ and with parametrizing sequence $\eta_k^{(n)} := \delta_{\{\overline{\zeta}_k^{(n)}\}}, k = 0, 1, \cdots, 2^nT$. Here, $x^{(n)}$ will be the discrete-time dynamics to be interpolated. We will at times slightly abuse notation by denoting the continuous-time objects with the same letters as their discrete-time counterparts, but will differentiate using time index t and k when referring to the continuous-time and discrete-time processes, respectively. For each $n \in \mathbb{N}$, we construct continuous-time processes $(X_t^{(n)})_{t\in[0,T]}$, $(a_t^{(n)})_{t\in[0,T]}$, and $(\eta_t^{(n)})_{t\in[0,T]}$ as follows. For each $t \in [0,T]$:

$$X_t^{(n)} := x_k^{(n)} - ca_k^{(n)}\Big(t - \frac{k}{2^n}\Big), \quad \eta_t^{(n)} := \eta_k^{(n)}, \quad \text{and} \quad a_t^{(n)} := a_k^{(n)} \quad \text{for} \quad t \in \Big[\frac{k}{2^n}, \frac{k+1}{2^n}\Big).$$

In words, $X^{(n)}$ linearly interpolates the drift, with a jump at multiples of $1/2^n$ whenever the discrete-time process jumps up. By construction, sample paths of $X^{(n)}$ are right-continuous with left limits for every $n \in \mathbb{N}$. By defining $m_t^{(n)} := \delta_{\{a_t^{(n)}\}}$ we obtain a relaxed control representation of the control process $a_t^{(n)}$. The following representation (which defines $N_t^{(n)}$) will be very useful:

$$X_t^{(n)} = X_0^{(n)} - c\int_0^t a_s^{(n)}ds + rN_t^{(n)} = x_0^{(n)} - c\int_0^t \int_{\mathcal{U}} am_s^{(n)}(da)ds + rN_t^{(n)}, \tag{3.2}$$

where $N_t^{(n)}$ is a jump process taking values in $\{0, 1, 2, 3, \cdots\}$, with unit jumps possible only on dyadic rationals of order n. Note that one can give an explicit definition of $N^{(n)}$ as

$$N_0^{(n)} = 0, \quad N_k^{(n)} = N_{k-1}^{(n)} + \mathbf{1}_{\left\{x_k^{(n)} = x_{k-1}^{(n)} - \frac{c}{2^n}a_{k-1}^{(n)} + r\right\}} \quad \text{for} \quad k \geq 1$$

and interpolate exactly as done for $a^{(n)}$ and $\eta^{(n)}$ in (3.2) to define the continuous-time process $(N_t^{(n)})_{t\in[0,T]}$ in (3.2). Note also that the random variables used to define the above continuous processes are a countable family and we assume them to be defined on a common probability space $(\Omega, \mathcal{F}, P)$. Finally, we define the process

$$\pi_t^{(n)} = m_t^{(n)} \otimes \eta_t^{(n)}.$$

Next, we extract limit points from the joint process laws $\mathcal{L}(X^{(n)}, \pi^{(n)}, N^{(n)})$ as $n \to \infty$. Note that the processes $X^{(n)}$ naturally take values in the space $D[0,T]$ (right-continuous functions with left limits on $[0,T]$ taking values in $\mathbb{R}$) which can be endowed with the Skorokhod J1 topology. Together with this topology, the space $D[0,T]$ is a Polish space. For a collection of probability measures on the Borel sets of a Polish space, tightness is equivalent to sequential compactness (i.e. every sequence of measures from the collection admits a further weakly convergent sub-sequence) as per Prokhorov's theorem. Establishing tightness will be made easy due to [24, Theorem 9.2.1], which is restated in the Appendix for convenience. Because the drift term is linear and the jump probability is scaled with n, it is almost immediate that the sequences of laws $\mathcal{L}(X^{(n)})_{n\in\mathbb{N}}$ and $\mathcal{L}(N^{(n)})_{n\in\mathbb{N}}$ satisfy the referenced tightness condition and are thus tight. Tightness of the laws $\mathcal{L}(\pi^{(n)})_{n=1}^{\infty}$ is immediate (see the discussion following Definition 3.1 of relaxed controls) since we have assumed a compact action space $\mathcal{U}$.

As we have checked tightness of each of the sequence of marginals, we conclude that

$$\mathcal{L}(X^{(n)}, \pi^{(n)}, N^{(n)})_{n=1}^{\infty}$$

is tight and admits a weakly converging sub-sequence. For simplicity of notation, we will dispense with the sub-sequence and assume that the convergence is as $n \to \infty$.

By the Skorokhod representation theorem (see [24, Theorem 9.1.7]), there exists some probability space $(\Omega, \mathcal{F}, P)$ supporting random variables $(\tilde{X}^{(n)}, \tilde{\pi}^{(n)}, \tilde{N}^{(n)})_{n=1}^{\infty}$ converging almost surely to (X, π, N) and such that $\mathcal{L}(\tilde{X}^{(n)}, \tilde{\pi}^{(n)}, \tilde{N}^{(n)}) = \mathcal{L}(X^{(n)}, \pi^{(n)}, N^{(n)})$. It is immediate that

$(\tilde{X}^{(n)}, \tilde{\pi}^{(n)}, \tilde{N}^{(n)})$ satisfies the representation (3.2), and hereafter we abuse notation by dropping tilde.

To summarize, we have that for P-almost every $\omega \in \Omega$, $X^{(n)}(\omega) \xrightarrow{n\to\infty} X(\omega)$ w.r.t. the Skorokhod topology on $D[0,T]$, $\pi^{(n)}(\omega) \xrightarrow{n\to\infty} \pi(\omega)$ in the weak topology, and redundantly, $N^{(n)}(\omega) \xrightarrow{n\to\infty} N(\omega)$ in the Skorokhod topology on $D[0,T]$. We define $m_t(\cdot) = \pi_t(\cdot \times \mathcal{U})$ and $\eta_t(\cdot) = \pi_t(\mathcal{U} \times \cdot)$ and also set

$$\mathcal{F}_t \coloneqq \sigma(X_s, \pi_s, N_s : s \le t), \quad \mathcal{F}_t^{(n)} \coloneqq \sigma(X_s^{(n)}, \pi_s^{(n)}, N_s^{(n)} : s \le t) \quad \text{for every } t \in [0,T], n \in \mathbb{N}.$$

It is important to note that we can always define a derivative $\pi_t(\omega)$ such that $\pi_t(\omega)(A)$ is $(\mathcal{F}_t)_{t\in[0,T]}$ adapted for each $t \in [0,T]$ and $A \in \mathcal{B}(\mathcal{U} \times \mathcal{U})$, and such that the disintegration $\pi(\omega)(da, dh, dt) = \pi_t(\omega)(da, dh)dt$ holds [24, Section 9.5]. Similarly we can also do this for approximations, with the Borel measurability of maps of the form $t \mapsto E[\pi_t(A \times \mathcal{U})] = \int \pi_t(\omega)(A \times \mathcal{U})dP(\omega)$ for $A \in \mathcal{B}(\mathcal{U})$ following from progressive measurability (see for example [15, Theorem 3.1]).

Step 2 (Characterizing the Limit Point): Using the almost sure convergence, it follows that for any continuity point $t \in [0,T]$ of the limit $X(\omega)$, (X, m, N) satisfies

$$\begin{aligned} X_t(\omega) = \lim_{n\to\infty} X_t^{(n)}(\omega) &= \lim_{n\to\infty} \Big(X_0^{(n)}(\omega) - c\int_0^t \int_{\mathcal{U}} a m_s^{(n)}(\omega)(da)ds + rN_t^{(n)}(\omega) \Big) \\ &= X_0(\omega) - c\int_0^t \int_{\mathcal{U}} a m_s(\omega)(da)ds + rN_t(\omega). \end{aligned}$$

Since the common continuity points of $X(\omega)$ and $N(\omega)$ are dense (since sample paths are cadlag), and both sides define cadlag functions of t, the equality holds on the entire interval $[0,T]$. Clearly, N is the law of a unit jump process. The fact that it is in fact an $\mathbb{F}$-jump process follows by arguments as in [24, Equation 10.1.8]. We verify that it has the correct (stochastic) intensity by fixing $0 \le s < t$ and denoting

$$J_{s,t} = \int_s^t \int_{\mathcal{U}} \int_{\mathcal{U}} \lambda(a,h)\pi_\rho(da, dh)d\rho.$$

We now verify the martingale property. Define

$$M_t \coloneqq N_t - J_{0,t} = N_t - \int_0^t \int_{\mathcal{U}^2} \lambda(a,h)\pi_\rho(da, dh)d\rho.$$

It suffices to show that for every bounded $\mathcal{F}_s$-measurable random variable H,

$$E[H(M_t - M_s)] = 0.$$

We first verify this identity for bounded continuous functionals of the stopped path. Let

$$\mathcal{E}_s \coloneqq D([0,s]; \mathbb{R}) \times \mathcal{R}_s \times D([0,s]; \mathbb{N}),$$

where $\mathcal{R}_s$ denotes the relaxed-measure space on $\mathcal{U}^2 \times [0,s]$ with Lebesgue time marginal. Let $F : \mathcal{E}_s \to \mathbb{R}$ be bounded and continuous. By the discrete-time construction,

$$E\left[F\left(X_{\cdot\wedge s}^{(n)}, \pi^{(n)}|_{[0,s]}, N_{\cdot\wedge s}^{(n)}\right)\left(N_t^{(n)} - N_s^{(n)} - \int_s^t \int_{\mathcal{U}^2} \lambda(a,h)\pi_\rho^{(n)}(da, dh)d\rho\right)\right] \xrightarrow{n\to\infty} 0, \quad (3.3)$$

with the difference between the discrete compensator sum and the displayed integral vanishing as the mesh-size goes to zero. By the almost sure convergence from the Skorokhod representation and the continuity of F, we have

$$F\left(X^{(n)}_{\cdot\wedge s}, \pi^{(n)}|_{[0,s]}, N^{(n)}_{\cdot\wedge s}\right) \to F\left(X_{\cdot\wedge s}, \pi|_{[0,s]}, N_{\cdot\wedge s}\right)$$

almost surely. Moreover, since λ is bounded and continuous and $\pi^{(n)} \Rightarrow \pi$ as relaxed measures,

$$\int_s^t \int_{\mathcal{U}^2} \lambda(a,h)\pi^{(n)}_\rho(da,dh)d\rho \to \int_s^t \int_{\mathcal{U}^2} \lambda(a,h)\pi_\rho(da,dh)d\rho$$

almost surely. Uniform integrability follows from the boundedness of λ, thus passing to the limit in (3.3) gives

$$E\left[F\left(X_{\cdot\wedge s}, \pi|_{[0,s]}, N_{\cdot\wedge s}\right)(M_t - M_s)\right] = 0.$$

Since bounded continuous functions on $\mathcal{E}_s$ generate its Borel σ-algebra, and since $\mathcal{F}_s$ is generated by the stopped path

$$(X_{\cdot\wedge s}, \pi|_{[0,s]}, N_{\cdot\wedge s}),$$

it follows from a monotone class argument that $E[H(M_t - M_s)] = 0$ for every bounded $\mathcal{F}_s$-measurable random variable H and hence M is an $\mathbb{F}$-martingale. We thus conclude that N has stochastic intensity

$$\lambda^\pi_t = \int_{\mathcal{U}^2} \lambda(a,h)\pi_t(da,dh).$$

We thus conclude that the limiting processes satisfy the controlled jump dynamics in Definition 3.2.

Step 3 (Consistency): We now verify consistency of the limit point with respect to the deterministic relaxation η of the average control, which we state as a Lemma.

Lemma 3.12. *(Consistency of the limit) For Lebesgue-almost every $t \in [0,T]$ we have that $E[\int_{\mathcal{U}} a\pi_t(da|h)] = h$ for every h in the support of $\eta_t(\cdot) = \pi_t(\mathcal{U} \times \cdot)$.*

Proof. Recall that for each n, by construction,

$$\pi^{(n)}_t = m^{(n)}_t \otimes \eta^{(n)}_t, \qquad \eta^{(n)}_t = \delta_{\{\bar\zeta^{(n)}_k\}} \quad \text{for } t \in [k2^{-n}, (k+1)2^{-n}),$$

and

$$\bar\zeta^{(n)}_k = \int_{\mathcal{U}} a\,\zeta^{(n)}_k(da) = E[a^{(n)}_k].$$

Since $m^{(n)}_t = \delta_{\{a^{(n)}_t\}}$, it follows that for every bounded continuous test function $\varphi : [0,T] \times \mathcal{U} \to \mathbb{R}$,

$$\begin{aligned}
E\left[\int_0^T \int_{\mathcal{U}\times\mathcal{U}} \varphi(t,h)a\,\pi^{(n)}_t(da,dh)\,dt\right] &= E\left[\int_0^T \varphi\left(t, \bar\zeta^{(n)}_{\lfloor 2^n t\rfloor}\right) a^{(n)}_t\,dt\right] \\
&= \int_0^T \varphi\left(t, \bar\zeta^{(n)}_{\lfloor 2^n t\rfloor}\right) E[a^{(n)}_t]\,dt \\
&= \int_0^T \int_{\mathcal{U}} \varphi(t,h)h\,\eta^{(n)}_t(dh)\,dt.
\end{aligned}$$

We now pass to the limit as $n \to \infty$ in the above expression to obtain

$$E\left[\int_0^T \int_{\mathcal{U}\times\mathcal{U}} \varphi(t,h) a\, \pi_t(da,dh)\, dt\right] = \int_0^T \int_{\mathcal{U}} \varphi(t,h) h\, \eta_t(dh)\, dt \tag{3.4}$$

for every bounded continuous $\varphi : [0,T]\times\mathcal{U} \to \mathbb{R}$. Disintegrating $\pi_t(\omega)$ with respect to its deterministic second marginal, (3.4) becomes

$$\int_0^T \int_{\mathcal{U}} \varphi(t,h) \left(E\left[\int_{\mathcal{U}} a\, \pi_t(da \mid h)\right] - h \right) \eta_t(dh)\, dt = 0.$$

By the arbitrariness of φ, it follows that

$$E\left[\int_{\mathcal{U}} a\, \pi_t(da \mid h)\right] = h \qquad (dt \otimes \eta_t(dh))\text{-a.e. on } [0,T]\times\mathcal{U}.$$

Equivalently, for Lebesgue-a.e. $t \in [0,T]$,

$$h = E\left[\int_{\mathcal{U}} a\, \pi_t(da \mid h)\right] \qquad \text{for } \eta_t\text{-a.e. } h.$$

This establishes the desired consistency condition. □

Step 4 (Establishing Optimality): For a fixed positive discretization integer n and sequence of measures $\eta^{(n)} := (\eta_k^{(n)})_{k=0}^{2^nT} \subseteq \mathcal{P}(\mathcal{U})$ parametrizing the dynamics, the expected reward of an admissible (in the sense of Definition 3.7) control/state pair $a^{(n)} := (a_k^{(n)})_{k=0}^{2^nT}$, $x^{(n)} := (x_k^{(n)})_{k=0}^{2^nT}$ under Scheme 1 dynamics in Definition 3.6 is denoted by

$$w^{(n)}(k,x,a^{(n)},\eta^{(n)}) := E[\varphi(x_{2^nT}^{(n)})|x_k^{(n)} = x],$$

and the value function is denoted by

$$v^{(n)}(k,x,\eta^{(n)}) := \sup_{a^{(n)},x^{(n)}} w^{(n)}(k,x,a^{(n)},\eta^{(n)})$$

where the supremum is taken over admissible control/state pairs, and where the underlying probability space may vary with the control/state. We use capital letters for the analogous continuous-time quantities, and define

$$W(t,x,\pi) := E[\varphi(X_T^{\pi})|X_t = x]$$

where π and X^{π} (along with an unnamed jump process, probability space, and filtration) satisfy Definition 3.2. Moreover, for a given flow of control measure $\eta := (\eta_t)_{t\in[0,T]} \subseteq \mathcal{P}(\mathcal{U})$, we define

$$V(t,x,\eta) = \sup_{\pi:\pi_s(\mathcal{U}\times\cdot)=\eta_s(\cdot),s\geq t} W(t,x,\pi),$$

where we allow the underlying probability space to vary with π in the supremum. The following lemma, inspired by the construction in [23, Theorem 3.5.2], will allow us to conclude that, given an approximating sequence $\eta^{(n)} \xrightarrow{\mathcal{L},n\to\infty} \eta$, one can construct processes $\pi^{(n)}$ with deterministic second marginal $\eta^{(n)}$ such that $W(t,x,\pi^{(n)}) \xrightarrow{n\to\infty} V(t,x,\eta)$. This will allow us to make the connection between discrete-time and continuous-time value functions and is a vital ingredient in proving the optimality of the limit process extracted from discrete-time chains.

Lemma 3.13 (Joint chattering lemma). *Consider the control problem from Definition 3.2, fix a map $\lambda : \mathcal{U}^2 \to \mathbb{R}_{\geq 0}$ and a deterministic flow of measures $(\eta_t)_{t\in[0,T]} \subseteq \mathcal{P}(\mathcal{U})$, and assume that we can represent the latter as a limit*

$$\eta^{(n)}(dh, dt) := \eta_t^{(n)}(dh)dt \Rightarrow \eta(dh, dt) := \eta_t(dh)dt$$

where $\eta_t^{(n)} = \delta_{\{f^n(t)\}}$ is a Dirac for every $t \in [0, T]$ and is constant on dyadic intervals of the form $[k2^{-n}, (k + 1)2^{-n})$. Let $(\Omega, \mathcal{F}, \mathbb{F}, P, \pi, N, m, \eta, X)$ be an arbitrary admissible tuple (not necessarily the limit of discrete-time processes obtained in Step 2) in the sense of Definition 3.2, and assume 1, 2, *and* 3 *from Assumption 3.9 hold. Then we may construct, on the same filtered space $(\Omega, \mathcal{F}, P, \mathbb{F})$, a process $(\pi_t^{(n,m)})_{t\in[0,T]}$ which is progressively measurable w.r.t. $\mathbb{F}$ and such that*

- $\pi_t^{(n,m)}(\mathcal{U} \times \cdot) = \eta_{t-2^{-m}}^{(n)}(\cdot)$ *for every $t \in [2^{-m}, T]$,*
- *for P-almost every $\omega \in \Omega$ the measure $\pi_t^{(n,m)}(\omega) \in \mathcal{P}(\mathcal{U} \times \mathcal{U})$ decomposes into a product structure for Lebesgue almost every t,*
- *there exists subsequences $(n_l)_{l=1}^{\infty}$ and $(m_l)_{l=1}^{\infty}$ along which, for any $0 \leq s < p \leq T$, we have that*
$$\int_s^p \int_{\mathcal{U}^2} \lambda(a, h)\, \pi_t^{(n_l,m_l)}(\omega)(da, dh)\, dt \xrightarrow{l\to\infty} \int_s^p \int_{\mathcal{U}^2} \lambda(a, h)\, \pi_t(\omega)(da, dh)\, dt$$
P-almost surely, with the convergence uniform over the set of ω's on which it occurs.

Proof. For simplicity, we will assume that s and p are dyadic rationals. At the end, we will discuss how the arguments generalize.
Step 1: Time and space discretization. For a given $m \in \mathbb{N}$ denote the dyadic intervals by

$$I_k^m := [k2^{-m}, (k + 1)2^{-m}), \qquad k = 0, \dots, T2^m - 1,$$

and for such a k define the averaged (random) measure

$$\pi_k^{(m)}(A \times B) := 2^m \int_{I_k^m} \pi_t(A \times B)\, dt, \qquad A, B \in \mathcal{B}(\mathcal{U}).$$

Then
$$\int_0^T \int_{\mathcal{U}^2} \lambda(a, h)\, \pi_t(da, dh)\, dt = 2^{-m} \sum_{k=0}^{2^m T-1} \int_{\mathcal{U}^2} \lambda(a, h)\, \pi_k^{(m)}(da, dh).$$

To discretize space, fix $\delta > 0$ and choose a finite partition of the closed interval $\mathcal{U}$

$$\mathcal{U} = \bigsqcup_{j=1}^{J} C_j$$

of Borel sets of diameter at most δ, along with representative point $h_j \in C_j$ for each j. For each k, j, define the deterministic (since the second marginal of π_k is deterministic) quantities

$$w_{k,j}^{(m)} := \pi_k^{(m)}(\mathcal{U} \times C_j) = 2^m \int_{I_k^m} \eta_t(C_j)dt.$$

If $w_{k,j}^{(m)} > 0$, define a (random) probability measure $\alpha_{k,j}^{(m)} \in \mathcal{P}(\mathcal{U})$ by

$$\alpha_{k,j}^{(m)}(A) := \frac{\pi_k^{(m)}(A \times C_j)}{w_{k,j}^{(m)}}, \qquad A \in \mathcal{B}(\mathcal{U})$$

and if $w_{k,j}^{(m)} = 0$, choose $\alpha_{k,j}^{(m)}$ arbitrarily. Since λ is uniformly continuous on the compact set $\mathcal{U}^2$, it admits a modulus of continuity ω_λ, thus for $h \in C_j$ we have $|\lambda(a,h) - \lambda(a,h_j)| \le \omega_\lambda(\delta)$ and it follows that surely

$$\left| \int_{\mathcal{U}^2} \lambda(a,h)\, \pi_k^{(m)}(da,dh) - \sum_{j=1}^{J} w_{k,j}^{(m)} \int_{\mathcal{U}} \lambda(a,h_j)\, \alpha_{k,j}^{(m)}(da) \right| \le \omega_\lambda(\delta).$$

Multiplying by 2^{-m}, summing over k, and recalling (3.2), we obtain

$$\left| \int_0^T \int_{\mathcal{U}^2} \lambda(a,h)\, \pi_t(da,dh)\, dt - 2^{-m} \sum_{k=0}^{T2^m-1} \sum_{j=1}^{J} w_{k,j}^{(m)} \int_{\mathcal{U}} \lambda(a,h_j)\, \alpha_{k,j}^{(m)}(da) \right| \le T\, \omega_\lambda(\delta). \tag{3.5}$$

Step 2: Chattering approximation of the second marginal. For integers $m < n$, define the (deterministic) sets

$$A_{k,j}^{(n,m)} := \{t \in I_k^m : \ f^{(n)}(t) \in C_j\}.$$

Then

$$|A_{k,j}^{(n,m)}| = \int_{I_k^m} \mathbf{1}_{\{f^{(n)}(t) \in C_j\}}\, dt = \int_{[0,T]\times\mathcal{U}} \mathbf{1}_{I_k^m}(t) \mathbf{1}_{C_j}(h)\, \eta^{(n)}(dh,dt).$$

By the weak convergence $\eta^{(n)}(dh,dt) \Rightarrow \eta(dh,dt)$, and after choosing the partition so that the boundaries of the sets $C_j \times I_k^m$ are $\eta(dh,dt)$-null, we obtain for a fixed $m \in \mathbb{N}$ that

$$|A_{k,j}^{(n,m)}| \xrightarrow{n\to\infty} \int_{I_k^m} \eta_t(C_j)\, dt = 2^{-m} w_{k,j}^{(m)}.$$

Define now the (random) measure

$$\tilde{\pi}_t^{(n,m)}(da,dh) := \sum_{k=0}^{T2^m-1} \sum_{j=1}^{J} \mathbf{1}_{A_{k,j}^{(n,m)}}(t)\, [\alpha_{k,j}^{(m)}(da) \otimes \delta_{\{f^{(n)}(t)\}}(dh)].$$

Equivalently, if $t \in I_k^m$ and $f^{(n)}(t) \in C_j$, then

$$\tilde{\pi}_t^{(n,m)}(da,dh) = \alpha_{k,j}^{(m)}(da) \otimes \delta_{\{f^{(n)}(t)\}}(dh).$$

Note that because $f^{(n)}$ are cadlag and piecewise constant on dyadics by assumption, the process $\tilde{\pi}^{(n,m)}$ is surely cadlag. By construction, the second marginal is exactly

$$\tilde{\pi}_t^{(n,m)}(\mathcal{U} \times dh) = \delta_{\{f^{(n)}(t)\}}(dh) = \eta_t^{(n)}(dh).$$

Step 3: Convergence of the intensity integral. Let $0 \le s < p \le T$ and assume for the moment that s and p are multiples of 2^{-q} for a positive integer q. Taking m and n larger than q, we have

$$\int_s^p \int_{\mathcal{U}^2} \lambda(a,h)\, \tilde{\pi}_t^{(n,m)}(da,dh)\, dt = \sum_{k,j: I_k^{(m)} \subseteq [s,p]} \int_{A_{k,j}^{(n,m)}} \int_{\mathcal{U}} \lambda(a, f^{(n)}(t))\, \alpha_{k,j}^{(m)}(da)\, dt.$$

Since $f^{(n)}(t) \in C_j$ on $A_{k,j}^{(n,m)}$, we also have

$$\left| \int_{\mathcal{U}} \lambda(a, f^{(n)}(t))\, \alpha_{k,j}^{(m)}(da) - \int_{\mathcal{U}} \lambda(a, h_j)\, \alpha_{k,j}^{(m)}(da) \right| \le \omega_\lambda(\delta).$$

Hence

$$\left| \int_s^p \int_{\mathcal{U}^2} \lambda(a,h)\, \tilde{\pi}_t^{(n,m)}(da,dh) dt - \sum_{\substack{k,j: \\ I_k^{(m)} \subseteq [s,p]}} |A_{k,j}^{(n,m)}| \int_{\mathcal{U}} \lambda(a, h_j)\, \alpha_{k,j}^{(m)}(da) \right| \le (p-s)\, \omega_\lambda(\delta).$$

Letting $n \to \infty$ with m fixed yields

$$\sum_{\substack{k,j: \\ I_k^{(m)} \subseteq [s,p]}} |A_{k,j}^{(n,m)}| \int_{\mathcal{U}} \lambda(a, h_j)\, \alpha_{k,j}^{(m)}(da) \xrightarrow{n\to\infty} \sum_{\substack{k,j: \\ I_k^{(m)} \subseteq [s,p]}} 2^{-m} w_{k,j}^{(m)} \int_{\mathcal{U}} \lambda(a, h_j)\, \alpha_{k,j}^{(m)}(da).$$

Combining the previous two equations with a modification of the bound in (3.5) for an integral from s to p, we obtain

$$\limsup_{n\to\infty} \left| \int_s^p \int_{\mathcal{U}^2} \lambda(a,h) [\tilde{\pi}_t^{(n,m)} - \pi_t](da,dh) dt \right| \le 2(p-s)\omega_\lambda(\delta).$$

Now let $\delta \downarrow 0$. Since λ is uniformly continuous, $\omega_\lambda(\delta) \to 0$. Thus, for each fixed m, one may choose $n = n(m)$ sufficiently large, and a $\delta(m)$ sufficiently small, so that

$$\left| \int_s^p \int_{\mathcal{U}^2} \lambda(a,h)\, \tilde{\pi}_t^{(n(m),m)}(da,dh)\, dt - \int_s^p \int_{\mathcal{U}^2} \lambda(a,h)\, \pi_t(da,dh)\, dt \right| \longrightarrow 0$$

as $m \to \infty$.

Step 4: Shift to preserve adaptedness. Finally, define

$$\pi_t^{(n(m),m)} := \begin{cases} \tilde{\pi}_{t-2^{-m}}^{(n(m),m)} & \text{for} \quad t \in [2^{-m}, T] \\ \delta_{\{0\}} \otimes \eta_t^{(n(m))} & \text{otherwise.} \end{cases}$$

Then for $t \ge 2^{-m}$

$$\pi_t^{(n(m),m)}(\mathcal{U} \times dh) = \delta_{f^{(n(m))}(t-2^{-m})}(dh),$$

and by construction $\pi_t^{(n(m),m)}$ is a product measure. Also,

$$\int_s^p \int_{\mathcal{U}^2} \lambda(a,h)\, \pi_t^{(n(m),m)}(da,dh)\, dt \longrightarrow \int_s^p \int_{\mathcal{U}^2} \lambda(a,h)\, \pi_t(da,dh)\, dt$$

uniformly in ω (since the bounds in step 3 are almost sure), where we note that the dyadic shift can alter the absolute difference by at most $2\|\lambda\|_\infty 2^{-m} \xrightarrow{m\to\infty} 0$. This proves the claim. Note that the process $(\pi_t^{(n(m),m)})_{t\in[0,T]}$ is defined on the original probability space $(\Omega, \mathcal{F}, P)$ and, due to the time shift, remains progressively measurable w.r.t. $\mathbb{F}$. □

Proposition 3.14. *Consider the setup of the previous lemma. Then for every $\gamma > 0$ and any sufficiently large n and m along the subsequence n_l, m_l from the statement of lemma 3.13 (with the threshold depending on γ), there exist a tuple $(\Omega^{(n,m)}, \mathcal{F}^{(n,m)}, \mathbb{F}^{(n,m)}, P^{(n,m)}, \pi^{(n,m)}, N^{(n,m)}, m^{(n,m)}, \eta^{(n,m)}, X^{(n,m)})$ satisfying Definition 3.2, such that $\pi^{(n,m)}$ is the process constructed in the previous lemma, and such that*

$$\left|W(0,x,\pi) - W(0,x,\pi^{(n,m)})\right| < \gamma.$$

In addition, using the sequence of measures $\eta_k^{(n)} = \pi_{k2^{-n}}^{(n,m)}(\mathcal{U} \times \cdot)$ as the parametrizing sequence in the discrete-time Scheme 2 system, we have that

$$\liminf_{n,m\to\infty}[v^{(n)}(0,x,(\eta_k^{(n)})_{k=0}^{T2^n}) - W(0,x,\pi^{(n,m)})] \geq 0$$

Proof. Recall that in lemma 3.13, $\pi^{(n,m)}$ was constructed on the same probability space as the original process π. Also, note that when taking $m, n \to \infty$, we will do so as in lemma 3.13 so that (3.2) holds. By possibly augmenting the probability space (recall Remark 3.3), let $N^{(n,m)}$ denote another jump process with stochastic intensity given by $\lambda_t^{(n,m)} = \int_{\mathcal{U}^2} \lambda(a,h)\pi_t^{(n,m)}(da,dh)$ and let $X^{(n,m)}$ and $N^{(n,m)}$ satisfy

$$X_t^{(n,m)} = x - \int_0^t \int_{\mathcal{U}} ca\pi_s^{(n,m)}(da \times \mathcal{U})ds + rN_t^{(n,m)}.$$

From the uniform almost-sure convergence in (3.2) we immediately conclude that $\mathcal{L}(X_T^{(n,m)}) \xrightarrow{\mathcal{L},n,m\to\infty} \mathcal{L}(X_T)$ which implies $E[\varphi(X_T^{(n,m)})] \xrightarrow{n,m\to\infty} E[\varphi(X_T)]$ as desired. Alternatively, if the construction is such that the probability measure is modified to obtain the new process $N^{(n,m)}$, then we argue as follows. By construction, the laws $\{\mathcal{L}(\pi^{(n,m)}, N^{(n,m)}) : n, m \in \mathbb{N}\}$ are tight and we let $\tilde{\pi}$ and $\tilde{N}$ denote a limit point. From the chattering construction, it follows that $\mathcal{L}(\tilde{\pi}) = \mathcal{L}(\pi)$, thus it follows that the limit point of the above is in fact the original law $\mathcal{L}(\pi, N)$ and we again conclude $E[\varphi(X_T^{(n,m)})] \xrightarrow{n,m\to\infty} E[\varphi(X_T)]$. For more details regarding convergence of values, particularly in the weak formulation, we refer the reader to [23, Theorem 3.5.2, Theorem 3.2.2]. We note that the chattering constructions and value approximation arguments are classical; our approach differs in that it constructs the approximations while respecting a prescribed marginal approximation, and ensuring that the disintegration measures $\pi_t^{(n,m)}$ admit a product measure structure.

Now we move on to the second assertion. On the probability space where $\pi^{(n,m)}$ is defined, define a random sequence of measures

$$a_k^{(n,m)}(\omega)(\cdot) = 2^n \int_{I_k^n} \pi_t^{(n,m)}(\cdot \times \mathcal{U})dt.$$

Now construct a discrete-time state process $x_k^{(a,n)}$ driven by this relaxed control, using the relaxed Scheme 2 dynamics. By the Poisson limit theorem, we have that $|E[\varphi(X_T^{(n,m)})] - E[\varphi(x_{2^nT}^{(a,n)})]| \xrightarrow{n,m\to\infty} 0$. Since relaxation does not improve the value function (recall Lemma 3.10), the result follows. Note that, abusing notation, we have used $\eta^{(n)}$ to refer to both the original discrete-time sequence obtained from the discrete-time MFG existence result, as well as the shifted approximating sequence used in the construction of $\pi^{(n,m)}$. This poses no issues, since the effect of the 2^{-m} time shift from lemma 3.13 vanishes as $m \to \infty$. We show this rigorously in Lemma 6.2. $\square$

We can finally conclude optimality, which we state as a proposition.

Proposition 3.15. *Consider the tuple obtained as a limit of discrete-time processes. Then amongst admissible processes which admit the second marginal η, the limit π obtained is optimal.*

Proof. The proof follows the following roadmap. We use fixed-n optimality of the discrete-time controls, approximate a competing continuous-time control by the joint chattering construction, pass values to the limit, and obtain a contradiction to discrete-time optimality.

Suppose now for a contradiction that π is not optimal for η, in the sense that there exists another process $\tilde{\pi}$ with second marginal equal to η such that its expected terminal wealth utility under the agent dynamics is strictly better than the one obtained by the process π. Then using the Chattering Lemma 3.13 there exists another admissible process $\pi^{(n,m)}$ (which in this case approximates $\tilde{\pi}$) such that its second marginal matches $\eta^{(n)}$ (up to a 2^{-m} dyadic shift) and with the property that $W(0,x,\pi) < W(0,x,\pi^{(n,m)})$. We sample $\pi^{(n,m)}$ for each n at n-th order dyadics to obtain a discrete control process $\tilde{a}^{(n,m)}$. Using (3.14) and (3.14) we obtain

$$\begin{aligned}\liminf_{n\to\infty} v^{(n)}(0,x,\eta^{(n)}) &\geq \liminf_{n,m\to\infty} w^{(n)}(0,x,\tilde{a}^{(n,m)},\eta^{(n)}) \\ &= W(0,x,\tilde{\pi}) > W(0,x,\pi) = \lim_{n\to\infty} w^{(n)}(0,x,a^{(n)},\eta^{(n)}),\end{aligned}$$

thus by taking m and n sufficiently large along the subsequence in Lemma 3.13, we obtain a contradiction to the optimality of $a^{(n)}$ for the fixed sequence of measures $\eta^{(n)}$. □

We have established the first assertion of the theorem. For the second assertion regarding the sharpness of the first marginal of the admissible process π, assume φ is strictly increasing and that $\lambda(\cdot,h)$ is strictly concave for each h. Let π denote a relaxed MFG equilibrium for the tuple. Fix $t\in[0,T]$, define $a_t^*(h) = \int a\pi_t(da|h)$ where $\pi_t(da,dh) = \pi_t(da|h)\eta_t(dh)$, and note that

$$\int_{\mathcal{U}} \lambda(a_t^*(h),h)\pi_t(\mathcal{U}\times dh) \geq \int_{\mathcal{U}}\int_{\mathcal{U}} \lambda(a,h)\pi_t(da|h)\eta_t(dh).$$

Thus, letting X' and X denote the wealth processes (possibly defined on distinct probability spaces) driven by the processes $\delta_{\{a_t^*(h)\}}(da)\eta_t(dh)$ and π_t respectively, we have that $E[\varphi(X'_T)] \geq E[\varphi(X_T)]$, with equality if and only if π_t disintegrates into a sharp kernel conditioned on the second marginal for every $t\in[0,T]$. By optimality of π, it follows immediately that π is itself (up to redefining on Lebesgue null sets) a sharp control. This concludes the proof. □

We conclude with a remark on how the dynamics and objective functionals can be generalized.

Remark 3.16. *The specific dynamics in Definition 3.2 were chosen to match the structure of cryptocurrency mining model considered below. The argument can be extended to more general controlled jump dynamics and to objective functionals with running rewards, provided the coefficients and rewards satisfy the compactness, continuity, and boundedness conditions needed for tightness, identification of the limiting martingale problem, and convergence of values. This extra generality was avoided to keep the presentation focused on the controlled-intensity structure used in [28].*

4 Cryptocurrency Mining MFG

We now return to the cryptocurrency mining MFG model from [28]. We provide MFG existence results for both the original continuous-time model and a discrete-time analog. For the latter, we numerically illustrate that damped fixed-point iterations of the discrete-time Bellman and Kolmogorov equations converge to an MFG equilibrium, and we recover the qualitative preferential attachment behavior obtained from the continuous-time model in [28] under CRRA wealth utility. We note that the numerical scheme utilized in this paper is distinct from the fixed point iterations used to solve the continuous-time MFG in [28]. There, finite difference methods are used to solve the HJB and Kolmogorov PDEs, whereas we exactly solve a space-discretized discrete-time game, and we need not scale the dynamics so that time steps correspond to small intervals of time.

4.1 Discrete-Time Model

We begin by considering a family of discrete-time models parametrized by $n \in \mathbb{N}$ and $\epsilon \geq 0$. Using Scheme 2 from Definition 3.6, we define

$$\lambda^{(\epsilon)}(a,h) \coloneqq \frac{a}{a+hM+\epsilon}\mathbb{1}_{\{a>0\}}, \tag{4.1}$$

for a constant $M > 0$ (which proxies for the number of miners, see [28] for details). Given a fixed sequence of probability measures $(\zeta_k^{(n,\epsilon)})_{k=0}^{2^nT-1} \subseteq \mathcal{P}(\mathcal{U})$, the transition dynamics for the wealth process $(x_k^{(n,\epsilon)})_{k=0}^{2^nT}$ are given by

$$x_{k+1}^{(n,\epsilon)} \sim \frac{\lambda^{(\epsilon)}(a_k^{(n,\epsilon)},\overline{\zeta}_k^{(n,\epsilon)})}{2^n}\delta_{\{x_k^{(n,\epsilon)}-(ca_k^{(n,\epsilon)}/2^n)+r\}} + \Big(1-\frac{\lambda^{(\epsilon)}(a_k^{(n,\epsilon)},\overline{\zeta}_k^{(n,\epsilon)})}{2^n}\Big)\delta_{\{x_k^{(n,\epsilon)}-(ca_k^{(n,\epsilon)}/2^n)\}}, \tag{4.2}$$

for $k = 0, 1, \cdots, T2^n - 1$. The control objective is the maximization of expected terminal utility of wealth, $E[\varphi(x_{T2^n}^{(n,\epsilon)})]$. Assuming that the initial wealth law μ_0 is compactly supported, that the action space is a compact interval $[0, L]$, and that the wealth utility φ is bounded (which is WLOG since agents can increase their wealth by at most 2^nTr over the course of the game), it is straightforward to check that, for $\epsilon > 0$, Theorem 2.4 applies, guaranteeing the existence of an MFG equilibrium for the game arising from the above dynamics. The condition $\epsilon > 0$ is required in order to ensure continuity of the transition kernel. For the remainder of this section, we fix $n \in \mathbb{N}$ and will extract an MFG equilibrium as a limit of equilibria by taking $\epsilon \to 0$.

Fix n and $\epsilon > 0$, and let $\nu^{(n,\epsilon)} \coloneqq (\nu_k^{(n,\epsilon)})_{k=0}^{2^nT}$ denote a sequence of laws on $\mathcal{X}\times\mathcal{U}$ which characterize an MFG equilibrium (recall that the control policy is obtained via the disintegration as in Section 2) for the game arising from the dynamics in (4.2). For ease of notation, let $\mu_k^{(n,\epsilon)} = \nu_{k,1}^{(n,\epsilon)}$ and $\eta_k^{(n,\epsilon)} = \nu_{k,2}^{(n,\epsilon)}$ where $\nu_{k,1}^{(n,\epsilon)}$ and $\nu_{k,2}^{(n,\epsilon)}$ denote the state and control marginals of $\nu_k^{(n,\epsilon)}$, respectively. In this specific setup, $\mathcal{X} = \mathbb{R}$ and $\mathcal{U} = [0, L]$. Recall also that the optimal control

$$a_k^{(n,\epsilon)} \sim \pi_k^{(n,\epsilon)}(\cdot|x_k^{(n,\epsilon)}) \quad \text{where} \quad \nu_k^{(n,\epsilon)}(dx,da) = \pi_k^{(n,\epsilon)}(da|x)\mu_k^{(n,\epsilon)}(dx).$$

We now extract a weak limit $\nu^{(n,\epsilon)} \to \nu^{(n)}$ as $\epsilon \to 0$ and will establish that $\nu^{(n)}$ is an MFG equilibrium for the cryptocurrency mining model (4.2) with $\epsilon = 0$. Since $n \in \mathbb{N}$ is fixed we suppress it from the superscript notation hereafter, and also write $\lambda = \lambda^{(0)}$ for simplicity. We can justify the limit extraction by noting that the joint state-action laws $\nu^{(\epsilon)}$ are elements of the compact set Ξ constructed in Section 2, which need not be ϵ-dependent since we can always pick the moment $w : \mathcal{X} \to [1,\infty)$ and constant α used in its construction so that, for all $\epsilon > 0$ sufficiently small, Assumption 2.3 holds.

We begin by observing that if

$$B \coloneqq \{k \in \{0, 1, \cdots, 2^nT - 1\} : \eta_k = \delta_{\{0\}}\}$$

is non-empty, then an optimal control for the original ($\epsilon = 0$) model with fixed background hash-rate $(\overline{\eta}_k)_{k\in\{0,1,2,\cdots,2^nT\}}$ will in general not exist: this follows because any non-zero hash rate will result in a unit jump intensity for the reward process, whereas zero hash-rate results in zero reward process intensity, hence zero hash-rate on B is sub-optimal for any reasonable choice of terminal utility/initial condition. On the other hand, any control process that is non-zero on B can be improved by making it even closer to zero while keeping it positive. Since $\mathcal{U}$ is a subset of the nonnegative reals, $\int h\, \eta_k(dh) = 0$ if and only if $\eta_k = \delta_0$; hence assigning positive mass to zero is not by itself problematic. Fortunately, assuming the following mild assumption, we can show that the limit hash-rate is indeed positive for every time k.

Assumption 4.1. *Assume that φ is non-decreasing and that the initial wealth distribution is not supported on the set of maximizers for the function φ, i.e.*

$$\mu_0\{\arg\max_{y\in\mathbb{R}} \varphi(y)\} < 1.$$

Proposition 4.2. *Under Assumption 4.1, the limit (as $\epsilon \to 0$ with n fixed but arbitrary) sequence of measures on the control space $(\eta_k)_{k=0}^{2^nT} = (\nu_{k,2})_{k=0}^{2^nT} \subseteq \mathcal{P}([0,L])$ is never a Dirac at zero. In fact, there exists some $d > 0$ such that $\overline{\eta}_k > d$ for every $k = 0, 1, \cdots, 2^nT - 1$ and we can find such a d that holds for any choice of $n \in \mathbb{N}$.*

Proof. Suppose for a contradiction that there is some k such that $\overline{\eta}_k = 0$. It follows that $\lim_{\epsilon\to 0} \overline{\eta}_k^{(\epsilon)} = 0$. Recall that $a^{(\epsilon)}$ is an optimal control for the fixed background hash-rate $\eta^{(\epsilon)}$. For a given ϵ, let $\nu(\epsilon) := \epsilon \vee \overline{\eta}_k^{(\epsilon)}$, and define on the same probability space a new control process given by

$$\tilde{a}_j^{(\epsilon)} = \begin{cases} a_j^{(\epsilon)} & j \neq k \\ \sqrt{\nu(\epsilon)} \vee a_j^{(\epsilon)} & j = k, \end{cases}$$

for $j = 0, 1, \cdots, 2^nT - 1$. We will complete the proof by showing that for sufficiently small ϵ, $\tilde{a}^{(\epsilon)}$ results in higher expected wealth utility than $a^{(\epsilon)}$. We only need to compare the controls at time k. The expected rate of the Bernouli reward at time k of the original control satisfies

$$E\Big[\frac{1}{2^n}\lambda^{(\epsilon)}(a_k^{(\epsilon)}, \overline{\eta}_k^{(\epsilon)})\Big] \leq \frac{1}{2^n}\lambda^{(\epsilon)}(E[a_k^{(\epsilon)}], \overline{\eta}_k^{(\epsilon)})dt = \frac{\overline{\eta}_k^{(\epsilon)}}{2^n((M+1)\overline{\eta}_k^{(\epsilon)} + \epsilon)} \leq \frac{1}{2^n(M+1)}.$$

where we have used MFG consistency in the equality. On the other hand, we also have that

$$E\Big[\frac{1}{2^n}\lambda^{(\epsilon)}(\tilde{a}_k^{(\epsilon)}, \overline{\eta}_k^{(\epsilon)})\Big] \geq \frac{\sqrt{\nu(\epsilon)}}{2^n(\sqrt{\nu(\epsilon)} + \overline{\eta}_k^{(\epsilon)}M + \epsilon)} \geq \frac{\sqrt{\nu(\epsilon)}}{2^n(\sqrt{\nu(\epsilon)} + \nu(\epsilon)(M+1))} \xrightarrow{\epsilon\to 0} 2^{-n}. =$$

Observe that the difference in costs between the two controls converge as $\epsilon \to 0$. Also, Assumption 4.1 implies that uniformly over $\epsilon \geq 0$, the choice of control processes (since the intensity is bounded), and the possible values of k, there is some positive probability that the wealth $x_k^{(n,\epsilon)}$ has not reached $\arg\max_{y\in\mathbb{R}} \varphi(y)$. As such, taking $\epsilon \to 0$ results in a converging of the expected costs of the controls $\tilde{a}^\epsilon$ and a^ϵ, whereas the number of rewards of the former is strictly higher than the latter, with the difference not converging as $\epsilon \to 0$. As such, the former control will eventually result in strictly higher expected terminal wealth utility compared to the latter for ϵ sufficiently small. This contradicts optimality of a^ϵ for the fixed sequence η^ϵ, completing the proof.

To prove the last statement regarding the uniformity of d over the choice of n, we make the dependency on n explicit. Note that the positive probability that the wealth $x_k^{(n,\epsilon)}$ has not reached the supremum of the terminal wealth utility is itself uniform in n; this follows using the Poisson limit theorem. Also, the dependency of n on the bound

$$\liminf_{\epsilon\to 0}\left(E\Big[\frac{1}{2^n}\lambda^{(\epsilon)}(\tilde{a}_k^{(n,\epsilon)},\overline{\eta}_k^{(n,\epsilon)})\Big]-E\Big[\frac{1}{2^n}\lambda^{(\epsilon)}(a_k^{(n,\epsilon)},\overline{\eta}_k^{(n,\epsilon)})\Big]\right)\geq\frac{M}{2^n(M+1)}$$

is offset by the fact that as n increases, the difference in cost of the controls $a_k^{(n,\epsilon)}$ and $\tilde{a}_k^{(n,\epsilon)}$ converges to zero faster (also with a $1/2^n$ factor) as $\epsilon \to 0$. □

We are now ready to establish the existence of an MFG equilibrium as the limit of equilibria for the model parametrized by ϵ for a fixed $n \in \mathbb{N}$.

Proposition 4.3. *Recall that ($n \in \mathbb{N}$ is fixed and suppressed from the notation) $\nu^{(\epsilon)}$ denotes (the laws of) an MFG equilibrium for the model parametrized by $\epsilon > 0$ and ν is a weak limit as $\epsilon \to 0$. This object constitutes an MFG equilibrium for the model with $\epsilon = 0$.*

Proof. We have shown that none of the measures in the sequence $(\eta_k)_{k=0}^{2^nT}$ are Dirac measures at zero. Let $a_k \sim \pi_k(\cdot|x_k)$ where $\nu_k(da,dx) = \pi_k(da|x)\nu_{k,1}(dx)$. Using the continuity and boundedness of the map

$$(\mathcal{X}\times\mathcal{U})^{2^nT+1}\ni(x_k,a_k)_{k=0}^{2^nT}\mapsto\varphi(x_{2^nT})\in\mathbb{R},$$

and the weak convergence $\nu^{(\epsilon)} \xrightarrow{\epsilon\to 0} \nu$, it follows that $E[\varphi(x_{2^nT}^{(\epsilon)})] \xrightarrow{\epsilon\to 0} E[\varphi(x_{2^nT})]$. Suppose now for a contradiction that there exists a control policy $\tilde{\pi} = (\tilde{\pi}_k)_{k=0}^{2^nT-1}$ which outperforms a under the fixed background sequence $(\eta_k)_{k=0}^{2^nT-1}$. Let $\tilde{x}^{(\epsilon,\eta^{(\epsilon)})}$ and $\tilde{a}^{(\epsilon,\eta^{(\epsilon)})}$ denote (each possibly defined on its own probability space) the resulting control and state processes under the policy $\tilde{\pi}$ using the reward probability function $\lambda^{(\epsilon)}$ and the fixed sequence of measures $\eta^{(\epsilon)}$. Using the fact that $\overline{\eta}_k^{(\epsilon)} \xrightarrow{\epsilon\to 0} \overline{\eta}_k > 0$ for every $k = 0, \cdots, 2^nT$, one obtains that $E[\varphi(\tilde{x}_{2^nT}^{(\epsilon,\eta^{(\epsilon)})})] \xrightarrow{\epsilon\to 0} E[\varphi(\tilde{x}_{2^nT}^{(0,\eta)})] > E[\varphi(x_{2^nT})]$ which, for all $\epsilon > 0$ sufficiently small, contradicts the optimality of the control policy obtained by disintegration of $\nu^{(\epsilon)}$ for the fixed background flow $\eta^{(\epsilon)}$ and probability reward function $\lambda^{(\epsilon)}$. Note that consistency of the limit measures follows because it holds for every $\epsilon > 0$. □

To summarize, we have shown that for any $n \in \mathbb{N}$ there exists an MFG equilibrium for the discrete-time game with $\epsilon = 0$. Moreover, the population control distributions are bounded away form zero, uniformly over the time step and the parameter $n \in \mathbb{N}$. Before moving to continuous-time, we solve the discrete-time game numerically.

4.2 Numerical Computation of Discrete-Time MFG

We solve the discrete-time MFG for parameters $n = 1$, $T = 300$, $M = 1000$ (players), take the terminal utility to be a CRRA utility $\varphi(x) = 2x^{1/2}$, and assume normally distributed initial wealth. The population hash-rate is initialized to be constant in time, and the following steps are then iterated until convergence of the population hash-rate is observed:

1. Compute the optimal control policy using dynamic programming, given the fixed background hash-rate from the previous step.

2. Compute the resulting hash-rate using the initial state law and the optimal control policy from step 1.
3. Update the new population hash-rate as a convex combination of that obtained in the last two steps.

Step 3 is a damping step which is required to attain convergence. The damping factor can be tuned to find a trade-off between convergence and speed. Solving with a damping factor of 0.9, we observe convergence from any choice of initial constant population hash-rate, and obtain numerical wealth distribution dynamics and optimal control policies for the representative agent (plotted below).

The results obtained are qualitatively similar to those obtained from the numerical PDE approach in [28], where we see that agents drop out of the game if their wealth falls below a (time-dependent) threshold and observe the preferential attachment phenomenon where a small percentage of the population becomes increasingly wealthy, with the majority of miners eventually dropping out of the game with comparatively little wealth. Figures 1 and 2 illustrate the discrete-time convergence results. Convergence speed is highly dependent on the choice of initial condition, in part due to the use of a large damping factor. The stability of the algorithm suggests the uniqueness of the equilibria.

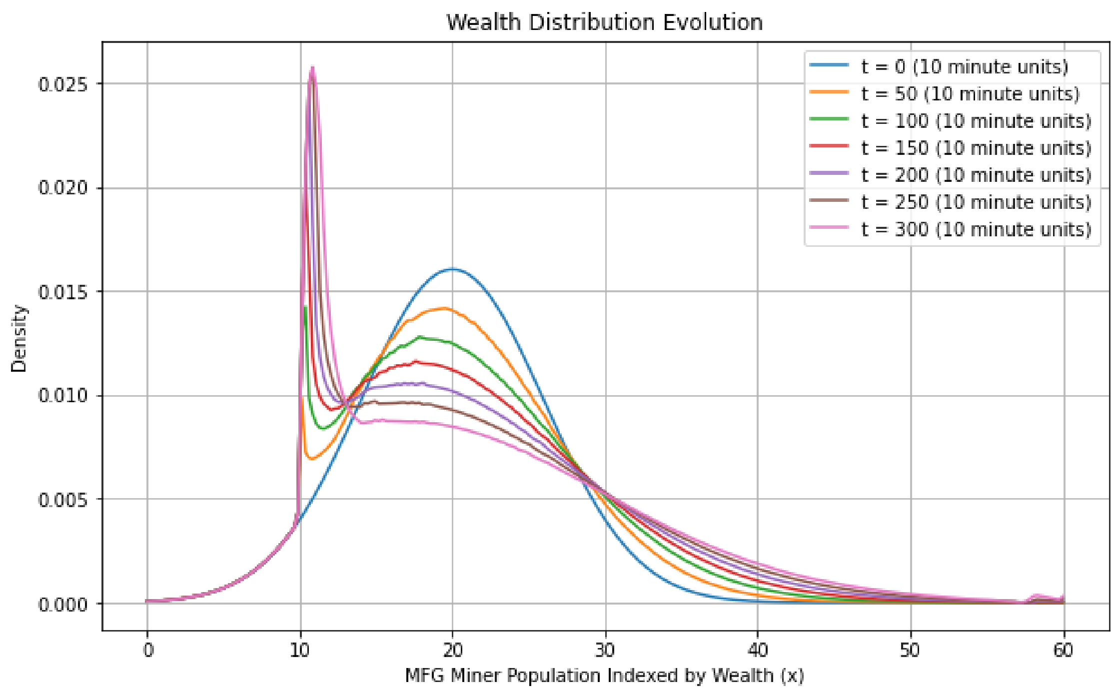

Figure 1: Evolution of Wealth Distribution

4.3 Continuous-Time Existence

To apply the convergence result Theorem 3.11 (and Theorem 2.4 for the discrete-time games), we again work with the intensity map $\lambda^{(\epsilon)}$ defined in (4.1). Taking $\mathcal{U} = [0, L]$ and assuming the terminal wealth utility is continuous and bounded and that the initial state law is compact, it is straightforward to check that for $\epsilon > 0$, Theorem 3.11 applies, guaranteeing the existence of a relaxed MFG equilibrium (of *a priori* relaxed controls) for the continuous-time game. As in the

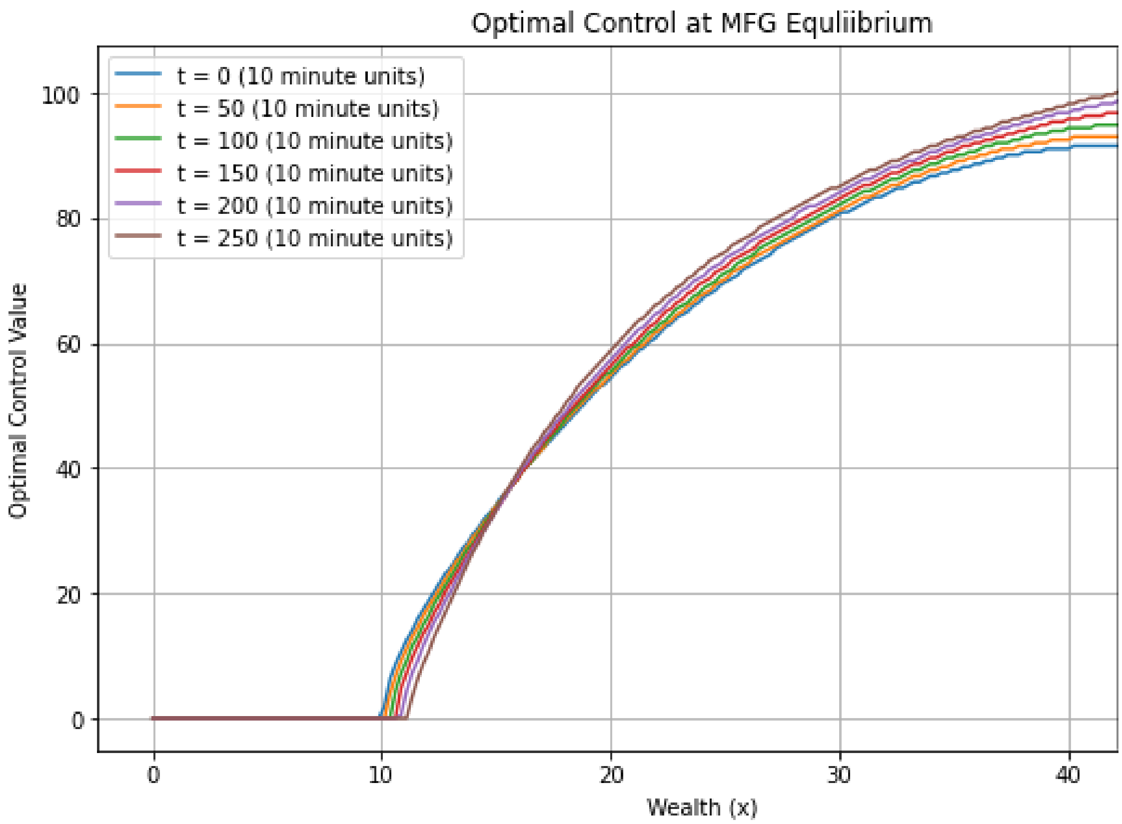

Figure 2: Optimal Control at Equilibrium

discrete-time case, we turn to the problem of extracting a limit for the original model by taking $\epsilon \to 0$. First, we note the following:

Proposition 4.4. *If the terminal wealth utility φ is strictly increasing, any MFG equilibrium for the continuous-time cryptocurrency model is of sharp controls.*

Proof. This follows from the second assertion of Theorem 3.11. □

Suppose now that for a given $\epsilon > 0$, $(X^{(\epsilon)}, N^{(\epsilon)}, \pi^{(\epsilon)})$ is an MFG equilibrium with the state, jump, and control processes defined on a common probability space (which may vary with ϵ). As in discrete-time, our aim is to extract a weak limit of process laws by taking $\epsilon \to 0$. Let (X, N, m) denote processes (defined on a possibly distinct probability space) such that

$$\mathcal{L}(X^{(\epsilon)}, N^{(\epsilon)}, \pi^{(\epsilon)}) \xrightarrow{\mathcal{L}, \epsilon \to 0} \mathcal{L}(X, N, \pi). \tag{4.3}$$

We need to establish that (X, N, π) is an MFG equilibrium for the cryptocurrency mining model with $\epsilon = 0$. Similarly to the discrete-time case, if the set $B := \{t \in [0,T] : \eta_t = \eta_{\{0\}}\}$ has positive Lebesgue measure, then an optimal control for the $\epsilon = 0$ model with fixed control measure flow $(\eta_t)_{t\in[0,T]}$ will in general not exist. Under Assumption 4.1, we will see that B has zero Lebesgue measure. As in Section 3, we recall the notation $m_t(\cdot) = \pi_t(\cdot \times \mathcal{U})$ and $\eta_t(\cdot) = \pi_t(\mathcal{U} \times \cdot)$.

Proposition 4.5. *The limit disintegration $(\eta_t)_{t\in[0,T]} \subseteq \mathcal{P}([0, L])$ is not a Dirac at zero except for possibly on a Lebesgue null set.*

Proof. Recall that the object η has been obtained by, for each ϵ, interpolating and then embedding $\overline{\eta}^{(n,\epsilon)} \in [0, L]^{n+1}$ into a space of measures, and taking $n \to \infty$ to obtain $\eta^{(\epsilon)} \in \mathcal{P}([0, T] \times [0, L])$. Then, we take $\epsilon \to 0$ to obtain η. By Proposition 4.2, however, we know that there is some $d > 0$ such that $\overline{\eta}_t^{(\epsilon)} > d$ for every $t \in [0, T]$ and every ϵ sufficiently small, from which the result follows. □

We conclude with the following proposition.

Proposition 4.6. *The tuple (X, N, m, η) defined above (4.3) constitutes an MFG equilibrium for the cryptocurrency mining model with $\epsilon = 0$. The equilibrium is relaxed in the sense of Definition 3.5, but is of sharp control.*

Proof. We have shown that the limit hash-rate $\overline{\eta}_t \geq d$ for every $t \in [0, T]$. We now show that the control m is optimal for the fixed measure flow η. Note the continuity (note that continuity on $D([0, T])$ is with respect to the Skorokhod topology for which the evaluation map at the endpoint T is continuous), and boundedness of the map

$$D([0, T]) \times \mathcal{P}([0, T] \times [0, L]) \ni (N, m) \mapsto \varphi\Big(rN_T - c \int_0^T \int_{[0,L]} am(dt, da)\Big).$$

Suppose now for a contradiction that there exists a measure-valued control process $\tilde{m}$ which outperforms m under the fixed background flow η and intensity function λ. Then we have that (expectations taken on possibly different probability spaces)

$$E[\varphi(rN_T - c \int_0^T m_t dt)] < E[\varphi(r\tilde{N}_T^{(0)} - c \int_0^T \overline{\tilde{m}}_t dt)] = \lim_{\epsilon \to 0} E[\varphi(r\tilde{N}_T^{(\epsilon)} - c \int_0^T \overline{\tilde{m}}_t dt)],$$

where $\tilde{N}^{(\epsilon)}$ denotes a unit jump process with stochastic intensity $\int \int \lambda^{(\epsilon)}(a, h)\tilde{m}_t(da)\eta_t^{(\epsilon)}(dh)$. The above limit is justified as follows. By Proposition 4.2, the limiting background hash rate is bounded away from zero: $\int_{\mathcal{U}} h \, \eta_t(dh) \geq d > 0$ for all $t \in [0, T]$. Hence, in the sharp limiting background flow considered here, the denominator is bounded away from zero uniformly in the competitor's control, and the singular point $(a, h) = (0, 0)$ of the unregularized intensity is not encountered. For any fixed admissible relaxed control $\tilde{m}$, the corresponding intensities

$$\int_{\mathcal{U}} \int_{\mathcal{U}} \lambda^{(\epsilon)}(a, h) \, \tilde{m}_t(da)\eta_t^{(\epsilon)}(dh)$$

therefore converge to the analogous intensity with $\epsilon = 0$. Since the intensities are uniformly bounded by one, the induced jump processes are tight, and the bounded continuous payoff functional converges along this sequence. This proves the required convergence of payoffs. For $\epsilon > 0$ sufficiently small, this contradicts the optimality of $m^{(\epsilon)}$ for the regularized background flow $\eta^{(\epsilon)}$. As such, we conclude that m is optimal for the hash-rate η as desired. Consistency follows by continuity, taking limits, and using the fact that consistency holds for every $\epsilon > 0$, and we thus conclude that the tuple (X, N, m, η) constitutes an MFG equilibrium for the cryptocurrency mining model with $\epsilon = 0$. The equilibrium is relaxed in the sense of Definition 3.5, but is of sharp control. □

4.4 Uniqueness and Sharpness of MFG Equilibrium

So far, we have not discussed the question of MFG uniqueness for the cryptocurrency MFG model. Numerically, one observes that, for reasonable wealth utilities, the same MFG solution is obtained from fixed-point iterations independently of starting conditions of the algorithm, suggesting uniqueness in these cases. Consider the discrete-time representative agent problem for a given $n \in \mathbb{N}$ with fixed population hash-rates $(\overline{\eta}_k^{(n)})_{k=0}^{2^n T}$ and assume an increasing and strictly concave utility function $v_{2^n T}^{(n)} = \varphi$. Economic intuition suggests that strict concavity and monotonicity extend to the value functions at all times. Under the strict concavity of the value functions, one can show that

$$a_k^{(n)*}(x) = \underset{\alpha\in[0,L]}{\arg\max}\{E[v_{k+1}^{(n)}(x_{k+1})|x_k^{(n)} = x, a_k^{(n)} = a]\}$$

is a singleton. This follows by observing that the objective function in the above maximization is strictly quasi-concave. Moreover, using [27, Theorem 1] one can see that the quantity $a_k^{(n)*}(x)$ is strictly decreasing in $\overline{\eta}_k^{(n)}$ everywhere except possibly on a small neighborhood of zero. Provided that one can show that any MFG equilibrium results in a population hash-rate that is not in this small neighborhood (which may follow in certain cases from arguments as in Proposition 4.2), then uniqueness of equilibrium follows from a simple contradiction argument by assuming two distinct equilibria. Note also that uniqueness of equilibrium for each $n \in \mathbb{N}$ allows one to conclude that the continuous-time equilibrium hash rate $(\eta_t)_{t\in[0,T]}$ is in fact $[0, L]-$valued (and not $\mathcal{P}([0, L])$ valued), that is, sharpness of MFG equilibrium.

Because the question of establishing conditions for uniqueness and sharpness of the MFG equilibrium are specific to the wealth utility and parameter choices, we leave this for future work.

5 Conclusion

In this paper, we have accomplished three tasks. First, a general discrete-time MFG existence theorem was established, involving general transition dynamics with mean-field interactions via both the states and controls, and influencing both the transition dynamics and costs. Second, the discrete-time result was used to obtain relaxed MFG equilibria existence results for models of controlled jump intensity with mean-field interaction via the controls, and affecting the intensity of the jump processes. Finally, the results were applied to provide existence guarantees for a cryptocurrency mining MFG model, and an alternative numerical scheme, motivated by the discrete-time to continuous-time convergence result was implemented. This scheme was shown to coincide (in the sense of obtaining similar qualitative agent behavior) with numerical solutions to the original continuous-time cryptocurrency mining MFG, which was solved by numerically solving coupled Kolmogorov and HJB PDEs.

6 Appendix: Tightness Theorem

The following result is restated here for convenience, and its proof can be found in [24, Theorem 9.2.1].

Theorem 6.1. *(Tightness Criteria for the space D) Consider an arbitrary collection (possibly uncountable) of processes $\{X^{(\alpha)} : \alpha \in I\}$ taking values in the space $D^k[0,\infty)$ ($\mathbb{R}^k$-valued cadlag functions on $[0,\infty)$) and defined on a common probability space $(\Omega, \mathcal{F}, P)$. Assume that for each $\delta > 0$ and rational $t \in [0,\infty) \cap \mathbb{Q}$ there exists a corresponding compact set $K_{\delta,t}$ such that*

$$\sup_{\alpha\in I} P(X_t^{(\alpha)} \notin K_{\delta,t}) \leq \delta.$$

Let now $\mathcal{F}_t^{(\alpha)} := \sigma\{X_s^{(\alpha)} : s \leq t\}$ *and let* $\mathcal{T}_T^{(\alpha)}$ *denote the set of* $\mathcal{F}_t^{(\alpha)}$ *stopping times that are bounded by* T*. Suppose now that for each* $T \in [0, \infty)$ *we have that*

$$\lim_{\delta \to 0} \sup_{\alpha \in I} \sup_{\tau \in \mathcal{T}_T^{(\alpha)}} E(\mathbb{1} \wedge |X_{\tau+\delta}^{(\alpha)} - X_\tau^{(\alpha)}|) = 0.$$

Then the family of laws $\mathcal{L}(X^{(\alpha)})_{\alpha \in I}$ *is tight.*

Lemma 6.2. *Let* $n > m$ *be two integers and consider a deterministic parametrizing sequence* $\eta^{(n)} = (\eta_k^{(n)})_{k=0}^{2^n T} \subseteq \mathcal{P}(\mathcal{U})$*. Define the shifted sequence*

$$\tilde{\eta}_k^{(n)} = \begin{cases} \eta_{k-m}^{(n)}, & k \geq m, \\ \delta_{\{0\}}, & k < m. \end{cases}$$

Assume that $n, m \to \infty$ *with* n *going sufficiently faster than* m*. Then, for the Scheme 2 dynamics from Definition 3.6, and under the regularit assumptions 1, 2, and 3 from Assumption 3.9*

$$\left| v^{(n)}(0, x, \eta^{(n)}) - v^{(n)}(0, x, \tilde{\eta}^{(n)}) \right| \to 0.$$

Proof. Write $\Delta_n = 2^{-n}$ and $N_n = T 2^n$. For a fixed deterministic sequence $\eta^{(n)}$, let $Q_k^{\eta^{(n)}}$ denote the one-step controlled transition operator associated with Scheme 2 at time k. Thus

$$v^{(n)}(N_n, x, \eta^{(n)}) = \varphi(x), \quad v^{(n)}(k, x, \eta^{(n)}) = \sup_{a \in \mathcal{U}} Q_k^{\eta^{(n)}} v^{(n)}(k+1, \cdot, \eta^{(n)})(x),$$

for $k = N_n - 1, \ldots, 0$, and where,

$$Q_k^{\eta^{(n)}} f(x) = \alpha_k^{(n)} f(x - ca\Delta_n + r) + (1 - \alpha_k^{(n)}) f(x - ca\Delta_n), \qquad \alpha_k^{(n)} = \Delta_n \lambda(a, \overline{\eta}_k^{(n)}).$$

We first note a short-time stability estimate. Let $\delta_n = m\Delta_n$. Since $\mathcal{U}$ is compact and λ is bounded, over any block of m time steps the total drift displacement is bounded by $cL\delta_n$ for some $L > 0$ (in our context, we have $\mathcal{U} = [0, L]$), and the probability of at least one jump is bounded by $\|\lambda\|_\infty \delta_n + o(\delta_n)$, uniformly over the parametrizing sequence and over admissible controls. Since μ_0 is compactly supported and the jump intensity is uniformly bounded, the state processes can be localized on compact sets with probability arbitrarily close to one, uniformly in n and in the parametrizing sequence. On such compact sets, φ is uniformly continuous. Therefore there exists a deterministic function $\rho(\delta)$, with $\rho(\delta) \to 0$ as $\delta \downarrow 0$, such that, uniformly over deterministic parametrizing sequences,

$$\left| v^{(n)}(k, x, \eta^{(n)}) - v^{(n)}(k+m, x, \eta^{(n)}) \right| \leq \rho(m\Delta_n)$$

whenever both sides are defined. The same estimate holds with $\eta^{(n)}$ replaced by $\tilde{\eta}^{(n)}$.

Now compare the dynamic programs corresponding to $\eta^{(n)}$ and $\tilde{\eta}^{(n)}$. By construction,

$$\tilde{\eta}_k^{(n)} = \eta_{k-m}^{(n)}, \qquad k \geq m.$$

Thus, after the first m steps, the shifted parametrizing sequence uses the same measures as the original sequence, only delayed by m time steps. Hence the difference between the two value functions is controlled by the value loss over the initial block of length m and the corresponding terminal truncation of length m. Applying the preceding short-time stability estimate to these two blocks gives

$$\left| v^{(n)}(0, x, \eta^{(n)}) - v^{(n)}(0, x, \tilde{\eta}^{(n)}) \right| \leq 2\rho(m\Delta_n).$$

Since $m\Delta_n = m2^{-n} \to 0$, the result follows. □

References

[1] Sachin Adlakha, Ramesh Johari, and Gabriel Y Weintraub. Equilibria of dynamic games with many players: Existence, approximation, and market structure. *Journal of Economic Theory*, 156:269–316, 2015.

[2] Charalambos D Aliprantis and Kim C Border. *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media, 2006.

[3] Erhan Bayraktar, Amarjit Budhiraja, and Asaf Cohen. A numerical scheme for a mean field game in some queueing systems based on markov chain approximation method. *SIAM Journal on Control and Optimization*, 56(6):4017–4044, 2018.

[4] Erhan Bayraktar, Amarjit Budhiraja, and Asaf Cohen. Rate control under heavy traffic with strategic servers. *Annals of Applied Probability*, 29(1):1–35, 2019.

[5] Erhan Bayraktar and Asaf Cohen. Analysis of a finite state many player game using its master equation. *SIAM Journal on Control and Optimization*, 56(5):3538–3568, 2018.

[6] Chiara Benazzoli, Luciano Campi, and Luca Di Persio. Mean field games with controlled jump–diffusion dynamics: Existence results and an illiquid interbank market model. *Stochastic Processes and their Applications*, 130(11):6927–6964, 2020.

[7] Dimitri Bertsekas and Steven E Shreve. *Stochastic optimal control: the discrete-time case*, volume 5. Athena Scientific, 1996.

[8] Charles Bertucci, Louis Bertucci, Jean-Michel Lasry, and Pierre-Louis Lions. A mean field game approach to bitcoin mining. *SIAM Journal on Financial Mathematics*, 15(3):960–987, 2024.

[9] Patrick Billingsley. *Convergence of Probability Measures*. Wiley, 2 edition, 1999.

[10] Anup Biswas. Mean field games with ergodic cost for discrete time Markov processes. *arXiv preprint arXiv:1510.08968*, 2015.

[11] Pierre Brémaud. *Point processes and queues*. Springer, 1981.

[12] René Carmona, François Delarue, et al. *Probabilistic theory of mean field games with applications I-II*. Springer, 2018.

[13] Mao Fabrice Djete. Mean field games of controls: on the convergence of Nash equilibria. *The Annals of Applied Probability*, 33(4):2824–2862, 2023.

[14] Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*. John Wiley & Sons, 2009.

[15] Liviu C Florescu and Christiane Godet-Thobie. *Young measures and compactness in measure spaces*. de Gruyter Berlin, 2012.

[16] Diogo A Gomes et al. Continuous time finite state mean field games. *Applied Mathematics & Optimization*, 68(1):99–143, 2013.

[17] Diogo A Gomes, Joana Mohr, and Rafael Rigao Souza. Discrete time, finite state space mean field games. *Journal de mathématiques pures et appliquées*, 93(3):308–328, 2010.

[18] Onésimo Hernández-Lerma and Jean B Lasserre. *Discrete-time Markov control processes: basic optimality criteria*, volume 30. Springer Science & Business Media, 2012.

[19] Werner Hildenbrand. Core and equilibria of a large economy (princeton, nj), 1974.

[20] K Hinderer. *Foundations of Non-stationary Dynamic Programming with Discrete Time Parameter*. Springer-Verlag, 1970.

[21] Minyi Huang, Roland P Malhamé, and Peter E Caines. Large population stochastic dynamic games: closed-loop Mckean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Informations and Systems*, 6(3):221–251, 2006.

[22] Boyan Jovanovic and Robert W Rosenthal. Anonymous sequential games. *Journal of Mathematical Economics*, 17(1):77–87, 1988.

[23] Harold Kushner. *Weak convergence methods and singularly perturbed stochastic control and filtering problems*. Springer Science & Business Media, 2012.

[24] Harold Joseph Kushner and Paul G Dupuis. *Numerical methods for stochastic control problems in continuous time*, volume 24. Springer Science & Business Media, 2001.

[25] Daniel Lacker. Mean field games via controlled martingale problems: existence of Markovian equilibria. *Stochastic Processes and their Applications*, 125(7):2856–2894, 2015.

[26] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese journal of mathematics*, 2(1):229–260, 2007.

[27] Jonathan Levin. Supermodular games. *Lectures Notes, Department of Economics, Stanford University*, 2003.

[28] Zongxi Li, A Max Reppen, and Ronnie Sircar. A mean field games model for cryptocurrency mining. *Management Science*, 70:2188–2208, 2024.

[29] Jun Moon and Tamer Başar. Discrete-time decentralized control using the risk-sensitive performance criterion in the large population regime: a mean field approach. In *2015 American Control Conference (ACC)*, pages 4779–4784. IEEE, 2015.

[30] Yuan-Hua Ni, Robert Elliott, and Xun Li. Discrete-time mean-field stochastic linear–quadratic optimal control problems, ii: Infinite horizon case. *Automatica*, 57:65–77, 2015.

[31] Mojtaba Nourian and Girish N Nair. Linear-quadratic-gaussian mean field games under high rate quantization. In *52nd IEEE Conference on Decision and Control*, pages 1898–1903. IEEE, 2013.

[32] Naci Saldi, Tamer Basar, and Maxim Raginsky. Markov–Nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization*, 56(6):4256–4287, 2018.

[33] Richard Serfozo. Convergence of lebesgue integrals with varying measures. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 380–402, 1982.

[34] Amirheckmat Taghizadeh, Hamed Kebriaei, and Dusit Niyato. Mean field game for equilibrium analysis of mining computational power in blockchains. *IEEE Internet of Things Journal*, 7(8):7625–7635, 2020.