
Learning Sparse Graphon Mean Field Games

Christian Fabian

Technische Universität Darmstadt

Kai Cui

Technische Universität Darmstadt

Heinz Koepl

Technische Universität Darmstadt

Abstract

Although the field of multi-agent reinforcement learning (MRL) has made considerable progress in the last years, solving systems with a large number of agents remains a hard challenge. Graphon mean field games (GMFGs) enable the scalable analysis of MRL problems that are otherwise intractable. By the mathematical structure of graphons, this approach is limited to dense graphs which are insufficient to describe many real-world networks such as power law graphs. Our paper introduces a novel formulation of GMFGs, called LPGMFGs, which leverages the graph theoretical concept of L^p graphons and provides a machine learning tool to efficiently and accurately approximate solutions for sparse network problems. This especially includes power law networks which are empirically observed in various application areas and cannot be captured by standard graphons. We derive theoretical existence and convergence guarantees and give empirical examples that demonstrate the accuracy of our learning approach for systems with many agents. Furthermore, we extend the Online Mirror Descent (OMD) learning algorithm to our setup to accelerate learning speed, empirically show its capabilities, and conduct a theoretical analysis using the novel concept of smoothed step graphons. In general, we provide a scalable, mathematically well-founded machine learning approach to a large class of otherwise intractable problems of great relevance in numerous research fields.

1 INTRODUCTION

In various research areas, scientists are confronted with systems of many interacting individuals or components.

Examples include neurons in the human brain (Avena-Koenigsberger et al. (2018), Bullmore and Sporns (2009), Bullmore and Sporns (2012), Sporns (2022)), people trading on a stock market (Bakker et al. (2010), Bian et al. (2016)) or the spreading of contagious diseases among citizens of a society (Newman (2002), Pastor-Satorras et al. (2015)). Due to their complexity, these systems are in general hard to model and are often controlled by using multi-agent reinforcement learning (MRL). In the last years, the field of MRL has experienced significant progress, see Canese et al. (2021) or Yang and Wang (2020) for an overview, but crucial open problems remain. While many approaches provide sound empirical results, they often lack a solid theoretical foundation (Zhang et al. (2021)). Furthermore, as the number of agents in the system increases, numerous MRL algorithms become computationally expensive and are thereby hardly scalable.

Mean field games (MFGs) (Carmona and Delarue (2018a), Carmona and Delarue (2018b)) have proven to be a valid approach for achieving both scalability and theoretical guarantees in multi-agent systems. Since they were introduced independently by Huang et al. (2006) and Lasry and Lions (2007) to address game theoretic challenges, they have become a major interest in various research fields. Extensions of the original MFG model include discrete-time formulations (Cui and Koepl (2022), Saldi et al. (2018)), variants with major and minor agents (Carmona and Zhu (2016), Firoozi et al. (2020), Nourian and Caines (2013)) as well as zero-sum games (Choutri and Djehiche (2019), Choutri et al. (2019)). MFGs are based on the weak interaction principle where each individual has a negligible influence on the whole system. Besides the numerical and theoretical benefits of this principle, MFGs provide the modelling framework for various applications, such as autonomous driving (Huang et al. (2020)), cyber security (Kolokoltsov and Malafeyev (2018)), big data architectures (Castiglione et al. (2014)), and systemic risk in financial markets (Carmona et al. (2015), Elie et al. (2020a)). There is also some work that aims to apply MFGs to real world tasks, e.g. social networks (Yang et al. (2018)) or swarm robotics (Elamvazhuthi and Berman (2019), Cui et al. (2022)), but this field largely remains to be developed. Although our paper is of theoretical nature, its goal is to make MFGs more realistic, as we discuss below.

Both from the classical equilibrium learning perspective and the reinforcement learning (RL) perspective, MFGs are able to provide solutions for numerous challenges where other equilibrium learning or MARL algorithms become computationally intractable. Here, learning refers to both classical computation of equilibria and RL – also known as approximate optimal control (Bertsekas (2019)) – with focus on solving complex control problems without knowing or using the model. Current RL research is addressing the approximation of Nash equilibria for stationary games (Subramanian and Mahajan (2019)), under the occurrence of noise (Carmona et al. (2019)), using entropy regularization (Anahtarci et al. (2020), Cui and Koepl (2021), Guo et al. (2022)), leveraging Fictitious Play (Cardaliaguet and Hadikhanloo (2017), Delarue and Vasileiadis (2021), Hadikhanloo and Silva (2019), Mguni et al. (2018), Perrin et al. (2021a), Perrin et al. (2021b), Perrin et al. (2020)), and increasing the robustness and efficiency of learning algorithms in general (Guo et al. (2019), Guo et al. (2020)). A learning scheme of particular interest for our paper is Online Mirror Descent (OMD) (Orabona et al. (2015), Srebro et al. (2011)). RL research has leveraged OMD to learn MFGs (Hadikhanloo (2017), Laurière et al. (2022), Perolat et al. (2021)) which ensures algorithmic scalability.

For learning applications, decision-making on graphs appears to be particularly interesting. Here, we refer to agents connected via graphical edges, as opposed to agents with states on a graph as considered e.g. in Li et al. (2019). Apart from direct graphical decompositions in MARL Qu et al. (2020), there has been recent research interest in MFGs on graphs. In classical MFGs each agent weakly interacts with all other agents at once which seems to be an unrealistic modelling assumption for many applications. To overcome this concern, graphon mean field games (Aurell et al. (2021), Caines and Huang (2019), Carmona et al. (2022), Gao and Caines (2017), Gao et al. (2021)) (GMFGs) provide a well-established tool to model games with a graphical structure. For example, Tangpi and Zhou (2022) apply GMFGs to model investment decisions in a financial market. Also based on GMFGs, Aurell et al. (2022) develop models on epidemics and provide the corresponding machine learning methods for outcome estimation. So far, most of the literature has focused on MFGs on dense graphs. To the best of our knowledge, there are only a few papers that consider sparse graphs. While Gkogkas and Kuehn (2022) focuses on Kuramoto-like models, Lacker and Soret (2022) is concerned with linear-quadratic stochastic differential games. Finally, Bayraktar et al. (2020) considers systems on not-so-dense graphs but without control and without leveraging L^p graphons. By utilizing L^p graphons (Borgs et al. (2018), Borgs et al. (2019)), our paper’s aim is to provide a general framework for learning MFGs on sparse graphs.

Sparse power law graphs (Barabási and Albert (1999), Barabási et al. (1999)) are of great interest for various

research applications such as social networks (Aparicio et al. (2015)), software engineering (Concas et al. (2007), Louridas et al. (2008), Wheeldon and Counsell (2003)), finance (D’Arcangelis and Rotundo (2016)), or biology (Nosonovsky and Roy (2020)). For more examples, see Newman (2018). Although there is strong empirical evidence for power law graphs in many research fields as mentioned above, GMFGs cannot capture these naturally sparse structures. LPGMFGs and the corresponding learning methods presented in our paper provide a novel ML tool to solve such real-world problems that are otherwise intractable. Our contributions can be summarized as follows: (i) introducing MFGs on L^p Graphons (LPGMFGs) which formalize MFGs on sparse graphs; (ii) conducting a theoretical analysis of LPGMFGs that includes the existence of equilibria as well as convergence guarantees; (iii) evaluating LPGMFGs on different examples empirically, especially in a multi-class agent setup; (iv) adapting the OMD learning scheme to our setup and thereby accelerating learning speed; (v) conducting both an empirical and a theoretical convergence analysis for the adapted OMD algorithm. Thus, our paper provides a scalable, mathematically well founded approach for learning MARL problems on sparse graphs on the theoretical side. On the practical side, different empirical examples demonstrate the scalability of the learning method and give an impression of how models on sparse graphs are often more realistic than dense networks.

2 L^p GRAPHONS

Central Concepts. In this section, we briefly introduce the concept of L^p graphons pioneered by Borgs et al. (2018) and Borgs et al. (2019) which provide more details. L^p graphons can be informally thought of as adjacency matrices for graphs with (almost) infinitely many nodes. Naturally, approximating sparse finite graphs by these L^p graphons leads to the loss of some topological information, see e.g. Borgs et al. (2018). Nevertheless, L^p graphons provide an expressive asymptotic approximation of the finite case which we show both theoretically and empirically on the next pages. In contrast to standard graphons which are limited to dense graphs, L^p graphons are more general and also apply to sparse graphs. An L^p graphon is a symmetric, integrable function $W: [0, 1]^2 \rightarrow \mathbb{R}$ with $\|W\|_p < \infty$ where the L^p norm on graphons is $\|W\|_p := \left(\int_{[0,1]^2} |W(x, y)|^p dx dy \right)^{1/p}$ for $1 \leq p < \infty$ and the essential supremum if $p = \infty$.

To quantify whether a graphon is a good approximation for a sequence of finite graphs, we associate every finite graph $G = (V(G), E(G))$ with a graphon W^G . For a graph G with N nodes, we partition the unit interval $[0, 1]$ into N intervals I_1, \dots, I_N of equal length. Then, the function W^G is assigned a constant value on each square $I_i \times I_j$ ($i, j \in V(G)$) which is equal to one if there is an edge between

the nodes i, j in G and zero otherwise. Thus, W^G is by construction a step-function and therefore often called step-graphon. To compare some graph G and graphon W , we can simply compare the graphons W and W^G in the space of graphons. Instead of W^G itself, we frequently consider the normalized associated graphon $W_G/\|G\|_1$ to derive results that also hold in the sparse case. To measure how close two graphons are, the cut norm is a natural candidate and possesses many useful properties, see e.g. Lovász (2012) for a detailed discussion. For a graphon $W: [0, 1]^2 \rightarrow \mathbb{R}$, define the cut norm by

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

where S and T range over the measurable subsets of $[0, 1]$.

Starting with a graphon, we can use the following well-established construction to generate sparse random graphs with N nodes where we assume without loss of generality that the vertices are labeled by $1, \dots, N$. We choose x_1, \dots, x_N i.i.d. uniformly at random in $[0, 1]$ and fix some $\rho > 0$. For all vertex pairs $1 \leq i < j \leq N$ there is an edge between i and j with probability $\min\{\rho W(x_i, x_j), 1\}$ which yields a sparse random graph $\mathbf{G}(N, W, \rho)$. A sequence of sparse random graphs generated by this method converges to the generating graphon W in the cut norm, see (Borgs et al., 2019, Theorem 2.14).

Assumption 1. *The sequence of normalized step-graphons $(W_N)_{N \in \mathbb{N}}$ converges in cut norm $\|\cdot\|_{\square}$ or equivalently in operator norm $\|\cdot\|_{L_{\infty} \rightarrow L_1}$ (see Lovász (2012)) as $N \rightarrow \infty$ to some graphon $W \in \mathcal{W}_0$, i.e.*

$$\|\rho_N^{-1} W_N - W\|_{\square} \rightarrow 0, \quad \|\rho_N^{-1} W_N - W\|_{L_{\infty} \rightarrow L_1} \rightarrow 0. \quad (1)$$

The limiting graphon in Assumption 1 is only guaranteed to exist for so-called L^p upper regular graph sequences, for details see Borgs et al. (2018), Borgs et al. (2019). This implicit assumption means that our approach does not apply to arbitrary sequences of sparse random graphs. If the average degree in the graph sequence does not tend to infinity as $N \rightarrow \infty$, L^p upper regularity is not fulfilled. Thus, for example, the asymptotic behavior of ultra-sparse graph sequences is beyond the scope of L^p graphons. Nevertheless, L^p graphons are the limit of crucial sparse graph sequences such as power law graphs which cannot be provided by standard graphons. In our paper, 'power law' refers to the tail of the distribution. Figure 1 shows the advantages of L^p graphons over standard graphons using an exemplary real-world network (data from Rozemberczki et al. (2019)). The examples in the next sections are usually based on power law graphons $W: [0, 1]^2 \rightarrow \mathbb{R}$ with $W(x, y) = (1 - \alpha)^2 (xy)^{-\alpha}$ where $\alpha \in (0, 1)$, see Borgs et al. (2018) for details.

Smoothing Step Graphons. For the theoretical analysis of the OMD algorithm in the next sections, we introduce the concept of smoothed step graphons which is new to the best of our knowledge. The basic idea is to smooth the borders of the steps. Then, the smoothed step graphon is Lipschitz continuous but still close to the original step graphon as we decrease the width ξ of the border regions.

Consider an arbitrary step graphon W_s on the unit interval partitioned into M subintervals $\mathcal{I}_1, \dots, \mathcal{I}_M$ of equal length $1/M$ such that $W_s(x, y) = w_{i,j} \geq 0$ for all $(x, y) \in \mathcal{I}_i \times \mathcal{I}_j, 1 \leq i, j \leq M$. Then, for an arbitrary but fixed $0 < \xi < 1/(2M)$, we define the corresponding smoothed step graphon $W_{s,\xi}$ as follows. For all $(x, y) \in \{(x, y) \in [0, 1]^2 : (x \leq \xi) \vee (x \geq 1 - \xi)\}$ and $(x, y) \in \{(y \leq \xi) \vee (y \geq 1 - \xi)\}$ we define $W_{s,\xi}(x, y) := W_s(x, y)$. The values of the two graphons are also defined to be identical for (x, y) with $(x, y) \in \tilde{\mathcal{I}}_i \times \tilde{\mathcal{I}}_j$ for some $1 \leq i, j \leq M$ where $\tilde{\mathcal{I}}_i := [(i-1)/M + \xi, i/M - \xi)$ for all $1 \leq i \leq M$. In contrast to that, if $x \in \tilde{\mathcal{I}}_i$ and $y \in [j/M - \xi, j/M + \xi)$ for some $1 \leq i, j \leq M$, we have $W_{s,\xi}(x, y) := (\frac{1}{2} - \frac{y-j/M}{2\xi})w_{i,j} + \frac{y-j/M+\xi}{2\xi}w_{i,j+1}$ and analogously for x and y with switched roles. Finally, if both $x \in [i/M - \xi, i/M + \xi)$ and $y \in [j/M - \xi, j/M + \xi)$ for some $1 \leq i, j \leq M$, we define

$$\begin{aligned} W_{s,\xi}(x, y) := & \left(\frac{1}{2} - \frac{x-i/M}{2\xi}\right) \left(\frac{1}{2} - \frac{y-j/M}{2\xi}\right) w_{i,j} \\ & + \left(\frac{1}{2} - \frac{x-i/M}{2\xi}\right) \frac{y-j/M+\xi}{2\xi} w_{i,j+1} \\ & + \frac{x-i/M+\xi}{2\xi} \left(\frac{1}{2} - \frac{y-j/M}{2\xi}\right) w_{i+1,j} \\ & + \frac{x-i/M+\xi}{2\xi} \cdot \frac{y-j/M+\xi}{2\xi} w_{i+1,j+1}. \end{aligned}$$

Note that $W_{s,\xi}$ is Lipschitz continuous and that by construction we have $\|W_{s,\xi} - W_s\|_{\square} \leq 4M\xi \cdot \max_{i,j} w_{i,j}$ which approaches zero as $\xi \rightarrow 0$.

3 MODEL

Finite Agent Model. For the finite case, we assume that there are N agents with finite state and action spaces \mathcal{X} and \mathcal{U} , respectively. The agents implement actions at time points $\mathcal{T} := \{0, \dots, T-1\}$ with terminal time point T . The interactions between individuals are modeled by a graph $G_N = (V_N, E_N)$ where each vertex represents one agent and each edge a connection between two agents. For an arbitrary finite set A we denote by $\mathcal{P}(A)$ the set of all probability measures on A and by $\mathcal{B}(A)$ the set of all bounded measures on A . Thus, the space of policies is defined as $\Pi := \mathcal{P}(\mathcal{U})^{\mathcal{T} \times \mathcal{X}}$ and a policy of agent i is denoted by $\pi^i = (\pi_t^i)_{t \in \mathcal{T}} \in \Pi$ correspondingly. Furthermore, agents in the model are assumed to implement Markovian feedback policies such that they only consider local state information.




Figure 1: Three networks and their empirical and mathematically expected degree distributions (DD): Erdős-Rényi graph (left), real-world Facebook network (middle, data from Rozemberczki et al. (2019)), power law graph (right, for expected DD see Bollobás et al. (2007)); highly connected nodes are larger and darker: the Facebook graph has some nodes with high degrees and small and medium degrees otherwise. The power law graph generated by an L^p graphon is qualitatively similar. All nodes in the ER graph generated by a standard graphon have degrees smaller than thirty which contradicts the real-world network. Other standard graphons, e.g. ranked attachment (Borgs et al. (2011)), yield similar results as the ER graph but are omitted for space reasons. Each graph consists of 3892 nodes and around 17500 edges to match the real data set.

Formally, this is captured by defining for all $t \in \mathcal{T}$ and $i \in V_N$ the model dynamics

$$U_t^i \sim \pi_t^i(\cdot | X_t^i), \quad X_{t+1}^i \sim P(\cdot | X_t^i, U_t^i, \mathbb{G}_t^i) \quad (2)$$

with $X_0^i \sim \mu_0$ for some transition kernel $P: \mathcal{X} \times \mathcal{U} \times \mathcal{B}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$ and the neighborhood state distribution

$$\mathbb{G}_t^i := \frac{1}{N\rho_N} \sum_{j \in V_N} \mathbf{1}_{\{ij \in E_N\}} \delta_{X_t^j} \quad (3)$$

for each agent i with δ being the Dirac measure and $\mathbb{G}_t^i \in \mathcal{B}(\mathcal{X})$ for all $i \leq N$ by definition. The normalization factor ρ_N ensures that the neighborhood distribution does not converge to a vector of zeros as N approaches infinity. Here, ρ_N is assumed to have the same asymptotic order as the edge density of G_N , i.e. $\rho_N = \Theta(|E_N|/N^2)$ as $N \rightarrow \infty$. Each agent faces a reward function $r: \mathcal{X} \times \mathcal{U} \times \mathcal{B}(\mathcal{X}) \rightarrow \mathbb{R}$ which yields her reward depending on her state, action, and the state distribution of her neighbors. Agents competitively try to maximize their expected sum of rewards

$$J_i^N(\pi^1, \dots, \pi^N) := \mathbb{E} \left[\sum_{t=0}^{T-1} r(X_t^i, U_t^i, \mathbb{G}_t^i) \right]. \quad (4)$$

Note that we can handle the infinite-horizon discounted objective case analogously, see e.g. Cui and Koepl (2022). Finding equilibria for this type of model requires a suitable equilibrium concept. Although the classical Nash equilibrium notion (see, e.g. Carmona and Delarue (2018a)) seems to be a natural candidate, its definition requires that no agent has an incentive to unilaterally deviate from the current policy. As we are primarily interested in an approximation via L^p graphons, this equilibrium concept is too strict. Even in

the limit $N \rightarrow \infty$ there can always occur (small) subgroups of agents whose graph connections deviate from the structure of the underlying graphon. Therefore, we work with the (ϵ, p) -Markov-Nash equilibrium (see, for example Carmona (2004), Elie et al. (2020b), Cui and Koepl (2022)) which only requires optimality for a fraction $1-p$ of all individuals. This fraction will increase, $(1-p) \rightarrow 1$ as $N \rightarrow \infty$.

Definition 1. An (ϵ, p) -Markov-Nash equilibrium (MNE) for $\epsilon, p > 0$ is defined as a tuple of policies $\pi = (\pi^1, \dots, \pi^N) \in \Pi^N$ such that for any $i \in \mathcal{W}_N$ we have

$$J_i^N(\pi) \geq \sup_{\pi \in \Pi} J_i^N(\pi^1, \dots, \pi^{i-1}, \pi, \pi^{i+1}, \dots, \pi^N) - \epsilon \quad (5)$$

for some $\mathcal{W}_N \subseteq V_N$ with $|\mathcal{W}_N| \geq \lfloor (1-p)N \rfloor$ such that \mathcal{W}_N contains at least $\lfloor (1-p)N \rfloor$ agents.

Mean Field Model. The L^p graphon mean field model (LPGMFG) constitutes the limit of the finite agent model as $N \rightarrow \infty$ and provides a reasonable approximation for the finite case. Before we formalize this claim and provide rigorous statements, we introduce the LPGMFG itself. The main difference to the N -agent model is that we now consider an infinite number of agents $\alpha \in \mathcal{I} := [0, 1]$. Thus, $\mathcal{M}_t := \mathcal{P}(\mathcal{X})^{\mathcal{I}}$ denotes the space of measurable state marginal ensembles at time t , and $\mathcal{M} := \mathcal{P}(\mathcal{X})^{\mathcal{I} \times \mathcal{T}}$ the space of measurable mean field ensembles. Here, measurable means that $\alpha \mapsto \mu_t^\alpha(x)$ is measurable for all $\mu \in \mathcal{M}, t \in \mathcal{T}, x \in \mathcal{X}$. Analogously, a space of uniformly Lipschitz, measurable policy ensembles $\Pi \subseteq \Pi^{\mathcal{I}}$ is defined such that $\alpha \mapsto \pi_t^\alpha(u|x)$ is measurable and L_Π -Lipschitz for any $\pi \in \Pi, t \in \mathcal{T}, u \in \mathcal{U}, x \in \mathcal{X}$. Intuitively, a policy ensemble $\pi \in \Pi$ includes an infinite number of policies

$\pi^\alpha \in \Pi$ where each policy is associated with one agent α . State and action variables are defined for all $(\alpha, t) \in \mathcal{I} \times \mathcal{T}$ as

$$U_t^\alpha \sim \pi_t^\alpha(\cdot | X_t^\alpha), \quad X_{t+1}^\alpha \sim P(\cdot | X_t^\alpha, U_t^\alpha, \mathbb{G}_t^\alpha), \quad (6)$$

with $X_0^\alpha \sim \mu_0$ where the deterministic neighborhood MF of agent α for some deterministic MF $\mu \in \mathcal{M}$ is

$$\mathbb{G}_t^\alpha := \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta d\beta \quad (7)$$

with $\mathbb{G}_t^\alpha \in \mathcal{B}(\mathcal{X})$ by definition. Each agent tries to competitively maximize her rewards given by

$$J_\alpha^\mu(\pi^\alpha) := \mathbb{E} \left[\sum_{t=0}^{T-1} r(X_t^\alpha, U_t^\alpha, \mathbb{G}_t^\alpha) \right]. \quad (8)$$

Now, it remains to adapt the Nash equilibrium concept to the LPGMFG case. Thus, we introduce two functions $\Psi: \Pi \rightarrow \mathcal{M}$ and $\Phi: \mathcal{M} \rightarrow 2^\Pi$. Ψ maps a policy ensemble $\pi \in \Pi$ to the mean field ensemble $\mu = \Psi(\pi) \in \mathcal{M}$ generated by π which is formalized by the recursive equation

$$\mu_{t+1}^\alpha(x) = \sum_{\substack{x' \in \mathcal{X} \\ u \in \mathcal{U}}} \mu_t^\alpha(x') \pi_t^\alpha(u|x') P(x|x', u, \mathbb{G}_t^\alpha) \quad (9)$$

for all $\alpha \in [0, 1]$ with $\mu_0^\alpha \equiv \mu_0$. The second map $\Phi: \mathcal{M} \rightarrow 2^\Pi$ takes a mean field ensemble $\mu \in \mathcal{M}$ and maps it to the set of policy ensembles $\Phi(\mu) \subseteq 2^\Pi$ that are optimal with respect to μ , i.e. $\pi^\alpha = \arg \max_{\pi \in \Pi} J_\alpha^\mu(\pi^\alpha)$ for all $\alpha \in [0, 1]$. With the above definitions, we can state the equilibrium concept for LPGMFGs, namely the L^p graphon mean field equilibrium (GMFE).

Definition 2. A GMFE is a tuple $(\mu, \pi) \in \Pi \times \mathcal{M}$ such that $\pi \in \Phi(\mu)$ and $\mu = \Psi(\pi)$.

We also refer to the policy part of a GMFE as its Nash Equilibrium (NE). Intuitively, a GMFE consists of a policy ensemble π and a MF ensemble μ such that π generates μ and is also an optimal response to the generated MF. We will frequently use a Lipschitz assumption common in the MFG literature (Bayraktar et al. (2020), Carmona and Delarue (2018a), Cui and Koepl (2022)) to enable the derivation of expressive theoretical results. The power law graphon, however, is not Lipschitz, so we derive a Lipschitz cutoff version in Appendix N. Since this cutoff power law graphon does not yield qualitatively different results compared to the power law graphon, it is omitted from the main text.

Assumption 2. Let r, P, W be Lipschitz continuous with Lipschitz constants $L_r, L_P, L_W > 0$, or alternatively there exist $L_W > 0$ and disjoint intervals $\{\mathcal{I}_1, \dots, \mathcal{I}_Q\}$, $\cup_i \mathcal{I}_i = \mathcal{I}$ s.t. $\forall i, j \leq Q$ and $\forall(x, y), (\tilde{x}, \tilde{y}) \in \mathcal{I}_i \times \mathcal{I}_j$,

$$|W(x, y) - W(\tilde{x}, \tilde{y})| \leq L_W(|x - \tilde{x}| + |y - \tilde{y}|). \quad (10)$$

Under Assumption 2, the model defined above has a GMFE which is formalized by the next theorem.

Theorem 1. Under Assumption 2 and for Lipschitz W , there exists a GMFE $(\pi, \mu) \in \Pi \times \mathcal{M}$.

Proof. The existence of a GMFE follows from (Saldi et al., 2018, Theorem 3.3) for the extended state space $\mathcal{X} \times [0, 1]$. See also (Cui and Koepl, 2022, Proof of Theorem 1). \square

Mean Field Approximation. The proofs of all theoretical results can be found in the Appendix. This paragraph relates the finite agent model to the MF model by showing that LPGMFGs yield an increasingly accurate approximation for the N -agent case as the number of agents grows. We emphasize that the LPGMFG yields an approximation for the N -agent game for all N at once. Both the theoretical results as well as the empirical findings show that the accuracy of this approximation increases with the number N of agents, see the next sections for details. As a consequence, the LPGMFG concept provides a scalable and increasingly accurate approximation for otherwise intractable multi agent problems with a large number of individuals. By Theorem 1, there exists a GMFE (π, μ) which yields an approximate NE for the N -agent problem through the map $\Gamma_N(\pi) := (\pi^1, \dots, \pi^N) \in \Pi^N$ defined by $\pi_t^i(u|x) := \pi_t^{\alpha_i}(u|x)$ for all $\alpha \in \mathcal{I}$, $t \in \mathcal{T}$, $x \in \mathcal{X}$, $u \in \mathcal{U}$ with $\alpha_i = i/N$.

For a theoretical comparison, we lift both the policies and empirical distributions in the finite agent model to the continuous domain \mathcal{I} . Thus, for an N -agent policy tuple $(\pi^1, \dots, \pi^N) \in \Pi^N$ the corresponding step policy ensemble $\pi^N \in \Pi$ and the random empirical step measure ensemble $\mu^N \in \mathcal{M}$ are defined by $\pi_t^{N, \alpha} := \sum_{i \in V_N} \mathbf{1}_{\{\alpha \in (\frac{i-1}{N}, \frac{i}{N}]\}} \cdot \pi_t^i$ and $\mu_t^{N, \alpha} := \sum_{i \in V_N} \mathbf{1}_{\{\alpha \in (\frac{i-1}{N}, \frac{i}{N}]\}} \cdot \delta_{X_t^i}$ for all $\alpha \in \mathcal{I}$ and $t \in \mathcal{T}$. For notational convenience, we furthermore define for any $f: \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ and state marginal ensemble $\mu_t \in \mathcal{M}_t$, $\mu_t(f) := \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} f(x, a) \mu_t^\alpha(x) d\alpha$. With these definitions in place, we state our first main theoretical result.

Theorem 2. Consider $\pi \in \Pi$ with $\mu = \Psi(\pi)$. Under Assumption 1 and the N -agent policy $(\pi^1, \dots, \pi^{i-1}, \hat{\pi}, \pi^{i+1}, \dots, \pi^N) \in \Pi^N$ with $(\pi^1, \pi^2, \dots, \pi^N) = \Gamma_N(\pi) \in \Pi^N$, $\hat{\pi} \in \Pi$, $t \in \mathcal{T}$, we have for all measurable functions $f: \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ uniformly bounded by some $M_f > 0$, that

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

uniformly over all possible deviations $\hat{\pi} \in \Pi$, $i \in V_N$. If the graphon convergence in Assumption 1 is up to rate $O(1/\sqrt{N})$, then this rate of convergence is the same.

Based on Theorem 2, we can derive a central result of this paper which formalizes the capability of LPGMFGs to approximate finite N -agent models. In contrast to prior work, proving Theorem 2 requires an additional mathematical effort which is discussed in Appendix A.

Theorem 3. Consider a GMFE $(\boldsymbol{\pi}, \boldsymbol{\mu})$ under Assumptions 1 and 2. For any $\varepsilon, p > 0$ there exists N' such that for all $N > N'$, the policy $\Gamma_N(\boldsymbol{\pi}) \in \Pi^N$ is an (ε, p) -MNE.

Intuitively, Theorem 3 states that the GMFE provides an increasingly accurate approximation of the N -agent problem as the number of agents goes up. Since the algorithmic computation of NE is in general intractable (Conitzer and Sandholm (2008), Papadimitriou (2001), Papadimitriou (2007)), the LPGMFGs approximation can overcome these difficulties by choosing ε and p in Theorem 3 close to zero when the number N of agents is sufficiently large.

4 LEARNING LPGMFGS

Equivalence Class Method. For learning equilibria in LPGMFGs, we introduce equivalence classes (Cui and Koepl (2022)). We discretize the continuous interval \mathcal{I} of agents by some finite number M of subintervals that form a partition of \mathcal{I} . For convenience, we usually assume that every subinterval has the same length. Then, all agents within one class, i.e. a subinterval, are approximated by the agent who is located at the center of the respective subinterval. Subsequently, we can solve the optimal control problem for each equivalence class separately by applying either backwards induction or RL. Although this formulation seems to resemble classical multi-population mean field games (MP MFGs) (Huang et al. (2006), Perolat et al. (2021)) at first, the crucial advantages of LPGMFGs are that they are on the one hand rigorously connected to finite agent games. On the other hand, they can handle an uncountable number of agent equivalence classes that cannot be captured by the standard multi-class model. Beyond that, the just described learning method for LPGMFGs does not just provide an approximation for some finite N -agent problem with a fixed N . Instead, it yields an estimation for the N -agent problem for all arbitrary, large enough N at once. The technical details of the approach can be found in Appendix I.

Online Mirror Descent (OMD). The discretized game generated by the equivalence class method can be interpreted as a MP MFG with M populations. In the literature, the concept of OMD is used to learn equilibria in such MP MFGs (Hadikhanloo (2017), Perolat et al. (2021)). Our paper leverages these concepts to learn LPGMFGs.

To prove convergence for the OMD algorithm, we have to ensure that a NE exists. Here, the discretized GMFG can be interpreted as a GMFG on the step-graphon W_s created by discretization. To facilitate the theoretical analysis of the OMD algorithm, we consider the corresponding smoothed step graphon $W_{s,\xi}$ and the smoothed GMFG given by the dynamics $\hat{U}_t^\alpha \sim \pi_t^\alpha(\cdot | \hat{X}_t^\alpha)$ and $\hat{X}_{t+1}^\alpha \sim P(\cdot | \hat{X}_t^\alpha, \hat{U}_t^\alpha, \hat{\mathbb{G}}_t^\alpha)$, with $\hat{X}_0^\alpha \sim \mu_0$ for all $(\alpha, t) \in \mathcal{I} \times \mathcal{T}$ where $\hat{\mathbb{G}}$ is the neighborhood state distribution for the smoothed step graphon. One advantage of this approach is that the existence of a

GMFE $(\boldsymbol{\mu}_{s,\xi}, \boldsymbol{\pi}_{s,\xi})$ is ensured by Theorem 1 for this GMFG. Also, for ξ close enough to zero, the smoothed step graphon converges to the original step graphon in the cut norm.

Theorem 4. Suppose that $(\boldsymbol{\mu}_{s,\xi}, \boldsymbol{\pi}_{s,\xi})$ is a GMFE in the smoothed version of the MP MFG on the step graphon W under Assumption 2. Then, for every $\varepsilon, p > 0$ there exists a $\xi' > 0$ such that for all $0 < \xi < \xi'$

$$\sup_{\boldsymbol{\pi} \in \Pi} J_{\alpha,W}(\boldsymbol{\pi}_{s,\xi}) - J_{\alpha,W}(\boldsymbol{\pi}) \leq \varepsilon \quad (12)$$

for all $\alpha \in \mathcal{J}$ for some $\mathcal{J} \subseteq \mathcal{I}$ with Lebesgue measure $\lambda(\mathcal{J}) \geq 1 - p$.

This means that a GMFE in the smoothed version of the GMFG is an (ε, p) -MNE for the discretized game. Combining this insight with existing results (Cui and Koepl, 2022, Theorem 5) indicates that the smoothed GMFE provides a good approximation for the finite agent case, but we leave a rigorous proof for future work. We call a smoothed MP MFG weakly monotone if for any $\boldsymbol{\pi}, \boldsymbol{\pi}' \in \Pi$ we have

$$\tilde{d}(\boldsymbol{\pi}, \boldsymbol{\pi}') := \int_{\mathcal{I}} \left[J_\alpha^\mu(\pi^\alpha) + J_\alpha^{\mu'}(\pi'^\alpha) - J_\alpha^\mu(\pi'^\alpha) - J_\alpha^{\mu'}(\pi^\alpha) \right] d\alpha \leq 0 \quad (13)$$

where $\boldsymbol{\mu} = \Psi(\boldsymbol{\pi})$ and $\boldsymbol{\mu}' = \Psi(\boldsymbol{\pi}')$ are the MFs associated with the respective policies. If the inequality is strict $\forall \boldsymbol{\pi} \neq \boldsymbol{\pi}'$, we call the MP MFG strictly weakly monotone.

Assumption 3. The smoothed MP MFG is strictly weakly monotone.

Weak monotonicity can be interpreted as agents preferring less crowded areas over crowded ones. Under Assumption 3, the NE guaranteed by Theorem 1 is unique.

Lemma 1. If the smoothed MP MFG satisfies Assumptions 2 and 3, it has a unique NE.

We define the OMD algorithm as in Perolat et al. (2021) and consider the continuous time case (CTOMD) where we denote the time of the algorithm by $\tau > 0$. Then, we obtain the following convergence result.

Theorem 5. If the smoothed MP MFG satisfies Assumptions 2 and 3 and the transition kernel does not depend on the MF, the sequence of policies $(\boldsymbol{\pi}_\tau)_{\tau \geq 0}$ generated by the CTOMD algorithm converges to the unique NE as $\tau \rightarrow \infty$.

Empirical evidence collected in our simulations suggests that a convergence guarantee as in Theorem 5 also holds for the case where the MF depends on the transition kernel, but we leave a rigorous proof for future work.

5 EXPERIMENTS

In our experiments, we use OMD with its hyperparameter γ set to 1, the power law L^p -graphon W , and discretize \mathcal{I}




Figure 2: Experimental results for OMD on the Cyber Security problem. (a): The exploitability ΔJ over iterations n of OMD; (b)-(e): The probability of choosing action $u = 1$ at graphon index α and time t under the final equilibrium policy in states DI, DS, UI, US respectively; (f): The probability (mean-field) of infected agents, visualized for each discretized α .

into $M = 25$ subintervals for the Cyber Security problem or $M = 10$ for the Beach Bar problem given as follows. Here, we emphasize that using L^p -graphons in the experiments is a key component of our LPGMFG approach. This allows us to model many realistic networks which are characterized by sparsity and power law degree distributions. As we discussed previously, standard GMFG approaches are conceptually unable to capture these networks.

Cyber Security. We modify an existing cyber security model (Carmona and Delarue (2018a), Kolokoltsov and Bensoussan (2016)) where a virus spreads to computers either directly by an attack, or by other nearby infected computers. In contrast to existing work, we use LPGMFGs to allow malware spread only by neighboring computers to increase the modelling accuracy. Each computer can be either infected (I) or susceptible (S), as well as defended (D) or unprotected (U), formally $\mathcal{X} := \{DI, DS, UI, US\}$. Agents may attempt to switch (with geometrically distributed delay) between defense states, $\mathcal{U} := \{0, 1\}$. The recovery and infection probabilities depend on the defense state and number of infected neighbors, while the reward function consists of costs for being defended or infected. Details can be found in Appendix O.1.

Heterogeneous Cyber Security. A natural extension of the cyber security model is the adaptation to a multi-class framework with heterogeneous agent classes. For illustrative purposes we focus on only two types of agents – private (Pri) and corporate (Cor), see Appendix O.2 for details.

Beach Bar Process. Introduced as the Santa Fe bar problem (Arthur (1994), Farago et al. (2002)), variations of the

Beach Bar Process are frequently used to demonstrate the capabilities of learning algorithms (Perolat et al. (2021), Perrin et al. (2020)). Agents can move their towels between locations and try to be close to the bar but also avoid crowded areas and neighbors in an underlying network. Formally, we consider a one-dimensional beach bar process with $|\mathcal{X}| = 10$ locations $\mathcal{X} = \{0, 1, \dots, |\mathcal{X}| - 1\}$ where a bar is located in the middle $B = |\mathcal{X}|/2$ of the beach. The N agents may move their towel between locations, $\mathcal{U} = \{-1, 0, 1\}$. Furthermore, the agents are connected by a power law network where connected agents try to avoid being close to each other. See Appendix O.3 for details.

Experimental Results. As seen in Figures 2 and 3, the approximate exploitability $\Delta J(\pi) = \int_{\mathcal{I}} \sup_{\pi^* \in \Pi} J_\alpha^\Psi(\pi^*)(\pi^*) - J_\alpha^\Psi(\pi)(\pi) d\alpha$ of a MF policy π quantifies the sub-optimality of the obtained equilibrium and quickly converges in the Cyber Security and Beach Bar scenario using OMD. We obtain near-identical results also for the Heterogeneous Cyber Security problem, which are omitted for space reasons. The algorithm converges to an equilibrium where, as expected, the agents with the most connections in the graph attempt to defend at fixed cost, as their expected cost from not defending is higher than for agents with fewer connections. The system quickly runs into an almost time-stationary state, where the costs of defending equilibrate with the expected cost of future infection, see Figure 2. Since we consider a finite-horizon however, the option of defending becomes increasingly unattractive as time runs out. The probability of an agent α being infected at any time shows an interesting behavior: At $\alpha = 0$, the probability is quite high due to the great number of connections. As α decreases, so does the probability of




Figure 3: Experimental results for OMD on the Beach problem. (a) The exploitability ΔJ over iterations n of OMD; (b) The final distribution over positions on the beach at time $t = T - 1$ for each discretized α ; (c) The evolution of distributions over time.

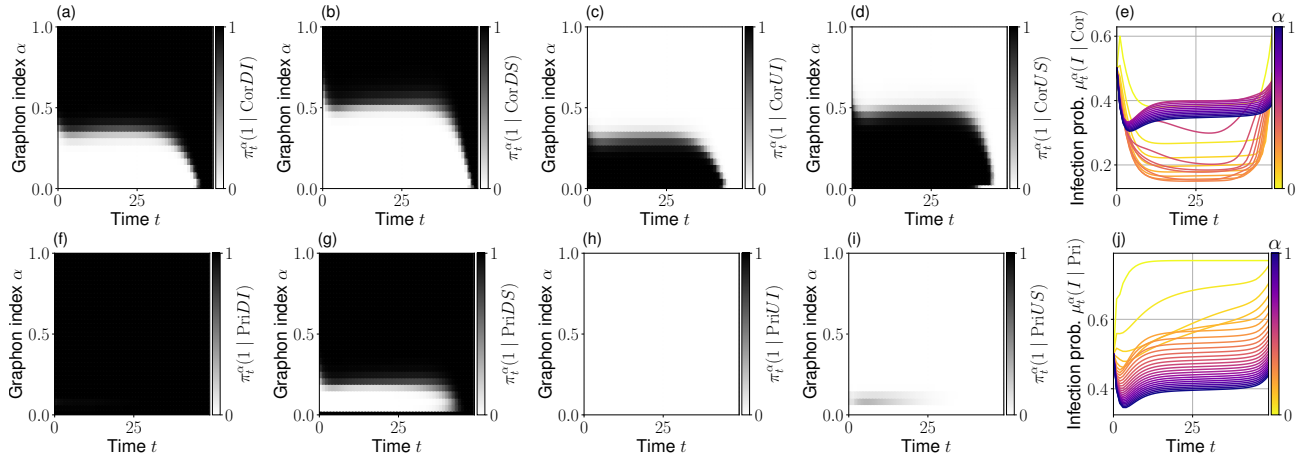


Figure 4: Experimental results for OMD on the heterogeneous Cyber Security problem. (a)-(d): The probability of action $u = 1$ at graphon index α and time t under the final equilibrium policy in states *CorDI*, *CorDS*, *CorUI*, *CorUS*; (e): The probability (MF) of infected *Cor* agents, visualized for each discretized α ; (f)-(j): Same as in (a)-(e) but for *Pri* agents.




Figure 5: The L_1 error between the empirical distribution and the limiting MF $\Delta\mu = \mathbb{E} \left[\sum_{t \in \mathcal{T}, x \in \mathcal{X}} \left| \frac{1}{N} \sum_i \delta_{X_t^i}(x) - \int_{\mathcal{I}} \mu_t^\alpha(x) d\alpha \right| \right]$ at $\beta = 0.51$ averaged over 100 randomly sampled graphs with N nodes and 68% confidence interval (shaded).

being infected at all times. However, as soon as α passes a threshold where defense is no longer individually worth it, the fraction of infected nodes jumps up.

In the heterogeneous case, as seen in Figure 4, we consider an additional class of nodes with partially similar behavior. For very high connectivity $\alpha \rightarrow 0$ however, we observe that *Pri* nodes will never defend themselves, since for the given problem parameters, the probability of infection will be very high regardless of the defense state. Otherwise, we can observe similar behavior as in the homogeneous case. Perhaps most interesting is the asymmetry between choosing to switch between defended and undefended. When agents are susceptible, some agents will opt to neither switch from defended to undefended, nor vice versa. This stems from the fact that agents switching in state *US* could likely jump to *UI* and *DI*, while in state *DS* likely jumps are *DS* and *US*, each of which may have different future returns. For the Beach Bar process in Figure 3, we see results similar to the classical ones in Perrin et al. (2020). By giving each agent




Figure 6: The L_1 error between the empirical distribution and the limiting MF as in Figure 5 over 50 uniformly spaced $\beta \in (0, 1)$ and $N \leq 100$ for the Cyber Security problem.

an incentive to avoid only their direct graphical neighbors, we obtain an equilibrium behavior where agents with many connections will stay further away from the bar, while agents with few connections will not mind many other agents.

Finally, in Figures 5 and 6 for $\rho_N = N^{-\beta}$ and sparsity parameter $\beta \in (0, 1)$, we observe convergence of the N -agent system objective to the MF objective, implying that sufficiently large finite systems are well-approximated by the LPGMFG. In Figure 6, for β close to 0 or 1, convergence slows down, as by (Borgs et al., 2019, Theorem 2.14) convergence is only guaranteed for $0 < \beta < 1$. Even though for $\beta = 0$, one would get the same model as in Cui and Koepl (2022), since the power law graphon is not $[0, 1]$ -valued, approximation guarantees fail for $\beta = 0$ and we observe increasingly slow convergence as we approach zero.

6 CONCLUSION

In this paper we have introduced LPGMFGs which enable the scalable, mathematically sound analysis of otherwise intractable MARL problems on large sparse graphs. We rigorously derived existence and convergence guarantees for LPGMFGs and provided learning schemes to find equilibria algorithmically where we adapted the OMD learning algorithm to the setting of LPGMFGs. Beyond that, we demonstrated the benefits of our approach empirically on different examples and showed that the practical results match the theory. As for the societal impact we foresee from our work, we believe that while our techniques remain very general, they could in the future lead to an analysis of self-interested agents on real graphs such as from social networks. This could find application e.g. in control strategies for future pandemics, or other interventions. Future work could extend our model in numerous ways such as considering continuous time, action, and state spaces or adding noise terms. A challenging task could also be to find similar learning concepts for ultra-sparse graphs where the degrees remain constant as the number of agents becomes large. For applications, it would be interesting to use LPGMFGs to solve real-world problems that occur in various research

fields. In general, we hope that our work contributes to the MARL literature and inspires future work on scalable learning methods on sparse graphs.

Acknowledgements

This work has been co-funded by the Hessian Ministry of Science and the Arts (HMWK) within the projects "The Third Wave of Artificial Intelligence - 3AI" and hessian.AI, and the LOEWE initiative (Hesse, Germany) within the emergenCITY center.

References

- Anahtarci, B., Kariksiz, C. D., and Saldi, N. (2020). Q-learning in regularized mean-field games. *arXiv preprint arXiv:2003.12151*.
- Aparicio, S., Villazón-Terrazas, J., and Álvarez, G. (2015). A model for scale-free networks: application to twitter. *Entropy*, 17(8):5848–5867.
- Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American Economic Review*, 84(2):406–411.
- Aurell, A., Carmona, R., Dayanıklı, G., and Laurière, M. (2022). Finite state graphon games with applications to epidemics. *Dynamic Games and Applications*, pages 1–33.
- Aurell, A., Carmona, R., and Lauriere, M. (2021). Stochastic graphon games: II. the linear-quadratic case. *arXiv preprint arXiv:2105.12320*.
- Avena-Koenigsberger, A., Misic, B., and Sporns, O. (2018). Communication dynamics in complex brain networks. *Nature Reviews Neuroscience*, 19(1):17–33.
- Bakker, L., Hare, W., Khosravi, H., and Ramadanovic, B. (2010). A social network model of investment behaviour in the stock market. *Physica A: Statistical Mechanics and its Applications*, 389(6):1223–1229.
- Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.
- Barabási, A.-L., Albert, R., and Jeong, H. (1999). Mean-field theory for scale-free random networks. *Physica A: Statistical Mechanics and its Applications*, 272(1-2):173–187.
- Bayraktar, E., Chakraborty, S., and Wu, R. (2020). Graphon mean field systems. *arXiv preprint arXiv:2003.13180*.
- Bertsekas, D. (2019). *Reinforcement learning and optimal control*. Athena Scientific.
- Bian, Y.-t., Xu, L., and Li, J.-s. (2016). Evolving dynamics of trading behavior based on coordination game in complex networks. *Physica A: Statistical Mechanics and its Applications*, 449:281–290.

- Bollobás, B., Janson, S., and Riordan, O. (2007). The phase transition in inhomogeneous random graphs. *Random Structures & Algorithms*, 31(1):3–122.
- Borgs, C., Chayes, J., Cohn, H., and Zhao, Y. (2018). An Lp theory of sparse graph convergence II: Ld convergence, quotients and right convergence. *The Annals of Probability*, 46(1):337–396.
- Borgs, C., Chayes, J., Cohn, H., and Zhao, Y. (2019). An Lp theory of sparse graph convergence I: Limits, sparse random graph models, and power law distributions. *Transactions of the American Mathematical Society*, 372(5):3019–3062.
- Borgs, C., Chayes, J., Lovász, L., Sós, V., and Vesztegombi, K. (2011). Limits of randomly grown graph sequences. *European Journal of Combinatorics*, 32(7):985–999.
- Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198.
- Bullmore, E. and Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336–349.
- Caines, P. E. and Huang, M. (2019). Graphon mean field games and the GMFG equations: ε -Nash equilibria. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 286–292. IEEE.
- Canese, L., Cardarilli, G. C., Di Nunzio, L., Fazzolari, R., Giardino, D., Re, M., and Spanò, S. (2021). Multi-agent reinforcement learning: A review of challenges and applications. *Applied Sciences*, 11(11):4948.
- Cardaliaguet, P. and Hadikhannoo, S. (2017). Learning in mean field games: the fictitious play. *ESAIM: Control, Optimisation and Calculus of Variations*, 23(2):569–591.
- Carmona, G. (2004). Nash equilibria of games with a continuum of players.
- Carmona, R., Cooney, D. B., Graves, C. V., and Lauriere, M. (2022). Stochastic graphon games: I. the static case. *Mathematics of Operations Research*, 47(1):750–778.
- Carmona, R. and Delarue, F. (2018a). *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games*, volume 83. Springer.
- Carmona, R. and Delarue, F. (2018b). *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*. Probability Theory and Stochastic Modelling. Springer International Publishing.
- Carmona, R., Fouque, J.-P., and Sun, L.-H. (2015). Mean field games and systemic risk. *Communications in Mathematical Sciences*, 13(4):911–933.
- Carmona, R., Laurière, M., and Tan, Z. (2019). Model-free mean-field reinforcement learning: mean-field mdp and mean-field q-learning. *arXiv preprint arXiv:1910.12802*.
- Carmona, R. and Zhu, X. (2016). A probabilistic approach to mean field games with major and minor players. *The Annals of Applied Probability*, 26(3):1535–1580.
- Castiglione, A., Gribaudo, M., Iacono, M., and Palmieri, F. (2014). Exploiting mean field analysis to model performances of big data architectures. *Future Generation Computer Systems*, 37:203–211.
- Choutri, S. E. and Djehiche, B. (2019). Mean-field risk sensitive control and zero-sum games for markov chains. *Bulletin des Sciences Mathématiques*, 152:1–39.
- Choutri, S. E., Djehiche, B., and Tembine, H. (2019). Optimal control and zero-sum games for markov chains of mean-field type. *Mathematical Control & Related Fields*, 9(3):571.
- Concas, G., Marchesi, M., Pinna, S., and Serra, N. (2007). Power-laws in a large object-oriented software system. *IEEE Transactions on Software Engineering*, 33(10):687–708.
- Conitzer, V. and Sandholm, T. (2008). New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641.
- Cui, K. and Koepl, H. (2021). Approximately solving mean field games via entropy-regularized deep reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1909–1917. PMLR.
- Cui, K. and Koepl, H. (2022). Learning graphon mean field games and approximate Nash equilibria. In *International Conference on Learning Representations*.
- Cui, K., Li, M., Fabian, C., and Koepl, H. (2022). Scalable task-driven robotic swarm control via collision avoidance and learning mean-field control. *arXiv preprint arXiv:2209.07420*.
- Delarue, F. and Vasileiadis, A. (2021). Exploration noise for learning linear-quadratic mean field games. *arXiv preprint arXiv:2107.00839*.
- D’Arcangelis, A. M. and Rotundo, G. (2016). Complex networks in finance. In *Complex Networks and Dynamics*, pages 209–235. Springer.
- Elamvazhuthi, K. and Berman, S. (2019). Mean-field models in swarm robotics: A survey. *Bioinspiration & Biomimetics*, 15(1):015001.
- Elie, R., Ichiba, T., and Laurière, M. (2020a). Large banking systems with default and recovery: A mean field game model. *arXiv preprint arXiv:2001.10206*.
- Elie, R., Perolat, J., Laurière, M., Geist, M., and Pietquin, O. (2020b). On the convergence of model free learning in mean field games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7143–7150.
- Farago, J., Greenwald, A., and Hall, K. (2002). Fair and efficient solutions to the santa fe bar problem. In *Grace Hopper Celebration of Women in Computing*. Citeseer.

- Firoozi, D., Jaimungal, S., and Caines, P. E. (2020). Convex analysis for lqg systems with applications to major–minor lqg mean–field game systems. *Systems & Control Letters*, 142:104734.
- Gao, S. and Caines, P. E. (2017). The control of arbitrary size networks of linear systems via graphon limits: An initial investigation. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1052–1057. IEEE.
- Gao, S., Tchuendom, R. F., and Caines, P. E. (2021). Linear quadratic graphon field games. *Communications in Information and Systems*, 21(3):341–369.
- Gkogkas, M. A. and Kuehn, C. (2022). Graphop mean-field limits for kuramoto-type models. *SIAM Journal on Applied Dynamical Systems*, 21(1):248–283.
- Guo, X., Hu, A., Xu, R., and Zhang, J. (2019). Learning mean-field games. In *Advances in Neural Information Processing Systems*, pages 4966–4976.
- Guo, X., Hu, A., Xu, R., and Zhang, J. (2020). A general framework for learning mean-field games. *arXiv preprint arXiv:2003.06069*.
- Guo, X., Xu, R., and Zariphopoulou, T. (2022). Entropy regularization for mean field games with learning. *Mathematics of Operations Research*.
- Hadikhanloo, S. (2017). Learning in anonymous nonatomic games with applications to first-order mean field games. *arXiv preprint arXiv:1704.00378*.
- Hadikhanloo, S. and Silva, F. J. (2019). Finite mean field games: fictitious play and convergence to a first order continuous mean field game. *Journal de Mathématiques Pures et Appliquées*, 132:369–397.
- Huang, K., Di, X., Du, Q., and Chen, X. (2020). A game-theoretic framework for autonomous vehicles velocity control: Bridging microscopic differential games and macroscopic mean field games. *Discrete & Continuous Dynamical Systems-B*, 25(12):4869.
- Huang, M., Malhamé, R. P., and Caines, P. E. (2006). Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3):221–252.
- Khalil, H. (2002). *Nonlinear Systems*. Pearson Education. Prentice Hall.
- Kolokoltsov, V. N. and Bensoussan, A. (2016). Mean-field-game model for botnet defense in cyber-security. *Applied Mathematics & Optimization*, 74(3):669–692.
- Kolokoltsov, V. N. and Malafeyev, O. A. (2018). Corruption and botnet defense: a mean field game approach. *International Journal of Game Theory*, 47(3):977–999.
- Lacker, D. and Soret, A. (2022). A case study on stochastic games on large graphs in mean field and sparse regimes. *Mathematics of Operations Research*, 47(2):1530–1565.
- Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260.
- Laurière, M., Perrin, S., Girgin, S., Muller, P., Jain, A., Cabannes, T., Piliouras, G., Pérolat, J., Élie, R., Pietquin, O., et al. (2022). Scalable deep reinforcement learning algorithms for mean field games. *arXiv preprint arXiv:2203.11973*.
- Li, M., Qin, Z., Jiao, Y., Yang, Y., Wang, J., Wang, C., Wu, G., and Ye, J. (2019). Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning. In *The world wide web conference*, pages 983–994.
- Louridas, P., Spinellis, D., and Vlachos, V. (2008). Power laws in software. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 18(1):1–26.
- Lovász, L. (2012). *Large networks and graph limits*, volume 60. American Mathematical Soc.
- Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. (2018). Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2703–2717. SIAM.
- Mguni, D., Jennings, J., and de Cote, E. M. (2018). Decentralised learning in systems with many, many strategic agents. *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Mondal, W. U., Agarwal, M., Aggarwal, V., and Ukkusuri, S. V. (2021). On the approximation of cooperative heterogeneous multi-agent reinforcement learning (MARL) using mean field control (MFC). *arXiv preprint arXiv:2109.04024*.
- Newman, M. (2005). Power laws, pareto distributions and zipf’s law. *Contemporary Physics*, 46(5):323–351.
- Newman, M. (2018). *Networks*. Oxford University Press.
- Newman, M. E. (2002). Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128.
- Nosonovsky, M. and Roy, P. (2020). Scaling in colloidal and biological networks. *Entropy*, 22(6):622.
- Nourian, M. and Caines, P. E. (2013). ϵ -Nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents. *SIAM Journal on Control and Optimization*, 51(4):3302–3331.
- Orabona, F., Crammer, K., and Cesa-Bianchi, N. (2015). A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435.
- Papadimitriou, C. (2001). Algorithms, games, and the internet. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, pages 749–753.
- Papadimitriou, C. H. (2007). The complexity of finding Nash equilibria. *Algorithmic Game Theory*, 2:30.

- Pastor-Satorras, R., Castellano, C., Van Mieghem, P., and Vespignani, A. (2015). Epidemic processes in complex networks. *Reviews of Modern Physics*, 87(3):925.
- Perolat, J., Perrin, S., Elie, R., Laurière, M., Piliouras, G., Geist, M., Tuyls, K., and Pietquin, O. (2021). Scaling up mean field games with online mirror descent. *arXiv preprint arXiv:2103.00623*.
- Perrin, S., Laurière, M., Pérolat, J., Élie, R., Geist, M., and Pietquin, O. (2021a). Generalization in mean field games by learning master policies. *arXiv preprint arXiv:2109.09717*.
- Perrin, S., Laurière, M., Pérolat, J., Geist, M., Élie, R., and Pietquin, O. (2021b). Mean field games flock! the reinforcement learning way. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 356–362. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Perrin, S., Pérolat, J., Laurière, M., Geist, M., Elie, R., and Pietquin, O. (2020). Fictitious play for mean field games: Continuous time analysis and applications. *Advances in Neural Information Processing Systems*, 33:13199–13213.
- Qu, G., Wierman, A., and Li, N. (2020). Scalable reinforcement learning of localized policies for multi-agent networked systems. In *Learning for Dynamics and Control*, pages 256–266. PMLR.
- Rozemberczki, B., Davies, R., Sarkar, R., and Sutton, C. (2019). GEMSEC: Graph embedding with self clustering. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 65–72. IEEE Computer Society.
- Saldi, N., Basar, T., and Raginsky, M. (2018). Markov–Nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization*, 56(6):4256–4287.
- Sporns, O. (2022). Structure and function of complex brain networks. *Dialogues in Clinical Neuroscience*.
- Srebro, N., Sridharan, K., and Tewari, A. (2011). On the universality of online mirror descent. *Advances in Neural Information Processing Systems*, 24.
- Subramanian, J. and Mahajan, A. (2019). Reinforcement learning in stationary mean-field games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 251–259.
- Tangpi, L. and Zhou, X. (2022). Optimal investment in a large population of competitive and heterogeneous agents. *arXiv preprint arXiv:2202.11314*.
- Wheeldon, R. and Counsell, S. (2003). Power law distributions in class relationships. In *Proceedings Third IEEE International Workshop on Source Code Analysis and Manipulation*, pages 45–54. IEEE.
- Yang, J., Ye, X., Trivedi, R., Xu, H., and Zha, H. (2018). Learning deep mean field games for modeling large population behavior. In *International Conference on Learning Representations*.
- Yang, Y. and Wang, J. (2020). An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583*.
- Zhang, K., Yang, Z., and Başar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of Reinforcement Learning and Control*, pages 321–384.

SUPPLEMENTARY MATERIAL

The following pages provide additional information and proofs for the statements in the main part of the paper "Learning Sparse Graphon Mean Field Games". In the sections A to D we give the proofs of Theorems 2 and 3. Subsequently, sections E to H contain the proofs of Theorem 4 and the necessary intermediate results. Furthermore, the sections I to M include the details of the learning methods as well as the proofs for the theoretical convergence guarantee (Theorem 5) of the OMD algorithm. We provide a derivation of the cutoff power law graphon in section N and conclude the supplementary materials with the problem details for our examples in section O.

A OVERVIEW

In this section we collect intermediate results and auxiliary definitions that help to prove Theorems 2 and 3. Proofs which require an extensive argumentation are deferred to the subsequent paragraphs. First, we define the α -neighborhood map $\mathbb{G}^\alpha: \mathcal{M}_t \rightarrow \mathcal{P}(\mathcal{X})$ of an agent $\alpha \in \mathcal{I}$ by

$$\mathbb{G}^\alpha(\boldsymbol{\mu}_t) := \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta d\beta$$

and similarly the empirical α -neighborhood map $\mathbb{G}_N^\alpha: \mathcal{M}_t \rightarrow \mathcal{P}(\mathcal{X})$ by

$$\mathbb{G}_N^\alpha(\boldsymbol{\mu}_t) := \int_{\mathcal{I}} \frac{W_N(\alpha, \beta)}{\rho_N} \mu_t^\beta d\beta.$$

Here, we point out that $\mathbb{G}_t^\alpha = \mathbb{G}^\alpha(\boldsymbol{\mu}_t)$ holds for the MF system and $\mathbb{G}_t^i = \mathbb{G}_N^i(\boldsymbol{\mu}_t^N)$ holds for the finite model. For the following proofs, we require the concept of an ensemble transition kernel operator $P_{t,\nu,W}^\pi$ which is defined by

$$(\nu P_{t,\nu,W}^\pi)^\alpha = \sum_{x \in \mathcal{X}} \nu^\alpha(x) \sum_{u \in \mathcal{U}} \pi_t^\alpha(u|x) \cdot P\left(\cdot \mid x, u, \int_{\mathcal{I}} W(\alpha, \beta) \nu'^\beta d\beta\right)$$

where W is a graphon, $\nu, \nu' \in \mathcal{M}_t$, and $\pi \in \Pi$. This definition especially implies the useful property $\boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t P_{t,\mu_t,W}^\pi$. For convenience, we also introduce the notation

$$(W\mu)(\alpha, x) := \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta(x) d\beta.$$

To prove Theorem 3, we have to establish the following result which is proven in a subsequent section. Note that this Lemma 2 is not required in the standard graphon theory but becomes necessary for our L^p graphon approach. Specifically, the proof of Lemma 3 will include functions of the form $f'_{N,i,x}(x', \beta) = \frac{1}{\rho_N} W_N(\frac{i}{N}, \beta) \cdot \mathbf{1}_{\{x=x'\}}$ for which a result like the one given by Theorem 2 is required. However, the factor $\frac{1}{\rho_N}$ introduced as part of the L^p graphon approach makes it impossible to bound functions such as $f'_{N,i,x}(x', \beta)$ uniformly over all N because for sparse graphs $\frac{1}{\rho_N}$ approaches infinity as N approaches infinity. Therefore, Theorem 2 is not applicable in this case which makes a new statement, i.e. Lemma 2, necessary. Put differently, the formulation and proof of Lemma 2 is a new mathematical contribution of our work.

Lemma 2. *Let $(c_N)_{N \geq 1}$ be a sequence of positive real numbers such that $c_N = o(1)$ and $(c_N)^{-1} = o(N)$ and let $B_N \subseteq \mathcal{V}_N$ be some sequence of sets with $|B_N| = O(N \cdot c_N)$. Then, we consider a sequence of measurable functions $f_N: \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ where for all $x \in \mathcal{X}$ we have $f_N(x, \alpha) = O(c_N^{-1})$ if $\alpha \in A_N := \bigcup_{i \in B_N} (\frac{i-1}{N}, \frac{i}{N}]$ and $f_N(x, \alpha) = 0$ otherwise. Consider Lipschitz continuous $\pi \in \Pi$ up to a finite number of discontinuities D_π , with associated mean field ensemble $\boldsymbol{\mu} = \Psi(\pi)$. Under Assumption 1 and the N -agent policy $(\pi^1, \dots, \pi^{i-1}, \hat{\pi}, \pi^{i+1}, \dots, \pi^N) \in \Pi^N$ with $(\pi^1, \pi^2, \dots, \pi^N) = \Gamma_N(\pi) \in \Pi^N$, $\hat{\pi} \in \Pi$, $t \in \mathcal{T}$, we have*

$$\mathbb{E} \left[\left| \boldsymbol{\mu}_t^N(f_N) - \boldsymbol{\mu}_t(f_N) \right| \right] \rightarrow 0 \tag{14}$$

uniformly over all possible deviations $\hat{\pi} \in \Pi$, $i \in \mathcal{V}_N$.

With Lemma 2 in place, we are able to provide another intermediate result which is a key element for the proof of Theorem 3.

Lemma 3. *Consider Lipschitz continuous $\pi \in \Pi$ up to a finite number of discontinuities D_π , with associated mean field ensemble $\boldsymbol{\mu} = \Psi(\pi)$. Under Assumptions 1 and 2 and the N -agent policy $(\pi^1, \dots, \pi^{i-1}, \hat{\pi}, \pi^{i+1}, \dots, \pi^N) \in \Pi^N$ where*

$(\pi^1, \pi^2, \dots, \pi^N) = \Gamma_N(\boldsymbol{\pi}) \in \Pi^N$, $\hat{\pi} \in \Pi$ arbitrary, for any uniformly bounded family of functions \mathcal{G} from \mathcal{X} to \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

$$\sup_{g \in \mathcal{G}} \left| \mathbb{E} [g(X_t^i)] - \mathbb{E} [g(\hat{X}_t^{\frac{i}{N}})] \right| < \varepsilon \quad (15)$$

uniformly over $\hat{\pi} \in \Pi, i \in \mathcal{W}_N$ for some $\mathcal{W}_N \subseteq V_N$ with $|\mathcal{W}_N| \geq \lfloor (1-p)N \rfloor$.

Similarly, for any uniformly Lipschitz, uniformly bounded family of measurable functions \mathcal{H} from $\mathcal{X} \times \mathcal{B}(\mathcal{X})$ to \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

$$\sup_{h \in \mathcal{H}} \left| \mathbb{E} \left[h(X_t^i, \mathbb{G}_N^{\frac{i}{N}}(\boldsymbol{\mu}_t^N)) \right] - \mathbb{E} \left[h(\hat{X}_t^{\frac{i}{N}}, \mathbb{G}_N^{\frac{i}{N}}(\boldsymbol{\mu}_t)) \right] \right| < \varepsilon \quad (16)$$

uniformly over $\hat{\pi} \in \Pi, i \in \mathcal{W}_N$ for some $\mathcal{W}_N \subseteq V_N$ with $|\mathcal{W}_N| \geq \lfloor (1-p)N \rfloor$.

Finally, keeping in mind the above statements, Theorem 3 can be proven.

Proof of Theorem 3. Leveraging Lemma 3, we can establish Theorem 3 with an argumentation as in (Cui and Koepl, 2022, proofs of Corollary A.1 and Theorem 3). \square

B PROOF OF THEOREM 2

Proof. We provide a proof by induction over t which is structurally similar to an argument in Cui and Koepl (2022). The case $t = 0$ follows from a law of large numbers argument. For the induction step we consider the inequality

$$\begin{aligned} \mathbb{E} \left[\left| \boldsymbol{\mu}_{t+1}^N(f) - \boldsymbol{\mu}_{t+1}(f) \right| \right] &\leq \mathbb{E} \left[\left| \boldsymbol{\mu}_{t+1}^N(f) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, \frac{W_N}{\rho_N}}^{\boldsymbol{\pi}^N}(f) \right| \right] \\ &\quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, \frac{W_N}{\rho_N}}^{\boldsymbol{\pi}^N}(f) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\boldsymbol{\pi}^N}(f) \right| \right] \\ &\quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\boldsymbol{\pi}^N}(f) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\boldsymbol{\pi}}(f) \right| \right] \\ &\quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\boldsymbol{\pi}}(f) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t, W}^{\boldsymbol{\pi}}(f) \right| \right] \\ &\quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t, W}^{\boldsymbol{\pi}}(f) - \boldsymbol{\mu}_{t+1}(f) \right| \right]. \end{aligned}$$

where, as before, $f: \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ is an arbitrary bounded function such that $|f| \leq M_f$. While the last three terms can be bounded as in Cui and Koepl (2022), the first two summands now include the factor ρ_N . Thus, the first term is bounded by

$$\begin{aligned} &\mathbb{E} \left[\left| \boldsymbol{\mu}_{t+1}^N(f) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, \frac{W_N}{\rho_N}}^{\boldsymbol{\pi}^N}(f) \right| \right] \\ &= \mathbb{E} \left[\left| \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} \mu_{t+1}^{N, \alpha}(x) f(x, \alpha) d\alpha - \int_{\mathcal{I}} \sum_{\substack{x, x' \in \mathcal{X} \\ u \in \mathcal{U}}} \mu_t^{N, \alpha}(x) \pi_t^{N, \alpha}(u | x) \cdot P \left(x' | x, u, \left(\frac{W_N}{\rho_N} \boldsymbol{\mu}_t^N \right) (\alpha) \right) f(x', \alpha) d\alpha \right| \right] \\ &= \mathbb{E} \left[\left| \sum_{i \in V_N} \left(\int_{(\frac{i-1}{N}, \frac{i}{N}] } f(X_{t+1}^i, \alpha) d\alpha - \mathbb{E} \left[\int_{(\frac{i-1}{N}, \frac{i}{N}] } f(X_{t+1}^i, \alpha) d\alpha \mid \mathbf{X}_t \right] \right) \right| \right] \\ &\leq \left(\mathbb{E} \left[\left(\sum_{i \in V_N} \left(\int_{(\frac{i-1}{N}, \frac{i}{N}] } f(X_{t+1}^i, \alpha) d\alpha - \mathbb{E} \left[\int_{(\frac{i-1}{N}, \frac{i}{N}] } f(X_{t+1}^i, \alpha) d\alpha \mid \mathbf{X}_t \right] \right) \right)^2 \right] \right)^{\frac{1}{2}} \\ &= \left(\sum_{i \in V_N} \mathbb{E} \left[\left(\int_{(\frac{i-1}{N}, \frac{i}{N}] } f(X_{t+1}^i, \alpha) d\alpha - \mathbb{E} \left[\int_{(\frac{i-1}{N}, \frac{i}{N}] } f(X_{t+1}^i, \alpha) d\alpha \mid \mathbf{X}_t \right] \right)^2 \right] \right)^{\frac{1}{2}} \leq \frac{2M_f}{\sqrt{N}} \end{aligned}$$

where we point out that the $\{X_{t+1}^i\}_{i \in V_N}$ are independent if conditioned on $\mathbf{X}_t \equiv \{X_t^i\}_{i \in V_N}$. Turning to the second summand, under Assumption 1 we can derive

$$\begin{aligned} & \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, \frac{W_N}{\rho_N}}^{\pi^N}(f) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\pi^N}(f) \right| \right] \\ & \leq |\mathcal{X}| M_f L_P \mathbb{E} \left[\int_{\mathcal{I}} \left\| \int_{\mathcal{I}} \frac{W_N(\alpha, \beta)}{\rho_N} \mu_t^{N, \beta} d\beta - \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^{N, \beta} d\beta \right\| d\alpha \right] \\ & \leq |\mathcal{X}|^2 M_f L_P \sup_{x \in \mathcal{X}} \mathbb{E} \left[\int_{\mathcal{I}} \left| \int_{\mathcal{I}} \frac{W_N(\alpha, \beta)}{\rho_N} \mu_t^{N, \beta}(x) - W(\alpha, \beta) \mu_t^{N, \beta}(x) d\beta \right| d\alpha \right] \rightarrow 0 \end{aligned}$$

where $\mu_t^{N, \beta}(x)$ as a probability is less than or equal to one. The convergence in the last line is at rate $O(1/\sqrt{N})$ if the convergence rate in Assumption 1 is also $O(1/\sqrt{N})$. \square

C PROOF OF LEMMA 2

Proof. The proof is by induction as follows. Note that it is similar to the proof of Theorem 2.

Base case. Starting with $t = 0$, we derive

$$\begin{aligned} & \mathbb{E} [|\boldsymbol{\mu}_0^N(f_N) - \boldsymbol{\mu}_0(f_N)|] \\ & = \mathbb{E} \left[\left| \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} (\mu_0^{N, \alpha}(x) - \mu_0^\alpha(x)) f_N(x, \alpha) d\alpha \right| \right] \\ & = \mathbb{E} \left[\left| \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} \left(\sum_{i \in V_N} \mathbf{1}_{\alpha \in (\frac{i-1}{N}, \frac{i}{N})} \delta_{X_0^i} \right) f_N(x, \alpha) - \sum_{x \in \mathcal{X}} \mu_0^\alpha(x) f_N(x, \alpha) d\alpha \right| \right] \\ & \stackrel{(Fubini)}{=} \mathbb{E} \left[\left| \sum_{i \in V_N} \int_{(\frac{i-1}{N}, \frac{i}{N})} f_N(X_0^i, \alpha) d\alpha - \mathbb{E} \left[\int_{(\frac{i-1}{N}, \frac{i}{N})} f_N(X_0^i, \alpha) d\alpha \right] \right| \right] \\ & \leq \left(\mathbb{E} \left[\left(\sum_{i \in B_N} \left(\int_{(\frac{i-1}{N}, \frac{i}{N})} f_N(X_0^i, \alpha) d\alpha - \mathbb{E} \left[\int_{(\frac{i-1}{N}, \frac{i}{N})} f_N(X_0^i, \alpha) d\alpha \right] \right)^2 \right) \right] \right)^{\frac{1}{2}} \\ & = \left(\sum_{i \in B_N} \mathbb{E} \left[\left(\int_{(\frac{i-1}{N}, \frac{i}{N})} f_N(X_0^i, \alpha) d\alpha - \mathbb{E} \left[\int_{(\frac{i-1}{N}, \frac{i}{N})} f_N(X_0^i, \alpha) d\alpha \right] \right)^2 \right] \right)^{\frac{1}{2}} \\ & = \left(\sum_{i \in B_N} \mathbb{E} \left[\left(O(c_N^{-1}) \cdot \frac{1}{N} \right)^2 \right] \right)^{\frac{1}{2}} = \left(O(N \cdot c_N) \cdot \left(O(c_N^{-2}) \cdot \frac{1}{N^2} \right) \right)^{\frac{1}{2}} = \frac{O(c_N^{-\frac{1}{2}})}{\sqrt{N}} = o(1) \end{aligned}$$

by exploiting the fact that the $\{X_0^i\}_{i \in V_N}$ are independent and $X_0^i \sim \mu_0 = \mu_0^\alpha$ holds $\forall i \in V_N, \alpha \in \mathcal{I}$.

Induction step. For performing the induction step, we start with the following reformulation

$$\begin{aligned} \mathbb{E} [|\boldsymbol{\mu}_{t+1}^N(f_N) - \boldsymbol{\mu}_{t+1}(f_N)|] & \leq \mathbb{E} \left[\left| \boldsymbol{\mu}_{t+1}^N(f_N) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, \frac{W_N}{\rho_N}}^{\pi^N}(f_N) \right| \right] \\ & \quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, \frac{W_N}{\rho_N}}^{\pi^N}(f_N) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\pi^N}(f_N) \right| \right] \\ & \quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t^N, W}^{\pi^N}(f_N) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t, W}^{\pi^N}(f_N) \right| \right] \\ & \quad + \mathbb{E} \left[\left| \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t, W}^{\pi^N}(f_N) - \boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t, W}^{\pi^N}(f_N) \right| \right] \\ & \quad + \mathbb{E} [|\boldsymbol{\mu}_t^N P_{t, \boldsymbol{\mu}_t, W}^{\pi^N}(f_N) - \boldsymbol{\mu}_{t+1}(f_N)|] \end{aligned}$$

and assume as usual that the induction assumption is fulfilled for t . Apart from the second summand, all terms converge to zero by arguments as in Theorem 2. For the second term, however, we require a different argument.

Before we consider the second term itself, we make the following useful observation. By Assumption 1, we know that

$$\left\| \frac{W_N}{\rho_N} - W \right\|_{L_\infty \rightarrow L_1} = \sup_{\|g\|_\infty \leq 1} \int_{\mathcal{I}} \left| \int_{\mathcal{I}} \left(\frac{W_N(\alpha, \beta)}{\rho_N} - W(\alpha, \beta) \right) g(\beta) d\beta \right| d\alpha$$

converges to zero which especially implies that for almost all $\alpha \in \mathcal{I}$ we have

$$\left| \int_{\mathcal{I}} \left(\frac{W_N(\alpha, \beta)}{\rho_N} - W(\alpha, \beta) \right) \mu_t^{N, \beta}(x) d\beta \right| = o(1)$$

for all $x \in \mathcal{X}$ since $\mu_t^{N, \beta}(x)$ is trivially bounded by 1. This in turn implies that for every positive real number $\epsilon > 0$ there exists a $N' \in \mathbb{N}$ such that for all $N > N'$ and $x \in \mathcal{X}$

$$\left| \int_{\mathcal{I}} \left(\frac{W_N(\alpha, \beta)}{\rho_N} - W(\alpha, \beta) \right) \mu_t^{N, \beta}(x) d\beta \right| < \epsilon \quad (17)$$

holds for almost all $\alpha \in \mathcal{I}$. For an arbitrary but fixed $\epsilon > 0$ and under the assumption that $N > N'$, we have

$$\begin{aligned} & \mathbb{E} \left[\left| \mu_t^N P_{t, \mu_t^N, \frac{W_N}{\rho_N}}^{\pi^N}(f_N) - \mu_t^N P_{t, \mu_t^N, W}^{\pi^N}(f_N) \right| \right] \\ &= \mathbb{E} \left[\left| \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} \mu_t^{N, \alpha}(x) \sum_{u \in \mathcal{U}} \pi_t^{N, \alpha}(u | x) \sum_{x' \in \mathcal{X}} f_N(x', \alpha) \left(P \left(x' | x, u, \left(\frac{W_N}{\rho_N} \mu_t^N \right) \right) - P \left(x' | x, u, (W \mu_t^N) \right) \right) d\alpha \right| \right] \\ &= \mathbb{E} \left[\left| \int_{A_N} \sum_{x \in \mathcal{X}} \mu_t^{N, \alpha}(x) \sum_{u \in \mathcal{U}} \pi_t^{N, \alpha}(u | x) \sum_{x' \in \mathcal{X}} f_N(x', \alpha) \left(P \left(x' | x, u, \left(\frac{W_N}{\rho_N} \mu_t^N \right) \right) - P \left(x' | x, u, (W \mu_t^N) \right) \right) d\alpha \right| \right] \\ &\leq O(c_N^{-1}) \cdot \sup_{x \in \mathcal{X}} \mathbb{E} \left[\int_{A_N} \left| \left(\frac{W_N}{\rho_N} \mu_t^N \right) (\alpha, x) - (W \mu_t^N) (\alpha, x) \right| d\alpha \right] \\ &\stackrel{\text{ineq. (17)}}{\leq} O(c_N^{-1}) \cdot \sup_{x \in \mathcal{X}} \mathbb{E} \left[\int_{A_N} \epsilon d\alpha \right] = O(1) \cdot \epsilon. \end{aligned}$$

Choosing ϵ arbitrarily close to zero yields the desired result, i.e.

$$\mathbb{E} \left[\left| \mu_t^N P_{t, \mu_t^N, \frac{W_N}{\rho_N}}^{\pi^N}(f_N) - \mu_t^N P_{t, \mu_t^N, W}^{\pi^N}(f_N) \right| \right] \rightarrow 0.$$

□

D PROOF OF LEMMA 3

Proof. Parts of this proof are built on arguments in (Cui and Koepl, 2022, Proof of Lemma A.1). Therefore, we focus on the differences and indicate parts that carry over from Cui and Koepl (2022) to our case. For establishing (15) we refer to Cui and Koepl (2022). Thus, it remains to show that (15) implies (16). As a starting point, let \mathcal{H} be a uniformly bounded (by M_h), uniformly Lipschitz (with Lipschitz constant L_h) family of measurable functions $h: \mathcal{X} \times \mathcal{B}(\mathcal{X}) \rightarrow \mathbb{R}$. The following reformulation holds for all $h \in \mathcal{H}$

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

Focusing on the first term, we obtain

$$\begin{aligned} \left| \mathbb{E} \left[h(X_t^i, \mathbb{G}_N^{\frac{i}{N}}(\mu_t^N)) \right] - \mathbb{E} \left[h(X_t^i, \mathbb{G}_N^{\frac{i}{N}}(\mu_t)) \right] \right| &\leq \mathbb{E} \left[\left| \mathbb{E} \left[h(X_t^i, \mathbb{G}_N^{\frac{i}{N}}(\mu_t^N)) h(X_t^i, \mathbb{G}_N^{\frac{i}{N}}(\mu_t)) \mid X_t^i \right] \right| \right] \\ &\leq L_h \mathbb{E} \left[\left\| \mathbb{G}_N^{\frac{i}{N}}(\mu_t^N) - \mathbb{G}_N^{\frac{i}{N}}(\mu_t) \right\| \right] \end{aligned}$$

$$= L_h \sum_{x' \in \mathcal{X}} \mathbb{E} \left[\left| \int_{\mathcal{I}} \frac{1}{\rho_N} W_N\left(\frac{i}{N}, \beta\right) (\mu_t^{N,\beta}(x) - \mu_t^\beta(x)) d\beta \right| \right]$$

which converges to zero. To see this, define the functions $f'_{N,i,x}(x', \beta) = \frac{1}{\rho_N} W_N\left(\frac{i}{N}, \beta\right) \cdot \mathbf{1}_{\{x=x'\}}$ and apply Lemma 2 to them. Combining our findings for the first term with results in Cui and Koepl (2022) for the second and third term concludes the proof. \square

E PROOF OF THEOREM 4

From now on, we use the notation \widehat{W} for the smoothed step graphon and W is the original step graphon. Ultimately, our goal is to establish Theorem 4. Therefore, we have to provide some intermediate results first for which the corresponding proofs can be found in the subsequent sections. Note that the structure and proofs of the intermediate results follow ideas in Cui and Koepl (2022) although the referenced paper does not consider smoothed step graphons.

Lemma 4. *Let $\pi \in \Pi$ be a policy ensemble generating $\hat{\mu} \in \mathcal{M}$ in the smoothed MP MFG and generating $\mu \in \mathcal{M}$ in the standard MP MFG under Assumption 2. Then, we have for all $t \in \mathcal{T}$ and all measurable functions $f : \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ uniformly bounded by some $M_f > 0$ that*

$$|\hat{\mu}_t(f) - \mu_t(f)| \rightarrow 0 \quad \text{as } \xi \rightarrow 0 \quad (18)$$

uniformly over all $\pi \in \Pi$.

This first lemma in turn allows us to provide the following result.

Lemma 5. *In the setup of Lemma 4, for any uniformly bounded family of functions \mathcal{G} from \mathcal{X} to \mathbb{R} and any $\varepsilon, p > 0, t \in \mathcal{T}$, there exists $\xi' > 0$ such that for all $\xi \in (0, \xi')$*

$$\sup_{g \in \mathcal{G}} \left| \mathbb{E}[g(X_t^\alpha)] - \mathbb{E}[g(\hat{X}_t^\alpha)] \right| < \varepsilon \quad (19)$$

holds for all $\alpha \in \mathcal{J} \subseteq \mathcal{I}$ where $\lambda(\mathcal{J}) \geq 1 - p$. Under the same conditions, we have that for any uniformly bounded, uniformly Lipschitz family of measurable functions \mathcal{H} from $\mathcal{X} \times \mathcal{B}_1(\mathcal{X})$ to \mathbb{R} and any $\varepsilon, p > 0, t \in \mathcal{T}$, there exists $\xi' > 0$ such that for all $\xi \in (0, \xi')$

$$\sup_{h \in \mathcal{H}} \left| \mathbb{E}[h(X_t^\alpha, \mathbb{G}_W^\alpha(\mu_t))] - \mathbb{E}[h(\hat{X}_t^\alpha, \mathbb{G}_W^\alpha(\hat{\mu}_t))] \right| < \varepsilon \quad (20)$$

holds for all $\alpha \in \mathcal{J} \subseteq \mathcal{I}$ where $\lambda(\mathcal{J}) \geq 1 - p$.

The next statement is a consequence of the just stated results.

Corollary 1. *In the setup of Lemma 4, for every $\varepsilon, p > 0$ there exists a $\xi' > 0$ such that for all $\xi \in (0, \xi')$*

$$\sup_{\pi \in \Pi} |J_{\alpha, W}^t(\pi) - J_{\alpha, \widehat{W}}^t(\pi)| < \varepsilon.$$

holds for all $\alpha \in \mathcal{J} \subseteq \mathcal{I}$ where $\lambda(\mathcal{J}) \geq 1 - p$.

With these statements in place, we can prove the theorem of interest.

Proof of Theorem 4. We consider the inequality

$$\begin{aligned} \sup_{\pi \in \Pi} J_{\alpha, W}(\pi_{\text{smo}}) - J_{\alpha, W}(\pi) &\leq \sup_{\pi \in \Pi} |J_{\alpha, W}(\pi_{\text{smo}}) - J_{\alpha, W_{\text{smo}}}(\pi_{\text{smo}})| \\ &\quad + \sup_{\pi \in \Pi} J_{\alpha, W_{\text{smo}}}(\pi_{\text{smo}}) - J_{\alpha, W_{\text{smo}}}(\pi) \\ &\quad + \sup_{\pi \in \Pi} |J_{\alpha, W_{\text{smo}}}(\pi) - J_{\alpha, W}(\pi)|. \end{aligned}$$

The second term is equal to zero because $(\pi_{\text{smo}}, \mu_{\text{smo}})$ is a GMFE for the smoothed game. The first and third term, on the other hand, can be bounded by $\varepsilon/2$ each by using Corollary 1 for all $\alpha \in \mathcal{J}$ for some $\mathcal{J} \subseteq \mathcal{I}$ with Lebesgue measure $\lambda(\mathcal{J}) \geq 1 - p$. Eventually, this yields

$$\sup_{\pi \in \Pi} |J_{\alpha, W}(\pi) - J_{\alpha, W}(\pi_{\text{smo}})| \leq \frac{\varepsilon}{2} + 0 + \frac{\varepsilon}{2} = \varepsilon$$

and thereby concludes the proof. \square

F PROOF OF LEMMA 4

We prove the statement via induction over t . The base case trivially holds. For the induction step, we consider the following reformulation

$$\begin{aligned} |\hat{\mu}_{t+1}(f) - \mu_{t+1}(f)| &\leq \left| \hat{\mu}_{t, \hat{\mu}_t, \widehat{W}}^{P^\pi}(f) - \hat{\mu}_{t, \hat{\mu}_t, W}^{P^\pi}(f) \right| \\ &\quad + \left| \hat{\mu}_{t, \hat{\mu}_t, W}^{P^\pi}(f) - \hat{\mu}_{t, \mu_t, W}^{P^\pi}(f) \right| \\ &\quad + \left| \hat{\mu}_{t, \mu_t, W}^{P^\pi}(f) - \mu_{t+1}(f) \right|. \end{aligned}$$

We consider the three summands separately.

First term. For the first term we note that

$$\begin{aligned} \left| \hat{\mu}_{t, \hat{\mu}_t, \widehat{W}}^{P^\pi}(f) - \hat{\mu}_{t, \hat{\mu}_t, W}^{P^\pi}(f) \right| &\leq |\mathcal{X}| M_f L_P \int_{\mathcal{I}} \left\| \int_{\mathcal{I}} \widehat{W}(\alpha, \beta) \hat{\mu}_t^\beta d\beta - \int_{\mathcal{I}} W(\alpha, \beta) \hat{\mu}_t^\beta d\beta \right\| d\alpha \\ &\leq |\mathcal{X}|^2 M_f L_P \sup_{x \in \mathcal{X}} \int_{\mathcal{I}} \left| \int_{\mathcal{I}} \widehat{W}(\alpha, \beta) \hat{\mu}_t^\beta(x) - W(\alpha, \beta) \hat{\mu}_t^\beta(x) d\beta \right| d\alpha \end{aligned}$$

converges to zero as $\xi \rightarrow 0$ by the construction of the smoothed step graphon \widehat{W} .

Second term. Consider the following reformulation

$$\begin{aligned} &\left| \hat{\mu}_{t, \hat{\mu}_t, W}^{P^\pi}(f) - \hat{\mu}_{t, \mu_t, W}^{P^\pi}(f) \right| \\ &\leq M_f |\mathcal{X}| \sup_{x, u} \int_{\mathcal{I}} \left| P \left(x' \mid x, u, \int_{\mathcal{I}} W(\alpha, \beta) \hat{\mu}_t^\beta d\beta \right) - P \left(x' \mid x, u, \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta d\beta \right) \right| d\alpha \\ &\leq M_f |\mathcal{X}| L_P \sum_{x' \in \mathcal{X}} \int_{\mathcal{I}} \left| \int_{\mathcal{I}} W(\alpha, \beta) \hat{\mu}_t^\beta(x') d\beta - \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta(x') d\beta \right| d\alpha. \end{aligned}$$

Then, the induction assumption can be applied to

$$\int_{\mathcal{I}} \left| \int_{\mathcal{I}} W(\alpha, \beta) \hat{\mu}_t^\beta(x') d\beta - \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta(x') d\beta \right| d\alpha = \int_{\mathcal{I}} |\hat{\mu}_t(f'_{\alpha, x'}) - \mu_t(f'_{\alpha, x'})| d\alpha \rightarrow 0$$

for $\xi \rightarrow 0$ where we choose $f'_{x', \alpha}(x, \beta) = W(\alpha, \beta) \cdot \mathbf{1}_{\{x=x'\}}$. This immediately implies

$$\left| \hat{\mu}_{t, \hat{\mu}_t, W}^{P^\pi}(f) - \hat{\mu}_{t, \mu_t, W}^{P^\pi}(f) \right| \rightarrow 0 \quad \text{as } \xi \rightarrow 0.$$

Third term. For the last term we have by the induction assumption

$$\begin{aligned} \left| \hat{\mu}_{t, \mu_t, W}^{P^\pi}(f) - \mu_{t+1}(f) \right| &= \left| \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} \hat{\mu}_t^\alpha(x) f''(x, \alpha) d\alpha - \int_{\mathcal{I}} \sum_{x \in \mathcal{X}} \mu_t^\alpha(x) f''(x, \alpha) d\alpha \right| \\ &= |\hat{\mu}_t(f'') - \mu_t(f'')| \rightarrow 0 \quad \text{as } \xi \rightarrow 0. \end{aligned}$$

Here we have applied the induction assumption to the function f'' defined by

$$f''(x, \alpha) = \sum_{\substack{u \in \mathcal{U} \\ x' \in \mathcal{X}}} \pi_t^\alpha(u \mid x) P \left(x' \mid x, u, \int_{\mathcal{I}} W(\alpha, \beta) \mu_t^\beta d\beta \right) f(x', \alpha).$$

G PROOF OF LEMMA 5

We start by showing that (19) implies (20). Subsequently, we establish (19) to complete the proof. First, note that

$$\begin{aligned} \left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_W^\alpha(\mu_t))] - \mathbb{E} [h(\hat{X}_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\mu}_t))] \right| &= \left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_W^\alpha(\mu_t))] - \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\mu_t))] \right| \\ &\quad + \left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\mu_t))] - \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\mu}_t))] \right| \\ &\quad + \left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\mu}_t))] - \mathbb{E} [h(\hat{X}_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\mu}_t))] \right|. \end{aligned}$$

Let us consider the three terms separately. We denote by L_h the Lipschitz constant and by M_h the uniform bound of \mathcal{H} .

First term. For the first term we have

$$\begin{aligned} \left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_W^\alpha(\boldsymbol{\mu}_t))] - \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\boldsymbol{\mu}_t))] \right| &\leq L_h \mathbb{E} \left[\left\| \mathbb{G}_W^\alpha(\boldsymbol{\mu}_t) - \mathbb{G}_{\widehat{W}}^\alpha(\boldsymbol{\mu}_t) \right\| \right] \\ &= L_h \sum_{x \in \mathcal{X}} \left| \int_{\mathcal{I}} (W(\alpha, \beta) - \widehat{W}(\alpha, \beta)) \mu_t^\beta(x) d\beta \right| \end{aligned}$$

which converges to 0 as $\xi \rightarrow 0$ for arbitrary fractions of the possible α since the set of points $(\alpha, \beta) \in \mathcal{I}^2$ where W and \widehat{W} differ becomes arbitrarily small by construction for ξ close enough to 0.

Second term. Using a similar reformulation as for the first term, we obtain

$$\left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\boldsymbol{\mu}_t))] - \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\boldsymbol{\mu}}_t))] \right| \leq L_h \sum_{x \in \mathcal{X}} \left| \int_{\mathcal{I}} (\mu_t^\beta(x) - \hat{\mu}_t^\beta(x)) \widehat{W}(\alpha, \beta) d\beta \right|$$

which converges to zero as $\xi \rightarrow 0$. This follows from applying Lemma 4 to the functions $f_{\alpha, x}(x', \beta) = \widehat{W}(\alpha, \beta) \mathbf{1}_{\{x=x'\}}$.

Third term. Finally, the last term

$$\left| \mathbb{E} [h(X_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\boldsymbol{\mu}}_t))] - \mathbb{E} [h(\hat{X}_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\boldsymbol{\mu}}_t))] \right|$$

converges to zero as $\xi \rightarrow 0$ by (19). Thus, we have established that (19) implies (20).

Now it remains to show (19). We prove (19) via induction over t . For $t = 0$, the statement trivially holds. For the induction step, we note that

$$\left| \mathbb{E} [g(X_{t+1}^\alpha)] - \mathbb{E} [g(\hat{X}_{t+1}^\alpha)] \right| = \left| \mathbb{E} [f'_t(X_t^\alpha, \mathbb{G}_W^\alpha(\boldsymbol{\mu}_t))] - \mathbb{E} [f'_t(\hat{X}_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\boldsymbol{\mu}}_t))] \right|$$

with

$$f'_t(x, \nu) = \sum_{u \in \mathcal{U}} \pi_t(u | x) \sum_{x' \in \mathcal{X}} P(x' | x, u, \nu) g(x').$$

Then we can apply (20) to

$$\left| \mathbb{E} [f'_t(X_t^\alpha, \mathbb{G}_W^\alpha(\boldsymbol{\mu}_t))] - \mathbb{E} [f'_t(\hat{X}_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\boldsymbol{\mu}}_t))] \right|$$

to obtain the desired result. Note that (20) holds by the induction assumption because we have already established that (19) implies (20). This concludes the proof.

H PROOF OF COROLLARY 1

The statement follows from Lemma 5. To see this, we define the functions

$$r_t(x, \nu) = \sum_{u \in \mathcal{U}} \pi_t(u | x) r(x, u, \nu)$$

which are uniformly Lipschitz and uniformly bounded by construction. Then, we obtain by Lemma 5

$$|J_{\alpha, W}^t(\pi) - J_{\alpha, \widehat{W}}^t(\pi)| \leq \sum_{t=0}^{T-1} \left| \mathbb{E} [r_t(X_t^\alpha, \mathbb{G}_W^\alpha(\boldsymbol{\mu}_t))] - \mathbb{E} [r_t(\hat{X}_t^\alpha, \mathbb{G}_{\widehat{W}}^\alpha(\hat{\boldsymbol{\mu}}_t))] \right| \leq \varepsilon$$

for all $\alpha \in \mathcal{J}$ for some $\mathcal{J} \subseteq \mathcal{I}$ with $\lambda(\mathcal{J}) \geq 1 - p$. This concludes the proof.

I DETAILS ON LEARNING LPGMFGS

Before we prove the theoretical results on learning LPGMFGs, we introduce some terminology and technical details that were not included in the main text for space reasons. Our implementation is based on Cui and Koepl (2022) (MIT license). For Figure 1, we used data from Rozemberczki et al. (2019) (GPL-3.0 license). The dataset from Rozemberczki et al. (2019) contains mutual likes among verified Facebook pages of TV shows that were obtained from public APIs. Thus, the data set does not contain any personal information of individuals and no person is identifiable. Our code is given in the supplemental material. The total amount of compute to reproduce our work amounts to approximately 2 days on a 64 core AMD Epyc processor and 64 GB of RAM. No GPUs were used in this research.

Equivalence Classes. For learning LPGMFGs, we apply the well-known equivalence class method, see for example Cui and Koepl (2022). Therefore, we discretize the continuous spectrum of agents \mathcal{I} into $M + 1$ disjoint subsets. Formally, we consider a grid of agents $(\alpha_m)_{0 \leq m \leq M}$ with $\alpha_m = m/100$ where each α_m is associated with a policy $\pi^{\alpha_m} \in \mathcal{P}(\mathcal{U})^T$ and a mean field $\mu^{\alpha_m} \in \mathcal{P}(\mathcal{X})^T$. Then, each agent $\alpha \in [0, 1]$ follows the policy π^{α_m} of the agent at grid point α_m closest to α . Thus, there are $M + 1$ intervals such that $\mathcal{I} = \bigcup_{m=0}^M \tilde{\mathcal{I}}_m$ with $\tilde{\mathcal{I}}_m = [\frac{m-1/2}{M}, \frac{m+1/2}{M})$ for $0 < m < M$, $\tilde{\mathcal{I}}_0 = [0, \frac{1}{2M}]$, and $\tilde{\mathcal{I}}_M = [\frac{M-1}{M}, 1]$. Note that all agents in an interval follow the same policy. For any fixed policy ensemble $\boldsymbol{\pi} \in \boldsymbol{\Pi}$ we define

$$\hat{\Psi}(\boldsymbol{\pi}) = \sum_{m=1}^M \mathbf{1}_{\{\alpha \in \tilde{\mathcal{I}}_m\}} \hat{\mu}^{\alpha_m}$$

where $\hat{\mu}$ is recursively defined by

$$\hat{\mu}_{t+1}^{\alpha_m}(x) = \sum_{\substack{x' \in \mathcal{X} \\ u \in \mathcal{U}}} \hat{\mu}_t^{\alpha_m}(x') \pi_t^{\alpha_m}(u|x') P(x|x', u, \hat{\mathbb{G}}_t^{\alpha_m})$$

with $\hat{\mu}_0^{\alpha_m}(x) = \mu_0(x)$ for $x \in \mathcal{X}$ and $m = 1, \dots, M$, and neighborhood mean fields

$$\hat{\mathbb{G}}_t^\alpha = \frac{1}{M} \sum_{m=1}^M W(\alpha, \alpha_m) \hat{\mu}_t^{\alpha_m}.$$

Similarly, we approximate the policy ensemble by

$$\hat{\Phi}(\boldsymbol{\mu}) = \sum_{m=1}^M \mathbf{1}_{\{\alpha \in \tilde{\mathcal{I}}_m\}} \pi^{\alpha_m}$$

where π^{α_m} is the optimal policy of α_m for any fixed $\boldsymbol{\mu}$.

Multi-Population Mean Field Games. In this paragraph, we briefly introduce some crucial definitions for the following proofs. For more details on MP MFGs, see for example Perolat et al. (2021). A player of population i who implements policy π^i while the mean field of players is distributed according to $\boldsymbol{\mu} \in \boldsymbol{\mathcal{M}}$ can expect the following sum of rewards

$$J_i^\mu(\pi^i) = \mathbb{E} \left[\sum_{t=0}^T r(x_t^i, u_t^i, \mu_t) \mid x_0^i \sim \mu_0^i, u_t^i \sim \pi_t^i(\cdot | x_t^i), x_{t+1}^i \sim P(\cdot | x_t^i, u_t^i, \mu_t) \right].$$

Similarly, we define the corresponding Q -function

$$Q_i^{\mu, \pi^i}(t, x^i, u^i) = \mathbb{E} \left[\sum_{k=t}^T r(x_k^i, u_k^i, \mu_k) \mid x_t^i = x^i, u_t^i = u^i, u_k^i \sim \pi_k^i(\cdot | x_k^i), x_{k+1}^i \sim P(\cdot | x_k^i, u_k^i, \mu_k) \right]$$

and the corresponding value function

$$V_i^{\mu, \pi^i}(t, x^i) = \mathbb{E} \left[\sum_{k=t}^T r(x_k^i, u_k^i, \mu_k) \mid x_t^i = x^i, u_k^i \sim \pi_k^i(\cdot | x_k^i), x_{k+1}^i \sim P(\cdot | x_k^i, u_k^i, \mu_k) \right].$$

The sequence of state distributions $\boldsymbol{\mu}^{\boldsymbol{\pi}, \boldsymbol{\mu}'}$ $\in \boldsymbol{\mathcal{M}}$ is defined by the forward equation

$$\mu_{t+1}^{i, \boldsymbol{\pi}, \boldsymbol{\mu}'}(x^i) := \sum_{(x^i, u^i) \in \mathcal{X} \times \mathcal{U}} \pi_t^i(u^i | x^i) \mu_t^{i, \boldsymbol{\pi}, \boldsymbol{\mu}'}(x^i) P(x^i | x^i, u^i, \mu_t^i)$$

for $(\boldsymbol{\pi}, \boldsymbol{\mu}) \in \boldsymbol{\Pi} \times \boldsymbol{\mathcal{M}}$. In the case of $\boldsymbol{\mu} = \boldsymbol{\mu}'$, we will often use the abbreviated notation $\boldsymbol{\mu}^\boldsymbol{\pi} = \boldsymbol{\mu}^{\boldsymbol{\pi}, \boldsymbol{\mu}'}$.

OMD Algorithm. As we will see in the subsequent statements and proofs, there is some terminology necessary for analyzing OMDs such as the concept of a regularizer $h: \mathcal{P}(\mathcal{U}) \rightarrow \mathbb{R}$. Here, we choose the entropy as the regularizer, i.e. $h(\pi) := -\sum_{u \in \mathcal{U}} \pi(u) \log(\pi(u))$. Then, the convex conjugate of h is defined as $h^*(y) = \max_{p \in \mathcal{P}(\mathcal{U})} (\langle y, p \rangle - h(\pi)) = \log(\sum_{u \in \mathcal{U}} \exp(y(u)))$. Finally, we define for (almost) every y

$$\Gamma(y) := \nabla h^*(y) = \arg \max_{p \in \mathcal{P}(\mathcal{U})} (\langle y, p \rangle - h(\pi)) = \frac{\exp(y)}{\sum_{u' \in \mathcal{U}} \exp(y(u'))}.$$

A more detailed introduction to these concepts can be found, for example, in Perolat et al. (2021). For our learning approach, we apply the OMD algorithm as formulated in Perolat et al. (2021), see Algorithm 1.

Proof Strategy The theoretical analysis is conducted by considering the continuous time analogue (Mertikopoulos et al. (2018)) of Algorithm 1 on the smoothed step graphon, which we refer to as smoothed CTOMD, and which is characterized by the equations

$$y_{t,\tau}^\alpha(x^\alpha, u^\alpha) = \int_0^\tau Q_{\alpha^s}^{\pi_{\alpha^s}, \mu^{\pi_{\alpha^s}}} (x^\alpha, u^\alpha) ds \quad (21)$$

and

$$\pi_{t,\tau}^\alpha(\cdot | x^\alpha) = \Gamma(y_{t,\tau}^\alpha(x^\alpha, \cdot)) \quad (22)$$

for all $\alpha \in \mathcal{I}$, $t \in \mathcal{T}$, and $\tau \in \mathbb{R}_+$ with initial values $y_{t,0}^\alpha = 0$. Taking the derivative with respect to τ on both sides of equation (21) yields

$$\frac{d}{d\tau} y_{t,\tau}^\alpha(x^\alpha, u^\alpha) = Q_{\alpha^\tau}^{\pi_{\alpha^\tau}, \mu^{\pi_{\alpha^\tau}}} (x^\alpha, u^\alpha)$$

and taking the derivative with respect to τ on both sides of equation (22) results in

$$\begin{aligned} \frac{d}{d\tau} \pi_{t,\tau}^\alpha(u^\alpha | x^\alpha) &= \frac{d}{d\tau} \Gamma(y_{t,\tau}^\alpha(x^\alpha, u^\alpha)) \\ &= \exp(y_{t,\tau}^\alpha(x^\alpha, u^\alpha)) \cdot Q_{\alpha^\tau}^{\pi_{\alpha^\tau}, \mu^{\pi_{\alpha^\tau}}} (x^\alpha, u^\alpha) \cdot \frac{\sum_{u' \in \mathcal{U}} \exp(y_{t,\tau}^\alpha(x^\alpha, u'^\alpha))}{(\sum_{u' \in \mathcal{U}} \exp(y_{t,\tau}^\alpha(x^\alpha, u'^\alpha)))^2} \\ &\quad - \left[\sum_{u' \in \mathcal{U}} Q_{\alpha^\tau}^{\pi_{\alpha^\tau}, \mu^{\pi_{\alpha^\tau}}} (x^\alpha, u'^\alpha) \exp(y_{t,\tau}^\alpha(x^\alpha, u'^\alpha)) \right] \cdot \frac{\exp(y_{t,\tau}^\alpha(x^\alpha, u^\alpha))}{(\sum_{u' \in \mathcal{U}} \exp(y_{t,\tau}^\alpha(x^\alpha, u'^\alpha)))^2} \\ &= \pi_{t,\tau}^\alpha(u^\alpha | x^\alpha) \cdot Q_{\alpha^\tau}^{\pi_{\alpha^\tau}, \mu^{\pi_{\alpha^\tau}}} - \pi_{t,\tau}^\alpha(u^\alpha | x^\alpha) \cdot \sum_{u' \in \mathcal{U}} Q_{\alpha^\tau}^{\pi_{\alpha^\tau}, \mu^{\pi_{\alpha^\tau}}} (x^\alpha, u'^\alpha) \pi_{t,\tau}^\alpha(u'^\alpha | x^\alpha) =: g(\pi_{t,\tau}^\alpha(u^\alpha | x^\alpha)). \end{aligned}$$

Following the above derivations, our goal is to solve the differential equation

$$\frac{d}{d\tau} \pi_{t,\tau}^\alpha(u^\alpha | x^\alpha) = g(\pi_{t,\tau}^\alpha(u^\alpha | x^\alpha)). \quad (23)$$

The strategy of this section is to eventually apply LaSalle's theorem (Khalil, 2002, Theorem 4.4) to equation (23) which will yield Theorem 5. First, we define a measure of similarity

$$H(y) := \int_{\mathcal{I}} \sum_{t=0}^T \sum_{x^i \in \mathcal{X}} \mu_t^{i, \pi^*} (x^i) \left[h^*(y_{t,\tau}^i(x^i, \cdot)) - h^*(y_t^{i,*}(x^i, \cdot)) - \left\langle \pi_t^{i,*}(\cdot | x^i), y_{t,\tau}^i(x^i, \cdot) - y_t^{i,*}(x^i, \cdot) \right\rangle \right] d\alpha \quad (24)$$

using the convention that $\Gamma(y_t^{i,*}(x^i, \cdot)) = \pi_t^{i,*}(\cdot | x^i)$, where π^* is the unique NE. Next, we make a rather technical observation.

Lemma 6. *The convex conjugate of the regularizer satisfies*

$$\frac{d}{d\tau} h^*(y_{t,\tau}^\alpha(x^\alpha, \cdot)) = \left\langle \pi_{t,\tau}^\alpha(\cdot | x^\alpha), Q_{\alpha^\tau}^{\mu^{\pi_{\alpha^\tau}}, \pi_{\alpha^\tau}}(t, x^\alpha, \cdot) \right\rangle.$$

Algorithm 1 Online Mirror Descent (Perolat et al. (2021))

- 1: **Input:** γ , number of iterations τ_{\max} , and $\forall i \leq M, t \in \mathcal{T} : y_{t,0}^i = 0$
 - 2: **for** $\tau = 1, \dots, \tau_{\max}$ **do**
 - 3: Forward update for all i : μ^{i, π_τ}
 - 4: Backward update for all i : $Q_i^{\pi_\tau, \mu^{\pi_\tau}}$
 - 5: Update for all i, t, x, u
 - 6: $y_{t, \tau+1}^i(x, u) = y_{t, \tau}^i(x, u) + \gamma \cdot Q_i^{\pi_\tau, \mu^{\pi_\tau}}(x, u)$
 - 7: $\pi_{t, \tau+1}^i(\cdot | x) = \Gamma(y_{t, \tau+1}^i(x, \cdot))$
 - 8: **end for**
-

Then, we can leverage this technical observation to derive the following result which will be crucial for proving Theorem 5 from the main text.

Lemma 7. *The similarity measure H defined in equation (24) satisfies*

$$\frac{d}{d\tau} H(y_\tau) = \Delta J(\pi_\tau, \pi^*) + \tilde{d}(\pi_\tau, \pi^*)$$

where $\Delta J(\pi_\tau, \pi^*) = \int_{\mathcal{I}} J_\alpha^{\mu^*}(\pi_\tau^\alpha) - J_\alpha^{\mu^*}(\pi^{\alpha, *}) d\alpha$ is always non-positive with $\mu^* = \Psi(\pi^*)$ and $\tilde{d}(\pi, \pi')$ defined as in equation (13).

Note that the proofs of Lemma 6 and Lemma 7 can be found in the subsequent sections.

J PROOF OF LEMMA 1

Note that the following proof is based on an argument structure in Perolat et al. (2021).

Proof. Theorem 1 ensures that a NE exists. So we prove uniqueness by contradiction. Therefore, assume there are two different Nash Equilibria π and π' . Their associated mean fields are denoted by $\mu = \Psi(\pi)$ and $\mu' = \Psi(\pi')$ as before. On the one hand, the strictly weak monotonicity of the mean field game implies

$$\int_{\mathcal{I}} \left[J_\alpha^\mu(\pi^\alpha) + J_\alpha^{\mu'}(\pi'^\alpha) - J_\alpha^\mu(\pi'^\alpha) - J_\alpha^{\mu'}(\pi^\alpha) \right] d\alpha < 0.$$

On the other hand, we know that both $\int_{\mathcal{I}} J_\alpha^\mu(\pi^\alpha) - J_\alpha^\mu(\pi'^\alpha) d\alpha \geq 0$ and $\int_{\mathcal{I}} J_\alpha^{\mu'}(\pi'^\alpha) - J_\alpha^{\mu'}(\pi^\alpha) d\alpha \geq 0$ hold by assuming that π and π' are Nash Equilibria. This in turn yields

$$\int_{\mathcal{I}} \left[J_\alpha^\mu(\pi^\alpha) + J_\alpha^{\mu'}(\pi'^\alpha) - J_\alpha^\mu(\pi'^\alpha) - J_\alpha^{\mu'}(\pi^\alpha) \right] d\alpha \geq 0.$$

which is an obvious contradiction to the strictly weak monotonicity. This concludes the proof. \square

K PROOF OF LEMMA 6

Proof. Recall that h^* is defined as a function $h^*: \mathbb{R}^{|\mathcal{U}|} \rightarrow \mathbb{R}$. For notational convenience, we define $h^*(a_1, \dots, a_{|\mathcal{U}|}) = z \in \mathbb{R}$ where $a_1, \dots, a_{|\mathcal{U}|} \in \mathbb{R}$ and each a_j represents the probability weight of one specific action $u_j \in \mathcal{U}$. Furthermore, we denote by $\frac{d}{d\tau}$ the total differential with respect to τ . These definitions allow us to state

$$\frac{d}{d\tau} h^*(y_{t, \tau}^\alpha(x^\alpha, \cdot)) = \sum_{u_j \in \mathcal{U}} \frac{\partial h^*}{\partial a_j}(y_{t, \tau}^\alpha(x^\alpha, u_j)) \frac{dy_{t, \tau}^\alpha(x^\alpha, u_j)}{d\tau}.$$

Keeping in mind the OMD algorithm and especially

$$y_{t, \tau}^\alpha(x^\alpha, u^\alpha) = \int_0^\tau Q_\alpha^{\mu^{\pi_s}, \pi_s^\alpha}(t, x^\alpha, u^\alpha) ds$$

yields

$$\frac{dy_{t,\tau}^\alpha(x^\alpha, u_j)}{d\tau} = Q_i^{\mu^{\pi_\tau}, \pi_\tau^\alpha}(t, x^\alpha, u^\alpha).$$

Additionally, by definition we have

$$\Gamma(y) = \nabla h^*(y) \quad \text{and} \quad \pi_{t,\tau}^\alpha(\cdot|x^\alpha) = \Gamma(y_{t,\tau}^\alpha(x^\alpha, \cdot))$$

which yields

$$\frac{\partial h^*}{\partial a_j}(y_{t,\tau}^\alpha(x^\alpha, u_j)) = \Gamma(y_{t,\tau}^\alpha(x^\alpha, u_j)) = \pi_{t,\tau}^\alpha(u_j|x^\alpha).$$

In summary, the above observations imply

$$\begin{aligned} \frac{d}{d\tau} h^*(y_{t,\tau}^\alpha(x^\alpha, \cdot)) &= \sum_{u_j \in \mathcal{U}} \frac{\partial h^*}{\partial a_j} \frac{dy_{t,\tau}^\alpha(x^\alpha, u_j)}{d\tau} \\ &= \sum_{u_j \in \mathcal{U}} \pi_{t,\tau}^\alpha(u_j|x^\alpha) Q_i^{\mu^{\pi_\tau}, \pi_\tau^\alpha}(t, x^\alpha, u^\alpha) = \left\langle \pi_{t,\tau}^\alpha(\cdot|x^\alpha), Q_i^{\mu^{\pi_\tau}, \pi_\tau^\alpha}(t, x^\alpha, \cdot) \right\rangle \end{aligned}$$

which yields the desired result. \square

L PROOF OF LEMMA 7

This section follows (Perolat et al., 2021, Appendix D).

Proof. Assume that $\pi^* \in \Pi$ is a NE. Furthermore, for all $\tau > 0$ define the mean field $\mu'_\tau = \Psi(\pi_\tau)$ and keep in mind that the transition kernel does not depend on the mean field by assumption. Then we have

$$\begin{aligned} &\frac{d}{d\tau} H(y_\tau) \\ &= \int_{\mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \frac{d}{d\tau} [h^*(y_{t,\tau}^\alpha(x^\alpha, \cdot)) - h^*(y_t^{\alpha,*}(x^\alpha, \cdot)) - \langle \pi_t^{\alpha,*}(\cdot|x^\alpha), y_{t,\tau}^\alpha(x^\alpha, \cdot) - y_t^{\alpha,*}(x^\alpha, \cdot) \rangle] d\alpha \\ &= \int_{\mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \cdot \left[\frac{d}{d\tau} h^*(y_{t,\tau}^\alpha(x^\alpha, \cdot)) - \left\langle \pi_t^{\alpha,*}(\cdot|x^\alpha), \frac{d}{d\tau} y_{t,\tau}^\alpha(x^\alpha, \cdot) \right\rangle \right] d\alpha \\ &= \int_{\mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \left[\left\langle \pi_{t,\tau}^\alpha(\cdot|x^\alpha), Q_\alpha^{\mu'_\tau, \pi_\tau^\alpha}(t, x^\alpha, \cdot) \right\rangle - \left\langle \pi_t^{\alpha,*}(\cdot|x^\alpha), Q_i^{\mu'_\tau, \pi_\tau^\alpha}(t, x^\alpha, \cdot) \right\rangle \right] d\alpha \\ &= \int_{\mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \left[V_\alpha^{\mu'_\tau, \pi_\tau^\alpha}(t, x^\alpha) - \langle \pi_t^{\alpha,*}(\cdot|x^\alpha), r(x^\alpha, \cdot, \mu'_{t,\tau}) \rangle \right] d\alpha \\ &\quad - \int_{\mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \left\langle \pi_t^{\alpha,*}(\cdot|x^\alpha), \sum_{x'^\alpha \in \mathcal{X}} P(x'^\alpha|x^\alpha, \cdot) V_\alpha^{\mu'_\tau, \pi_\tau^\alpha}(t+1, x'^\alpha) \right\rangle d\alpha \\ &= \int_{\mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \left[\mu_t^{\alpha, \pi^*}(x^\alpha) V_\alpha^{\mu'_\tau, \pi_\tau^\alpha}(t, x^\alpha) \right] - \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \langle \pi_t^{\alpha,*}(\cdot|x^\alpha), r(x^\alpha, \cdot, \mu'_{t,\tau}) \rangle \\ &\quad - \sum_{t \in \mathcal{T}} \sum_{x'^\alpha \in \mathcal{X}} V_\alpha^{\mu'_\tau, \pi_\tau^\alpha}(t+1, x'^\alpha) \mu_{t+1}^{\alpha, \pi^*}(x'^\alpha) d\alpha \\ &= \int_{\mathcal{I}} \left(\sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \left[\mu_t^{\alpha, \pi^*}(x^\alpha) V_\alpha^{\mu'_\tau, \pi_\tau^\alpha}(t, x^\alpha) \right] - \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \left[\mu_{t+1}^{\alpha, \pi^*}(x^\alpha) V_i^{\mu'_\tau, \pi_\tau^\alpha}(t+1, x^\alpha) \right] \right. \\ &\quad \left. - \sum_{t \in \mathcal{T}} \sum_{x^\alpha \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^\alpha) \langle \pi_t^{\alpha,*}(\cdot|x^\alpha), r(x^\alpha, \cdot, \mu'_{t,\tau}) \rangle \right) d\alpha \end{aligned}$$

$$\begin{aligned}
 &= \int_{\mathcal{I}} \left[J_{\alpha}^{\mu'_{\tau}}(\pi_{\tau}^{\alpha}) - J_{\alpha}^{\mu'_{\tau}}(\pi^{\alpha,*}) \right] d\alpha \\
 &= \int_{\mathcal{I}} \left[J_{\alpha}^{\mu'_{\tau}}(\pi_{\tau}^{\alpha}) - J_{\alpha}^{\mu'_{\tau}}(\pi^{\alpha,*}) + J_i^{\mu^*}(\pi^{\alpha,*}) - J_{\alpha}^{\mu^*}(\pi_{\tau}^{\alpha}) - J_{\alpha}^{\mu^*}(\pi^{\alpha,*}) + J_{\alpha}^{\mu^*}(\pi_{\tau}^{\alpha}) \right] d\alpha \\
 &= \Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*)
 \end{aligned}$$

where the third equality follows from Lemma 6. \square

M PROOF OF THEOREM 5

The following proof is based on an idea in (Perolat et al., 2021, Appendix G).

Proof. Define the function $f : \Pi \rightarrow \mathbb{R}$ given by

$$f(\pi) = \int_{\mathcal{I}} \sum_{t=0}^T \sum_{x^{\alpha} \in \mathcal{X}} \mu_t^{\alpha, \pi^*}(x^{\alpha}) D_{\text{KL}}(\pi_t^{\alpha, \pi^*}(\cdot | x^{\alpha}), \pi_t^{\alpha}(\cdot | x^{\alpha})) d\alpha$$

where π^* is the NE of the smoothed MP-MFG which is unique by Lemma 1. We point out that this immediately implies that $f(\pi) = 0$ if and only if $\pi = \pi^*$ (almost everywhere) and $f(\pi) > 0$ otherwise by the basic properties of the Kullback-Leibler divergence.

The just defined function f and the function H defined in (24) are closely related. To see that, we start with some calculations

$$\begin{aligned}
 &h^*(y_{t,\tau}^i(x^i, \cdot)) - h^*(y_t^{i,*}(x^i, \cdot)) - \left\langle \pi_t^{i,*}(\cdot | x^i), y_{t,\tau}^i(x^i, \cdot) - y_t^{i,*}(x^i, \cdot) \right\rangle \\
 &= \log \left(\sum_{u \in \mathcal{U}} \exp(y_{t,\tau}^i(x^i, u)) \right) - \log \left(\sum_{u \in \mathcal{U}} \exp(y_t^{i,*}(x^i, u)) \right) - \left\langle \pi_t^{i,*}(\cdot | x^i), y_{t,\tau}^i(x^i, \cdot) - y_t^{i,*}(x^i, \cdot) \right\rangle \\
 &= \log \left(\sum_{u \in \mathcal{U}} \exp(y_{t,\tau}^i(x^i, u)) \right) - \left\langle \pi_t^{i,*}(\cdot | x^i), y_{t,\tau}^i(x^i, \cdot) \right\rangle + \sum_{u \in \mathcal{U}} \pi_t^{i,*}(u | x^i) \log \left(\frac{\log(y_t^{i,*}(x^i, u))}{\sum_{u' \in \mathcal{U}} \exp(y_t^{i,*}(x^i, u'))} \right) \\
 &= \log \left(\sum_{u' \in \mathcal{U}} \exp(y_{t,\tau}^i(x^i, u')) \right) - \left\langle \pi_t^{i,*}(\cdot | x^i), y_{t,\tau}^i(x^i, \cdot) \right\rangle + \sum_{u \in \mathcal{U}} \pi_t^{i,*}(u | x^i) \cdot \log \left(\pi_t^{i,*}(u | x^i) \right) \\
 &= \sum_{u \in \mathcal{U}} \pi_t^{i,*}(u | x^i) \cdot \left[\log \left(\pi_t^{i,*}(u | x^i) \right) + \log \left(\frac{\sum_{u' \in \mathcal{U}} \exp(y_{t,\tau}^i(x^i, u'))}{\exp(y_{t,\tau}^i(x^i, u))} \right) \right] \\
 &= \sum_{u \in \mathcal{U}} \pi_t^{i,*}(u | x^i) \cdot \log \left(\frac{\pi_t^{i,*}(u | x^i)}{\Gamma(y_{t,\tau}^i(x^i, u))} \right) \\
 &= D_{\text{KL}} \left(\pi_t^{i,*}(\cdot | x^i), \Gamma(y_{t,\tau}^i(x^i, \cdot)) \right).
 \end{aligned}$$

Therefore, we have $H(y) = f(\pi)$ for all pairs y, π for which $\Gamma(y) = \pi$ holds. This, in turn implies together with Lemma 6 that for an arbitrary $\pi_{\tau} \in \pi$ we have

$$\frac{d}{d\tau} f(\pi_{\tau}) = \Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) \quad (25)$$

Now, we claim that

$$\frac{d}{d\tau} f(\pi_{\tau}) = \Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) \leq 0 \quad (26)$$

is true for all $\pi_{\tau} \in \pi$ and that it is equal to zero if and only if π_{τ} is the NE, i.e. $\pi_{\tau} = \pi^*$. To see this, consider an arbitrary but fixed $\pi_{\tau} \in \pi$. The term $\Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*)$ can either be negative or equal to zero. If $\Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) = 0$ holds, this implies $\tilde{d}(\pi_{\tau}, \pi^*) = 0$ which in turn yields $\mu^{\pi_{\tau}} = \mu^{\pi^*}$. Besides that, $\Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) = 0$ also means that $\Delta J(\pi_{\tau}, \pi^*) = 0$ has to hold. Reformulating the equation, we immediately obtain $\int_{\mathcal{I}} J_{\alpha}^{\mu^*}(\pi_{\tau}^{\alpha}) d\alpha =$

$\int_{\mathcal{I}} J_{\alpha}^{\mu^*}(\pi^{\alpha,*}) d\alpha$. In combination with the previous observation $\mu^{\pi_{\tau}} = \mu^{\pi^*}$, this implies that π_{τ} is a NE, i.e. $\pi^* = \pi_{\tau}$. To see this, recall that π^* is a NE by assumption and thereby a best response to μ^{π^*} (and $\mu^{\pi_{\tau}}$ by the first argument). Then, $\int_{\mathcal{I}} J_{\alpha}^{\mu^*}(\pi_{\tau}^{\alpha}) d\alpha = \int_{\mathcal{I}} J_{\alpha}^{\mu^*}(\pi^{\alpha,*}) d\alpha$ means that both π^* and π_{τ} yield the same expected reward given $\mu^{\pi_{\tau}}$. Since π^* is a best response to $\mu^{\pi_{\tau}}$, π_{τ} also has to be a best response to $\mu^{\pi_{\tau}}$ and therefore a NE.

However, if $\Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) < 0$ holds, π cannot be a NE. Assume, by contradiction, that π is a NE. Then, we have

$$\tilde{d}(\pi_{\tau}, \pi^*) = \int_{\mathcal{I}} \left[J_{\alpha}^{\mu_{\tau}}(\pi_{\tau}^{\alpha}) - J_{\alpha}^{\mu_{\tau}}(\pi^{\alpha,*}) + J_{\alpha}^{\mu^*}(\pi^{\alpha,*}) - J_{\alpha}^{\mu^*}(\pi_{\tau}^{\alpha}) \right] d\alpha \geq 0$$

because π_{τ} and π^* which implies $\int_{\mathcal{I}} J_{\alpha}^{\mu_{\tau}}(\pi_{\tau}^{\alpha}) - J_{\alpha}^{\mu_{\tau}}(\pi^{\alpha,*}) d\alpha \geq 0$ and $\int_{\mathcal{I}} J_{\alpha}^{\mu^*}(\pi^{\alpha,*}) - J_{\alpha}^{\mu^*}(\pi_{\tau}^{\alpha}) d\alpha \geq 0$, respectively. This observation combined with inequality (13) yields $\tilde{d}(\pi_{\tau}, \pi^*) = 0$ and thereby $\mu^{\pi_{\tau}} = \mu^{\pi^*}$. Thus, π_{τ} and π^* are both best responses to μ^{π^*} which means that $\Delta J(\pi_{\tau}, \pi^*) = 0$. This is a contradiction to the assumption $\Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) < 0$. Hence, $\Delta J(\pi_{\tau}, \pi^*) + \tilde{d}(\pi_{\tau}, \pi^*) < 0$ implies that π_{τ} is not a NE.

With the above arguments in place, we can apply LaSalle's theorem (Khalil, 2002, Theorem 4.4) to solve the differential equation (23). Choosing $f(\mathbf{y})$ as the function V in (Khalil, 2002, Theorem 4.1) with $\mathbf{\Pi}$ as the compact set, we point out that the unique NE π^* is the only element in $\mathbf{\Pi}$ with $\frac{d}{dt}f(\pi^*) = 0$ as we have established in equation (26). Therefore, the OMD algorithm converges to π^* for every starting point $\pi_0 \in \mathbf{\Pi}$. \square

N CUTOFF POWER LAW GRAPHON

In this paragraph we consider graphs G_n with cutoff power law degree distributions. The benefits of adding a cutoff are twofold. On the one hand, a technical advantage of the modified degree distribution is that it is both integrable and Lipschitz continuous. On the other hand, the cutoff version of the power-law turns out to be a more realistic modelling option in numerous real-world applications. In practice, the expression 'power law distribution' frequently refers only to the tail of the distribution. This accounts for the fact that pure power laws often do not describe the empirical observations accurately, see for example Newman (2005). The interested reader is also referred to Newman (2018) and the references therein for an overview on the topic.

The cutoff power-law distribution can be constructed by connecting two vertices $i, j \in \{1, \dots, n\}$ with probability $\min\left(1, n^{\beta} (\max\{i, c \cdot n\} \cdot \max\{j, c \cdot n\})^{-\alpha}\right)$ independently of all other edges where $0 < c < 1$, $0 < \alpha < 1$, and $0 \leq \beta < 2\alpha$. Furthermore, the assumption $\beta > 2\alpha - 1$ implies that the expected number of edges is super-linear. We note that $n^{\beta} (c \cdot n \cdot c \cdot n)^{-\alpha} = c^{-2\alpha} n^{\beta-2\alpha} = o(1)$ is implied by the assumption $\beta < 2\alpha$. This in turn allows us to drop the min term in the edge probability since for $n \rightarrow \infty$ we have $n^{\beta} (\max\{i, c \cdot n\} \cdot \max\{j, c \cdot n\})^{-\alpha} = o(1)$. The following calculation provides the expected edge density which is a key element for the subsequent arguments.

$$\begin{aligned} \mathbb{E}[\|G_n\|_1] &= n^{-2} \sum_{i,j \in V(G_n)} \mathbb{E}[\mathbf{1}_{(ij) \in E(G_n)}] \\ &= \sum_{i,j \in V(G_n)} \frac{\min\left(1, n^{\beta} (\max\{i, cn\} \cdot \max\{j, cn\})^{-\alpha}\right)}{n^2} \\ &= n^{-2} \sum_{i,j \in V(G_n)} n^{\beta} (\max\{i, c \cdot n\} \cdot \max\{j, c \cdot n\})^{-\alpha} + o(n^{-2}) \\ &= n^{\beta-2} \left(\sum_{i,j \in V(G_n)} (c \cdot n)^{-2\alpha} \mathbf{1}_{\{i,j \leq c \cdot n\}} \right) + o(n^{-2}) + n^{\beta-2} \left(\sum_{i,j \in V(G_n)} (ij)^{-\alpha} \mathbf{1}_{\{i,j > c \cdot n\}} \right) \\ &\quad + 2n^{\beta-2} \left(\sum_{i,j \in V(G_n)} (i \cdot c \cdot n)^{-\alpha} \mathbf{1}_{\{i > c \cdot n \geq j\}} \right) \\ &= c^{-2\alpha} n^{\beta-2-2\alpha} (cn)^2 + o(n^{-2}) + n^{\beta-2} \left(\sum_{i,j \in V(G_n)} (ij)^{-\alpha} \mathbf{1}_{\{i,j > cn\}} \right) + 2c^{-\alpha} \cdot n^{\beta-2-\alpha} cn \left(\sum_{i \in V(G_n)} i^{-\alpha} \mathbf{1}_{\{i > cn\}} \right) \\ &\sim c^{2-2\alpha} n^{\beta-2\alpha} + n^{\beta-2} \left(\int_{c \cdot n}^n i^{-\alpha} di \right) \left(\int_{c \cdot n}^n j^{-\alpha} dj \right) + 2c^{1-\alpha} n^{\beta-1-\alpha} \int_{c \cdot n}^n i^{-\alpha} di \end{aligned}$$

$$\begin{aligned}
 &= c^{2-2\alpha} n^{\beta-2\alpha} + n^{\beta-2} (1-\alpha)^{-2} (n^{1-\alpha} - (cn)^{1-\alpha})^2 \\
 &\quad + 2c^{1-\alpha} n^{\beta-1-\alpha} (1-\alpha)^{-1} (n^{1-\alpha} - (cn)^{1-\alpha}) \\
 &= c^{2-2\alpha} n^{\beta-2\alpha} + n^{\beta-2\alpha} (1-\alpha)^{-2} (1-c^{1-\alpha})^2 + 2c^{1-\alpha} n^{\beta-2\alpha} (1-\alpha)^{-1} (1-c^{1-\alpha}) \\
 &= n^{\beta-2\alpha} c^{2-2\alpha} + n^{\beta-2\alpha} \frac{1-c^{1-\alpha}}{(1-\alpha)^2} [1-c^{1-\alpha} + 2(1-\alpha)c^{1-\alpha}] \\
 &= (1-\alpha)^{-2} n^{\beta-2\alpha} (c^{2-2\alpha} (1-\alpha)^2 + 1 - c^{2-2\alpha} - 2\alpha c^{1-\alpha} (1-c^{1-\alpha})) \\
 &= (1-\alpha)^{-2} n^{\beta-2\alpha} (\alpha^2 c^{2-2\alpha} + 1 - 2\alpha c^{1-\alpha}) \\
 &= (1-\alpha)^{-2} n^{\beta-2\alpha} (1 - \alpha c^{1-\alpha})^2
 \end{aligned}$$

Knowing the expected edge density, we are now able to determine the limiting graphon of interest. The required mathematical framework can be found in Borgs et al. (2018) which also provides the idea for our proof. Therefore, we just give the key definitions and refer to Borgs et al. (2018) for more details. In general, consider a weighted graph G and a partition \mathcal{P} of G into q parts, i.e. $\mathcal{P} := \{V_1, \dots, V_q\}$ where some sets of the partition can be empty. Alternatively, this partition can be expressed in terms of a map $\phi: V(G) \rightarrow [q]$ where $\phi(x) = i$ if and only if $x \in V_i$. This map ϕ gives rise to the definition of the quotient G/ϕ which consists of a pair $(\alpha(G/\phi), \beta(G/\phi))$ where $\alpha(G/\phi) \in \mathbb{R}^q$ and $\beta(G/\phi) \in \mathbb{R}^{q \times q}$. The entries of $\alpha(G/\phi)$ and $\beta(G/\phi)$ are defined as

$$\alpha_i(G/\phi) := \frac{\alpha_{V_i}(G)}{\alpha_G}$$

and

$$\beta_{ij}(G/\phi) := \frac{1}{\|G\|_1} \sum_{(u,v) \in V_i \times V_j} \frac{\alpha_u(G)}{\alpha_G} \frac{\alpha_v(G)}{\alpha_G} \beta_{uv}(G)$$

for all $i, j \in \{1, \dots, q\}$. To obtain analogous concepts for graphons, we define a fractional partition of $[0, 1]$ into q classes as a q -tuple of measurable functions $\rho_1, \dots, \rho_q: [0, 1] \rightarrow [0, 1]$ with $\rho_1(x) + \dots + \rho_q(x) = 1$ for all $x \in [0, 1]$. Then, the pair $(\alpha_i(W/\rho), \beta(W/\rho))$ with entries

$$\begin{aligned}
 \alpha_i(W/\rho) &:= \alpha_i(\rho) := \int_0^1 \rho_i(x) dx \\
 \beta_{ij}(W/\rho) &:= \int_{[0,1]^2} \rho_i(x) \rho_j(y) W(x, y) dx dy
 \end{aligned}$$

defines the quotient W/ρ . With the above definitions in place, we turn to the proof that the graph sequence $(G_n)_{n \in \mathbb{N}}$ converges to the L^p graphon $W(x, y) = (1-\alpha)^2 (1-\alpha c^{1-\alpha})^{-2} (\max\{x, c\} \cdot \max\{y, c\})^{-\alpha}$. For each V_i we define

$$A_i := \bigcup_{x \in V_i} (x-1, x] \quad \text{and} \quad B_i := \bigcup_{x \in V_i} \left(\frac{x-1}{n}, \frac{x}{n} \right]$$

which implies both $A_i \subset [0, n]$ and $B_i \subset [0, 1]$. Turn to the expectation

$$\begin{aligned}
 \mathbb{E}[\beta_{ij}(G_n/\phi)] &\sim \frac{1}{n^2 \mathbb{E}[\|G\|_1]} \sum_{(u,v) \in V_i \times V_j} \min\left(1, n^\beta (\max\{u, c \cdot n\} \cdot \max\{v, c \cdot n\})^{-\alpha}\right) \\
 &\sim \frac{(1-\alpha)^2}{n^{\beta-2\alpha} (1-\alpha c^{1-\alpha})^2} \sum_{(u,v) \in V_i \times V_j} n^{\beta-2} (\max\{u, c \cdot n\} \cdot \max\{v, c \cdot n\})^{-\alpha} \\
 &= (1-\alpha)^2 n^{2\alpha-2} (1-\alpha c^{1-\alpha})^{-2} \sum_{(u,v) \in V_i \times V_j} (\max\{u, c \cdot n\} \cdot \max\{v, c \cdot n\})^{-\alpha} \\
 &\sim \left(\frac{1-\alpha}{1-\alpha c^{1-\alpha}} \right)^2 n^{2\alpha-2} \cdot \int_{[0,n]^2} \mathbf{1}_{\{x \in A_i\}} \mathbf{1}_{\{y \in A_j\}} \\
 &\quad \cdot (\max\{x, c \cdot n\} \cdot \max\{y, c \cdot n\})^{-\alpha} dx dy
 \end{aligned}$$

$$\begin{aligned}
 &= \left(\frac{1-\alpha}{1-\alpha c^{1-\alpha}} \right)^2 n^{-2} \cdot \int_{[0,n]^2} \mathbf{1}_{\{x \in A_i\}} \mathbf{1}_{\{y \in A_j\}} (\max\{x/n, c\} \cdot \max\{y/n, c\})^{-\alpha} dx dy \\
 &= \int_{[0,1]^2} \mathbf{1}_{\{x \in B_i\}} \mathbf{1}_{\{y \in B_j\}} (1-\alpha)^2 (1-\alpha c^{1-\alpha})^{-2} (\max\{x, c\} \cdot \max\{y, c\})^{-\alpha} dx dy \\
 &= \int_{[0,1]^2} \rho_i(x) \rho_j(y) W(x, y) dx dy
 \end{aligned}$$

where

$$W(x, y) = \left(\frac{1-\alpha}{1-\alpha c^{1-\alpha}} \right)^2 (\max\{x, c\} \cdot \max\{y, c\})^{-\alpha}.$$

Eventually, arguments as in (Borgs et al., 2018, Section 3.3.3) yield the desired convergence result.

O PROBLEM DETAILS

In this section, we give more comprehensive descriptions of the problems considered in our work.

O.1 Cyber Security

Let us start with fundamental components of the cyber security model. Computers are either infected or susceptible to infection. Furthermore, they can be defended and are otherwise called unprotected. This leads to a total number of four states, namely *DI* (defended infected), *DS* (defended and susceptible to infection), *UI* (unprotected infected), and *US* (unprotected and susceptible to infection). Formally, we define the state space as $\mathcal{X} := \{DI, DS, UI, US\}$. Each owner of a computer in the network can choose between two actions at each time step, i.e. $\mathcal{U} := \{0, 1\}$. Here, action 0 means that the computer owner is satisfied with its current defense status (defended or unprotected) and does not try to adjust it. Conversely, action 1 captures the fact that the owner is trying to change the current defense status to the opposing one. It is important to note that a geometrically distributed (with parameter $0 < \lambda \leq 1$) waiting time passes before the desired adjustment takes place. This brings us to specifying the network structure of interest. In contrast to Carmona and Delarue (2018a) we assume that the connections in the computer network are characterized by a not-so-dense graph structure. More specifically, we choose sequence of graphs that converge to a power-law L^p -graphon W (see the previous chapters for details) when the number of computers in the network approaches infinity. This gives rise to the usual empirical neighborhood mean field \mathbb{G}_t^α of agent α at time t .¹

Already infected computers have a recovery rate which depends on their protection level, meaning there is one recovery probability q_{rec}^D for defended computers and one for unprotected ones, i.e. q_{rec}^U . Computers can become infected either directly by the hacker or by other infected computers. Formally, we define v_H to be the intensity of the hacker's attacks. Then, the probability for direct infection of a defended computer is $v_H z_{\text{inf}}^D$ while it is $v_H z_{\text{inf}}^U$ for an unprotected one. For the second way of infection through other infected computers, we assume that computers can only infect each other if they share a direct connection in the network. Put differently, the nodes depicting the respective computers in the graph have to be connected by an edge. This concept provides the foundation for the transition probabilities we have not specified yet. If α is an unprotected susceptible computer, it's probability to be infected by an infected and defended computer at time t is $\beta_{DU} \mathbb{G}_t^\alpha \{DI\}$. Similarly, the unprotected susceptible computer α has a probability of $\beta_{UU} \mathbb{G}_t^\alpha \{UI\}$ to be infected by an unprotected and infected computer.

If we assume α to be a defended susceptible computer, the transition probabilities have to be adapted accordingly. The probability for α to be infected by a defended and infected computer at time t is then given by $\beta_{DD} \mathbb{G}_t^\alpha \{DI\}$, and it's probability to be infected by an unprotected infected computer is defined as $\beta_{UD} \mathbb{G}_t^\alpha \{UI\}$. Summing up the probabilities of infection by the hacker and infection by another computer yields the overall probability of being infected, i.e.

$$\begin{aligned}
 q_{\text{inf}}^D &:= v_H z_{\text{inf}}^D + \beta_{DD} \tilde{\mathbb{G}}_t^\alpha \{DI\} + \beta_{UD} \tilde{\mathbb{G}}_t^\alpha \{UI\} - v_H z_{\text{inf}}^D \beta_{DD} \tilde{\mathbb{G}}_t^\alpha \{DI\} - v_H z_{\text{inf}}^D \beta_{UD} \tilde{\mathbb{G}}_t^\alpha \{UI\} \\
 &\quad - \beta_{DD} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UD} \tilde{\mathbb{G}}_t^\alpha \{UI\} + v_H z_{\text{inf}}^D \beta_{DD} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UD} \tilde{\mathbb{G}}_t^\alpha \{UI\}
 \end{aligned}$$

¹Carmona and Delarue (2018a) define the model for a continuous time interval in contrast to our discrete time approach. Consequently, we make some adjustments such as using geometrically distributed waiting times instead of exponentially distributed ones.

and similarly for undefended computers

$$q_{\text{inf}}^U := v_H z_{\text{inf}}^U + \beta_{DU} \tilde{\mathbb{G}}_t^\alpha \{DI\} + \beta_{UU} \tilde{\mathbb{G}}_t^\alpha \{UI\} - v_H z_{\text{inf}}^U \beta_{DU} \tilde{\mathbb{G}}_t^\alpha \{DI\} - v_H z_{\text{inf}}^U \beta_{UU} \tilde{\mathbb{G}}_t^\alpha \{UI\} \\ - \beta_{DU} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UU} \tilde{\mathbb{G}}_t^\alpha \{UI\} + v_H z_{\text{inf}}^U \beta_{DU} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UU} \tilde{\mathbb{G}}_t^\alpha \{UI\}$$

where $\tilde{\mathbb{G}}_t^\alpha := \min(1, \mathbb{G}_t^\alpha)$, and the superscript (D or U) indicates whether the respective computer is defended or unprotected at the current time point. For convenience of notation, we do not explicitly indicate the dependence of q_{inf}^D and q_{inf}^U on α and t . Thus, we obtain the following state-transition matrix which specifies all transition probabilities

$$M(\alpha, t, u) = \begin{array}{c} \\ \\ \\ \\ \end{array} \begin{array}{cccc} DI & DS & UI & US \\ \begin{pmatrix} \bar{\lambda} \bar{q}_{\text{rec}}^D & \bar{\lambda} \bar{q}_{\text{rec}}^D & u \lambda \bar{q}_{\text{rec}}^D & u \lambda \bar{q}_{\text{rec}}^D \\ \bar{\lambda} \bar{q}_{\text{inf}}^D & \bar{\lambda} \bar{q}_{\text{inf}}^D & u \lambda \bar{q}_{\text{inf}}^D & u \lambda \bar{q}_{\text{inf}}^D \\ u \lambda \bar{q}_{\text{rec}}^U & u \lambda \bar{q}_{\text{rec}}^U & \bar{\lambda} \bar{q}_{\text{rec}}^U & \bar{\lambda} \bar{q}_{\text{rec}}^U \\ u \lambda \bar{q}_{\text{inf}}^U & u \lambda \bar{q}_{\text{inf}}^U & \bar{\lambda} \bar{q}_{\text{inf}}^U & \bar{\lambda} \bar{q}_{\text{inf}}^U \end{pmatrix} & & & \end{array} \quad (27)$$

where we have introduced $\bar{\lambda} := 1 - u\lambda$ and $\bar{q} := 1 - q$ for all q with the different respective subscripts and superscripts. By choosing their controls $(u)_{t=0, \dots, T-1} \in \{0, 1\}^T$, computer owners try to maximize their rewards given by the reward function $r(X_t^i, U_t^i, \mathbb{G}_t^i) = -k_D \mathbf{1}_{\{X_t^i \in D\}} - k_I \mathbf{1}_{\{X_t^i \in I\}}$ where $D := \{DI, DS\}$, $I := \{DI, UI\}$, and $k_D \geq 0$ ($k_I \geq 0$) is the cost associated with being defended (infected) for one computer at one time step.

In our experiments, we use the parameters $T = 50$, $\mu_0 = [0.25, 0.25, 0.25, 0.25]$, $q_{\text{rec}}^D = 0.3$, $q_{\text{rec}}^U = 0.2$, $\lambda = 0.3$, $v_H = 0.1$, $z_{\text{inf}}^D = 0.05$, $z_{\text{inf}}^U = 0.1$, $\beta_{DD} = 0.1$, $\beta_{UD} = 0.2$, $\beta_{DU} = 0.7$, $\beta_{UU} = 0.8$, $k_D = 0.7$, $k_I = 2$.

O.2 Heterogeneous Cyber Security

We assume existence of two classes of computers, such as the ones in private ownership (Pri) and the ones bought by corporations for business purposes. Let γ_{Pri} and γ_{Cor} be the fraction of private and corporate computers in the network, respectively, with $\gamma_{\text{Pri}} + \gamma_{\text{Cor}} = 1$. This distinction between classes can bring the model closer to reality for several reasons. First of all, the accessible level of protection as well as the associated costs may vary drastically between private users and corporations. While private computer owners usually have access to average protection measures at affordable price levels, corporations can decide to invest into large-scale protection solutions, raising both the level of protection and the costs. Second, the costs in case of an infection may differ between private and commercial users. For the private user, costs may consist of inconveniences, i.e. the inability to use one's computer properly, or financial losses such as hacked bank accounts. In contrast to that, infected corporate computers may lead to immense financial and economic damages while personal inconveniences are secondary in this case. Third, the infection probabilities might vary between commercial and private computers. This can either be caused by changes in the owner behavior due to the different usage context, i.e. doing work versus casual activities (listening to music, social media, etc.). Or it might be the case that the hacker chooses one class of computer owners as her primary target, e.g. because corporate computers promise higher financial gains.

The implementation of multiple classes into our model is a rather straightforward extension of the basic setup. We implement the two different agent types as part of the state, see also for example Cui and Koepl (2022) or Mondal et al. (2021). Then, the new state space is given by $\mathcal{X} := \mathcal{X}_P \cup \mathcal{X}_C$ with $\mathcal{X}_P := \{\text{Pri}DI, \text{Pri}DS, \text{Pri}UI, \text{Pri}US\}$ and $\mathcal{X}_C := \{\text{Cor}DI, \text{Cor}DS, \text{Cor}UI, \text{Cor}US\}$ while the action space $\mathcal{U} := \{0, 1\}$ remains the same. As the base version of the model, the multi-class setup is also built on the assumption that the computer network structure follows a graph with power-law degree distribution. Coming to the transition probabilities, the multi-class approach requires two versions of the transition matrix in (27).

For various possible reasons, which we have discussed above, the infection and recovery probabilities for commercially used computers do not have to be the same as for privately owned ones. Formalizing this observation, we introduce the recovery probabilities for defended private computers $q_{\text{rec}}^{\text{Pri}D}$ and unprotected private computers $q_{\text{rec}}^{\text{Pri}U}$. These and the following definitions are easily extended to the corporate case. However, we will not write them down explicitly in order to keep this paragraph clear and brief. Similarly to the basic version of the model, a privately owned defended computer is directly infected by the hacker with probability $v_H z_{\text{inf}}^{\text{Pri}D}$ and a private unprotected one is directly infected with probability $v_H z_{\text{inf}}^{\text{Pri}U}$ where v_H denotes the effort of the hacker as before. Besides direct infection, computers can also be infected by already infected, neighboring computers. To account for the potential differences of the two classes of private and corporate computers, the probability for an defended private computer α to be infected by an unprotected neighbor is denoted as

$\beta_{UPriD} \mathbb{G}_t^\alpha \{UI\}$ where \mathbb{G}_t^α denotes the empirical neighborhood mean field of agent α at time t as before. Accordingly, we define the other network infection probabilities by $\beta_{DPriD} \mathbb{G}_t^\alpha \{UI\}$, $\beta_{DPriU} \mathbb{G}_t^\alpha \{UI\}$, and $\beta_{UPriU} \mathbb{G}_t^\alpha \{UI\}$. Combining both the direct infection probabilities and those through the network, we obtain the overall infection probabilities

$$q_{\text{inf}}^{\text{Pri}D} := v_H z_{\text{inf}}^{\text{Pri}D} + \beta_{DPriD} \tilde{\mathbb{G}}_t^\alpha \{DI\} + \beta_{UPriD} \tilde{\mathbb{G}}_t^\alpha \{UI\} - v_H z_{\text{inf}}^{\text{Pri}D} \beta_{DPriD} \tilde{\mathbb{G}}_t^\alpha \{DI\} - v_H z_{\text{inf}}^{\text{Pri}D} \beta_{UPriD} \tilde{\mathbb{G}}_t^\alpha \{UI\} \\ - \beta_{DPriD} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UPriD} \tilde{\mathbb{G}}_t^\alpha \{UI\} + v_H z_{\text{inf}}^{\text{Pri}D} \beta_{DPriD} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UPriD} \tilde{\mathbb{G}}_t^\alpha \{UI\}$$

and

$$q_{\text{inf}}^{\text{Pri}U} := v_H z_{\text{inf}}^{\text{Pri}U} + \beta_{DPriU} \tilde{\mathbb{G}}_t^\alpha \{DI\} + \beta_{UPriU} \tilde{\mathbb{G}}_t^\alpha \{UI\} - v_H z_{\text{inf}}^{\text{Pri}U} \beta_{DPriU} \tilde{\mathbb{G}}_t^\alpha \{DI\} - v_H z_{\text{inf}}^{\text{Pri}U} \beta_{UPriU} \tilde{\mathbb{G}}_t^\alpha \{UI\} \\ - \beta_{DPriU} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UPriU} \tilde{\mathbb{G}}_t^\alpha \{UI\} + v_H z_{\text{inf}}^{\text{Pri}U} \beta_{DPriU} \tilde{\mathbb{G}}_t^\alpha \{DI\} \beta_{UPriU} \tilde{\mathbb{G}}_t^\alpha \{UI\}$$

where $\tilde{\mathbb{G}}_t^\alpha := \min(1, \mathbb{G}_t^\alpha)$. With the above definitions, we provide the state-transition matrix $M_{\text{Pri}}(\alpha, t, u)$ for the 'Pri' class and similarly $M_{\text{Cor}}(\alpha, t, u)$ for the 'Cor' class analogously to the homogeneous problem. Eventually, the state-transition matrix M for the whole state space \mathcal{X} is given by the definition

$$M(\alpha, t, u) := \begin{pmatrix} M_{\text{Pri}}(\alpha, t, u) & \mathbf{0} \\ \mathbf{0} & M_{\text{Cor}}(\alpha, t, u) \end{pmatrix}$$

where each $\mathbf{0}$ is a 4×4 matrix with all entries equal to zero. Since there is no possibility of changing ownership in our setup, the transition probabilities we have not specified yet all turn out to be zero. Computer owners try to maximize their rewards $r_{\text{Pri}}(X_t^i, U_t^i, \mathbb{G}_t^i) = -k_{\text{Pri}D} \mathbf{1}_{\{X_t^i \in D\}} - k_{\text{Pri}I} \mathbf{1}_{\{X_t^i \in I\}}$ or $r_{\text{Cor}}(X_t^i, U_t^i, \mathbb{G}_t^i) = -k_{\text{Cor}D} \mathbf{1}_{\{X_t^i \in D\}} - k_{\text{Cor}I} \mathbf{1}_{\{X_t^i \in I\}}$, depending on the class of ownership. Here, $k_{\text{Pri}D}$ denotes the cost to defend a privately owned computer and $k_{\text{Pri}I}$ the cost associated with a privately owned, infected computer. Note that both costs are measured in terms of one time period and that $k_{\text{Cor}D}$ and $k_{\text{Cor}I}$ are defined analogously.

In our experiments, we use slightly adjusted parameters $T = 50$, $\mu_0 = [0.125, 0.125, \dots, 0.125]$, $q_{\text{rec}}^{\text{Cor}D} = 0.4$, $q_{\text{rec}}^{\text{Cor}U} = 0.3$, $\lambda = 0.3$, $v_H = 0.1$, $z_{\text{inf}}^{\text{Cor}D} = 0.05$, $z_{\text{inf}}^{\text{Cor}U} = 0.1$, $\beta_{D\text{Cor}D} = 0.1$, $\beta_{U\text{Cor}D} = 0.2$, $\beta_{D\text{Cor}U} = 0.7$, $\beta_{U\text{Cor}U} = 0.8$, $k_{\text{Cor}D} = 0.7$, $k_{\text{Cor}I} = 2$, and parameters for the Pri agents $q_{\text{rec}}^{\text{Pri}D} = 0.4$, $q_{\text{rec}}^{\text{Pri}U} = 0.3$, $\lambda = 0.3$, $v_H = 0.1$, $z_{\text{inf}}^{\text{Pri}D} = 0.05$, $z_{\text{inf}}^{\text{Pri}U} = 0.1$, $\beta_{D\text{Pri}D} = 0.2$, $\beta_{U\text{Pri}D} = 0.3$, $\beta_{D\text{Pri}U} = 0.9$, $\beta_{U\text{Pri}U} = 1.0$, $k_{\text{Pri}D} = 0.6$, $k_{\text{Pri}I} = 2$.

O.3 Beach Bar

For implementation of the model described in the main text, we define a reward function $r(X_t^i, U_t^i, \mathbb{G}_t^i) = \frac{2}{|\mathcal{X}|} |B - X_t^i| + \frac{2}{|\mathcal{X}|} |U_t^i| - 3\mathbb{G}_t^i(x)$ and dynamics $X_{t+1}^i = X_t^i + U_t^i + \epsilon_t^i$, where ϵ_t^i is a random noise variable equal to -1 or 1 with probabilities 0.05 each, and 0 otherwise.