

# Dynamic Population Games: A Tractable Intersection of Mean-Field Games and Population Games

Ezzat Elokda, Saverio Bolognani, Andrea Censi, Florian Dörfler and Emilio Frazzoli

**Abstract**—In many real-world large-scale decision problems, self-interested agents have individual dynamics and optimize their own long-term payoffs. Important examples include the competitive access to shared resources (e.g., roads, energy, or bandwidth) but also non-engineering domains like epidemic propagation and control. These problems are natural to model as *mean-field games*. However, existing mathematical formulations of mean field games have had limited applicability in practice, since they require solving non-standard initial-terminal-value problems that are tractable only in limited special cases. In this letter, we propose a novel formulation, along with computational tools, for a practically relevant class of *Dynamic Population Games (DPGs)*, which correspond to discrete-time, finite-state-and-action, stationary mean-field games. Our main contribution is a mathematical reduction of *Stationary Nash Equilibria (SNE)* in DPGs to standard *Nash Equilibria (NE)* in static population games. This reduction is leveraged to guarantee the existence of a SNE, develop an *evolutionary dynamics*-based SNE computation algorithm, and derive simple conditions that guarantee stability and uniqueness of the SNE. Additionally, DPGs enable us to tractably incorporate multiple agent types, which is of particular importance to assess fairness concerns in resource allocation problems. We demonstrate our results by computing the SNE in two complex application examples: fair resource allocation with heterogeneous agents and control of epidemic propagation.

Open source software for SNE computation: <https://gitlab.ethz.ch/elokdae/dynamic-population-games>

## I. INTRODUCTION

*Mean field games*, introduced under that name in [1], [2], and otherwise referred to as *anonymous sequential games* [3], [4], model competitive settings in large populations of selfish agents, who have individual state dynamics and optimize individual long-term average or discounted payoffs. The microscopic state dynamics and payoffs of different agents are coupled through the *mean field*, that is, the macroscopic distribution of states (and in some works, the distribution of states and actions [3], [4], [5]). Mean field games occur naturally in many important real-world problems. For example, in an epidemic outbreak, agents have individual infection states, and the probability of getting infected depends on the proportion of infected agents, as well as whether the agents adhere to social distancing policies [6], [7]. In a congested power grid, an Electric Vehicle (EV) with individual state of charge experiences different charging rates depending on the proportion of EV's charging at

the same time [8], [9], [10]. Moreover, recent works have employed mean field games for the design of fair, non-monetary mechanisms in shared resource allocation settings, e.g., for traffic congestion management [11]. In this case, the individual state is a private, stochastic urgency for the shared resource, as well as a scarce budget of tokens used to bid for the resource, whose dynamics depend on the macroscopic bidding behavior of others [12], [13].

In all of these applications, the availability of practical computation tools to predict the *Mean Field Equilibrium (MFE)* is instrumental to study the effect of complex control or policy design choices. However, existing formulations of mean field games face a computational bottleneck as they require solving two coupled partial differential equations (PDEs): a Hamilton-Jacobi-Bellman (HJB) equation characterizing the optimality of the agents' policy; and a Fokker-Planck-Kolmogorov (FPK) equation ensuring that the mean field is consistent with the state dynamics under the optimal policy. The resulting *initial-terminal-value problem* is highly non-standard, and computation tools are not yet fully mature, since the existence (and uniqueness) of solutions is guaranteed under abstract assumptions that are difficult to verify in practice [1], [2], [14], [15], [16], [17]. In practice, these assumptions typically manifest in linear dynamics with quadratic costs [8], [9], [10], [18], and/or a weak dependency on the mean field (e.g., through the average state or action only [8], [10], [19]).

On the other hand, the well-understood class of static *population games* provides a rich set of *Nash Equilibrium (NE)* analysis and computation tools [20]. In population games, a large population of agents optimize immediate payoffs over a finite set of actions, and the agents *do not* have individual state dynamics. The macroscopic evolution of actions is modelled by *evolutionary dynamics*, which include the well-known replicator, best response, Smith, and projection dynamics [20]. Evolutionary dynamics are useful for NE computation as their stationary points coincide with NE, and they are guaranteed to converge in the well-studied class of *stable population games* [20]. Moreover, recent works have enriched the classical evolutionary dynamics with passivity-based and higher-order learning techniques that extend the classes of games in which convergence is guaranteed [21], [22], [23], [24], [25].

Population games are trivially mean field games, in which agents have only one state (with trivial dynamics) and finite actions. The main contribution of this letter is to show that there is a larger and meaningful intersection: many practically relevant mean field games *are* population games.

All authors are with ETH Zurich, 8092 Zurich, Switzerland. {elokdae, ccenedese, kenzhang, acensi, bsaverio, efrazzoli}@ethz.ch

Research supported by NCCR Automation, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 180545).

Those are the class of *discrete-time, finite-state-and-action, stationary mean-field games*, which we call *Dynamic Population Games (DPGs)* for short. A natural solution concept in this class of games is the *Stationary Nash Equilibrium (SNE)* [4], [26], also referred to as the *stationary mean field equilibrium* [5], [27], [28], in which microscopically, each agent follows an optimal policy, and macroscopically, the state distribution is stationary. We show a mathematical reduction of SNE to standard NE in population games: For each DPG, there exists a suitably defined static population game whose NE coincide with SNE of the DPG.

The reduction of SNE in DPGs to NE in population games unlocks a rich set of tools for SNE analysis and computation. We leverage it to guarantee the existence of a SNE, develop an *evolutionary dynamics*-based SNE computation algorithm, and provide simple conditions that guarantee stability and uniqueness of the SNE. Finally, the versatility of our approach is demonstrated in two complex application examples involving fair resource allocation to heterogeneous agents, as well as epidemic propagation and control. Open source software for SNE computation is provided in: <https://gitlab.ethz.ch/elokdae/dynamic-population-games>.

The remainder of the letter is organized as follows: Section II introduces notation and definitions for DPGs; Section III presents our main result regarding the reduction of SNE to NE in population games; Section IV discusses important consequences of the reduction; and Section V presents numerical results from two application examples. The letter is concluded in Section VI.

### A. Notation

Let  $a, d \in D \subseteq \mathbb{N}$  and let  $c \in C \subseteq \mathbb{R}^n$ , then for a function  $f : D \times C \rightarrow \mathbb{R}$ , we distinguish discrete and continuous arguments through the notation  $f[d](c)$ . Alternatively, we write  $f : C \rightarrow \mathbb{R}^{|D|}$  as the vector-valued function  $f(c)$ , with  $f[d](c)$  denoting its  $d$ -th element. Similarly,  $p[a | d](c)$  denotes the conditional probability of  $a$  given  $d$  and  $c$ . Specifically,  $p[d^+ | d](c)$  denotes one-step transition probabilities. For  $m \in \mathbb{R}_+$  denoting a measure of total mass, we denote by  $\delta \in m \Delta(D) := \left\{ \sigma \in \mathbb{R}_+^{|D|} \mid \sum_{d \in D} \sigma[d] = m \right\}$  a distribution over the elements of  $D$ , with  $\delta[d] \in [0, m]$  denoting the mass of element  $d$  (if  $m = 1$  then  $\delta$  is a probability distribution). Finally, when considering heterogeneous agent types, we denote by  $x_\tau$  a quantity associated to type  $\tau$ .

## II. DYNAMIC POPULATION GAMES (DPGs)

In this section, we introduce the definition of a *Dynamic Population Game (DPG)*, along with its solution concept, the *Stationary Nash Equilibrium (SNE)*. This class of games has been previously referred to as discrete-time, finite state-and-action, stationary mean-field games [5], [27], [28].

### A. Dynamic Population Game (DPG) Setting

A DPG, denoted by  $\mathcal{G}$ , consists of a population  $\mathcal{N} = \{1, \dots, N\}$  of anonymous agents, with  $N$  large such that the agents approximately form a *continuum of mass*. Each agent has one of a finite number of private static *types*

$\tau \in \mathcal{T} = \{1, \dots, N_\tau\}$ , and the distribution of types is the exogenous parameter  $g \in \Delta(\mathcal{T})$ . In addition to types, each agent has a private *state*  $x \in \mathcal{X} = \{1, \dots, N_x\}$ , which follows a dynamic process to be described hereafter. The *state distribution* of type  $\tau$  agents is given by  $d_\tau \in \mathcal{D}_\tau = g_\tau \Delta(\mathcal{X})$ , and the joint *type-state distribution* is the concatenation of state distributions  $d = (d_1, \dots, d_{N_\tau}) \in \mathcal{D} = \prod_{\tau \in \mathcal{T}} \mathcal{D}_\tau$ .

At discrete time steps, agents participate in a game in which each agent of type-state  $[\tau, x]$  plays one of a finite number of actions  $a \in \mathcal{A}_\tau[x]$ . Type  $\tau$  agents pick their action according to the stationary homogeneous *policy*  $\pi_\tau : \mathcal{X} \rightarrow \Delta(\mathcal{A}_\tau[x])$ , which maps the state to a probability distribution over the actions. The probability to play action  $a$  when in state  $x$  is denoted by  $\pi_\tau[a | x] \in [0, 1]$ . The set of policies of type  $\tau$  is denoted by  $\Pi_\tau$ , and the concatenation of the policies is denoted by  $\pi = (\pi_1, \dots, \pi_{N_\tau}) \in \Pi = \prod_{\tau \in \mathcal{T}} \Pi_\tau$ .

The pair  $(d, \pi) \in \mathcal{D} \times \Pi$  is referred to as either the *mean field* (in the tradition of mean field games), or the *social state* (in the tradition of population games; a formal equivalency between mean fields and social states is shown in Section III), since  $(d, \pi)$  gives a macroscopic description of the distribution of the agents' types and states, as well as how they behave.

The discrete time state dynamics of type  $\tau$  agents are governed by the *state transition function*

$$p_\tau[x^+ | x, a](d, \pi), \quad (1)$$

which is the probability of the state to transition from  $x$  to  $x^+$  after playing  $a$ , as a function of the mean field  $(d, \pi)$ . The agent receives an *immediate payoff* for playing  $a$  when in state  $x$ , which is also a function of  $(d, \pi)$  given by

$$r_\tau[x, a](d, \pi). \quad (2)$$

A DPG  $\mathcal{G}$  is thus specified by the pair of state transition and immediate payoff functions (1)–(2).

### B. Solution Concept: Stationary Nash Equilibrium (SNE)

Before introducing the solution concept in DPGs, we must introduce definitions pertaining to the strategic problem faced by the population agents. Each agent faces a *discounted Markov Decision Process (MDP)*, i.e., it aims to maximize its discounted infinite horizon payoff, with the type-dependant discount factor  $\alpha_\tau \in [0, 1]$ . The individual MDP of each agent is coupled to other agents through the mean field  $(d, \pi)$ . Namely, when a type  $\tau$  agent follows policy  $\pi_\tau$ , its *expected immediate payoff* and *state transition matrix* are respectively given by,

$$R_\tau[x](d, \pi) = \sum_{a \in \mathcal{A}_\tau[x]} \pi_\tau[a | x] r_\tau[x, a](d, \pi), \quad (3)$$

$$P_\tau[x^+ | x](d, \pi) = \sum_{a \in \mathcal{A}_\tau[x]} \pi_\tau[a | x] p_\tau[x^+ | x, a](d, \pi), \quad (4)$$

and its *expected infinite horizon payoff*, also known as the *value function*, is recursively defined as

$$\begin{aligned} V_\tau[x](d, \pi) \\ = R_\tau[x](d, \pi) + \alpha_\tau \sum_{x^+ \in \mathcal{X}} P_\tau[x^+ | x](d, \pi) V_\tau[x^+](d, \pi). \end{aligned} \quad (5)$$

Equation (5) is the Bellman recursion characterizing the value function associated with policy  $\pi_\tau$ . Moreover, the *single stage deviation payoff*, also known as the *state-action value function* or *Q-function*, is given by

$$\begin{aligned} Q_\tau[x, a](d, \pi) \\ = r_\tau[x, a](d, \pi) + \alpha_\tau \sum_{x^+ \in \mathcal{X}} p_\tau[x^+ | x, a](d, \pi) V_\tau[x^+](d, \pi). \end{aligned} \quad (6)$$

Equation (6) is the expected infinite horizon payoff if the type  $\tau$  agent in state  $x$  deviates from  $\pi_\tau$  for a single stage by playing  $a$ , then follows  $\pi_\tau$  in the future.

It is well-known that a policy attains the maximum value function (5) if and only if it maximizes the state-action value function (6) for all states  $x \in \mathcal{X}$  [29, Proposition 1.3.4]. This motivates the definition of the *state-dependent best response correspondence*

$$B_\tau[x](d, \pi) \in \left\{ \sigma \in \Delta(\mathcal{A}_\tau[x]) \left| \begin{array}{l} \forall \sigma' \in \Delta(\mathcal{A}_\tau[x]), \\ \sum_{a \in \mathcal{A}_\tau[x]} (\sigma[a] - \sigma'[a]) Q_\tau[x, a](d, \pi) \geq 0 \end{array} \right. \right\}, \quad (\text{BR})$$

which is the set of probability distributions over the actions maximizing the state-action value function for a given type-state  $[\tau, x] \in \mathcal{T} \times \mathcal{X}$ . With this, we are ready to define the solution concept in DPGs.

**Definition 1 (Stationary Nash Equilibrium):** A *Stationary Nash Equilibrium (SNE)* is a mean field  $(d^*, \pi^*) \in \mathcal{D} \times \Pi$  which satisfies, for all  $[\tau, x] \in \mathcal{T} \times \mathcal{X}$ ,

$$d_\tau^*[x] = \sum_{x^- \in \mathcal{X}} d_\tau^*[x^-] P_\tau[x | x^-](d^*, \pi^*), \quad (\text{SNE.1})$$

$$\pi_\tau^*[\cdot | x] \in B_\tau[x](d^*, \pi^*). \quad (\text{SNE.2})$$

The SNE is similar to a classical Nash Equilibrium (NE) in that agents have no incentive to unilaterally deviate from the optimal policy  $\pi^*$  (SNE.2). However, in addition, it requires that the type-state distribution  $d^*$  is *stationary* under the stochastic process characterized by  $P_\tau(d^*, \pi^*)$  (SNE.1). The stationarity of  $d^*$  ensures that the MDP faced by each agent is time invariant at the equilibrium.

### III. MAIN RESULT: REDUCTION TO POPULATION GAME

In this section, we present the main result of this letter, which is a reduction of the SNE in DPGs to a standard NE in static population games. In order to state our result, we must first briefly recall the definition of population games [20]. A population game is defined by a *payoff vector*  $F(s) = (F_1(s), \dots, F_{N_\rho}(s))$ , where  $F_\rho(s)$  is the payoff vector of population  $\rho \in \mathcal{P} = \{1, \dots, N_\rho\}$ . Each entry of  $F_\rho(s)$ , denoted by  $F_\rho[a](s)$ , gives the payoff of a population

$\rho$  agent for playing one of its finite actions  $a \in \mathcal{A}_\rho$ . The *social state*  $s = (s_1, \dots, s_{N_\rho}) \in \mathcal{S}$  gives the distribution of actions in the population. Namely,  $s_\rho \in g_\rho \Delta(\mathcal{A}_\rho)$  is the distribution of actions in population  $\rho$ , where  $g_\rho \in \mathbb{R}_+$  is the total mass of agents in that population. In the language of DPGs, the populations  $\rho$  can be viewed as the types  $\tau$ , and there are no time varying states  $x$  (or, equivalently, the set of states  $\mathcal{X}$  is a singleton). The *best response correspondence* of population  $\rho$  is given by

$$B_\rho(s) \in \left\{ \sigma \in g_\rho \Delta(\mathcal{A}_\rho) \left| \begin{array}{l} \forall \sigma' \in g_\rho \Delta(\mathcal{A}_\rho), \\ \sum_{a \in \mathcal{A}_\rho} (\sigma[a] - \sigma'[a]) F_\rho[a](s) \geq 0 \end{array} \right. \right\}, \quad (\text{PG-BR})$$

and the *Nash Equilibrium (NE)* is defined as a social state  $s^*$  that satisfies  $s_\rho^* \in B_\rho(s^*)$  for all  $\rho \in \mathcal{P}$ .

**Theorem 1 (Reduction of SNE):** For every DPG  $\mathcal{G}$ , there exists a static population game  $F^\mathcal{G}$  whose Nash Equilibria (NE) coincide with the Stationary Nash Equilibria (SNE) of  $\mathcal{G}$ , i.e., it holds that  $(d^*, \pi^*)$  is a SNE of  $\mathcal{G}$  if and only if  $s^* = (d^*, \pi^*)$  is a NE of  $F^\mathcal{G}$ .

**Proof:** The proof proceeds by construction. We construct a population game  $F^\mathcal{G}$  with  $N_\rho = N_\tau (1 + N_x)$  populations. The first  $N_\tau$  populations correspond to the different types, and we denote these populations as  $\rho = \tau \in \mathcal{T}$ . The remaining  $N_\tau N_x$  populations correspond to the *type-state pairs*, and we denote these populations as  $\rho = [\tau, x] \in \mathcal{T} \times \mathcal{X}$ . The action space of populations  $\rho = \tau$  is the state space  $\mathcal{X}$ , i.e.,  $\mathcal{A}_\tau = \mathcal{X}$ , whereas the action space of populations  $\rho = [\tau, x]$  is respectively  $\mathcal{A}_{[\tau, x]} = \mathcal{A}_\tau[x]$ . The mass of populations  $\rho = \tau$  is respectively  $g_\tau$ , whereas the mass of populations  $\rho = [\tau, x]$  is uniformly  $g_{[\tau, x]} = 1$ . With this construction, it holds that  $\mathcal{S} = \mathcal{D} \times \Pi$ , i.e., each social state  $s \in \mathcal{S}$  under  $F^\mathcal{G}$  is formally equivalent to a mean field  $(d, \pi) \in \mathcal{D} \times \Pi$  under  $\mathcal{G}$ . We thus write  $s = (d, \pi)$  and define the payoffs  $F^\mathcal{G}(d, \pi)$  for both groups of populations as

$$F_\tau^\mathcal{G}[x](d, \pi) = \sum_{x^- \in \mathcal{X}} d_\tau[x^-] P_\tau[x | x^-](d, \pi) - d_\tau[x], \quad (7)$$

$$F_{[\tau, x]}^\mathcal{G}[a](d, \pi) = Q_\tau[x, a](d, \pi). \quad (8)$$

The payoffs of populations  $\rho = \tau$  in Equation (7) are motivated by condition (SNE.1). Intuitively, these populations can be interpreted as *agents of nature playing the dynamics*, i.e., they seek to bring the state distributions  $d_\tau$  to stationarity. On the other hand, the payoffs of populations  $\rho = [\tau, x]$  in Equation (8) coincide with the state-action value functions, as these are the quantities maximized by the agents in the DPG  $\mathcal{G}$ .

To prove the theorem, we now show that the NE condition of  $F^\mathcal{G}$ , i.e.,  $s_\rho^* \in B_\rho(s^*)$  for all  $\rho \in \mathcal{P}$ , is formally equivalent to conditions (SNE.1)–(SNE.2), with  $s^* = (d^*, \pi^*)$ . By construction, for all populations  $\rho = [\tau, x] \in \mathcal{T} \times \mathcal{X}$ , with  $s_{[\tau, x]}^* = \pi_\tau^*[\cdot | x]$ , condition  $\pi_\tau^*[\cdot | x] \in B_{[\tau, x]}(d^*, \pi^*)$  is formally equivalent to condition (SNE.2). Therefore, we must show that for the remaining populations  $\rho = \tau \in \mathcal{T}$ , with  $s_\tau^* = d_\tau^*$ ,  $d_\tau^* \in B_\tau(d^*, \pi^*)$

if and only if  $d_\tau^*[x]$  satisfies condition (SNE.1) for all  $x \in \mathcal{X}$ . One direction of the implication is trivial: if  $d_\tau^*[x]$  satisfies condition (SNE.1) for all  $[\tau, x] \in \mathcal{T} \times \mathcal{X}$ , then  $F_\tau^\mathcal{G}[x](d^*, \pi^*) = 0$ , and all state distributions satisfy  $d_\tau \in B_\tau(d^*, \pi^*)$ , including  $d_\tau^*$ . For the other direction of the implication, suppose, for the sake of contradiction, that  $d_\tau^* \in B_\tau(d^*, \pi^*)$  for all  $\tau \in \mathcal{T}$ , but there is some type-state  $[\tilde{\tau}, \tilde{x}] \in \mathcal{T} \times \mathcal{X}$  for which  $F_{\tilde{\tau}}^\mathcal{G}[\tilde{x}](d^*, \pi^*) \neq 0$ . Notice that  $\sum_{x \in \mathcal{X}} F_\tau^\mathcal{G}[x](d, \pi) = 0$  for all  $\tau \in \mathcal{T}$  and  $(d, \pi) \in \mathcal{D} \times \mathcal{X}$ ; therefore, if  $F_{\tilde{\tau}}^\mathcal{G}[\tilde{x}](d^*, \pi^*) > 0$  (respectively,  $< 0$ ), there must be another state  $\hat{x} \in \mathcal{X}$  for which  $F_{\tilde{\tau}}^\mathcal{G}[\hat{x}](d^*, \pi^*) < 0$  (respectively,  $> 0$ ). Suppose without loss of generality that  $F_{\tilde{\tau}}^\mathcal{G}[\tilde{x}](d^*, \pi^*) > 0$  and  $F_{\tilde{\tau}}^\mathcal{G}[\hat{x}](d^*, \pi^*) < 0$ . Since  $d_{\tilde{\tau}}^* \in B_{\tilde{\tau}}(d^*, \pi^*)$ , it must be that  $d_{\tilde{\tau}}^*[\hat{x}] = 0$ . But this contradicts  $F_{\tilde{\tau}}^\mathcal{G}[\hat{x}](d^*, \pi^*) < 0$ , in which  $-d_{\tilde{\tau}}^*[\hat{x}]$  is the only negative term. We thus showed that if  $d_\tau^* \in B_\tau(d^*, \pi^*)$  for all  $\tau \in \mathcal{T}$ , it must be that  $F_\tau^\mathcal{G}[x](d^*, \pi^*) = 0$  for all  $[\tau, x] \in \mathcal{T} \times \mathcal{X}$ , i.e.,  $d_\tau^*$  satisfies condition (SNE.1) for all  $[\tau, x] \in \mathcal{T} \times \mathcal{X}$ . ■

#### IV. CONSEQUENCES OF THE REDUCTION

In this section, we discuss important consequences of our main result, Theorem 1, which establishes that SNE in DPGs can be reduced to standard NE in static population games. In general, all existing results pertaining to the NE of the *equivalent population game*  $F^\mathcal{G}$ , constructed in the proof of Theorem 1, can be adopted to the SNE of DPG  $\mathcal{G}$ . We specialize a few of these results in this section.

##### A. Existence of SNE

A SNE is guaranteed to exist under the following mild assumption. This result can be proven directly using a fixed point argument [3], [4], but in this letter we present a simpler proof based on Theorem 1.

*Assumption 1 (Continuity):* The state transition function  $p_\tau(d, \pi)$  and the immediate payoff function  $r_\tau(d, \pi)$  are continuous in the mean field  $(d, \pi)$ .

*Proposition 1 (Existence of SNE):* Let Assumption 1 hold for DPG  $\mathcal{G}$ . Then  $\mathcal{G}$  is guaranteed to have at least one SNE  $(d^*, \pi^*)$ .

*Proof:* Consider the equivalent population game  $F^\mathcal{G}(d, \pi)$  whose payoffs are given in Equations (7)–(8). Assumption 1 implies that  $F^\mathcal{G}(d, \pi)$  is continuous in  $(d, \pi)$ : continuity of (7) is immediate; and continuity of (8) follows from noticing that the value function (5) can be equivalently written in vector form as  $V_\tau(d, \pi) = (I - \alpha_\tau P_\tau(d, \pi))^{-1} R_\tau(d, \pi)$ , which is continuous in  $(d, \pi)$  for all  $\alpha_\tau \in [0, 1]$ . It follows that at least one NE  $(d^*, \pi^*)$  is guaranteed to exist in  $F^\mathcal{G}$  [20, Theorem 2.1.1], which coincides with a SNE of  $\mathcal{G}$  by Theorem 1. ■

##### B. Evolutionary Dynamics for SNE Computation

Theorem 1 moreover inspires the following *evolutionary state-policy dynamics* model for SNE computation

$$\dot{d}_\tau[x] = \sum_{x^- \in \mathcal{X}} d_\tau[x^-] P_\tau[x | x^-](d, \pi) - d_\tau[x], \quad (\text{EV.1})$$

$$\dot{\pi}_\tau[\cdot | x] = \eta_\tau H_\tau(Q_\tau[x, \cdot](\pi, d), \pi_\tau[\cdot | x]), \quad (\text{EV.2})$$

where  $H_\tau : \mathbb{R}^{|\mathcal{A}_\tau[x]|} \times \Delta(\mathcal{A}_\tau[x]) \rightarrow \mathbb{R}^{|\mathcal{A}_\tau[x]|}$  is one of the familiar *mean dynamics* in population games, such as the replicator, best response, Smith, or projection dynamic [20, Chapters 4–6], and  $\eta_\tau$  is a rate parameter that controls the speed of policy dynamics (EV.2) with respect to state dynamics (EV.1). The dynamical system (EV.1)–(EV.2) is derived by applying evolutionary dynamics to the payoffs of the equivalent population game  $F^\mathcal{G}(d, \pi)$  given in Equations (7)–(8). In particular, the *projection dynamic* is applied to (7) to yield (EV.1)<sup>1</sup>. This choice is interpretable since it yields Laplacian diffusion dynamics in (EV.1), however other mean dynamics could be used for the purpose of SNE computation. On the other hand, (EV.2) corresponds to applying mean dynamic  $H_\tau$  to (8). For computation purposes, it is natural to choose a mean dynamic that satisfies the well-known condition of *Nash stationarity* [20, Section 5.2] for  $H_\tau$ , to guarantee that stationary points of  $H_\tau$  coincide with NE of the underlying population game. Mean dynamics satisfying Nash stationarity include the best response, Smith, and projection dynamic (but not the replicator dynamic) [20, Table 5.1]. The following proposition is then immediate from Theorem 1.

*Proposition 2:* Suppose that mean dynamics  $H_\tau$  satisfy Nash stationarity for all  $\tau \in \mathcal{T}$ . Then stationary points of the dynamical system (EV.1)–(EV.2) coincide with the SNE of the DPG  $\mathcal{G}$ .

A SNE can thus be computed by forward simulating dynamical system (EV.1)–(EV.2). If the system reaches a stationary point, a SNE is found. Notice that in comparison to existing mean field game approaches, the SNE computation is reduced to solving an ordinary differential equation (ODE) rather than an initial-terminal-value problem. Notice also that  $\eta_\tau$ , as well as the choice of the mean dynamic  $H_\tau$  itself, are tuning parameters that can assist with the convergence to a SNE. For example, the *Smith dynamic* is better behaved than the *best response dynamic* in many examples of games [20]. Theorem 1 is also useful to provide formal convergence guarantees for (EV.1)–(EV.2), as discussed next.

##### C. Stability and Uniqueness of SNE

As a final consequence of Theorem 1, we introduce the following criterion for the stability and uniqueness of a SNE.

*Definition 2 (Stable DPG):* DPG  $\mathcal{G}$  is a *stable DPG* if, for all  $s = (d, \pi)$ ,  $s' = (d', \pi') \in \mathcal{D} \times \Pi$ ,

$$\begin{aligned} & \sum_{\tau \in \mathcal{T}} \left[ (d_\tau - d'_\tau)^\top \left( (P_\tau(s)^\top - I) d_\tau - (P_\tau(s')^\top - I) d'_\tau \right) \right. \\ & \left. + \sum_{x \in \mathcal{X}} (\pi_\tau[\cdot | x] - \pi'_\tau[\cdot | x])^\top (Q_\tau[x, \cdot](s) - Q_\tau[x, \cdot](s')) \right] \leq 0. \end{aligned} \quad (9)$$

Moreover, if inequality (9) holds strictly whenever  $s' \neq s$ , then  $\mathcal{G}$  is a *strictly stable DPG*.

<sup>1</sup>In the projection dynamic, the social state is updated simply in the direction of the payoffs, which is projected onto the set of feasible directions keeping the social state in the simplex. Notice that the projection operation is redundant in this case.

Definition 2 is an immediate application of the definition of *stable population games* [20, Section 3.3] to the equivalent population game  $F^G$ . Thus, the following proposition follows immediately from Theorem 1 and [20, Theorem 3.3.10, 7.2.1, 7.2.4, 7.2.7, 7.2.9].

**Proposition 3:** Let Assumption 1 hold for DPG  $G$ . If  $G$  is a *stable DPG*, then its SNE are guaranteed to form a convex set. Moreover, the set of SNE is globally asymptotically stable under the best response and Smith dynamics. If instead  $G$  is a *strictly stable DPG*, then it is guaranteed to have a unique SNE, which is globally asymptotically stable under the replicator, best response, Smith and projection dynamics.

## V. APPLICATIONS

We present two application examples that demonstrate the versatility of our DPG formulation in tractably modelling complex real-world problems. The examples are adopted from recent application-specific works that have utilized the general results presented in this letter [7], [11], [13], [30], [31], [32].

### A. Fair Resource Allocation

Recent works have shown that introducing dynamics in a resource allocation problem enables the design of fair, non-monetary incentive schemes [11], [13], [31]. This is particularly important in contexts involving scarce, shared resources, e.g., for traffic congestion management. Instead of money, a non-tradable token, called *karma*, is issued to the agents, who may use their karma to bid for shared resources. Karma then flows from agents that acquire resources to those who yield resources, in accordance to exchange rules that are up to the policy maker to design. The agents must be strategic and future sighted in their bidding, in order to spare karma for times of highest *urgency*. Thus, karma naturally induces a DPG in which an agent's state is its current urgency and karma, and the immediate payoff for a given karma bid, as well as the karma dynamics, depend on the mean field through a potentially complex set of karma exchange rules.

Fortunately, the DPG formulation has enabled devising tractable computational models of such a *self-contained karma economy*. In [13], computational insight on the SNE of the karma DPG is provided for different karma exchange rules, including first vs. second price auctions, peer-to-peer payments vs. payments based on societal redistribution, and even exchanges involving karma taxation. Moreover, the fairness of the karma-based scheme with respect to heterogeneity in the agents' future awareness and urgency processes is investigated by incorporating multiple agent types. In [11], the self-contained karma economy is combined with a highway congestion model, where, in addition to bidding karma to access a congestion-free priority lane, agents must also choose when to travel in order to avoid the peak morning congestion, leading to even more complex and high-dimensional SNE. The model of [11] is further extended in [31], where additionally an elastic demand of commuters is introduced as a new agent type.




Fig. 1. SNE in karma DPG with heterogeneous discount factors.

In all of the above works, the SNE is computed using the evolutionary dynamics model (EV.1)–(EV.1), with the *perturbed best response dynamic* [20] as the mean dynamic  $H_\tau$ . Figure 1 shows two representative examples of karma SNE adopted from [13], which correspond to a case where there are two agent types with different discount factors  $\alpha_1 = 0.7$ ,  $\alpha_2 = 0.99$ . The top and bottom row of Figure 1 allow to assess the effect of a small progressive karma tax (collected from each agent and redistributed) on the SNE. The left column of Figure 1 shows the mean bid of the SNE policy  $\pi^*$  ( $y$ -axis) per type and karma state ( $x$ -axis); the center column shows the stationary karma distribution  $d^*$  per type; and the right column shows the stationary proportion of times each type gains access to the shared resource. With no tax, the future sighted agents bid slightly less (Figure 1 top-left) and end up holding more karma on average (Figure 1 top-center) than the short sighted agents; as a consequence, the future sighted agents are better off in terms of average access to the resource (Figure 1 top-right). The karma tax equalizes the field by redistributing karma to the short sighted agents. To the extent of our knowledge, such complex and high-dimensional SNE have not been computationally tractable before.

### B. Epidemic Modelling and Control

Recent works have furthermore emphasized the importance of epidemic modelling and control as a mean field game, in which agents properly account for the future risk of infection rather than behave myopically [6], [7], [30], [32]. Among these works, [7], [30], [32] utilize the DPG formulation to investigate the effect of complex policy interventions under minimal assumptions. The evolutionary dynamics model (EV.1)–(EV.1) is used not only for SNE computation, but also as a model of the coupled evolution of infection states and population behaviors, which is needed to predict important off-equilibrium quantities such as the peak infections. This computational approach, enabled by the results presented in this letter, allowed to study the effect of targeted lockdown restrictions in [30] and the interplay between cost and availability of vaccines and testing kits in [7]. Figure 2 shows a representative example adopted from [7], which demonstrates that decreasing the cost of




Fig. 2. Effect of vaccine cost and availability on epidemic spread.

vaccines without increasing its availability could lead to a surprising increase in peak infections (Figure 2 right vs. left). The top of Figure 2 shows the policy evolution (i.e., agents' behavior): decreasing the cost of vaccines causes more asymptomatic agents to attempt to vaccinate rather than test for infection. When they fail to do so due to limited vaccine availability, higher peak infections are observed (Figure 2 bottom-right).

The availability of a tractable formulation and efficient numerical tools for this class of games allowed to perform extensive comparative studies of multiple interventions despite the complexity of this socio-technical system.

## VI. CONCLUSION

We show that Stationary Nash Equilibria (SNE) in Dynamic Population Games (DPGs) can be reduced to standard Nash Equilibria (NE) in static population games. This result unlocks a rich set of existing tools for SNE analysis and computation, which makes the DPG formulation favorable in comparison to other mean-field game formulations. We demonstrate the versatility of the result in two complex application examples. An important future topic to investigate is the decentralized computation of the SNE by the agents.

## REFERENCES

- [1] M. Huang, R. P. Malhamé, P. E. Caines *et al.*, "Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle," *Communications in Information & Systems*, vol. 6, no. 3, pp. 221–252, 2006.
- [2] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese journal of mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
- [3] B. Jovanovic and R. W. Rosenthal, "Anonymous sequential games," *Journal of Mathematical Economics*, vol. 17, no. 1, pp. 77–87, 1988.
- [4] S. Adlakha, R. Johari, and G. Y. Weintraub, "Equilibria of dynamic games with many players: Existence, approximation, and market structure," *Journal of Economic Theory*, vol. 156, pp. 269–316, 2015.
- [5] X. Guo, A. Hu, R. Xu, and J. Zhang, "A general framework for learning mean-field games," *Mathematics of Operations Research*, vol. 48, no. 2, pp. 656–686, 2023.
- [6] V. Petrakova and O. Krivorotko, "Mean field game for modeling of covid-19 spread," *Journal of Mathematical Analysis and Applications*, vol. 514, no. 1, p. 126271, 2022.
- [7] A. R. Hota, U. Maitra, E. Elokda, and S. Bolognani, "Learning to mitigate epidemic risks: A dynamic population game approach," *Dynamic Games and Applications*, vol. 13, pp. 1106–1129, 2023.

- [8] F. Parise, M. Colombino, S. Grammatico, and J. Lygeros, "Mean field constrained charging policy for large populations of plug-in electric vehicles," in *53rd IEEE Conference on Decision and Control*, 2014.
- [9] D. Bauso, "Dynamic demand and mean-field games," *IEEE Transactions on Automatic Control*, vol. 62, no. 12, pp. 6310–6323, 2017.
- [10] M. A. Tajeddini and H. Kebriaei, "A mean-field game method for decentralized charging coordination of a large population of plug-in electric vehicles," *IEEE Systems Journal*, vol. 13, no. 1, pp. 854–863, 2018.
- [11] E. Elokda, C. Cenedese, K. Zhang, J. Lygeros, and F. Dörfler, "Carma: Fair and efficient bottleneck congestion management with karma," *arXiv preprint arXiv:2208.07113*, 2022.
- [12] S. R. Balseiro, O. Besbes, and G. Y. Weintraub, "Repeated auctions with budgets in ad exchanges: Approximations and design," *Management Science*, vol. 61, no. 4, pp. 864–884, 2015.
- [13] E. Elokda, S. Bolognani, A. Censi, F. Dörfler, and E. Frazzoli, "A self-contained karma economy for the dynamic allocation of common resources," *Dynamic Games and Applications*, pp. 1–33, 2023.
- [14] D. A. Gomes, J. Mohr, and R. R. Souza, "Discrete time, finite state space mean field games," *Journal de mathématiques pures et appliquées*, vol. 93, no. 3, pp. 308–328, 2010.
- [15] —, "Continuous time finite state mean field games," *Applied Mathematics & Optimization*, vol. 68, no. 1, pp. 99–143, 2013.
- [16] N. Almulla, R. Ferreira, and D. Gomes, "Two numerical approaches to stationary mean-field games," *Dynamic Games and Applications*, vol. 7, pp. 657–682, 2017.
- [17] N. Saldi, T. Basar, and M. Raginsky, "Markov-nash equilibria in mean-field games with discounted cost," *SIAM Journal on Control and Optimization*, vol. 56, no. 6, pp. 4256–4287, 2018.
- [18] A. Bensoussan, K. Sung, S. C. P. Yam, and S.-P. Yung, "Linear-quadratic mean field games," *Journal of Optimization Theory and Applications*, vol. 169, pp. 496–529, 2016.
- [19] M. Aziz and P. E. Caines, "A mean field game computational methodology for decentralized cellular network optimization," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 2, pp. 563–576, 2016.
- [20] W. H. Sandholm, *Population games and evolutionary dynamics*. MIT press, 2010.
- [21] M. J. Fox and J. S. Shamma, "Population games, stable games, and passivity," *Games*, vol. 4, no. 4, pp. 561–583, 2013.
- [22] M. Arcak and N. C. Martins, "Dissipativity tools for convergence to nash equilibria in population games," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 1, pp. 39–50, 2020.
- [23] B. Gao and L. Pavel, "On passivity, reinforcement learning, and higher order learning in multiagent finite games," *IEEE Transactions on Automatic Control*, vol. 66, no. 1, pp. 121–136, 2020.
- [24] S. Park and N. E. Leonard, "KI divergence regularized learning model for multi-agent decision making," in *American Control Conference (ACC)*, 2021.
- [25] S. Toonsi and J. Shamma, "Higher-order uncoupled dynamics do not lead to nash equilibrium-except when they do," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [26] B. A. Neumann, "Stationary equilibria of mean field games with finite state and action space," *Dynamic Games and Applications*, vol. 10, p. 845–871, 2020.
- [27] J. Subramanian and A. Mahajan, "Reinforcement learning in stationary mean-field games," in *18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019.
- [28] B. Yardim, S. Cayci, M. Geist, and N. He, "Policy mirror ascent for efficient and independent learning in mean field games," in *International Conference on Machine Learning*, 2023.
- [29] D. Bertsekas, *Dynamic programming and optimal control*. Belmont MA: Athena Scientific, 2007, vol. 2.
- [30] E. Elokda, S. Bolognani, and A. R. Hota, "A dynamic population model of strategic interaction and migration under epidemic risk," in *IEEE Conference on Decision and Control (CDC)*, 2021.
- [31] E. Elokda, C. Cenedese, K. Zhang, A. Censi, S. Bolognani, and E. Frazzoli, "A dynamic population game model of non-monetary bottleneck congestion management under elastic demand using karma," in *62nd IEEE Conference on Decision and Control (CDC)*, 2023.
- [32] U. Maitra, A. R. Hota, and V. Srivastava, "Sis epidemic propagation under strategic non-myopic protection: A dynamic population game approach," *IEEE Control Systems Letters*, 2023.