

# Mean Field Games on Weighted and Directed Graphs via Colored Digraphons

Christian Fabian<sup>1</sup>, Kai Cui<sup>1</sup>, *Graduate Student Member, IEEE*, and Heinz Koepll

**Abstract**—Multi-agent systems are in general hard to model and control due to their complex nature involving many individuals. Numerous approaches focus on empirical and algorithmic aspects of approximating outcomes and behavior in multi-agent systems and lack a rigorous theoretical foundation. Graphon mean field games (GMFGs) on the other hand provide a mathematically well-founded and numerically scalable framework for a large number of connected agents. In standard GMFGs, the connections between agents are undirected, unweighted and invariant over time. Our paper introduces colored digraphon mean field games (CDMFGs) which allow for weighted and directed links between agents that are also adaptive over time. Thus, CDMFGs are able to model more complex connections than standard GMFGs. Besides a rigorous theoretical analysis including both existence and convergence guarantees, we employ the online mirror descent algorithm to learn equilibria. To conclude, we illustrate our findings with an epidemics model and a model of the systemic risk in financial markets.

**Index Terms**—Mean field games, agents-based systems, control of networks, machine learning.

## I. INTRODUCTION

IN NUMEROUS research fields such as logistics [1] or robotics [2], scientists are confronted with research questions that involve a large number of individuals or units interacting with each other. Making predictions for these situations is often challenging due to the high quantity of involved agents. Specifically, many of the existing approaches and models either lack a solid theoretical foundation or they can hardly be solved numerically as the number of agents increases.

A concept that aims to resolve both of the just mentioned challenges is the one of mean field games (MFG) which was introduced independently by [3] and [4]. MFGs, in general, model games where a large number of indistinguishable agents

interact with each other such that each single agent has a negligible influence on the other agents. The crucial idea of MFGs is that each individual only interacts with the other agents through the mean field (MF) of agents, i.e., a probability distribution over some state space.

Although MFGs facilitate the mathematical and algorithmic solution of many otherwise intractable models (see [5] for an overview), their initial assumptions appear to be too rigid for various applications. This especially includes the fact that in many real world scenarios, individuals do not interact with the entity of all other agents but with a rather small subgroup. To mathematically model such more complicated setups, one can utilize graphon mean field games (GMFGs) [6], [7].

Intuitively, GMFGs are similar to MFGs but allow agents to be connected by a graph such that only neighbors influence each other. This enables a more realistic analysis of crucial problems such as modelling the dynamics of an epidemic [8]. A shortcoming of this technique is, that the underlying standard graphons [9] are only suitable limiting objects for undirected and unweighted graphs. Therefore, scenarios that contain directed interactions, e.g., sending and receiving messages in a communication network [10], or weighted interactions such as monetary transactions in a financial market [11], cannot be depicted properly by standard GMFGs.

To address these challenges, our paper proposes the concept of colored digraphon mean field games (CDMFGs). We build our modelling approach on the mathematical concept of colored digraphons [12] which are a suitable limiting object for sequences of weighted and directed graphs. Furthermore, our examples also outline how CDMFGs are able to depict adaptive connections between agents, i.e., when neighbors of an agent vary over time. Our contributions can be summarized as follows: i) We introduce the novel concept of CDMFGs which provides a modelling framework for MFGs on weighted, directed, and adaptive graphs; ii) we conduct a rigorous theoretical analysis of CDMFGs and give both existence and convergence guarantees; iii) we employ the online mirror descent algorithm [13] to determine equilibria in empirical applications; and iv) we demonstrate the empirical capabilities of our approach with examples on the systemic risk in financial markets and the spread of contagious diseases.

## II. COLORED DIGRAPHONS

This section provides a short introduction to  $k$ -colored digraphons. For more details, see e.g., [12].

Manuscript received 8 September 2022; revised 4 November 2022; accepted 22 November 2022. Date of publication 7 December 2022; date of current version 20 December 2022. This work was supported in part by the Hessian Ministry of Science and the Arts (HMMWK) within the Projects “The Third Wave of Artificial Intelligence—3AI” and hessian.AI; in part by the LOEWE initiative (Hesse, Germany) within the emergenCITY center; and in part by the German Research Foundation (DFG) via the Collaborative Research Center (CRC) 1053-MAKI. Recommended by Senior Editor L. Zhang. (*Corresponding author: Christian Fabian.*)

The authors are with the Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, 64287 Darmstadt, Germany (e-mail: christian.fabian@tu-darmstadt.de; kai.cui@tu-darmstadt.de; heinz.koepll@tu-darmstadt.de).

Digital Object Identifier 10.1109/LCSYS.2022.3227453

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see <https://creativecommons.org/licenses/by/4.0/>

The general idea of graphons [9] is to provide a limiting object for growing sequences of graphs. Then, one can avoid large adjacency matrices and instead work with the more amenable graphons which contain the asymptotic structure of the graph sequence. A standard graphon  $W$  is formally defined as a symmetric, measurable function  $W : [0, 1]^2 \rightarrow [0, 1]$  which yields the limit for sequences of undirected and unweighted graphs. Some papers such as [7] and [14] interpret the probability  $W(x, y)$  of an edge between nodes  $x$  and  $y$  as the ‘‘weight’’ of this edge. However, each realization of this random graph yields an unweighted graph. The colored digraphon on the other hand allows us to specify edge weights and separately provide the probabilities with which they occur such that we can sample weighted and directed graphs. Thus, standard graphons provide probability weights but are in general not the limiting objects of weighted and directed graph sequences. Consequently, it becomes necessary to provide a more general concept, namely  $k$ -colored digraphons [12].

In order to depict directed graphs, often referred to as digraphs, we drop the symmetry condition which yields measurable functions  $W : [0, 1]^2 \rightarrow [0, 1]$  called digraphons. To describe  $k$ -colored digraphs, i.e., digraphs where each edge is assigned one of  $k$  different colors, we extend the previous definition to a tuple of  $k$  digraphons  $\mathbf{W} = (W^1, \dots, W^k)$  such that  $\sum_{i=1}^k W^i = 1$  and where  $\mathbf{W}$  is called a  $k$ -digraphon. Note that as usual, the  $k$  edge colors can be equivalently thought of as edge weights.

We can identify any  $k$ -colored digraph  $G$  with  $N$  nodes with a  $k$ -digraphon as follows. Let  $I_1, \dots, I_N$  be intervals of equal length that form a partition of  $[0, 1]$ . Then, the graph  $G$  can be associated with a  $k$ -digraphon  $\mathbf{W}_G = (W_G^1, \dots, W_G^k)$  defined by  $W_G^h(x, y) = 1$  if  $(x, y) \in I_i \times I_j$ ,  $h \in \{1, \dots, k\}$  and  $h$  is the color of edge  $ij$ , and  $W_G^h(x, y) = 0$  otherwise.

On the other hand, we can sample finite, random  $k$ -colored digraphs from a given  $k$ -digraphon  $\mathbf{W}$ . Assuming that our sampled graph is supposed to have  $N$  nodes, we draw points  $x_1, \dots, x_N$  i.i.d. uniformly from  $[0, 1]$  at random and interpret each value as a node. Subsequently, we color the edge between vertices  $1 \leq i, j \leq N$  with color  $h$  with probability  $W^h(x_i, x_j)$ , independently of all other edges. One can show that for every convergent sequence of  $k$ -colored digraphs there exists a limiting  $k$ -digraphon [12, Proposition 2.1].

To compare  $k$ -digraphons, we first recall the cut norm for standard graphons  $W : [0, 1]^2 \rightarrow [0, 1]$ , i.e.,

$$\|W\|_{\square} := \sup_{S, T \subseteq [0, 1]} \left| \int_{S \times T} W(x, y) dx dy \right|, \quad (1)$$

where  $S$  and  $T$  range over the measurable subsets of  $[0, 1]$ . The cut norm can be extended to  $k$ -digraphons, i.e., for a  $k$ -digraphon  $\mathbf{W} = (W^1, \dots, W^k)$  we define

$$\|\mathbf{W}\|_{\square} := \sum_{h=1}^k \|W^h\|_{\square}. \quad (2)$$

For the proofs in this letter, a different formulation of convergence is useful. For a standard graphon  $W : [0, 1]^2 \rightarrow [0, 1]$  and a converging sequence of standard graphons  $(W_N)_N$  with  $\|W - W_N\|_{\square} \rightarrow 0$  we recall that this is equivalent to convergence in the operator norm  $\|\cdot\|_{L_{\infty} \rightarrow L_1}$ , defined by

$$\begin{aligned} & \|W - W_N\|_{L_{\infty} \rightarrow L_1} \\ &= \sup_{\|f\|_{\infty} \leq 1} \int_{\mathcal{I}} \left| \int_{\mathcal{I}} (W(x, y) - W_N(x, y)) f(y) dy \right| dx \rightarrow 0, \quad (3) \end{aligned}$$

see [9, Lemma 8.11] for details.

*Assumption 1:* The sequence of  $k$ -digraphons  $(\mathbf{W}_{G_N})_{N \in \mathbb{N}}$  generated by the graph sequence  $(G_N)_{N \in \mathbb{N}}$  converges in cut norm  $\|\cdot\|_{\square}$  as  $N \rightarrow \infty$  to some  $k$ -digraphon  $\mathbf{W}$ , i.e.,

$$\sum_{h=1}^k \left\| W_{G_N}^h - W^h \right\|_{\square} \rightarrow 0 \quad (4)$$

which implies  $\sum_{h=1}^k \|W_{G_N}^h - W^h\|_{L_{\infty} \rightarrow L_1} \rightarrow 0$ .

### III. FINITE AGENT MODEL

In the following, if  $Y$  is some finite set,  $\mathcal{P}(Y)$  denotes the set of all probability measures on  $Y$  and  $\mathcal{B}(Y)$  the set of all bounded measures. The finite agent model (or  $N$  agent model) consists of  $N$  agents who can choose actions from some finite action set  $\mathcal{U}$  while they are in one of the states defined by the finite state space  $\mathcal{X}$ . Furthermore, the game is setup on some  $k$ -colored digraph  $G_N = (V_N, E_N)$  with  $N$  nodes where the vertices represent agents while the (weighted and directed) edges indicate connections between pairs of agents. The game starts at time  $t = 0$  and ends at the terminal time  $T \in \mathbb{N}$  where agents are allowed to act at each of the time points  $t \in \mathcal{T} := \{0, \dots, T - 1\}$ . Formally, agent  $i$  can choose to implement a policy  $\pi^i \in \Pi := \mathcal{P}(\mathcal{U})^{\mathcal{X} \times \mathcal{T}}$  where  $\pi_t^i$  denotes the policy of agent  $i$  at time  $t$ .

In contrast to models on simple graphs, a finite agent model on a  $k$ -colored digraph includes  $2k$  unnormalized neighborhood state distributions for each agent  $i$ , i.e.,

$$\mathbb{G}_{t,h,\text{out}}^i := \frac{1}{N} \sum_{j \in V_N} \mathbf{1}_{\{ij \text{ has color } h\}} \delta_{X_j^i} \quad (5)$$

$$\mathbb{G}_{t,h,\text{in}}^i := \frac{1}{N} \sum_{j \in V_N} \mathbf{1}_{\{ji \text{ has color } h\}} \delta_{X_j^i} \quad (6)$$

where  $h$  can be any of the  $k$  colors of the graph and  $\delta$  denotes the Dirac measure. Here, ‘unnormalized’ means that the distribution weights do not sum up to one in general. For notational convenience, we define  $\mathbb{G}_t^i := (\mathbb{G}_{t,1,\text{out}}^i, \dots, \mathbb{G}_{t,k,\text{out}}^i, \mathbb{G}_{t,1,\text{in}}^i, \dots, \mathbb{G}_{t,k,\text{in}}^i)$  such that  $\mathbb{G}_t^i \in \mathcal{B}(\mathcal{X})^{2k}$ . Agents follow policies that only depend on their current state, i.e., local Markovian feedback policies. This is taken into account by defining the model dynamics

$$U_t^i \sim \pi_t^i(\cdot | X_t^i) \quad \text{and} \quad X_{t+1}^i \sim P(\cdot | X_t^i, U_t^i, \mathbb{G}_t^i) \quad (7)$$

for all  $t \leq T$ ,  $i \in V_N$  and  $X_0^i \sim \mu_0$  where  $\mu_0$  is the initial state distribution and  $P : \mathcal{X} \times \mathcal{U} \times \mathcal{B}(\mathcal{X})^{2k} \rightarrow \mathcal{P}(\mathcal{X})$  is some transition kernel. Finally, each agent’s rewards are given by the function  $r : \mathcal{X} \times \mathcal{U} \times \mathcal{B}(\mathcal{X})^{2k} \rightarrow \mathbb{R}$ . Agents aim to competitively maximize their total expected rewards given by  $J_N^i(\pi^1, \dots, \pi^N) := \mathbb{E}[\sum_{t \in \mathcal{T}} r(X_t^i, U_t^i, \mathbb{G}_t^i)]$ .

To analyze the above model, we recall the  $(\epsilon, p)$ -Markov-Nash equilibrium ( $(\epsilon, p)$ -MNE) concept, see e.g., [15], [16].

*Definition 1:* For  $\epsilon, p > 0$ , a  $(\epsilon, p)$ -MNE is a tuple  $\pi = (\pi^1, \dots, \pi^N) \in \Pi^N$  such that  $\forall i \in V_N \subseteq V_N$

$$\sup_{\pi \in \Pi} J_N^i(\pi^1, \dots, \pi^{i-1}, \pi, \pi^{i+1}, \dots, \pi^N) - J_N^i(\pi) \leq \epsilon$$

for some subset  $V'_N$  of nodes with  $|V'_N| \geq \lfloor (1-p)N \rfloor$ .

The advantage of the  $(\epsilon, p)$ -MNE is that it allows relatively small groups of agents to be out-of-equilibrium. Even in the asymptotic case, the network can contain local structures that deviate from the colored digraphon. However, taking  $p \rightarrow 0$ , brings us arbitrarily close to an equilibrium for all agents.

#### IV. COLORED DIGRAPHON MEAN FIELD GAME

In CDMFGs, an infinite number of agents, i.e., the mean field, try to competitively maximize their expected rewards. In contrast to standard GMFGs, CDMFGs allow for directed and weighted connections between agents. This also enables the modelling of adaptive network structures where connections change over time. Specifically, the edge colors can carry the information of when an edge is active, e.g., color 1 means active in the first half of  $\mathcal{T}$ , color 2 indicates an edge being active in the second half of  $\mathcal{T}$ . CDMFGs are closely related to the finite agent models defined above and provide an increasingly accurate approximation for the  $N$  agent case as  $N$  increases, for details see the next sections.

The MF of agents is represented by the interval  $\mathcal{I} := [0, 1]$  where each  $\alpha \in \mathcal{I}$  represents one agent. Then, we denote the space of measurable MF ensembles by  $\mathcal{M} := \mathcal{P}(\mathcal{X})^{\mathcal{I} \times \mathcal{T}}$  such that for any  $(t, x, \mu) \in \mathcal{T} \times \mathcal{X} \times \mathcal{M}$  the map  $\alpha \mapsto \mu_t^\alpha(x)$  is measurable. Similarly,  $\mathbf{\Pi} \subseteq \Pi^{\mathcal{I}}$  is the space of measurable policy ensembles where measurable refers to  $\alpha \mapsto \pi_t^\alpha(u | x)$  being measurable for all  $(t, x, u, \pi) \in \mathcal{T} \times \mathcal{X} \times \mathcal{U} \times \mathbf{\Pi}$ .

*Assumption 2:*  $\pi \in \mathbf{\Pi}$  is Lipschitz continuous up to a finite number of discontinuities.

As in the  $N$  agent case, a CDMFG on a  $k$ -colored digraphon has  $2k$  unnormalized neighborhood state distributions

$$\mathbb{G}_{t,h,\text{out}}^{\alpha,\mu,W} := \int_{\mathcal{I}} W^h(\alpha, \beta) \mu_t^\beta d\beta \quad (8)$$

$$\mathbb{G}_{t,h,\text{in}}^{\alpha,\mu,W} := \int_{\mathcal{I}} W^h(\beta, \alpha) \mu_t^\beta d\beta \quad (9)$$

for each agent  $\alpha$  where  $h$  can be any of the  $k$  colors. As before, we let  $\mathbb{G}_t^{\alpha,\mu,W}$  denote the vector of the  $2k$  neighborhood state distributions. We usually drop the dependence on  $\mu$  and  $W$  if it is clear from the context. Thus, the model dynamics are defined by  $U_t^\alpha \sim \pi_t^\alpha(\cdot | X_t^\alpha)$  and  $X_{t+1}^\alpha \sim P(\cdot | X_t^\alpha, U_t^\alpha, \mathbb{G}_t^\alpha)$  with  $X_0^\alpha \sim \mu_0$ . Agents try to competitively maximize their rewards  $J_\alpha^\mu(\pi^\alpha) := \mathbb{E}[\sum_{t \in \mathcal{T}} r(X_t^\alpha, U_t^\alpha, \mathbb{G}_t^\alpha)]$ .

To state an equilibrium concept for CDMFGs, we require two functions  $\Psi : \mathbf{\Pi} \rightarrow \mathcal{M}$  and  $\Phi : \mathcal{M} \rightarrow 2^{\mathbf{\Pi}}$ .  $\Phi$  maps  $\mu \in \mathcal{M}$  to the set of policies with  $\pi^\alpha = \arg \max_{\pi \in \mathbf{\Pi}} J_\alpha^\mu(\pi^\alpha)$  for all  $\alpha \in [0, 1]$ . Intuitively, this corresponds to the optimal response policies given the respective mean field  $\mu$ .  $\Psi$  assigns  $\pi \in \mathbf{\Pi}$  to the  $\mu = \Psi(\pi) \in \mathcal{M}$  which is induced for all  $\alpha \in [0, 1]$  through the recursive equation

$$\mu_{t+1}^\alpha(x) = \sum_{x' \in \mathcal{X}} \mu_t^\alpha(x') \sum_{u \in \mathcal{U}} \pi_t^\alpha(u | x') P(x | x', u, \mathbb{G}_t^\alpha) \quad (10)$$

with  $\mu_0^\alpha = \mu_0$ . Next, we state the concept of the colored digraphon mean field equilibrium (CDMFE) which partially transfers existing work on MFGs [17] to colored digraphons.

*Definition 2:* A pair  $(\mu, \pi) \in \mathcal{M} \times \mathbf{\Pi}$  is a CDMFE if and only if both  $\pi \in \Phi(\mu)$  and  $\mu = \Psi(\pi)$  hold.

To ensure the existence of a CDMFE, we make a Lipschitz assumption commonly used in the literature, see Sections VI and VII for examples fulfilling the Lipschitz assumption.

*Assumption 3:*  $r, P, W^1, \dots, W^k$  are Lipschitz continuous with Lipschitz constants  $L_r, L_P, L_{W^1}, \dots, L_{W^k} > 0$ .

*Lemma 1:* Under Assumption 3 there exists a CDMFE.

*Proof:* Identical to [18, Proof of Theorem 1].  $\blacksquare$

*Remark 1:* Lemma 1 includes the case where the set of agents  $[0, 1]$  consists of  $2 \leq \kappa < \infty$  disjoint subsets  $A_1, \dots, A_\kappa$  and each subset  $A_j$  and the corresponding agents have their own state space  $\mathcal{X}_j$ . To see this, define the state space  $\mathcal{X}' = \bigcup_j \mathcal{X}_j$  such that Lemma 1 can be applied to  $\mathcal{X}'$ .

#### V. APPROXIMATION VIA CDMFGS

A crucial benefit of CDMFGs is that they yield an accurate approximation of finite agent models with sufficiently many agents. To relate the two model classes, we discretize a policy ensemble  $\pi \in \mathbf{\Pi}$  using the function  $\Gamma_N : \mathbf{\Pi} \rightarrow \Pi^N$  defined by  $\Gamma_N(\pi) = (\pi_1^{\alpha_1}, \dots, \pi_N^{\alpha_N})$  with  $\alpha_i = i/N$ . Conversely, an  $N$ -agent policy  $(\pi_1, \dots, \pi_N) \in \Pi^N$  can be transferred to the CDMFG setup by defining  $\mu^N \in \mathcal{M}$  and  $\pi^N \in \mathbf{\Pi}$  through  $\mu_t^{N,\alpha} := \sum_{i \in V_N} \mathbf{1}_{\{\alpha \in (\frac{i-1}{N}, \frac{i}{N}]\}} \cdot \delta_{X_t^i}$  and  $\pi_t^{N,\alpha} := \sum_{i \in V_N} \mathbf{1}_{\{\alpha \in (\frac{i-1}{N}, \frac{i}{N}]\}} \cdot \pi_t^i$  for all  $(\alpha, t) \in \mathcal{T} \times \mathcal{I}$ .

To abbreviate the following formal expressions, we define  $\mu_t(f) := \int_{\mathcal{T}} \sum_{x \in \mathcal{X}} f(x, \alpha) \mu_t^\alpha(x) d\alpha$  for all functions  $f : \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$  which averages  $f$  over the mean fields of all indices  $\alpha$ . Furthermore, we introduce the policy ensemble  $\Gamma_N(\pi, i, \hat{\pi}) \in \Pi^N$  for all  $i \leq N$  and  $\hat{\pi} \in \mathbf{\Pi}$  which is equal to  $\Gamma_N(\pi)$  except that  $\pi_i^{\alpha_i}$  is substituted by  $\hat{\pi}$ .

*Theorem 1:* Assume that Assumptions 1 and 2 hold with  $\pi \in \mathbf{\Pi}$  and  $\mu = \Psi(\pi)$ . Then, we have for the  $N$ -agent policy  $\Gamma_N(\pi, i, \hat{\pi}) \in \Pi^N$ , for all measurable functions  $f : \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$  uniformly bounded by some  $M_f$  that for all  $t \in \mathcal{T}$

$$\mathbb{E}[|\mu_t^N(f) - \mu_t(f)|] \rightarrow 0 \quad (11)$$

uniformly over all  $\hat{\pi} \in \mathbf{\Pi}, i \in V_N$ .

*Proof:* The transition operator  $P_{t,\mu',W}^\pi$  is given by

$$(\mu P_{t,\mu',W}^\pi)^\alpha := \sum_{(x,u) \in \mathcal{X} \times \mathcal{U}} \mu^\alpha(x) \pi_t^\alpha(u | x) P(\cdot | x, u, \mathbb{G}_t^{\alpha,\mu',W})$$

and formalizes the transition of MF  $\mu$  from time point  $t$  to  $t+1$  under the  $k$ -digraphon  $W$ , the MF  $\mu'$  and the policy ensemble  $\pi \in \mathbf{\Pi}$ . We define  $(W\mu)(\alpha, x) := \int_{\mathcal{T}} W(\alpha, \beta) \mu^\beta(x) d\beta$  and  $(\mu W)(\alpha, x) := \int_{\mathcal{T}} W(\beta, \alpha) \mu^\beta(x) d\beta$  with  $(W\mu - W'\mu')(\alpha, x) := (W\mu)(\alpha, x) - (W'\mu')(\alpha, x)$  and  $(\mu W - \mu'W')$  analogously.

We prove Theorem 1 via induction similar to [18, Th. 2]. The induction start follows by a law of large numbers (LLN) argument. For the induction step we have

$$\begin{aligned} & \mathbb{E}[|\mu_{t+1}^N(f) - \mu_{t+1}(f)|] \\ & \leq \mathbb{E}\left[|\mu_{t+1}^N(f) - \mu_t^N P_{t,\mu_t^N,W_N}^{\pi^N}(f)|\right] \\ & \quad + \mathbb{E}\left[|\mu_t^N P_{t,\mu_t^N,W_N}^{\pi^N}(f) - \mu_t^N P_{t,\mu_t^N,W}^{\pi^N}(f)|\right] \\ & \quad + \mathbb{E}\left[|\mu_t^N P_{t,\mu_t^N,W}^{\pi^N}(f) - \mu_t^N P_{t,\mu_t^N,W}^\pi(f)|\right] \\ & \quad + \mathbb{E}\left[|\mu_t^N P_{t,\mu_t^N,W}^\pi(f) - \mu_t^N P_{t,\mu_t,W}^\pi(f)|\right] \\ & \quad + \mathbb{E}\left[|\mu_t^N P_{t,\mu_t,W}^\pi(f) - \mu_{t+1}(f)|\right] \end{aligned}$$

where, the first, third, and fifth term are bounded as in [18] and uniformly converge to zero as  $N \rightarrow \infty$ . Specifically,  $\mathbb{E} [|\mu_{t+1}^N(f) - \mu_{t,\mu_t^N,W_N}^N(f)|] \rightarrow 0$  follows from a LLN argument,  $\mathbb{E} [|\mu_t^N P_{t,\mu_t^N,W}^{\pi}(f) - \mu_t^N P_{t,\mu_t^N,W}^{\pi}(f)|]$  converges to zero by Assumption 2, and by the induction assumption we have  $\mathbb{E} [|\mu_t^N P_{t,\mu_t^N,W}^{\pi}(f) - \mu_{t+1}^N(f)|] \rightarrow 0$ .

For the second and fourth term, however, we require an additional mathematical effort. We start with the second summand. Leveraging Assumption 1 and keeping in mind that  $0 \leq \mu_t^{N,\beta}(x) \leq 1$  by definition, we obtain

$$\begin{aligned} & \mathbb{E} \left[ \left| \mu_t^N P_{t,\mu_t^N,W_N}^{\pi}(f) - \mu_t^N P_{t,\mu_t^N,W}^{\pi}(f) \right| \right] \\ & \leq \xi |\mathcal{X}| \sup_{x \in \mathcal{X}} \mathbb{E} \left[ \int_{\mathcal{I}} \max_{h \in \{1, \dots, k\}} \left| (W_N^h \mu_t^N - W^h \mu_t^N)(\alpha, x) \right| \right. \\ & \quad \left. + \max_{h \in \{1, \dots, k\}} \left| (\mu_t^N W_N^h - \mu_t^{N,\beta} W^h)(\alpha, x) \right| d\alpha \right] \rightarrow 0 \end{aligned}$$

where  $\xi := k|\mathcal{X}|M_f L_P$ . We make a useful observation by defining the uniformly bounded functions  $f'_{x',\alpha}(x, \beta) = W^h(\alpha, \beta) \cdot \mathbf{1}_{x=x'}$  for any  $(x', \alpha, h) \in \mathcal{X} \times \mathcal{I} \times \{1, \dots, k\}$ . Applying the induction assumption to them yields

$$\begin{aligned} & \mathbb{E} \left[ \int_{\mathcal{I}} \left| (W^h \mu_t^N - W^h \mu_t)(\alpha, x') \right| d\alpha \right] \\ & = \int_{\mathcal{I}} \mathbb{E} \left[ \left| \mu_t^N(f'_{x',\alpha}) - \mu_t(f'_{x',\alpha}) \right| \right] d\alpha \rightarrow 0. \end{aligned}$$

Note that an analogous argument yields  $\mathbb{E} [ \int_{\mathcal{I}} |(\mu_t^N W^h - \mu_t W^h)(\alpha, x')| d\alpha ] \rightarrow 0$ . Then, we can prove convergence of the fourth summand, i.e.,

$$\begin{aligned} & \mathbb{E} \left[ \left| \mu_t^N P_{t,\mu_t^N,W}^{\pi}(f) - \mu_t^N P_{t,\mu_t^N,W}^{\pi}(f) \right| \right] \\ & \leq \xi |\mathcal{X}| \sup_{x \in \mathcal{X}} \mathbb{E} \left[ \int_{\mathcal{I}} \max_{h \in \{1, \dots, k\}} \left| (W^h \mu_t^N - W^h \mu_t)(\alpha, x) \right| \right. \\ & \quad \left. + \max_{h \in \{1, \dots, k\}} \left| (\mu_t^N W^h)(\alpha, x) (\mu_t W^h(\alpha, x)) \right| d\alpha \right] \rightarrow 0. \end{aligned}$$

This concludes the induction and the proof.  $\blacksquare$

Theorem 1 can be leveraged to show that an equilibrium in a CDMFG yields an approximate equilibrium in the corresponding finite agent model. Moreover, the estimation's accuracy increases with the number of agents in the finite model which is formalized by the next theorem.

*Theorem 2:* Let  $(\pi, \mu)$  be a CDMFE and assume that Assumptions 1, 2, and 3 hold. Then, for any  $\varepsilon, p > 0$  there exists  $N'$  such that for all  $N > N'$ , the policy  $\Gamma_N(\pi) \in \Pi^N$  is an  $(\varepsilon, p)$ -Markov Nash equilibrium.

*Proof:* Follows from Corollary 1 in the Appendix.  $\blacksquare$

## VI. EXAMPLES

To verify our model numerically, we give a number of example tasks fulfilling Assumption 3. We consider  $\mathbb{G}_{t,\text{out}}^i := \sum_{h=1}^k c_{h,\text{out}} \cdot \mathbb{G}_{t,h,\text{out}}^i$  and  $\mathbb{G}_{t,\text{in}}^i := \sum_{h=1}^k c_{h,\text{in}} \cdot \mathbb{G}_{t,h,\text{in}}^i$  for our examples where the coefficients  $c_{h,\text{out}}$  and  $c_{h,\text{in}}$  can be understood as the strength of the respective edge color.

a) *Epidemics:* In the SIS model [19] agents may protect themselves ( $D$ ) from costly infection ( $I$ ) to stay susceptible ( $S$ ). We endow the model with additional directionality of virus spread along (colorless) directed edges. Therefore,

formally  $\mathcal{X} = \{S, I\}$ ,  $\mathcal{U} = \{\bar{D}, D\}$ ,  $\mu_0(I) = 0.5$ ,  $T = 50$ ,  $r(X_t^i, U_t^i, \mathbb{G}_t^i) = -2 \cdot \mathbf{1}_{\{X_t^i=I\}} - 0.5 \cdot \mathbf{1}_{\{U_t^i=D\}}$  and  $P(S | I, U_t^i, \mathbb{G}_t^i) = 0.2$ ,  $P(I | S, U_t^i, \mathbb{G}_t^i) = 0.8 \mathbb{G}_{t,\text{in}}^i(I) \cdot \mathbf{1}_{\{U_t^i=\bar{D}\}}$ .

b) *Beach:* In the adjusted Beach example (e.g., [20]), many people on the beach adjust their position to be both close to a bar and away from (incoming) disliked neighbors. There are 10 locations  $\mathcal{X} = \{0, 1, \dots, 9\}$  with bar location  $B = 5$ , movement actions  $\mathcal{U} = \{-1, 0, 1\}$ ,  $\mu_0 = \text{Unif}(\mathcal{X})$ ,  $T = 30$ ,  $r(X_t^i, U_t^i, \mathbb{G}_t^i) = \frac{1}{5}|B - X_t^i| + \frac{1}{5}|U_t^i| - \mathbb{G}_{t,\text{in}}^i(x)$  and  $X_{t+1}^i = X_t^i + U_t^i + \epsilon_t^i$ , where the random noise  $\epsilon_t^i$  is equal to  $-1$  or  $1$  with probabilities 0.05, and otherwise 0.

c) *Systemic Risk:* In our interbanking loan market model, the weights are  $c_{h,\text{out}} = c_{h,\text{in}} = h - 1$  such that the edge color  $h$  depicts the financial volume of a loan. An agent, i.e., a bank, can be in one of three states  $\mathcal{X} := \{m, \ell, b\}$  meaning that its capital reserves are medium ( $m$ ), low ( $\ell$ ), or it is bankrupt ( $b$ ). The bank's state  $x \in \mathcal{X}$  as well as its loan contracts with other banks characterized by the neighborhood distributions determine its liquidity

$$c_x := \xi_x + \left( \sum_{x' \in \mathcal{X}} \mathbb{G}_{t,\text{in}}^i(x') \cdot w_{x'} - \mathbb{G}_{t,\text{out}}^i(x') \cdot (3 - w_{x'}) \right)$$

where each state  $x$  is assigned an initial level of capital endowment  $\xi_x \in \mathbb{R}$ . For notational convenience we define  $\sigma_{c,x} := \max\{-1, \min\{1, c_x\}\}$  for all states  $x \in \mathcal{X}$ . Banks can choose to keep ( $k$ ) their capital at its current level, raise ( $r$ ) or decrease ( $d$ ) it which is formalized by  $\mathcal{U} := \{k, r, d\}$ . We model negative external shocks, e.g., wars, by including a parameter  $1 > \beta > 0$  in the transition probabilities. Intuitively, a bank is more likely to be in a high liquidity state in  $t + 1$  if it already is there in  $t$  or if it decides to raise its liquidity.


Thus, for  $(i, j) \in \{m, \ell, b\}^2$  define a  $3 \times 3$  transition matrix  $M(\alpha, t, k)$  for choosing action  $u = k$  by  $M_{i,m}(\alpha, t, k) := (1 + \sigma_{c,i})/4$  and  $M_{i,\ell}(\alpha, t, k) := (1 - 0.25\sigma_{c,i} + \beta \cdot \mathbf{1}_{\{i=m\}})/4$  and  $M_{i,b}(\alpha, t, k) := (2 - 0.75\sigma_{c,i} - \beta \cdot \mathbf{1}_{\{i=m\}})/4$ .

If a bank chooses to raise its capital, the respective transition matrix  $M(\alpha, t, r)$  is defined by  $M_{m,j}(\alpha, t, r) := M_{m,j}(\alpha, t, k)$  and  $M_{\ell,j}(\alpha, t, r) := (1 - \lambda)M_{\ell,j}(\alpha, t, k) + \lambda(M_{m,j}(\alpha, t, k) - \beta \mathbf{1}_{\{j=\ell\}} + \beta \mathbf{1}_{\{j=b\}})$  and  $M_{b,j}(\alpha, t, r) := (1 - \lambda)M_{b,j}(\alpha, t, k) + \lambda M_{\ell,j}(\alpha, t, k)$ .


Finally, if a bank implements action  $d$ , the transition matrix  $M(\alpha, t, d)$  is defined by  $M_{m,j}(\alpha, t, d) := (1 - \lambda)M_{m,j}(\alpha, t, k) + \lambda(M_{\ell,j}(\alpha, t, d) + \beta \mathbf{1}_{\{j=\ell\}} - \beta \mathbf{1}_{\{j=b\}})$  and  $M_{\ell,j}(\alpha, t, d) := (1 - \lambda)M_{\ell,j}(\alpha, t, k) + \lambda M_{b,j}(\alpha, t, k)$  and  $M_{b,j}(\alpha, t, d) := M_{b,j}(\alpha, t, k)$ .

Banks maximize  $J(u) := \mathbb{E}[-\sum_{t=0}^T k_m \mathbf{1}_m + k_b \mathbf{1}_b]$  where  $k_b \gg k_m$  are the costs for being in the respective state. In our experiments, we use  $\beta = 0.6$ ,  $\mu_0(m) = \mu_0(I) = 0.5$ ,  $T = 50$ ,  $\lambda = 0.1$ ,  $k_b = 10k_m = 0.4$ ,  $(\xi_m, \xi_\ell, \xi_b) = (1, 0, -1)$  and  $(w_m, w_\ell, w_b) = (1, 0, -1)$ .


d) *Adaptive Versions:* We also introduce adaptive versions of SIS and Systemic Risk by changing the underlying graph structure (for simplicity at  $t^* = T/2$ ) via time-dependent coefficients  $c_{h,\text{out}} = c_{h,\text{in}} = \mathbf{1}_{\{t \in ((h-2)T/2, (h-1)T/2)\}}$  for  $h > 1$  in the definitions of  $\mathbb{G}_{t,\text{out}}^i$  and  $\mathbb{G}_{t,\text{in}}^i$ . This means that edges with color  $h = 2$  are active within the time interval  $(0, T/2]$  while edges with color  $h = 3$  are active during  $(T/2, T]$ . For adaptive SIS, we adjust  $P(I | S, U_t^i, \mathbb{G}_t^i) = \min(1, 1.6 \mathbb{G}_{t,\text{in}}^i(I) \cdot \mathbf{1}_{\{U_t^i=\bar{D}\}})$ .




**Fig. 1.** The (approximate) exploitability of policies  $\Delta J(\pi) := \int_{\mathcal{I}} \sup_{\pi^* \in \Pi} J_{\alpha}^{\Psi(\pi)}(\pi^*) - J_{\alpha}^{\Psi(\pi)}(\pi) d\alpha$  at each iteration  $n$  of the online mirror descent algorithm is decreasing. (a): The rotated uniform attachment digraphon scenario; (b): The ranked attachment version.



**Fig. 2.** The limiting mean-field (state marginals) for various graphon indices  $\alpha \in [0, 1]$  in the combined uniform ranked attachment scenario. (a): SIS; (b): Beach; (c-d): Systemic Risk.



**Fig. 3.** The estimated deviation between empirical and limiting mean field  $\Delta \mu := \mathbb{E}[\sum_{t \in \mathcal{T}, x \in \mathcal{X}} |\frac{1}{N} \sum_i \delta_{X_t^i}(x) - \int_{\mathcal{I}} \mu_t^{\alpha}(x) d\alpha|]$  and its 95% confidence interval over the number of nodes  $N$ , for each averaged over 100 random samples. (a): The rotated uniform attachment digraphon scenario; (b): The ranked attachment version.



**Fig. 4.** The limiting mean-field (state marginals) for various graphon indices  $\alpha \in [0, 1]$  in the adaptive graph case. (a): The SIS example in the combined uniform ranked attachment digraphon scenario; (b): The Systemic Risk example in the double rotated uniform attachment digraphon scenario.

## VII. EXPERIMENTS

To verify the usefulness of our proposed models, we use the online mirror descent algorithm [13] and discretization of the graphon index space  $[0, 1]$  [18] to numerically find CDMFE for the aforementioned examples and the following

digraphons. We consider the rotated uniform attachment digraphon scenario resulting from the uniform attachment process (see e.g., [9]) with added directionality

$$W^1 = 1 - W^2, W^2(x, y) = W_{r\text{-unif}}(x, y) := 1 - x(1 - y),$$

the double rotated uniform attachment digraphon scenario

$$W^1 = 1 - W^2 - W^3, W^2 = \frac{1}{2} W_{r\text{-unif}}, W^3 = \frac{1}{2} W_{l\text{-unif}}$$

where  $W_{l\text{-unif}}(x, y) := 1 - (1 - x)y$ , and the combined uniform ranked attachment digraphon scenario

$$W^1 = 1 - W^2 - W^3, W^2 = \frac{1}{2} W_{r\text{-unif}}, W^3 = \frac{1}{2} W_{r\text{-rank}}$$

where  $W_{r\text{-rank}} := 1 - \max(x, (1 - y))$ , with weights  $c_{h,\text{out}} = c_{h,\text{in}} = h - 1$  unless stated otherwise.

As can be seen in Fig. 1, the numerical approach converges to zero in approximate exploitability [18], indicating good convergence to a CDMFE. The resulting behavior is intuitive as shown in Fig. 2: For the SIS problem, agents with less incoming edges (low  $\alpha$ ) are less likely to protect themselves due to the lower perceived risk, leading to infections; In the Beach scenario, agents with more incoming connections will spread out further away from the bar to avoid other agents, while agents with little to no incoming connections will stay near the bar; Finally, in the systemic risk scenario, banks borrowing more money from other banks are more likely to keep low capital reserves. In Fig. 3, we can see that the difference between the empirical mean-field and the limiting mean-field converges to zero as the finite graph system becomes sufficiently big, supporting Theorem 1. Lastly, in Fig. 4, for the adaptive scenario we can see analogous behavior to the one seen in Fig. 2 with qualitative reversals in behavior at  $t = T/2$  when the set of active edges shifts.

## VIII. CONCLUSION

In this letter we have constructed and analyzed CDMFGs by providing both rigorous mathematical results as well as illustrating our theory on different examples. CDMFGs extend the existing GMFG theory to weighted and directed graphs and thereby enable the scalable analysis of many problems that are outside the scope of standard GMFGs. Furthermore, CDMFGs yield a framework to depict adaptive graph structures, as we have demonstrated in the experiments. For future work, possible extensions of CDMFGs are manifold and could include partial observability, continuous time, state, and action spaces as well as using CDMFGs in applications. We hope that our work inspires future MFG research and facilitates the analysis of many real-world problems.

## APPENDIX

Define  $\forall t \in \mathcal{T}$  the dynamics  $\hat{U}_{N_t}^i \sim \hat{\pi}_t(\cdot | \hat{X}_{N_t}^i)$ ,  $\hat{X}_{N_{t+1}}^i \sim P(\cdot | \hat{X}_{N_t}^i, \hat{U}_{N_t}^i, \mathbb{G}_{N_t}^i)$ , and  $\hat{X}_{N_0}^i \sim \mu_0$  for all agents  $i$ .

**Lemma 2:** Assume that Assumptions 1, 2, and 3 hold with  $\pi \in \Pi$  and  $\mu = \Psi(\pi)$ . For  $\Gamma_N(\pi)$ ,  $\Gamma(\pi, i, \hat{\pi}) \in \Pi^N$  with arbitrary  $\hat{\pi} \in \Pi$ , and for any uniformly bounded family  $\mathcal{G}$  of functions  $g : \mathcal{X} \rightarrow \mathbb{R}$  and any  $\varepsilon, p > 0$ ,  $t \in \mathcal{T}$ , there exists  $N' \in \mathbb{N}$  such that for all  $N > N'$  we have

$$\sup_{g \in \mathcal{G}} \left| \mathbb{E}[g(X_t^i)] - \mathbb{E}\left[g\left(\hat{X}_{N_t}^i\right)\right] \right| < \varepsilon \quad (12)$$

uniformly over  $\hat{\pi} \in \Pi$ ,  $i \in V'_N$  for some  $V'_N \subseteq V_N$  with  $|V'_N| \geq \lfloor (1-p)N \rfloor$ . Under the same conditions,

$$\sup_{h \in \mathcal{H}} \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t^N, \mathbf{W}_N} \right) - h \left( \hat{X}_t^{\frac{i}{N}}, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}} \right) \right] \right| < \varepsilon \quad (13)$$

holds for any uniformly Lipschitz, uniformly bounded family  $\mathcal{H}$  of functions  $h : \mathcal{X} \times (B_1(\mathcal{X}))^{2k} \rightarrow \mathbb{R}$ .

*Proof:* The first statement (12) can be proven with an argument in [18, Proof of Lemma A.1]. To show that (13) also holds, we will establish that (12) implies (13). Thus, let  $\mathcal{H}$  be defined as in Lemma 2 with uniform Lipschitz constant  $L_h$  and bound  $M_h$ . For any  $h \in \mathcal{H}$ , we have

$$\begin{aligned} & \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t^N, \mathbf{W}_N} \right) \right] - \mathbb{E} \left[ h \left( \hat{X}_t^{\frac{i}{N}}, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}} \right) \right] \right| \\ & \leq \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t^N, \mathbf{W}_N} \right) \right] - \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}_N} \right) \right] \right| \\ & \quad + \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}_N} \right) \right] - \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}} \right) \right] \right| \\ & \quad + \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}} \right) \right] - \mathbb{E} \left[ h \left( \hat{X}_t^{\frac{i}{N}}, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}} \right) \right] \right|. \end{aligned}$$

While the third term can be bounded as in [18, Proof of Lemma A.1], the first two summands converge to zero by the following argument. Starting with the first term, we obtain

$$\begin{aligned} & \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t^N, \mathbf{W}_N} \right) \right] - \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}_N} \right) \right] \right| \\ & \leq \mathbb{E} \left[ \mathbb{E} \left[ \left| h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t^N, \mathbf{W}_N} \right) - h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}_N} \right) \right| \middle| X_t^i \right] \right] \\ & \leq kL_h \sum_{x \in \mathcal{X}} \mathbb{E} \left[ \max_{j \in \{1, \dots, k\}} \left| (W_N^j \mu_t^N - W_N^j \mu_t) \left( \frac{i}{N}, x \right) \right| \right. \\ & \quad \left. + \max_{j \in \{1, \dots, k\}} \left| (\mu_t^N W_N^j - \mu_t W_N^j) \left( \frac{i}{N}, x \right) \right| \right] = o(1) \end{aligned}$$

where we apply Theorem 1 to the uniformly bounded functions  $f'_{N,i,x}(x', \beta) = W_N^j(\frac{i}{N}, \beta) \cdot \mathbf{1}_{x=x'}$  and  $f''_{N,i,x}(x', \beta) = W_N^j(\beta, \frac{i}{N}) \cdot \mathbf{1}_{x=x'}$ . For the second summand, we have

$$\begin{aligned} & \left| \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}_N} \right) \right] - \mathbb{E} \left[ h \left( X_t^i, \mathbb{G}_t^{\frac{i}{N}, \mu_t, \mathbf{W}} \right) \right] \right| \\ & \leq kL_h \sum_{x \in \mathcal{X}} \max_{j \in \{1, \dots, k\}} \left| (W_N^j \mu_t - W^j \mu_t) \left( \frac{i}{N}, x \right) \right| \\ & \quad + kL_h \sum_{x \in \mathcal{X}} \max_{j \in \{1, \dots, k\}} \left| (\mu_t W_N^j - \mu_t W^j) \left( \frac{i}{N}, x \right) \right| \end{aligned}$$

where each of the two summands converges to zero by an argument as in [18, Proof of Lemma A.1]. ■

Thus, we obtain an  $(\varepsilon, p)$ -MNE from a CDMFE.

*Corollary 1:* Assume that Assumptions 1, 2, and 3 hold with  $\boldsymbol{\pi} \in \Pi$  and  $\boldsymbol{\mu} = \Psi(\boldsymbol{\pi})$ . For  $\Gamma_N(\boldsymbol{\pi})$ ,  $\Gamma(\boldsymbol{\pi}, i, \hat{\pi}) \in \Pi^N$  with arbitrary  $\hat{\pi} \in \Pi$ , for any  $\varepsilon, p > 0$ , there exists  $N' \in \mathbb{N}$  such that  $\forall N > N'$  we have uniformly over  $\hat{\pi} \in \Pi$ ,  $i \in V'_N$

$$\left| J_i^N(\pi^1, \dots, \pi^{i-1}, \hat{\pi}, \pi^{i+1}, \dots, \pi^N) - J_i^\mu(\hat{\pi}) \right| < \varepsilon \quad (14)$$

for some  $V'_N \subseteq V_N$  with  $|V'_N| \geq \lfloor (1-p)N \rfloor$ .

*Proof:* Define functions  $r_{\hat{\pi}} := \sum_{u \in \mathcal{U}} r(x, u, v) \hat{\pi}_t(u|s)$  which are uniformly bounded and uniformly Lipschitz. By Lemma 2 there exists  $N' \in \mathbb{N}$  such that for all  $N > N'$

$$\begin{aligned} & \left| J_i^N(\pi^1, \dots, \pi^{i-1}, \hat{\pi}, \pi^{i+1}, \dots, \pi^N) - J_i^\mu(\hat{\pi}) \right| \\ & \leq \sum_{t=0}^{T-1} \left| \mathbb{E} \left[ r_{\hat{\pi}_t}(X_t^i, \mathbb{G}_t^{\frac{i}{N}}(\boldsymbol{\mu}_t)) - r_{\hat{\pi}_t}(\hat{X}_t^{\frac{i}{N}}, \mathbb{G}_t^{\frac{i}{N}}(\boldsymbol{\mu}_t)) \right] \right| < \varepsilon \end{aligned}$$

uniformly over  $\hat{\pi} \in \Pi$ ,  $i \in V'_N$  for some  $V'_N \subseteq V_N$  with  $|V'_N| \geq \lfloor (1-p)N \rfloor$ . See [18, Proof of Corollary A.1]. ■

## REFERENCES

- [1] R. Dominguez and S. Cannella, "Insights on multi-agent systems applications for supply chain management," *Sustainability*, vol. 12, no. 5, p. 1935, 2020.
- [2] S. Vorotnikov, K. Ermishin, A. Nazarova, and A. Yuschenko, "Multi-agent robotic systems in collaborative robotics," in *Int. Conf. Interact. Collab. Robot.* Springer, 2018, pp. 270–279.
- [3] 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," *Commun. Inf. Syst.*, vol. 6, no. 3, pp. 221–252, 2006.
- [4] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese J. Math.*, vol. 2, no. 1, pp. 229–260, 2007.
- [5] P. E. Caines, "Mean field games," in *Encycl. Syst. Control.* Springer, 2021, pp. 1197–1202.
- [6] P. E. Caines and M. Huang, "Graphon mean field games and the GMFG equations," in *IEEE Conf. Decis. Control*, 2018, pp. 4129–4134.
- [7] P. E. Caines and M. Huang, "Graphon mean field games and the GMFG equations:  $\varepsilon$ -Nash equilibria," in *IEEE Conf. Decis. Control*, 2019, pp. 286–292.
- [8] R. Vizuete, P. Frasca, and F. Garin, "Graphon-based sensitivity analysis of SIS epidemics," *IEEE Control Syst. Lett.*, vol. 4, no. 3, pp. 542–547, 2020.
- [9] L. Lovász, *Large Networks and Graph Limits*. Am. Math. Soc., 2012, vol. 60.
- [10] H. Zhang, F. L. Lewis, and Z. Qu, "Lyapunov, adaptive, and optimal design techniques for cooperative systems on directed communication graphs," *IEEE Trans. Ind. Electron.*, vol. 59, no. 7, pp. 3026–3041, 2011.
- [11] D. Y. Kenett, T. Preis, G. Gur-Gershgoren, and E. Ben-Jacob, "Dependency network and node influence: Application to the study of financial markets," *Int. J. Bifurc. Chaos*, vol. 22, no. 07, p. 1250181, 2012.
- [12] L. Lovász and K. Vesztegombi, "Non-deterministic graph property testing," *Comb. Probab. Comput.*, vol. 22, no. 5, pp. 749–762, 2013.
- [13] J. Pérolat et al., "Scaling mean field games by online mirror descent," in *Proc. AAMAS*, 2022, pp. 1028–1037.
- [14] A. Aurell, R. Carmona, G. Dayanikli, and M. Laurière, "Finite state graphon games with applications to epidemics," *Dyn. Games Appl.*, pp. 1–33, 2022.
- [15] R. Elie, J. Perolat, M. Laurière, M. Geist, and O. Pietquin, "On the convergence of model free learning in mean field games," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, 2020, pp. 7143–7150.
- [16] G. Carmona, "Nash equilibria of games with a continuum of players," 2004.
- [17] N. Saldi, T. Basar, and M. Raginsky, "Markov–Nash equilibria in mean-field games with discounted cost," *SIAM J. Control Optim.*, vol. 56, no. 6, pp. 4256–4287, 2018.
- [18] K. Cui and H. Koeppl, "Learning graphon mean field games and approximate Nash equilibria," in *Proc. ICLR*, 2022, pp. 1–31.
- [19] K. Cui and H. Koeppl, "Approximately solving mean field games via entropy-regularized deep reinforcement learning," in *Proc. AISTATS*, 2021, pp. 1909–1917.
- [20] S. Perrin, J. Pérolat, M. Laurière, M. Geist, R. Elie, and O. Pietquin, "Fictitious play for mean field games: Continuous time analysis and applications," *Proc. NeurIPS*, vol. 33, pp. 13199–13213, 2020.