Tohoku University Policy Design Lab Discussion Paper

TUPD-2021-002

Evolutionary dynamics in heterogeneous populations: a general framework for an arbitrary type distribution

Dai ZUSAI
Policy Design Lab and Graduate School of Economics and Management
Tohoku University

May 2021

Evolutionary dynamics in heterogeneous populations: a general framework for an arbitrary type distribution

Dai ZUSAI*

May 11, 2021

Abstract

We present a general framework of evolutionary dynamics under persistent heterogeneity in payoff functions and revision protocols, allowing continuously many types in a game with finitely many strategies. Unlike the preceding literature, we do not assume anonymity of the game or aggregability of the dynamic. The dynamic is rigorously formulated as a differential equation of a joint probability measure of types and strategies. To establish a foundation of this framework, we clarify regularity assumptions on the revision protocol, the game and the type distribution to guarantee the existence of a unique solution trajectory as well as those to guarantee the existence of an equilibrium in a heterogeneous population game. We further verify equilibrium stationarity in general and stability in potential games under admissible dynamics. Our framework exhibits a wide range of possible applications, including equilibrium selection in Bayesian games and spatial evolution.
Keywords: evolutionary dynamics; heterogeneity; continuous space; potential games
JEL classification: C73, C62, C61.

1 Introduction

Evolutionary dynamics formulate off-equilibrium adjustment processes of agents' choices in games, allowing various decision rules (revision protocols) such as exact optimization, better reply based on pairwise comparison of payoffs, imitation, etc. Despite a wide range of applications to social and economic problems and also a potential role to challenge a conventional equilibrium-based approach, evolutionary dynamics have not fully captured one common staple of mathematical models of the economy/society: that is, heterogeneity of agents. It is a common practice in applied or empirical studies to assume continuous types of agents—especially, in many of applied economic models (e.g. auctions, aggregate demand[1]), in econometric estimation of discrete choice models (e.g. logit regression) and in theoretical

*Graduate School of Economics and Management and Policy Design Lab, Tohoku University. E-mail: ZusaiDPublic@gmail.com.

[1] Dynamic demand of myopic consumers is considered in the literature on dynamic monopoly pricing: Rohlfs (1974); Dhebar and Oren (1985, 1986) are seminal papers. They assume a continuous type distribution to define a continuous dynamic of the aggregate demand, though they implicitly assume aggregability. Employing the aggregability result in Ely and Sandholm (2005), Zusai (2015) justifies the aggregate demand dynamic as an aggregate obtained from the standard best response dynamic.

investigations of game experiments (quantal response equilibria). To embed heterogeneity to evolutionary dynamics, we typically assume that there are only finitely many types so they can be formulated as distinct populations (or genes); it requires some technical twists for discrete approximation of a continuous type space and also leaves non-negligible impacts of each individual type on others.

There are a few studies that deal with a continuous range of payoff heterogeneity in evolutionary dynamics. But, these studies focus on *anonymous games*—payoffs depend on others' choices only up to the aggregate strategy (Example 1)—and rely on *aggregability* of the dynamic—the change in the aggregate strategy is wholly determined from the current state of the aggregate distribution alone, independently of the underlying correlation between strategy choices and payoff types.[2] Aggregability may be assumed as in Blonski (1999) or may be derived from some specific form of the agents' strategy revision processes as in Ely and Sandholm (2005). Anyway, aggregability is a demanding restriction for games and dynamics; heterogeneous choices of agents cannot have an impact on payoffs or dynamics through something beyond their average, for example through the variance or distribution of strategies over different types. It is virtually the same as having just *one* "representative/average" type of agents and thus cannot capture impacts persistent heterogeneity among agents (or "fixed effects" in discrete choice regression) on evolution of their strategies.

In this paper, we provide a general framework to extend evolutionary dynamics to heterogeneous population games without requiring aggregability or restricting to a finite type space. We allow agents not only to have different payoff functions but also to follow different decision rules. Besides, our framework does not require anonymity and thus covers a wider range of games such as Bayesian games (Example 2) and spatial evolution (Example 3). To allow continuously many types in our framework, we face technical difficulty in dealing with continuous dimensions. The state of an evolutionary dynamic is the strategy distribution over different types; the dimension of the dynamic is just as large as the number of types. Without averaging off heterogeneity or assuming a finite type space, we need to deal with a dynamic system on infinite dimension. Therefore, we start from carefully defining evolutionary dynamics with a measure theoretic formulation of the state space, following the literature on evolution in games with continuously many *strategies*, especially Oechssler and Riedel (2001, 2002) and Cheung (2014).[3]

Even the unique existence of a solution trajectory cannot be simply granted for infinite dimensional dynamics. We clarify the regularity conditions on games and individual decision rules to assure it (Theorem 1): if individual agents respond to changes in payoffs in a Lipschitz-continuous way (L-continuous revision protocols in Definition 1) and the switching rates are uniformly bounded over all types (Assumption 2), the dynamic has a unique solution trajec-

[2]Hummel and McAfee (2018) apply (a generalized version of) replicator dynamics to formulate the demand dynamic in the monopoly pricing problem, as argued in footnote 1. While the replicator dynamic is not aggregable as argued in Zusai (2017), they obtain an explicit solution for the differential equation that represent the demand dynamic, thanks to their specification of functions (especially in their Lemma 1). Since the demand dynamic is only a part of the monopolist's dynamic optimization, equilibrium stationarity or stability is not discussed in their paper. (Actually, terms like 'equilibrium' or 'stability' do not appear in their paper, except the bibliography in their paper.)

[3]To name a few more, see also Hofbauer et al. (2009), Friedman and Ostrov (2013), Lahkar and Seymour (2013), Lahkar and Riedel (2015) and Cheung (2016).

tory from an arbitrary initial state in the heterogeneous setting. If an agent takes only the exact best response strategy (exact optimization protocols in Definition 2) just as in the best response dynamic, the individual revision protocol exhibits discontinuity when the transition of the strategy distribution triggers a switch of the agent's best response strategy through changes in payoffs. To mitigate discontinuity at the individual level and retain the unique existence of a solution trajectory, we additionally impose a kind of Lipschitz continuity on the distribution of the types whose best response strategies change with such a transition (Assumption 3).

We then confirm that standard properties of evolutionary dynamics can be extended from the homogeneous setting to the heterogeneous setting. First, if the individual decision rule assures stationarity of Nash equilibrium in the homogeneous setting, it also assures equilibrium stationarity in heterogeneous population games (Theorem 3). We also obtain the condition for the existence of an equilibrium (Theorem 4). Combining them, we can guarantee the existence of a stationary state in heterogeneous evolutionary dynamics. While stability of equilibrium is not granted generally even in a homogeneous population game, it is known that potential games assure equilibrium stability over a wide range of homogeneous evolutionary dynamics. With a rigorous formulation of heterogeneous potential games (Definition 5), we verify that equilibrium stability is extended to the heterogeneous setting (Theorem 5). In particular, a local maximum of the potential function is locally stable under *any* admissible dynamics. Provided that the equilibrium is isolated, the converse is true: once we find a locally stable equilibrium in a potential game under *some* particular admissible dynamic, it is a local maximum of the potential of the game and thus the local stability carries over *any* admissible dynamics (Corollary 3). Furthermore, we consider perturbation of a game by introducing payoff heterogeneity (Example 1), incomplete information (Example 2) and an uneven spatial structure of interactions (Example 3)). We confirm that the potential function of a base game can be naturally extended under such modifications (Theorem 6) and local stability in the base game is robust (Corollary 4).

In the next section, we define a heterogeneous population game and then build a heterogeneous evolutionary dynamic from an individual agent's revision protocol. Next we present our main results. In Section 3, we study the regularity conditions to guarantee the existence of a unique solution path. In Section 4, we extend equilibrium stationarity in general and equilibrium stability of potential games to the heterogeneous setting. Until this section, we consider heterogeneity only in payoff functions and focus on non-observational evolutionary dynamics, in which an agent's switching rate depends only on the payoff vector for the agent but not on other agents' strategies. In Section 5, we consider heterogeneity in revision protocols and observational dynamics such as imitative dynamics and excess payoff dynamics; we confirm that the theorems in this paper are robust to these extensions. We conclude the paper in Section 6 with a summary of the positive results in this paper and discussion on their implications and limitations. Appendices provide the proofs and a few technical details on the measure-theoretic construction of heterogeneous dynamics.

2 The base model

2.1 Heterogeneous population games

We first set up the game played in a heterogeneous population; here we quickly introduce essential components for our analysis, while we provide a complete illustration of the measure-theoretic formulation in Appendix A.1.

The society consists of a continuous population of agents, each of whom chooses a strategy from the same strategy set $\mathcal{S} = \{1, \cdots, S\}$. Each agent is assigned to type $\theta \in \Theta$, where type space Θ is closed in $\mathbb{R}$.[4] Types may represent heterogeneity in assessments of payoffs (possibly due to private information) as we focus in this base model, or heterogeneity in revision protocols as we discuss in Section 5, or both. If there are only finitely many types, these "types" could be formulated as different populations (or species in a biological context) in a conventional approach; however, we may have continuously many types in our model. Let $\mathcal{B}$ the set of Borel sets over Θ, and μ be the measure over the type space Θ: for any Borel set $B \in \mathcal{B}$ of types, $\mu(B)$ is the mass of agents whose types belong to B. We assume that the total mass of agents in the society is 1, i.e., $\mu(\Theta) = 1$; so, μ is a probability measure.

The social state is described by the **strategy distribution** $\mathbf{X} = (X_s)_{s \in \mathcal{S}}$, a joint distribution of strategies and types such that the marginal distribution of types coincides with μ. For each strategy $s \in \mathcal{S}$ and each Borel set $B \in \mathcal{B}$ of types, $X_s(B)$ is a mass of strategy-s players whose types belong to B. For each B, the strategy distribution $\mathbf{X} = (X_s)_{s \in \mathcal{S}}$ must satisfy $\sum_{s \in \mathcal{S}} X_s(B) = \mu(B)$. Denote by $\mathcal{X}$ the space of strategy distributions.

Since $\mathbf{X}$ satisfies $X_s(B) \leq \mu(B)$ for each $s \in \mathcal{S}$, each X_s is absolutely continuous with respect to μ; see (A.1) in Appendix A.1. We denote this relationship of the absolute continuity by $\mu \gg \mathbf{X}$. By Radon-Nikodym theorem, the absolute continuity guarantees the existence of a density function $x_s : \Theta \to \mathbb{R}_+$ of X_s such that $X_s(B) = \int_B x_s d\mu$. Then, **strategy density function** $\mathbf{x} = (x_s)_{s \in \mathcal{S}}$ is defined by collecting the density functions x_s over all $s \in \mathcal{S}$;[5] we abbreviate the relationship between $\mathbf{X}$ and its density $\mathbf{x}$ as $\mathbf{X} = \int \mathbf{x} d\mu$. Notice $\mathbf{x}(\theta) \in \Delta^{\mathcal{S}} := \{\mathbf{z} \in \mathbb{R}^S_+ : \sum_{s \in \mathcal{S}} z_s = 1\}$ for each type $\theta \in \Theta$.[6] The density $x_s(\theta) \in [0, 1]$ can be interpreted as the population share of strategy-s players in the subpopulation of type-θ agents. Denote by $\mathcal{F}_{\mathcal{X}}$ the set of strategy density functions.[7] Strategy density function $\mathbf{x}$ is (μ-almost) uniquely determined from strategy distribution $\mathbf{X}$ by Radon-Nikodym theorem, and vice versa. In this sense, we can regard $\mathcal{X}$ as equivalent to $\mathcal{F}_{\mathcal{X}}$.

Let $F_s[\mathbf{X}](\theta)$ be a type θ-agent's payoff from strategy s when the strategy distribution is $\mathbf{X}$. Thus, $\mathbf{F}[\mathbf{X}](\theta) = (F_s[\mathbf{X}](\theta))_{s \in \mathcal{S}} \in \mathbb{R}^S$ is the payoff vector for type θ given strategy distribution

---

[4] This is just to simplify exposition in the main body. All the theorems are applicable to any type space Θ as long as it is Polish (complete, separable, and metrizable).

[5] In an incomplete information game with a finite number of players, $\mathbf{X}$ is essentially a distributional strategy and $\mathbf{x}$ is a behavioral strategy in Milgrom and Weber (1985). Ely and Sandholm (2005) call $\mathbf{x}$ a Bayesian strategy.

[6] We denote $\mathbb{R}_+ = [0, +\infty)$ and $\mathbb{R}_{++} = (0, +\infty)$. Consider a $|\mathscr{U}|$-dimensional real space, each of whose coordinate is labeled with one element of $\mathscr{U} = \{1, \ldots, |\mathscr{U}|\}$. For set $S \subset \mathscr{U}$, we define an $|S|$-dimensional simplex $\Delta^{|\mathscr{U}|}(S)$ as $\Delta^{\mathscr{U}}(S) := \left\{\mathbf{x} \in \mathbb{R}_+^{|\mathscr{U}|} : \sum_{k \in S} x_k = 1 \text{ and } x_l = 0 \text{ for any } l \in \mathscr{U} \setminus S\right\}$. When S is the whole space $\mathscr{U}$ itself, we omit $|\mathscr{U}|$ and denote it by $\Delta^{\mathscr{U}}$.

[7] Two strategy density functions $\mathbf{x}, \mathbf{x}' \in \mathcal{F}_{\mathcal{X}}$ are considered as identical, i.e., $\mathbf{x} = \mathbf{x}'$, if $\mathbf{x}(\theta) = \mathbf{x}'(\theta)$ for μ-almost all $\theta \in \Theta$. They indeed yield the same strategy distribution.

$\mathbf{X}$. Given $\mathbf{X}$, $\mathbf{F}[\mathbf{X}] : \Theta \to \mathbb{R}^S$ specifies the payoff vector $\mathbf{F}[\mathbf{X}](\theta)$ for each type $\theta \in \Theta$; thus we call $\mathbf{F}[\mathbf{X}]$ the payoff vector profile. We assume that $\mathbf{F}[\mathbf{X}]$ belongs to $\mathcal{C}$, the set of continuous functions from Θ to $\mathbb{R}^S$. Payoff function $\mathbf{F} : \mathcal{X} \to \mathcal{C}$ maps a strategy distribution $\mathbf{X} \in \mathcal{X}$ to a payoff vector profile $\mathbf{F}[\mathbf{X}] \in \mathcal{C}$. A **heterogeneous population game** is defined by $(\mathcal{S}, (\Theta, \mathcal{B}, \mu), \mathbf{F})$, which we represent by $\mathbf{F}$.

Let $\mathcal{S}_{\mathrm{BR}}(\boldsymbol{\pi}^0) \subset \mathcal{S}$ be the set of best response strategies given payoff vector $\boldsymbol{\pi}^0 = (\pi^0_s)_{s\in\mathcal{S}} \in \mathbb{R}^S$: i.e., $\mathcal{S}_{\mathrm{BR}}(\boldsymbol{\pi}^0) := \operatorname{argmax}_{s\in\mathcal{S}} \pi^0_s$. Denote by $\Delta(\mathcal{S}_{\mathrm{BR}}(\boldsymbol{\pi}^0))$ the set of strategy distributions that assign positive probabilities only to the best response strategies given $\boldsymbol{\pi}^0$: i.e., $\Delta(\mathcal{S}_{\mathrm{BR}}(\boldsymbol{\pi}^0)) = \{\mathbf{y} \in \Delta^S : y_s > 0 \Rightarrow s \in \mathcal{S}_{\mathrm{BR}}(\boldsymbol{\pi}^0)\}$.

In heterogeneous population game $\mathbf{F}$, $\mathcal{S}^{\mathbf{F}}_{\mathrm{BR}}[\mathbf{X}](\theta) := \mathcal{S}_{\mathrm{BR}}(\mathbf{F}[\mathbf{X}](\theta))$ collects the best response strategies given payoff vector $\mathbf{F}[\mathbf{X}](\theta)$ for type θ; namely, it is the set of type-θ's best response strategies to $\mathbf{X}$ in game $\mathbf{F}$. Let $\Theta^{\mathbf{F}}_{s\in\mathrm{BR}}[\mathbf{X}]$ be the set of types for which strategy s is a best response to $\mathbf{X}$, and $\Theta^{\mathbf{F}}_{s=\mathrm{uniqBR}}[\mathbf{X}]$ the set of types for which strategy s is the *unique* best response to $\mathbf{X}$: i.e.,

$$\Theta^{\mathbf{F}}_{s\in\mathrm{BR}}[\mathbf{X}] := \{\theta \in \Theta : s \in \mathcal{S}^{\mathbf{F}}_{\mathrm{BR}}[\mathbf{X}](\theta)\} \quad \supset \quad \Theta^{\mathbf{F}}_{s=\mathrm{uniqBR}}[\mathbf{X}] := \left\{\theta \in \Theta \ : \ \{s\} = \mathcal{S}^{\mathbf{F}}_{\mathrm{BR}}[\mathbf{X}](\theta)\right\}.$$

In a Nash equilibrium, (almost) every agent correctly predicts strategy distribution $\mathbf{X}$ and takes the best response to it. Correspondingly, strategy distribution $\mathbf{X} \in \mathcal{X}$ with density $\mathbf{x} \in \mathcal{F}_{\mathcal{X}}$ is an **equilibrium strategy distribution** in game $\mathbf{F}$, if

$$\mathbf{x}(\theta) \in \Delta(\mathcal{S}^{\mathbf{F}}_{\mathrm{BR}}[\mathbf{X}](\theta)) \qquad \text{for } \mu\text{-almost all } \theta \in \Theta, \tag{1}$$

or equivalently,

$$x_s(\theta) = \begin{cases} 1 & \text{if } \theta \in \Theta^{\mathbf{F}}_{s=\mathrm{uniqBR}}[\mathbf{X}] \\ 0 & \text{if } \theta \notin \Theta^{\mathbf{F}}_{s\in\mathrm{BR}}[\mathbf{X}] \end{cases} \qquad \text{for all } s \in \mathcal{S} \text{ and } \mu\text{-almost all } \theta \in \Theta. \tag{1'}$$

That is, if s is the unique best response for type θ, (almost) all the agents of this type should take it; if s is not a best response, (almost) none of these agents should take it. We leave indeterminacy of $x_s(\theta)$ in an equilibrium when there are multiple best response strategies for θ and s is just one of them. Note that (1) is equivalent to

$$\mu(\Theta^{\mathbf{F}}_{s=\mathrm{uniqBR}}[\mathbf{X}] \cap B) \le X_s(B) \le \mu(\Theta^{\mathbf{F}}_{s\in\mathrm{BR}}[\mathbf{X}] \cap B) \qquad \text{for all } s \in \mathcal{S} \text{ and } B \in \mathcal{B}. \tag{2}$$

Among types in B, all those who have s as the *unique* best response *must* choose this strategy s in equilibrium; thus $X_s(B)$ must be at least $\mu(\Theta^{\mathbf{F}}_{s=\mathrm{uniqBR}}[\mathbf{X}] \cap B)$. On the other hand, those who have s as *one* of the best responses *may or may not* add to strategy-s players and thus $X_s(B)$ is at most $\mu(\Theta^{\mathbf{F}}_{s\in\mathrm{BR}}[\mathbf{X}] \cap B)$.

Examples of heterogeneous population games

Example 1 (Anonymous game). Denote by $\bar{x}_s := X_s(\Theta) = \int x_s d\mu \in [0,1]$ the mass of agents who take strategy $s \in \mathcal{S}$ in the entire population over all types in Θ. We call $\bar{\mathbf{x}} := (\bar{x}_s)_{s\in\mathcal{S}} \in \Delta^S$ the **aggregate strategy**. If each type's payoff function $\mathbf{F}(\theta) : \mathcal{X} \to \mathbb{R}^S$ depends only on aggregate strategy, that is, $\mathbf{F}$ satisfies $\mathbf{F}[\mathbf{X}](\theta) = \mathbf{F}[\mathbf{X}'](\theta)$ for any type $\theta \in \Theta$ under any pair of two strategy distributions $\mathbf{X}, \mathbf{X}' \in \mathcal{X}$ that yields the same aggregate strategy $\mathbf{X}(\Theta) = \mathbf{X}(\Theta')$, then we call the game an **anonymous game**.[8]

Especially, in the context of discrete choice models such as in Anderson et al. (1992), it is common to introduce payoff heterogeneity in an additively separable manner. That is, the payoff function is additively separated to the common part and the idiosyncratic part: with type space $\Theta \subset \mathbb{R}^S$, type $\boldsymbol{\theta} = (\theta_s)_{s\in\mathcal{S}} \in \mathbb{R}^S$ is defined as the idiosyncratic payoff vector for this type, which varies among agents but does not change over time regardless of the state of the population. Given aggregate strategy $\bar{\mathbf{x}}$, $\mathbf{F}^0(\bar{\mathbf{x}}) = (F_s^0(\bar{\mathbf{x}}))_{s\in\mathcal{S}} \in \mathbb{R}^S$ is the common payoff vector, shared by all the agents in the entire population. Thus, at each strategy distribution $\mathbf{X} \in \mathcal{X}$, the payoff vector for a type-$\boldsymbol{\theta}$ agent is

$$\mathbf{F}[\mathbf{X}](\boldsymbol{\theta}) = \mathbf{F}^0(\mathbf{X}(\Theta)) + \boldsymbol{\theta}. \tag{3}$$

We call an anonymous game with such additively separable idiosyncratic payoffs an **additively separable anonymous game (ASAG)**. We can regard an ASAG as an extension of a homogeneous population game $\mathbf{F}^0$ to a heterogeneous setting. ∎

Example 2 (Bayesian game). A Bayesian game can be fit into our framework. Let Σ be the set of possible states and $\mathbb{P}_\Sigma$ be the prior belief over Σ with $\mathcal{B}_\Sigma$ the set of measurable sets. State $\sigma \in \Sigma$ determines the distribution $\mathbb{P}_{\Theta|\sigma}$ of types (signals) and payoff function $\mathbf{F}^\sigma = (F_s^\sigma)_{s\in\mathcal{S}} : \mathcal{X} \to \mathbb{R}^S$, while the strategy (action) set $\mathcal{S}$ is common over all states.

Receiving signal $\theta \in \Theta$, an agent forms the posterior belief $\mathbb{P}_{\Sigma|\theta}$ such as

$$\mathbb{P}_\Sigma(B_\Sigma|\theta) = \frac{\int_{\sigma\in B_\Sigma} \mathbb{P}_\Theta(d\theta|\sigma)\mathbb{P}_\Sigma(d\sigma)}{\int_{\sigma\in\Sigma} \mathbb{P}_\Theta(d\theta|\sigma)\mathbb{P}_\Sigma(d\sigma)} \qquad \text{for each } B_\Sigma \in \mathcal{B}_\Sigma.$$

Based on this, the type-θ agent assesses the expected payoff from action $s \in \mathcal{S}$ given strategy distribution $\mathbf{X} \in \mathcal{X}$ as

$$F_s[\mathbf{X}](\theta) = \int_{\sigma\in\Sigma} F_s^\sigma(\mathbf{X})\mathbb{P}_\Sigma(d\sigma|\theta).$$

Note that $\mathbf{x}(\theta)$ indicates agents' choices of strategies s conditional on receiving signal θ; thus $\mathbf{x}$ corresponds to a Bayesian strategy. The prior distribution of signals $\mathbb{P}_\Theta$ such as $\mathbb{P}_\Theta(B) = \int_B \mathbb{P}_{\Theta|\sigma}(d\theta|\sigma)\mathbb{P}_\Sigma(d\sigma)$ is regarded as the type distribution μ. ∎

Example 3 (Structured population game). We could interpret a type just as a "population" in a conventional model in evolutionary game theory, while we allow continuously many populations. Then, a type represents an affiliation to a certain subgroup of agents in the society; so, Θ is a set of subgroups. Let a base game be a two-population game $\mathbf{F}^0 : \Delta^{\mathcal{S}} \times \Delta^{\mathcal{S}} \to \mathbb{R}^S$; an agent chooses a strategy, say s, from $\mathcal{S}$ and then receives payoff $F_s^0(\mathbf{x}, \mathbf{x}')$ given the strategy distribution (density) in the agent's own population $\mathbf{x} \in \Delta^{\mathcal{S}}$ and that in the opponent's population $\mathbf{x}' \in \Delta^{\mathcal{S}}$. When the society is divided into many subgroups, their connections may not be uniform. Say, an agent in subgroup θ assigns weight $g(\theta, \theta') \in \mathbb{R}$ to the game with subgroup θ'. For example, the society may be geographically split to subgroups by locations; then, θ represents a location and $g(\theta, \theta')$ may be determined from the distance or commuting

[8] Notice the difference from an *aggregative* game (Corchón, 1994; Jensen, 2018). The payoff depends only on the population-weighted sum of strategies $\sum_{s\in\mathcal{S}} s\bar{x}_s$ in a linearly aggregative game, or a scalar-valued summary $g(\bar{\mathbf{x}}) \in \mathbb{R}$ in a generalized aggregative game; to make sense, strategies must be some quantities. Cournot competition where strategy s is the quantity of production is a canonical example of an aggregative game. An aggregative game is a special case of anonymous games, since the latter does not require the aggregate strategy $\bar{\mathbf{x}} \in \mathbb{R}^S$ to reduce to a scalar.

cost between two locations θ and θ' (e.g. Hwang et al. (2013)). Or, a subgroup may be defined by racial or social identity, which may exhibit continuous gradation. Then, $g(\theta, \theta')$ represents the frequency of interactions between the identity groups θ and θ' or the subjective weight for θ on interactions with θ'.[9]

Assuming that an agent must apply the same strategy to any opponent subgroups, the total payoff for an agent in subgroup θ from strategy s given the strategy distribution $\mathbf{X}$ is

$$F_s[\mathbf{X}](\theta) := \int_\Theta F_s^0(\mathbf{x}(\theta), \mathbf{x}(\theta'))g(\theta, \theta')\mu(d\theta').$$

This defines a population game $\mathbf{F}$, which Wu and Zusai (2019) call a *structured population game*[10]

$\blacksquare$

2.2 Evolutionary dynamics

In an evolutionary dynamic, an agent occasionally changes the strategy over a continuous time horizon $\mathbb{R}_+$, following a Poisson process. The timing of a switch and the choice of which strategy to switch to are determined by **revision protocol** $\boldsymbol{\rho} = (\rho_{ss'})_{s,s' \in \mathcal{S}} : \mathbb{R}^S \to \mathbb{R}_+^{S \times S}$, a collection of switching rate functions $\rho_{ss'} : \mathbb{R}^S \to \mathbb{R}_+$ over all the pairs $(s, s') \in \mathcal{S} \times \mathcal{S}$ of two strategies. An economic agent should base the switching decision on the payoff vector that the agent is facing. Let $\boldsymbol{\pi}^0 \in \mathbb{R}^S$ be the payoff vector for the agent. The switching rate $\rho_{ss'}(\boldsymbol{\pi}^0) \in \mathbb{R}_+$ is a Poisson arrival rate at which this agent switches to strategy $s' \in \mathcal{S}$ conditional on that the agent has been taking strategy $s \in \mathcal{S}$ so far and currently faces payoff vector $\boldsymbol{\pi}^0$. The analysis in this paper is applicable to *observational dynamics*, in which the switching rates also depend on the strategy distribution; e.g. the replicator dynamic and excess payoff dynamics. In addition, all our theorems hold even when different types of agents follow different revision protocols. We confirm applicability to these extensions in Section 5, while we focus on heterogeneity only in payoff functions and thus assume that all the types of agents share the same revision protocol $\boldsymbol{\rho}$ until that section.

In the heterogeneous setting, different types of agents may face different payoff vectors. Let $\boldsymbol{\pi} : \Theta \to \mathbb{R}^S$ be a payoff vector profile that specifies payoff vector $\boldsymbol{\pi}(\theta)$ of each type θ. From revision protocol $\boldsymbol{\rho} : \mathbb{R}^S \to \mathbb{R}_+^{S \times S}$, we construct an evolutionary dynamic of strategy density function $\mathbf{x} \in \mathcal{F}_\mathcal{X}$ with function $\mathbf{v} = (v_s)_{s \in \mathcal{S}} : \mathbb{R}^S \times \Delta^S \to \mathbb{R}^S$ as

$$\dot{x}_s(\theta) = v_s(\boldsymbol{\pi}(\theta), \mathbf{x}(\theta)) := \sum_{s' \in \mathcal{S}} x_{s'}(\theta)\rho_{s's}(\boldsymbol{\pi}(\theta)) - x_s(\theta) \sum_{s' \in \mathcal{S}} \rho_{ss'}(\boldsymbol{\pi}(\theta)) \tag{4}$$

for each type $\theta \in \Theta$ and each strategy $s \in \mathcal{S}$, i.e., $\dot{\mathbf{x}}(\theta) = \mathbf{v}(\boldsymbol{\pi}(\theta), \mathbf{x}(\theta))$. In an infinitesimal length of time $dt \in \mathbb{R}$, $\sum_{s' \in \mathcal{S}} x_{s'}(\theta)\rho_{s's}(\boldsymbol{\pi}(\theta))dt$ is approximately the mass of type-θ agents who switch to strategy s from other strategies $s' \in \mathcal{S}$, namely, the gross inflow to $x_s(\theta)$; similarly, $x_s(\theta) \sum_{s' \in \mathcal{S}} \rho_{ss'}(\boldsymbol{\pi}(\theta))dt$ is the gross outflow from $x_s(\theta)$. Thus, $v_s(\boldsymbol{\pi}(\theta), \mathbf{x}(\theta))dt$ is the net flow

[9]The weight $g(\theta, \theta')$ can be negative, which implies that an agent has a reversed preference in interactions with subgroup θ'. For example, if a base game is a coordination game, an agent may want to coordinate to the same action with a 'friend'; but, with agents in an 'enemy' subgroup, the agent wants to take a different action. See Example 1 in Wu and Zusai (2019).

[10]They restrict attention to finitely many subgroups of agents who play a linear game (with no influence of the own population) such as $\mathbf{F}^0(\mathbf{x}, \mathbf{x}') = \mathbf{U}^0\mathbf{x}'$ with an $S \times S$ matrix $\mathbf{U}^0$, while they consider both the "medium run" dynamic where an agent's affiliation is fixed exogenously (as in our model) and the "long run" dynamic where an agent can change both strategy and affiliation (not covered in this present paper).

to $x_s(\theta)$ in this period of time dt.

Embedding a heterogeneous population game $\mathbf{F} : \mathbf{x} \mapsto \boldsymbol{\pi}$ into the evolutionary dynamic $\mathbf{v} : (\boldsymbol{\pi}, \mathbf{x}) \mapsto \dot{\mathbf{x}}$, we obtain an autonomous dynamic $\mathbf{v}^{\mathbf{F}} : \mathbf{x} \mapsto \dot{\mathbf{x}}$ of strategy density function $\mathbf{x} \in \mathcal{F}_{\mathcal{X}}$ by

$$\dot{\mathbf{x}}(\theta) = \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) := \mathbf{v}(\mathbf{F}[\mathbf{X}](\theta), \mathbf{x}(\theta)) \in \mathbb{R}^S \quad \text{for each type } \theta \in \Theta, \text{ where } \mathbf{X} = \int \mathbf{x} d\mu.$$

By collecting $\mathbf{v}^{\mathbf{F}}[\mathbf{x}]$ over types, we can further define the **heterogeneous dynamic** $\mathbf{V}^{\mathbf{F}}$ of strategy distribution in $\mathcal{X}$ as

$$\dot{\mathbf{X}}(B) = \mathbf{V}^{\mathbf{F}}[\mathbf{X}](B) := \int_B \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta)\mu(d\theta) \qquad \text{for each } B \in \mathcal{B}. \tag{5}$$

When we distinguish $\mathbf{v}^{\mathbf{F}}$ (or $\mathbf{V}^{\mathbf{F}}$) from $\mathbf{v}$, we call the former a **combined dynamic**, i.e., a dynamic obtained from combination of $\mathbf{v}$ and $\mathbf{F}$. (See Footnote 22.) Note that (4) implies the forward invariance of $\mathcal{X}$, i.e., $\mathbf{X}_0 \in \mathcal{X} \;\Rightarrow\; \mathbf{X}_t \in \mathcal{X}\ \forall t > 0$ in any solution trajectory of $\mathbf{V}^{\mathbf{F}}$ with any $\mathbf{F}$, since $\sum_{s \in \mathcal{S}} v_s(\boldsymbol{\pi}(\theta), \mathbf{x}(\theta)) = 0$ and $[x_s(\theta) = 0 \;\Rightarrow\; v_s(\boldsymbol{\pi}(\theta), \mathbf{x}(\theta)) \geq 0]$ for any $\boldsymbol{\pi}, \mathbf{x}$. Note that an agent's type θ is *persistently fixed* over time: each agent draws its type θ from Θ at time 0 and keeps it forever.

Examples of evolutionary dynamics

To make a concrete image of revision protocols, here we review major evolutionary dynamics.[11] In particular, we separate the dynamics based on optimization from others because they need different regularity conditions to guarantee the existence of a unique solution trajectory.

L-continuous revision protocols. Under an L-continuous revision protocol $\boldsymbol{\rho}$, the switching rate function $\rho_{ss'}$ is a Lipschitz continuous function of the payoff vector.

Definition 1 (L-continuous revision protocols)**.** In an **L-continuous revision protocol $\boldsymbol{\rho}$**, the switching rate function $\rho_{ss'} : \mathbb{R}^S \to \mathbb{R}_+$ of each pair of strategies $s, s' \in \mathcal{S}$ is Lipschitz continuous:[12] there exists $L_\rho > 0$ such that

$$|\rho_{ss'}(\boldsymbol{\pi}) - \rho_{ss'}(\boldsymbol{\pi}')| \leq L_\rho |\boldsymbol{\pi} - \boldsymbol{\pi}'| \qquad \text{for any } s, s' \in \mathcal{S}, \boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathbb{R}^S.$$

Example 4. In a class of **pairwise comparison dynamics**, the switching rate $\rho_{ss'}(\boldsymbol{\pi})$ increases with the payoff difference $\pi_{s'} - \pi_s$. In particular, the revision protocol $\rho_{ss'}(\boldsymbol{\pi}) = [\pi_{s'} - \pi_s]_+$ defines the **Smith dynamic** (Smith, 1984).[13] ∎

Example 5. Because of continuity of a switching rate function, we see **smooth best response dynamics** (Fudenberg and Kreps, 1993) as constructed from continuous revision protocols. For example, the **logit dynamic** (Fudenberg and Levine, 1998) is constructed from $\rho_{ss'}(\boldsymbol{\pi}) = \exp(\mu^{-1}\pi_{s'}) \;/\; \sum_{s'' \in \mathcal{S}} \exp(\mu^{-1}\pi_{s''})$ with noise level $\mu > 0$.

This revision protocol can be obtained from perturbed optimization: upon the receipt of each revision opportunity, an agent draws each random perturbation in each strategy s's pay-

[11] Readers who are familiar with major evolutionary dynamics may just scan this subsection quickly and jump to Definitions 1 and 2.

[12] We adopt the L^1- norm as a norm on a finite-dimensional real space, which we denote by $|\cdot|$: for vector $\mathbf{v} = (v_i)_{i=1}^I \in \mathbb{R}^I$, $|\mathbf{v}| := \sum_{i=1}^I |v_i|$.

[13] $[\cdot]_+$ is an operator to truncate the negative part of a number: i.e., $[\check{\pi}]_+$ is $\check{\pi}$ if $\check{\pi} \geq 0$ and 0 otherwise.

off ε_s from a double exponential distribution[14] and then switches to the strategy that maximizes $\pi_s + \varepsilon_s$ among all strategies $s \in \mathcal{S}$. In general, a smooth best response dynamic can be constructed from such perturbed optimization under some admissibility condition: see Hofbauer and Sandholm (2002); Hofbauer et al. (2007). Note that, upon the receipt of a revision opportunity and a draw of $\boldsymbol{\varepsilon} \in \mathbb{R}^S$, an agent *always* switches to the best response strategy, however small the payoff gain by this switch is.

Note that payoff perturbation $\boldsymbol{\varepsilon} = (\varepsilon_s)_{s\in\mathcal{S}}$ is *transient* : a different value of $\boldsymbol{\varepsilon}$ will be drawn at each revision opportunity from an i.i.d. distribution. So, there is no (ex ante) heterogeneity in $\boldsymbol{\varepsilon}$. In contrast, the idiosyncratic payoff type θ in our heterogeneous dynamics is *persistent*. ∎

Exact optimization protocols. In an exact optimization protocol, an agent switches *only* to the best response given the current payoff vector: if strategy s' does not yield the maximal payoff among $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_S)$, then $\rho_{ss'}(\boldsymbol{\pi}) = 0$ regardless of the agent's current strategy s. We allow the switching rate to a best response strategy to vary with $\boldsymbol{\pi}$ and $s, s' \in \mathcal{S}$. Denote by $Q_{ss'}(\boldsymbol{\pi})$ the *conditional* switching rate from s to s', *provided* that s' is already designated as the new strategy. In the definition below, we extend the domain of $Q_{ss'}$ to $\mathbb{R}^S$ while assuming its continuity over the whole domain. The actual switching rate $\rho_{ss'}$ is defined as the truncation of $Q_{ss'}$ when s' is not a best response; the truncation causes discontinuity.

Definition 2 (Exact optimization protocols)**.** In an **exact optimization protocol**, the switching rate function $\rho_{ss'} : \mathbb{R}^S \to \mathbb{R}_+$ of each pair of strategies $s, s' \in \mathcal{S}$ is expressed as[15]

$$\rho_{ss'}(\boldsymbol{\pi}) = \begin{cases} 0 & \text{if } s' \notin \operatorname{argmax}_{s''\in\mathcal{S}} \pi_{s''}, \\ Q_{ss'}(\boldsymbol{\pi}) & \text{if } \{s'\} = \operatorname{argmax}_{s''\in\mathcal{S}} \pi_{s''}, \end{cases}$$

with a Lipschitz continuous function $Q_{ss'} : \mathbb{R}^S \to \mathbb{R}_+$.

Example 6. In the **standard best response dynamic (BRD)** as defined by Hofbauer (1995b); Gilboa and Matsui (1991), a revising agent always switches to the best response strategy that maximizes the current payoff with probability 1, however small the payoff gain by this optimization is. That is, the standard BRD is constructed from an exact optimization dynamic with $Q_{ss'} \equiv 1$. The heterogeneous version is considered in Ely and Sandholm (2005); they prove that the *aggregate* strategy distribution in the heterogeneous standard BRD follows a homogenized smooth BRD, i.e., the BRD of a single representative population of homogeneous agents whose payoff types θ is transient. ∎

Example 7. Consider a version of BRD in which the switching rate to the unique best response $Q_{ss'}$ depends on the payoff difference (the **payoff deficit**) between the current strategy s and the best response s', i.e., $Q_{ss'}(\boldsymbol{\pi}) = Q(\pi_{s'} - \pi_s)$ whenever $s' \in \operatorname{argmax}_{s''\in\mathcal{S}} \pi_{s''}$. Function $Q : \mathbb{R}_+ \to [0, 1]$ is called a *tempering function* and assumed to be continuously differentiable and satisfy $Q(0) = 0$ and $Q(q) > 0$ whenever $q > 0$. Then this revision protocol yields

[14] Given the noise level μ, the cumulative distribution function of the double exponential distribution is $\mathbb{P}(\varepsilon_s \leq c) = \exp(-\exp(-\mu^{-1}c - \gamma))$ where $\gamma \approx 0.5772$ is Euler's constant.

[15] When $\operatorname{argmax}_{s''\in\mathcal{S}} \pi_{s''}$ is not a singleton, this definition does not specify $\rho_{ss'}(\boldsymbol{\pi})$ for a best response $s' \in \operatorname{argmax}_{s''\in\mathcal{S}} \pi_{s''}$. However, we later imposes Assumption 3 and this implies the uniqueness of the best response for almost all types.

the **tempered BRD**; Zusai (2018) constructs this revision protocol from optimization with a stochastic switching cost whose cumulative distribution function is Q. ∎

3 Existence of a unique solution trajectory

We verify Lipschitz continuity of a heterogeneous dynamic to guarantee the existence of a unique solution trajectory from an arbitrary initial strategy distribution by using a version of Picard-Lindelöf theorem (Theorem 2). To apply this theorem, we prove the Lipschitz continuity of heterogeneous dynamic $\mathbf{V}^{\mathbf{F}}$ with respect to the variational norm $\|\cdot\|_\infty$ (see Appendix A.2), as in (6) in Appendix B.

For this, we assume that payoff function $\mathbf{F}$ is Lipschitz continuous and switching rate function ρ is bounded.

Assumption 1 (Lipschitz continuity of the payoff function)**.** For μ-almost every type $\theta \in \Theta$, $\mathbf{F}(\theta) : \mathcal{X} \to \mathbb{R}^S$ is Lipschitz continuous with Lipschitz constant $L_{\mathbf{F}}(\theta)$:

$$|\mathbf{F}[\mathbf{X}](\theta) - \mathbf{F}[\mathbf{X}'](\theta)| \leq L_{\mathbf{F}}(\theta)\|\mathbf{X} - \mathbf{X}'\|_\infty \qquad \text{for any } \mathbf{X}, \mathbf{X}' \in \mathcal{X}.$$

In addition, $\bar{L}_{\mathbf{F}} := \int_\Theta L_{\mathbf{F}}(\theta)\mu(d\theta) < \infty$.[16]

Assumption 2 (Uniformly bounded switching rates)**.** There exists $\bar{\rho} \in \mathbb{R}_+$ such that[17]

$$\rho_{ss'}(\mathbf{F}[\mathbf{X}](\theta)) \leq \bar{\rho} \quad \text{ for any } s, s' \in \mathcal{S} \text{ and } \mu\text{-almost all } \theta \in \Theta, \text{ and any } \mathbf{X} \in \mathcal{X}.$$

Note that the keys in these assumptions specifically for a heterogeneous setting are boundedness of the Lipschitz constant $L(\theta)$ and uniformness of $\bar{\rho}$ over all types. If there are only finitely many types, these assumptions are trivially satisfied as long as each type θ's payoff function $\mathbf{F}(\theta)$ is Lipschitz continuous.[18]

One might attempt to merge Assumption 2 into Assumption 1 by strengthening the latter to impose a common Lipschitz constant $\hat{L}_{\mathbf{F}}$ over all types θ. But this does not imply Assumption 2. For example, consider a binary ASAG with $\mathcal{S} = \{I, O\}$ such that $F_I[\mathbf{X}](\theta) = F_I^0(\mathbf{X}(\Theta))$ and $F_O[\mathbf{X}](\theta) = \theta$ for each $\mathbf{X} \in \mathcal{X}$. Assumption 1 holds as long as F_I^0 is Lipschitz continuous. But, F_O is not bounded unless Θ is bounded; hence $\rho_{ss'}(\mathbf{F}[\mathbf{X}](\theta))$ may not be bounded, for example, if $\rho_{ss'}(\boldsymbol{\pi})$ grows unboundedly with the payoff difference $|\boldsymbol{\pi}_I - \boldsymbol{\pi}_O|$; e.g. $\rho_{ss'}(\boldsymbol{\pi}) = [\pi_{s'} - \pi_s]_+$.

These two assumptions are sufficient for the Lipschitz continuity of $\mathbf{V}^{\mathbf{F}}$ if all types follow an L-continuous revision protocol thanks to continuity of the protocol itself. However, an exact optimization protocol essentially involves discontinuity. One may recall that, in a homogeneous setting (i.e., a single type), the standard BRD is a discontinuous dynamic (and thus formulated as a differential inclusion, not a differential equation) and its solution trajectory is typically not unique. Actually, it is rather natural to assume a *continuous* type distribution for

[16]Note that we do not need a uniform bound on $L_{\mathbf{F}}(\theta)$. Assumption 1 is satisfied in an ASAG, as long as the common payoff function $\mathbf{F}^0 : \mathbb{R}^S \to \mathbb{R}^S$ is Lipschitz continuous.

[17]Assumption 2 is satisfied in an ASAG, if the type distribution μ has a bounded support and the common payoff function $\mathbf{F}^0$ is continuous, even if the switching rate function itself is not bounded over the whole domain $\mathbb{R}^S$ like the Smith dynamic.

[18]Then, the population-wide Lipschitz constant $\bar{L}_{\mathbf{F}}$ is finite and thus Assumption 1 holds. For Assumption 2, recall continuity of $\rho_{ss'}$ for L-continuous revision protocols and $Q_{ss'}$ for exact optimization revision protocols.

smoothing out the discontinuity. Now we specify about what aspect of the type distribution we should assume continuity.

For this, we start from clarifying the cause of discontinuity in an exact optimization protocol. We have indeed assumed Lipschitz continuity of $Q_{ss'}$. Why cannot it guarantee Lipschitz continuity of revision protocol $\rho_{ss'}$? It is due to truncation when the best response strategy changes. The continuity of $Q_{ss'}$ assures continuous change in switching rate $\rho_{ss'}$ with the payoff vector, when strategy s' remains to be the unique best response. However, payoff changes may cause changes in the best responses, which trigger discontinuous changes in the switching rates: the switching rate $\rho_{ss'}$ to the new best response strategy changes from zero to some positive rate $Q_{ss'}$ and the switching rate to the old one changes from positive to zero.

In the next assumption we consider a change in the strategy distribution from $\mathbf{X}$ to $\mathbf{X}'$. We look at agents whose types belong to both $\Theta^{\mathbf{F}}_{s\in\mathrm{BR}}(\mathbf{X})$ and $\Theta^{\mathbf{F}}_{s'\in\mathrm{BR}}(\mathbf{X}')$. These agents' best responses change from s to s'. The left hand side $\mu(\ldots)$ in the assumption collects the mass of such agents. The assumption requires that this mass grows only (at rapidest) proportionally with the distance between the old and new strategy distributions $\|\mathbf{X}-\mathbf{X}'\|_\infty$. This assumption implies that, despite discontinuous changes in individual agents' switching rates, the sum of these changes over all the agents is continuous.

Assumption 3 (Continuous changes in best responses)**.** *If* revision protocol $\boldsymbol{\rho} : \mathbb{R}^S \to \mathbb{R}^{S\times S}_+$ is an exact optimization protocol, then there exists $L_{\mathrm{BR}} \in \mathbb{R}_+$ such that

$$\mu(\Theta^{\mathbf{F}}_{s\in\mathrm{BR}}(\mathbf{X}) \cap \Theta^{\mathbf{F}}_{s'\in\mathrm{BR}}(\mathbf{X}')) \leq L_{\mathrm{BR}}\|\mathbf{X}-\mathbf{X}'\|_\infty$$

for any two distinct strategies $s, s' \in \mathcal{S}$ such that $s \neq s'$ and any $\mathbf{X}, \mathbf{X}' \in \mathcal{X}$.

Remark 1. In an ASAG, Assumption 3 is satisfied if the cumulative distribution of differences in idiosyncratic payoffs between every two strategies satisfies a Lipschitz-like continuity in the following sense: let $\Delta_{ss'}(d) = \mu(\{\theta \in \Theta : \theta_{s'} - \theta_s \leq d\})$, i.e., $\Delta_{ss'}$ is a c.d.f. of the difference in idiosyncratic payoffs between strategies s and s'. Then, there exists $\bar{m} \in \mathbb{R}$ such that $|\Delta_{ss'}(d') - \Delta_{ss'}(d)| \leq \bar{m}|d'-d|$ for any $s, s' \in \mathcal{S}$ and any $d, d' \in \mathbb{R}$. ∎

Assumption 3 implies that the best response is unique for μ-almost all types (let $\mathbf{X} = \mathbf{X}'$). Note that this assumption imposes the condition on the type distribution only if the revision protocol is an exact optimization protocol; L-continuous revision protocols do not need any such assumption on the type distribution for the existence of a unique solution trajectory.

Theorem 1 (Lipschitz continuity of $\mathbf{V}^{\mathbf{F}}$)**.** *Consider a heterogeneous dynamic $\mathbf{V}^{\mathbf{F}}$ in a population game $\mathbf{F}$. Then, function $\mathbf{V}^{\mathbf{F}}$ is Lipschitz continuous in variational norm $\|\cdot\|_\infty$, if (i) $\mathbf{V}^{\mathbf{F}}$ is built upon an L-continuous revision protocol and satisfies Assumptions 1 and 2, or (ii) it is upon an exact optimization protocol and satisfies Assumption 3 as well as Assumptions 1 and 2.*

Corollary 1 (Existence of a unique solution trajectory)**.** *If $\mathbf{V}^{\mathbf{F}}$ satisfies the assumptions for Theorem 1, then there exists a unique solution trajectory $\{\mathbf{X}_t\}_{t\in\mathbb{R}_+} \subset \mathcal{X}$ of $\dot{\mathbf{X}}_t = \mathbf{V}^{\mathbf{F}}[\mathbf{X}]$ from any initial strategy distribution $\mathbf{X}_0 \in \mathcal{X}$.*

For (homogeneous) evolutionary dynamics on a continuous strategy space, Oechssler and Riedel (2001) and Cheung (2014) prove the existence of a solution trajectory by applying Picard-Lindelöf theorem as well. While both ours and theirs deal with the dynamic of probability measure on a (possibly) continuous space, our assumptions and details in the proof are unique to our mathematical problem that involves continuity in a *type* space and also aims to offer a general framework to cover various dynamics/revision protocols, especially exact optimization protocols. We discuss these differences below.

Remark 2. One of the differences comes from the essential defining nature of heterogeneous dynamics that each agent is born with a certain type θ and posses it persistently; so $X_s(B)$ can never exceed $\mu(B)$ for any $B \in \mathcal{B}, s \in \mathcal{S}$. As argued in Section 2, this assures $\mu \gg \mathbf{X}$, i.e., the absolute continuity of $\mathbf{X}$ with respect to μ. This enables us to obtain a strategy density function $\mathbf{x}$ as a density of $\mathbf{X}$ w.r.t. μ and to interpret $x_s(\theta)$ as a proportion of strategy-s players among type-θ agents. Since an agent's strategy revision crucially depends on the own type, it is natural to construct dynamic $\mathbf{v}$ of strategy density function $\mathbf{x}(\theta)$ at each θ, as in (4); then, dynamic $\mathbf{V}$ of strategy distribution $\mathbf{X}$ is just derived from $\mathbf{v}$, as in (5).

On the other hand, in continuous-*strategy* evolutionary dynamics, an agent is assumed to be homogeneous and thus has no persistent characteristic. When these dynamics need a distribution that dominates the strategy distribution to obtain absolute continuity, they *create* some *ad hoc* distribution artificially from the strategy distribution.[19] A continuous-strategy evolutionary dynamic typically defines the transition of the *measure* (the mass of players in a Borel *set* of strategies) directly; they obtain a density only to prove Lipschitz continuity. Thus, a dominating distribution for absolute continuity is only an artificial addition to continuous-*strategy* evolutionary dynamics, not an essential component of games or dynamics. ∎

Remark 3. Another difference is that we cover exact optimization dynamics such as the standard and tempered BRDs, whose revision protocols $\boldsymbol{\rho}$ are discontinuous. As far as the author is aware of, the studies on continuous-strategy evolutionary dynamics focus on L-continuous revision protocols: imitative dynamics (Oechssler and Riedel, 2001; Cheung, 2016), the BNN dynamic (Hofbauer et al., 2009), the gradient dynamic (Friedman and Ostrov, 2013), payoff comparison dynamics (Cheung, 2014) and the logit dynamic (Lahkar and Riedel, 2015). ∎

Sketch of the proof of Theorem 1

For the existence of a unique solution path, we use a version of Picard-Lindelöf theorem as below.

Theorem 2 (adopted from Ely and Sandholm 2005, Theorem A.3.). *Suppose that a dynamic* $\dot{\mathbf{x}}_t = \mathbf{v}^{\mathbf{F}}[\mathbf{x}_t]$ *with* $\mathbf{v}^{\mathbf{F}} : \mathcal{F}_{\mathcal{X}} \to T\mathcal{F}_{\mathcal{X}} := \{\dot{\mathbf{x}} : \Theta \to \mathbb{R}^A\}$ *satisfies*

- *Lipschitz continuity in* L^1 *norm on* $\mathcal{F}_{\mathcal{X}}$;
- *forward invariance of* $\mathcal{F}_{\mathcal{X}}$, *i.e.,* $\sum_{s\in\mathcal{S}} v_s^{\mathbf{F}}[\mathbf{x}](\theta) = 0$ *and* $[x_s(\theta) = 0 \;\Rightarrow\; v_s^{\mathbf{F}}[\mathbf{x}](\theta) \geq 0]$ *for any* $\mathbf{x} \in \mathcal{F}_{\mathcal{X}}$ *and* $\theta \in \Theta$; *and*
- *uniform boundedness of* $\mathbf{v}^{\mathbf{F}}$, *i.e., there exists* $M > 0$ *s.t.* $|v^{\mathbf{F}}[\mathbf{x}](\theta)| \leq M$ *for any* $\mathbf{x} \in \mathcal{F}_{\mathcal{X}}$ *and* $\theta \in \Theta$.

[19] For example, see Oechssler and Riedel (2001, p.159) and Cheung (2014, p.2 in Online Appendix).

Consider a Lipschitz continuous extension of $v^{\mathbf{F}}$ from $\mathcal{F}_{\mathcal{X}}$ to an affine space $\tilde{\mathcal{F}}_{\mathcal{X}} := \{\mathbf{x} : \Theta \to \mathbb{R}^A : \sum_{s\in\mathcal{S}} x_s(\theta) = 1$ for any $\theta \in \Theta\}$. Then, there exists a unique solution path $\{\mathbf{x}_t\}_{t\in\mathbb{R}_+}$ from any initial state $\mathbf{x}_0 \in \mathcal{F}_{\mathcal{X}}$. It is Lipschitz continuous with respect to $\mathbf{x}_0$ and remains in $\mathcal{F}_{\mathcal{X}}$ for all time $t \in \mathbb{R}_+$.

Forward invariance is guaranteed from our construction of an evolutionary dynamic as in (4). Given $\mathbf{x}(\theta) \in \Delta^A$, it is easy to see Assumption 2 guarantees a uniform bound on $\mathbf{v}^{\mathbf{F}}$. The remaining is Lipschitz continuity. Since the variational norm on $\mathcal{X}$ is equivalent to the L^1 norm on $\mathcal{F}_{\mathcal{X}}$, Lipschitz continuity of $\mathbf{v}^{\mathbf{F}}$ on $\mathcal{F}_{\mathcal{X}}$ is equivalent to that of $\mathbf{V}^{\mathbf{F}}$ on $\mathcal{X}$.[20] Similarly, the conclusion of this theorem reduces to the existence of a unique solution path of $\dot{\mathbf{X}} = \mathbf{V}^{\mathbf{F}}[\mathbf{X}]$ and its Lipschitz continuity in $(\mathcal{X}, \|\cdot\|_\infty)$.

In Appendix B, we prove Lipschitz continuity of $\mathbf{V}^{\mathbf{F}}$ by finding $L_V^{\mathbf{F}} > 0$ such that

$$\|\mathbf{V}^{\mathbf{F}}[\mathbf{X}] - \mathbf{V}^{\mathbf{F}}[\mathbf{X}']\|_\infty \leq L_V^{\mathbf{F}}\|\mathbf{X} - \mathbf{X}'\|_\infty \qquad \text{for any } \mathbf{X}, \mathbf{X}' \in \mathcal{X}. \tag{6}$$

For an L-continuous revision protocol, the Lipschitz continuity of $\mathbf{V}^{\mathbf{F}}$ is a natural consequence of the Lipschitz continuity of switching rate function $\boldsymbol{\rho}^\theta$ and of payoff function $\mathbf{F}$.

On the other hand, an exact optimization protocol is discontinuous. If the best response strategies for some type of agents have changed by a change in the strategy distribution from $\mathbf{X}$ to $\mathbf{X}'$, these agents should experience discontinuous changes in the switching rates. However, these discontinuous changes in their switching rates are bounded thanks to Assumption 3. The assumption also implies that the mass of agents who belong to such types increases only Lipschitz-continuously with the change in the strategy distribution.[21] As a result, the aggregate change in their switching rates grows only continuously. This mitigation of discontinuity in an exact optimization protocol by continuity of the type distribution marks the second difference from the preceding studies on continuous-*strategy* evolutionary dynamics.

4 Equilibrium stationarity and stability

Our heterogeneous dynamics could be seen as an extension of evolutionary dynamics in a single homogeneous population to (possibly) continuously many heterogeneous subpopulations, though the existence of a unique solution trajectory requires careful formulation of the state space. It is natural to expect that stationarity and stability of Nash equilibria are extended to *equilibrium strategy distributions* in the heterogeneous setting.

We first define the properties of evolutionary dynamic $\mathbf{v}$ that induce stationarity and stability of equilibria, separately from the population game.[22] This separation is useful because both homogeneous and heterogeneous dynamics stem from the same evolutionary dynamic

[20] Since we will argue the measure of types who experience changes in the best response strategies, it is indeed clearer to work with measure-based $\mathbf{V}^{\mathbf{F}}$ on $\mathcal{X}$ than density-based $\mathbf{v}^{\mathbf{F}}$ on $\mathcal{F}_{\mathcal{X}}$.

[21] Note that this assumption also restricts the mass of types who have multiple best responses to a null set (zero measure) in μ.

[22] This separation accords with the view proposed by Sandholm (2010). Sandholm (2010, especially, Sec. 1.2.2 and Ch.4) proposes to construct an evolutionary dynamic $\mathbf{v}$ from agents' revision protocol $\boldsymbol{\rho}$, separately from a game $\mathbf{F}$, and thus has guided our attention to individual decision rules behind the collective population dynamic. Note that, to clarify that Definition 3 defines a property of $\mathbf{V}$, independently of specification of a game (especially whether the population is homogeneous or heterogeneous), we name it best response stationarity; Nash stationarity (Sandholm, 2010), which refers to stationarity of $\mathbf{V}^{\mathbf{F}}$ at Nash equilibria, makes sense only after specifying a game.

$\mathbf{v}$ (constructed from the same revision protocol $\boldsymbol{\rho}$). Their difference lies only in the difference in the population game played by agents, namely the difference between $\mathbf{F} : \mathcal{X} \to \mathcal{C}$ and $\mathbf{F}^0 : \Delta^S \to \mathbb{R}^S$.

In the homogeneous setting, the stationarity of a Nash equilibrium under $\mathbf{v}^{\mathbf{F}^0}$ is an immediate consequence of *best response stationarity* under $\mathbf{v}$; the evolutionary dynamic stays at a strategy distribution if and only if agents are taking the best response to the current payoffs.

Definition 3 (BR stationarity of evolutionary dynamic)**.** Evolutionary dynamic $\mathbf{v} : \Delta^S \times \mathbb{R}^S \to \mathbb{R}^S$ satisfies **best response (BR) stationarity** if, for any $\boldsymbol{\pi}^0 \in \mathbb{R}^S, \mathbf{x}^0 \in \Delta^S$,

$$\mathbf{v}(\boldsymbol{\pi}^0, \mathbf{x}^0) = \mathbf{0} \qquad \Longleftrightarrow \qquad \mathbf{x}^0 \in \Delta(\mathcal{S}_{\mathrm{BR}}(\boldsymbol{\pi}^0)). \tag{7}$$

All the evolutionary dynamics mentioned in Section 2.2, except smooth BRDs, satisfy BR stationarity.[23] In a homogeneous population game, the best response stationarity implies the stationarity of a Nash equilibrium and non-stationarity of non-equilibrium states.

The key property of evolutionary dynamics for equilibrium stability is *positive correlation (PC)*: each strategy's payoff and the net increase in the mass of the strategy's players are positively correlated and the correlation is strictly positive unless the strategy distribution is unchanged. Major evolutionary dynamics, except smooth BRDs, satisfy PC.

Definition 4 (Positive correlation of evolutionary dynamic)**.** Evolutionary dynamic $\mathbf{v} : \Delta^S \times \mathbb{R}^S \to \mathbb{R}^S$ satisfies **positive correlation (PC)** if

$$\boldsymbol{\pi}^0 \cdot \mathbf{v}(\boldsymbol{\pi}^0, \mathbf{x}^0) \begin{cases} \geq 0 & \text{for any } \boldsymbol{\pi}^0 \in \mathbb{R}^S, \mathbf{x}^0 \in \Delta^S; \\ > 0 & \text{if } \mathbf{v}(\boldsymbol{\pi}^0, \mathbf{x}^0) \neq \mathbf{0}. \end{cases} \tag{8}$$

While stability of a Nash equilibrium is not generally guaranteed even in the homogeneous setting, it is assured for potential games under a wide class of evolutionary dynamics. In the homogeneous setting, population game $\mathbf{F}^0 : \Delta^{\mathcal{S}} \to \mathbb{R}^S$ is a potential game if there is a differentiable function $f^0 : \Delta^{\mathcal{S}} \to \mathbb{R}$, called a potential function, such that $\nabla f^0 \equiv \mathbf{F}^0$.[24] PC immediately implies that the value of f^0 increases over time until it reaches a stationary point, since the definition of a potential function implies $\dot{f}^0(\mathbf{x}^0) = \nabla f^0(\mathbf{x}^0) \cdot \dot{\mathbf{x}}^0 = \mathbf{F}^0(\mathbf{x}^0) \cdot \dot{\mathbf{x}}^0 = \mathbf{F}^0(\mathbf{x}^0) \cdot \mathbf{v}(\mathbf{F}^0(\mathbf{x}^0), \mathbf{x}^0)$. Hence, the homogeneous *potential* function f^0 works as a Lyapunov function commonly in these evolutionary dynamics and thus PC assures stability of local maxima of f^0 (Sandholm, 2001).

If a dynamic satisfies BR stationarity and PC, we call it an **admissible** dynamic. Pairwise comparison dynamics and exact optimization dynamics are admissible dynamics.[25]

[23] In the homogeneous version of exact optimization dynamics, BR stationarity needs to assume $\rho_{ss'}(\boldsymbol{\pi}) = 0$ when the current strategy s is a best response to $\boldsymbol{\pi}$; this was not assumed in our definition in cases of multiple best responses. In the heterogeneous setting, this concern on multiple best responses is eliminated by Assumption 3. Hence, this assumption replaces the assumption of $\rho_{ss'}(\boldsymbol{\pi}) = 0$ for best response s to $\boldsymbol{\pi}$.

[24] Having a potential function is equivalent to externality symmetry: the change in the payoff of a strategy by a change in the mass of another strategy's players is symmetric between these two strategies. The class of potential games includes random matching in common interest games, binary games and nonatomic congestion games. Sandholm (2010, Chapter 3) provides further explanation and examples.

[25] Smooth BRDs satisfy analogous properties of Nash stationarity and PC for perturbed payoffs; see Sandholm (2010, §6.2.4). For observational dynamics such as the replicator dynamic and excess payoff dynamics, see Section 5.1. See Sandholm (2010, Chapter 5) for summary of the relationship between dynamics and the two properties in this section.

4.1 Equilibrium stationarity in general

In the heterogeneous setting, the best response stationarity applies to each type: the strategy distribution of a particular type θ remains unchanged if and only if almost all agents of this type choose the best response to the current payoff for this type. Thus, it is straightforward that the best response stationarity implies the stationarity of an equilibrium strategy distribution and non-stationarity of non-equilibrium strategy distributions.

Theorem 3 (Equilibrium stationarity in a population game)**.** *Suppose that evolutionary dynamic* $\mathbf{v}$ *satisfies BR stationarity* (7)*. Then, in any heterogeneous population game* $\mathbf{F}$*, an equilibrium strategy distribution is stationary under the heterogeneous combined dynamic* $\mathbf{V^F}$ *derived from these* $\mathbf{v}$ *and* $\mathbf{F}$*, and vice versa:*[26]

$$\mathbf{V^F}[\mathbf{X}] = \mathbf{O} \qquad \Longleftrightarrow \qquad \mathbf{X} \text{ is an equilibrium strategy distribution in } \mathbf{F}. \tag{9}$$

This theorem implies that the existence of a stationary point is equivalent to that of an equilibrium state. Following the outline of the proof for the existence of a distributional equilibrium in an incomplete information game of finitely many players by Milgrom and Weber (1985), we can guarantee the existence of an equilibrium strategy distribution in a heterogeneous population game that exhibits a kind of uniform continuity and boundedness over types of the payoff function.

Theorem 4 (Existence of equilibrium in a population game)**.** *Suppose that* $\mathbf{F} : \mathcal{X} \to \mathcal{C}$ *satisfies Assumption 1,* ***equicontinuity over types*** *(with respect to the weak topology metrized by Prokhorov metric* d^{Pr}*),*[27] *i.e., for any* $\mathbf{X} \in \mathcal{X}, \varepsilon > 0$*, there exists* $\delta_{Ct}[\mathbf{X}] > 0$ *such that*

$$d^{Pr}(\mathbf{X}, \mathbf{X}') < \delta_{Ct}[\mathbf{X}] \quad \Longrightarrow \quad \left[|F_s[\mathbf{X}](\theta) - F_s[\mathbf{X}'](\theta)| < \varepsilon \quad \text{for any } s \in \mathcal{S} \text{ and } \mu\text{-almost all } \theta \in \Theta \right],$$

and ***near-boundedness over types****, i.e., for any* $\mathbf{X} \in \mathcal{X}, \varepsilon > 0$*, there exists* $\bar{F}[\mathbf{X}] \geq 0$ *such that*

$$\int_\Theta \sum_{s \in \mathcal{S}} [|F_s[\mathbf{X}](\theta)| - \bar{F}[\mathbf{X}]]_+ \cdot \mu(d\theta) < \varepsilon.$$

Then, there exists an equilibrium strategy distribution in the heterogeneous population game $\mathbf{F}$*.*

Corollary 2. *Under the assumptions for Theorems 3 and 4, heterogeneous dynamic* $\mathbf{V^F}$ *has a stationary state.*

If $\mathbf{F}[\mathbf{X}] : \Theta \to \mathbb{R}^S$ is bounded over Θ, i.e., there exists $\bar{F}[\mathbf{X}] \in \mathbb{R}_+$ such that $|F_s[\mathbf{X}](\theta)| \leq \bar{F}[\mathbf{X}]$ for any $s \in \mathcal{S}$ and μ-almost any type θ, then it is nearly bounded over types. The *near*-boundedness allows $\mathbf{F}[\mathbf{X}]$ to be unbounded as long as $|F_s[\mathbf{X}](\theta)|$ exceeds $\bar{F}[\mathbf{X}]$ only in a sufficiently small mass of agents. It is satisfied as long as the expected value of $|F_s[\mathbf{X}]|$ over types exists, i.e., $\int |F_s[\mathbf{X}](\theta)|\mu(d\theta) < \infty$ for each $s \in \mathcal{S}$. In an ASAG, this reduces to $\int |\theta_s| d\mu < \infty$; so, it holds, for example, if $\boldsymbol{\theta}$ follows the double-exponential distribution just as assumed in the logit choice model.

Just like Milgrom and Weber (1985), we use Glicksberg's fixed point theorem (Aliprantis and Border, 2006, Corollary 17.55), since an equilibrium strategy distribution can be formulated as a fixed point of the "distributional" best response correspondence (check $B[\mathbf{X}]$ in Appendix C.2). But the objective function in the best response correspondence is different from

[26] $\mathbf{O} = (O_s)_{s \in \mathcal{S}} \in \mathcal{M}$ denotes a zero measure such as $O_s(B) = 0$ for any $B \in \mathcal{B}, s \in \mathcal{S}$.

[27] See Appendix A.1 for the weak topology and Prokhorov metric d^{Pr}.

theirs because of difference in the player set and in the strategy set. Thus, we need to prove continuity of the objective function specifically for our population-game setting.

4.2 Equilibrium stability in potential games

Heterogeneous (weighted) potential games

For a game played in large population, a potential game is defined as a game in which payoff vector can be derived as the derivative of some scalar-valued function, i.e., a potential function. By generalizing this idea to a function defined on the (possibly infinite-dimensional) space of strategy distributions, we define a heterogeneous potential game.

Definition 5 (Heterogeneous potential game)**.** Heterogeneous population game $\mathbf{F}: \mathcal{S} \to \mathcal{C}$ is called a **heterogeneous weighted potential game** if there are a continuous function $w: \Theta \to \mathbb{R}_{++}$ and a scalar-valued Fréchet-differentiable function $f: \mathcal{X} \to \mathbb{R}$ that is continuous in the weak topology on $\mathcal{X}$ and whose Fréchet derivative coincides with $w\mathbf{F}$: at each strategy distribution $\mathbf{X} \in \mathcal{X}$, the payoff vector function $\mathbf{F}[\mathbf{X}] \in \mathcal{C}$ satisfies[28]

$$f(\mathbf{X}') = f(\mathbf{X}) + \langle w\mathbf{F}[\mathbf{X}], \mathbf{X}' - \mathbf{X}\rangle + o(\|\mathbf{X}' - \mathbf{X}\|_\infty) \qquad \text{for any } \mathbf{X}' \in \mathcal{X}.$$

We call f a (heterogeneous) w-weighted potential function for $\mathbf{F}$. If $w \equiv 1$, f is called an (exact) potential function and $\mathbf{F}$ an (exact) potential game.

Both in the homogeneous and heterogeneous settings, all local maxima and interior local minima of a potential function, and indeed all the solutions of the Karush-Kuhn-Tucker first-order condition for maxima are equilibria in a potential game; see Sandholm (2001) for the proof for Nash equilibria in a homogeneous potential game and Sandholm (2005, Appendix A.3) for equilibrium strategy distributions in a heterogeneous potential game.

Stability and potential maximization

In the heterogeneous setting, PC of $\mathbf{v}$ implies a positive correlation between the payoffs and the strategy distribution among *each type* of agents. Thus, by the same token as in a homogeneous potential game, this guarantees that the heterogeneous potential function f works as a Lyapunov function for equilibrium stability in a heterogeneous potential game $\mathbf{F}$. Hence, part i) in the theorem below is a natural extension of equilibrium stability in a homogeneous potential game (Sandholm, 2001) to a heterogeneous setting. Part ii) verifies the opposite relationship; that is, stability of an equilibrium in any admissible dynamic implies local maximum of the potential. Our proof applies to a homogeneous setting, though the result has not been mentioned in the literature even in a homogeneous setting as far as the author is aware of.

[28] Here, operator $\langle \cdot, \cdot \rangle$ is defined as $\langle \boldsymbol{\pi}, \Delta\mathbf{X}\rangle = \int_\Theta \boldsymbol{\pi}(\theta) \cdot \Delta\mathbf{X}(d\theta)$ and $w\mathbf{F}: \mathcal{X} \to \mathcal{C}$ is defined as $(w\mathbf{F}[\mathbf{X}])(\theta) = w(\theta)\mathbf{F}[\mathbf{X}](\theta)$. The norm $\|\cdot\|_\infty$ is the variational norm on $\mathcal{X}$ to metrize the strong topology; see Appendix A.2. Fréchet differentiability is defined for the strong topology and thus continuity in the weak topology (stronger than continuity in the strong topology) is additionally required. Note that Fréchet differentiation is a generalization of total differentiation, while another differentiation in a Banach space, Gateaux differentiation, generalizes directional differentiation. Since $\dot{\mathbf{x}}$ may not stay in the same direction, we need total differentiation.

Theorem 5 (Equilibrium stability of heterogeneous potential games)**.** *Suppose that evolutionary dynamic* $\mathbf{v}$ *satisfies PC* (8) *as well as Assumptions 1 to 3. Then, in any heterogeneous potential game* $\mathbf{F}$, *the following holds.*

i) The set of stationary strategy distributions $\{\mathbf{X} \in \mathcal{X} : \mathbf{V}^{\mathbf{F}}[\mathbf{X}] = \mathbf{O}\}$ *is globally attracting under* $\mathbf{V}^{\mathbf{F}}$. *A local maximum (local strict maximum, resp.) of* f *is Lyapunov stable (asymptotically stable, resp.).*

ii) Let $\mathbf{X}^*$ *be an isolated stationary strategy distribution in the sense that, in a neighborhood* $\mathcal{X}^*$ *of* $\mathbf{X}^*$ *in the space* $\mathcal{X}$, *there is no other stationary strategy distribution than* $\mathbf{X}^*$. *a) If it is (locally) asymptotically stable, then it is a local strict maximum of* f. *b) Further assume that* $\gamma : \mathcal{X} \to \mathbb{R}$ *defined as* $\gamma(\mathbf{X}) = \langle w\mathbf{F}[\mathbf{X}], \mathbf{V}^{\mathbf{F}}[\mathbf{X}]\rangle$ *is continuous in weak topology.*[29] *If* $\mathbf{X}^*$ *is Lyapunov stable, then it is a local maximum of* f.

Since equilibria and potential maximizers can be found solely from $\mathbf{F}$ independently of the dynamic $\mathbf{v}$, Theorems 3 and 5 imply that the set of stationary states and the set of locally stable states are common to all admissible dynamics.

Corollary 3. *Consider admissible dynamic* $\mathbf{V}$ *in a heterogeneous population game* $\mathbf{F}$ *that satisfies Assumptions 1 to 3.*[30]

i) Strategy distribution $\mathbf{X}^*$ *is an equilibrium in* $\mathbf{F}$, *if and only if it is stationary in any those dynamics.*

ii) Further, suppose that $\mathbf{F}$ *is a heterogeneous potential game. Then, the set of equilibrium strategy distributions is globally attracting under any those dynamics. Isolated equilibrium strategy distribution* $\mathbf{X}^*$ *is a local maximum of potential function* f, *if and only if it is Lyapunov stable under any those dynamics; strictness of a local maximum is equivalent to asymptotic stability.*

That is, just like in the homogeneous setting, specification of evolutionary dynamics do not matter in a heterogeneous potential game to assure global stability of the Nash equilibrium set or to tell which equilibrium is locally stable. This commonality of equilibrium stability is appealing to applications; see Example 9.

Naturally in our canonical three examples, if the base game is a potential game, then its heterogeneous versions are also potential games.

Example 1′. Recall Example 1. Assume that the base homogeneous game $\mathbf{F}^0$ is a potential game with potential function $f^0 : \Delta^{\mathcal{S}} \to \mathbb{R}$ such that $\nabla f^0 \equiv \mathbf{F}^0$. Then, the ASAG is a heterogeneous potential game, with potential function $f : \mathcal{X} \to \mathbb{R}$ such as[31]

$$f(\mathbf{X}) = f^0(\mathbf{X}(\Theta)) + \int_\Theta \boldsymbol{\theta} \cdot \mathbf{X}(d\boldsymbol{\theta}) \qquad \text{for each } \mathbf{X} \in \mathcal{X}. \qquad \blacksquare$$

[29] Continuity of γ in weak topology is assured if continuity of $\mathbf{F}$ (Assumption 1) and $\mathbf{V}$ (Definition 1 and assumption 3) are strengthened to that in weak topology; assumptions in Footnote 16 and assumption 3 suffice it in an ASAG. γ is continuous in strong topology if $\mathbf{F}$ and $\mathbf{V}^{\mathbf{F}}$ are continuous in strong topology, which is guaranteed by Assumption 1 and Theorem 1. But continuity in weak topology is stronger than that in strong topology, since convergence in the former is weaker than that in the latter (and then the value of a "continuous" function must approach arbitrarily close to a limit).

[30] In each of the two claims i,ii), the former condition (equilibrium/potential maximum) is sufficient for the latter (stationarity/stability) to hold under *all* admissible heterogeneous dynamics, while the latter under *any* (single) admissible dynamic is sufficient for the former.

[31] This function f appears in the study of evolutionary implementation by Sandholm (2005, Appendix A.3). But it was used there only to characterize an equilibrium as a solution of the KKT condition for local maxima and minima of f.

Example 2′. Recall Example 2.[32] Assume that, with continuous weighting function $w : \Theta \to \mathbb{R}_{++}$ common over all states $\sigma \in \Sigma$, $f^\sigma : \mathcal{X} \to \mathbb{R}$ is a w-weighted potential function for base game $\mathbf{F}^\sigma : \mathcal{X} \to \mathbb{R}^S$ in each state $\sigma \in \Sigma$ in the sense that, for any $\mathbf{X} = \int \mathbf{x} d\mu, \mathbf{X}' = \int \mathbf{x}' d\mu$,

$$f^\sigma(\mathbf{X}') = f^\sigma(\mathbf{X}) + \int w(\theta)\mathbf{F}^\sigma[\mathbf{X}] \cdot \{\mathbf{x}'(\theta) - \mathbf{x}(\theta)\} d\mathbb{P}_{\Theta|\sigma}(d\theta|\sigma) + \|\mathbf{X}' - \mathbf{X}\|_\infty.$$

Then, the Bayesian game $\mathbf{F}$ is a w-weighted potential game with w-weighted potential function $f : \mathcal{X} \to \mathbb{R}$ such that

$$f(\mathbf{X}) := \int_{\sigma\in\Sigma} f^\sigma(\mathbf{X})\mathbb{P}_\Sigma(d\sigma) \qquad \text{for each } \mathbf{X} \in \mathcal{X}. \quad ▮$$

Example 3′. Recall Example 3. Now assume that the base game is a weighted potential game with potential function $f^0 : \Delta^{\mathcal{S}} \times \Delta^{\mathcal{S}} \to \mathbb{R}$, i.e., $\nabla_1 f^0(\mathbf{x},\mathbf{x}') = w_1\mathbf{F}^0(\mathbf{x},\mathbf{x}'), \nabla_2 f^0(\mathbf{x},\mathbf{x}') = w_2\mathbf{F}^0(\mathbf{x}',\mathbf{x})$, where $w_i \in \mathbb{R}_{++}$ is a positive constant and $\nabla_i f^0$ is the gradient vector of f^0 with respect to the strategy distribution in the i-th argument. (Recall the first argument is the strategy distribution in the own population and the second is that in the opponent.) Further, assume that the weight function g is symmetric: $g(\theta,\theta') = g(\theta',\theta)$. Then, the structured population game is an (exact) potential game, with potential function $f : \mathcal{X} \to \mathbb{R}$ such that

$$f(\mathbf{X}) = \frac{1}{w_1 + w_2} \int_{(\theta,\theta')\in\Theta^2} f^0(\mathbf{x}(\theta),\mathbf{x}(\theta'))g(\theta,\theta')\mu(d\theta)\mu(d\theta') \qquad \text{for each } \mathbf{X} = \int \mathbf{x} d\mu \in \mathcal{X}. \quad ▮$$

Theorem 6. *The heterogeneous population games in the above examples are potential games.*

The following corollary makes a bridge between stability in a base game and that in a modified heterogeneous game. The claim in each part can be immediately proven from the above construction of the potential functions.

Corollary 4. *i) Consider Example 2′. Assume that f^0 is locally concave around $\bar{\mathbf{x}}^*$: that is, f^0 is concave in some neighborhood of $\bar{\mathbf{x}}^*$. Then f is locally concave around any $\mathbf{X}^*$ such that $\mathbf{X}^*(\Theta) = \bar{\mathbf{x}}^*$.*

ii) Consider Example 2′. Assume that $\mathbf{X}^*$ attains a local (locally strict, resp.) maximum of f^σ in base game $\mathbf{F}^\sigma$ commonly over all the states σ. Then, $\mathbf{X}^*$ attains a local (locally strict, resp.) maximum of f in the Bayesian game $\mathbf{F}$.

iii) Consider Example 3′.Suppose that symmetric strategy profile $(\bar{\mathbf{x}}^*, \bar{\mathbf{x}}^*)$ is a local (locally strict, resp.) maximizer of f^0 in the two-population base game $\mathbf{F}^0$. Let $\mathbf{X}^*$ be constructed from strategy density function $\mathbf{x}^* \in \mathcal{F}_\mathcal{X}$ such as $\mathbf{x}(\theta) = \bar{\mathbf{x}}^*$ for μ-almost all θ. Then, $\mathbf{X}^*$ attains a local (locally strict, resp.) maximum of f in the structured potential game $\mathbf{F}$.

This corollary suggests that local stability is robust to the introduction of payoff perturbation, incomplete information and uneven interaction structure as long as the base game exhibits a potential function and the dynamic is admissible. Parts ii) and iii) straightforwardly suggests that local stability of an equilibrium in the base game carries over to that of the corresponding equilibrium in the modified game. About part i), assume that $\bar{\mathbf{x}}^*$ in the claim is indeed a Nash equilibrium in the base game; under the local concavity assumption in the

[32] van Heumen et al. (1996) define a Bayesian potential game in a normal form with finitely many players. See Ui (2009) for examples of such games, including team production problems. As suggested by Ui, the continuous-population game studied by Angeletos and Pavan (2007) is a Bayesian potential game.

proposition, it must be a local maximum of $\mathbf{f}^0$ and thus Lyapunov stable. Further, consider a *neutral* payoff perturbation in the sense that the ASAG $\mathbf{F}$ under the perturbation keeps an equilibrium strategy distribution $\mathbf{X}^*$ that aggregates to $\bar{\mathbf{x}}^*$. Then, part ii) implies that $\mathbf{X}^*$ is also Lyapunov stable. Furthermore, if $\bar{\mathbf{x}}^*$ is asymptotically stable and thus a strict maximum of $\mathbf{f}^0$, then the set of such equilibrium strategy distributions that aggregates to $\bar{\mathbf{x}}^*$ is also asymptotically stable.

Applications

If $\mathbf{X}^*$ attains a local strict maximum, then it must be asymptotically stable in any admissible dynamics. Thus, once we establish asymptotic stability of $\mathbf{X}^*$ under some particular admissible dynamic in a potential game, then it holds robustly over specifications of the dynamic, as long as agents' choices meet the two intuitive assumptions, i.e., best response stationarity and positive correlation. We see below applications of this positive result to ASAGs.

Example 8 (Convergence to a free-entry equilibrium.). Consider a binary homogeneous game $\mathcal{S} = \{I, O\}$ with *negative* externality: $F_I^0(\bar{x}_I)$ decreases with $\bar{x}_I \in [0,1]$ and $F_O^0 \equiv 0$. Then, the potential function $f^0 : [0,1] \to \mathbb{R}$ is given by $f^0(\bar{x}_I) = \int_0^{\bar{x}_I} F_I^0(\bar{y})d\bar{y}$ and strictly concave. With the boundedness of the domain $[0,1]$, the strict concavity of f^0 implies that the global maximum exists uniquely and there is no other local maximum of f^0. The global maximum of f^0 is the only equilibrium of this game.

For an example in microeconomic theory to fall into this class of games, consider an entry-exit game played by suppliers in a particular industry. To make entry and exit symmetric, it is conventionally assumed that fixed costs exist but they are not sunk: fixed costs are paid only to maintain production capacities and they are revocable when the supplier becomes inactive. Further, the choice of entry or exit is conventionally regarded as a "long run" decision while the choice of the quantity supplied is a "short run" decision (as well as the underlying consumers' decisions on the demand side); thus, it is commonly assumed that the market is settled to a market equilibrium (the state where the demand equals to the total supply) at each moment of time, given the mass (number) of active suppliers at the moment. A free-entry or so-called "long run" equilibrium is characterized in the homogeneous setting as a state in which the gross profit for an active producer is equal to the fixed cost.

One may want to introduce heterogeneity in the suppliers' fixed costs; it not only sounds realistic but also eliminates indeterminacy of individual choices of entry or exit at a free-entry equilibrium. Under heterogeneity in fixed costs, a free-entry equilibrium should be redefined as a state in which all the active producers have smaller fixed costs than the gross profit and all the inactive ones have greater fixed costs.

Under perfect competition in a standard setting as in Mas-Colell et al. (1995, Section 10.F), the instantaneous market-equilibrium profit of an active supplier decreases with the number of active suppliers. We can regard $F_I^0(\bar{x}_I)$ as the gross profit at this instantaneous competitive equilibrium given the current mass $\bar{x}_I$ of active suppliers and $\theta_O(\omega)$ as the fixed costs of supplier ω, while setting $F_O^0 \equiv 0$ and $\theta_I \equiv 0$ for all agents; then, the choice between entry and exit in perfect competition falls into an ASAG with negative externality.

Thanks to our stability result, we can justify the free-entry equilibrium as the globally stable state in an evolutionary dynamic; indeed it is so strengthened to be stable in any admissible dynamics. As argued in Zusai (2018), the tempered BRD can be regarded as a version of the BRD in which a revising agent pays a stochastic switching cost.[33] Thus, the stability of the free-entry equilibrium under the tempered BRD suggests in this context that, even if entry and exit incur sunk costs to build or scrap the production capacity, the "long-run" equilibrium is indeed the long-run limit state under such an entry-exit dynamic.

By the same token, we can justify a free-entry equilibrium in the standard (static) monopolistic competition model such as Dixit and Stiglitz (1977) as a dynamically stable state under an arbitrary admissible dynamic. ∎

Example 9 (Dynamic implementation of the social optimum.). Imagine a central planner whose goal is to maximize the total payoff of agents in an ASAG:

$$\int_{\Theta} \mathbf{F}[\mathbf{X}](\boldsymbol{\theta}) \cdot \mathbf{X}(d\boldsymbol{\theta}) = \mathbf{F}^0(\bar{\mathbf{x}}) \cdot \bar{\mathbf{x}} + \int_{\Theta} \boldsymbol{\theta} \cdot \mathbf{X}(d\boldsymbol{\theta}) \quad \text{with } \bar{\mathbf{x}} = \mathbf{X}(\Theta).$$

To help the central planner achieve this goal, we introduce a monetary transfer to the agent's payoff: now a type-θ agent's payoff from strategy $s \in \mathcal{S}$ is $\tilde{F}^{\mathbf{T}}_s[\mathbf{x}](\theta) := F_s[\mathbf{x}](\theta) - T_s[\bar{\mathbf{x}}]$, where function $\mathbf{T} = (T_s)_{s\in\mathcal{S}} : \Delta^{\mathcal{S}} \to \mathbb{R}^S$ is a pricing scheme to determine the amount of the monetary transfer (in terms of payoff) from the agent to the planner for taking each strategy given aggregate strategy $\mathbf{x} \in \Delta^{\mathcal{S}}$.

Sandholm (2002, 2005) proposes the dynamic Pigouvian pricing scheme such as

$$T_s[\bar{\mathbf{x}}] = -\sum_{s'\in\mathcal{S}} \bar{x}_{s'} \frac{\partial F^0_{s'}}{\partial \bar{x}_s}(\bar{\mathbf{x}}) \quad \text{for each } \bar{\mathbf{x}} \in \Delta^S.$$

Notice that this pricing scheme does not require the central planner to know agents' revision protocols, the type distribution, or even the current strategy distribution; the observation of aggregate strategy $\bar{\mathbf{x}}$ is enough for the planner to update T_s.

Strictly speaking, in a setting where there are *finitely many* payoff types, Sandholm (2002) verify that, with $\mathbf{T}$ being the above dynamic Pigouvian pricing scheme, $\tilde{\mathbf{F}}^{\mathbf{T}}$ has a potential function $\tilde{f}^{\mathbf{T}}$ being the total payoff:

$$\tilde{f}^{\mathbf{T}}(\mathbf{X}) = \int_{\Theta} \mathbf{F}[\mathbf{X}](\boldsymbol{\theta}) \cdot \mathbf{X}(d\boldsymbol{\theta}).$$

In particular, if the common payoff function $\mathbf{F}^0$ exhibits negative externality, $\tilde{f}^{\mathbf{T}}$ is concave and thus the unique social optimum is achieved in the long run through this pricing scheme regardless of the initial state. Thanks to Theorem 5, now we can extend this claim to the games with *infinitely many* payoff types.[34] ∎

[33] For this, we regard Q as a cumulative distribution function of switching costs (with Q scaled to meet $\lim_{q\to\infty} Q(q) = 1$). Then, the revision protocol of the tempered BRD is obtained by having a revising agent first draw a stochastic switching cost q from Q and then switch to the best response strategy only if the payoff gain from the switch, the difference between the maximal payoff and the current payoff, exceeds q.

[34] Sandholm (2005, p.903) speculated it by referring to Ely and Sandholm (2005), which allows us to reduce the heterogeneous standard BRD to a homogeneous smooth BRD of agents by treating persistent heterogeneity in the former as transitory perturbation, i.e., letting an agent draw a new θ from μ at each revision opportunity. However, Zusai (2017) finds that heterogeneous evolutionary dynamics generically cannot reduce to a homogeneous dynamic, except the standard BRD and the smooth BRDs.

5 Extensions

5.1 Observational dynamics

In some of major evolutionary dynamics, an agent observes other agents' strategies and the observation influences the agents' switching decision. For example, an agent may imitate other agents' strategies or the switching rate may depend on the relative payoffs compared to the average payoff of the observed population. We can generalize these dynamics as *observational dynamics* by having the strategy distribution among observed agents $\tilde{\mathbf{x}} \in \Delta^{\mathcal{S}}$, not only payoff vector $\boldsymbol{\pi} \in \mathbb{R}^S$, in the argument of revision protocol $\boldsymbol{\rho}$.

Example 10. With an **excess payoff protocol**, a revising agent calculates the average payoff $\tilde{\mathbf{x}} \cdot \boldsymbol{\pi}$ and switches to strategy s' with the rate that increases with the excess payoff of the new strategy $\pi_{s'} - \tilde{\mathbf{x}} \cdot \boldsymbol{\pi}$. In particular, the revision protocol $\rho_{ss'}(\boldsymbol{\pi}, \tilde{\mathbf{x}}) = [\pi_{s'} - \tilde{\mathbf{x}} \cdot \boldsymbol{\pi}]_+$ defines the **Brown-von Neumann-Nash (BNN) dynamic** (Hofbauer, 2001).[35] ▮

Example 11. With an **imitative protocol**, a revising agent randomly samples another agent's strategy s' according to $\tilde{\mathbf{x}}$ and switches to it with the rate $I_{ss'}(\boldsymbol{\pi}) \in \mathbb{R}_+$: the overall switching rate is $\rho_{ss'}(\boldsymbol{\pi}, \tilde{\mathbf{x}}) = \tilde{x}_{s'} I_{ss'}(\boldsymbol{\pi})$. There are several imitative protocols that yield the **replicator dynamic** (Taylor and Jonker, 1978): imitative pairwise comparison $I_{ss'} = [\pi_{s'} - \pi_s]_+$ (Schlag, 1998), imitation driven by dissatisfaction $I_{ss'} = \bar{\pi} - \pi_s$ with constant $\bar{\pi} \in \mathbb{R}$ (Björnerstedt and Weibull, 1996), and imitation of success $I_{ss'} = \pi_{s'} - \underline{\pi}$ with constant $\underline{\pi} \in \mathbb{R}$ (Hofbauer, 1995a). ▮

They fall into the class of L-continuous revision protocols and satisfy Assumption 2.[36] (Note that Assumption 3 is not needed for L-continuous revision protocols.) We can readily extend all the positive results, i.e., Theorems 1, 3 and 5, to observational dynamics, if we assume that an agent observes the strategy distribution of the same type: a type-θ agent observes $\mathbf{x}(\theta) \in \Delta^{\mathcal{S}}$.[37] This assumption of within-type observability matches with an assumption on imitative dynamics in the society of finitely many subpopulations where a member of each subpopulation imitates the behavior of those in the same subpopulation; for example, Hummel and McAfee (2018) adopt this assumption.[38] The proofs of these theorems in the appendix are indeed written explicitly to include $\mathbf{x}(\theta)$ as an argument of revision protocol $\boldsymbol{\rho}$.

To maintain the existence of a unique solution trajectory (Theorem 1) and stationarity of an equilibrium strategy distribution (Theorem 3), this assumption of within-type observability

[35] Excess payoff dynamics allow for innovation of a new strategy, while imitative dynamics do not. Thus, stationarity and stability of equilibria are restricted to the interior of the strategy space for the latter dynamics while they are not for the former.

[36] Precisely for observational dynamics, $\bar{\rho}$ is an upper bound on $\rho_{ss'}(\mathbf{F}[\mathbf{X}](\theta), \mathbf{x}(\theta))$. As ρ has two arguments, its Lipschitz continuity should mean the existence of L_ρ such as $|\rho_{ss'}(\boldsymbol{\pi}, \tilde{\mathbf{x}}) - \rho_{ss'}(\boldsymbol{\pi}', \tilde{\mathbf{x}}')| \leq L_\rho |(\boldsymbol{\pi}, \tilde{\mathbf{x}}) - (\boldsymbol{\pi}', \tilde{\mathbf{x}}')| = L_\rho(|\boldsymbol{\pi} - \boldsymbol{\pi}'| + |\tilde{\mathbf{x}} - \tilde{\mathbf{x}}'|)$ for any $s, s' \in \mathcal{S}, \boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathbb{R}^S, \tilde{\mathbf{x}}, \tilde{\mathbf{x}}' \in \Delta^{\mathcal{S}}$. (See footnote 12 for our choice of the L^1 norm of finite-dimensional vectors.) Technically, we can allow $\rho_{\cdot\cdot}$ to depend on $\mathbf{X} \in \mathcal{X}$, not only on $\tilde{\mathbf{x}} \in \Delta^{\mathcal{S}}$. Then, this Lipschitz continuity condition is simply generalized as $|\rho_{ss'}(\boldsymbol{\pi}, \mathbf{X}) - \rho_{ss'}(\boldsymbol{\pi}', \mathbf{X}')| \leq L_\rho(|\boldsymbol{\pi} - \boldsymbol{\pi}'| + \|\mathbf{X} - \mathbf{X}'\|_\infty)$. We can easily confirm that the proof for Theorem 1 needs no change, just by glancing over calculation in the proof.

[37] Corollary 3 holds for excess payoff dynamics. In the homogeneous setting, imitative dynamics such as the replicator dynamic satisfy the best response stationarity only if $\mathbf{x}^0$ is in the interior of $\Delta^{\mathcal{S}}$; thus Corollary 3 holds for a solution trajectory of an imitative dynamic starting from the interior of $\mathcal{X}$.

[38] However, Hummel and McAfee (2018) simply apply a general formula of the replicator dynamic and do not construct the dynamic from a revision protocol; thus they do not explicitly discuss imitation.

can be replaced with an alternative assumption that an agent of any type θ observes the aggregate strategy $\bar{\mathbf{x}} = \mathbf{X}(\Theta)$, instead of $\mathbf{x}(\theta)$, as $\tilde{\mathbf{x}}(\theta)$ and the agent's own current payoff vector $\mathbf{F}[\mathbf{X}](\theta)$ as $\boldsymbol{\pi}(\theta)$. But, then PC may not be extended from the homogeneous setting to the heterogeneous setting. If observations are sampled from the entire population, stability analysis becomes essentially different from how we have investigated stability in this paper.[39]

5.2 Heterogeneity in revision protocols

All of our results are robust to heterogeneity in revision protocols. Now, let Θ be the type space and each type θ of agents not only have its peculiar payoff function $\mathbf{F}(\theta)$ but also follow its own revision protocol $\boldsymbol{\rho}^\theta$.[40] In the case of an exact optimization protocol, this should be constructed from the conditional switching rate function $(Q^\theta_{ss'})_{(i,j)\in\mathcal{S}^2}$. The mean dynamic $\mathbf{v}^\theta$: $\mathbb{R}^S \times \Delta^S \to \Delta^S$ is defined by tagging θ to (4) as

$$\dot{x}_s(\theta) = v^\theta_s(\boldsymbol{\pi}(\theta), \mathbf{x}(\theta)) := \sum_{s'\in\mathcal{S}} x_{s'}(\theta)\rho^\theta_{s's}(\boldsymbol{\pi}(\theta)) - x_s(\theta) \sum_{s'\in\mathcal{S}} \rho^\theta_{ss'}(\boldsymbol{\pi}(\theta)) \qquad \text{for each } s \in \mathcal{S}.$$

Then, heterogeneous dynamic $\mathbf{v}^{\mathbf{F}}$ is defined in the same fashion as $\dot{\mathbf{x}}(\theta) = \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) := \mathbf{v}^\theta(\mathbf{F}[\mathbf{X}](\theta), \mathbf{x}(\theta))$. Again, these notations are explicitly shown in the proofs in the appendix. Theorems 3 and 5 hold as long as the assumptions in each theorem are satisfied with $\mathbf{v}^\theta$ of (almost) every type $\theta \in \Theta$.

The existence of a unique solution trajectory (Theorem 1) is also guaranteed, though we should clarify what modifications of the assumptions are needed to include heterogeneous revision protocols. For this, let Θ_C be the set of types that adopt any of L-continuous revision protocols and Θ_E be the set of those who use exact optimization protocols with any conditional switching rate functions. Then, the assumptions for Theorem 1 should read as follows.

Definition 1 There should be a *common* Lipschitz constant $\bar{L}_\rho$ of the switching rate function over almost all the types in Θ_C: $|\rho^\theta_{ss'}(\boldsymbol{\pi}) - \rho^\theta_{ss'}(\boldsymbol{\pi}')| \le \bar{L}_\rho|\boldsymbol{\pi} - \boldsymbol{\pi}'|$ for any $s, s' \in \mathcal{S}, \boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathbb{R}^S$ and μ-almost all $\theta \in \Theta_C$.

Definition 2 There should be a *common* Lipschitz constants $\bar{L}_Q$ of the conditional switching rate functions $Q^\theta_{ss'}$ over almost all the types in Θ_E: $|Q^\theta_{ss'}(\boldsymbol{\pi}) - Q^\theta_{ss'}(\boldsymbol{\pi}')| \le \bar{L}_Q|\boldsymbol{\pi} - \boldsymbol{\pi}'|$ for any $s, s' \in \mathcal{S}, \boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathbb{R}^S$ and μ-almost all $\theta \in \Theta_E$.

Assumption 1 Since this is about payoff function $\mathbf{F}$, this needs no modification.

Assumption 2 There should be a *common* upper bound $\bar{\rho}$ on the switching rate functions $\rho^\theta_{ss'}$ over almost all the types: $\rho^\theta_{ss'}(\mathbf{F}[\mathbf{X}](\theta)) \le \bar{\rho}$ for any strategy distribution $\mathbf{X} \in \mathcal{X}$, any $s, s' \in \mathcal{S}$ and, μ-almost all $\theta \in \Theta$.

[39] About unobservable heterogeneity in aspiration levels in imitative dynamics (i.e., $\bar{\pi}$ or $\underline{\pi}$ in Example 11), Sawa and Zusai (2014) verify that, although the dynamic becomes more complicated and basic properties such as PC do not hold, long-run outcomes are robust to the unobservable heterogeneity. For this, they verify that the difference in the aggregate strategy between under the heterogeneous dynamic and under the homogeneous dynamic vanishes in the long run in any game, whether the dynamic converges to equilibrium or not.

[40] As noted in footnote 4, all our theorems hold for any Polish (complete, separable, and metrizable) space as the type space Θ. A revision protocol $\boldsymbol{\rho}^\theta$ is mathematically a function ρ^θ from payoff vector $\boldsymbol{\pi} \in \mathbb{R}^S$ to switching rates between two strategies $\boldsymbol{\rho}^\theta(\boldsymbol{\pi}) \in \mathbb{R}^{S\times S}_+$. Then, Θ is a space of "admissible" revision protocols, where admissibility is defined so as to make Θ Polish; e.g. continuity of $\boldsymbol{\rho}^\theta$ with respect to $\boldsymbol{\pi}$.

Assumption 3 This is an assumption about $\mathbf{F}$ and μ, not about $\boldsymbol{\rho}^\theta$ or $\mathbf{v}^\theta$. It is needed as long as $\mu(\Theta_E) > 0$; otherwise, it is not needed.

Of course, this extension to heterogeneous revision protocols cover observational dynamics. While we assume within-type observability for Nash stationarity and stability, it is not needed for the existence of a unique solution trajectory. Thus, our existence theorem would provide the most fundamental starting point to study the effect of both observable and unobservable heterogeneity of revision protocols on population dynamics and equilibrium stability.

Example 12. Under the modificaction of Definitions 1 and 2 and assumptions 1 to 3 as above, Theorem 1 holds when different types of agents follow different revision protocols. ∎

6 Concluding remarks

In this paper, we extend evolutionary dynamics to allow for (possibly) continuously many types under persistent heterogeneity in payoff functions and revision protocols. With a rigorous formulation of a heterogeneous evolutionary dynamic as a differential equation over the space of probability measures, we clarify the regularity conditions on the revision protocol, the game and the type distribution to guarantee the existence of a unique solution path from an arbitrary initial state. We extend equilibrium stationarity in general and equilibrium stability in potential games from the homogeneous setting to the heterogeneous setting.[41] This study establishes the foundation to study evolution in heterogeneous populations and opens up a wide field of applications, including spatial evolution and Bayesian games.

Our result on extension of equilibrium stability in potential games suggests that any admissible dynamics share global stability of the equilibrium *set* and also local stability of each local potential-maximizing equilibrium. In contrast, different admissible dynamics may yield different basins of attraction to each locally stable equilibrium and thus they may converge to different locally stable equilibria when starting from the same initial state. Especially, in anonymous games, the preceding studies (Ely and Sandholm, 2005; Blonski, 1999) assume *aggregability* in the sense that the dynamic of aggregate state is completely predictable from its current state, independently of the underlying strategy distribution over different types. However, Zusai (2017) argues that evolutionary dynamics are generically not aggregable, except the standard and smoothed BRDs, even in anonymous games. On the positive side, Zusai (2017) proposes to use nonaggregability to select equilibria by requiring robustness of stability to any distortion in the underlying strategy distribution under nonaggregable dynamics.

In an application to dynamic implementation of the social optimum, the dependency of the aggregate transition on the underlying strategy distribution suggests that a bang-bang control results in excessive instability generally in the heterogeneous setting, though it achieves the fastest convergence in the homogeneous setting. Yet, the dynamic Pigouvian pricing, proposed by Sandholm (2002, 2005), still guarantees convergence to the social optimum, while not

[41] Zusai (2020) provides a universal (and economically intuitive) proof for equilibrium stability of a stable game under a wide range of "(cost-benefit) rationalizable" dynamics (see the paper for its definition) by discovering a universal formula of a Lyapunov function; for a stable game, we need to create a Lyapunov function for each dynamic, as listed in Hofbauer and Sandholm (2009). The universal proof suggests that the stability holds robustly under heterogeneous populations, though Zusai (2020) restricts attention to finitely many populations.

requiring any ex-ante information about the underlying dynamic or type distribution. Nevertheless, there might be a better pricing scheme that lies between the bang-bang control and the dynamic Pigovian pricing and achieves faster convergence than the Pigovian pricing without requiring too much information. Actually, nonaggregability also suggests that the direction of the transition in the aggregate strategy distribution is related with the underlying strategy distribution. If we can find a way to extract the information of the strategy distribution from the transition of the aggregate state, it could be used to improve the pricing scheme.[42]

Acknowledgment

I greatly acknowledge Olena Berchuk, Man Wah Cheung, Dimitrios Diamantaras, Ian Dobson, Shota Fujishima, Hon Ho Kwok, Daisuke Oyama, Moritz Ritter, Bill Sandholm, Ryoji Sawa, Ricardo Serrano-Padial, Marciano Siniscalchi, Marek Weretka, Noah Williams, Jiabin Wu and Dao Zhi-Zeng for their comments.

A Appendix to Section 2

A.1 Measure-theoretic definition of strategy distribution

This subsection provides a mathematically rigorous definition of a strategy distribution based on measure theory. Let $\Omega := [0,1] \subset \mathbb{R}$ be the set (population) of agents. We define a (probability) measure $\mu_\Omega : \mathcal{B}_\Omega \to [0,1]$ as the Lebesgue measure so $\mu_\Omega(\Omega) = 1$. Denote by $\mathcal{B}_\Omega$ the Lebesgue σ-field over Ω. Let $\mathfrak{s}(\omega) \in \mathcal{S}$ denote the strategy taken by agent ω. We restrict strategy profile $\mathfrak{s} : \Omega \to \mathcal{S}$ to a $\mathcal{B}_\Omega$-measurable function.

Let $\theta(\omega) \in \Theta$ be the type of agent $\omega \in \Omega$; assume that type space Θ is a Polish (separable completely metrizable) space, more general than in the main text. Denote by $\mathcal{B}$ be the Borel σ-field on this space. Agents' type profile $\theta : \Omega \to \mathbb{R}^T$ is assumed to be measurable with respect to $\mathcal{B}_\Omega$. Then, it induces probability measure $\mu : \mathcal{B} \to [0,1]$ by $\mu(B) := \mu_\Omega(\{\omega \in \Omega : \theta(\omega) \in B\})$ for each $B \in \mathcal{B}$.

Combination of strategy profile $\mathfrak{s} : \Omega \to \mathcal{S}$ and type profile $\theta : \Omega \to \Theta$ generates a finite measure $X_s : \mathcal{B} \to \mathbb{R}_+$ for each $s \in \mathcal{S}$ from $\mathbb{P}_\Omega$:

$$X_s(B) := \mu_\Omega(\{\omega \in \Omega \;:\; \mathfrak{s}(\omega) = s \text{ and } \theta(\omega) \in B\}) \qquad \text{for each } B \in \mathcal{B}.$$

$X_s(B)$ represents the mass of strategy-s players whose types belong to set B. The **strategy distribution X** is a collection of these measures X_s, i.e., $\mathbf{X} = (X_s)_{s\in\mathcal{S}}$. We can see this vector measure as a joint probability measure over the product space $\mathcal{S} \times \Theta$. [43] The space of strategy distributions $\mathcal{X}$ is thus the set of probability measures over $\mathcal{S} \times \Theta$ such that the marginal

[42]**?** considers a heterogeneous congestion game where payoff heterogeneity is not additively separable and the social planner does not exactly know its distribution, and proposes a modified Pigouvian pricing that is combined with estimation of the distribution.

[43]Abusing notation, we could say that **X** defines a measure of a Borel set $B_{\mathcal{S}\Theta}$ on the product space $\mathcal{S} \times \Theta$ by

$$\mathbf{X}(B_{\mathcal{S}\Theta}) := \sum_{s\in\mathcal{S}} X_s(\{\theta \in \Theta : (a,\theta) \in B_{\mathcal{S}\Theta}\}) = \mu_\Omega(\{\omega \in \Omega : (\mathfrak{s}(\omega),\theta(\omega)) \in B_{\mathcal{S}\Theta}\}).$$

distribution of types coincides with μ, i.e., $\sum_{s\in\mathcal{S}} X_s(B) = \mu(B)$ for each $B \in \mathcal{B}$. Let $\mathcal{B}_{\mathcal{S}\Theta}$ be the Borel σ-field on the product space $\mathcal{S} \times \Theta$.

Since $\mathbf{X}$ must satisfy $X_s(B) \leq \mu(B)$ for each $s \in \mathcal{S}$, X_s is dominated by μ in the sense that

$$\mu(B) = 0 \quad \Longrightarrow \quad X_s(B) = 0 \qquad \text{for each } B \in \mathcal{B}. \tag{A.1}$$

Denote by $X_s \ll \mu$ this dominance relation, i.e., absolute continuity of X_s with respect to μ. It follows by Radon-Nikodym theorem that there exists a $\mathcal{B}$-measurable nonnegative function $x_s : \Theta \to \mathbb{R}_+$ such that $X_s(B) = \int_B x_s(\theta)\mu(d\theta)$ for any $B \in \mathcal{B}$. x_s is the density function of measure X_s. The density is determined uniquely in the sense that, if another measurable function x'_s satisfies $X_s(B) = \int_B x'_s(\theta)\mu(d\theta)$ for all $B \in \mathcal{B}$, then $x'_s(\theta) = x_s(\theta)$ for μ-almost all $\theta \in \Theta$.

$\mathbf{X}$ is dominated by μ in the sense that $X_s \ll \mu$ for all $s \in \mathcal{S}$; we abuse notation to denote this domination by $\mathbf{X} \ll \mu$. The dominance of strategy distribution $\mathbf{X}$ by the type distribution μ is peculiar to heterogeneous dynamics, making a difference in the proof of Lipschitz continuity of the dynamic from the one for continuous strategy dynamics. See Remark 2 in Section 3.

The collection of Radon-Nikodym densities $\mathbf{x} = (x_s)_{s\in\mathcal{S}} : \Theta \to \mathbb{R}^S_+$ is the **strategy density function** corresponding to $\mathbf{X}$. From the fact that $\sum_{s\in\mathcal{S}} X_s(B) = \mu(B)$ and $X_s(B) \geq 0$ for any $B \in \mathcal{B}$ and $s \in \mathcal{S}$, we can confirm that $\mathbf{x}(\theta)$ is a probability vector, i.e., $\mathbf{x}(\theta) \in \Delta^S$, for μ-almost all types $\theta \in \Theta$. Thus, the space of strategy density functions $\mathcal{F}_{\mathcal{X}}$ is a set of $\mathcal{B}$-measurable vector functions from Θ to Δ^S.

A.2 Topology of the space of strategy distributions

Choice of a topology is a sensitive issue when we argue dynamics of a measure over a continuous space. We follow the convention in the literature on evolutionary dynamics over a continuous strategy space, such as in Cheung (2014). That is, we use the *strong topology* to prove the existence of a unique solution path with Picard-Lindelöf theorem (Theorem 2) and the *weak topology* to obtain stability of equilibrium strategy distribution with Lyapunov stability theorem (Theorem 6). See Cheung (2014, Section 4) for a detailed explanation on the strong and weak topology in evolutionary dynamics on a continuous space.

Below we define these two topologies on the space of finite signed measures $\mathcal{M}$ over $\mathcal{S} \times \Theta$. $\mathcal{X}$ and its tangent space $T\mathcal{X}$ are subsets of this space. The **strong topology** is metrized by the variational norm $\|\cdot\|_\infty$ defined as $\|\mathbf{X}\|_\infty = \sup_{\mathbf{g}} \left|\int_{\theta\in\Theta} \mathbf{g}(\theta) \cdot \mathbf{X}(d\theta)\right|$ where the sup is taken over the set of measurable functions $\mathbf{g} = (g_s)_{s\in\mathcal{S}}$ on $(\mathcal{S} \times \Theta, \mathcal{B}_{\mathcal{S}\Theta})$ that are bounded by 1, i.e., $\sup_{(a,\theta)\in\mathcal{S}\times\Theta} |g_s(\theta)| \leq 1$. According to Oechssler and Riedel (2001, Theorem 5), the variational norm on $\mathcal{X}$ is equivalent to the L^1-norm on $\mathcal{F}_{\mathcal{X}}$ in the sense that, for any strategy distribution $\mathbf{X} = \int \mathbf{x} d\mu \in \mathcal{X}$,[44]

$$\|\mathbf{X}\|_\infty = \int_\Theta \sum_{s\in\mathcal{S}} |x_s(\theta)| d\mu(\theta).$$

Under the **weak topology** on the set of measures over space $\mathcal{S} \times \Theta$, a mapping $\mathcal{M} \to \mathbb{R}$ such as $\mu \mapsto \int_{\mathcal{S}} \mathbf{g} \cdot d\mathbf{X}$ is continuous for any bounded and continuous function $\mathbf{g} : \Theta \to \mathbb{R}^S$. Note that convergence in strong topology implies that in weak topology. The product space

[44] This formula extends to any absolutely continuous finite signed measures.

$\mathcal{S} \times \Theta$ is separable with metric $d_{\mathcal{S}\Theta} : (\mathcal{S} \times \Theta)^2 \to \mathbb{R}_+$ given by[45]

$$d_{\mathcal{S}\Theta}((s,\theta),(s',\theta')) := \mathbb{1}\{s \neq s'\} + d(\theta,\theta'),$$

where d is a metric under which Θ is complete. Then, the weak topology is metrized by Prokhorov metric $d^{\mathrm{Pr}} : \mathcal{M} \times \mathcal{M} \to \mathbb{R}_+$ such that[46]

$$\begin{aligned} d^{\mathrm{Pr}}(\mathbf{X},\mathbf{X}') := \inf\{\varepsilon > 0 \ : \ & \mathbf{X}(B_{\mathcal{S}\Theta}) \le \mathbf{X}'(B^{\varepsilon}_{\mathcal{S}\Theta}) + \varepsilon \\ & \text{and } \mathbf{X}'(B_{\mathcal{S}\Theta}) \le \mathbf{X}(B^{\varepsilon}_{\mathcal{S}\Theta}) + \varepsilon \quad \text{for all } B_{\mathcal{S}\Theta} \in \mathcal{B}_{\mathcal{S}\Theta}\}, \end{aligned}$$

where $B^{\varepsilon}_{\mathcal{S}\Theta}$ is defined from $B_{\mathcal{S}\Theta}$ as $B^{\varepsilon}_{\mathcal{S}\Theta} := \{(s,\theta) \in \mathcal{S} \times \Theta \ : \ d_{\mathcal{S}\Theta}((s,\theta),(s',\theta')) < \varepsilon \text{ with some } (s',\theta') \in B_{\mathcal{S}\Theta}\}$.[47] Under the weak topology, the space of probability measures, i.e., the space of strategy distributions becomes compact. Then, we can apply the Lyapunov stability theorem, as in Cheung (2014, Thm. 6). See Theorem 6 in Appendix C.3.

B Appendix to Section 3

Proof of Theorem 1. Below we consider two strategy distributions $\mathbf{X} = \int \mathbf{x} d\mu, \mathbf{X}' = \int \mathbf{x}' d\mu \in \mathcal{X}$. Let $\rho^{\mathbf{F}}_{s's}(\theta) := \rho^{\theta}_{s's}(\mathbf{F}[\mathbf{X}](\theta), \mathbf{x}(\theta))$ and ${\rho'}^{\mathbf{F}}_{s's}(\theta) := \rho^{\theta}_{s's}(\mathbf{F}[\mathbf{X}'](\theta), \mathbf{x}'(\theta))$. Denote $\Delta v^{\mathbf{F}}_s(\theta) := v^{\mathbf{F}}_s[\mathbf{X}](\theta) - v^{\mathbf{F}}_s[\mathbf{X}'](\theta)$. We divide Θ by the two classes of revision protocols: let Θ_C be the set of the types of agents who follow L-continuous revision protocols and Θ_E be the set of the types who follow exact optimization protocols; we have $\Theta_C \cup \Theta_E = \Theta$ and $\Theta_C \cap \Theta_E = \emptyset$. First we prove Lipschitz continuity of the transition of the strategy distribution in each of these sets of types. Then, we merge them to get Lipschitz continuity of the transition of the entire strategy distribution.

1°: L-continuous revision protocols. Now we focus on Θ_C. Let $\bar{L}_\rho > 0$ be the upper bound on the Lipschitz constants of functions $\rho^{\theta}_{s's}$ over all pairs of strategies $s, s' \in \mathcal{S}$ and all types $\theta \in \Theta_C$. The Lipschitz continuity of $\rho^{\theta}_{s's}$ (Definition 1) and $\mathbf{F}$ (Assumption 1) implies

$$\begin{aligned} |\rho^{\mathbf{F}}_{s's}(\theta) - {\rho'}^{\mathbf{F}}_{s's}(\theta)| &\le \bar{L}_\rho |(\mathbf{F}[\mathbf{X}](\theta), \mathbf{x}(\theta)) - (\mathbf{F}[\mathbf{X}'](\theta), \mathbf{x}'(\theta))| \\ &= \bar{L}_\rho \left\{ |\mathbf{F}[\mathbf{X}](\theta) - \mathbf{F}[\mathbf{X}'](\theta)| + |\mathbf{x}(\theta) - \mathbf{x}'(\theta)| \right\} \\ &\le \bar{L}_\rho \left(L_{\mathbf{F}}(\theta) \|\mathbf{X} - \mathbf{X}'\|_\infty + |\mathbf{x}(\theta) - \mathbf{x}'(\theta)| \right). \end{aligned} \tag{B.2}$$

From the definition of $v^{\mathbf{F}+}_s$, we have

$$\begin{aligned} |v^{\mathbf{F}+}_s[\mathbf{X}](\theta) - v^{\mathbf{F}+}_s[\mathbf{X}'](\theta)| &\le \sum_{s' \in \mathcal{S}} |\rho^{\mathbf{F}}_{s's}(\theta) x_{s'}(\theta) - {\rho'}^{\mathbf{F}}_{s's}(\theta) x'_{s'}(\theta)| \\ \le \sum_{s' \in \mathcal{S}} &\left\{ |\rho^{\mathbf{F}}_{s's}(\theta) - {\rho'}^{\mathbf{F}}_{s's}(\theta)| \, |x_{s'}(\theta)| + |{\rho'}^{\mathbf{F}}_{s's}(\theta)| \cdot |x_{s'}(\theta) - x'_{s'}(\theta)| \right\} \\ \le \sum_{s' \in \mathcal{S}} &\left[\bar{L}_\rho \left(L_{\mathbf{F}}(\theta) \|\mathbf{X} - \mathbf{X}'\|_\infty| + |\mathbf{x}(\theta) - \mathbf{x}'(\theta)| \right) + \bar{\rho} |\mathbf{x}(\theta) - \mathbf{x}'(\theta)| \right] \end{aligned}$$

[45]The metric $d_{\mathcal{S}\Theta}$ is a product metric constructed from the discrete norm on $\mathcal{S}$ and metric d on Θ. Notice $S < \infty$ and Θ is separable; so the product metric $d_{\mathcal{S}\Theta}$ makes $\mathcal{S} \times \Theta$ separable. Here $\mathbb{1}\{s \neq s'\}$ is an indicator function and takes 1 if $s \neq s'$ and 0 otherwise.

[46]If there is no payoff heterogeneity, i.e., $\Theta = \{\theta_0\}$, then strategy distribution $\mathbf{X}$ can be simply represented by an S-dimensional vector $(\bar{x}_s)_{s \in \mathcal{S}} \in \mathbb{R}^S$ such that $\bar{x}_s = X_s(\{\theta_0\})$. Then, $d^{\mathrm{Pr}}(\mathbf{X}, \mathbf{X}') = \varepsilon$ is equivalent to $\sup_{s \in \mathcal{S}} |\bar{x}_s - \bar{x}'_s| = \varepsilon$. So, the metric d^{Pr} reduces to the sup norm on $\mathbb{R}^S$.

[47]If $\varepsilon < 1$, the condition for $(s,\theta) \in B^{\varepsilon}_{\mathcal{S}\Theta}$ is equivalent to the existence of $\theta' \in \Theta$ such that $d(\theta,\theta') < \varepsilon$ and $(s,\theta') \in B_{\mathcal{S}\Theta}$. Thus, provided that $B_{\mathcal{S}\Theta} = \cup_{s \in \mathcal{S}} \{s\} \times B_{s\Theta}$ with each $B_{s\Theta} \in \mathcal{B}$, we have $B^{\varepsilon}_{\mathcal{S}\Theta} = \cup_{s \in \mathcal{S}} \{s\} \times B^{\varepsilon}_{s\Theta}$ with each $B^{\varepsilon}_{s\Theta} = \{\theta \in \Theta : d(\theta,\theta') < \varepsilon \text{ for some } \theta \in B_{s\Theta}\}$ for each $s \in \mathcal{S}$.

$$\leq S\left\{\bar{L}_\rho L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_\rho+\bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right\} \tag{B.3}$$

Here the third inequality comes from (B.2), Assumption 2 and $|x_{s'}(\cdot)| \leq 1$. Similarly, we get

$$|v_s^{\mathbf{F}-}[\mathbf{X}](\theta) - v_s^{\mathbf{F}-}[\mathbf{X}'](\theta)| \leq S\left\{\bar{L}_\rho L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_\rho+\bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right\}.$$

Therefore, we have

$$\int_{\Theta_C}\sum_{s\in\mathcal{S}}|\Delta v_s^{\mathbf{F}}(\theta)|\mu(d\theta) \leq \int_{\Theta_C}\sum_{s\in\mathcal{S}}\left(|v_s^{\mathbf{F}+}[\mathbf{X}](\theta)-v_s^{\mathbf{F}+}[\mathbf{X}'](\theta)| + |v_s^{\mathbf{F}-}[\mathbf{X}](\theta)-v_s^{\mathbf{F}-}[\mathbf{X}'](\theta)|\right)\mu(d\theta)$$

$$\leq \int_{\Theta_C}\left[\sum_{s\in\mathcal{S}}2A\left\{\bar{L}_\rho L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_\rho+\bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right\}\right]\mu(d\theta)$$

$$=2S^2\cdot\bar{L}_\rho\|\mathbf{X}-\mathbf{X}'\|_\infty\int_{\Theta_C}L_{\mathbf{F}}(\theta)\mu(d\theta) + 2S^2(\bar{L}_\rho+\bar{\rho})\int_{\Theta_C}|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\mu(d\theta)$$

$$\leq 2S^2(\bar{L}_\rho\bar{L}_{\mathbf{F}}+\bar{L}_\rho+\bar{\rho})\|\mathbf{X}-\mathbf{X}'\|_\infty. \tag{B.4}$$

The last inequality comes from $\int_{\Theta_C}L_{\mathbf{F}}(\theta)\mu(d\theta) \leq \mathbb{E}L_{\mathbf{F}} = \bar{L}_{\mathbf{F}}$ by $L_{\mathbf{F}}(\theta) \geq 0$.

2°: exact optimization protocols. Now we focus on Θ_E. In an exact optimization protocol, the dynamic reduces as the following: if strategy b is the unique maximizer of $F_s[\mathbf{X}](\theta)$ among all strategies $s \in \mathcal{S}$, i.e., the unique best response to $\mathbf{X}$ for type θ, then

$$v_b(\theta)[\mathbf{X}] = \sum_{s'\in\mathcal{S}\setminus\{b\}} Q_{s'b}(\mathbf{F}[\mathbf{X}](\theta))x_{s'}(\theta), \qquad v_s(\theta)[\mathbf{X}] = -Q_{sb}(\mathbf{F}[\mathbf{X}](\theta))x_s(\theta) \quad \text{for any } s \in \mathcal{S}\setminus\{b\}.$$

Since Assumption 3 implies that the best response is unique for almost every type, this determines the heterogeneous dynamic without ambiguity.

Let $\bar{L}_Q > 0$ be an upper bound on Lipschitz constants of functions $Q^\theta_{ss'}$ over all pairs of strategies $s, s' \in \mathcal{S}$ and all the types $\theta \in \Theta_E$. Let N be the set of types who have multiple best responses to either $\mathbf{X}$ or $\mathbf{X}'$ or both. Assumption 3 implies $\mu(N) = 0$. Define partitions of $\Theta \setminus N$ by

$$\cap\beta_b := \Theta^{\mathbf{F}}_{b=\text{uniqBR}}(\mathbf{X}) \cap \Theta^{\mathbf{F}}_{b=\text{uniqBR}}(\mathbf{X}') \cap \Theta_E, \qquad \Delta\beta_{bb'} := \Theta^{\mathbf{F}}_{b=\text{uniqBR}}(\mathbf{X}) \cap \Theta^{\mathbf{F}}_{b'=\text{uniqBR}}(\mathbf{X}') \cap \Theta_E$$

for each $b, b' \in \mathcal{S}$ with $b' \neq b$. Let $\cap\beta := \bigcup_{b\in\mathcal{S}}\cap\beta_b$ and $\Delta\beta := \bigcup_{b\in\mathcal{S}}\bigcup_{b'\in\mathcal{S}\setminus\{b\}}\Delta\beta_{bb'}$.

Denote $Q^{\mathbf{F}}_{s's}(\theta) := Q^\theta_{s's}(\mathbf{F}[\mathbf{X}](\theta))$ and $Q'^{\mathbf{F}}_{s's}(\theta) := Q^\theta_{s's}(\mathbf{F}[\mathbf{X}'](\theta))$. Similarly to (B.2), the Lipschitz continuity of $Q^\theta_{s's}$ (Definition 2) and $\mathbf{F}$ (Assumption 1) implies

$$|Q^{\mathbf{F}}_{s's}(\theta) - Q'^{\mathbf{F}}_{s's}(\theta)| \leq \bar{L}_Q\left(L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + |\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right) \tag{B.5}$$

for all $s, s' \in \mathcal{S}$, and μ-almost all $\theta \in \Theta_E$. Note that, if $\theta \in \Theta^{\mathbf{F}}_{b=\text{uniqBR}}(\mathbf{X})$, then Assumption 2 assures the existence of an upper bound $\bar{\rho}$ on $Q^\theta_{\cdot b}(\theta)$ such as $Q^{\mathbf{F}}_{sb}(\theta) = \rho^\theta_{sb}(\mathbf{F}[\mathbf{X}](\theta)) \leq \bar{\rho}$ for any $s \in \mathcal{S}$ and $\theta \in \Theta_E$.

i) Consider $\cap\beta_b$ for an arbitrary $b \in \mathcal{S}$. Fix $\theta \in \cap\beta_b$: strategy b is the best response strategy for this type θ both in the state $\mathbf{X}$ and the state $\mathbf{X}'$. Then, similarly to (B.3), Assumption 2 and (B.5) imply

$$|\Delta v_b(\theta)| \leq \sum_{s'\in\mathcal{S}\setminus\{b\}} |Q^{\mathbf{F}}_{s'b}(\theta)x_{s'}(\theta) - Q'^{\mathbf{F}}_{s'b}(\theta)x'_{s'}(\theta)|$$

$$\leq (S-1)\left\{\bar{L}_Q L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_Q+\bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right\}.$$

For strategy $s \neq b$,

$$\Delta v_s(\theta) = \left(-Q^{\mathbf{F}}_{sb}(\theta)x_s(\theta)\right) - \left(-Q'^{\mathbf{F}}_{sb}(\theta)x_s(\theta)\right) = -\{Q^{\mathbf{F}}_{sb}(\theta) - Q'^{\mathbf{F}}_{sb}(\theta)\}x_s(\theta) - Q'^{\mathbf{F}}_{sb}(\theta)\{x_s(\theta) - x'_s(\theta)\}.$$

(B.5) implies

$$\begin{aligned}|\Delta v_s(\theta)| &\leq \bar{L}_Q\left(L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + |\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right)|x_s(\theta)| + Q'^{\mathbf{F}}_{sb}(\theta)|\mathbf{x}(\theta)-\mathbf{x}'(\theta)| \\ &\leq \bar{L}_Q L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_Q + \bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|.\end{aligned}$$

The second inequality comes from boundeness of Q^θ.

Therefore, we have

$$\sum_{s\in\mathcal{S}} |\Delta v_s(\theta)| \leq 2(S-1)\left\{\bar{L}_Q L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_Q+\bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right\}$$

and thus

$$\begin{aligned}&\int_{\cap\beta}\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)|\mu(d\theta) \leq 2(S-1)\int_{\cap\beta}\left\{\bar{L}_Q L_{\mathbf{F}}(\theta)\|\mathbf{X}-\mathbf{X}'\|_\infty + (\bar{L}_Q+\bar{\rho})|\mathbf{x}(\theta)-\mathbf{x}'(\theta)|\right\}\mu(d\theta) \\ &\leq 2(S-1)(\bar{L}_Q\bar{L}_{\mathbf{F}} + \bar{L}_Q + \bar{\rho})\|\mathbf{X}-\mathbf{X}'\|_\infty. \qquad \text{(B.6)}\end{aligned}$$

The second inequality comes from $\mu(\cap\beta) \leq \mu(\Theta) = 1$, $\int_{\cap\beta} L_{\mathbf{F}} d\mu \leq \int L_{\mathbf{F}} d\mu = \bar{L}_{\mathbf{F}}$, and $\int_{\cap\beta}|\mathbf{x} - \mathbf{x}'|d\mu \leq \|\mathbf{X}-\mathbf{X}'\|_\infty$.

ii) Consider $\Delta\beta_{bb'}$ for two arbitrary distinct strategies $b, b' \in \mathcal{S}$ with $b \neq b'$. Fix $\theta \in \Delta\beta_{bc}$: strategy b is the best response strategy for this type θ in the state $\mathbf{X}$ and c is the optimal in the state $\mathbf{X}'$. Then,

$$0 \leq \sum_{s'\in\mathcal{S}\setminus\{b\}} Q^{\mathbf{F}}_{s'b}(\theta)x_{s'}(\theta) - \left(-Q'^{\mathbf{F}}_{bb'}(\theta)x'_b(\theta)\right) = \Delta v_b(\theta) \leq \sum_{s'\in\mathcal{S}\setminus\{b\}} \bar{\rho}\cdot + \bar{\rho}\cdot = S\bar{\rho}.$$

Similarly, we have $0 \geq \Delta v_{b'}(\theta) \geq -S\bar{\rho}$. For $s \neq b, b'$,

$$\Delta v_s(\theta) = \left(-Q^{\mathbf{F}}_{sb}(\theta)x_s(\theta)\right) - \left(-Q'^{\mathbf{F}}_{sb}(\theta)x'_s(\theta)\right).$$

Since $Q^\theta(\cdot) \in [0, \bar{\rho}]$ and $x_\cdot(\cdot) \in [0, 1]$, we have

$$|\Delta v_s(\theta)| \leq |Q^{\mathbf{F}}_{sb}(\theta)x_s(\theta)| + |Q'^{\mathbf{F}}_{sb}(\theta)x'_s(\theta)| \leq 2\bar{\rho}.$$

Therefore,

$$\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)| \leq 2S\bar{\rho} + (S-2)\cdot 2\bar{\rho} = 4(S-1)\bar{\rho}.$$

By Assumption 3, we have

$$\int_{\Delta\beta_{bb'}}\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)|\mu(d\theta) \leq 4(S-1)\bar{\rho}\mu(\Delta\beta_{bb'}) \leq 4(S-1)\bar{\rho}L_{\mathrm{BR}}\|\mathbf{X}-\mathbf{X}'\|_\infty$$

and thus

$$\int_{\Delta\beta}\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)|\mu(d\theta) = \sum_{b\in\mathcal{S}}\sum_{b'\in\mathcal{S}\setminus\{b\}}|\Delta v_s(\theta)|\mu(d\theta) \leq 4S(S-1)^2\bar{\rho}L_{\mathrm{BR}}\|\mathbf{X}-\mathbf{X}'\|_\infty.$$

3°: Merge them. Since Θ is a union of $\Theta_C, \cap\beta, \Delta\beta$ and N and $\mu(N) = 0$, we have

$$\|\mathbf{V}[\mathbf{X}] - \mathbf{V}[\mathbf{X}']\|_\infty = \int_{\Theta_E}\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)|\mu(d\theta) + \int_{\cap\beta}\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)|\mu(d\theta) + \int_{\Delta\beta}\sum_{s\in\mathcal{S}}|\Delta v_s(\theta)|\mu(d\theta).$$

(B.4), (B.6) and (B.7) imply (6), namely $\|\mathbf{V}[\mathbf{X}] - \mathbf{V}[\mathbf{X}']\|_\infty \le L_V \|\mathbf{X} - \mathbf{X}'\|_\infty$ with

$$L_V := 2S^2(\bar{L}_\rho \bar{L}_\mathbf{F} + \bar{L}_\rho + \bar{\rho}) + 2(S-1)\{\bar{L}_Q \bar{L}_\mathbf{F} + (\bar{L}_Q + \bar{\rho})\} + 4S(S-1)^2 \bar{\rho} L_{\mathrm{BR}}. \qquad \square$$

Proof of Corollary 1. We leave only the boundedness of the dynamic; it comes from Assumption 2. Since it implies $|v_s^\mathbf{F}[\mathbf{X}](\theta)| \le 2(S-1)\bar{\rho}$, we obtain $\|\mathbf{V}^\mathbf{F}[\mathbf{X}]\|_\infty = \sum_{s\in\mathcal{S}} \int_\Theta |v_s^\mathbf{F}[\mathbf{X}]| d\mu \le 2S(S-1)\bar{\rho}$ for all $\mathbf{X} \in \mathcal{X}$. Then, Theorem 2 implies the existence of a unique solution path of the dynamic on $\mathcal{M}$. Notice that $\mathcal{X}$ is forward invariant under $\mathbf{V}^\mathbf{F}$. Therefore, if the initial state $\mathbf{X}_0$ lies in $\mathcal{X} \subset \mathcal{M}$, then the unique solution that passes $\mathbf{X}_0$ at time 0 should remain in $\mathcal{X}$. $\square$

C Appendix to Section 4

C.1 Proof of Theorem 3

Proof. First of all, strategy distribution $\mathbf{X} = \int \mathbf{x} d\mu$ being an equilibrium (1) means $\mathbf{x}(\theta) \in \Delta\mathcal{S}^\mathbf{F}_{\mathrm{BR}}[\mathbf{X}](\theta)$ for μ-almost all types θ. Then, for such θ, $\mathbf{x}(\theta) \in \Delta\mathcal{S}^\mathbf{F}_{\mathrm{BR}}[\mathbf{X}](\theta)$ is equivalent to $\mathbf{v}^\mathbf{F}[\mathbf{x}](\theta) = \mathbf{0}$ by (7). It holds for μ-almost all types θ, which means the stationarity of strategy density function $\mathbf{x}$. This is equivalent to stationarity of strategy distribution $\mathbf{X}$, i.e., $\mathbf{V}^\mathbf{F}[\mathbf{X}] = \mathbf{O}$. $\square$

C.2 Proof of Theorem 4

First of all, notice that an equilibrium strategy distribution is a fixed point of the "distributional strategy" best response correspondence $B : \mathcal{X} \rightrightarrows \mathcal{X}$ defined as

$$B[\mathbf{X}] := \operatorname*{argmax}_{\mathbf{Y}\in\mathcal{X}} \int_\Theta \mathbf{F}[\mathbf{X}](\theta) \cdot \mathbf{Y}(d\theta) \qquad \text{for each } \mathbf{X} \in \mathcal{X}.$$

Below we prove that the assumptions in Theorem 4 assures that the maximized function $\int_\Theta \mathbf{F}[\mathbf{X}] \cdot d\mathbf{Y}$ is continuous in $(\mathbf{X}, \mathbf{Y}) \in \mathcal{X}^2$.

Proof of continuity of $\int_\Theta \mathbf{F}[\mathbf{X}] \cdot d\mathbf{Y}$. Fix $\varepsilon > 0$ and $(\mathbf{X}, \mathbf{Y}) \in \mathcal{X}^2$ arbitrarily. By equicontinuity of $\mathbf{F}$, we have some $\delta_{\mathrm{Ct}}[\mathbf{X}] > 0$ such that, whenever $d^{\mathrm{Pr}}(\mathbf{X}', \mathbf{X}) < \delta_{\mathrm{Ct}}[\mathbf{X}]$, we have $|F_s[\mathbf{X}'](\theta) - F_s[\mathbf{X}](\theta)| < 0.5\varepsilon$ for any $s \in \mathcal{S}$ and μ-almost all θ. The latter statement implies, for any $\mathbf{Y}' \in \mathcal{X}$,

$$\left| \int_\Theta \left(\mathbf{F}[\mathbf{X}'](\theta) - \mathbf{F}[\mathbf{X}](\theta) \right) \cdot \mathbf{Y}'(d\theta) \right| \le \int_\Theta \sum_{s\in\mathcal{S}} \left| F_s[\mathbf{X}'](\theta) - F_s[\mathbf{X}](\theta) \right| Y'_s(d\theta)$$
$$< \int_\Theta \sum_{s\in\mathcal{S}} 0.5\varepsilon Y'_s(d\theta) = 0.5\varepsilon \int_\Theta \mathbb{P}(d\theta) = 0.5\varepsilon.$$

Near-boundedness of $\mathbf{F}$ implies that, for any $\varepsilon > 0$, there exists a combination of $\bar{F}[\mathbf{X}] \ge 0$ and $\delta^1_{\mathrm{Bd}}[\mathbf{X}] > 0$ such that for any $\mathbf{Y}, \mathbf{Y}' \in \mathcal{X}$[48]

$$d^{\mathrm{Pr}}(\mathbf{Y}, \mathbf{Y}') < \delta^1_{\mathrm{Bd}}[\mathbf{X}] \quad \Longrightarrow \quad \left| \int_\Theta \sum_{s\in\mathcal{S}} [|F_s[\mathbf{X}](\theta)| - \bar{F}[\mathbf{X}]]_+ \cdot (Y'_s(d\theta) - Y_s(d\theta)) \right| < 0.25\varepsilon.$$

For each $s \in \mathcal{S}$, define set $\Theta^\varepsilon_s \subset \Theta$ and function $g^\varepsilon_s[\mathbf{X}] : \Theta \to \mathbb{R}_+$ by

$$\Theta^\varepsilon_s := \{\theta : |F_s[\mathbf{X}](\theta)| > \bar{F}[\mathbf{X}]\}, \qquad g^\varepsilon_s[\mathbf{X}](\theta) := \mathbb{1}\{\theta \in \Theta^\varepsilon_s\}\bar{F}[\mathbf{X}] + \mathbb{1}\{\theta \notin \Theta^\varepsilon_s\}|F_s[\mathbf{X}](\theta)|.$$

[48] Here $[\cdot]_+$ is an operator such as $[z]_+ = \max\{0, z\}$.

Notice that $g^{\varepsilon}_s[\mathbf{X}](\theta) + [|F_s[\mathbf{X}](\theta)| - \bar{F}[\mathbf{X}]]_+ \equiv |F_s[\mathbf{X}](\theta)|$ for any $\theta \in \Theta$. Since $F_s[\mathbf{X}] : \Theta \to \mathbb{R}$ is measurable, Θ^{ε}_s is a measurable set and $g^{\varepsilon}_s[\mathbf{X}]$ is a measurable function on Θ. This function is bounded by definition and also continuous; so is the vector-valued function $\mathbf{g}^{\varepsilon}[\mathbf{X}] = (g^{\varepsilon}_s[\mathbf{X}])_{s\in\mathcal{S}} : \Theta \to \mathbb{R}^S_+$. Hence, under the weak topology, $\int_\Theta \mathbf{g}^{\varepsilon}[\mathbf{X}](\theta) \cdot \Delta\mathbf{Y}(d\theta)$ is a continuous function of finite signed measure $\Delta\mathbf{Y} \in \mathcal{M}$; thus, there exists $\delta^2_{\text{Bd}}[\mathbf{X}] > 0$ such that

$$d^{\text{Pr}}(\mathbf{Y}', \mathbf{Y}) < \delta^2_{\text{Bd}}[\mathbf{X}] \quad \Longrightarrow \quad \left| \int_\Theta \sum_{s\in\mathcal{S}} g^{\varepsilon}_s[\mathbf{X}](\theta) \left(Y'_s(d\theta) - Y_s(d\theta)\right) \right| < 0.25\varepsilon.$$

Therefore, if $\mathbf{Y}'$ satisfies $d^{\text{Pr}}(\mathbf{Y}', \mathbf{Y}) < \delta_{\text{Bd}}[\mathbf{X}] := \min\{\delta^1_{\text{Bd}}[\mathbf{X}], \delta^2_{\text{Bd}}[\mathbf{X}]\}$, then

$$\begin{aligned}
&\left| \int_\Theta \mathbf{F}[\mathbf{X}](\theta) \cdot \left(\mathbf{Y}'(d\theta) - \mathbf{Y}(d\theta)\right) \right| = \left| \int_\Theta \sum_{s\in\mathcal{S}} F_s[\mathbf{X}](\theta) \left(Y'_s(d\theta) - Y_s(d\theta)\right) \right| \\
\leq &\left| \int_\Theta \sum_{s\in\mathcal{S}} |F_s[\mathbf{X}](\theta)| \left(Y'_s(d\theta) - Y_s(d\theta)\right) \right| = \left| \int_\Theta \sum_{s\in\mathcal{S}} \left(g^{\varepsilon}_s[\mathbf{X}](\theta) + [|F_s[\mathbf{X}](\theta)| - \bar{F}[\mathbf{X}]]_+\right) \left(Y'_s(d\theta) - Y_s(d\theta)\right) \right| \\
\leq &\left| \int_\Theta \sum_{s\in\mathcal{S}} g^{\varepsilon}_s[\mathbf{X}](\theta) \left(Y'_s(d\theta) - Y_s(d\theta)\right) \right| + \left| \int_\Theta \sum_{s\in\mathcal{S}} [|F_s[\mathbf{X}](\theta)| - \bar{F}[\mathbf{X}]]_+ \cdot (Y'_s(d\theta) - Y_s(d\theta)) \right| \\
\leq &0.25\varepsilon + 0.25\varepsilon = 0.5\varepsilon.
\end{aligned}$$

In sum, if a pair of $(\mathbf{X}', \mathbf{Y}') \in \mathcal{X}^2$ satisfies $d^{\text{Pr}}(\mathbf{X}', \mathbf{X}) < \delta_{\text{Ct}}[\mathbf{X}]$ and $d^{\text{Pr}}(\mathbf{Y}', \mathbf{Y}) < \delta_{\text{Bd}}[\mathbf{X}]$

$$\begin{aligned}
&\left| \int_\Theta \mathbf{F}[\mathbf{X}'](\theta) \cdot \mathbf{Y}'(d\theta) - \int_\Theta \mathbf{F}[\mathbf{X}](\theta) \cdot \mathbf{Y}(d\theta) \right| \\
\leq &\left| \int_\Theta \left(\mathbf{F}[\mathbf{X}'](\theta) - \mathbf{F}[\mathbf{X}](\theta)\right) \cdot \mathbf{Y}'(d\theta) \right| + \left| \int_\Theta \mathbf{F}[\mathbf{X}](\theta) \cdot \left(\mathbf{Y}'(d\theta) - \mathbf{Y}(d\theta)\right) \right| \\
< &0.5\varepsilon + 0.5\varepsilon = \varepsilon.
\end{aligned}$$

That is, $\int_\Theta \mathbf{F}[\mathbf{X}] \cdot d\mathbf{Y}$ is continuous with respect to weak topology in $\mathcal{X}^2$ at each $(\mathbf{X}, \mathbf{Y}) \in \mathcal{X}^2$. □

Proof of Theorem 4. To show the existence of a fixed point of the "mixed strategy" best response correspondence $B : \mathcal{X} \rightrightarrows \mathcal{X}$, we use Glicksberg's fixed point theorem (Aliprantis and Border, 2006, henceforth AP; Corollary 17.55). First, we confirm the assumptions on domain $\mathcal{X}$. The type space Θ is a complete separable metric space. Since the domain $\mathcal{X}$ is regarded as the set of distributional strategies over the product of this type space Θ and the finite strategy space $\mathcal{S}$, we can borrow the result in Milgrom and Weber (1985) about $\mathcal{X}$:[49] $\mathcal{X}$ is a nonempty, compact and convex subspace of $\mathcal{M}$, which is convex and Hausdorff under the weak topology.

With nonemptiness and compactness of $\mathcal{X}$, continuity of the maximized function $\int_\Theta \mathbf{F}[\mathbf{X}] \cdot d\mathbf{Y}$ implies by Berge's maximum theorem (AP, Theorem 17.31) that B is nonempty, compact-valued and upper hemicontinuous. In the Hausdorff metric space, this further implies by AP Theorem 17.10 that B has a closed graph. Since $\int_\Theta \mathbf{F}[\mathbf{X}] \cdot d\mathbf{Y}$ is a linear function of $\mathbf{Y}$, B is convex-valued. From the aforementioned properties of $\mathcal{X}$ and these properties of B, Glicksberg's theorem guarantees the existence of a fixed point of B (as well as compactness of the set of fixed points). □

[49] We could interpret $\int_\Theta \mathbf{F}[\mathbf{X}] \cdot d\mathbf{Y}$ as the "payoff" from distributional strategy $\mathbf{X}$ in their finite-player model. However, this is different from their payoff function, which is constructed from a normal-form game and thus is bilinear in $\mathbf{X}$ and $\mathbf{Y}$.

C.3 Proof of Theorem 5

For stability, we use the weak topology and apply the Lyapunov stability theorem as below.

Theorem 6 (Cheung, 2014: Theorems 5–6, Corollary 2)**.** *Let $Z \subset \mathcal{X}$ be a closed set and let $Y \subset \mathcal{X}$ be a neighborhood of Z in the weak topology on $\mathcal{X}$. Let $\mathcal{L} : Y \to \mathbb{R}$ be a decreasing Lyapunov function for dynamic $\mathbf{V}$: that is, $\mathcal{L}$ is continuous with respect to the weak topology and Fréchet-differentiable with $\dot{\mathcal{L}}(\mathbf{X}) = \langle \nabla \mathcal{L}(\mathbf{X}), \mathbf{V}[\mathbf{X}] \rangle \leq 0$ for all $\mathbf{X} \in Y$. Then, the following holds.*

i) Any solution path starting from Y converges to the set $\{\mathbf{X} \in Y \ : \ \dot{\mathcal{L}}(\mathbf{X}) = 0\}$ with respect to the weak topology; i.e., this set is attracting under $\mathbf{V}$.

ii) If $\mathcal{L}^{-1}(0) = Z$, Z is Lyapunov stable under $\mathbf{V}$ with respect to the weak topology. Furthermore, if $\dot{\mathcal{L}}(\mathbf{X}) < 0$ whenever $\mathbf{X} \in Y \setminus Z$, then Z is asymptotically stable under $\mathbf{V}$.

Part i) holds for an increasing Lyapunov function; part ii) is retained by defining Z as an isolated set of local maxima.

Proof of Theorem 5. **i)** Since f is a w-weighted potential function for $\mathbf{F}$, we have

$$\dot{f}(\mathbf{X}) = \langle \nabla f(\mathbf{X}), \dot{\mathbf{X}} \rangle = \langle w\mathbf{F}[\mathbf{X}], \mathbf{V}^{\mathbf{F}}[\mathbf{X}] \rangle = \int_\Theta w(\theta)\mathbf{F}[\mathbf{X}](\theta) \cdot \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta)\mu(d\theta),$$

for any $\mathbf{X} = \int \mathbf{x} d\mu \in \mathcal{X}$.

Since $\mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) = \mathbf{v}^\theta(\mathbf{F}[\mathbf{X}](\theta), \mathbf{x}(\theta))$, the first part of PC (8) implies $\mathbf{F}[\mathbf{X}](\theta) \cdot \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) \geq 0$ for all θ and thus

$$\dot{f}(\mathbf{X}) = \int_\Theta w(\theta)\mathbf{F}[\mathbf{X}](\theta) \cdot \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta)\mu(d\theta) \geq 0.$$

Suppose $\mathbf{X}$ is not a stationarity state under dynamic $\mathbf{V}^{\mathbf{F}}$, which is equivalent to $\mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) \neq \mathbf{0}$ for μ-almost all types θ . For a type with $\mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) \neq \mathbf{0}$, the second part of PC (8) implies $\mathbf{F}[\bar{\mathbf{x}}](\theta) \cdot \mathbf{v}^{\mathbf{F}}[\mathbf{x}](\theta) > 0$. Since this holds for a positive mass of types and $w(\theta) \in \mathbb{R}_{++}$ for all θ, we find that the above equation on $\dot{f}(\mathbf{X})$ holds with a strict inequality.

Therefore, f is a strictly increasing Lyapunov function and the set $\{\mathbf{X} \in \mathcal{X} : \dot{f}(\mathbf{X}) = 0\}$ is the set of stationary states, i.e.,$\{\mathbf{X} \in \mathcal{X} : \mathbf{V}^{\mathbf{F}}[\mathbf{X}] = \mathbf{O}\}$. By Theorem 6, this implies that the set of stationary states is globally attracting; a local maximum (local strict maximum, resp.) of f is Lyapunov stable (asymptotically stable, resp.).

ii) a) Suppose that the corresponding isolated stationary strategy distribution $\mathbf{X}^*$ is asymptotically stable, with a (nonempty) basin of attraction $\mathcal{X}^0 \subset \mathcal{X}^*$. Take an arbitrary strategy distribution $\mathbf{X}_0 \neq \mathbf{X}^*$ from $\mathcal{X}^0$ and let $\{\mathbf{X}_t\}_{t \in \mathbb{R}_+}$ be a solution trajectory under the heterogeneous dynamic $\mathbf{V}^{\mathbf{F}}$ from $\mathbf{X}_0$. Since f is a strictly increasing Lyapunov function, it must be the case that $\dot{f}(\mathbf{X}_t) > 0$ as long as $\mathbf{X}_t$ has not reached exactly $\mathbf{X}^*$. Thus, $f(\mathbf{X}^*) = f(\mathbf{X}_0) + \int_0^\infty \dot{f}(\mathbf{X}_t)dt > f(\mathbf{X}_0)$. Since $\mathbf{X}_0$ is taken arbitrarily from $\mathcal{X}^0$, this verifies that $\mathbf{X}^*$ strictly maximizes f in this neighborhood $\mathcal{X}^0$.

b) We prove the claim by contradiction. Assume that, while $\mathbf{X}^*$ is Lyapunov stable, $\mathbf{X}^*$ is not a local maximum of f. Take an open neighborhood $\tilde{\mathcal{X}}^*$ of $\mathbf{X}^*$ such that $\mathrm{cl}\, \tilde{\mathcal{X}}^* \subset \mathcal{X}^*$. Let $\mathcal{X}^1 := f^{-1}((f(\mathbf{X}^*) - h, f(\mathbf{X}^*) + h)) \cap \tilde{\mathcal{X}}^*$ with an arbitrarily fixed constant $h > 0$; since f is continuous, $\mathcal{X}^1$ is an open neighborhood of $\mathbf{X}^*$. By Lyapunov stability of $\mathbf{X}^*$, there exists another open neighborhood $\mathcal{X}^0 \subset \mathcal{X}^1$ of $\mathbf{X}^*$ such that any solution trajectory starting from $\mathcal{X}^0$ stays in $\mathcal{X}^1$ at any moment of time. As $\mathbf{X}^*$ is not a local maximum, there exists another strategy distribution $\mathbf{X}^\dagger \in \mathcal{X}^0$ such that $f(\mathbf{X}^\dagger) > f(\mathbf{X}^*)$; note that $f(\mathbf{X}^\dagger) < f(\mathbf{X}^*) + h$ since $\mathbf{X}^\dagger \in \mathcal{X}^0 \subset \mathcal{X}^1$.

Let $\bar{\mathcal{X}}^2 = f^{-1}([f(\mathbf{X}^*) - h, f(\mathbf{X}^*) + h]) \cap \mathrm{cl}\, \tilde{\mathcal{X}}^*$; $\mathcal{X}^1 \subset \bar{\mathcal{X}}^2$ and this is closed in $\mathcal{X}$ and thus compact, as $\mathcal{X}$ is compact. Consider a solution trajectory $\{\mathbf{X}^t\}_{t\in\mathbb{R}_+}$ starting from $\mathbf{X}^\dagger$; it stays in $\mathcal{X}^1$ and thus in $\bar{\mathcal{X}}^2$. By PC, $f(\mathbf{X}^t)$ is weakly increasing in t and thus $f(\mathbf{X}^t) \geq f(\mathbf{X}^0) = f(\mathbf{X}^\dagger) > f(\mathbf{X}^*)$. As $\mathbf{X}^*$ is the only stationary point in $\mathcal{X}^* \supset \bar{\mathcal{X}}^2$, this implies that the trajectory $\{\mathbf{X}^t\}$ does not arrive at or even converge to a stationary point; thus, PC further implies $f(\mathbf{X}^t)$ *strictly* increases in t. Since $f(\mathbf{X}^t) \in [f(\mathbf{X}^\dagger), f(\mathbf{X}^*) + h]$ (the upper bound obtained by $\mathbf{X}^t \in \bar{\mathcal{X}}^2$) for all t, $f(\mathbf{X}^t)$ must converge to some finite value $\bar{f} \in [f(\mathbf{X}^\dagger), f(\mathbf{X}^*) + h]$.

On the other hand, the value of the potential $f(\mathbf{X}^t)$ is expressed as

$$f(\mathbf{X}^t) = \int_0^t \dot{f}(\mathbf{X}_t)dt = \int_0^t \gamma(\mathbf{X}_t)dt$$

with $\gamma(\mathbf{X}) = \langle w\mathbf{F}[\mathbf{X}], \mathbf{V}^{\mathbf{F}}[\mathbf{X}]\rangle$.

Define $\tilde{\gamma} : \mathbb{R}_+ \to \mathbb{R}$ by $\tilde{\gamma}(t) = \gamma(\mathbf{X}^t)$. As $\gamma(\mathbf{X})$ is continuous in $\mathbf{X}$ and $\mathbf{X}^t$ is continuous in t, $\gamma(\mathbf{X}^t)$ is continuous in t. With the convergence of $f(\mathbf{X}^t) = \int_0^t \tilde{\gamma}(\tau)d\tau$, this implies there exists a sequence $\{t_n\}_{n\in\mathbb{N}}$ such that $t_n \to \infty$ and $\tilde{\gamma}(t_n) \to 0$.[50] As $\{\mathbf{X}^{t_n}\}_{n\in\mathbb{N}}$ is contained in compact set $\bar{\mathcal{X}}^2$, there further exists a subsequence $\{t'_m\}_{m'\in\mathbb{N}} \subset \{t_n\}_{n\in\mathbb{N}}$ such that $\mathbf{X}^{t'_m}$ converges to some $\mathbf{X}^\infty \in \bar{\mathcal{X}}^2 \subset \mathcal{X}^*$ as $m \to \infty$. Since $\tilde{\gamma}(t'_m) \to 0$ and γ is continuous, we have $\gamma(\mathbf{X}^\infty) = 0$, which implies by PC of $\mathbf{v}$ that $\mathbf{X}^\infty \in \mathcal{X}^*$ is a stationary point. Since $\mathbf{X}^*$ is the only stationary point in $\mathcal{X}^*$, this limit point $\mathbf{X}^\infty$ must be $\mathbf{X}^*$. However, as $f(\mathbf{X}^{t'_m}) \to \bar{f}$ as $t'_m \to \infty$ and f is also continuous, we have $f(\mathbf{X}^\infty) = \bar{f} > f(\mathbf{X}^*)$ and thus $\mathbf{X}^\infty \neq \mathbf{X}^*$, a contradiction. □

C.4 Proof of Theorem 6

Proof for Example 1′. As $f(\mathbf{X}) = f^0(\mathbf{X}(\Theta)) + \int_\Theta \boldsymbol{\theta} \cdot \mathbf{X}(d\boldsymbol{\theta})$, weak continuity of f is obtained from continuity of f^0 and the dominated convergence theorem. By applying the definition of f to $\mathbf{X} + \Delta\mathbf{X}$, we have

$$\begin{aligned} &f(\mathbf{X} + \Delta\mathbf{X}) = f^0(\bar{\mathbf{x}} + \Delta\bar{\mathbf{x}}) + \int_\Theta \boldsymbol{\theta} \cdot (\mathbf{X} + \Delta\mathbf{X})(d\boldsymbol{\theta}) \\ =& \left\{ f^0(\bar{\mathbf{x}}) + \nabla f^0(\bar{\mathbf{x}}) \cdot \Delta\bar{\mathbf{x}} + o(|\Delta\bar{\mathbf{x}}|) \right\} + \left\{ \int_\Theta \boldsymbol{\theta} \cdot \mathbf{X}(d\boldsymbol{\theta}) + \int_\Theta \boldsymbol{\theta} \cdot \Delta\mathbf{X}(d\boldsymbol{\theta}) \right\} \\ =& f(\mathbf{X}) + \mathbf{F}^0(\bar{\mathbf{x}}) \cdot \Delta\bar{\mathbf{x}} + \int_\Theta \boldsymbol{\theta} \cdot \Delta\mathbf{X}(d\boldsymbol{\theta}) + o(|\Delta\bar{\mathbf{x}}|) = f(\mathbf{X}) + \int_\Theta (\mathbf{F}^0(\bar{\mathbf{x}}) + \boldsymbol{\theta}) \cdot \Delta\mathbf{X}(d\boldsymbol{\theta}) + o(|\Delta\bar{\mathbf{x}}|). \end{aligned}$$

Here $\bar{\mathbf{x}} = \mathbf{X}(\Theta)$ and $\Delta\bar{\mathbf{x}} = \Delta\mathbf{X}(\Theta)$. The second equality comes from differentiability of f^0; the third is from the assumption that f^0 is a potential function of $\mathbf{F}^0$ and the definition of f applied

[50] Consider continuous function $g : \mathbb{R}_+ \to \mathbb{R}_+$ that always takes a non-negative value. Improper integral $\int_0^\infty g(\tau)d\tau$ is well-defined (i.e., converges to a finite real number) if and only if, for any $\varepsilon > 0$, there exists $T \geq 0$ such that $|\int_t^{t'} g(\tau)d\tau| < \varepsilon$ for any $t' > t > T$; fix T to the one that satisfies this for $\varepsilon = 1$. Define a sequence $\{t_n\}_{n\in\mathbb{N}}$ by letting t_n be the moment of time to attain the minimum value of g in close interval of time $[T + 2^n + 1, T + 2^{n+1}]$, i.e., $g(t_n) = \min_{\tau\in[T+2^n+1, T+2^{n+1}]} g(\tau)$; this minimum exists and $t_n \to \infty$, since g is continuous and we have $t_n \leq T + 2^{n+1} \leq T + 2^{n+1} + 1 \leq t_{n+1}$. For each $n \in \mathbb{N}$, this t_n satisfies $0 \leq (2^{n+1} - 2^n - 1)g(t_n) \leq \int_{T+2^n+1}^{T+2^{n+1}} g(\tau)d\tau = |\int_{T+2^n+1}^{T+2^{n+1}} g(\tau)d\tau| < 1$; hence, $0 \leq g(t_n) \leq 1/(2^{n+1} - 2^n - 1) = 1/(2^n - 1)$. This implies $g(t_n) \to 0$ as $n \to 0$ and thus $t_n \to \infty$.

to $\mathbf{X}$. Then, we should recall $\mathbf{F}[\mathbf{X}(\Theta)](\boldsymbol{\theta}) = \mathbf{F}^0(\mathbf{X}(\Theta)) + \boldsymbol{\theta}$. So the second term is $\langle \mathbf{F}[\mathbf{X}(\Theta)], \Delta\mathbf{X}\rangle$. About the third error term, note that $|\Delta\bar{\mathbf{x}}| = |\Delta\mathbf{X}(\Theta)| \leq \|\Delta\mathbf{X}\|$. Therefore, we obtain

$$f(\mathbf{X}+\Delta\mathbf{X}) = f(\mathbf{X}) + \langle \mathbf{F}[\mathbf{X}(\Theta)], \Delta\mathbf{X}\rangle + o(\|\Delta\mathbf{X}\|).$$

Thus, f is (Fréchet) differentiable with derivative $\nabla f(\mathbf{X}) \equiv \mathbf{F}[\mathbf{X}(\Theta)]$. So, we have verified that f is a potential function of the game $\mathbf{F}$ defined on $\mathcal{X}$. □

Proof for Example 2′. Let $\Delta\mathbf{X} = \int_\Theta \Delta\mathbf{x} d\mu$ and $\Delta\mathbf{X}' = \int_\Theta \Delta\mathbf{x}' d\mu$. Then, we have

$$\begin{aligned}
f(\mathbf{X}+\Delta\mathbf{X}) - f(\mathbf{X}) &= \int_\Sigma \{f^\sigma(\mathbf{X}+\Delta\mathbf{X}) - f^\sigma(\mathbf{X})\}\mathbb{P}_\Sigma(d\sigma) \\
&= \int_\Sigma \left\{\int_\Theta w(\theta)\mathbf{F}^\sigma[\mathbf{X}]\cdot\Delta\mathbf{x}(\theta)\mathbb{P}(d\theta|\sigma) + o(\|\Delta\mathbf{X}\|_\infty)\right\}\mathbb{P}_\Sigma(d\sigma) \\
&= \int_\Theta\int_\Sigma w(\theta)\mathbf{F}^\sigma[\mathbf{X}]\cdot\Delta\mathbf{x}(\theta)\mathbb{P}_\Sigma(d\sigma|\theta)\mathbb{P}(d\theta) + o(\|\Delta\mathbf{X}\|_\infty) \\
&= \int_\Theta w(\theta)\left(\int_\Sigma \mathbf{F}^\sigma[\mathbf{X}]\mathbb{P}_\Sigma(d\sigma|\theta)\right)\cdot\Delta\mathbf{x}(\theta)\mathbb{P}(d\theta) + o(\|\Delta\mathbf{X}\|_\infty)
\end{aligned}$$

Since $\mathbf{F}[\mathbf{X}](\theta) = \int_\Sigma \mathbf{F}^\sigma[\mathbf{X}]\mathbb{P}_\Sigma(d\sigma|\theta)$, this verifies that $\mathbf{F}[\mathbf{X}]$ is the Fréchet derivative of f and thus f is a potential function of $\mathbf{F}$. □

Proof for Example 3′. Let $\Delta\mathbf{X} = \int_\Theta \Delta\mathbf{x} d\mu$ and $\Delta\mathbf{X}' = \int_\Theta \Delta\mathbf{x}' d\mu$. Then, since f^0 is the potential function for $\mathbf{F}$ and $g(\theta_1,\theta_2) = g(\theta_2,\theta_1)$ for any $(\theta_1,\theta_2)\in\Theta^2$, we have

$$\begin{aligned}
&\int_{\Theta^2} \nabla_i f^0(\mathbf{x}(\theta_1),\mathbf{x}(\theta_2))\Delta\mathbf{x}(\theta_i)g(\theta_1,\theta_2)\mu(d\theta_1)\mu(d\theta_2) \\
=&\int_{\Theta^2} w_i\mathbf{F}^0(\mathbf{x}(\theta_i),\mathbf{x}(\theta_j))\Delta\mathbf{x}(\theta_i)g(\theta_1,\theta_2)\mu(d\theta_1)\mu(d\theta_2) \\
=&w_i\int_{\Theta^2} \mathbf{F}^0(\mathbf{x}(\theta_i),\mathbf{x}(\theta_j))g(\theta_i,\theta_j)\mu(d\theta_j)\Delta\mathbf{x}(\theta_i)\mu(d\theta_i) \\
=&w_i\int_\Theta \mathbf{F}[\mathbf{X}](\theta_i)\cdot\Delta\mathbf{x}(\theta_i)\mu(d\theta_i) = w_i\int_\Theta \mathbf{F}[\mathbf{X}](\theta_i)\cdot\Delta\mathbf{X}(d\theta_i) \qquad \text{for each } i\in\{1,2\}, j\neq i.
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
&f(\mathbf{X}+\Delta\mathbf{X}) - f(\mathbf{X}) \\
=&\frac{1}{w_1+w_2}\int_{\Theta^2}\left\{\nabla_1 f^0(\mathbf{x}(\theta),\mathbf{x}(\theta'))\Delta\mathbf{x}(\theta) + \nabla_2 f^0(\mathbf{x}(\theta),\mathbf{x}(\theta'))\Delta\mathbf{x}(\theta')\right\}g(\theta,\theta')\mu(d\theta)\mu(d\theta') + o(\|\Delta\mathbf{X}\|_\infty) \\
=&\frac{1}{w_1+w_2}\left(w_1\int_\Theta \mathbf{F}[\mathbf{X}](\theta)\cdot\Delta\mathbf{X}(d\theta) + w_2\int_\Theta \mathbf{F}[\mathbf{X}](\theta')\cdot\Delta\mathbf{X}(d\theta')\right) + o(\|\Delta\mathbf{X}\|_\infty) \\
=&\int_\Theta \mathbf{F}[\mathbf{X}](\theta)\cdot\Delta\mathbf{X}(d\theta) + o(\|\Delta\mathbf{X}\|_\infty) = \langle\mathbf{F}[\mathbf{X}],\Delta\mathbf{X}\rangle + o(\|\Delta\mathbf{X}\|_\infty).
\end{aligned}$$

That is, $\mathbf{F}[\mathbf{X}]$ is the Fréchet derivative of f and thus f is a potential function of $\mathbf{F}$. □

References

ALIPRANTIS, C. D. AND K. C. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer, 3 ed.

ANDERSON, S. P., A. DE PALMA, AND J. F. THISSE (1992): *Discrete choice theory of product differentiation*, MIT Press.

ANGELETOS, G.-M. AND A. PAVAN (2007): "Efficient Use of Information and Social Value of Information," *Econometrica*, 75, 1103–1142.

BJÖRNERSTEDT, J. AND J. W. WEIBULL (1996): "Nash equilibrium and evolution by imitation," in *The rational foundations of economic behavior*, ed. by K. Arrow, E. Colombatto, M. Perlman, and C. Schmidt, New York, NY: St. Martin's Press, 155–181.

BLONSKI, M. (1999): "Anonymous games with binary actions," *Games and Economic Behavior*, 28, 171–180.

CHEUNG, M.-W. (2014): "Pairwise Comparison Dynamics for Games with Continuous Strategy Space," *Journal of Economic Theory*, 153, 344–375.

——— (2016): "Imitative dynamics for games with continuous strategy space," *Games and Economic Behavior*, 99, 206 – 223.

CORCHÓN, L. C. (1994): "Comparative statics for aggregative games the strong concavity case," *Mathematical Social Sciences*, 28, 151 – 165.

DHEBAR, A. AND S. S. OREN (1985): "Optimal Dynamic Pricing for Expanding Networks," *Marketing Science*, 4, 336–351.

——— (1986): "Optimal Dynamic Pricing for Expanding Networks," *Operations Research*, 34, 384–394.

DIXIT, A. K. AND J. E. STIGLITZ (1977): "Monopolistic Competition and Optimum Product Diversity," *The American Economic Review*, 67, 297–308.

ELY, J. AND W. H. SANDHOLM (2005): "Evolution in Bayesian Games I: Theory," *Games and Economic Behavior*, 53, 83–109.

FOX, M. J. AND J. S. SHAMMA (2013): "Population games, stable games, and passivity," *Games*, 4, 561–583.

FRIEDMAN, D. AND D. N. OSTROV (2013): "Evolutionary dynamics over continuous action spaces for population games that arise from symmetric two-player games," *Journal of Economic Theory*, 148, 743–777.

FUDENBERG, D. AND D. M. KREPS (1993): "Learning Mixed Equilibria," *Games and Economic Behavior*, 5, 320–367.

FUDENBERG, D. AND D. K. LEVINE (1998): *The theory of learning in games*, vol. 2, MIT press.

GILBOA, I. AND A. MATSUI (1991): "Social Stability and Equilibrium," *Econometrica*, 59, 859–867.

HOFBAUER, J. (1995a): "Imitation Dynamics for Games," Mineo, University of Vienna.

——— (1995b): "Stability for the best response dynamics," Mimeo, University of Vienna.

——— (2001): "From Nash and Brown to Maynard Smith: Equilibria, Dynamics and ESS," *Selection*, 1, 81–88.

HOFBAUER, J., J. OECHSSLER, AND F. RIEDEL (2009): "Brown–von Neumann–Nash dynamics: The continuous strategy case," *Games and Economic Behavior*, 65, 406–429.

HOFBAUER, J. AND W. H. SANDHOLM (2002): "On the global convergence of stochastic fictitious play," *Econometrica*, 70, 2265–2294.

——— (2009): "Stable games and their dynamics," *Journal of Economic Theory*, 144, 1665–1693.

HOFBAUER, J., W. H. SANDHOLM, J. HOFBAUER, AND W. H. SANDHOLM (2007): "Evolution in games with randomly disturbed payoffs," *Journal of Economic Theory*, 132, 47–69.

HUMMEL, P. AND R. P. MCAFEE (2018): "Evolutionary consumers imply monopolies exit," *International Economic Review*, 59, 1733–1746.

HWANG, S.-H., M. KATSOULAKIS, AND L. REY-BELLET (2013): "Deterministic equations for stochastic spatial evolutionary games," *Theoretical Economics*, 8, 829–874.

JENSEN, M. K. (2018): "Aggregative games," in *Handbook of Game Theory and Industrial Organization*, ed. by L. C. Corchón and M. A. Marini, Edward Elgar Publishing, vol. 1, chap. 4, 66–92.

LAHKAR, R. AND F. RIEDEL (2015): "The logit dynamic for games with continuous strategy sets," *Games and Economic Behavior*, 91, 268–282.

LAHKAR, R. AND R. M. SEYMOUR (2013): "Reinforcement learning in population games," *Games and Economic Behavior*, 80, 10–38.

MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic theory*, Oxford University Press.

MILGROM, P. R. AND R. J. WEBER (1985): "Distributional strategies for games with incomplete information," *Mathematics of operations research*, 10, 619–632.

OECHSSLER, J. AND F. RIEDEL (2001): "Evolutionary dynamics on infinite strategy spaces," *Economic Theory*, 17, 141–162.

——— (2002): "On the dynamic foundation of evolutionary stability in continuous models," *Journal of Economic Theory*, 107, 223–252.

ROHLFS, J. (1974): "A Theory of Interdependent Demand for a Communications Service," *Bell Journal of Economics and Management Science*, 5, 16–37.

SANDHOLM, W. H. (2001): "Potential Games with Continuous Player Sets," *Journal of Economic Theory*, 97, 81–108.

——— (2002): "Evolutionary Implementation and Congestion Pricing," *Review of Economic Studies*, 69, 667–689.

——— (2005): "Negative Externalities and Evolutionary Implementation," *Review of Economic Studies*, 72, 885–915.

——— (2010): *Population games and evolutionary dynamics*, MIT Press.

SAWA, R. AND D. ZUSAI (2014): "Evolutionary imitative dynamics with population-varying aspiration levels," *Journal of Economic Theory*, 154, 562–577.

SCHLAG, K. H. (1998): "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits," *Journal of Economic Theory*, 78, 130–56.

SMITH, M. J. (1984): "The Stability of a Dynamic Model of Traffic Assignment: An Application of a Method of Lyapunov," *Transportation Science*, 18, 245–252.

TAYLOR, P. AND L. JONKER (1978): "Evolutionarily Stable Strategies and Game Dynamics," *Mathematical Biosciences*, 40, 145–156.

UI, T. (2009): "Bayesian potentials and information structures: Team decision problems revisited," *International Journal of Economic Theory*, 5, 271–291.

VAN HEUMEN, R., B. PELEG, S. TIJS, AND P. BORM (1996): "Axiomatic characterizations of solutions for Bayesian games," *Theory and Decision*, 40, 103–129.

WU, J. AND D. ZUSAI (2019): "A potential game approach to modelling evolution in a connected society," *Nature Human Behaviour*, 3, 604–610.

ZUSAI, D. (2015): "Market size effects on long-run demand of a network good," *Economics Bulletin*, 35, 2768–2775.

——— (2017): "Distributional stability and deterministic equilibrium selection under heterogeneous evolutionary dynamics," Mimeo, `https://arxiv.org/abs/1805.04895`.

——— (2018): "Tempered best response dynamics," *International Journal of Game Theory*, 47, 1–34.

——— (2020): "Gains in evolutionary dynamics: A unifying and intuitive approach to linking static and dynamic stability," Mimeo, `https://arxiv.org/abs/1805.04898`.