

Deep Backward and Galerkin Methods for the Finite State Master Equation

Asaf Cohen

*Department of Mathematics
University of Michigan
Ann Arbor, MI 48104, USA*

AS AFC@UMICH.EDU

Mathieu Laurière

*Shanghai Frontiers Science Center of Artificial Intelligence and Deep Learning
NYU-ECNU Institute of Mathematical Sciences
New York University Shanghai
Shanghai, 200000, China*

MATHIEU.LAURIERE@NYU.EDU

Ethan Zell

*Department of Mathematics
University of Michigan
Ann Arbor, MI 48104, USA*

EZELL@UMICH.EDU

Editor: Maxim Raginsky

Abstract

This paper proposes and analyzes two neural network methods to solve the master equation for finite-state mean field games (MFGs). Solving MFGs provides approximate Nash equilibria for stochastic, differential games with finite but large populations of agents. The master equation is a partial differential equation (PDE) whose solution characterizes MFG equilibria for any possible initial distribution. The first method we propose relies on backward induction in a time component while the second method directly tackles the PDE without discretizing time. For both approaches, we prove two types of results: there exist neural networks that make the algorithms' loss functions arbitrarily small and conversely, if the losses are small, then the neural networks are good approximations of the master equation's solution. We conclude the paper with numerical experiments on benchmark problems from the literature up to dimension 15, and a comparison with solutions computed by a classical method for fixed initial distributions.

Keywords: Mean field game, deep Galerkin, deep backward, master equation, neural network, Nash equilibrium, PDE, stochastic differential game.

1. Introduction

1.1 Introduction to MFGs

In 1950, John Nash introduced the concept of a Nash equilibrium, an idea that describes when players trying to minimize an individual cost have no incentive to change their strategy (Nash, 1951). Nash proved that, under mild conditions, a Nash equilibrium always exists in mixed strategies. The question remained however, even when a Nash equilibrium

was known to exist, whether there is an algorithmic way to solve for it. It turns out that, in general, computing a Nash equilibrium is a very difficult task (Daskalakis et al., 2009).

The origins of limiting models for many-agent games can be traced back to Aumann (1964) and to Schmeidler (1973), both addressing one-shot games. In the realm of stochastic dynamical games, mean field game (MFG) theory was pioneered by Lasry and Lions (2006, 2007), along with an independent development by Huang et al. (2006). MFG paradigms study situations in which a very large number of agents interact in a strategic manner. Specifically, MFGs provide one way to approximate Nash equilibria of large, anonymous, symmetric finite-player games.

In characterizing an equilibrium in MFGs, we introduce a *representative player* who, instead of reacting to n players, responds to a flow of measures. The *MFG equilibrium* is consequently defined by a fixed point, ensuring that the distribution of the dynamics of the representative player under optimality aligns with the flow of measures to which they respond. It is worth noting that Nash equilibria in finite-player games are similarly described using fixed points.

Returning to the motivating question for n -player games, an MFG equilibrium provides an approximate Nash equilibrium in a corresponding finite-player game: if all the n players use the strategy determined by the MFG, each player can be at most ε_n better off by choosing another strategy, and ε_n converges to zero as $n \rightarrow \infty$. This approximation is valid for games in which players are homogeneous and the interactions are symmetric, in the sense that they occur through the players' empirical distribution. While static games are prevalent in game theory for the sake of simplicity, dynamic games are more realistic for applications. In this work, we consider the MFG studied by Bayraktar and Cohen (2018); Cecchin and Pelino (2019); Gomes et al. (2010) which approximates an n -player stochastic differential game with a finite number of states.

Classically, an MFG equilibrium can be described by a coupled forward-backward system of equations, called the *MFG system*. The backward equation is a Bellman-type equation describing the evolution of a representative player's *value function* (the optimal cost), from which the optimal strategy can be deduced. The forward equation describes the evolution of the population's distribution, which coincides with the distribution of one representative player over the state space. We refer to the books by Bensoussan et al. (2013) and Carmona and Delarue (2018) for more details on this approach using partial differential equations (PDEs) or stochastic differential equations (SDEs). In finite-state MFGs as we will consider in this work, the forward-backward system is a system of ordinary differential equations (ODEs); the interested reader is referred to Carmona and Delarue (2018, Chapter 7.2).

The solution to the forward-backward system depends on the initial distribution of the population. The representative player's *value function* and control depend only *implicitly* on the population distribution. This dependence can be expressed *explicitly* through the notion of a decoupling field. The MFG equilibrium can be described by the so-called *master equation*, introduced by Lions (2014). The interested reader is referred to the monograph by Cardaliaguet et al. (2019) for a detailed analysis of the master equation and the question of convergence of finite-player Nash equilibria to MFG equilibria. The solution to this equation is a function of the time, the representative player's state, and the population distribution, which will be denoted respectively by t , x and η . The master equation derives from the MFG system (see Proposition 4 below), whose structure is crucial for our analysis. Besides

fully characterizing the MFG equilibrium for any initial distribution, the master equation is also used to study the convergence of the finite player equilibrium to its MFG counterpart (Cardaliaguet et al., 2019; Delarue et al., 2019, 2020). In this paper, we are interested in the finite state, finite horizon master equation studied by Bayraktar and Cohen (2018); Cecchin and Pelino (2019):

$$\begin{cases} \partial_t U(t, x, \eta) + \sum_{y, z \in [d]} \eta_y D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y U(t, \cdot, \eta)) + \bar{H}(x, \eta, \Delta_x U(t, \cdot, \eta)) = 0, \\ U(T, x, \eta) = g(x, \eta), \quad (t, x, \eta) \in [0, T) \times [d] \times \mathcal{P}([d]). \end{cases} \quad (1)$$

Above, $[d] := \{1, \dots, d\}$ is the state space, T is the terminal time of the game, \bar{H} is the Hamiltonian (defined in Section 2), and g is a terminal condition. The function γ^* denotes the minimizer of the Hamiltonian, which under regularity assumptions is equal to the gradient of the Hamiltonian. The variables are the time t , the state x , and the distribution η , which lies in the $(d-1)$ -dimensional simplex $\mathcal{P}([d])$. We use the notation $\Delta_x U(t, \cdot, \eta) := (U(t, y, \eta) - U(t, x, \eta))_{y \in [d]}$ for each $x \in [d]$. Derivatives on the simplex, denoted by $D_{yz}^\eta U(t, x, \eta)$, can be thought of as usual directional derivatives and will be explicitly defined in Section 2.

Note that (1) is, in general, a nonlinear PDE, for which there is no analytical solution and the question of well-posedness is challenging (for an overview, see the book by Evans (2010)). In fact, even computing numerically an approximate solution is a daunting problem. When $d \leq 3$, traditional methods such as finite difference schemes can be employed. However, the computational cost of these methods becomes prohibitive for dimensions greater than 3. For example, in grid-based methods, the number of points grows exponentially quickly in d . Other methods include Markov chain approximations, see Bayraktar et al. (2017). We refer to Knabner and Angermann (2021) for an overview of classical methods for PDEs and Achdou and Laurière (2020); Laurière (2021) for an overview of classical methods for MFGs. Besides, many applications require $d > 3$. And additionally, Bertucci and Cecchin (2022) recently demonstrated that MFGs with a continuous state space can be approximated by MFGs with finite, but large, state spaces. This approximation relies on the convergence of the solution of the associated master equations as the number of states tends to infinity. Neural networks provide one avenue to mitigate the curse of dimensionality, and this has triggered interest for machine learning-based methods, as we discuss in the next section.

The value function under the unique Nash equilibrium in the n -player game converges to the solution to the master equation up to an error of order $\mathcal{O}(n^{-1})$ (and so do the optimal controls) (Bayraktar and Cohen, 2018; Cecchin and Pelino, 2019). Then, we prove that the two algorithms we propose, the deep backward master equation (DBME) and deep Galerkin master equation (DGME) methods, approximate the master equation which, as discussed in the prior paragraph, is not easy to solve analytically. So while the primary object of interest for this paper is the master equation, our main contribution completes the picture for how to provably approximate the solution to the n -player game.

1.2 Overview of Machine Learning for PDEs

In this work, we propose and rigorously analyze deep learning methods to tackle the master equation (1) for finite-state MFGs beyond what is possible with traditional numerical methods. Our numerical methods build upon the recent development of deep learning ap-

proaches for PDEs. In the past few years, numerous methods have been proposed along these lines, such as the deep BSDE method (E et al., 2017), the deep Galerkin method (Sirignano and Spiliopoulos, 2018), physics informed neural networks (Raissi et al., 2019) and deep backward dynamic programming (Pham et al., 2021). In the context of MFGs and mean field control problems, deep learning methods have been applied using the probabilistic approach (Carmona and Laurière, 2022; Fouque and Zhang, 2020; Germain et al., 2022b) or the analytical approach (Al-Arabi et al., 2022; Cao et al., 2020; Carmona and Laurière, 2021; Ruthotto et al., 2020). The aforementioned works consider neural networks whose inputs are only the individual state and possibly time. A few works have considered neural networks taking as inputs a representation of the mean field. The paper by Germain et al. (2022a) proposed a deep backward method for mean field control problems. As with the first method (called DBME) we present below, it is based on a discrete-time, backward scheme. The rate of convergence in terms of number of particles was studied in Germain et al. (2022c). However, this method solves the Bellman equation for mean field control problems while we solve the master equation for MFGs, which fails to satisfy a comparison principle. Furthermore, Germain et al. (2022a) considered a continuous space model using particle-based approximations while we focus on a finite space model. Work by Dayanikli et al. (2023) solves continuous space mean field control problems using population-dependent controls by directly learning the control function instead of using a backward dynamic programming equation. As for MFG with population-dependent controls, Perrin et al. (2022) proposed a deep reinforcement learning method for finite-state MFGs based on a fictitious play and deep Q-networks. In contrast, our two methods do not rely on reinforcement learning and we provide a detailed analysis of the algorithms' errors. We refer to the overview of Hu and Laurière (2023) for a review of deep learning methods in optimal control and games.

In this work, we will focus on two types of methods. The first approach finds its roots in the pioneering work of E et al. (2017), which exploits the connection between SDEs and PDEs: using Feynman–Kac type formulas, the PDE solution satisfies a backward SDE, that is, they satisfy a terminal condition. Corresponding to the SDE, the PDE has a terminal condition. Such SDEs and PDEs crop up often in mathematical finance, most iconically in the Black–Scholes equation for pricing European options (Black and Scholes, 1973). While the original method of E et al. (2017) consists in replacing the BSDE by a shooting method for a forward SDE, subsequent works exploit the backward structure differently. In particular, deep backward methods were explored in control problems in the nonlinear and fully-nonlinear cases (Huré et al., 2020; Pham et al., 2021). The Deep Backward Dynamic Programming algorithm (DBDP) from Huré et al. (2020) (more specifically their DBDP2 algorithm) uses a sequence of neural networks on a discretization of $[0, T]$. Training is done by starting from the terminal time and, going backward in time, the neural network for a given time step is trained to minimize a loss function which involves the neural network at the following time step.

The second method we propose is the DGME, introduced as the deep Galerkin method by Sirignano and Spiliopoulos (2018). The PDE solution is approximated by a neural network, which is a function of the time and space variables. Training is done by minimizing an empirical loss which captures the residual of the PDE, as well as possible boundary conditions. The loss is computed over points sampled inside the domain and, if needed, on

its boundary. In this way, the method does not require any discretization of space or time. The solution is learned over the whole domain thanks to the generalization capability of neural networks.

For these methods, most theoretical results rely on universal approximation theorems for suitable classes of functions related to the regularity of PDE solutions. In that sense, other classes of approximators could be used instead of neural networks. However, neural networks seem to be a very suitable choice to solve PDEs. Two of the main limitations of neural networks are the lack of explainability and the large amount of data required for training, but these are not issues when solving PDEs with the aforementioned methods. Indeed, explainability is not a priority when approximating PDE solutions, and data points are obtained by sampling over the PDE domain as many times as desired.

1.3 Contributions and Challenges

While the well-posedness of (1) was studied by Bayraktar and Cohen (2018) and Cecchin and Pelino (2019), no analytical solution is known. Therefore we provide two algorithms, we prove their correctness, and numerically solve the master equation (1) for some examples.

1.3.1 DBME

The first method we propose, called deep backward master equation (DBME), is inspired by the DBDP of Huré et al. (2020), although there are key differences. An important disparity lies in the corresponding stochastic model for each setting: the diffusion dynamics found in the underlying DBDP problem are substituted with jump dynamics in the finite state MFG. Essentially, in the DBDP the target function is twice differentiable; this allows the use of correspondence between certain martingale and loss terms that would otherwise be hard to bound. Unfortunately in the case of the master equation, it is not necessarily true that the solution U is twice differentiable. Another crucial distinction lies in the fact that the DBDP is designed to fit Hamilton–Jacobi–Bellman equations originating from optimal control problems. In contrast, the MFG solution manifests as a fixed point of a best response mapping within a control problem and consequently, the master equation lacks a so-called comparison principle, which leads us to design distinct algorithms and proofs of convergence. To establish the convergence, we leverage the transformation of the master equation into the coupled, forward-backward MFG system for a fixed initial distribution. At a technical level, another difference lies in the fact that the DBME minimizes a maximum-based error, while the DBDP minimizes an L^2 -error. All these differences necessitate a different algorithm and a different analysis. To overcome these difficulties, the analysis for the DBME, like that of the DGME, makes use of the MFG system directly. Our approach contrasts with that of the DBDP which uses propagation by SDEs.

1.3.2 DGME

The second algorithm we consider builds upon the DGM of Sirignano and Spiliopoulos (2018) and adapts it to solve the master equation (1), as explored by Laurière (2021, Algorithm 7). That DGM adaptation is recalled here in Algorithm 2, with a modification. As used by Laurière (2021), the DGM minimizes the expectation of an L^2 -error. In our formulation, the DGME minimizes a worst-case loss function formulated as a maximum

that is, in practice, sampled and not a true maximum over the entire state space. While *applying* the DGM to the master equation is not novel, the finite state master equation is a different form of PDE than that studied by Sirignano and Spiliopoulos (2018) and so the convergence theorem (Theorem 7.3) of their work does not apply. Moreover, no proofs of DGM convergence were offered by Laurière (2021).

We prove the convergence of our DGME algorithm to the unique master equation solution in Theorem 20. Our result is analogous to that of Sirignano and Spiliopoulos (2018, Theorem 7.3), but as previously mentioned, since the PDE structure differs, our proof requires entirely different methods. Namely, recall that the master equation can be constructed from the MFG system and our proof relies on the structure of the MFG system. Using the neural network trained by the DGME, we construct an approximate MFG system and use MFG techniques (specifically, MFG duality) to compare the approximate solution and the true solution. As a consequence of our modification and the structure of the master equation, we obtain convergence of the neural network approximation to the true solution in the supremum norm, which is different from and arguably stronger than the L^2 -convergence obtained by Sirignano and Spiliopoulos (2018). Broadly speaking, this choice allows us to deal with the maximum over all states instead of an expectation and hence our analysis is state-agnostic. One more subtle difference is that Sirignano and Spiliopoulos (2018) tackle the case of a neural network with one-hidden layer for simplicity’s sake, while the proofs we offer for the DGME apply to feedforward networks of arbitrary depth. This generality comes from our use of the universal approximation theorem (specifically Hornik et al., 1990, Theorem 3.1).

1.3.3 THE CONVERGENCE RESULTS

For each algorithm, we have two main types of results. The first type of result, found in Theorems 14 and 19, builds on universal approximation properties of neural networks. We show that there exist neural networks that approximate the master equation solution and that when this is the case, the algorithms’ loss functions are small.

The second type of result, found in Theorems 15 and 20, asserts that when the algorithm results in a neural network (or in the case of the DBME, a family of neural networks) with small loss, then the obtained network is in fact close to the true solution in supremum norm. The rate of convergence in each case depends on the loss value and, in the case of the DBME only, the size of the time partition. Note that the empirical loss achievable in practice by the algorithm depends on the depth and width of the neural network, the number of epochs it is trained for, the amount of training data, and so on. This type of bound is morally different from that of Huré et al. (2020), which uses a constant obtained from universal approximation.

1.4 Structure of the Paper

We begin by defining notation in Section 2 before providing a description of the MFG model and recalling the major results concerning the MFG and the master equation in Section 3. In Section 4 we outline the architecture of the neural networks we consider. Then in Sections 5 and 6, we describe and analyze the DBME and DGME, respectively. Finally, in Section 7, we present numerical results for two examples using these algorithms.

2. Notation

We let $[d] := \{1, \dots, d\}$ be the finite set of states. The finite difference notation for a vector $b \in \mathbb{R}^d$ is $\Delta_x b := (b_y - b_x)_{y \in [d]}$, $x \in [d]$. Generally, $x, y, z \in [d]$ unless otherwise stated. A probability measure over $[d]$ is identified with the $(d-1)$ -dimensional simplex in \mathbb{R}^d , which we denote $\mathcal{P}([d])$; by this we mean that:

$$\mathcal{P}([d]) := \left\{ \eta \in \mathbb{R}^d : \sum_{x \in [d]} \eta_x = 1, \quad \eta_x \geq 0 \right\}.$$

Unless otherwise stated, $\eta, \mu \in \mathcal{P}([d])$. For a vector in \mathbb{R}^k , $k \in \mathbb{N}$, $|\cdot|$ denotes the usual Euclidean norm. When A is a finite set, $|A|$ denotes the cardinality of the set. For any Euclidean spaces $\mathcal{E}_1, \dots, \mathcal{E}_{k_1}$ and any measurable function $\varphi : \prod_{j=1}^{k_1} \mathcal{E}_j \rightarrow \mathbb{R}^{k_2}$, $k_1, k_2 \in \mathbb{N}$, we denote by $\|\varphi\|_\infty$ the supremum norm of φ . In the case that some arguments of φ are specified, the essential supremum is taken over only the unspecified arguments. To alleviate the notations, we will denote by C_φ the supremum norm of a bounded function φ , and by $C_{L,\varphi}$ the Lipschitz constant of a Lipschitz function φ .

The measure derivative on the simplex $\mathcal{P}([d])$ is written D^η (where the direction of the derivative is specified below). Since the simplex is finite-dimensional, one can think of D^η as the usual directional derivative; that is, for each $\eta \in \mathcal{P}([d])$ and for each $y \in [d]$, such that $\eta_y > 0$, we define:

$$D_{yz}^\eta \varphi(\eta) := \lim_{h \searrow 0} \frac{\varphi(\eta + e_{yz}h) - \varphi(\eta)}{h}, \quad (2)$$

where e_{yz} is the z -th standard basis vector of \mathbb{R}^d minus the y -th, and $\varphi : \mathcal{P}([d]) \rightarrow \mathbb{R}$ is a measurable function such that the limit exists. Also, $D_y^\eta \varphi := (D_{yz}^\eta \varphi)_{z \in [d]}$ is a vector. Note that by definition $D_{yy}^\eta \varphi(\eta) = 0$. Furthermore, in equation (1) this derivative is multiplied by η_y , so its definition at $\eta_y = 0$ does not affect the expression. Additionally, D_{yy}^η in (1) is multiplied by $\gamma_y^*(y, \Delta_y U(t, \cdot, \eta))$, which does not represent a rate of transition. Alternatively, we could define the derivatives for $y \neq z$ only and use the double sum in (1) as $\sum_{y \in [d]} \sum_{z \in [d], z \neq y}$.

Let $\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ be the set of \mathbb{R} -valued functions defined on $[0, T] \times \mathcal{P}([d])$, where the time derivative is continuous and whose measure derivative D^η is Lipschitz. The space $\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ can be thought of as combining the Banach space of continuously differentiable, real-valued functions $\mathcal{C}^1([0, T])$ with the Hölder space $\mathcal{C}^{1,1}(\mathcal{P}([d]))$ as defined in (Evans, 2010, Section 5.1). Let \mathcal{E} be a Euclidean space and, for a given differentiable function $R : \mathcal{E} \rightarrow \mathbb{R}$ with derivative $D^\mathcal{E} R$, we define the $(1, 1)$ -Hölder seminorm as:

$$[R]_{\mathcal{C}^{1,1}} := \sup_{p \neq q} \left\{ \frac{|D^\mathcal{E} R(p) - D^\mathcal{E} R(q)|}{|p - q|} \right\}.$$

The first “1” in the notation $\mathcal{C}^{1,1}$ refers to the order of the derivative, while the second refers to the modulus of continuity. Then, the $(1, 1)$ -Hölder norm is:

$$\|R\|_{\mathcal{C}^{1,1}} := [R]_{\mathcal{C}^{1,1}} + \|D^\mathcal{E} R\|_\infty + \|R\|_\infty.$$

With this in mind, we endow the space $\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ with the norm:

$$\|g\|_{\mathcal{C}^{1,1}} := \|g\|_{\infty} + \|\partial_t g\|_{\infty} + \|D_1^\eta g\|_{\infty} + \sup_{t \in [0, T]} [g(t, \cdot)]_{\mathcal{C}^{1,1}}, \quad (3)$$

where D_{yz}^η is defined in (2), and recall that, by definition of $\|\cdot\|_{\infty}$, the supremum is taken over all $(t, \eta) \in \mathcal{P}([d]) \times [0, T]$. Note that $\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ is a Banach space with $\|\cdot\|_{\mathcal{C}^{1,1}}$.

A few times in the paper we will mention the space of continuously differentiable functions $\mathcal{C}^1(X, Y)$ where X is either $[t_0, T]$ or $[t_0, T] \times [d]$ with $t_0 \in [0, T]$ and where Y is some subset of euclidean space. We endow this space with the norm:

$$\|g\|_{\mathcal{C}^1(X, Y)} := \|g\|_{\infty} + \|\partial_t g\|_{\infty},$$

making it Banach.

3. The Finite State Master Equation and the MFG

In this section, we continue with structure and the notation adopted by Bayraktar and Cohen (2018); Cecchin and Pelino (2019) with only minor additions. For readability, we recall the notation here. Throughout we fix $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P})$, a filtered probability space in the background satisfying the usual conditions.

3.1 The MFG

We consider a representative player that uses Markovian controls taking values in a set of rates $\mathbb{A} \subseteq \mathbb{R}_+ := [0, +\infty)$. We will consider rates bounded in a compact set; that is, $\mathbb{A} := [\mathbf{a}_l, \mathbf{a}_u]$, with $0 \leq \mathbf{a}_l \leq \mathbf{a}_u < +\infty$. A *Markovian control* is a measurable function $\alpha : \mathbb{R}_+ \times [d] \rightarrow \mathbb{A}_{[d]}^d$ where $\mathbb{A}_{[d]}^d = \cup_{x \in [d]} \mathbb{A}_{-x}^d$, and:

$$\mathbb{A}_{-x}^d := \left\{ a \in \mathbb{R}^d \mid \forall y \neq x, \quad a_y \in \mathbb{A}, \quad a_x = - \sum_{y \neq x} a_y \right\}.$$

Here, for any $y \neq x$, $\alpha_y(t, x) := \alpha(t, x)_y \in \mathbb{A}$ is the rate of transition at time t to move from state x to state y and as usual with continuous-time Markov chains, $\alpha_x(t, x) = - \sum_{y, y \neq x} \alpha_y(t, x)$. We denote by $\mathcal{Q}[\mathbb{A}]$ the set of $d \times d$ transition-rate matrices with rates in \mathbb{A} . We denote by \mathcal{A} the set of measurable mappings $\alpha : [0, T] \rightarrow \mathcal{Q}[\mathbb{A}]$.

Fix $t_0 \in [0, T]$ to be the initial time. The jump dynamics of the representative player's state \mathcal{X} are given by:

$$\mathcal{X}_t = \mathcal{X}_{t_0} + \int_{t_0}^t \int_{\mathbb{A}^d} \sum_{y \in [d]} (y - \mathcal{X}_{s^-}) \mathbf{1}_{\{\xi_y \in (0, \alpha_y(s, \mathcal{X}_{s^-}))\}} \mathcal{N}(ds, d\xi), \quad t \in [t_0, T]. \quad (4)$$

where \mathcal{N} is a Poisson random measure with intensity measure ν given by

$$\nu(E) := \sum_{y \in [d]} \text{Leb}(E \cap \mathbb{A}_y^d), \quad (5)$$

with $\mathbb{A}_y^d := \{a \in \mathbb{A}^d \mid a_x = 0 \text{ for all } x \neq y\}$, where Leb is the Lebesgue measure on \mathbb{R} , and E is any Borel measurable set. We can think of \mathbb{A}_y^d as the y -projection of \mathbb{A}^d , the set of rates for each state transition; thus, \mathbb{A}_y^d is the set of transition rates from state $y \in [d]$.

In order to set up the MFG equilibrium, we describe the cost structure of the game with respect to a representative player, who selects a Markovian control α in order to play against a smooth flow of measures $\mu \in \mathcal{C}^1([t_0, T], \mathcal{P}([d]))$. We emphasize that the representative player's cost includes the flow of measures μ . In this work, we make use of the regularity result from Cecchin and Pelino (2019), which assumes that the *total running cost* is separable in the form $f + F$. The running cost is $f : [d] \times \mathbb{A}_{[d]}^d \rightarrow \mathbb{R}$, a function of the current state and control, while the mean field cost is $F : [d] \times \mathcal{P}([d]) \rightarrow \mathbb{R}$, a function of the state and mean field distribution. Recall that $\alpha_x(t, x)$ does not represent a rate, thus $f(x, a)$ is assumed independent of a_x as a technicality. Given a flow $\mu \in \mathcal{C}^1([t_0, T], \mathcal{P}([d]))$ of distributions, the representative agent's control problem is to minimize:

$$J(t_0, x, (\mu(s))_{s \in [t_0, T]}, \alpha) = \mathbb{E}_{(x, \eta)} \left[g(\mathcal{X}_T, \mu(T)) + \int_{t_0}^T [f(\mathcal{X}_s, \alpha(s, \mathcal{X}_s)) + F(\mathcal{X}_s, \mu(s))] ds \right], \quad (6)$$

over all $\alpha \in \mathcal{A}$ where $(\mathcal{X}_{t_0}, \mu(t_0)) = (x, \eta) \in [d] \times \mathcal{P}([d])$. Given the flow of measure μ and an initial state $\mathcal{X}_{t_0} = x \in [d]$, the distribution (or law) of the representative player, denoted $(\mathcal{L}(\mathcal{X}_t))_{t \in [t_0, T]}$, is a fully deterministic flow of measures on $\mathcal{P}([d])$. Formally, when $(\mathcal{L}(\mathcal{X}_t))_{t \in [t_0, T]} = (\mu(t))_{t \in [t_0, T]}$ and α minimizes the cost:

$$J(t_0, x, (\mu(t))_{t \in [t_0, T]}, \alpha),$$

we say that we have a *mean field game equilibrium* and we refer to this optimal cost as the *value function*. To be mathematically precise, the MFG equilibrium is a fixed point, as described next.

Definition 1 (MFG equilibrium) Fix $t_0 \in [0, T]$ and $\mu_0 \in \mathcal{P}([d])$. A pair $(\hat{\alpha}, \hat{\mu}) \in \mathcal{A} \times \mathcal{C}^1([t_0, T], \mathcal{P}([d]))$ is a MFG equilibrium over the interval $[t_0, T]$ for the initial distribution μ_0 if the following two conditions are satisfied. First, $\hat{\alpha}$ is an optimal control for $\mathbb{E}_{x \sim \mu_0} [J(t_0, x, \hat{\mu}, \hat{\alpha})]$ and second, for every $t \in [t_0, T]$, $\hat{\mu}(t) = \mathcal{L}(\mathcal{X}_t^{\hat{\alpha}})$, where $\mathcal{X}_t^{\hat{\alpha}}$ solves (4) controlled by $\hat{\alpha}$.

Crucially, notice that the equilibrium is defined for a fixed initial distribution μ_0 . For different initial distributions, we generally obtain different equilibrium controls and mean field flows. The master equation characterizes MFG equilibria for *any* initial distribution.

As explained below in Section 3.3, the equilibrium can be characterized using optimality conditions which involve the Hamiltonian, defined as:

$$\bar{H}(x, \eta, b) := \min_{a \in \mathbb{A}_{-x}^d} \left\{ f(x, a) + F(x, \eta) + \sum_{y \neq x} a_y b_y \right\} = H(x, b) + F(x, \eta),$$

where,

$$H(x, b) := \min_{a \in \mathbb{A}_{-x}^d} \left\{ f(x, a) + \sum_{y \neq x} a_y b_y \right\}.$$

At times, using \bar{H} rather than $H + F$ will simplify the presentation of the PDE.

Remark 2 The separability of the cost and the Hamiltonian is a standard assumption in the MFG literature. This, combined with the monotonicity assumption on functions F and

G , leads to the uniqueness of the MFG solution and the corresponding master equation solution, see e.g., Cardaliaguet et al. (2019); Bayraktar and Cohen (2018); Cecchin and Pelino (2019) as well as (Carmona and Delarue, 2018, P. 652). It is important to note that while these are the prevalent assumptions, there are alternative conditions in the literature, such as alternative monotone conditions, see e.g., Gangbo and Mészáros (2022); Gangbo et al. (2022); Graber and Mészáros (2023), that imply the uniqueness of solutions to MFG systems and the well-posedness of the corresponding master equations; or the anti-monotonicity condition, which can also establish global well-posedness for mean field game master equations with nonseparable Hamiltonians, see Mou and Zhang (2022).

We emphasize that the goal of this paper is to leverage the established results from previous works to solve the master equation, rather than to propose new conditions for achieving uniqueness in the solution of the master equation.

3.2 Assumptions

In this section, we list the assumptions made on our MFG model. Such assumptions are standard and appear in previous work by Bayraktar and Cohen (2018); Cecchin and Pelino (2019); Cohen and Zell (2023).

Recall that the action set is $\mathbb{A} := [\mathbf{a}_l, \mathbf{a}_u]$ where $0 \leq \mathbf{a}_l \leq \mathbf{a}_u < \infty$.¹

Assumption 1 *The Hamiltonian has a unique minimizer to which we refer to as the optimal rate selector and is denoted:*

$$\gamma^*(x, b) := \arg \min_{a \in \mathbb{A}_{-x}^d} \{f(x, a) + a \cdot b\}. \quad (7)$$

The optimal rate selector γ^* is a measurable function.

We note that since γ^* is given by (7), one sufficient condition to guarantee this is when f is strictly convex in a .

As for the mean field cost F , and the terminal cost g , we assume

Assumption 2 *The functions F and g are continuously differentiable in η with Lipschitz derivative; namely, for any $x \in [d]$ $F(x, \cdot) \in \mathcal{C}^1(\mathcal{P}([d]))$ and $D_{1z}^\eta F(x, \cdot)$ is Lipschitz. Moreover, F and g are Lasry–Lions monotone, meaning: for both $\phi = F, g$,*

$$\sum_{x \in [d]} (\phi(x, \eta) - \phi(x, \hat{\eta}))(\eta_x - \hat{\eta}_x) \geq 0, \quad \forall \eta, \hat{\eta} \in \mathcal{P}([d]). \quad (8)$$

Intuitively, the Lasry–Lions monotonicity condition means that a representative player would prefer to avoid congestion to decrease their cost.

In the following assumption, we specify the regularity of H , which, since it is defined through f implicitly defines most of the necessary regularity for f . Moreover, since F is defined on a compact set, its regularity assumptions imply it is bounded.

1. In the papers by Bayraktar and Cohen (2018); Cohen and Zell (2023) it is assumed that $\mathbf{a}_l > 0$; however in his thesis, Cecchin managed to relax the assumption to allow for $\mathbf{a}_l = 0$ (Cecchin and Pelino, 2019).

Assumption 3 *Assume f is bounded. We let*

$$W := \sqrt{2d}[T(C_f + C_F) + C_g] + 1, \quad (9)$$

and we assume that, on the compact set $[-W, W]$, the derivatives $D_{pp}^2 H$ and $D_p H$ exist and are Lipschitz in p . Moreover, there exists a positive constant $C_{2,H}$ such that:

$$D_{pp}^2 H(x, p) \leq -C_{2,H}. \quad (10)$$

Note that when H is differentiable, Gomes et al. (2013, Proposition 1) proved:

$$\gamma^*(x, p) = D_p H(x, p). \quad (11)$$

In order to keep this level of generality for the regularity of H (namely, that these additional assumptions hold on $[-W, W]$ but not necessarily globally), we will on occasion prove that a particular p -argument of H is uniformly bounded by W . For the purpose of this paper, the assumption on W in (9) is not too restrictive. Computations with the Hamiltonian only appear with the argument $\Delta_x U$ or $\Delta_x \mathcal{U}$, where \mathcal{U} is a neural network specified in the following sections that approximates U . As we explain in detail in Remark 7, both of these arguments are bounded by W from (9) and hence the Hamiltonian will always be regular enough when it is required.

Remark 3 *From the above assumption on H , we have that γ^* is locally Lipschitz. Note that a sufficient condition is that f is uniformly convex in a .*

Throughout the paper C denotes a generic, positive constant that depends only on the problem's parameters (that is, the parameters introduced above). To alleviate excessive notation, C 's value may change from one line to another.

3.3 Known Results

Next up, we recall some established results from the study of finite-state, finite-horizon MFGs and master equations.

The functions $u^{t_0, \eta}$ and $\mu^{t_0, \eta}$ are defined in a moment but we first describe their meanings. The measure $\mu^{t_0, \eta}$ is the evolving MFG equilibrium as in Definition 1. The value function $u^{t_0, \eta}$ is the value function (optimal cost) of the MFG starting at time t_0 with initial distribution $\eta \in \mathcal{P}([d])$ along the MFG equilibrium. So, the value $u^{t, \eta}(t, x)$ is the remaining optimal cost to a player in state x at time $t \geq t_0$ until the game ends, at time T .

The MFG system is:

$$\begin{aligned} \frac{d}{dt} u^{t_0, \eta}(t, x) + \bar{H}(x, \mu^{t_0, \eta}(t), \Delta_x u^{t_0, \eta}(t, \cdot)) &= 0, & (t, x) \in [t_0, T] \times [d], \\ \frac{d}{dt} \mu^{t_0, \eta}(t, x) &= \sum_{y \in [d]} \mu^{t_0, \eta}(t, y) \gamma_x^*(y, \Delta_y u^{t_0, \eta}(t, \cdot)), & (t, x) \in [t_0, T] \times [d], \\ \mu^{t_0, \eta}(t_0, x) &= \eta(x), & x \in [d], \\ u^{t_0, \eta}(T, x) &= g(x, \mu^{t_0, \eta}(T)), & x \in [d]. \end{aligned} \quad (12)$$

Note that the system is composed of a forward equation and a backward equation. The above ODE system characterizes the MFG equilibrium for a given initial condition (t_0, η) . However, our goal is to compute the MFG equilibrium for every initial condition. This would require solving a continuum of ODE systems, which is not feasible. Instead, we will focus on the master equation. Informally, it encompasses the continuum of ODE systems and its solution captures the dependence of the value function on the mean field. More precisely, the master equation's solution U is defined through its MFG system. The following proposition is a combination of the results from Cecchin and Pelino (2019, Proposition 1, Proposition 5, Theorem 6)² and (Cardaliaguet et al., 2019, Section 1.2.4).

Proposition 4 *There exists a unique solution, denoted by $(u^{t_0, \eta}, \mu^{t_0, \eta})$, in $\mathcal{C}^1([t_0, T] \times [d], \mathbb{R}) \times \mathcal{C}^1([t_0, T] \times [d], \mathcal{P}([d]))$ to (12). Let U be defined by:*

$$U(t_0, x, \eta) := u^{t_0, \eta}(t_0, x). \quad (13)$$

Then U is the unique classical solution to the master equation (1). Moreover, the consistency relation holds: for all $t_0 \in [0, T]$,

$$U(t, x, \mu^{t_0, \eta}(t)) = u^{t, \mu^{t_0, \eta}(t)}(t) = u^{t_0, \eta}(t), \quad (t, x, \eta) \in [t_0, T] \times [d] \times \mathcal{P}([d]). \quad (14)$$

Additionally, $U(\cdot, x, \cdot) \in \mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ for every $x \in [d]$. Moreover,

$$\begin{aligned} U(t_0, x, \eta) &= J(t_0, x, (\mu^{t_0, \eta}(s))_{s \in [t_0, T]}, (\gamma^*(\cdot, \Delta_x U(s, \cdot, \mu^{t_0, \eta}(s))))_{s \in [t_0, T]}) \\ &= \mathbb{E}_{(x, \eta)} \left[g(\mathcal{X}_T, \mu^{t_0, \eta}(T)) \right. \\ &\quad \left. + \int_{t_0}^T [f(\mathcal{X}_s, \gamma^*(\mathcal{X}_s, \Delta_{\mathcal{X}_s} U(s, \cdot, \mu^{t_0, \eta}(s)))) + F(\mathcal{X}_s, \mu^{t_0, \eta}(s))] ds \right], \end{aligned} \quad (15)$$

where the second line is by definition of J , see (6). As a consequence of (15),

$$|U(t_0, x, \eta)| \leq C_g + T(C_f + C_F) =: \tilde{T}. \quad (16)$$

Recalling the definition of W in (9), note that $W = \sqrt{2d\tilde{T}} + 1$.

Remark 5 By (15), we can think of $U(t_0, x, \eta)$ as the value of the MFG equilibrium when the initial time is t_0 , the initial state is x , and the initial distribution is η . Moreover, $(\gamma_y^*(x, \Delta_x U(s, \cdot, \mu^{t_0, \eta}(s))))_{x, y \in [d]}$ is the rate matrix under the mean field equilibrium; namely, these are the rates that dictate how quickly the representative player transitions from one state to another under this equilibrium. By (14), we have $\gamma_y^*(x, \Delta_x U(s, \cdot, \mu^{t_0, \eta}(s))) = \gamma_y^*(x, \Delta_x u^{t_0, \eta}(s, \cdot))$, which are precisely the transition rates in the Kolmogorov equation for $\mu^{t_0, \eta}(t)$ from (12).

Remark 6 Recall that the motivation for investigating the master equation in the first place was to solve the corresponding n -player game. Per Bayraktar and Cohen (2018, Theorem 2.1) or Cecchin and Pelino (2019, Theorem 1), the n -player game's value function (cost under the unique Nash equilibrium) converges to the master equation's solution with rate n^{-1} . Moreover by Cecchin and Pelino (2019, Theorem 2, Theorem 4), the n -player Nash equilibrium results in an empirical distribution of players whose trajectory converges to the MFG equilibrium, with rate depending on the notion of convergence.

2. In fact, we reuse the argument from Proposition 1, but applied to the MFG system.

4. Neural Networks

With the major points from the MFG literature in hand, we can switch to the neural network side. We define the neural networks to be used as function approximators and outline their parameters. From a proof perspective, the DBME and DGME algorithms require networks with slightly different levels of regularity, so here we only describe what features the networks used for each algorithm have in common.

Throughout the paper, the use of neural networks is justified by the Universal Approximation Theorem, specifically the version from Hornik et al. (1990, Theorem 3.1). This version of the Universal Approximation Theorem asserts that any differentiable function defined on a compact domain in Euclidean space is arbitrarily well-approximated (and has derivative arbitrarily well-approximated) by a neural network (and its derivative). This is relevant in our case since we are concerned with the master equation solution, $U : [0, T] \times [d] \times \mathcal{P}([d]) \rightarrow \mathbb{R}$, and we may identify $\mathcal{P}([d])$ with a compact subset of \mathbb{R}^{d-1} .

We will denote the total number of dense hidden layers of a deep, feed-forward, fully-connected neural network as L , with the number of parameters in layer ℓ as δ_ℓ for all $\ell \leq L$. By convention, the first layer will always be the input layer. For any such neural network, define $\bar{\delta}$ as the number of parameters. Throughout, the parameters are represented by a trainable vector $\theta \in \mathbb{R}^{\bar{\delta}}$.

Depending on the algorithm at hand, the input dimension of the neural networks will differ. For the DBME method, the input dimension is $d + 1$ since the time component is discretized as part of the algorithm. In the DGME, the input vector is $(t, x, \eta) \in [0, T] \times [d] \times \mathcal{P}([d])$, which is of dimension $d + 2$; so δ_1 is either $d + 4$ or $d + 3$ when accounting for the affine transformation parameters in the neural network. We will give a concrete example later in (20). In either case, the output dimension is 1.

For simplicity, the activation function will be the same for all layers. We will denote it by ϕ . So that necessary theoretical results are accessible, ϕ must be smooth and nonlinear so the neural networks we consider are smooth universal approximators; for example, the hyperbolic tangent function or the sigmoid function are good choices. We use \mathcal{NN} to denote the set of all such neural networks using activation function ϕ . We denote the subset:

$$\mathcal{NN}(C_0, C_1) := \{\mathcal{U} \in \mathcal{NN} : |\mathcal{U}| \leq C_0, |D^n \mathcal{U}| \leq C_1\}. \quad (17)$$

Note that what we mean by a neural network is not simply an architecture but an architecture with fixed values for the parameters. The constraints in (17) are constraints on the parameters of the neural network.

The restriction in (17) on $|D^n \mathcal{U}|$ implies that we are interested in Lipschitz neural networks. Recent machine learning research emphasizes that Lipschitz neural networks ensure robustness of a solution to new or unforeseen data (Fazlyab et al., 2019; Gouk et al., 2021; Pauli et al., 2021; Tseng, 1998).

Using the version provided by Hornik et al. (1990, Theorem 3.1), the Universal Approximation Theorem asserts that \mathcal{NN} is dense in $\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ with the norm $\|\cdot\|_{\mathcal{C}^{1,1}}$ defined in (3). We also define:

$$\mathcal{C}^{1,1}([d] \times \mathcal{P}([d]); C_0, C_1) := \{\mathcal{V} \in \mathcal{C}^{1,1}([d] \times \mathcal{P}([d])) : |\mathcal{V}| \leq C_0, |D^n \mathcal{V}| \leq C_1\}. \quad (18)$$

Since we are only concerned with the master equation's solution, which is a Lipschitz function with Lipschitz derivatives (see Proposition 4), it will be appropriate at times to deal with this restricted set.

5. The DBME Algorithm and Convergence Results

Having given the necessary background, we now present a deep-learning-based backward algorithm for learning the master equation termed Deep Backward Master Equation (DBME).

5.1 The DBME Algorithm 1

First, we introduce some notation relevant to the algorithm. Then, we justify some assumptions made on the neural networks under consideration, and introduce an auxiliary Kolmogorov forward ODE, before introducing the DBME algorithm itself.

The time interval $[0, T]$ will be partitioned into $\pi := \{t_0 = 0, t_1, \dots, t_{N-1}, t_N = T\}$, with increments denoted $\Delta t_i := t_{i+1} - t_i$. Define

$$|\pi| := \max_{i \in [N]} \Delta t_i.$$

Throughout we assume that there exist $\check{c}, \hat{c} > 0$ such that:

$$\frac{\check{c}}{N} \leq |\pi| \leq \frac{\hat{c}}{N}; \tag{19}$$

that is, $|\pi|^{-1}$ is roughly proportional to N . For functions and random variables at time $t_i \in \pi$, we will often use only a subscript i as opposed to t_i for brevity.

For the time $t_i \in \pi$, we approximate the master equation's solution $U(t_i, x, \eta)$ by $\mathcal{U}_i \in \mathcal{NN}$ with input dimension $\delta_1 = d + 2$. Its parameters are denoted by θ^i and we use the notation $\mathcal{U}_i(x, \eta; \theta^i)$. We will sometimes use the shorthand notation $\hat{\mathcal{U}}_i(x, \eta) := \mathcal{U}_i(x, \eta; \hat{\theta}^i)$, where the vector of parameters $\hat{\theta}^i$ optimizes a loss function defined in the sequel. We will not use a neural network at the last time in the partition and simply set $\hat{\mathcal{U}}_N(x, \eta) := g(x, \eta)$.

Let $L = 2$ for instance and $\bar{\delta} = (\delta_2(d + 3))$. For the vector of parameters $\theta \in \mathbb{R}^{\bar{\delta}}$, denote the parameters from the scaling inside the activation function by $\check{\theta}$. We can write a neural network $\mathcal{U}_i(x, \eta; \theta) \in \mathcal{NN}$ as:

$$\mathcal{U}_i(x, \eta; \theta) := \sum_{k \in [\delta_2]} b_k \phi(a^{0,k} x + a^k \cdot \eta + c^k), \tag{20}$$

where $a^k := (a^{y,k})_{y \in [d]} \in \mathbb{R}^d$, $a^{0,k} \in \mathbb{R}$, $b \in \mathbb{R}^{\delta_2}$, $c^k \in \mathbb{R}$ for all $k = 1, \dots, \delta_2$, and these parameters make up the entries in $\theta \in \mathbb{R}^{\bar{\delta}}$ and $\check{\theta} := (a^k)_{k=1}^{\delta_2}$. In the case $L = 2$, we note that $\check{\theta}$ is a $(\delta_2 \times 2)$ -matrix with real entries and denote $\|\cdot\|_2$ as the matrix operator 2-norm. Per Gouk et al. (2021, Section 4), we know that when $\|\check{\theta}\|_2 < K$, some positive real number, the Lipschitz constant of the network with parameters $\theta \in \mathbb{R}^{\bar{\delta}}$ is no more than K (possibly after modifying K to account for the Lipschitz constant of the activation function ϕ). So, for the rest of this section, which focuses on the DBME algorithm, we define the following set of parameters:

$$\Theta := \{\theta \in \mathbb{R}^{\bar{\delta}} \mid \|\check{\theta}\|_2 \leq 2C_{L,U}\}, \tag{21}$$

where recall that $C_{L,U}$ is the Lipschitz constant of U depending only on the problem parameters that can be computed from Cecchin and Pelino (2019, Section 3). Thus any network $\mathcal{U}(\cdot, \cdot; \theta)$ where $\theta \in \Theta$ is Lipschitz with constant no more than $2C_{L,U}$.

Remark 7 *Previously in (17), we introduced the class of neural networks $\mathcal{NN}(C_0, C_1)$ as appropriate for approximating U . In this remark, we elaborate further. By (16), the master equation solution is uniformly bounded by the constant \tilde{T} , given in (16). So, we are justified in truncating the neural network \mathcal{U}_i by replacing \mathcal{U}_i with $\mathcal{U}_i \wedge \tilde{T}$. This ensures that \mathcal{U}_i will always be uniformly bounded. Using a result of Cohen and Zell (2023, Lemma 3.1) and the boundedness of U implies:*

$$|\Delta_x U(t, x, \eta)| \leq \sqrt{2d} \max_{y \in [d]} |U(t, y, \eta)| \leq \sqrt{2d} \tilde{T}.$$

Because of the restriction of \mathcal{U}_i to $\mathcal{U}_i \wedge \tilde{T}$, the same kind of computation holds for \mathcal{U}_i as in the previous display. Consequently, (9) holds and the local regularity assumptions on the Hamiltonian H given in (10) are in force with the neural network as an argument; namely the derivatives $D_p H$ and $D_{pp}^2 H$ are Lipschitz and there exists a positive constant $C > 0$ such that:

$$D_{pp}^2 H(x, \Delta_x \mathcal{U}_i(\cdot, \eta)) \leq -C,$$

for all $(x, \eta) \in [d] \times \mathcal{P}([d])$ and all $i \in [N]$. The local regularity assumptions will be useful in the proof of algorithmic convergence. The upshot of this remark is that for any $\theta^i \in \Theta$,

$$\mathcal{U}_i(\cdot, \cdot; \theta^i) \wedge \tilde{T} \in \mathcal{NN}(\tilde{T}, 2C_{L,U}).$$

One final element to introduce before presenting the algorithm is the solution of a forward Kolmogorov's equation. This equation mirrors the flow of measures μ in the MFG equilibrium. However, instead of utilizing the master equation's solution U in the dynamics, we employ a neural network. This solution is a key element in the algorithm's loss function. In details: for any $\kappa \in \mathcal{P}([d])$, any $i \in \{0, \dots, N-1\}$, and any $\theta \in \Theta$, define $M_i^\theta(\kappa) := \rho_i^{\theta, \kappa}(t_{i+1})$, where $\rho_i^\theta(t) := (\rho_i^\theta(t, x))_{x \in [d]}$ uniquely solves the (forward) ODE:

$$\begin{aligned} \frac{d}{dt} \rho_i^{\theta, \kappa}(t, x) &= \sum_{y \in [d]} \rho_i^{\theta, \kappa}(t, y) \gamma_x^*(y, \Delta_y \mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(t); \theta)), & (t, x) \in [t_i, t_{i+1}] \times [d], \\ \rho_i^{\theta, \kappa}(t_i) &= \kappa. \end{aligned} \tag{22}$$

The previous display is Kolmogorov's forward equation over the time interval $[t_i, t_{i+1}]$ with initial distribution $\kappa \in \mathcal{P}([d])$, and which corresponds to a non-linear Markov chain with rate matrix $\gamma_x^*(y, \Delta_y \mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(t); \theta))$. Note that the control's input is the difference $\Delta_y \mathcal{U}_i$ of the neural network \mathcal{U}_i , rather than that of U ; in light of (14), this is similar to the MFG system (12). Moreover, $\rho_i^{\theta, \kappa}$ appears within the rate matrix, which makes the Markov chain non-linear. The equation (22) is well-posed because the neural network and γ^* are Lipschitz. The notation $M_i^\theta(\kappa)$ can be thought of as the terminal distribution of this non-linear Markov chain at time t_{i+1} ; that is, the effect of propagating the initial distribution κ by a control using neural network values, parameterized by $\theta \in \Theta$.

The next remark motivates the algorithm's structure and the form of the loss function (23) in the algorithm.

Algorithm 1 DBME

- 1: **Input:** A vector of initial parameters $\theta := (\theta^i)_{i=0}^{N-1}$.
- 2: **Output:** A grid of neural networks $(\mathcal{U}_i)_{i=0,\dots,N}$ approximating the solution to (1) on π .
- 3: $\hat{\mathcal{U}}_N \leftarrow g$
- 4: $i \leftarrow N - 1$
- 5: \mathcal{U}_i initialized via θ^i .
- 6: **for** i from $N - 1$ to 0 **do**
- 7: Recalling (22), compute: $\hat{\theta}^i \in \arg \min_{\theta^i \in \Theta} L_i(\theta^i)$
- 8: where

$$L_i(\theta^i) := \max_{(x,\kappa) \in [d] \times \mathcal{P}([d])} \left| \hat{\mathcal{U}}_{i+1}(x, M_i^{\theta^i}(\kappa)) - \mathcal{U}_i(x, \kappa; \theta^i) + (\Delta t_i) \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i(\cdot, \kappa; \theta^i)) \right| \quad (23)$$

- 9: $\hat{\mathcal{U}}_i(\cdot, \cdot) \leftarrow \mathcal{U}_i(\cdot, \cdot; \hat{\theta}^i) \wedge \tilde{T}$
 - 10: **end for**
-

Remark 8 From (12) and (13)–(14) it follows that:

$$U(t_{i+1}, x, \mu^{t_i, \kappa}(t_{i+1})) - U(t_i, x, \kappa) + (t_{i+1} - t_i) \bar{H}(x, \kappa, \Delta_x U(t_i, \cdot, \kappa)) \approx 0$$

when $\Delta t_i = t_{i+1} - t_i$ is small. Recall that the dynamics of $(\mu^{t_i, \kappa}(t))_{t \in [t_i, t_{i+1}]}$, the flow of measures under equilibrium from (12), depends on the master equation solution, making it unsuitable for the loss function. To address this, in the loss function, we replace $\mu^{t_i, \kappa}(t_{i+1})$ by $M_i^{\theta^i}(\kappa)$, defined through equation (22), which is similar to the equation for $\mu^{t_i, \kappa}(t)$ in (12) but with the neural network appearing in its dynamics instead of the function U . Notice that, while the loss uses the forward ODE (22), it does not explicitly make use of the backward ODE. This is because the master equation captures the value function, see (14).

Next, note that the algorithm goes backward in time. The value at terminal time T is known. Then at time t_i , assuming that we have an estimation for $U(t_{i+1}, x, \eta)$ for any $(x, \eta) \in [d] \times \mathcal{P}([d])$, we go backward in time. In the analysis to follow, we use the structure of the loss function and (22) to define a forward-backward system that resembles the MFG system (12) up to a small error that emerges from the loss function. It is the comparison between these two systems that yields the convergence result of Theorem 15.

5.2 Remarks for the Practitioner

We now make a few remarks on the algorithm that are relevant for its implementation.

Remark 9 The reader may wonder how to ensure that the vector of parameters θ^i obtained at each step in Algorithm 1 lies in the set Θ defined in (21). One solution proposed by Gouk et al. (2021, Algorithm 2) is to project the weights at each step of stochastic gradient descent. Projection guarantees $\theta^i \in \Theta$ and that the network obtained is Lipschitz. However, given a large enough upper bound $kC_{L,U}$, where previously in (20) we chose $k = 2$, the projection will not occur in practice and the algorithm can essentially run normal SGD. The result from Gouk et al. (2021) is preferred here because of its theoretical guarantees; however

many other methods may be used to improve the regularity of a neural network, for instance dropout layers, learning rate schedulers, random initialization, and so on.

Remark 10 *In practice, we cannot compute a maximum over infinitely many elements $\kappa \in \mathcal{P}([d])$ as we write in Algorithm 1. Therefore in the implementation, we uniformly generate samples κ in $\mathcal{P}([d])$ and compute the propagation $M_i^\theta(\kappa)$ for each before minimizing the loss by SGD. In what follows we remark on the order of this error due to the approximation. Denote the set of all samples by \mathcal{K} . Define:*

$$\tilde{L}_i(\kappa; \theta^i) := \max_{x \in [d]} |\hat{\mathcal{U}}_{i+1}(x, M_i^\theta(\kappa)) - \mathcal{U}_i(x, \kappa; \theta^i) + (\Delta t_i) \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i(\cdot, \kappa; \theta^i))|,$$

so that \tilde{L}_i is related to the algorithm's loss L_i as:

$$L_i(\theta^i) = \max_{\kappa \in \mathcal{P}([d])} \tilde{L}_i(\kappa; \theta^i).$$

For any $i = 0, \dots, N$, and θ^i , choose the following:

$$\hat{\kappa} \in \arg \max_{\kappa \in \mathcal{K}} \tilde{L}_i(\kappa; \theta^i).$$

Lemma 11 (Approximation of Sampled Loss to Theoretical) *With $\hat{\kappa}$, the sampled maximum, as defined above, denote the cardinality of \mathcal{K} by $|\mathcal{K}|$. Then, there exists a constant $C_d > 0$ depending on the parameters of the problem and the dimension d such that:*

$$\mathbb{E}|L_i(\theta^i) - \tilde{L}_i(\hat{\kappa}; \theta^i)| \leq C_d |\mathcal{K}|^{-1/(d-1)}.$$

The proof, together with an explicit expression for C_d , can be found in the appendix.

Remark 12 *Note that the maximum is in general not smooth and yet we aim to eventually take the gradient of the loss function which requires the loss to be differentiable. In practice, the maximum is computed as a smooth maximum; that is, a smooth function that approximates a maximum. For prior work applying stochastic gradient descent to a maximized loss function, and comments on its robustness, see Shalev-Shwartz and Wexler (2016).*

Remark 13 *It is important to note that computing \bar{H} explicitly is only possible when the Hamiltonian H is computable or at least approximable. This is true for instance in both numerical examples that we consider later in Section 7. In greater generality, when the running cost function is convex in its control argument, this amounts to solving a convex optimization problem at each time step.*

5.3 DBME Main Results

Recalling the definition of $\hat{\theta}^i$ from (23), define:

$$\varepsilon_i := L_i(\hat{\theta}^i) \quad \text{and} \quad \varepsilon^{\mathcal{U}} := \max_{i \in [N]} \varepsilon_i. \tag{24}$$

Note that $\varepsilon_N = 0$ because in the algorithm, we set $\mathcal{U}_N := g$. Nonetheless we include ε_N in (24) in order to use the notation $[N]$ instead of $[N - 1]$.

We can now state our two main results with respect to the DBME. Each result is, loosely speaking, converse to the other. The first theorem says that there exist neural networks close to the true solution in our desired class, and when a neural network is close to the true function, the loss from the algorithm is small.

We recall that \tilde{T} is defined in (16) and that C_g, C_f and C_F are bounds on the maximum norms of g, f and F respectively, and $C_{L,U}$ is a bound on the Lipschitz constant of U .

Theorem 14 (DBME Approximation) *Set:*

$$C_0 := 2\tilde{T} \quad \text{and} \quad C_1 := 2C_{L,U}. \quad (25)$$

For every $\varepsilon > 0$, there exist neural networks $(\mathcal{U}_i^\varepsilon)_{i \in [N]} \subset \mathcal{NN}(C_0, C_1)$ with parameters $(\theta^{i,\varepsilon})_{i \in [N]}$ such that for all $(i, x, \kappa) \in [N] \times [d] \times \mathcal{P}([d])$:

$$|U(t_i, x, \kappa) - \mathcal{U}_i^\varepsilon(x, \kappa; \theta^{i,\varepsilon})| < \varepsilon. \quad (26)$$

Moreover, for any neural network satisfying (26), the DBME algorithm's loss is bounded as:

$$\max_{i \in [N]} L_i(\theta^i) < 2\varepsilon + C|\pi|, \quad (27)$$

where we recall that C is a constant depending only on the model parameters.

The second theorem provides an upper bound on the difference between the neural networks and the true master equation solution in terms of the algorithm's error. Intuitively, it says that any network trained by the algorithm that minimizes the loss is in fact close to the true master equation solution.

Theorem 15 (DBME Convergence) *There exist $N_0 > 0$ and $C > 0$ depending only on the problem's data with the following property. For every $N > N_0$ and for every $i = 0, \dots, N-1$, letting $\hat{\mathcal{U}}_i = \mathcal{U}_i(\cdot, \cdot; \hat{\theta}^i)$ denote a neural network with parameters $\hat{\theta}^i$ minimizing the loss in the DBME Algorithm 1, we have:*

$$\max_{(i,x,\eta) \in [N] \times [d] \times \mathcal{P}([d])} |U(t_i, x, \eta) - \hat{\mathcal{U}}_i(x, \eta)| \leq C \left(\frac{1}{N} + N\varepsilon^{\mathcal{U}} \right).$$

Remark 16 *We note that the order of this error—that is, $\mathcal{O}(N^{-1})$ in one term and $\mathcal{O}(N\varepsilon^{\mathcal{U}})$ in the other—is like the error bound from Huré et al. (2020), but with a smaller order of N in the second term. The larger order of N in Huré et al. (2020) likely comes from their use of random propagation in the algorithm and corresponding fixed point system. In our case, the propagation is deterministic and the error is defined differently.*

Remark 17 *To minimize the right hand side of the theorem, one first selects the number N of time steps large enough to minimize the left error term. Then one selects the number $\bar{\delta}$ of parameters large enough so that $\varepsilon^{\mathcal{U}}$ is small enough. Notice that the definition of $\varepsilon^{\mathcal{U}}$ takes into account the approximation error but not the training error (which stems from the optimization method used, such as SGD or Adam). However, we can expect that with a large enough number of training steps, the realized error can also be made small enough.*

5.4 Auxiliary Lemma and Proof

In this section we prove a measure bound lemma that comes into play in the proofs of Theorem 14 and Theorem 15. For every $i \in \{0, \dots, N-1\}$, we denote $(u_i^\kappa, \mu_i^\kappa) = (u^{t_i, \kappa}, \mu^{t_i, \kappa})$ as the solution to (12) on $[t_i, T]$ with $\mu_i^\kappa(t_i) = \kappa \in \mathcal{P}([d])$. In the following sections, we write $\|\cdot\|_{i, \infty}$ to mean the sup-max norm taken over all $t \in [t_i, t_{i+1}]$ and all $x \in [d]$, that is: $\|\varphi\|_{i, \infty} = \sup_{t \in [t_i, t_{i+1}]} \max_{x \in [d]} |\varphi(t, x)|$, where $\varphi : [0, T] \times [d] \rightarrow \mathbb{R}$.

Lemma 18 (Measure Bound) *There exists $C > 0$ depending only on the problem data such that for all $|\pi|$ small enough and every $\theta \in \Theta$ and $(i, \kappa) \in [N] \times \mathcal{P}([d])$,*

$$\|\rho_i^{\theta, \kappa} - \mu_i^\kappa\|_{i, \infty} \leq \frac{C|\pi|}{1 - C|\pi|} \|\mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(\cdot); \theta) - U(\cdot, \cdot, \mu_i^\kappa(\cdot))\|_{i, \infty}. \quad (28)$$

Proof Recall the definition of the loss function in (23). For every $\theta \in \Theta$, define the mapping $e_i(\cdot, \cdot; \theta) : [d] \times \mathcal{P}([d]) \rightarrow \mathbb{R}$ by:

$$e_i(x, \kappa; \theta) := -\hat{\mathcal{U}}_{i+1}(x, M_i^\theta(\kappa)) + \mathcal{U}_i(x, \kappa; \theta) - (\Delta t_i) \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i(\cdot, \kappa; \theta)). \quad (29)$$

Recalling (22), $M_i^\theta(\kappa) = \rho_i^{\theta, \kappa}(t_{i+1})$, (29) rewrites:

$$\hat{\mathcal{U}}_{i+1}(x, \rho_i^{\theta, \kappa}(t_{i+1})) - \mathcal{U}_i(x, \kappa; \theta) + (\Delta t_i) \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i(\cdot, \kappa; \theta)) + e_i(x, \kappa; \theta) = 0. \quad (30)$$

Note that we initialized the DBME algorithm with $\hat{\mathcal{U}}_N(x, \eta) = g(x, \eta)$; also, going backwards with $i = N-1, \dots, 0$, on each time interval $[t_i, t_{i+1}]$, we may think about (30) as a backward finite-difference equation, with $\hat{\mathcal{U}}_{i+1}(x, \rho_i^{\theta, \kappa}(t_{i+1}))$ as the terminal condition. We refer to (30) and (22) together as a forward-backward system.

The difference $\rho_i - \mu^\kappa$ solves an ODE that after integrating yields: for every $s \in [t_i, t_{i+1}]$,

$$\begin{aligned} \rho_i^{\theta, \kappa}(s, x) - \mu_i^\kappa(s, x) &= \int_{t_i}^s \left[\sum_{y \in [d]} (\rho_i^{\theta, \kappa}(t, y) - \mu_i^\kappa(t, y)) \gamma_x^*(y, \Delta_y \mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(t); \theta)) \right. \\ &\quad \left. + \sum_{y \in [d]} \mu_i^\kappa(t, y) [\gamma_x^*(y, \Delta_y \mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(t); \theta)) - \gamma_x^*(y, \Delta_y u_i^\kappa(t, \cdot))] \right] dt. \end{aligned}$$

Since γ^* is bounded and locally Lipschitz (see Remark 3), for every $s \in [t_i, t_{i+1}]$

$$\begin{aligned} |\rho_i^{\theta, \kappa}(s, x) - \mu_i^\kappa(s, x)| &\leq \int_{t_i}^s \left[dC_{\gamma^*} \max_{y \in [d]} |\rho_i^{\theta, \kappa}(t, y) - \mu_i^\kappa(t, y)| \right. \\ &\quad \left. + dC_{L, \gamma^*} \max_{y \in [d]} |\Delta_y (\mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(t); \theta) - u_i^\kappa(t, \cdot))| \right] dt, \end{aligned}$$

where we recall that $C_{\gamma^*}, C_{L, \gamma^*}$ depend only on the model parameters and not on the neural network or the time discretization. Taking the supremum over all $(t, x) \in [t_i, t_{i+1}] \times [d]$ for the last term in the above display and a maximum over the state in the left hand side yields:

$$\begin{aligned} \max_{y \in [d]} |\rho_i^{\theta, \kappa}(s, y) - \mu_i^\kappa(s, y)| &\leq C \int_{t_i}^s \max_{y \in [d]} |\rho_i^{\theta, \kappa}(t, y) - \mu_i^\kappa(t, y)| dt \\ &\quad + C|\pi| \|\mathcal{U}_i(\cdot, \rho_i^{\theta, \kappa}(\cdot); \theta) - u_i^\kappa(\cdot, \cdot)\|_{i, \infty}. \end{aligned}$$

Again taking another supremum bound with respect to s and combining like terms, and then using the consistency property (14), we obtain (28). ■

5.5 Proof of Theorem 14

We recall that C is a generic, positive constant that depends only on the problem's parameters and in particular is independent of ε and $|\pi|$. We allow C to change from one line to the next as needed.

Recall the parameters C_0 and C_1 from (25). We notice that $U \in \mathcal{C}^{1,1}([d] \times \mathcal{P}([d]); C_0/2, C_1/2)$ (see (18)). Indeed, by (16), $\|U\|_\infty \leq \tilde{T} = C_0/2$. Due to Proposition 4, we know that U is Lipschitz with Lipschitz constant $C_{L,U}$ (in fact, $D^\eta U$ is Lipschitz). So, $\|D^\eta U\|_\infty \leq C_1/2$. Now, by (Hornik et al., 1990, Theorem 3.1), the Universal Approximation Theorem asserts that \mathcal{NN} is dense in $\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$: for any $\varepsilon > 0$, there exists $\mathcal{U} \in \mathcal{NN}$ such that $\|U - \mathcal{U}\|_{\mathcal{C}^{1,1}} < \varepsilon$. This neural network satisfies: $\|\mathcal{U}\|_\infty \leq C_0/2 + \varepsilon$ and $\|D^\eta \mathcal{U}\|_\infty \leq C_1/2 + \varepsilon$. For $\varepsilon < \min\{C_0/2, C_1/2\}$, we deduce that $\|\mathcal{U}\|_\infty \leq C_0$ and $\|D^\eta \mathcal{U}\|_\infty \leq C_1$, so it belongs to the class $\mathcal{NN}(C_0, C_1)$.

Next, we show (27). Since Lemma 18 is proved in greater generality, we will use it here for $\theta = \theta^{i,\varepsilon}$. We will denote

$$e_i^\varepsilon(x, \kappa) := e_i(x, \kappa; \theta^{i,\varepsilon}). \quad (31)$$

Let $(x, \kappa) \in [d] \times \mathcal{P}([d])$. Let $s \in [t_i, t_{i+1}]$. By integrating (12) and by (30), along with (14),

$$\begin{aligned} U(s, x, \mu_i^\kappa(s)) - \mathcal{U}_i^\varepsilon(x, \kappa) &= U(t_{i+1}, x, \mu_i^\kappa(t_{i+1})) - \mathcal{U}_{i+1}^\varepsilon(x, \rho_i^{\theta^{i,\varepsilon}, \kappa}(t_{i+1})) \\ &\quad + \int_s^{t_{i+1}} \bar{H}(x, \mu_i^\kappa(t), \Delta_x U(t, \cdot, \mu_i^\kappa(t))) dt \\ &\quad - \Delta t_i \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i^\varepsilon(\cdot, \kappa)) \\ &\quad + e_i^\varepsilon(x, \kappa). \end{aligned} \quad (32)$$

Solving for e_i^ε above, and using triangle inequality, yields:

$$\begin{aligned} |e_i^\varepsilon(x, \kappa)| &\leq |U(s, x, \mu_i^\kappa(s)) - \mathcal{U}_i^\varepsilon(x, \kappa)| \\ &\quad + |\mathcal{U}_{i+1}^\varepsilon(x, \rho_i^{\theta^{i,\varepsilon}, \kappa}(t_{i+1})) - U(t_{i+1}, x, \mu_i^\kappa(t_{i+1}))| \\ &\quad + \left| \int_s^{t_{i+1}} [\bar{H}(x, \mu_i^\kappa(t), \Delta_x U(t, \cdot, \mu_i^\kappa(t)))] dt - \Delta t_i \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i^\varepsilon(\cdot, \kappa)) \right|. \end{aligned}$$

We are going to bound each term of the right hand side, using the Lipschitz continuity of $\mathcal{U}_{i+1}^\varepsilon$, the triangle inequality again, Lemma 18 and the Lipschitz continuity of $\mathcal{U}_i^\varepsilon$, the fact (by integrating in time the forward equation and using the fact that U is bounded) that $|\mu_i^\kappa(s) - \kappa| \leq C|\pi|$, the boundedness of \bar{H} , and (26).

For the first term,

$$\begin{aligned} |U(s, x, \mu_i^\kappa(s)) - \mathcal{U}_i^\varepsilon(x, \kappa)| &\leq |U(s, x, \mu_i^\kappa(s)) - \mathcal{U}_i^\varepsilon(x, \mu_i^\kappa(s))| + |\mathcal{U}_i^\varepsilon(x, \mu_i^\kappa(s)) - \mathcal{U}_i^\varepsilon(x, \kappa)| \\ &\leq \varepsilon + CC_{L, \mathcal{U}_i^\varepsilon} |\pi|. \end{aligned}$$

For the second term,

$$\begin{aligned}
 & |\mathcal{U}_{i+1}^\varepsilon(x, \rho_i^{\theta^i, \varepsilon, \kappa}(t_{i+1})) - U(t_{i+1}, x, \mu_i^\kappa(t_{i+1}))| \\
 & \leq |\mathcal{U}_{i+1}^\varepsilon(x, \rho_i^{\theta^i, \varepsilon, \kappa}(t_{i+1})) - \mathcal{U}_{i+1}^\varepsilon(x, \mu_i^\kappa(t_{i+1}))| + |\mathcal{U}_{i+1}^\varepsilon(x, \mu_i^\kappa(t_{i+1})) - U(t_{i+1}, x, \mu_i^\kappa(t_{i+1}))| \\
 & \leq C_{L, \mathcal{U}_{i+1}^\varepsilon} |\rho_i^{\theta^i, \varepsilon, \kappa}(t_{i+1}) - \mu_i^\kappa(t_{i+1})| + |\mathcal{U}_{i+1}^\varepsilon(x, \mu_i^\kappa(t_{i+1})) - U(t_{i+1}, x, \mu_i^\kappa(t_{i+1}))| \\
 & \leq C_{L, \mathcal{U}_{i+1}^\varepsilon} |\rho_i^{\theta^i, \varepsilon, \kappa}(t_{i+1}) - \mu_i^\kappa(t_{i+1})| + |\mathcal{U}_{i+1}^\varepsilon(x, \mu_i^\kappa(t_{i+1})) - U(t_{i+1}, x, \mu_i^\kappa(t_{i+1}))| \\
 & \leq CC_{L, \mathcal{U}_{i+1}^\varepsilon} |\pi| + \varepsilon.
 \end{aligned}$$

For the third term,

$$\left| \int_s^{t_{i+1}} [\bar{H}(x, \mu_i^\kappa(t), \Delta_x U(t, \cdot, \mu_i^\kappa(t)))] dt - \Delta t_i \bar{H}(x, \kappa, \Delta_x \mathcal{U}_i^\varepsilon(\cdot, \kappa)) \right| \leq 2C_{\bar{H}} |\pi|.$$

Combining these bounds, we obtain:

$$|e_i^\varepsilon(x, \kappa)| \leq C(\varepsilon + |\pi|).$$

Now we can take a maximum over $(x, \kappa) \in [d] \times \mathcal{P}([d])$, and with (23), this yields (27). ■

5.6 Proof of Theorem 15

Fix $\eta \in \mathcal{P}([d])$. Recall (22). For any $i \in [N]$ and any $\hat{\theta}^i$ minimizing L_i in (23), set

$$\rho_i := \rho_i^{\hat{\theta}^i, \mu^\eta(t_i)}.$$

Note that for any θ^i ,

$$\begin{aligned}
 L_i(\theta^i) &= \max_{x, \kappa} e_i(x, \kappa; \theta^i) \\
 &\geq \max_{x \in [d]} e_i(x, \mu^\eta(t_i); \theta^i) \\
 &= \max_{x \in [d]} |\hat{\mathcal{U}}_{i+1}(x, M_i^{\theta^i}(\mu^\eta(t_i))) - \mathcal{U}_i(x, \mu^\eta(t_i); \theta^i) + (\Delta t_i) \bar{H}(x, \mu^\eta(t_i), \Delta_x \mathcal{U}_i(\cdot, \mu^\eta(t_i); \theta^i))|.
 \end{aligned}$$

So when $\theta^i = \hat{\theta}^i$,

$$\varepsilon_i \geq \max_{x \in [d]} |\hat{\mathcal{U}}_{i+1}(x, M_i^{\hat{\theta}^i}(\mu^\eta(t_i))) - \hat{\mathcal{U}}_i(x, \mu^\eta(t_i)) + (\Delta t_i) \bar{H}(x, \mu^\eta(t_i), \Delta_x \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i)))|, \quad (33)$$

where we recall ε_i is defined in (24). Like in (33), by the definition of the loss (29) and (31), $\|e_i^\varepsilon\|_\infty = \varepsilon_i$.

Once more integrating (12), by (30), and using (14),

$$\begin{aligned}
 U(s, x, \mu^\eta(s)) - \hat{\mathcal{U}}_i(x, \mu^\eta(t_i)) &= U(t_{i+1}, x, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(x, \rho_i(t_{i+1})) \\
 &\quad + \int_s^{t_{i+1}} [\bar{H}(x, \mu^\eta(t), \Delta_x U(t, \cdot, \mu^\eta(t)))] dt \\
 &\quad - \Delta t_i \bar{H}(x, \mu^\eta(t_i), \Delta_x \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))) \\
 &\quad + \hat{e}_i(x, \mu^\eta(t_i)).
 \end{aligned} \quad (34)$$

Furthremore, by the Lipschitz continuity of H ,

$$\begin{aligned} & \left| \int_s^{t_{i+1}} [\bar{H}(x, \mu^\eta(t), \Delta_x U(t, \cdot, \mu^\eta(t)))] dt - \Delta t_i \bar{H}(x, \mu^\eta(t_i), \Delta_x \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))) \right| \\ & \leq C_{L, \bar{H}} \Delta t_i \|\mu^\eta - \mu^\eta(t_i)\|_{i, \infty} + C_{L, \bar{H}} C \Delta t_i \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty}. \end{aligned}$$

Going back to (34), taking the supremum over $(s, x) \in [t_i, t_{i+1}] \times [d]$, using the fact that $\|\mu^\eta(\cdot) - \mu^\eta(t_i)\|_{i, \infty} \leq C \Delta t_i$, and combining like terms,

$$\begin{aligned} & \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty} \\ & \leq \frac{1}{1-C|\pi|} \left[\max_{y \in [d]} |U(t_{i+1}, y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \rho_i(t_{i+1}))| + C|\pi|^2 + \max_{y \in [d]} |\hat{e}_i(y, \mu^\eta(t_i))| \right]. \end{aligned}$$

Note that the above inequality requires $C|\pi| < 1$, which is true for N large enough, as we assume in the statement. Using the triangle inequality,

$$\begin{aligned} & \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty} \\ & \leq \frac{1}{1-C|\pi|} \left[C|\pi|^2 + \max_{y \in [d]} |U(t_{i+1}, y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1}))| \right. \\ & \quad \left. + \max_{y \in [d]} |\hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \rho_i(t_{i+1}))| + \max_{y \in [d]} |\hat{e}_i(y, \mu^\eta(t_i))| \right]. \end{aligned} \tag{35}$$

Next, we will use the fact that $\hat{\mathcal{U}}_{i+1}$ is Lipschitz, (28), triangle inequality, the fact that $\hat{\mathcal{U}}_i$ is Lipschitz, that $\|\rho_i(\cdot) - \mu^\eta(\cdot)\|_{i, \infty} \leq C|\pi|$, and the fact that

$$\text{the map} \quad (0, |\pi|^{-1}) \ni C \mapsto \frac{C}{1-C|\pi|} \quad \text{is nondecreasing,}$$

in order to update the generic constant $C > 0$.

We note that:

$$\begin{aligned} & \max_{y \in [d]} |\hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \rho_i(t_{i+1}))| \\ & \leq C_{L, \hat{\mathcal{U}}_{i+1}} |\mu^\eta(t_{i+1}) - \rho_i(t_{i+1})| \\ & \leq \frac{C|\pi|}{1-C|\pi|} \|\hat{\mathcal{U}}_i(\cdot, \rho_i(\cdot)) - U(\cdot, \cdot, \mu^\eta(\cdot))\|_{i, \infty} \\ & \leq \frac{C|\pi|}{1-C|\pi|} \|\hat{\mathcal{U}}_i(\cdot, \rho_i(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty} + \frac{C|\pi|}{1-C|\pi|} \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty} \\ & \leq \frac{C|\pi|^2}{1-C|\pi|} C_{L, \hat{\mathcal{U}}_i} + \frac{C|\pi|}{1-C|\pi|} \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty}, \end{aligned}$$

where we used the fact that $C_{L, \hat{\mathcal{U}}_i}$ is bounded by a constant depending only on the model's parameters.

Going back to (35) and using the above bound, we obtain:

$$\begin{aligned}
 & \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i,\infty} \\
 & \leq \frac{1}{1-C|\pi|} \left[C|\pi|^2 + \max_{y \in [d]} |U(t_{i+1}, y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1}))| \right. \\
 & \quad \left. + \max_{y \in [d]} |\hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \rho_i(t_{i+1}))| + \max_{y \in [d]} |\hat{e}_i(y, \mu^\eta(t_i))| \right] \\
 & \leq \frac{C}{(1-C|\pi|)^2} \left[|\pi|^2 + |\pi| \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i,\infty} \right] \\
 & \quad + \frac{1}{1-C|\pi|} \left[\max_{y \in [d]} |\hat{e}_i(y, \mu^\eta(t_i))| + \max_{y \in [d]} |U(t_{i+1}, y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1}))| \right].
 \end{aligned}$$

For the rest of the proof, we denote by Γ_0 the constant C appearing in the above right-hand side. That is, we are fixing Γ_0 to be C at this particular point in the proof, but C later on will be generic as usual. Rearranging terms, there exists $\Gamma_1 > 0$ depending only on Γ_0 and π such that:

$$\begin{aligned}
 & \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i,\infty} \\
 & \leq \left(\frac{1}{1-\Gamma_0|\pi|} \right) \left(\frac{(1-\Gamma_1|\pi|)^2}{1-3\Gamma_1|\pi|} \right) \left[\Gamma_0|\pi|^2 + \max_{y \in [d]} |\hat{e}_i(y, \mu^\eta(t_i))| \right. \\
 & \quad \left. + \max_{y \in [d]} |U(t_{i+1}, y, \mu^\eta(t_{i+1})) - \hat{\mathcal{U}}_{i+1}(y, \mu^\eta(t_{i+1}))| \right]. \tag{36}
 \end{aligned}$$

Taking $C_2 = 3\Gamma_1$, we have for all $|\pi|$ small enough that:

$$\left(\frac{1}{1-\Gamma_0|\pi|} \right) \left(\frac{(1-\Gamma_1|\pi|)^2}{1-3\Gamma_1|\pi|} \right) \leq \left(\frac{1-\Gamma_1|\pi|}{1-C_2|\pi|} \right)^2 =: Q.$$

Note that:

$$\begin{aligned}
 \sum_{j=1}^{N-1} Q^j &= \frac{Q^N - Q}{Q-1} \\
 &= \left[\left(\frac{1-\Gamma_1|\pi|}{1-C_2|\pi|} \right)^{2N} - Q \right] \frac{(1-C_2|\pi|)^2}{2|\pi|(C_2-\Gamma_1)+|\pi|^2(\Gamma_1^2-C_2^2)} \\
 &\leq \left(\frac{1-\Gamma_1|\pi|}{1-C_2|\pi|} \right)^{2N} \cdot \frac{(1-C_2|\pi|)^2}{2|\pi|(C_2-\Gamma_1)+|\pi|^2(\Gamma_1^2-C_2^2)}.
 \end{aligned}$$

Recalling (19), and noting that:

$$\left(\frac{1-\Gamma_1|\pi|}{1-C_2|\pi|} \right)^{2N} \rightarrow e^{2C_2-2\Gamma_1} = e^{4\Gamma_1},$$

monotonically as $N \rightarrow \infty$, we have for all $|\pi| > 0$ small enough that:

$$\sum_{j=1}^{N-1} Q^j \leq CN e^{4\Gamma_1}. \tag{37}$$

Returning to (36), we know that for $i = N-1$, the terminal conditions agree and therefore,

$$\|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_{N-1}(\cdot, \mu^\eta(t_{N-1}))\|_{N-1,\infty} \leq Q \left[C|\pi|^2 + \max_{y \in [d]} |\hat{e}_{N-1}(y, \mu^\eta(t_i))| \right].$$

By backward induction with (36) and then using (37),

$$\begin{aligned} \max_{i \in [N]} \|U(\cdot, \cdot, \mu^\eta(\cdot)) - \hat{\mathcal{U}}_i(\cdot, \mu^\eta(t_i))\|_{i, \infty} &\leq \left(\sum_{j=1}^{N-1} Q^j \right) [C|\pi|^2 + \max_{i \in [N]} \|\hat{e}_i\|_\infty] \\ &\leq C|\pi| + CN\varepsilon^\mathcal{U}. \end{aligned}$$

This inequality holds for all $\eta \in \mathcal{P}([d])$ and therefore we can take a maximum over all $\eta \in \mathcal{P}([d])$ to conclude the proof. ■

6. The DGME Algorithm and Convergence Results

In what follows, we describe the DGME adaptation of the DGM to solve the master equation, and we then provide two novel theoretical results. Recall that U denotes the unique classical solution to (1) and that $U(\cdot, x, \cdot) \in \mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d]))$ for all $x \in [d]$. By some previous remarks and from Hornik et al. (1990, Theorem 3.1), it is known that \mathcal{NN} is dense in $(\mathcal{C}^{1,1}([0, T] \times \mathcal{P}([d])), \|\cdot\|_{\mathcal{C}^{1,1}})$.

In Algorithm 2, we introduce the DGME, modified with the loss to a maximum from the expectation of a squared quantity. We also note that unlike the DBME, in which we restricted our attention to $\theta \in \Theta$, in the DGME we can take $\theta \in \mathbb{R}^{\bar{d}}$ without restriction.

Algorithm 2 DGME

- 1: **Input:** An initial vector θ .
- 2: **Output:** A trained vector $\hat{\theta}$ such that $\mathcal{U}(\cdot, \cdot, \cdot; \hat{\theta})$ approximately solves (1).
- 3: Compute: $\hat{\theta} \in \arg \min_{\theta \in \mathbb{R}^{\bar{d}}} L(\theta)$
- 4: where

$$\begin{aligned} L(\theta) := \max_{(t, x, \eta) \in [0, T] \times [d] \times \mathcal{P}([d])} &\left\{ \left| \partial_t \mathcal{U}(t, x, \eta; \theta) + \bar{H}(x, \eta, \Delta_x \mathcal{U}(t, \cdot, \eta; \theta)) \right. \right. \\ &+ \left. \sum_{y \in [d]} \eta_y D_y^\eta \mathcal{U}(t, x, \eta; \theta) \cdot \gamma^*(y, \Delta_y \mathcal{U}(t, \cdot, \eta; \theta)) \right| \quad (38) \\ &\left. + |\mathcal{U}(T, x, \eta; \theta) - g(x, \eta)| \right\} \end{aligned}$$

As we noted in Remark 10, we cannot, in practice, compute the maximum over the infinitely many elements in $[0, T] \times [d] \times \mathcal{P}([d])$. So, like for the DBME and the original DGM formulation, we randomly sample the space instead. As in Remark 10, we note that uniformly sampling the space results in an estimator for the maximum where for \mathcal{K} samples, we get $\mathcal{O}(|\mathcal{K}|^{-1/(d-1)})$ error for the estimator. The master equation we study in this paper only involves first-order derivatives with respect to the inputs t and η . Hence computing the gradient with respect to the neural network parameters of the loss function (38) only requires computing second-order derivatives, which can be done by automatic differentiation.³

3. If a second-order term was involved in the PDE, then the computation of the gradient of the loss would require third-order derivatives. Using automatic differentiation would probably be prohibitive in

The following two theorems establish convergence of the DGME and are analogous to results from Sirignano and Spiliopoulos (2018, Theorem 7.1, Theorem 7.3), but these results do not apply to the master equation due to its distinct form. Instead, we use the structure of the MFG system itself in order to obtain useful bounds. Yet another difference is that our results are given in the supremum norm while those of Sirignano and Spiliopoulos (2018) are in L^2 norm.

In the theorem below, we note that the density of the neural networks implies there always exists a neural network uniformly close to the master equation solution U ; then, we write that any close network satisfies the master equation (1) in an approximate sense. As we mentioned in the introduction, the two main results of each section are, loosely speaking, converse to one another. The first result deals with existence of networks close to the true solution and shows that when a network is close, the corresponding algorithm's loss is small.

Theorem 19 (DGME Approximation) *For every $\varepsilon > 0$, there exists a neural network $\hat{\mathcal{U}} \in \mathcal{NN}$ such that:*

$$\|\hat{\mathcal{U}} - U\|_{\mathcal{C}^{1,1}} < \varepsilon. \quad (39)$$

Moreover, let $e^{\hat{\mathcal{U}}} : [0, T] \times [d] \times \mathcal{P}([d]) \rightarrow \mathbb{R}$ be defined by:

$$\begin{aligned} e^{\hat{\mathcal{U}}}(t, x, \eta) &:= \sum_{y, z \in [d]} \eta_y D_{yz}^\eta \hat{\mathcal{U}}(t, x, \eta) \gamma_z^*(y, \eta, \Delta_y \hat{\mathcal{U}}(t, \cdot, \eta)) + H(x, \eta, \Delta_x \hat{\mathcal{U}}(t, \cdot, \eta)) + \partial_t \hat{\mathcal{U}}(t, x, \eta), \\ e^{\hat{\mathcal{U}}}(T, x, \eta) &:= \hat{\mathcal{U}}(T, x, \eta) - g(x, \eta). \end{aligned} \quad (40)$$

Then $e^{\hat{\mathcal{U}}}$ is measurable and there exists a positive constant C_{DGME} , depending only on the problem data and independent of ε , such that:

$$\|e^{\hat{\mathcal{U}}}\|_\infty < C_{DGME} \varepsilon. \quad (41)$$

We stress that Theorem 19 has two parts. The first part, namely (39), holds by universal approximation theorem and is straightforward given the (known) regularity of the master equation solution. It is not directly related to the algorithm. The second part, namely (41), provides a bound on the residual of the PDE (including the terminal condition) when using a neural network which gives (39). This provides an upper bound on the loss function of Algorithm 2, which is simply the PDE residual.

We now present the main DGME result concerning convergence. This theorem asserts that if the DGME finds a network with small error in its loss function, then that network must be uniformly close to the master equation's solution U .

Theorem 20 (DGME Convergence) *Let $0 < \varepsilon < 1$. Let $\hat{\mathcal{U}} : [0, T] \times [d] \times \mathcal{P}([d]) \rightarrow \mathbb{R}$ be a neural network obtained by the DGME such that (40)–(41) hold. Then, there exists $C > 0$ depending only on the problem data and independent of ε such that:*

$$\|\hat{\mathcal{U}} - U\|_\infty < C\varepsilon. \quad (42)$$

high dimension. In such cases, our method could perhaps be modified, using approximate second-order derivative computations, as proposed in (Sirignano and Spiliopoulos, 2018, Section 3).

Consequently, any sequence of neural networks $(\hat{U}_n)_{n \in \mathbb{N}}$, with corresponding errors $(\varepsilon_n)_{n \in \mathbb{N}}$ such that $\varepsilon_n \rightarrow 0$, converges to U uniformly.

The proof is based on a careful analysis of the master equation, and its connection with forward-backward ODE systems characterizing the MFG equilibrium for each initial condition. To the best of our knowledge, the result cannot be deduced from the literature. In particular, the analysis of the DGM in Sirignano and Spiliopoulos (2018) is only for a specific class of (quasilinear, parabolic) PDEs, which does not cover ours. Furthermore, the master equation, despite some similarities, is not an HJB equation and does not enjoy a comparison principle.

Remark 21 *The attentive reader would observe that in the proof of Theorems 19 and 20, we effectively establish the existence of two constants, $\underline{C}, \bar{C} > 0$, such that for any $\hat{U} \in \mathcal{NN}$, one has $\underline{C}\|e^{\hat{U}}\|_\infty \leq \|\hat{U} - U\|_\infty \leq \bar{C}\|e^{\hat{U}}\|_\infty$.*

6.1 Proof of Theorem 19

We note that (39) is immediate from the universal approximation result of Hornik et al. (1990, Theorem 3.1). It suffices to prove (40)–(41). Recall that $C > 0$ is a generic constant independent of t and η , and may change from one line to the next.

From (39), for any $\varepsilon > 0$ and any $x \in [d]$, there exists a neural network $\hat{U}(\cdot, x, \cdot) \in \mathcal{NN}$ such that:

$$\|U(\cdot, x, \cdot) - \hat{U}(\cdot, x, \cdot)\|_{\mathcal{C}^{1,1}} < \varepsilon, \quad \text{for all } x \in [d].$$

Since (39) holds for all $x \in [d]$, we may write $\|\hat{U} - U\|_{\mathcal{C}^{1,1}}$ to mean $\max_{x \in [d]} \|\hat{U}(\cdot, x, \cdot) - U(\cdot, x, \cdot)\|_{\mathcal{C}^{1,1}}$ without ambiguity. Note that:

$$\|\hat{U}\|_\infty \leq \|U\|_\infty + \varepsilon \leq \tilde{T} + \varepsilon,$$

so \hat{U} is bounded above by the right hand side of (9). When ε is small enough, H is restricted to the interval $[-W, W]$ and note that H is locally Lipschitz on $[-W, W]$ by assumption. We may therefore treat H as Lipschitz.

Recall the definition of $e^{\hat{U}}(t, x, \eta)$ from (40). For brevity, we write $e^{\hat{U}}$ as e during the proof. Since U solves (1), we find that:

$$\begin{aligned} e(t, x, \eta) &= \partial_t \hat{U}(t, x, \eta) - \partial_t U(t, x, \eta) + \bar{H}(x, \eta, \Delta_x \hat{U}(t, \cdot, \eta)) - \bar{H}(x, \eta, \Delta_x U(t, \cdot, \eta)) \\ &\quad + \sum_{y, z \in [d]} \left[\eta_y D_{yz}^\eta \hat{U}(t, x, \eta) \gamma_z^*(y, \Delta_y \hat{U}(t, \cdot, \eta)) - \eta_y D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y U(t, \cdot, \eta)) \right]. \end{aligned} \tag{43}$$

We can check that, for the last line of (43):

$$\begin{aligned} &\left| \sum_{y, z \in [d]} \left[\eta_y D_{yz}^\eta \hat{U}(t, x, \eta) \gamma_z^*(y, \Delta_y \hat{U}(t, \cdot, \eta)) - \eta_y D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y U(t, \cdot, \eta)) \right] \right| \\ &\leq d^2 \max_{y, z \in [d]} |D_{yz}^\eta \hat{U}(t, x, \eta) \gamma_z^*(y, \Delta_y \hat{U}(t, \cdot, \eta)) - D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y U(t, \cdot, \eta))| \\ &\leq d^2 \max_{y, z \in [d]} |D_{yz}^\eta \hat{U}(t, x, \eta) \gamma_z^*(y, \Delta_y \hat{U}(t, \cdot, \eta)) - D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y \hat{U}(t, \cdot, \eta))| \\ &\quad + d^2 \max_{y, z \in [d]} |D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y \hat{U}(t, \cdot, \eta)) - D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y U(t, \cdot, \eta))|. \end{aligned}$$

Using the facts that γ^* is bounded and Lipschitz, that the classical master equation solution U has $D^\eta U$ bounded, and (39), we can bound the above quantity by:

$$d^2 C_\gamma \max_{y,z \in [d]} |D_{yz}^\eta (\hat{U}(t, x, \eta) - U(t, x, \eta))| + d^2 C_{D^\eta U} C_{L,\gamma} \max_{y \in [d]} |\Delta_y (\hat{U}(t, \cdot, \eta) - U(t, \cdot, \eta))| < C\varepsilon,$$

where C is independent of t, η . Going back to (43), using the Lipschitz continuity of H and of F , and again (39), we obtain: $|e(t, x, \eta)| < C\varepsilon$. In other words, \hat{U} satisfies (40)–(41). ■

In the following section, we prove a lemma that is at the heart of Theorem 20. The lemma allows us to circumvent the unfortunate fact that the master equation has no comparison principle; thus we introduce an approximate MFG system and use duality.

6.2 Approximate MFG System

In the following lemma, we define the function v through the network \hat{U} and we use the fact that it satisfies an approximate Hamilton–Jacobi–Bellman equation. Its evolution is coupled with a function λ , which solves a forward Kolmogorov equation of a non-linear Markov chain. That is, define λ as the solution to:

$$\frac{d}{dt} \lambda_x(t) = \sum_{y \in [d]} \lambda_y(t) \gamma_x^*(y, \Delta_y \hat{U}(t, \cdot, \lambda(t))), \quad \lambda(t_0) = \eta,$$

on $t \in [t_0, T]$. Note that λ is well-defined by the Picard–Lindelöf Theorem.

Lemma 22 (Approximate MFG System and Network Bound) *Let $0 < \varepsilon < 1$. Let $\hat{U} : [0, T] \times [d] \times \mathcal{P}([d]) \rightarrow \mathbb{R}$ be a neural network obtained by the DGME. Let the error function e be defined as in (40). Assume (41) holds. Then, by setting $v_x(t) := \hat{U}(t, x, \lambda(t))$, we get that (v, λ) solves the approximate MFG system on $(t, x) \in [t_0, T] \times [d]$:*

$$\begin{aligned} \frac{d}{dt} v_x(t) + \bar{H}(x, \lambda(t), \Delta_x v(t)) &= e(t, x, \lambda(t)), \\ \frac{d}{dt} \lambda_x(t) &= \sum_{y \in [d]} \lambda_y(t) \gamma_x^*(y, \Delta_y v(t)), \\ \lambda(t_0) &= \eta, \\ v_x(T) &= g(x, \lambda(T)) + e(T, x, \lambda(T)). \end{aligned} \tag{44}$$

Moreover, for all $(t, x, \eta) \in [t_0, T] \times [d] \times \mathcal{P}([d])$ and when ε is small enough:

$$|\hat{U}(t, x, \eta)| \leq T[C_f + C_F + C_{DGME}\varepsilon] + C_g + C_{DGME}\varepsilon \leq W,$$

where C_{DGME} is defined in (41) and W is defined in (9). Hence, the regularity assumptions on H given in (10) apply.

Remark 23 *The careful reader may wonder why the bound on $|\hat{U}|$ above does not immediately follow from Theorem 19. This is because Lemma 22 and Theorem 20 deal with a*

network obtained from the DGME, while Theorem 19 does not. Theorem 19 shows that running the DGME is not hopeless since there exist neural networks with the properties we are seeking. Offhand, we only know that the network considered in Lemma 22 and Theorem 20 satisfies the master equation in some approximate way; therefore, we use only this fact to prove $|\hat{\mathcal{U}}|$ is bounded.

6.3 Proof of Lemma 22

We use the structure of the MFG system to obtain important estimates. That is, we fix $t_0 \in [0, T)$ and $\eta \in \mathcal{P}([d])$ and set λ to be the unique solution to the Kolmogorov equation in (44) on $[t_0, T]$.

For $v_x(t) := \hat{\mathcal{U}}(t, x, \lambda(t))$ defined on $[t_0, T]$, we use the fact that $\hat{\mathcal{U}}$ satisfies (40) to get:

$$\begin{aligned} \frac{d}{dt}v_x(t) &= \frac{d}{dt}\hat{\mathcal{U}}(t, x, \lambda(t)) \\ &= \partial_t \hat{\mathcal{U}}(t, x, \lambda(t)) + \sum_{y \in [d]} \lambda_y(t) D_y^\eta \hat{\mathcal{U}}(t, x, \lambda(t)) \cdot \gamma^*(y, \Delta_y \hat{\mathcal{U}}(t, \cdot, \lambda(t))) \\ &= e(t, x, \lambda(t)) - \bar{H}(x, \lambda(t), \Delta_x v(t)), \end{aligned}$$

where the error function e is defined in (40). So together, (v, λ) solves (44). It remains to show the bound on $\hat{\mathcal{U}}$.

Let $\tilde{\mathcal{X}}$ be a jump process on $[d]$ with rate matrix given by $\gamma_x^*(y, \Delta_y v(t))$; namely, its infinitesimal generator is given by $[\gamma_x^*(y, \Delta_y v(t))]_{x, y \in [d]}$. Set $\tilde{\mathcal{X}}_{t_0} = z \in [d]$. By an application of Itô's lemma, and using the boundary conditions with (40), we have for any $t \in [t_0, T]$,

$$\begin{aligned} &\mathbb{E}[g(\tilde{\mathcal{X}}_T, \lambda(T)) + e(T, \tilde{\mathcal{X}}_T, \lambda(T)) - v_z(t_0)] \\ &= \mathbb{E} \int_{t_0}^T [\partial_t v_{\tilde{\mathcal{X}}_s}(s) + \gamma^*(\tilde{\mathcal{X}}_s, \Delta_{\tilde{\mathcal{X}}_s} v(s)) \cdot \Delta_{\tilde{\mathcal{X}}_s} v(s)] ds \\ &= \mathbb{E} \int_{t_0}^T [-f(\tilde{\mathcal{X}}_s, \gamma^*(\tilde{\mathcal{X}}_s, \Delta_{\tilde{\mathcal{X}}_s} v(s))) - F(\tilde{\mathcal{X}}_s, \lambda(s)) + e(s, \tilde{\mathcal{X}}_s, \lambda(s))] ds. \end{aligned}$$

Then,

$$v_z(t_0) = J(t_0, x, \lambda, \gamma^*) + \mathbb{E} \left[\int_{t_0}^T e(t, \tilde{\mathcal{X}}_t, \lambda(t)) dt + e(T, \tilde{\mathcal{X}}_T, \lambda(T)) \right].$$

So by definition of v , by (41), since f and F are bounded, and since t_0 was arbitrary,

$$|\hat{\mathcal{U}}(t, x, \eta)| \leq T[C_f + C_F + C_{DGME}\varepsilon] + C_g + C_{DGME}\varepsilon. \quad (45)$$

By (45), $\hat{\mathcal{U}}$ is uniformly bounded. Moreover, since $\varepsilon < 1$, $\hat{\mathcal{U}}(t, x, \eta) \in [-W, W]$ for all (t, x, η) (W is defined in (9)). On the input $\hat{\mathcal{U}}(t, x, \eta)$ then, (10) holds. ■

In particular, from Lemma 22, we have that for all $(t, x, \eta) \in [0, T] \times [d] \times \mathcal{P}([d])$:

$$D_{pp}^2 H(x, \Delta_x \hat{\mathcal{U}}(t, \cdot, \eta)) \leq -C_{2,H}.$$

This observation will be useful in the following proof.

6.4 Proof of Theorem 20

The proof relies on an approximate MFG system on the time interval $[t_0, T]$ with initial distribution $\eta \in \mathcal{P}([d])$. We compare its solution to the true solution using MFG duality, that is, integrating the backward equation's solution u as a test function in the forward equation to obtain estimates on the solution. The MFG solution is related to the master equation as discussed in Section 3, and since t_0 and η are arbitrary, we obtain convergence to the master equation solution.

The proof proceeds in three steps. Already we formulated the approximate MFG system (44) defined through a neural network and meant to mimic the MFG. We then proved a bound on the solution to the approximate system. In the subsequent steps, we will take advantage of the fact that the approximate MFG system is structured like the MFG system. In Step 1, we use the duality of the MFG and approximate MFG systems to obtain several estimates. That is, we treat the approximate MFG value v as a test function to integrate against, in order to derive our estimates. In Steps 2 and 3, we integrate the differences of solution and combine these estimates with duality to finish the proof.

STEP 1: DUALITY

Fix $\eta \in \mathcal{P}([d])$ and let $(u^{t_0, \eta}, \mu^{t_0, \eta})$ solve (12) on $[t_0, T]$ with $\mu^{t_0, \eta}(t_0) = \eta$. For ease of notation, we suppress the superscript and simply write (u, μ) instead. Define $\mathcal{Q} := v - u$ and $\mathcal{M} := \lambda - \mu$. Using (12) and (44), $(\mathcal{Q}, \mathcal{M})$ solves:

$$\begin{aligned} -\frac{d\mathcal{Q}_x}{dt}(t) &= \bar{H}(x, \lambda(t), \Delta_x v(t)) - \bar{H}(x, \mu(t), \Delta_x u(t)) - e(t, x, \lambda(t)), \\ \frac{d\mathcal{M}_x}{dt}(t) &= \sum_{y \in [d]} [\lambda_y(t) \gamma_x^*(y, \Delta_y v(t)) - \mu_y(t) \gamma_x^*(y, \Delta_y u(t))], \\ \mathcal{Q}_x(T) &= g(x, \lambda(T)) - g(x, \mu(T)) + e(T, x, \lambda(T)), \\ \mathcal{M}(t_0) &= 0. \end{aligned} \tag{46}$$

Using the product rule and (46):

$$\begin{aligned} \sum_{x \in [d]} \mathcal{Q}_x(T) \mathcal{M}_x(T) &= \sum_{x \in [d]} \int_{t_0}^T \left\{ \mathcal{M}_x(t) [-\bar{H}(x, \lambda(t), \Delta_x v(t)) + \bar{H}(x, \mu(t), \Delta_x u(t))] \right. \\ &\quad + \mathcal{Q}_x(t) \sum_{y \in [d]} [\lambda_y(t) \gamma_x^*(y, \Delta_y v(t)) - \mu_y(t) \gamma_x^*(y, \Delta_y u(t))] \\ &\quad \left. - \mathcal{M}_x(t) e(t, x, \lambda(t)) \right\} dt. \end{aligned}$$

We use the definition of \bar{H} , the terminal condition on \mathcal{Q} , and the fact that the rates sum to zero (namely, for any $x \in [d]$ and $p \in \mathbb{R}^d$, $\sum_y \gamma_y^*(x, \Delta_x p) = 0$) to get:

$$\begin{aligned} & \sum_{x \in [d]} (g(x, \lambda(T)) - g(x, \mu(T))) \mathcal{M}_x(T) + \sum_{x \in [d]} \int_{t_0}^T [F(x, \lambda(t)) - F(x, \mu(t))] \mathcal{M}_x(t) dt \\ &= \int_{t_0}^T \left[\sum_{x \in [d]} [H(x, \Delta_x u(t)) - H(x, \Delta_x v(t))] \mathcal{M}_x(t) \right. \\ & \quad \left. + \Delta_x \mathcal{Q}(t) \cdot [\lambda_x(t) \gamma_y^*(x, \Delta_x v(t)) - \mu_x(t) \gamma_y^*(x, \Delta_x u(t))]_{y \in [d]} \right] dt \\ & \quad - \int_{t_0}^T e(t, \cdot, \lambda(t)) \cdot \mathcal{M}(t) dt - e(T, \cdot, \lambda(T)) \cdot \mathcal{M}(T). \end{aligned}$$

By assumption, g and F are Lasry–Lions monotone and hence the left-hand side of the above display is non-negative. Hence,

$$\begin{aligned} 0 \leq & \int_{t_0}^T \left[\sum_{x \in [d]} [H(x, \Delta_x u(t)) - H(x, \Delta_x v(t))] \mathcal{M}_x(t) \right. \\ & \quad \left. + \Delta_x \mathcal{Q}(t) \cdot [\lambda_x(t) \gamma_y^*(x, \Delta_x v(t)) - \mu_x(t) \gamma_y^*(x, \Delta_x u(t))]_{y \in [d]} \right] dt \\ & - \int_{t_0}^T e(t, \cdot, \lambda(t)) \cdot \mathcal{M}(t) dt - e(T, \cdot, \lambda(T)) \cdot \mathcal{M}(T). \end{aligned} \tag{47}$$

By Lemma 22, we may apply (11) and (10). That is, there exists $C_{2,H} > 0$ such that for all $t \in [t_0, T]$ and all $x \in [d]$:

$$\begin{aligned} H(x, \Delta_x v(t)) - H(x, \Delta_x u(t)) - \Delta_x \mathcal{Q}(t) \cdot \gamma^*(x, \Delta_x v(t)) &\leq -C_{2,H} |\Delta_x \mathcal{Q}(t)|^2, \\ H(x, \Delta_x u(t)) - H(x, \Delta_x v(t)) + \Delta_x \mathcal{Q}(t) \cdot \gamma^*(x, \Delta_x u(t)) &\leq -C_{2,H} |\Delta_x \mathcal{Q}(t)|^2. \end{aligned}$$

Using the above display in (47) and rearranging terms,

$$\begin{aligned} & C \int_{t_0}^T \left[\sum_{x \in [d]} |\Delta_x \mathcal{Q}(t)|^2 (\lambda_x(t) + \mu_x(t)) \right] dt \\ & \leq - \int_{t_0}^T e(t, \cdot, \lambda(t)) \cdot \mathcal{M}(t) dt - e(T, \cdot, \lambda(T)) \cdot \mathcal{M}(T), \end{aligned} \tag{48}$$

where recall that C denotes a generic positive constant and its value may change from one line to the next.

STEP 2: INTEGRATING THE BACKWARD EQUATION FROM (46)

Integrating the backward equation from (46), using the Lipschitz continuity of F and H (the latter uses Lemma 22, again), using the terminal condition for \mathcal{Q} , using the Lipschitz

continuity of g , and taking the supremum norm:

$$\begin{aligned}
 |\mathcal{Q}_x(t_0)| &\leq |\mathcal{Q}_x(T)| + \int_{t_0}^T [|F(x, \lambda(t)) - F(x, \mu(t))| + |H(x, \Delta_x v(t)) - H(x, \Delta_x u(t))| \\
 &\quad + |e(t, x, \lambda(t))|] dt \\
 &\leq |\mathcal{Q}_x(T)| + C \int_{t_0}^T [|\lambda(t) - \mu(t)| + |\Delta_x(v(t) - u(t))| + |e(t, x, \lambda(t))|] dt \\
 &= |\mathcal{Q}_x(T)| + C \int_{t_0}^T [|\mathcal{M}(t)| + |\Delta_x \mathcal{Q}(t)| + |e(t, x, \lambda(t))|] dt \\
 &\leq |g(x, \lambda(T)) - g(x, \mu(T))| + |e(T, x, \lambda(T))| \\
 &\quad + C \int_{t_0}^T [|\mathcal{M}(t)| + |\Delta_x \mathcal{Q}(t)| + |e(t, x, \lambda(t))|] dt \\
 &\leq C \left(\|\mathcal{M}\|_\infty + \|e\|_\infty + \int_{t_0}^T \max_{z \in [d]} |\mathcal{Q}_z(t)| dt \right).
 \end{aligned}$$

Then, by Gronwall's inequality and taking the supremum over all $t_0 \leq T$:

$$\|\mathcal{Q}\|_\infty \leq C(\|\mathcal{M}\|_\infty + \|e\|_\infty). \quad (49)$$

STEP 3: INTEGRATING THE FORWARD EQUATION FROM (46)

Integrating the measure equation from (46) on $[t_0, t_1]$, adding and subtracting a term, and using the boundedness of γ^* , and the Lischitz continuity of γ^* (over the range to which the inputs belongs):

$$\begin{aligned}
 |\mathcal{M}_x(t_1)| &\leq \sum_{y \in [d]} \int_{t_0}^{t_1} [|\mathcal{M}_y(t) \gamma_x^*(y, \Delta_y v(t))| + |\mu_y(t) (\gamma_x^*(y, \Delta_y v(t)) - \gamma_x^*(y, \Delta_y u(t)))|] dt \\
 &\leq C \sum_{y \in [d]} \int_{t_0}^{t_1} [|\mathcal{M}_y(t)| + \mu_y(t) |\Delta_y \mathcal{Q}(t)|] dt \\
 &\leq C \int_{t_0}^{t_1} \max_{y \in [d]} |\mathcal{M}_y(t)| dt + C \sum_{y \in [d]} \int_{t_0}^{t_1} \mu_y(t) |\Delta_y \mathcal{Q}(t)| dt.
 \end{aligned}$$

Using Jensen's inequality and Gronwall's inequality:

$$\|\mathcal{M}\|_\infty \leq C \int_{t_0}^{t_1} \sqrt{\sum_{y \in [d]} \mu(t, y) |\Delta_y \mathcal{Q}(t, \cdot)|^2} dt. \quad (50)$$

In the above display, using Jensen’s inequality, (48), and a supremum bound imply,

$$\begin{aligned}
 \|\mathcal{M}\|_\infty &\leq C \int_{t_0}^{t_1} \sqrt{\sum_{y \in [d]} \mu_y(t) |\Delta_y \mathcal{Q}(t)|^2} dt \\
 &\leq C \left(\int_{t_0}^{t_1} \sum_{y \in [d]} \mu_y(t) |\Delta_y \mathcal{Q}(t)|^2 dt \right)^{1/2} \\
 &\leq C \left(\int_{t_0}^T |e(t, \cdot, v(t)) \cdot \mathcal{M}(t)| dt + |e(T, \cdot, v(T)) \cdot \mathcal{M}(T)| \right)^{1/2} \\
 &\leq C (\|e\|_\infty \|\mathcal{M}\|_\infty + \|e\|_\infty \|\mathcal{M}\|_\infty)^{1/2} \\
 &\leq C \|e\|_\infty^{1/2} \|\mathcal{M}\|_\infty^{1/2}.
 \end{aligned}$$

So by Young’s inequality,

$$\|\mathcal{M}\|_\infty \leq C \|e\|_\infty.$$

Combining this estimate with (49), and using (41), we obtain:

$$\|\mathcal{Q}\|_\infty + \|\mathcal{M}\|_\infty \leq C \|e\|_\infty.$$

The above reasoning holds for every (t_0, η) . Recalling the definitions of \mathcal{Q} , v , and (13), we obtain (42). ■

7. Numerical experiments

In this section, both the DBME and DGME algorithms were implemented in Python using TensorFlow 2. Both programs were run on the Great Lakes computing cluster, a high-performance computing cluster available for University of Michigan research. All algorithms were run on the cluster’s standard nodes, each of which consists of thirty-six cores per node. We expect further improvement in run-time given a practitioner switches to a specialized, GPU-enabled node. We also note that, in the original formulation of the DGM, Sirignano and Spiliopoulos (2018) use an LSTM-like architecture while we are using a vanilla, fully-connected, feed-forward architecture. For all networks featured, we used four layers of sixty nodes each, with sigmoid activation function, excluding the input and output layers. The output layers used ELU. Code for the DGME and DBME algorithms, data for models used, and code used to create the visualizations in this section can be found on GitHub at <https://github.com/ethanzell/DGME-and-DBME-Algorithms>.

7.1 Example of Quadratic Cost

A classical example in MFG literature is that of quadratic costs. In this section, we adapt the examples of Cecchin and Pelino (2019, Example 1) and Bayraktar and Cohen (2018, Example 3.1) to solve the corresponding master equation using the DGME and DBME algorithms.

Namely, we set the running costs and terminal cost, respectively:

$$f(x, a) := b \sum_{y \neq x} (a_y - 2)^2, \quad F(x, \eta) := \eta_x, \quad g(x, \eta) \equiv 0,$$

where $b > 0$ will be chosen later. Fix also $\mathbb{A} := [1, 3]$. In our case, the Hamiltonian is:

$$H(x, p) := \min_{a_x \in \mathbb{A}_{-x}^d} \left\{ \sum_{y \neq x} b(a_{xy} - 2)^2 + a_{xy} p_y \right\}.$$

So when the minimum is attained in the interior of \mathbb{A}_{-x}^d , we have for all $y \neq x$ that:

$$\gamma_y^*(x, p) = 2 - \frac{p_y}{2b}, \quad H(x, p) = \sum_{y \neq x} \left(2p_y - \frac{p_y^2}{4b} \right), \quad (51)$$

and so:

$$\begin{aligned} \gamma_y^*(x, \Delta_x U(t_0, \cdot, \eta)) &= 2 + \frac{U(t_0, x, \eta) - U(t_0, y, \eta)}{2b}, \\ H(x, \Delta_x U(t_0, \cdot, \eta)) &= \sum_{y \neq x} \left(2[U(t_0, y, \eta) - U(t_0, x, \eta)] - \frac{[U(t_0, y, \eta) - U(t_0, x, \eta)]^2}{4b} \right) \end{aligned}$$

Note that, should (51) hold, the Hamiltonian satisfies the regularity we assume at the very beginning of the paper in (10). We will now explicitly verify the form of H . Along the way, we will show that this form for H is not guaranteed for all problem parameter choices; however, we will prove that certain parameter choices will guarantee this form of H and thus also the regularity required for the class of problems we consider. First, we will show why (51) holds when $|p| \leq 2b$.

The Hamiltonian has minimizer $a_x^* := (a_{xy}^*)_{y \in [d]}$:

$$a_{xy}^* = \begin{cases} 3 & 2 - \frac{p_y}{2b} \geq 3, \\ 2 - \frac{p_y}{2b} & 2 - \frac{p_y}{2b} \in \mathbb{A}, \\ 1 & 2 - \frac{p_y}{2b} \leq 1. \end{cases}$$

Note that: $2 - \frac{p_y}{2b} \in \mathbb{A}$ occurs if and only if:

$$|p_y| \leq 2b. \quad (52)$$

With this observation in mind, we are now interested in finding values of b, T such that

$$|[\Delta_x U(t_0, \cdot, \eta)]_y| \leq 2b,$$

as this will lead to the desired form of the Hamiltonian in (51).

The p -argument of H will always be $\Delta_x U(t_0, \cdot, \eta)$ for some $(t_0, x, \eta) \in [0, T] \times [d] \times \mathcal{P}([d])$ and so we are interested in a bound on this quantity. We will derive a bound for $|[\Delta_x U(t_0, \cdot, \eta)]_y|$ in terms of T . Then, by selecting T and b in concert with one another, we may guarantee (52) and therefore confirm (51) holds for all inputs to H .




Figure 1: The master equation’s solution over time for $d = 2$.

Let β be the control that always chooses rate 2. Using (15), the fact that γ^* is a minimizer, the choice of f , and the fact that $|F| \leq 1$, for $y \neq x$ we have:

$$\begin{aligned}
 |[\Delta_x U(t_0, \cdot, \eta)]_y| &= \left| \mathbb{E}_{(y, \eta)} \left(\int_{t_0}^T f(\mathcal{X}_t, \gamma^*(t, \mathcal{X}_t)) + F(\mathcal{X}_t, \mathcal{L}(\mathcal{X}_t)) dt \right) \right. \\
 &\quad \left. + \mathbb{E}_{(x, \eta)} \left(\int_{t_0}^T f(\mathcal{X}_t, \gamma^*(t, \mathcal{X}_t)) + F(\mathcal{X}_t, \mathcal{L}(\mathcal{X}_t)) dt \right) \right| \\
 &\leq \left| \mathbb{E}_{(y, \eta)} \left(\int_{t_0}^T f(\mathcal{X}_t^\beta, \beta) + F(\mathcal{X}_t^\beta, \mathcal{L}(\mathcal{X}_t^\beta)) dt \right) \right| \\
 &\quad + \left| \mathbb{E}_{(x, \eta)} \left(\int_{t_0}^T f(\mathcal{X}_t^\beta, \beta) + F(\mathcal{X}_t^\beta, \mathcal{L}(\mathcal{X}_t^\beta)) dt \right) \right| \\
 &\leq 2T.
 \end{aligned}$$

So we can guarantee that $|[\Delta_x U(t_0, \cdot, \eta)]_y| \leq 2b$ by requiring that: $T \leq b$. We thus choose $T = 1/2$ and $b = 4$ and this guarantees the smoothness of the Hamiltonian for any value of the master equation solution and also for any number of states d . While we could have chosen other values, these choices result in an easily interpretable picture as a consequence of the strong cost of rate selection and more prompt horizon.

7.2 Quadratic Cost—Low Dimensional Results

Now we present numerical results for the low-dimensional cases $d = 2, 3$ corresponding to this example. We begin with the case $d = 2$; the following graphs are from the DGME.

7.2.1 THE MASTER EQUATION’S SOLUTION FOR $d = 2$, VISUALIZATIONS VIA DGME

Here, we take an in-depth look at Figure 1. The horizontal axis of the graph indicates $\mu_1(t)$, the fraction of μ in state 1 at time t . The vertical axis is the value of the approximated master equation solution. The neural network is evaluated at a series of time steps and each time step is given a distinct color, as indicated by the legend on the right.

While there is no known explicit formula to sanity check our solution, we can make several qualitative judgements at the start. At the terminal time $T = 1/2$, the graph

indicates that the function is approximately zero; this agrees with the fact that the terminal condition is $g \equiv 0$ and $U(T, x, \eta) = g(x, \eta) = 0$. Second, the graph appears to respect the monotonicity of the cost. That is, for any fixed time $t_0 \in [0, T]$ the graph satisfies:

$$U(t_0, 1, \eta) \leq U(t_0, 1, \eta + (1, -1)\varepsilon),$$

for all η in the interior of $\mathcal{P}([d])$, and whenever $\varepsilon \geq 0$ is small enough so that U is defined. The reverse inequality is true for $x = 2$. We expect this cost monotonicity to come from the mean field cost F since a representative player experiences a lower cost when a smaller fraction from the distribution μ are in his or her own state. Finally, by (15) and definition of J , we expect $U(t, x, \eta) \rightarrow 0$ monotonically as $t \rightarrow T$ and indeed this is the case.

7.2.2 MEASURING DGME AND DBME LOSSES

In Figure 2 we give loss graphs for the DGME and DBME algorithms. The DGME algorithm trains one neural network to the entire master equation and so its loss graph is straightforward: for every epoch, we average the losses over the thirty steps in that epoch and present the losses. The DBME however is the result of neural networks being trained on the output of a neural network at a future time step, and so the DBME returns N networks. Moreover, recall that the final DBME “neural network” is set to be the terminal cost and so the first neural network that is trained with an actual loss function is the penultimate network \mathcal{U}_{N-1} . In practice, we found that training \mathcal{U}_{N-1} for more epochs than the other networks resulted in better results for networks closer to the initial time. Therefore, Figure 2 also presents three graphs for the DBME trained with a constant-sized partition of $|\pi| = 0.01$. The first DBME graph depicts the losses of \mathcal{U}_{49} by itself plotted on a log scale, while the second graph gives the losses for the networks $\mathcal{U}_0, \mathcal{U}_{10}, \mathcal{U}_{20}, \mathcal{U}_{30}$, and \mathcal{U}_{40} .

7.2.3 AGREEMENT BETWEEN THE DGME AND DBME ALGORITHMS

The result from the DBME appears similar given the partition step size is small enough. In Figure 3 we graph the error between the prior DGME model and its DBME counterpart over all times in the partition $t_i \in \pi$. The error is computed as the average difference over the state space $[d] \times \mathcal{P}([d])$ and the horizontal axis is the difference on $[0, T]$.

In Figure 3, for the partitions π such that $|\pi| \leq 0.01$, the DGME and DBME algorithms agree on the value of U on average. For larger values of $|\pi|$, there is some error that propagates backward and worsens the estimate at the initial time $t = 0$.

7.2.4 THE MASTER EQUATION’S SOLUTION FOR $d = 3$, VISUALIZATIONS VIA DGME

Here, we break down the result in Figure 4. It is qualitatively similar to Figure 1 (for $d = 2$) but we present it differently. In Figure 4, each sub-plot corresponds to a different time. The simplex $\mathcal{P}([3])$ is identified with a triangle in dimension 2, where the x-axis corresponds to the density in state 2 and the y-axis corresponds to the density in state 1. The horizontal axis is $\mu_2(t)$ while the vertical axis is $\mu_1(t)$. The color intensity gives the value of the approximate master equation solution. For clarification, the bottom left corner is the point in $\mathcal{P}([d])$ where $\mu_3(t) = 1$. In the first image of Figure 4 we notice the same cost monotonicity as in the two-dimensional case. Meanwhile, we note that the final image (where $t = 0.5$) is approximately zero as desired.

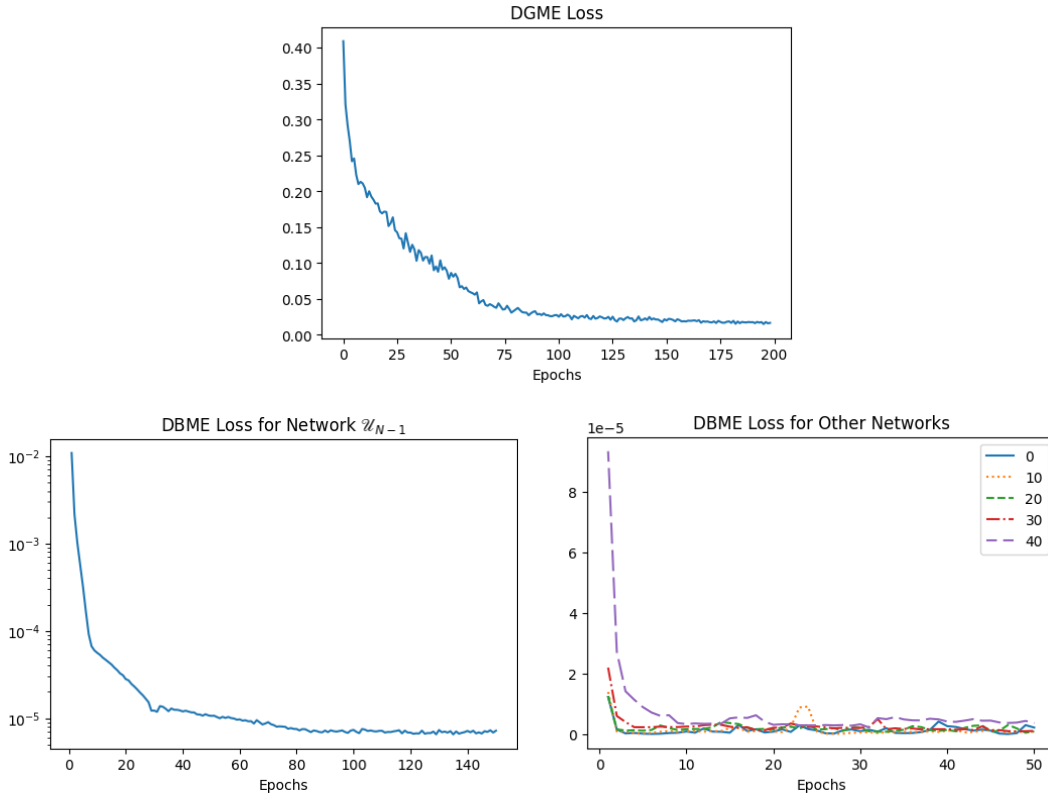


Figure 2: At the top: DGME losses for $d = 2$, averaged over thirty steps per epoch. At the bottom left, the average DBME losses for the penultimate network plotted on a log scale, and at the bottom right, the average losses for some DBME networks.




Figure 3: For $d = 2$, the average difference between the DBME and DGME solutions as the partition step decreases.

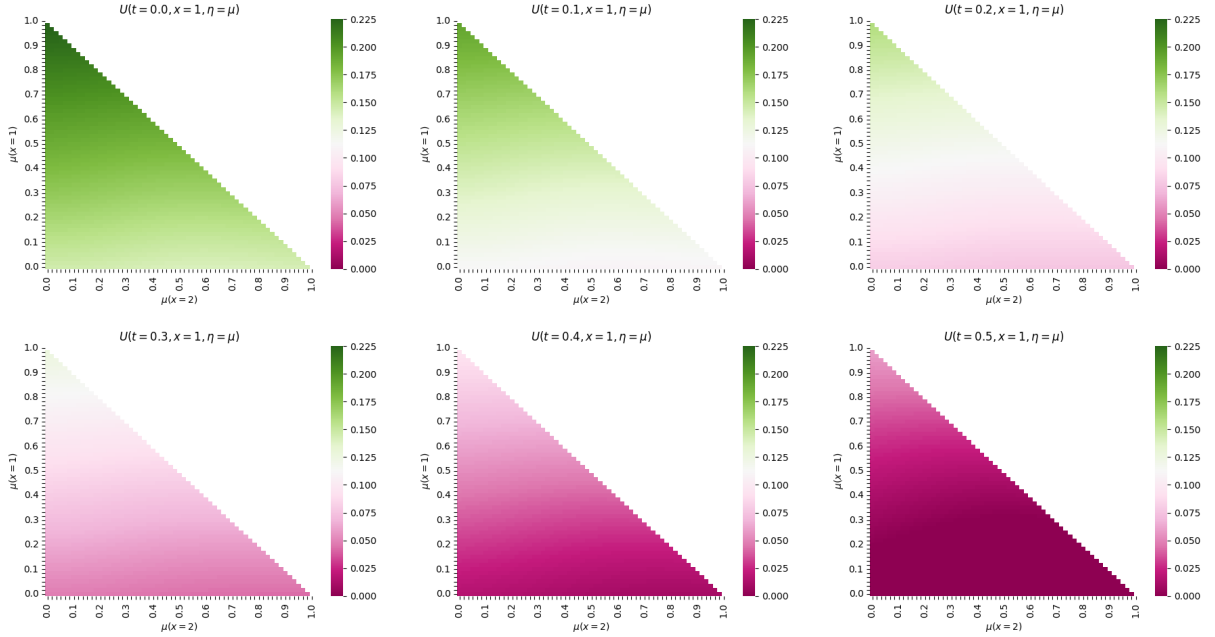


Figure 4: The master equation’s solution on the simplex $\mathcal{P}([3])$ over time.




Figure 5: Two ways of visualizing the MFG equilibrium $\mu^{0, (.7, .2, .1)}$.

Once we have approximately learned the master equation solution, we may compute the corresponding control by plugging the solution into the rate selector γ^* . Then, we are able to plot the trajectory of the player population over time. In Figure 5, we plot the trajectory as a proportion for a chosen initial state, and then on the simplex $\mathcal{P}([3])$. On $\mathcal{P}([3])$, the dark purple end of the line corresponds to the initial state $\mu(t = 0)$, while the light yellow end corresponds to the terminal state $\mu(t = T)$.




Figure 6: Comparing the increase in run-time of each algorithm as dimension increases, and when requiring the same loss.

7.3 Quadratic Cost—High Dimensional Results

How the DGME and DBME deal with the curse of dimensionality is of particular interest. To study the run-time efficiency as dimension increases, we compute the relative increase in run-time for each dimension by dividing the corresponding running time by the running time required for the previous dimension.

In Figure 6, we note that the DBME run-time increases at around the rate of increase in dimension; meanwhile, the DGME run-time increases faster with an equivalent increase in dimension. The lower relative increase for the DBME run-time indicates that it may be better suited to high-dimensional problems. Moreover, the rate of relative increase for the DBME agrees with other known deep backward schemes; see for example the deep backward scheme for an optimal switching problem studied in (Bayraktar et al., 2022, Figure 2).

We also note that the size of our state space increases quadratically in d . Namely, we are considering $[d] \times \mathcal{P}([d])$. The simplex $\mathcal{P}([d])$ is $(d-1)$ -dimensional and so the size of our state space is $d^2 - d = \mathcal{O}(d^2)$.

7.4 Cybersecurity Example

In this section, we study another popular MFG model that represents the cybersecurity of a network of computers. The cybersecurity model was introduced by Kolokoltsov and Bensoussan (2016), revisited by Carmona and Delarue (2018, Section 7.2.3), and computed numerically by Laurière (2021, Section 7.2). We recall the problem here, detail its master equation, and solve using the DGME and DBME methods.

In the cybersecurity MFG model, a continuum of agents try to minimize their costs that come from two sources: being infected with a computer virus, and paying for computer security software. An agent’s computer can be infected by another infected computer or by a hacker (considered exogenous to the model). Intuitively, the mean field interaction comes through the infection of a computer by other infected computers. Once infected, a

computer can recover back to its original state. Agents that elect to defend their computers with the software are infected at a slower rate than undefended computers.

So, every agent has two internal states: defended/undefended (D, U) and susceptible/infected (S, I) . Hence there are four states $(d = 4)$: $[d] := \{DS, US, DI, UI\}$. Every agent pays a cost while infected, $k_I > 0$, and a cost while defending, $k_D > 0$. The agent pays the most cost per second when both infected and defending. Specifically, the running cost f is given by:

$$f(x, a) := k_D \mathbb{1}_{\{DS, DI\}}(x) + k_I \mathbb{1}_{\{DI, UI\}}(x).$$

The mean field cost and terminal cost are both set to zero, $F = g \equiv 0$. The agent's only choice is to switch his or her defended or undefended status; this occurs with rate ρ , and the corresponding action set is $\mathbb{A} = \{0, 1\}$.

The rate a hacker hacks into any computer is determined by the parameter v_H which for defended or undefended computers is augmented by q_H^D and q_H^U , respectively. Correspondingly, the recovery parameters are q_R^D and q_R^U . Infections from defended and undefended computers are augmented by the parameters $\beta_{DD}, \beta_{UU}, \beta_{UD}, \beta_{DU} > 0$ in the following rate matrix. Namely, agents transition according to the rate matrix:

$$(M(x, y; a))_{x, y \in [d]} := \begin{bmatrix} * & \rho a & P_{DS, DI}^\eta & 0 \\ \rho a & * & 0 & P_{US, UI}^\eta \\ q_R^D & 0 & * & \rho a \\ 0 & q_R^U & \rho a & * \end{bmatrix},$$

where

$$\begin{aligned} P_{DS, DI}^\eta &:= v_H q_H^D + \beta_{DD} \eta(DI) + \beta_{UD} \eta(UI), \\ P_{US, UI}^\eta &:= v_H q_H^U + \beta_{UU} \eta(UI) + \beta_{DU} \eta(DI), \end{aligned}$$

and where the asterisk denotes the negative of the sum of the elements along that row. Note that agents can only determine their rate matrix by choosing either $a = 0$ or 1 and all other parameters are fixed. That said, $P_{DS, DI}^\eta$ and $P_{US, UI}^\eta$ change based on the proportion of infected undefended and defended agents' computers, respectively; it is through these terms that the agents experience a mean field-type interaction.

With this said, we can write down the pre-Hamiltonians, the functions inside the Hamiltonian's minimum, for each state:

$$\begin{aligned} \tilde{H}(DS, \Delta_{DS}U(t, \cdot, \eta); a) &:= k_D + P_{DS, DI}^\eta(U(t, DI, \eta) - U(t, DS, \eta)) \\ &\quad + \rho a(U(t, US, \eta) - U(t, DS, \eta)), \\ \tilde{H}(US, \Delta_{US}U(t, \cdot, \eta); a) &:= \rho a(U(t, DS, \eta) - U(t, US, \eta)) + P_{US, UI}^\eta(U(t, UI, \eta) - U(t, US, \eta)), \\ \tilde{H}(DI, \Delta_{DI}U(t, \cdot, \eta); a) &:= k_D + k_I + \rho a(U(t, UI, \eta) - U(t, DI, \eta)) \\ &\quad + q_R^D(U(t, DS, \eta) - U(t, DI, \eta)), \\ \tilde{H}(UI, \Delta_{UI}U(t, \cdot, \eta); a) &:= k_I + \rho a(U(t, DI, \eta) - U(t, UI, \eta)) + q_R^U(U(t, US, \eta) - U(t, UI, \eta)). \end{aligned}$$

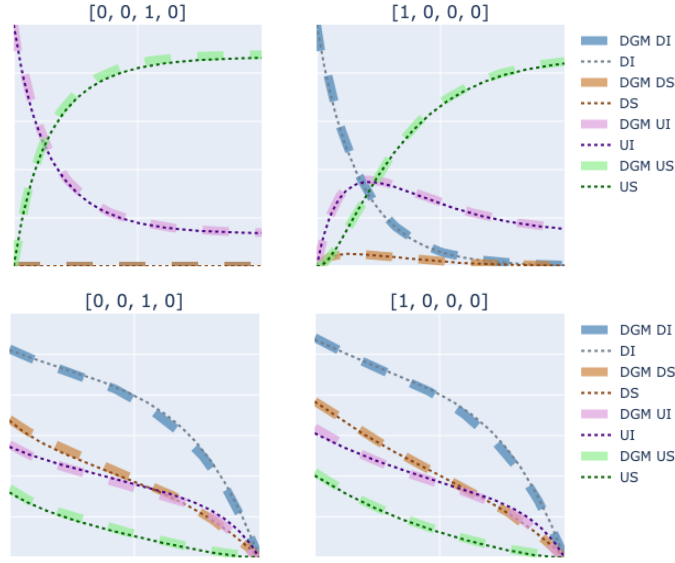


Figure 7: The top row gives two examples of μ given different initial conditions. The bottom row gives the corresponding u for these measures.

Taking the minimum over $a \in \{0, 1\}$ yields the usual Hamiltonian. We can then write the optimal rate matrix:

$$(\gamma_y^*(x, \Delta_x U(t, \cdot, \eta)))_{x, y \in [d]} := \begin{bmatrix} * & \rho \mathbb{1}_{\{U(t, US, \eta) < U(t, DS, \eta)\}} & P_{DS, DI}^\eta & 0 \\ \rho \mathbb{1}_{\{U(t, DS, \eta) < U(t, US, \eta)\}} & * & 0 & P_{US, UI}^\eta \\ q_R^D & 0 & * & \rho \mathbb{1}_{\{U(t, UI, \eta) < U(t, DI, \eta)\}} \\ 0 & q_R^U & \rho \mathbb{1}_{\{U(t, DI, \eta) < U(t, UI, \eta)\}} & * \end{bmatrix}.$$

So, we can write the master equation as: for all $t \in [0, T]$, $x \in [d]$, and $\eta \in \mathcal{P}([d])$,

$$\begin{aligned}
 \partial_t U(t, x, \eta) + \sum_{y, z \in [d]} \eta_y D_{yz}^\eta U(t, x, \eta) \gamma_z^*(y, \Delta_y U(t, \cdot, \eta)) + H(x, \Delta_x U(t, \cdot, \eta)) &= 0, \\
 U(T, x, \eta) &= 0.
 \end{aligned}$$

By running either algorithm then, we can reconstruct the optimal control and hence the trajectory of the measure. Recall the MFG system (12) and its solution (u, μ) (removing the superscripts). We demonstrate the measure trajectory μ in the first row of Figure 7, using the solution computed by the DGME. The corresponding cost u is displayed in the second row of Figure 7. In Figure 7, the legend depicts the four states DS , US , DI , and UI ; the prefix of DGM shows the DGME's result against a finite difference method solver. Further results for other initial distributions are provided in Appendix 8.

8. Conclusion

Motivated by the computation of Nash equilibria in games with many players, we proposed two algorithms, the DBME and the DGME, that can numerically solve the master equation of finite-state MFGs. For each method, we proved two results: the loss can be made arbitrarily small by some neural network, and conversely, if a neural network makes the loss small, then this neural network is close to the real solution of the master equation. Besides the theoretical analysis, we provided two in-depth numerical examples applying the methods and compare them at different dimensions. We found that the relative increase in computation time is superior for the DBME. Our numerical tests show that the proposed methods allow us to learn the master equation solution for any initial distribution, in contrast with traditional forward-backward methods, which work only for a single fixed initial distribution. Hence our methods open new potential applications of MFGs, when the initial distribution is not known in advance or when the population distribution may deviate from its expected trajectory.

Several directions are left for future work. First, our methods have the same limitations as the DBDP and the DGM, and will benefit from progress made on these methods. For instance, performance might be improved with a better understanding of how to choose the neural network architecture depending on the PDE under consideration. It should also be possible to improve the training speed by a suitable choice of sampling method. In particular, we leave for future work the incorporation of active sampling to the proposed methods.

On the theoretical side, we have mostly focused on the approximation error. It would be interesting to obtain an error rate in terms of the number of neurons and the dimension, that is, the number of states of the MFG. At this stage, numerical evidence suggests that our approaches work well even with a large number of states. However, as for many other methods, such as the DGM and the DBDP in particular, it is not clear how to prove rigorously that neural networks break the curse of dimensionality. Furthermore, it would also be interesting to study other error sources, whether from sampling regimes or gradient descent convergence. Ultimately, the goal would be to comprehensively understand the generalization error of neural network-based methods for master equations.

Last, real-world applications are beyond the scope of the present paper. This would involve studying in detail the link between the master equation in MFGs and systems of Bellman PDEs in finite-player games. Such systems can also be solved by deep learning. When the number of players is small, it is doable and more accurate to solve directly the PDE system. But as the number of players grows, it becomes easier to solve the master equation, and the MFG provides a good approximation of the finite-player game equilibrium. However, the precise trade-off depends on the model under consideration, which is left for future work in a case-by-case study.

References

Yves Achdou and Mathieu Laurière. Mean field games and applications: Numerical aspects. *Mean Field Games*, page 249, 2020.

- Ali Al-Aradi, Adolfo Correia, Gabriel Jardim, Danilo de Freitas Naiff, and Yuri Saporito. Extensions of the deep galerkin method. Applied Mathematics and Computation, 430: 127287, 2022.
- Robert J. Aumann. Markets with a continuum of traders. Econometrica, 32:39–50, 1964. ISSN 0012-9682.
- E. Bayraktar and A. Cohen. Analysis of a finite state many player game using its master equation. SIAM Journal on Control and Optimization, 56(5):3538–3568, 2018.
- E. Bayraktar, A. Budhiraja, and A. Cohen. A numerical scheme for a mean field game in some queueing systems based on Markov chain approximation method. SIAM Journal on Control and Optimization, 56, 2017.
- Erhan Bayraktar, Asaf Cohen, and April Nellis. A neural network approach to high-dimensional optimal switching problems with jumps in energy markets, 2022.
- Alain Bensoussan, Jens Frehse, and Phillip Yam. Mean field games and mean field type control theory. SpringerBriefs in Mathematics. Springer, New York, 2013.
- Charles Bertucci and Alekos Cecchin. Mean field games master equations: from discrete to continuous state space. arXiv e-prints, art. arXiv:2207.03191, July 2022.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. J. Polit. Econ., 81(3):637–654, 1973. ISSN 0022-3808,1537-534X.
- Haoyang Cao, Xin Guo, and Mathieu Laurière. Connecting GANs, MFGs, and OT. arXiv preprint arXiv:2002.04112, 2020.
- Pierre Cardaliaguet, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. The master equation and the convergence problem in mean field games, volume 201 of Annals of Mathematics Studies. Princeton University Press, 2019.
- René Carmona and François Delarue. Probabilistic theory of mean field games with applications. I, volume 83 of Probability Theory and Stochastic Modelling. Springer, Cham, 2018.
- René Carmona and Mathieu Laurière. Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games I: The ergodic case. SIAM Journal on Numerical Analysis, 59(3):1455–1485, 2021.
- René Carmona and Mathieu Laurière. Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: II—the finite horizon case. The Annals of Applied Probability, 32(6):4065–4105, 2022.
- Alekos Cecchin and Guglielmo Pelino. Convergence, fluctuations and large deviations for finite state mean field games via the master equation. Stochastic Processes and their Applications, 129(11):4510 – 4555, 2019.
- A. Cohen and E. Zell. Analysis of the Finite-State Ergodic Master Equation. Applied Mathematics and Optimization, 87, 2023.

- Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. Communications of the ACM, 52(2):89–97, 2009.
- Gokce Dayanikli, Mathieu Lauriere, and Jiacheng Zhang. Deep learning for population-dependent controls in mean field control problems. Accepted for an extended abstract to AAMAS'23 (arXiv preprint arXiv:2306.04788), 2023.
- François Delarue, Daniel Lacker, and Kavita Ramanan. From the master equation to mean field game limit theory: a central limit theorem. Electron. J. Probab., 24, 2019.
- François Delarue, Daniel Lacker, and Kavita Ramanan. From the master equation to mean field game limit theory: Large deviations and concentration of measure. Annals of Probability, 48(1):211–263, 2020.
- Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. Commun. Math. Stat., 5(4):349–380, 2017.
- Lawrence C. Evans. Partial differential equations, volume 19 of Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, second edition, 2010.
- Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. Advances in Neural Information Processing Systems, 32, 2019.
- Jean-Pierre Fouque and Zhaoyu Zhang. Deep learning methods for mean field control problems with delay. Frontiers in Applied Mathematics and Statistics, 6(11), 2020.
- Wilfrid Gangbo and Alpár R. Mészáros. Global well-posedness of master equations for deterministic displacement convex potential mean field games. Comm. Pure Appl. Math., 75(12):2685–2801, 2022. ISSN 0010-3640,1097-0312.
- Wilfrid Gangbo, Alpár R. Mészáros, Chenchen Mou, and Jianfeng Zhang. Mean field games master equations with nonseparable Hamiltonians and displacement monotonicity. Ann. Probab., 50(6):2178–2217, 2022. ISSN 0091-1798,2168-894X. doi: 10.1214/22-aop1580. URL <https://doi.org/10.1214/22-aop1580>.
- Maximilien Germain, Mathieu Laurière, Huyên Pham, and Xavier Warin. Deepsets and their derivative networks for solving symmetric PDEs. Journal of Scientific Computing, 91(2):63, 2022a.
- Maximilien Germain, Joseph Mikael, and Xavier Warin. Numerical resolution of McKean-Vlasov FBSDEs using neural networks. Methodology and Computing in Applied Probability, 24(4):2557–2586, 2022b.
- Maximilien Germain, Huyên Pham, and Xavier Warin. Rate of convergence for particle approximation of PDEs in wasserstein space. Journal of Applied Probability, 59(4):992–1008, 2022c.

- Diogo A. Gomes, Joana Mohr, and Rafael Rigão Souza. Discrete time, finite state space mean field games. J. Math. Pures Appl. (9), 93(3):308–328, 2010.
- Diogo A. Gomes, Joana Mohr, and Rafael Rigão Souza. Continuous time finite state mean field games. Appl. Math. Optim., 68(1):99–143, 2013.
- Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J. Cree. Regularisation of neural networks by enforcing Lipschitz continuity. Mach. Learn., 110(2):393–416, 2021.
- P. Jameson Graber and Alpár R. Mészáros. On monotonicity conditions for mean field games. J. Funct. Anal., 285(9):Paper No. 110095, 45, 2023. ISSN 0022-1236,1096-0783. doi: 10.1016/j.jfa.2023.110095. URL <https://doi.org/10.1016/j.jfa.2023.110095>.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. Neural Networks, 3:551–560, 1990.
- Ruimeng Hu and Mathieu Laurière. Recent developments in machine learning methods for stochastic control and games. arXiv preprint arXiv:2303.10257, 2023.
- Minyi Huang, Roland P. Malhamé, and Peter E. Caines. Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. Commun. Inf. Syst., 6(3):221–251, 2006.
- Côme Huré, Huyên Pham, and Xavier Warin. Deep backward schemes for high-dimensional nonlinear PDEs. Math. Comp., 89(324):1547–1579, 2020.
- Peter Knabner and Lutz Angermann. Numerical methods for elliptic and parabolic partial differential equations, volume 44 of Texts in Applied Mathematics. Springer, Cham, extended edition, 2021. With contributions by Andreas Rupp.
- V. N. Kolokoltsov and A. Bensoussan. Mean-field-game model for botnet defense in cybersecurity. Appl. Math. Optim., 74(3):669–692, 2016.
- Jean-Michel Lasry and Pierre-Louis Lions. Jeux à champ moyen. II. Horizon fini et contrôle optimal. C. R. Math. Acad. Sci. Paris, 343(10):679–684, 2006.
- Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. Jpn. J. Math., 2(1):229–260, 2007.
- Mathieu Laurière. Numerical methods for mean field games and mean field type control. In Mean field games, volume 78 of Proc. Sympos. Appl. Math., pages 221–282. Amer. Math. Soc., Providence, RI, 2021.
- Pierre-Louis Lions. Estimées nouvelles pour les équations quasilinéaires. Seminar in Applied Mathematics at the Collège de France., 2014.
- Chenchen Mou and Jianfeng Zhang. Mean Field Game Master Equations with Antimonotonicity Conditions. arXiv e-prints, art. arXiv:2201.10762, January 2022. doi: 10.48550/arXiv.2201.10762.

- John Nash. Non-cooperative games. Annals of mathematics, pages 286–295, 1951.
- Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust neural networks using lipschitz bounds. IEEE Control Systems Letters, 6:121–126, 2021.
- Sarah Perrin, Mathieu Laurière, Julien Pérolat, Romuald Élie, Matthieu Geist, and Olivier Pietquin. Generalization in mean field games by learning master policies. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, pages 9413–9421, 2022.
- Huyên Pham, Xavier Warin, and Maximilien Germain. Neural networks-based backward scheme for fully nonlinear PDEs. Partial Differ. Equ. Appl., 2(1):Paper No. 16, 24, 2021.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational physics, 378:686–707, 2019.
- Lars Ruthotto, Stanley J Osher, Wuchen Li, Levon Nurbekyan, and Samy Wu Fung. A machine learning framework for solving high-dimensional mean field game and mean field control problems. Proceedings of the National Academy of Sciences, 117(17):9183–9193, 2020.
- David Schmeidler. Equilibrium points of nonatomic games. J. Statist. Phys., 7:295–300, 1973. ISSN 0022-4715.
- Shai Shalev-Shwartz and Yonatan Wexler. Minimizing the maximal loss: How and why. In Proceedings of The 33rd International Conference on Machine Learning, volume 48 of Proceedings of Machine Learning Research, pages 793–801. PMLR, 2016.
- Justin Sirignano and Konstantinos Spiliopoulos. DGM: a deep learning algorithm for solving partial differential equations. J. Comput. Phys., 375:1339–1364, 2018.
- Paul Tseng. An incremental gradient(-projection) method with momentum term and adaptive stepsize rule. SIAM J. Optim., 8(2):506–531, 1998.

Appendix A. Proof of Lemma 11

Let $|\cdot|_1$ denote the L^1 -norm in \mathbb{R}^d and define $A := |\mathcal{K}|$. Choose:

$$\bar{\kappa} \in \arg \max_{\kappa \in \mathcal{P}([d])} \tilde{L}_i(\kappa; \theta^i),$$

so that:

$$\tilde{L}_i(\bar{\kappa}; \theta^i) = L_i(\theta^i).$$

Since $\mathcal{K} \subsetneq \mathcal{P}([d])$,

$$0 \leq \tilde{L}_i(\bar{\kappa}; \theta^i) - \tilde{L}_i(\hat{\kappa}; \theta^i).$$

Pick $\kappa_0 \in \mathcal{K}$ such that:

$$\kappa_0 \in \arg \min_{\kappa \in \mathcal{K}} |\kappa - \bar{\kappa}|.$$

Then,

$$0 \leq \tilde{L}_i(\bar{\kappa}; \theta^i) - \tilde{L}_i(\hat{\kappa}; \theta^i) \leq \tilde{L}_i(\bar{\kappa}; \theta^i) - \tilde{L}_i(\kappa_0; \theta^i).$$

Using the above display, the definition of $\bar{\kappa}$, and the assumption that the neural networks are uniformly Lipschitz, as well as the Lipschitz continuity of \bar{H} , there exists $C > 0$ depending on these Lipschitz constants such that:

$$|L_i(\theta^i) - \tilde{L}_i(\hat{\kappa}; \theta^i)| \leq C|\bar{\kappa} - \kappa_0|. \quad (53)$$

So it suffices to bound the right hand side of the above display in expectation. Let $p : \mathbb{R}^d \rightarrow \mathbb{R}^{d-1}$ be the projection onto the first $(d-1)$ -coordinates of a vector in \mathbb{R}^d . Using that $\bar{\kappa}, \kappa_0 \in \mathcal{P}([d])$ and the Cauchy–Schwartz inequality:

$$\begin{aligned} |\bar{\kappa} - \kappa_0|^2 &= \sum_{x \in [d]} (\bar{\kappa}_x - \kappa_{0,x})^2 \\ &= \sum_{x=1}^{d-1} (\bar{\kappa}_x - \kappa_{0,x})^2 + \left(1 - \sum_{x=1}^{d-1} \bar{\kappa}_x - 1 + \sum_{x=1}^{d-1} \kappa_{0,x}\right)^2 \\ &= \sum_{x=1}^{d-1} (\bar{\kappa}_x - \kappa_{0,x})^2 + \left(\sum_{x=1}^{d-1} (\kappa_{0,x} - \bar{\kappa}_x)\right)^2 \\ &\leq \sum_{x=1}^{d-1} (\bar{\kappa}_x - \kappa_{0,x})^2 + (d-1) \sum_{x=1}^{d-1} (\bar{\kappa}_x - \kappa_{0,x})^2 \\ &= d \sum_{x=1}^{d-1} (\bar{\kappa}_x - \kappa_{0,x})^2 \\ &= d|p(\bar{\kappa} - \kappa_0)|^2. \end{aligned}$$

Plugging this result into (53),

$$|L_i(\theta^i) - \tilde{L}_i(\hat{\kappa}; \theta^i)| \leq C\sqrt{d}|p(\bar{\kappa} - \kappa_0)|.$$

Recalling that κ_0 is a minimum, noting that the difference is bounded above by the case when $p(\bar{\kappa}) = 0$, then since $\kappa_x \leq 1$ for all $x \in [d]$, we have:

$$\mathbb{E}|p(\bar{\kappa} - \kappa_0)| = \mathbb{E}\left[\min_{\kappa \in \mathcal{K}} |p(\bar{\kappa} - \kappa)|\right] \leq \mathbb{E}\left[\min_{\kappa \in \mathcal{K}} |p(\kappa)|\right] \leq \mathbb{E}\left[\min_{\kappa \in \mathcal{K}} |p(\kappa)|_1\right].$$

Combining the two previous displays,

$$\mathbb{E}|L_i(\theta^i) - \tilde{L}_i(\hat{\kappa}; \theta^i)| \leq C\sqrt{d}\mathbb{E}\left[\min_{\kappa \in \mathcal{K}} |p(\kappa)|_1\right]. \quad (54)$$

Recall that sampling κ uniformly on the simplex $\mathcal{P}([d])$ is equivalent to creating d exponentially distributed random variables U_1, \dots, U_d (with rate 1), and then setting each component of $\kappa := (\kappa_x)_{x \in [d]}$:

$$\kappa_x := \frac{U_x}{\sum_{y \in [d]} U_y}.$$

Therefore,

$$|p(\kappa)|_1 = \sum_{x=1}^{d-1} \kappa_x \sim \beta(d-1, 1),$$

where by the last part of the expression we mean that $p(\kappa)$ has beta distribution with parameters $d-1$ and 1. So in terms of the incomplete beta function $B(m; d-1, 1)$ and the beta function $B(d-1, 1)$, we can state the cumulative distribution function as:

$$\mathbb{P}(|p(\kappa)|_1 \leq m) = \frac{B(m; d-1, 1)}{B(d-1, 1)} = \frac{m^{d-1}}{d-1}.$$

Using this fact, an integral identity, the fact that the κ 's are sampled i.i.d.:

$$\mathbb{E} \left[\min_{\kappa \in \mathcal{K}} |p(\kappa)|_1 \right] = \int_0^1 \mathbb{P}(|p(\kappa)|_1 \geq m)^A dm = \int_0^1 \left(1 - \frac{m^{d-1}}{d-1}\right)^A dm. \quad (55)$$

Performing the substitution $\ell := m^{d-1}/(d-1)$, we get:

$$\int_0^1 \left(1 - \frac{m^{d-1}}{d-1}\right)^A dm = (d-1)^{1/(d-1)-1} \int_0^{1/(d-1)} (1-\ell)^A \ell^{\frac{1}{d-1}-1} d\ell.$$

Notice that the integrand on the right hand side is the unnormalized probability density function of a beta distribution with parameters $1/(d-1)$ and $A+1$. Therefore the previous display is bounded above by extending the integral, then using Stirling's approximation:

$$\begin{aligned} (d-1)^{1/(d-1)-1} \int_0^1 (1-\ell)^A \ell^{\frac{1}{d-1}-1} d\ell &= (d-1)^{1/(d-1)-1} \frac{\Gamma((d-1)^{-1})\Gamma(A+1)}{\Gamma(A+1+(d-1)^{-1})} \\ &\approx (d-1)^{1/(d-1)-1} \Gamma((d-1)^{-1}) \sqrt{\frac{A+1}{A+1+(d-1)^{-1}}} \\ &\quad \times \left(\frac{e}{A+1+(d-1)^{-1}}\right)^{\frac{1}{d-1}} \left(\frac{A+1}{A+1+(d-1)^{-1}}\right)^{A+1} \\ &\leq \tilde{C}_d \left(\frac{e}{A+1+(d-1)^{-1}}\right)^{\frac{1}{d-1}} \\ &= \tilde{C}_d \mathcal{O}\left(\frac{1}{A^{1/(d-1)}}\right), \end{aligned} \quad (56)$$

where:

$$\tilde{C}_d := (d-1)^{1/(d-1)-1} \Gamma((d-1)^{-1}),$$

and where we used the facts that:

$$\left(\frac{A+1}{A+1+(d-1)^{-1}}\right)^{A+1} \leq 1 \quad \text{and} \quad \sqrt{\frac{A+1}{A+1+(d-1)^{-1}}} \leq 1.$$

So plugging in (55) into (54), using (56), and setting $C_d := C\tilde{C}_d\sqrt{d}$, we get:

$$\mathbb{E}|L_i(\theta^i) - \tilde{L}(\hat{\kappa}; \theta^i)| \leq \frac{C_d}{A^{1/(d-1)}}.$$

■

Appendix B. Additional Results for the Cybersecurity Example

In Figures 8 and 9, we provide the plots for the cybersecurity example, but for many more initial conditions than offered earlier. The same labeling convention as in Figure 7 is adopted: the legend depicts the four states DS , US , DI , and UI ; the prefix of DGM shows the DGME's result against a finite difference method solver. In each figure, we provide a plot for every initial condition $\eta \in \mathcal{P}([4])$ such that $\eta_x \in \{0, 1/4, 1/2, 3/4, 1\}$ for all $x \in [4]$. Figure 8 displays the equilibrium trajectory $\mu^\eta(t)$ according to the η specified while Figure 9 displays the cost along that equilibrium, $U(t, \cdot, \mu^\eta(t))$. The shape of the cost shows that the costliest situation for every initial condition and every time occurs when a player is both defended and infected. The cheapest situation is undefended and susceptible. For all initial conditions, there is a point in time when it is cheaper to be defended and susceptible than undefended and infected. Intuitively, since being infected carries a steep cost, one may want to pay for the computer's cybersecurity software to have a lower rate of incoming infection; if infected, transitioning to susceptible is more costly toward the end of the scenario because paying for the security software is costly and, without paying for the software, the transition to susceptible is slow. That said, the difference in total cost between defended susceptible and undefended infected in the latter part of the game is small, and all costs eventually converge to zero.

Acknowledgments

We thank the Action Editor, Maxim Raginsky, and the anonymous referee for their insightful comments, which helped us improve our paper.

TWO DEEP METHODS FOR THE MASTER EQUATION

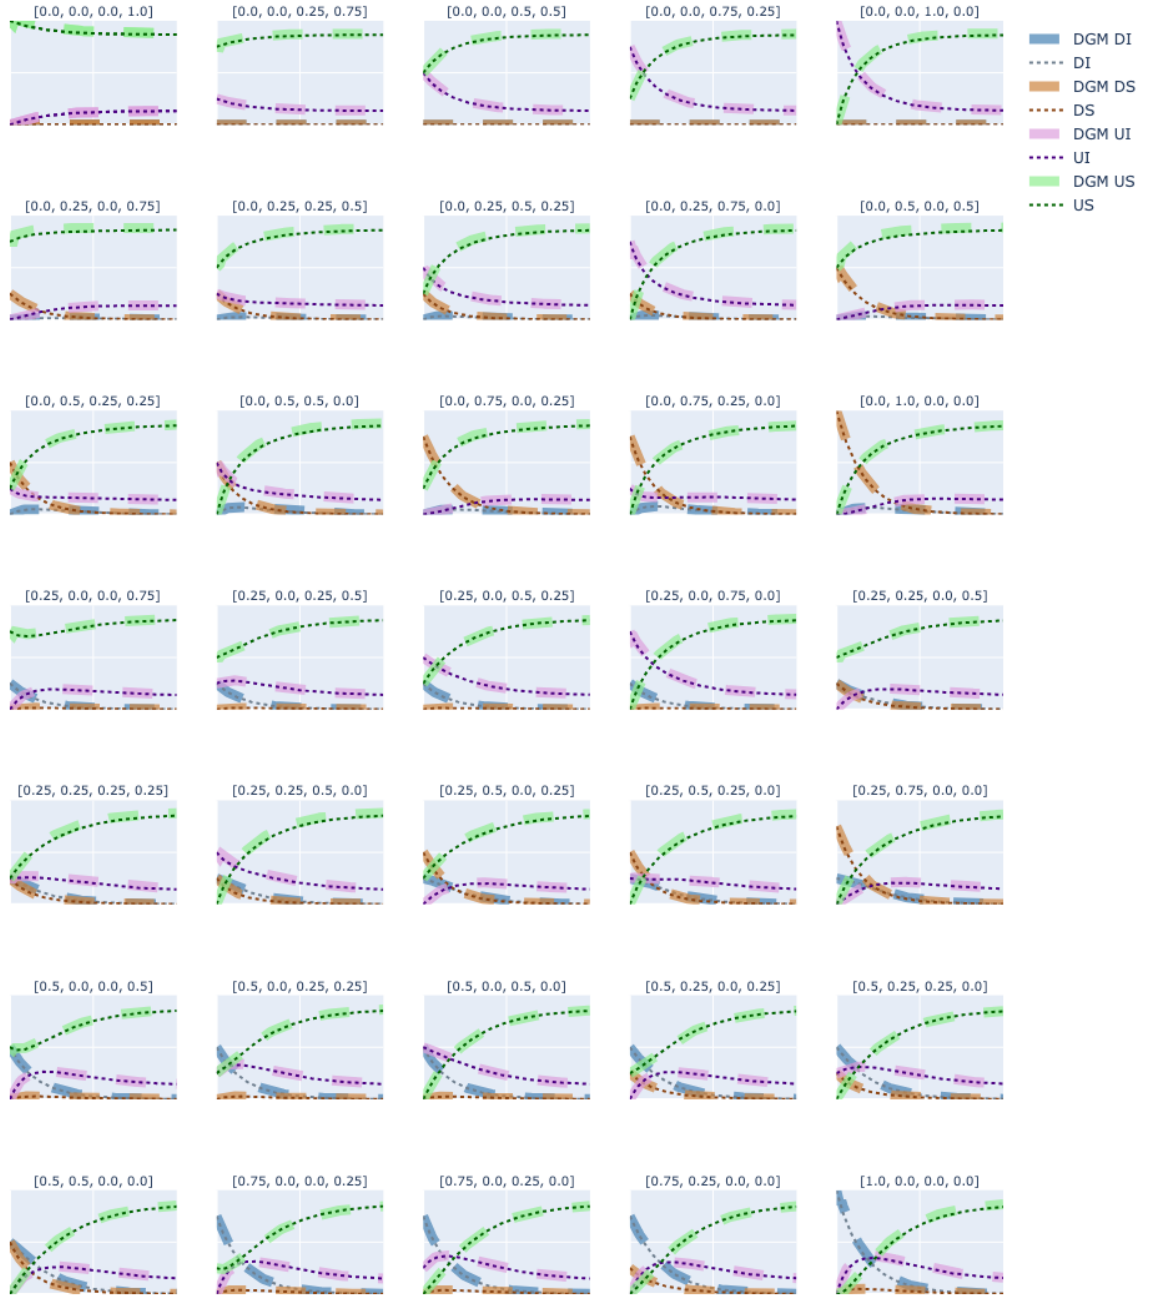


Figure 8: MFG equilibria μ for various initial conditions

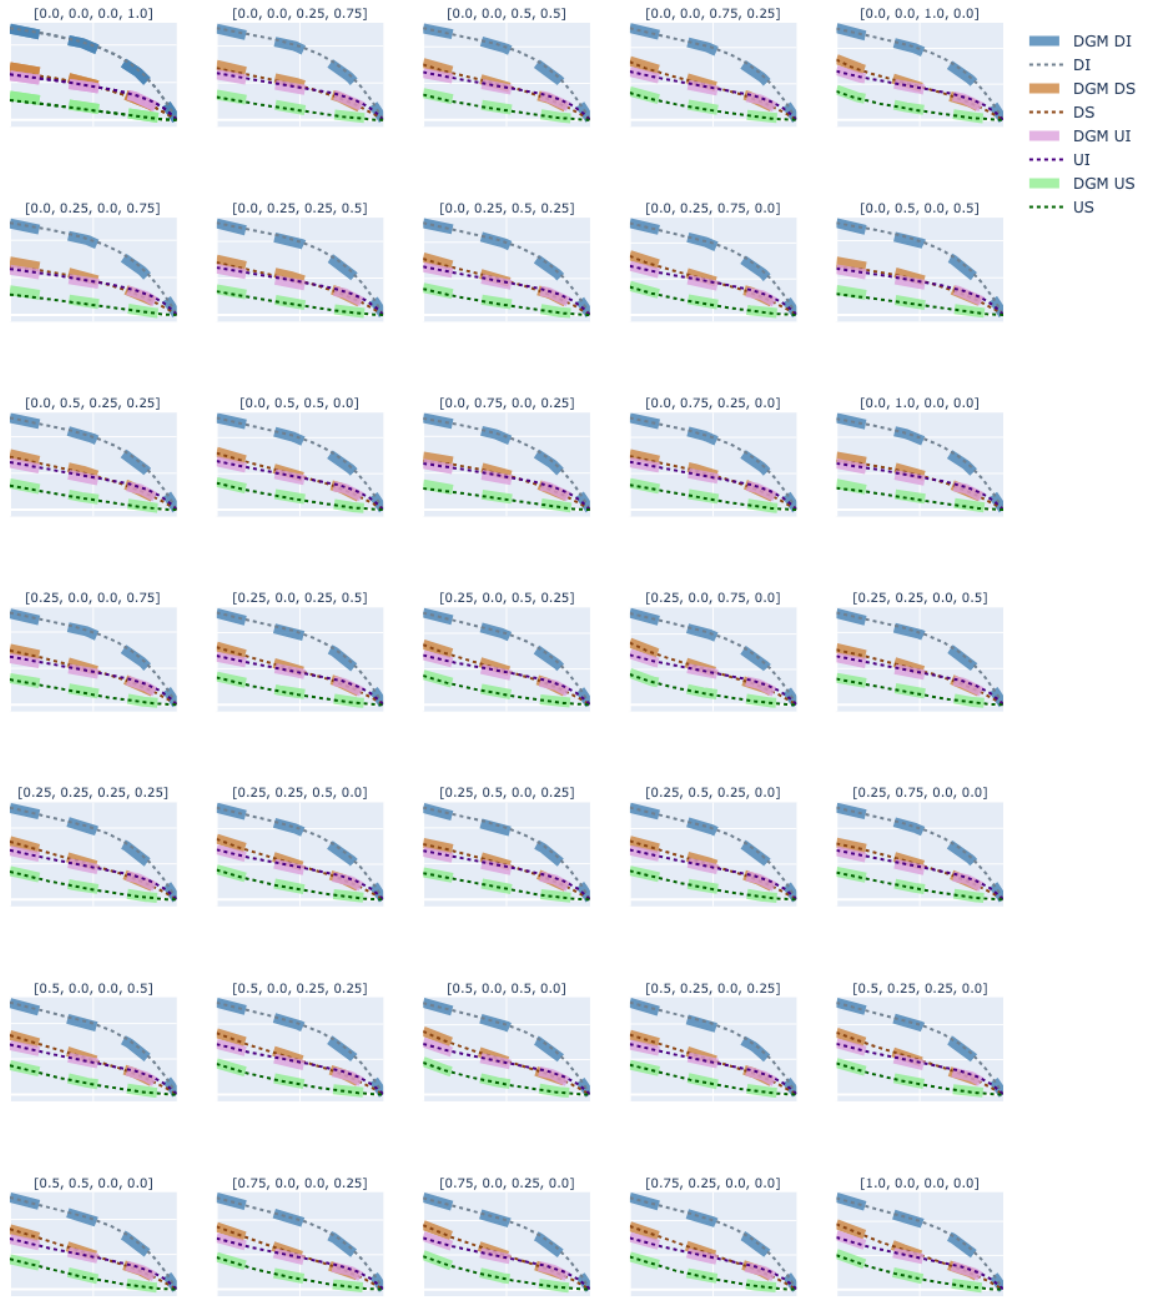


Figure 9: Cost trajectories u for various initial conditions