

Mean Field Games of Control and Cryptocurrency Mining

Nicolás García^{*†}

Ronnie Sircar^{*}

H. Mete Soner^{*‡}

April 21, 2025

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.

Contents

1	Introduction	2
2	Mean Field Games in Discrete Time	3
2.1	<i>N</i> -Player Game Formulation	3
2.2	Mean Field Game Formulation	4
2.3	Model Assumptions	5
2.4	Existence of MFG Equilibria	6
3	Continuous-Time MFG Equilibria as Discrete-Time Limits	9
3.1	MFG of Controlled Intensity	9
3.2	Time-Discretized MFG & Convergence Theorem	11
4	Cryptocurrency Mining MFG	17
4.1	Discrete-Time Model	18
4.2	Numerical Computation of Discrete-Time MFG	20
4.3	Continuous Time Existence	21
4.4	Uniqueness and Sharpness of MFG Equilibrium	23
5	Conclusion	24

^{*}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.

A Tightness Theorem

24

1 Introduction

The study of Mean Field Games (MFGs) dates back to the foundational papers [16] and [21], and we refer to [9] 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 [13], which considers a finite state and horizon setup where agents control their transition probabilities, [7] 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, [25], [24], and [26] consider linear dynamics in a variety of discrete-time settings, and [27] 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 [23], we study a class of problems set both in discrete 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 [3, 14], 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. A probabilistic treatment of MFGs with jumps is given in [4], 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 [20] for diffusion MFGs of state and [10] for diffusion MFGs of control. Our results complement [6, 28], and directly apply to the cryptocurrency MFG model from [23] 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. Using the methodology from [27] which in turn was inspired by [17], 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 [27], 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 [11, 18, 19].

In Section 4, we apply our results from Sections 2 and 3 to establish MFG existence for the cryptocurrency mining MFG model of [23] 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 [23] using a finite difference scheme for the associated coupled PDEs. We remark on uniqueness of equilibrium in Section 4.4.

Notation: For an integer d , let \mathcal{S}^d denote the d -dimensional simplex, $[d] := \{1, 2, \dots, d\}$, and define $\Delta_i : \mathbb{R}^d \rightarrow \mathbb{R}^d$ by $\Delta_i x := (x_1 - x_i, x_2 - x_i, \dots, x_d - x_i)$. Given a Polish space \mathcal{X} , let $\mathcal{B}(\mathcal{X})$ denote its Borel sets and let $\mathcal{P}(\mathcal{X})$ denote the set of probability measures on $\mathcal{B}(\mathcal{X})$. We endow this space with the weak topology, i.e., $\mu_n \rightarrow \mu$ if $\lim_{n \rightarrow \infty} \int_{\mathcal{X}} f d\mu_n = \int_{\mathcal{X}} f d\mu$ for any continuous and bounded $f : \mathcal{X} \rightarrow \mathbb{R}$. To emphasize convergence in the weak sense, we write $\mu_n \xrightarrow{\mathcal{L}, n \rightarrow \infty} \mu$. Let $\mathbb{R}_{\geq 0}$ denote the non-negative reals. For a random variable X , $\mathcal{L}(X)$ denotes its distribution. For $\mu \in \mathcal{P}(\mathbb{R})$, we write $\bar{\mu} := \int_{\mathcal{X}} x \mu(dx)$, and $\delta_{\{a\}}$ denotes the Dirac measure located at a .

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 $t = 0, 1, 2, \dots, T$. We denote by x_t^i the state of agent $i \in [N] := \{1, 2, \dots, N\}$ at time t .
- The initial states x_0^i are i.i.d. drawn from a measure $\mu_0 \in \mathcal{P}(\mathcal{X})$.
- The action space \mathcal{A} is assumed to be Polish, $a_t^i \in \mathcal{A}$ denotes the action of agent $i \in [N]$ at time t , and $e_t := (e_t^c, e_t^s) \in \mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X})$ denotes the empirical distribution of the agents' controls and states at time t .
- The transition dynamics of each agent are given by a Markov transition kernel

$$\rho : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X}) \quad \text{so that} \quad x_{t+1}^i \sim \rho(\cdot | x_t^i, a_t^i, e_t^c, e_t^s)$$

for each $t = 0, 1, 2, \dots, T-1$.

A control policy for player i consists of a sequence $(\pi_t^i)_{t=0}^{T-1}$ of $\mathcal{P}(\mathcal{A})$ -valued random variables adapted to the filtration

$$\mathcal{F}_0^i := \sigma(x_0^i, e_0^s), \quad \mathcal{F}_t^i := \sigma\left(\mathcal{F}_{t-1}^i \cup \sigma(x_t^i, a_{t-1}^i, e_{t-1}^c, e_t^s)\right), \quad \text{for all } t = 0, 1, \dots, T-1.$$

Conditioned on \mathcal{F}_t , the action a_t^i of agent i is drawn randomly (and independent of any other random quantity) from the distribution π_t^i (i.e. $a_t^i \sim \pi_t^i$).

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

$$c : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X}) \rightarrow [0, \infty), \quad \varphi : \mathcal{X} \times \mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X}) \rightarrow [0, \infty).$$

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

$$J^i(\pi^{(N)}) := J^i(\pi^{N,i}, \pi^{N,-i}) := \mathbb{E}^{\pi^{(N)}} \left[\sum_{t=0}^{T-1} c(x_t^i, a_t^i, e_t) + \varphi(x_T^i, e_T) \right],$$

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

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

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

2.2 Mean Field Game Formulation

We proceed with the reference agent problem which characterizes the corresponding MFG. We denote the state and control of the reference agent at time t by $x_t \in \mathcal{X}$ and $a_t \in \mathcal{A}$, respectively. In this case, the reference agent has identical dynamics as in the N -player game, but the sequences of empirical measures $(e_t^s)_{t=0}^T$ and $(e_t^c)_{t=0}^T$ are taken to be deterministic. Denoting $\delta_t = (e_t^c, e_t^s) \in \mathcal{P}(\mathcal{A} \times \mathcal{X})$, the dynamics follow

$$x_{t+1} \sim \rho(\cdot | x_t, a_t, \delta_t), \quad a_t \sim \pi_t(\cdot | x_t), \quad t = 0, 1, \dots, T, \quad (2.1)$$

where, because $\delta := (\delta_t)_{t=0}^T$ 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:

Definition 2.1. *A control policy is called Markov if it is a sequence $(\pi_t)_{t=0}^{T-1}$ where each π_t is measurable w.r.t. $\sigma(x_t)$. Under such a policy, the measure $\pi_t(\cdot | x_t)$ used to generate the control at a given time t 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 [27, Proposition 3.2]). Given a fixed sequence of probability measures $\delta := (\delta_t)_{t=0}^T \subseteq \mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X})$, the reference agent's control problem consists of determining a Markov policy π^* such that

$$J(\pi^*, \delta) = \inf_{\pi \in \mathbb{M}} J(\pi, \delta) \quad \text{for} \quad J(\pi, \delta) := E^\pi \left[\sum_{t=0}^{T-1} c(x_t, a_t, \delta_t) + \varphi(x_T, \delta_T) \right],$$

with dynamics given by (2.1).

For convenience in defining an MFG equilibrium, we introduce

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

We additionally define a map $\Lambda : \Pi \rightarrow (\mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X}))^{T+1}$ by constructing $\Lambda(\pi) := (\delta_t^c, \delta_t^s)_{t=0}^T$ iteratively using the initial state law μ_0 as follows: First, set $\delta_0^s = \mu_0$.

For all other $t \geq 0$, define

$$\delta_t^c(\cdot) = \int_{\mathcal{X}} \mathcal{P}_t^\pi(\cdot | x) \delta_t^s(dx), \quad \delta_{t+1}^s(\cdot) = \int_{\mathcal{X}} \int_{\mathcal{A}} \rho(\cdot | x, a, \delta_t) \mathcal{P}_t^\pi(da | x) \delta_t^s(dx),$$

where $\mathcal{P}_t^\pi(\cdot | x)$ denotes the conditional law of a_t given the event $\{x_t = x\}$ 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 action 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 $(\pi, \delta) \in \mathbb{M} \times (\mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{X}))^{T-1}$ is called an *MFG equilibrium* if and only if

$$\pi \in \Phi(\delta) \quad \text{and} \quad \delta = \Lambda(\pi).$$

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. As we will rely on some technical arguments therein, we follow the notation and structure in the proof of [27, 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 existence of an MFG equilibrium. Fix a continuous moment $w : \mathcal{X} \rightarrow [1, \infty)$ 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 $\lim_{n \rightarrow \infty} \inf_{x \in \mathcal{X} \setminus H_n} w(x) = \infty$, and which satisfy $w(\cdot) \geq 1 + d_{\mathcal{X}}(\cdot, x_0)^p$ 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. We proceed by defining the v -norm of a map $g : \mathcal{X} \rightarrow \mathbb{R}$ by

$$\|g\|_v := \sup_{x \in \mathcal{X}} \frac{|g(x)|}{v(x)}.$$

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

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

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{A} is compact and \mathcal{X} is locally compact (every point contains a compact neighborhood).
- (iii) There exists a constant $\alpha \geq 1$ such that

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

- (iv) The stochastic kernel ρ is weakly continuous in the sense that if $(x_n, a_n, \delta_n^c, \delta_n^s) \xrightarrow{n \rightarrow \infty} (x, a, \delta^c, \delta^s)$, then $\rho(\cdot|(x_n, a_n, \delta_n^c, \delta_n^s)) \xrightarrow{n \rightarrow \infty} \rho(\cdot|(x, a, \delta^c, \delta^s))$ 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 $R \in \mathbb{R}$ satisfying

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

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 for the finite-horizon model (2.1).*

To prove this result (proof at the end of this section), we first construct a set-valued map whose fixed points correspond to equilibria. We then use the Kakutani fixed point theorem to complete the proof. We begin with a definition.

Definition 2.5. (Value Function) *Consider the non-homogeneous Markov Decision Process given by the transition kernel $\rho_t(\cdot|x, a, \delta_t)_{t=0}^N$ for the fixed flow of measures δ and costs as defined above. We let $V_t^\delta : \mathcal{X} \rightarrow \mathbb{R}$ denote the value function at time $t = 1, 2, \dots, T$. In other words, we define $V_T^\delta := \varphi$ and iteratively define*

$$V_t^\delta(x) := \min_{a \in \mathcal{A}} \left(c(x, a, \delta_t) + \int_{\mathcal{X}} V_{t+1}^\delta(y) \rho_t(dy|x, a, \delta_t) \right)$$

for $t = 0, 1, 2, \dots, T-1$. Furthermore, let $V^\delta := (V_t^\delta)_{t=0}^T$.

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

$$\begin{aligned} \mathcal{P}_v^t(\mathcal{X}) &:= \{\mu \in \mathcal{P}_v(\mathcal{X}) : \int_{\mathcal{X}} w(x) \mu(dx) \leq \alpha^t M\}, \\ \mathcal{P}_v^t(\mathcal{X} \times \mathcal{A}) &:= \{\mu \in \mathcal{P}(\mathcal{X} \times \mathcal{A}) : \mu_1 \in \mathcal{P}_v^t(\mathcal{X})\}, \end{aligned}$$

where in the second set, μ_1 denotes the state marginal of μ (i.e. $\mu_1(\cdot) = \mu(\cdot \times \mathcal{A})$). Taking R as in Assumption 2.3, define

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

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

$$C_v^t(\mathcal{X}) := \{u \in C_v(\mathcal{X}) : \|u\|_v \leq L_t\}, \quad C := \prod_{t=0}^T C_v^t(\mathcal{X}), \quad \Xi := \prod_{t=0}^T \mathcal{P}_v^t(\mathcal{X} \times \mathcal{A}).$$

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

$$\begin{aligned} \frac{1}{v(x)} V_t^\nu(x) &= \min_{a \in \mathcal{A}} \left[\frac{1}{v(x)} c(x, a, \nu_t) + \int_{\mathcal{X}} \frac{1}{v(x)} V_{t+1}^\nu(y) \rho_t(dy|x, a, \nu_t) \right] \\ &\leq \min_{a \in \mathcal{A}} \left[R + \frac{1}{v(x)} \int_{\mathcal{X}} L_{t+1} v(y) \rho_t(dy|x, a, \nu_t) \right] \\ &\leq R + \alpha L_{t+1} = L_t. \end{aligned}$$

By assumption, φ is continuous and thus inductively, and using [5, Proposition 7.32], which establishes that the Bellman operator (see the following definition) preserves continuity, continuity of all the value functions follows. \square

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

$$T_t^\nu u(x) := \min_{a \in \mathcal{A}} \left[c(x, a, \nu_t) + \int_{\mathcal{X}} u(y) \rho_t(dy|x, a, \nu_t) \right] \quad \text{for } t = 0, 1, 2, \dots, N,$$

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

$$(T^\nu u)_t := \begin{cases} T_t^\nu u_{t+1} & \text{for } t = 0, 1, \dots, T-1 \\ \varphi & \text{for } t = T. \end{cases}$$

Observe that

$$(TV^\nu)_t(x) = (T_t V_{t+1}^\nu)(x) = \begin{cases} \min_{a \in \mathcal{A}} \left[c(x, a, \nu_t) + \int_{\mathcal{X}} J_{t+1}^\nu(y) \rho_t(dy|x, a, \nu_t) \right] = V_t^\nu(x) & \text{for } t < T \\ \varphi(x) = V_T^\nu(x) & \text{for } t = T. \end{cases}$$

Thus 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 $t \geq 0$ the operator T_t^ν maps $C_v^t(\mathcal{X})$ into $C_v^{t+1}(\mathcal{X})$.

Proof. Let $u \in C_v^t(\mathcal{X})$. Since the Bellman operator preserves continuity (see [5, Proposition 7.32]) we immediately have that $T_t^\nu u$ is continuous. Because

$$\begin{aligned} \|T_t^\nu u\|_v &= \sup_{x \in \mathcal{X}} \frac{|(T_t^\nu u)(x)|}{v(x)} = \sup_{x \in \mathcal{X}} \frac{\left| \min_{a \in \mathcal{A}} \left[c(x, a, \nu_t) + \int_{\mathcal{X}} u(y) \rho(dy|x, a, \nu_t) \right] \right|}{v(x)} \\ &\leq \sup_{(x,a) \in \mathcal{X} \times \mathcal{A}} \frac{c(x, a, \nu_t) + \int_{\mathcal{X}} |u(y)| \rho(dy|x, a, \nu_t)}{v(x)} \\ &\leq \sup_{(x,a) \in \mathcal{X} \times \mathcal{A}} \frac{c(x, a, \nu_t) + L_t \int_{\mathcal{X}} v(y) \rho(dy|x, a, \nu_t)}{v(x)} \\ &\leq \sup_{(x,a) \in \mathcal{X} \times \mathcal{A}} \frac{Rv(x) + \alpha L_t v(x)}{v(x)} = L_{t+1}, \end{aligned}$$

it follows that $T_t^\nu u \in C_v^{t+1}(\mathcal{X})$, as claimed. \square

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

$$C(\nu) := \left\{ \nu' \in \mathcal{P}(\mathcal{X} \times \mathcal{A})^{T+1} : \nu'_{t+1,1}(\cdot) = \int_{\mathcal{X} \times \mathcal{A}} \rho(\cdot | x, a, \nu_t) \nu_t(dx, da) \text{ for every } t, \quad \nu'_{0,1} = \mu_0 \right\},$$

$$B(\nu) := \left\{ \nu' \in \mathcal{P}(\mathcal{X} \times \mathcal{A})^{T+1} : \right.$$

$$\left. \nu'_t \left(\left\{ (x, a) : c(x, a, \nu_t) + \int_{\mathcal{X}} V_{t+1}^\nu(y) \rho(dy | x, a, \nu_t) = T_t^\nu V_{t+1}^\nu(x) \right\} \right) = 1 \text{ for every } t \geq 0 \right\}.$$

Lemma 2.9. *For any fixed point ν of Γ , there exists a corresponding MFG equilibrium.*

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

$$\nu_t(dx, da) = \pi_t(da | x) \nu_t^s(dx) \quad \text{for every } t = 0, 1, 2, \dots, T.$$

Writing $\pi := (\pi_t)_{t=0}^{T-1}$, 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 maximizers of the Bellman operator (see [15, Theorem 17.1] for a proof of this result). \square

We are now ready to complete the existence proof which uses the restrictions placed on the state-component marginals in the definition of Ξ .

Proof of Theorem 2.4. To show that Γ admits a fixed point we make use of [2, 17.55 Corollary (Kakutani–Fan–Glicksberg)]. 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 $\mathcal{X} \times \mathcal{A}$.

We first show 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 $\nu'_t \in \mathcal{P}_v^t(\mathcal{X} \times \mathcal{A})$ for every $t \geq 0$. In other words, we need to check that $\nu'_{t,1} \in \mathcal{P}_v^t(\mathcal{X})$ for all $t \geq 0$. By definition, we know that $\nu'_{0,1} = \mu_0$ thus we have that

$$\int_{\mathcal{X}} w(x) \nu_0^{s'}(dx) = \int_{\mathcal{X}} w(x) \mu_0(dx) = \alpha^0 M,$$

and so by definition, $\nu'_{0,1} \in \mathcal{P}_v^0(\mathcal{X})$. For any other t , simply observe that by definition of $C(\nu)$, we have

$$\begin{aligned} \int_{\mathcal{X}} w(y) \nu'_{t+1,1}(dy) &= \int_{\mathcal{X} \times \mathcal{A}} \int_{\mathcal{X}} w(y) \rho(dy | x, a, \nu_t) \nu_t(dx, da) \\ &\leq \int_{\mathcal{X} \times \mathcal{A}} \alpha w(x) \nu_t(dx, da) = \int_{\mathcal{X}} w(x) \alpha \nu_t^s(dx) \leq \alpha^{t+1} M, \end{aligned}$$

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 $B(\nu)$.

We omit the remaining details here, as they follow by a simplification (since we are working with a finite instead of infinite product space) of the arguments in [27, Proposition 3.9 and Proposition 3.10], and we have used the notation therein. In particular, compactness and closedness of Ξ rely on the use of the continuous moment function w , and the closedness of the graph of Γ follow using analytic arguments. \square

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. 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 [23], with equilibria conjectured to exist from the convergence of numerical schemes. Secondly, 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 [23].

3.1 MFG of Controlled Intensity

Although we prove our convergence result for 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 makes use of the convergence methods used for stochastic control problems considered in detail in [19]. We will work with a weak formulation and make use of relaxed controls and compactness arguments.

Definition 3.1. (*Relaxed Control*) Let \mathcal{U} denote a compact subset of \mathbb{R} . A relaxed control taking values in \mathcal{U} is a random measure $m \in \mathcal{P}(\mathcal{U} \times [0, T])$ 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 taking values in $\mathcal{P}(\mathcal{U} \times [0, T])$ (each possibly defined on its own probability space), then their process laws admit a weak limit. This follows because $\mathcal{P}(\mathcal{U} \times [0, T])$, endowed with the topology of weak convergence of measures, is a compact Polish space. Thus $(\mathcal{L}(m^{(n)}))_{n=1}^\infty \subseteq \mathcal{P}(\mathcal{P}(\mathcal{U} \times [0, T]))$ is a tight family and, by Prokhorov's theorem, admits a weak limit. 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*) First, fix a flow of measure on the control space $\eta = (\eta_t)_{t \in [0, T]} \subseteq \mathcal{P}(\mathcal{U})$, an intensity function $\lambda : \mathcal{U}^2 \rightarrow \mathbb{R}_{\geq 0}$, and constants $c > 0$ and $r > 0$. We say that state, control, and jump processes X, m, N defined on a common filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, P)$

satisfy the dynamics considered herein iff

$$X_t = X_0 - \int_0^t \int_{\mathcal{U}} c \alpha m_t(d\alpha) dt + r N_t, \quad (3.1)$$

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

$$\lambda_t^{m,\eta} := \int_{\mathcal{U}} \int_{\mathcal{U}} \lambda(\alpha, h) m_t(d\alpha) \eta_t(dh),$$

and where $m = (m_t)_{t \in [0, T]}$ is a relaxed admissible control.

Remark 3.3. (*Intensity Control Representation*) Unlike the more classical controlled diffusion setup where a Brownian filtration is fixed apriori, 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 [18] 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 [8, 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 [23], 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 reference agent wealth process. In their model, a representative miner of a cryptocurrency under the proof-of-work protocol, such as Bitcoin, hashes at rate $\alpha \geq 0$ to affect the probability of success, namely discovery, which comes as a Poisson-type arrival with 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 [8, T9 Theorem] (due to Watanabe) of a stochastic intensity Poisson process will be very useful throughout.

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

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

is a martingale.

For a given intensity function $\lambda : \mathcal{U}^2 \rightarrow \mathbb{R}_{\geq 0}$, let \mathcal{T} denote a tuple $(\Omega, \mathcal{F}, \mathbb{F}, P, \eta, N, m, X)$, where

- (Ω, \mathcal{F}, P) is a probability space equipped with a filtration $\mathbb{F} := (\mathcal{F}_t)_{t \in [0, T]}$;
- $\eta := (\eta_t)_{t \in [0, T]} \subseteq \mathcal{P}(\mathcal{U})$ is a deterministic measure flow on the action space $\mathcal{U} \subseteq \mathbb{R}$;
- and $m = (m_t)_{t \in [0, T]}$ is a relaxed admissible control with corresponding state and jump processes X and N satisfying the dynamics in Definition 3.2.

Given a measure $\eta_t \in \mathcal{P}(\mathcal{A})$, we denote its first moment by $\bar{\eta}_t = \int h \eta_t(dh)$.

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

Definition 3.5. A relaxed MFG of controlled jump intensity for a given intensity function $\lambda : \mathcal{U}^2 \rightarrow \mathbb{R}_{\geq 0}$, terminal reward function φ and constants $c > 0$ and $r > 0$ is a tuple \mathcal{T} such that

- the control is optimal in the sense that for any other tuple $(\Omega', \mathcal{F}', \mathbb{P}', P', \eta, N', m', X')$ satisfying the previous conditions, we have that $E^{P'}[\varphi(X_T)] \geq E^{P'}[\varphi(X'_T)]$;
- and consistency holds, i.e. $E[\bar{m}_t] = \bar{\eta}_t$ for Lebesgue almost every $t \in [0, T]$.

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 $\omega \in \mathcal{F}$ the measure $m_t(\omega)$ is a Dirac for Lebesgue almost every $t \in [0, T]$, we say that the MFG equilibria is of sharp controls.

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) discretizations approximations, parameterized by $n \in \mathbb{N}$.

Definition 3.6. (Discretization Scheme) For each $n \in \mathbb{N}$, the discrete-time, finite horizon state process $(x_k^{(n)})_{k=0}^{T2^n}$ follows

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

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

$$\text{Scheme 1: } \alpha_k^{(n)} = \frac{1}{2^n} \int \lambda(a_k^{(n)}, dh) \eta_k^{(n)}(dh), \quad \text{Scheme 2: } \alpha_k^{(n)} = \frac{1}{2^n} \lambda(a_k^{(n)}, \bar{\eta}_k^{(n)}).$$

It is assumed that $x_0^{(n)} \sim \mu_0$. The optimality criterion is the maximization of $E[\varphi(x_{T2^n}^{(n)})]$.

In either scheme, $(\eta_k)_{k=0}^{2^n T} \subseteq \mathcal{P}(\mathcal{U})$, 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 . The convergence theorem will hold under different subsets of the following assumptions.

Assumption 3.7. 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 continuous and bounded; and $\lambda(\cdot, h)$ is Lipschitz uniformly over the choice of $h \in \mathcal{U}$.

4. φ is strictly increasing and $\lambda(\cdot, h)$ is strictly concave for fixed $h \in \mathcal{U}$.

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

Theorem 3.8. Suppose that 1, 2, and 3 from Assumption 3.7 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 equilibria is of sharp control.

The proof is organized in four steps. Along the way, we prove four technical lemmas. 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 measures. 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 measures 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. 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 $E[\varphi(y_{T2^n}^{(n)})]$.

We now define new discretizations under Definition 3.6 Scheme 1 with state process denoted by $x^{(n)}$ and with parametrizing sequence $\eta_k^{(n)} := \delta_{\{\zeta_k^{(n)}\}}$, $k = 0, 1, \dots, 2^n T - 1$. Here, $x^{(n)}$ will be the discrete-time dynamics to be interpolated. We will slightly abuse notation by denoting the continuous-time control and background flows with the same letters as in discrete-time, 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 $t \in [0, T]$:

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

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:

$$X_t^{(n)} = X_0^{(n)} - c \int_0^t a_s^{(n)} ds + r N_t^{(n)} = x_0^{(n)} - c \int_0^t \int_{\mathcal{U}} a m_s^{(n)}(da) ds + r N_t^{(n)}, \quad (3.2)$$

where $N_t^{(n)}$ is a jump process taking values in $\{0, 1, 2, 3, \dots\}$, with unit jumps possible only on dyadic rationals of order n . Note that the random variables used to define the above continuous processes are a countable family. As such, we assume them to be defined on a common probability space (Ω, \mathcal{F}, P) .

Next, we extract limit points from the joint process laws $\mathcal{L}(X^{(n)}, m^{(n)}, N^{(n)})$ as $n \rightarrow \infty$. 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 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 [19, 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 $\mathcal{L}(m^{(n)})_{n=1}^\infty$ 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

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

is tight and admits a weakly converging sub-sequence. Similarly, we can extract a weak limit from the sequence $(\eta^{(n)})_{n=1}^\infty$ (since these are measures on the compact space $[0, T] \times \mathcal{U}$), and thus we can always find a sub-sequence along which both sequences converge. For simplicity of notation, we will dispense with this sub-sequence and assume that the convergence is as $n \rightarrow \infty$.

By the Skorokhod representation theorem (see [19, Theorem 9.1.7]), there exists some probability space (Ω, \mathcal{F}, P) supporting random variables $(\tilde{X}^{(n)}, \tilde{m}^{(n)}, \tilde{N}^{(n)})_{n=1}^\infty$ converging almost surely to (X, m, N) and such that $\mathcal{L}(\tilde{X}^{(n)}, \tilde{m}^{(n)}, \tilde{N}^{(n)}) = \mathcal{L}(X^{(n)}, m^{(n)}, N^{(n)})$. It is immediate that $(\tilde{X}^{(n)}, \tilde{m}^{(n)}, \tilde{N}^{(n)})$ satisfies the representation (3.2), and hereafter we abuse notation by dropping tilde.

To summarize, we have that for P -almost every ω , $X^{(n)}(\omega) \xrightarrow{n \rightarrow \infty} X(\omega)$ w.r.t the Skorokhod topology on $D[0, T]$, $m^{(n)}(\omega) \xrightarrow{n \rightarrow \infty} m(\omega)$ in the weak topology, and redundantly, $N^{(n)}(\omega) \xrightarrow{n \rightarrow \infty} N$ in the Skorokhod topology on $D[0, T]$. The $\eta^{(n)}$'s are deterministic and $\eta^{(n)} \xrightarrow{\mathcal{L}, n \rightarrow \infty} \eta$.

We define:

$$\begin{aligned} \mathcal{F}_t &:= \sigma(X_s, m_s, N_s : s \leq t) \quad \text{for every } t \in [0, T], \\ \mathcal{F}_t^{(n)} &:= \sigma(X_s^{(n)}, m_s^{(n)}, N_s^{(n)} : s \leq t) \quad \text{for every } t \in [0, T], n \in \mathbb{N}. \end{aligned}$$

It is important to note that we can always define a derivative $m_t(\omega)$ such that $m_t(\omega)(A)$ is $(\mathcal{F}_t)_{t \in [0, T]}$ adapted for each $t \in [0, T]$ and $A \in \mathcal{B}([0, L])$, and such that the disintegration $m(\omega)(da, dt) = m_t(\omega)(da)dt$ holds [19, Section 9.5]. Similarly we can also do this for approximations, with the Borel measurability of maps of the form $t \mapsto E[\bar{m}_t] = \int \alpha m_t(\omega)(d\alpha) dP(\omega)$ following from progressive measurability (see for example [12, Theorem 3.1]).

Step 2 (Characterizing the Limit Point): Using the almost sure convergence, it follows that (X, m, N) satisfy

$$\begin{aligned} X_t(\omega) &= \lim_{n \rightarrow \infty} X_t^{(n)}(\omega) = \lim_{n \rightarrow \infty} \left(X_0^{(n)}(\omega) - c \int_0^t \int_{\mathcal{U}} a m_s^{(n)}(\omega)(da) ds + r N_t^{(n)}(\omega) \right) \\ &= X_0(\omega) - c \int_0^t \int_{\mathcal{U}} a m_s(\omega)(da) ds + r N_t(\omega). \end{aligned}$$

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 [19, Equation 10.1.8]. We verify that it has the correct (stochastic) intensity by fixing $0 < s < t$ and denoting

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

Then we have

$$\mathbb{E} \left[N_t - J_{0,t} \middle| \mathcal{F}_s \right] = \mathbb{E} \left[N_t - N_s - J_{s,t} \middle| \mathcal{F}_s \right] + N_s - J_{0,s}$$

$$\begin{aligned}
&= \lim_n E \left[N_t^{(n)} - N_s^{(n)} - \int_s^t \int_{\mathcal{U}} \int_{\mathcal{U}} \lambda(a, h) m_\rho^{(n)}(da) \eta_\rho^{(n)}(dh) d\rho \middle| \mathcal{F}_s \right] + N_s - J_{0,s} \\
&= N_s - J_{0,s},
\end{aligned}$$

where the limit can be seen to equal zero from the construction of the approximations.

Step 3 (Consistency): We state consistency as a Lemma.

Lemma 3.9. (*Consistency of the limit*) For Lebesgue-almost every $t \in [0, T]$ we have that $E[\bar{m}_t] = \bar{\eta}_t$.

Proof. As discussed earlier, Borel measurability of the map $t \mapsto E[\bar{m}_t]$ follows from admissibility of the control process m . Observe that for any Borel set $B \subseteq [0, T]$ we can find a sequence $C([0, T]; \mathbb{R}) \ni f_l \xrightarrow{l \rightarrow \infty} \mathbb{1}_B$ in L^1 with $\|f_l\|_\infty \leq 1$ for every l , from which, using the triangle inequality, we see that

$$\begin{aligned}
\left| \int_B E[\bar{m}_t] dt - \int_B \bar{\eta}_t dt \right| &\leq \left| \int \mathbb{1}_B(t) E[\bar{m}_t] dt - \int f_l(t) E[\bar{m}_t] dt \right| \\
&\quad + \left| \int f_l(t) E[\bar{m}_t] dt - \int f_l(t) E[\bar{m}_t^{(n)}] dt \right| \\
&\quad + \left| \int f_l(t) E[\bar{m}_t^{(n)}] dt - \int \int f_l(t) a \eta_t^{(n)}(da) dt \right| \\
&\quad + \left| \int \int f_l(t) a \eta_t^{(n)}(da) dt - \int \int f_l(t) a \eta_t(da) dt \right| \\
&\quad + \left| \int \int f_l(t) a \eta_t(da) dt - \int \int \mathbb{1}_B(t) a \eta_t(da) dt \right| \\
&\xrightarrow{n, l \rightarrow \infty} 0,
\end{aligned}$$

where terms one and five converge due to the L^1 approximation, term three is identically zero by consistency of the discrete-time chains, and term four converges to zero by definition of weak convergence. For the second term, we rewrite as

$$\begin{aligned}
&\left| \int f_l(t) \int_\Omega \int am_t(\omega)(da) dP(\omega) dt - \int f_l(t) \int_\Omega \int am_t^{(n)}(\omega)(da) dP(\omega) dt \right| \\
&= \left| \int_\Omega \left[\int \int f_l(t) am_t(\omega)(da) dt - \int \int f_l(t) am_t^{(n)}(\omega)(da) dt \right] dP(\omega) \right| \\
&= \left| \int_\Omega \left[\int f_l(t) am(\omega)(da, dt) - \int f_l(t) am^{(n)}(\omega)(da, dt) \right] dP(\omega) \right|,
\end{aligned}$$

from which we see that the integrand converges to zero for almost every ω since $m_t^{(n)}(\omega) \xrightarrow{n \rightarrow \infty} m_t(\omega)$ for P -almost every ω in the topology of weak convergence. An application of the Dominated Convergence Theorem yields that term two also converges to zero. Consistency therefore follows. \square

Step 4 (Establishing Optimality): For a fixed discretization integer n and sequence of measures $\eta^{(n)} := (\eta^{(n)})_{k=0}^{T2^n-1}$, the expected cost of an adapted control chain $a := (a_k)_{k=0}^{2^n T}$ and the value function of the discrete-time system (dynamics given by Scheme 1 in Definition 3.6 with $\eta^{(n)}$ the parameterizing sequence) are denoted by

$$w^{(n)}(k, x, a, \eta^{(n)}) = E[\varphi(x_{2^n T}^{(n), a}) | x_k^{(n), a} = x], \quad v^{(n)}(k, x, \eta^{(n)}) = \sup_a w^{(n)}(k, x, a, \eta^{(n)}),$$

where the superscript a on the state emphasizes that a is the driving control. Analogously, we use capital letters for the equivalent continuous time quantities

$$W(t, x, m, \eta) = E[\varphi(X_T^m) | X_t = x], \quad V(t, x, \eta) = \sup_m W(t, x, m, \eta),$$

where this time $m := (m_t)_{t=0}^T$ denotes a relaxed admissible (possibly random) control process and the underlying probability space may vary with m . The dynamics of X are of course constrained to be a jump process of controlled jump intensity from Definition 3.5. The following lemma (which follows the construction for controlled diffusion in [18, Theorem 3.5.2]) establishes that allowing for relaxed controls does not improve the value function. (We will of course be interested in the map $t \mapsto \lambda(t, \eta_t)$ for a fixed flow of measure $(\eta_t)_{t \in [0, T]}$, where the objective is the maximization of expectation of continuous terminal wealth utility).

Lemma 3.10. (*Chattering Lemma*) *Consider the continuous-time intensity control problem from Definition 3.5 with compact action space \mathcal{U} and time-varying intensity function $f_t : \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$, assumed continuous for each $t \in [0, T]$. For any $\gamma > 0$, there exists a finite set $\mathcal{U}^\gamma := \{\alpha_1^\gamma, \dots, \alpha_k^\gamma\} \subseteq \mathcal{U}$ such that, for any admissible relaxed control m and associated jump and state processes N and X all defined on some probability space (Ω, \mathcal{F}, P) , there exists a piecewise constant control u_m^γ , with discontinuities only at dyadic rationals of a fixed order (and possibly defined on another probability space), such that*

$$|W(t, x, m) - W(t, x, u_m^\gamma)| < \gamma.$$

Proof. We first prove it without the dyadic rational requirement. First, fix $\rho > 0$ and let $B_1^\rho, B_2^\rho, \dots, B_{N_\rho}^\rho$ denote a disjoint partition of \mathcal{U} with the diameter of each of the sets not exceeding ρ . For a stochastic admissible control m , a fixed $\Delta > 0$, and an integer i , define the following (random) numbers:

$$r_{i,j}^{\rho,\Delta}(\omega) := \int_{i\Delta}^{(i+1)\Delta} m_s(\omega)(B_j^\rho) ds = m(\omega)(B_j^\rho \times [i\Delta, (i+1)\Delta]) \quad \text{for } j = 1, 2, \dots, N_\rho.$$

Observe that, by construction, we have $\Delta = \sum_{j=1}^{N_\rho} r_{i,j}^{\rho,\Delta}$. Next, subdivide the next interval (to preserve causality) $[(i+1)\Delta, (i+2)\Delta]$ into intervals of length given by the $r_{i,j}^{\rho,\Delta}$'s. On the interval of length $r_{i,j}^{\rho,\Delta}$ define $u_m^{\rho,\Delta}$ to take on the value α_j^ρ . Note that here, the r 's are dependent on ω and so the resulting ordinary control is itself stochastic. Let $m^{\rho,\Delta}$ denote the relaxed representation of this control. Note that the piecewise constant control is defined on the original probability space (Ω, \mathcal{F}, P) . By possibly augmenting the probability space or modifying the probability measure P (recall Remark 3.3), let $N^{\rho,\Delta}$ denote another jump process with stochastic intensity given by $\lambda_t^{\rho,\Delta} = \int_{\mathcal{U}} f_t(\alpha) m_t^{\rho,\Delta}(d\alpha)$ and let $X^{\rho,\Delta}$ and $N^{\rho,\Delta}$ satisfy

$$X_t^{\rho,\Delta} = X_0 - \int_0^t c \alpha m_s^{\rho,\Delta}(d\alpha) ds + r N_t^{\rho,\Delta}.$$

By construction, the laws $\{\mathcal{L}(m^{\rho,\Delta}, N^{\rho,\Delta}) : \rho > 0, \Delta > 0\}$ are tight and we let \tilde{m} and \tilde{N} denote a limit point. Clearly, $\mathcal{L}(\tilde{m}) = \mathcal{L}(m)$ by construction, thus it follows that the limit point of the above is in fact the original law $\mathcal{L}(m, N)$. From here, we simply observe that for any fixed ρ and Δ , we can define an integer $z(\rho, \Delta)$ and round the jump times of the control up to the nearest $z(\rho, \Delta)$ 'th order dyadic rational (i.e. an element in $0, 1/2^{z(\rho,\Delta)}, 2/2^{z(\rho,\Delta)}, 3/2^{z(\rho,\Delta)}, \dots$). Taking $z(\rho, \Delta) \xrightarrow{\rho, \Delta \rightarrow 0} \infty$ sufficiently fast, the arguments above remain true. \square

It follows from the above that there is no gain in expanding the class of controls in terms of improving the value function. We prove another useful lemma:

Lemma 3.11. *Recall that $V(t, x, \eta)$ denotes the value function for the system with a fixed measure flow $\eta_{t \in [0, T]}$. If $\eta^{(n)} \xrightarrow{\mathcal{L}, n \rightarrow \infty} \eta$, then we have that $V(\cdot, \cdot, \eta^{(n)}) \xrightarrow{n \rightarrow \infty} V(\cdot, \cdot, \eta)$ point-wise.*

Proof. Let $(a_t)_{t \in [0, T]}$ be a control process defined on some probability space (Ω, \mathcal{F}, P) which supports two jump processes with stochastic rates $\int \lambda(a_t, h) \eta_t(dh)$ and $\int \lambda(a_t, h) \eta_t^{(n)}(dh)$. Fix $\omega \in \Omega$ and let $C([0, T]; [0, L]) \ni f_k \xrightarrow{L^1, k \rightarrow \infty} a(\omega)$. For any $t \in [0, T]$

$$\begin{aligned} & \left| \int_t^T \int \lambda(a_s(\omega), h) \eta_s(dh) ds - \int_t^T \int \lambda(a_s(\omega), h) \eta_s^{(n)}(dh) ds \right| \\ & \leq \left| \int_t^T \int \lambda(a_s(\omega), h) \eta_s(dh) ds - \int_t^T \int \lambda(f_k(s), h) \eta_s(dh) ds \right| \\ & \quad + \left| \int_t^T \int \lambda(f_k(s), h) \eta_s(dh) ds - \int_t^T \int \lambda(f_k(s), h) \eta_s^{(n)}(dh) ds \right| \\ & \quad + \left| \int_t^T \int \lambda(f_k(s), h) \eta_s^{(n)}(dh) ds - \int_t^T \int \lambda(a_s(\omega), h) \eta_s^{(n)}(dh) ds \right|, \end{aligned}$$

with terms one and three going to zero using the uniform Lipschitz condition on $\lambda(\cdot, h)$ and the L^1 convergence, and the second term going to zero by definition of weak convergence. Because the intensities converge to one another almost surely, so does the expected terminal reward, establishing the convergence of the claim. \square

We now connect the value functions of discrete time (considering Scheme 1 from Definition 3.6) and continuous time systems. When using a flow $(\eta_t)_{t \in [0, T]}$ of control measure to specify transition dynamics of the discrete time systems, we discretize the measures by averaging, as follows:

$$\eta(n)_k(B) = 2^n \int_{k/2^n}^{(k+1)/2^n} \eta_t(B) dt \text{ for } k = 0, 1, \dots, 2^n T - 1, \quad \eta(0)^{(n)} = \eta_0. \quad (3.3)$$

Lemma 3.12. *Let $(a_t)_{t=0}^T$ denote a control process which is piece-wise constant, right-continuous with left limits, and jumps only on dyadic rationals of order n' . Then letting $a^{(n)}$ denote its discretization (i.e. sampling) on dyadic rationals of order n we have that*

$$\left| w^{(n)}(a^{(n)}, \eta) - w^{(n)}(a^{(n)}, \eta^{(n)}) \right| \xrightarrow{n \rightarrow \infty} 0,$$

and

$$|w^{(n)}(a^{(n)}, \eta) - W(a, \eta)| \xrightarrow{n \rightarrow \infty} 0.$$

Proof. For the first limit, we complete the proof by bounding the difference between the expected rewards earner corresponding to the a continuous-time interval of the form $[l/2^{n'}, (l+1)/2^{n'}]$. Note that for a given ω , $a_{l/2^{n-n'}}^{(n)}(\omega)$ is constant in n . For $n > n'$ we have

$$\left| \sum_{k=0}^{2^{n-n'}-1} \frac{1}{2^n} \int_{\mathcal{U}} \lambda(a_{l/2^{n-n'}}^{(n)}, h) \eta(n)_{l/2^{n-n'}+k}(dh) - \sum_{k=0}^{2^{n-n'}-1} \frac{1}{2^n} \int_{\mathcal{U}} \lambda(a_{l/2^{n-n'}}^{(n)}, h) \eta^{(n)}(n)_{l/2^{n-n'}+k}(dh) \right|$$

$$\begin{aligned}
&= \left| \sum_{k=0}^{2^{n-n'}-1} \int_{l/2^{n'}+k/2^n}^{l/2^{n'}+(k+1)/2^n} \int \lambda(a_{l2^{n-n'}}, h) \eta_t(dh) dt - \sum_{k=0}^{2^{n-n'}-1} \int_{l/2^{n'}+k/2^n}^{l/2^{n'}+(k+1)/2^n} \int_{\mathcal{U}} \lambda(a_{l2^{n-n'}}, h) \eta_t^{(n)}(dh) dt \right| \\
&= \left| \int_{l/2^{n'}}^{(l+1)/2^{n'}} \int \lambda(a_{l2^{n-n'}}, h) \eta_t^{(n)}(dh) dt - \int_{l/2^{n'}}^{(l+1)/2^{n'}} \int_{\mathcal{U}} \lambda(a_{l2^{n-n'}}, h) \eta_t(dh) dt \right| \xrightarrow{n \rightarrow \infty} 0,
\end{aligned}$$

from which the result follows. The second limit follows immediately from the Poisson limit theorem. \square

Proof of Theorem 3.8. We can finally conclude optimality. Recall that η and m denote the limit flow of measure resp. control obtained by taking a weak limit of the interpolated discrete-time approximations. Suppose now for a contradiction that m is not optimal for η . Then using the Chattering Lemma 3.10 there exists a piece-wise constant control \tilde{a} (with jumps on finite dyadics no finer than of order n') such that $W(m, \eta) < W(\tilde{a}, \eta)$. We sample \tilde{a} for each n at n -th order dyadics to obtain a discrete control process $\tilde{a}^{(n)}$. Combining the two limits established in Lemma 3.12, we have that

$$\lim_{n \rightarrow \infty} w^{(n)}(\tilde{a}^{(n)}, \eta^{(n)}) = W(\tilde{a}, \eta) > W(m, \eta) = \lim_{n \rightarrow \infty} w^{(n)}(a^{(n)}, \eta^{(n)}),$$

which for all n sufficiently large, contradicts the optimality of $a^{(n)}$ for the fixed sequence of measures $\eta^{(n)}$ (note that the discretization scheme from (3.3) applied to $(\eta_t^{(n)})_{t \in [0, T]}$ recovers the original sequence $(\eta_k^{(n)})_{k=0}^{2^n T-1}$ that was interpolated to obtain $(\eta_t^{(n)})_{t \in [0, T]}$). As such, we conclude that m is indeed optimal for η .

We have established the first assertion of the theorem. For the second assertion regarding sharpness of the limit control m , assume φ is strictly increasing and that $\lambda(\cdot, h)$ is strictly concave for each h . Let (m, η) denote a relaxed MFG equilibrium. Fix $t \in [0, T]$, define $a_t^* = \int a m_t(da)$ where $dm = dm_t dt$, and note that

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

Thus, letting X^{a^*} and X^{m^*} denote the wealth processes (possibly defined on distinct probability spaces) driven by the controls $\delta_{\{a_t^*\}}$ and m respectively, we have that $E[\varphi(X_T^{a^*})] \geq E[\varphi(X_T^{m^*})]$, with equality if and only if m_t is a Dirac mass for Lebesgue almost every $t \in [0, T]$ almost surely. By optimality of m , it follows immediately that m is itself (up to redefining on Lebesgue null sets) a sharp control. This concludes the proof. \square

4 Cryptocurrency Mining MFG

We now return to the cryptocurrency mining MFG model from [23]. 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 [23] 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 [23]. 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

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

for a constant $M > 0$ (which proxies for the number of miners, see [23] for details). Given a fixed sequence of probability measures $(\zeta_k^{(n, \epsilon)})_{k=0}^{2^n T - 1} \subseteq \mathcal{P}(\mathcal{A})$, the transition dynamics for the wealth process $(x_k^{(n, \epsilon)})_{k=0}^{2^n T}$ are given by

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

for $k = 0, 1, \dots, 2^n T - 1$. The control objective is the maximization of expected terminal utility of wealth, $E[\varphi(x_{2^n T}^{(n, \epsilon)})]$. 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 $2^n T r$ 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 $\epsilon \rightarrow 0$.

Fix n and $\epsilon > 0$, and let $\nu^{(n, \epsilon)} := (\nu_k^{(n, \epsilon)})_{k=0}^{2^n T}$ denote a sequence of laws on $\mathcal{X} \times \mathcal{A}$ 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 $\mathcal{A} = [0, L]$. 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 $\nu^{(n, \epsilon)} \rightarrow \nu^{(n)}$ as $\epsilon \rightarrow 0$ 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 begin by observing that if

$$B := \{k \in \{0, 1, \dots, 2^n T - 1\} : \eta_k = \delta_{\{0\}}\}$$

is non-empty, then an optimal control for the original ($\epsilon = 0$) model with fixed background hash-rate $(\bar{\eta}_k)_{k \in \{0, 1, 2, \dots, 2^n T\}}$ 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. 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 $\epsilon \rightarrow 0$ with n fixed but arbitrary) sequence of measures on the control space $(\eta_k)_{k=0}^{2^n T} = (\nu_{k,2})_{k=0}^{2^n T} \subseteq \mathcal{P}([0, L])$ is never a Dirac at zero. In fact, there exists some $d > 0$ such that $\bar{\eta}_k > d$ for every $k = 0, 1, \dots, 2^n T - 1$ 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 $\bar{\eta}_k = 0$. It follows that $\lim_{\epsilon \rightarrow 0} \bar{\eta}_k^{(\epsilon)} = 0$. Recall that $a^{(\epsilon)}$ is an optimal control for the fixed background hash-rate $\eta^{(\epsilon)}$. For a given ϵ , let $\nu(\epsilon) := \epsilon \vee \bar{\eta}_k^{(\epsilon)}$, 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 $j = 0, 1, \dots, 2^n T - 1$. 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

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

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

$$E\left[\frac{1}{2^n} \lambda^{(\epsilon)}(\tilde{a}_k^{(\epsilon)}, \bar{\eta}_k^{(\epsilon)})\right] \geq \frac{\sqrt{\nu(\epsilon)}}{2^n(\sqrt{\nu(\epsilon)} + \bar{\eta}_k^{(\epsilon)} M + \epsilon)} \geq \frac{\sqrt{\nu(\epsilon)}}{2^n(\sqrt{\nu(\epsilon)} + \nu(\epsilon)(M+1))} \xrightarrow{\epsilon \rightarrow 0} 2^{-n}.$$

Observe that the difference in costs between the two controls converge as $\epsilon \rightarrow 0$. 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 $\epsilon \rightarrow 0$ results in a converging of the expected costs of the controls \tilde{a}^ϵ and a^ϵ , whereas the number of rewards of the former is strictly higher than the latter, with the difference not converging as $\epsilon \rightarrow 0$. 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

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

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 $\epsilon \rightarrow 0$. \square

We are now ready to establish 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 $\epsilon \rightarrow 0$. This object constitutes an MFG equilibrium for the model with $\epsilon = 0$.*

Proof. We have shown that none of the measures in the sequence $(\eta_k)_{k=0}^{2^n T}$ 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

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

and the weak convergence $\nu^{(\epsilon)} \xrightarrow{\epsilon \rightarrow 0} \nu$, it follows that $E[\varphi(x_{2^n T}^{(\epsilon)})] \xrightarrow{\epsilon \rightarrow 0} E[\varphi(x_{2^n T})]$. Suppose now for a contradiction that there exists a control policy $\tilde{\pi} = (\tilde{\pi}_k)_{k=0}^{2^n T-1}$ which outperforms a under the fixed background sequence $(\eta_k)_{k=0}^{2^n T-1}$. 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 $\bar{\eta}_k^{(\epsilon)} \xrightarrow{\epsilon \rightarrow 0} \bar{\eta}_k > 0$ for every $k = 0, \dots, 2^n T$, one obtains that $E[\varphi(\tilde{x}_{2^n T}^{(\epsilon, \eta^{(\epsilon)})})] \xrightarrow{\epsilon \rightarrow 0} E[\varphi(\tilde{x}_{2^n T}^{(0, \eta)})] > E[\varphi(x_{2^n T})]$ 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 $\epsilon > 0$. \square

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 from 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 reference agent (plotted below).

The results obtained are qualitatively similar to those obtained from the numerical PDE approach in [23], 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 uniqueness of the equilibria.




Figure 1: Evolution of Wealth Distribution

4.3 Continuous Time Existence

To apply the convergence result Theorem 3.8 (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 3.8 applies, guaranteeing the existence of a relaxed MFG equilibrium (of *a priori* relaxed controls) for the continuous-time game. As in the discrete-time case, we turn to the problem of extracting a limit for the original model by taking $\epsilon \rightarrow 0$. 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. Follows from the second assertion of Theorem 3.8. □

Suppose now that for a given $\epsilon > 0$, $(X^{(\epsilon)}, N^{(\epsilon)}, m^{(\epsilon)}, \eta^{(\epsilon)})$ 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 $\epsilon \rightarrow 0$. Let (X, N, m)

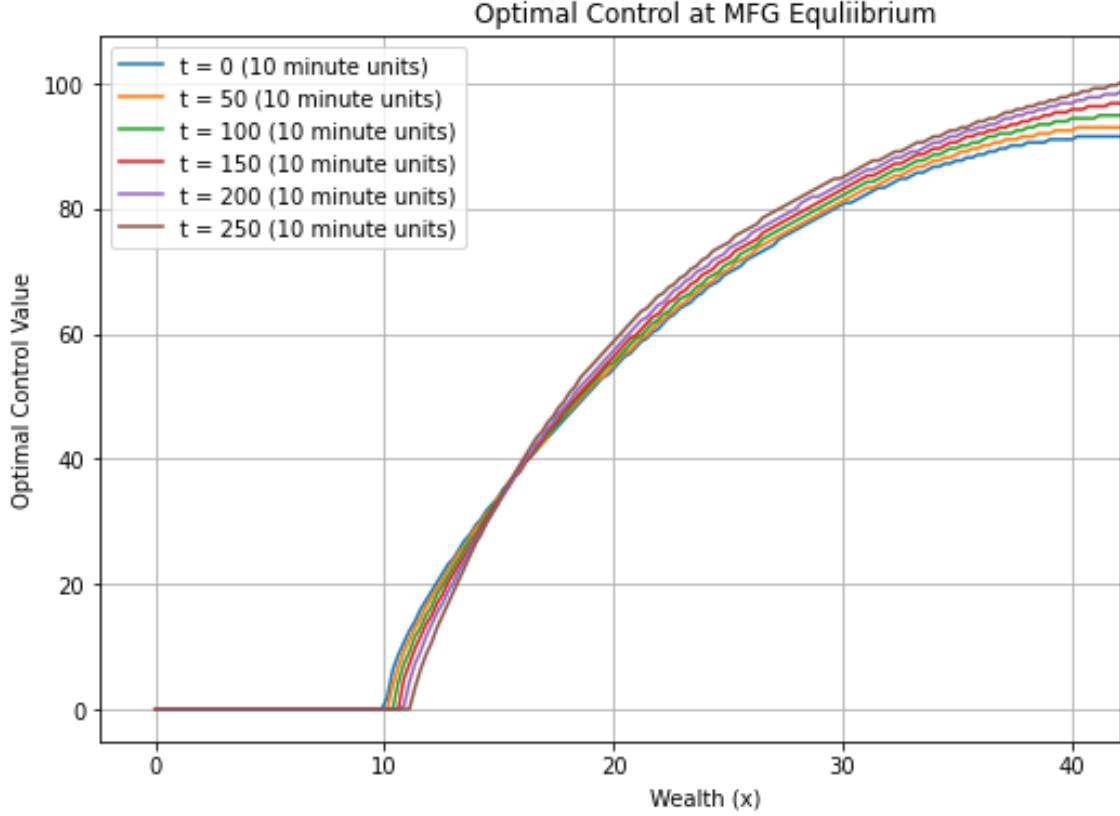


Figure 2: Optimal Control at Equilibrium

denote processes (defined on a possibly distinct probability space) such that

$$\mathcal{L}(X^{(\epsilon)}, N^{(\epsilon)}, m^{(\epsilon)}) \xrightarrow{\mathcal{L}, \epsilon \rightarrow 0} \mathcal{L}(X, N, m). \quad (4.3)$$

We need to establish that (X, N, m, η) 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.

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 $\bar{\eta}^{(n, \epsilon)} \in [0, L]^{n+1}$ into a space of measures, and taking $n \rightarrow \infty$ to obtain $\eta^{(\epsilon)} \in \mathcal{P}([0, T] \times [0, L])$. Then, we take $\epsilon \rightarrow 0$ to obtain η . By Proposition 4.2, however, we know that there is some $d > 0$ such that $\bar{\eta}_t^\epsilon > 0$ for every $t \in [0, T]$ and every ϵ sufficiently small, from which the result follows. \square

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 $\bar{\eta}_t \geq d$ 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

$$D([0, T]) \times \mathcal{P}([0, T] \times [0, L]) \ni (N, m) \mapsto \varphi\left(rN_T - c \int_0^T am(dt, da)\right). \quad (4.4)$$

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)

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

where $\tilde{N}^{(\epsilon)}$ denotes a unit jump process with stochastic intensity $\int \lambda^{(\epsilon)}(a, h) \tilde{m}_t(da) \eta_t^{(\epsilon)}(dh)$. The above limit is justified via the continuity and boundedness of the map in (4.4) combined with the fact the fact that η_t is never a Dirac at zero. For $\epsilon > 0$ sufficiently small, this contradicts the optimality of $m^{(\epsilon)}$ for the hash-rate $\eta^{(\epsilon)}$. 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. \square

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 reference agent problem for a given $n \in \mathbb{N}$ with fixed population hash-rates $(\bar{\eta}_k^{(n)})_{k=0}^{2^n T}$ 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 strict concavity of the value functions, one can show that

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

is a singleton. This follows by observing that the objective function in the above maximization is strictly quasi-concave. Moreover, using [22, Theorem 1] one can see that the quantity $a_k^{(n)*}(x)$ is strictly decreasing in $\bar{\eta}_k^{(n)}$ 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 behaviour) with numerical solutions to the original continuous-time cryptocurrency mining MFG which was solved by numerically solving coupled Kolmogorov and HJB PDEs.

A Tightness Theorem

The following result is restated here for convenience, and its proof can be found in [19, Theorem 9.2.1].

Theorem A.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 (Ω, \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

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

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

References

- [1] S. Adlakha, R. Johari, and G. Y. Weintraub. Equilibria of dynamic games with many players: Existence, approximation, and market structure. *Journal of Economic Theory*, 156:269–316, 2015.
- [2] C. D. Aliprantis and K. C. Border. *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media, 2006.
- [3] E. Bayraktar and A. Cohen. Analysis of a finite state many player game using its master equation. *SIAM Journal on Control and Optimization*, 56(5):3538–3568, 2018.

- [4] C. Benazzoli, L. Campi, and L. 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.
- [5] D. Bertsekas and S. E. Shreve. *Stochastic optimal control: the discrete-time case*, volume 5. Athena Scientific, 1996.
- [6] C. Bertucci, L. Bertucci, J.-M. Lasry, and P.-L. Lions. A mean field game approach to bitcoin mining. *SIAM Journal on Financial Mathematics*, 15(3):960–987, 2024.
- [7] A. Biswas. Mean field games with ergodic cost for discrete time Markov processes. *arXiv preprint arXiv:1510.08968*, 2015.
- [8] P. Brémaud. *Point processes and queues*. Springer, 1981.
- [9] R. Carmona and F. Delarue. *Probabilistic theory of mean field games with applications I-II*. Springer, 2018.
- [10] M. F. Djete. Mean field games of controls: on the convergence of Nash equilibria. *The Annals of Applied Probability*, 33(4):2824–2862, 2023.
- [11] S. N. Ethier and T. G. Kurtz. *Markov processes: characterization and convergence*. John Wiley & Sons, 2009.
- [12] L. C. Florescu and C. Godet-Thobie. *Young measures and compactness in measure spaces*. de Gruyter Berlin, 2012.
- [13] D. A. Gomes, J. Mohr, and R. R. Souza. Discrete time, finite state space mean field games. *Journal de mathématiques pures et appliquées*, 93(3):308–328, 2010.
- [14] D. A. Gomes et al. Continuous time finite state mean field games. *Applied Mathematics & Optimization*, 68(1):99–143, 2013.
- [15] K. Hinderer. *Foundations of Non-stationary Dynamic Programming with Discrete Time Parameter*. Springer-Verlag, 1970.
- [16] M. Huang, R. P. Malhamé, and P. 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.
- [17] B. Jovanovic and R. W. Rosenthal. Anonymous sequential games. *Journal of Mathematical Economics*, 17(1):77–87, 1988.
- [18] H. Kushner. *Weak convergence methods and singularly perturbed stochastic control and filtering problems*. Springer Science & Business Media, 2012.
- [19] H. J. Kushner and P. G. Dupuis. *Numerical methods for stochastic control problems in continuous time*, volume 24. Springer Science & Business Media, 2001.
- [20] D. Lacker. Mean field games via controlled martingale problems: existence of Markovian equilibria. *Stochastic Processes and their Applications*, 125(7):2856–2894, 2015.
- [21] J.-M. Lasry and P.-L. Lions. Mean field games. *Japanese journal of mathematics*, 2(1):229–260, 2007.

- [22] J. Levin. Supermodular games. *Lectures Notes, Department of Economics, Stanford University*, 2003.
- [23] Z. Li, A. M. Reppen, and R. Sircar. A mean field games model for cryptocurrency mining. *Management Science*, 70:2188–2208, 2024.
- [24] J. Moon and T. 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.
- [25] Y.-H. Ni, R. Elliott, and X. Li. Discrete-time mean-field stochastic linear–quadratic optimal control problems, ii: Infinite horizon case. *Automatica*, 57:65–77, 2015.
- [26] M. Nourian and G. 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.
- [27] N. Saldi, T. Basar, and M. Raginsky. Markov–Nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization*, 56(6):4256–4287, 2018.
- [28] A. Taghizadeh, H. Kebriaei, and D. Niyato. Mean field game for equilibrium analysis of mining computational power in blockchains. *IEEE Internet of Things Journal*, 7(8):7625–7635, 2020.