TCER Working Paper Series

Statistical Inference in Evolutionary Dynamics

Ryoji Sawa

March 2022

Working Paper E-170
https://www.tcer.or.jp/wp/pdf/e170.pdf

TOKYO CENTER FOR ECONOMIC RESEARCH
1-7-10-703 Iidabashi, Chiyoda-ku, Tokyo 102-0072, Japan

Abstract

We introduce evolutionary dynamics for two-action games where agents with diverse preferences use statistical inference to guide their behavior. We show that the dynamic converges to a Bayesian sampling equilibrium with statistical inference (SESI) and the set of Bayesian SESIs is globally asymptotically stable. We discuss the global convergence to a unique Bayesian SESI in anti-coordination games, a welfare-improving tax scheme, equilibrium selection in coordination games, an application to the diffusion of behavior on networks, and the extension of heterogeneity to the inference procedures.

Ryoji Sawa
University of Tsukuba
Department of Policy and Planning Sciences
Tennoudai 1-1-1, Tsukuba, Ibaraki, Japan
rsawa@sk.tsukuba.ac.jp

Statistical Inference in Evolutionary Dynamics*

Ryoji Sawa and Jiabin Wu†

December 7, 2021

Abstract

We introduce evolutionary dynamics for two-action games where agents with diverse preferences use statistical inference to guide their behavior. We show that the dynamic converges to a Bayesian sampling equilibrium with statistical inference (SESI) and the set of Bayesian SESIs is globally asymptotically stable. We discuss the global convergence to a unique Bayesian SESI in anti-coordination games, a welfare-improving tax scheme, equilibrium selection in coordination games, an application to the diffusion of behavior on networks, and the extension of heterogeneity to the inference procedures.

Keywords: Statistical inference, Bounded rationality, Deterministic evolutionary game theory, Sampling best response, Network games.
JEL codes: C44, C73, D80.

---

*We thank Daisuke Oyama, Satoru Takahashi and Dai Zusai for their constructive suggestions. We also thank seminar participants at Osaka University and Tohoku University for their helpful comments. RS acknowledges financial support from the Japan Society for the Promotion of Science: JSPS KAKENHI Grants-in-Aid for Early-Career Scientists No. 18K12740, Nomura Foundation (Grants in Social Science), KIER Foundation, and Tokyo Center for Economic Research (TCER). We dedicate this research to the memory of William H. Sandholm.

†Sawa: Department of Policy and Planning Sciences, University of Tsukuba, email: rsawa@sk.tsukuba.ac.jp. Wu: Department of Economics, University of Oregon, email: jwu5@uoregon.edu. Please direct correspondence to either author.

1 Introduction

Evolutionary game theory considers the dynamics of behavior in a large population of agents who choose actions in response to the action distribution in the population. For example, a consumer's choice of purchasing a status good or a fashion good depends on the number of consumers buying the same good. However, the exact action distribution is not always available and agents may only have limited information on others' actions. Sandholm [2001] and Oyama et al. [2015] consider evolutionary dynamics for such an imperfect information environment where agents draw a random sample of others' actions and best-respond to it. This paper takes a step further by assuming that the agents are able to use some statistical inference procedure to form an estimate of the action distribution.

Salant and Cherry [2020] pioneer the idea of incorporating statistical inference based on random sampling into games. They introduce a solution concept called the *sampling equilibrium with statistical inference* (SESI), prove its existence, and provide a dynamic justification for it. It is a distribution of actions such that sampling from this distribution and best-responding to the estimate using statistical inference result in the same distribution of actions. Based on their insights, we further explore the dynamic implications of statistical inference by describing population dynamics at the individual level. Agents with heterogeneous preferences play a two-action population game, and periodically revise actions in discrete time. The revising agents obtain a random sample of others' actions, use statistical inference to estimate the action distribution, and best-respond to the estimate.

Since agents' preferences are diverse, we propose a solution concept called the Bayesian SESI, in which each agent best-responds to the estimate derived from the drawn sample using statistical inference. Though it is conceptually equivalent to the original SESI in Salant and Cherry [2020], it provides a formal mathematical expression for the action distribution for each preference. Then we propose a *Bayesian k-sampling statistical inference best response dynamic* (Bayesian dynamic, for short) based on the evolutionary model with heterogeneous preferences (Ely and Sandholm 2005).[1] The Bayesian dynamic gives rise to a simple aggregate dynamic with one variable, the proportion of agents playing one of the actions, which coin-

[1]See Zusai [2018, 2019a,b] for general treatments of heterogeneity and aggregability in evolutionary dynamics.

cides with the dynamic in Salant and Cherry [2020]. Our first main result shows that the Bayesian dynamic converges to a Bayesian SESI if and only if the aggregate dynamic converges to the corresponding SESI proportion. Hence, we can utilize the aggregate dynamic to study the convergence of the Bayesian dynamic. Our second main result shows the global asymptotic stability of the set of Bayesian SESIs. That is, the Bayesian dynamic must converge to a Bayesian SESI. The two results clarify that there must be a process at the level of individual strategy adjustment that underlies the aggregate dynamic, and that nothing essential is lost by working directly with the aggregates in terms of convergence.

We examine the implications of our model in two classes of games, anti-coordination and coordination games. In anti-coordination games, there exists a unique Bayesian SESI, and the Bayesian dynamic globally converges to it. This supports the convergence result on the aggregate level shown in Salant and Cherry [2020] from the individual level perspective. Nevertheless, the Bayesian dynamic in general converges more slowly than the aggregate dynamic does. This discrepancy between the two dynamics suggests fluctuations in social welfare even after the aggregate behavior is stabilized. On improving social welfare, we extend the model to a form of an implementation problem. Equilibria in anti-coordination games are, in general, inefficient since the agents do not take into account the negative externality. We characterize a simple condition under which a tax scheme can improve social welfare. In coordination games with strict Nash equilibria, the dynamic has a sharp equilibrium selection result: for unbiased statistical inferences, the dynamic almost globally converges to the $1/k$-dominant equilibrium when the sample size is larger than 1 and at most k. This extends the results of Sandholm [2001] and Oyama et al. [2015] to games with heterogeneous preferences and sampling best response dynamics with statistical inference.

We consider two extensions. First, we extend the (almost) global convergence result in coordination games to network games. Network games assume an arbitrary structure for the interactions among agents. Hence, they are different from population games. Nevertheless, with a slight modification, our evolutionary dynamics can describe a particular diffusion process on networks, considered in Jackson and Yariv [2007], with agents using statistical inference. The global convergence result for coordination games can be extended almost directly. Second, we allow the agents to adopt different inference procedures and different sample sizes.

The two main theorems, namely the simultaneous convergence of the aggregate and Bayesian dynamics and the global asymptotic stability of the Bayesian SESIs, are robust against such heterogeneity.

The idea of sampling in games has long been considered in the literature. Young [1993, 1998] consider stochastic learning in games where the agents sample from the recent history of play. Kosfeld et al. [2002] consider an adjustment process in multi-player normal form games where the agents sample (pure-) action profiles from their opponents' mixed strategies. Kreindler and Young [2013] examine the expected convergence time in a logit stochastic evolution model where the agents obtain a random sample of other agents' actions as in Sandholm [2001] and Oyama et al. [2015]. Heller and Mohlin [2018] consider settings where agents stochastically choose the size of the sample in every period, and show that the mean of sample size greatly determines the dependence of future behavior on initial behavior. The procedurally rational agents of Osborne and Rubinstein [1998] sample the payoffs of each action and choose the one with the highest realized payoffs. This is called the payoff sampling approach, and its application is studied by Spiegler [2006a,b], Miekisz and Ramsza [2012], and Mantilla et al. [2020], for example. Evolutionary dynamics with such agents are considered by Sethi [2000, 2021], Sandholm et al. [2019, 2020], and Arigapudi et al. [2021]. Osborne and Rubinstein [2003] consider sampling equilibrium in static games similar to Salant and Cherry [2020] but without the consideration of statistical inference. To our limited knowledge, our paper is the first to systematically investigate statistical inference based on random sampling in evolutionary dynamics.

Examples of standard statistical inferences include Bayesian updating, maximum likelihood estimation, among others. It is also becoming increasingly apparent that people may adopt certain heuristics in their inference procedures. We just name a few related studies. Prospect theory of Kahneman and Tversky [1979] posits that people transform objective probabilities non-linearly. Roughly speaking, people weigh unlikely events more than their objective probabilities.[2] Prospect theory fits experimental data well and is further examined and developed by many studies, Bernheim and Sprenger [2020], Tversky and Kahneman [1992], Wu and Gonzalez [1996], for example.

Network games and network diffusion have been important research topics ex-

[2]See Example 2 for a model that incorporates such a non-linear transformation.

plored by a number of studies, including Arieli et al. [2020], Galeotti et al. [2010], Jackson and Yariv [2007], Kobayashi and Onaga [2021], Morris [2000], Newton and Angus [2015], Oyama and Takahashi [2015], for example. Our extension follows the network diffusion model in Jackson and Yariv [2007]. In contrast to the extant literature, we assume that agents use statistical inference to estimate a state variable of the network, and we characterize a sufficient condition for contagion with such agents.

The paper is organized as follows. Section 2 introduces the population game and our equilibrium concept. Section 3 introduces the evolutionary model, establishes the key equivalence result between the aggregate dynamic and the dynamic of the subpopulation of agents with a given preference, and prove the global asymptotic stability of the set of Bayesian SESIs. Sections 4 and 5 consider anti-coordination games and coordination games, respectively. Section 6 discusses an application to network games and the extension of heterogeneous inference procedures. Section 7 concludes.

2 Model

2.1 Population game

There is a single unit-mass population of agents who play a two-action game. Let $\mathcal{S} = \{A, B\}$ denote the set of actions. The utility from action B is 0. The utility from action A is $u(\theta, \alpha) = \theta - f(\alpha)$, where θ is an agent's idiosyncratic preference, and $f(\alpha)$ is the cost incurred by an agent taking action A when the proportion of agents taking A is α. The function f is continuous on $[0, 1]$. The distribution of preference θ is described by a probability measure λ on the set $\Theta = [0, 1]$. The probability measure λ is absolutely continuous with respect to the Lebesgue measure on Θ. Compared with Salant and Cherry [2020], who consider the uniform distribution, we allow for general preference distribution. We call the fraction of agents choosing action A the *aggregate population state*. Let $\mathcal{A} = [0, 1]$ be the set of aggregate population states. The game's payoffs can be described by the continuous function $F : \Theta \times \mathcal{A} \to \mathbb{R}^{\mathcal{S}}$. Let $F^{\theta}_{s}(\alpha)$ denote the payoff obtained at state α by agents with preference θ choosing $s \in S$, i.e. $F^{\theta}_{A}(\alpha) = \theta - f(\alpha)$ and $F^{\theta}_{B}(\alpha) = 0$.

The pure best response correspondence to the aggregate population state is denoted by b^θ:

$$b^\theta(\alpha) = \operatorname*{argmax}_{s \in \mathcal{S}} F^\theta_s(\alpha). \tag{1}$$

We assume that the agents choose action B when there is a tie. For each aggregate population state $\alpha \in \mathcal{A}$, the set of agents who are indifferent between the actions has measure zero. Thus, our result does not depend on the tie-breaking rule.

2.2 Statistical inference and equilibrium concept

The agents know their own payoff functions but do not have precise information on the population state α.[3] They are assumed to choose action using a statistical inference. Each agent randomly samples k agents, uses an inference procedure, which is an analogue of an estimator in statistics, to estimate the probability distribution over action distributions, and best-responds to it. Let z denote the sample mean of action A, that is, $z = j/k$ where j is the number of action-A agents in the sample. An *inference procedure* $\mathcal{G} = \{\mathcal{G}_{k,z}\}$ assigns a cumulative distribution function $\mathcal{G}_{k,z}$ to every sample (k, z) such that $\mathcal{G}_{k,\hat{z}}$ strictly first-order stochastically dominates $\mathcal{G}_{k,z}$ when $\hat{z} > z$. We introduce a few examples of inference procedures. Examples 1 and 3 are Examples 2 and 4 in Salant and Cherry [2020].

Example 1 (Maximum Likelihood Estimation (MLE))**.** The agents use the maximum likelihood estimation method (MLE) to estimate the most likely parameters to generate the sample. That is, the agents observing z solve for α that maximizes $\alpha^{kz}(1-\alpha)^{k(1-z)}$. The solution is $\alpha = z$. The inference procedure with MLE is

$$\mathcal{G}^{MLE}_{k,z}(\alpha) = 1_{\alpha \geq z} = \begin{cases} 0 & \text{if } \alpha < z, \\ 1 & \text{if } \alpha \geq z. \end{cases}$$

Hence, the agents consider the distribution of the sample as the true distribution.

Example 2 (Overweighting Low Probabilities (OLP))**.** This is a modified version of the nonlinear transformation of probabilities in Tversky and Kahneman [1992].

[3] Our results do not depend on whether the distribution λ is known to the agents.

The agents observing z use the following inference procedure:

$$\mathcal{G}_{k,z}^{OLP}(\alpha) = 1_{\alpha \geq \tilde{z}} = \begin{cases} 0 & \text{if } \alpha < \tilde{z}, \\ 1 & \text{if } \alpha \geq \tilde{z}, \end{cases} \quad \text{where } \tilde{z} = \frac{z^\delta}{z^\delta + (1-z)^\delta} \text{ for some } \delta \in (0,1].$$

If $\delta = 1$, then this inference procedure coincides with the MLE. For $\delta < 1$, this inference procedure exhibits the property of overweighting low probabilities and underweighting high probabilities.

For example, suppose $\delta = 0.6$. When $z = 0.1$,

$$\tilde{z} = \frac{0.1^{0.6}}{0.1^{0.6} + 0.9^{0.6}} \approx 0.211 > 0.1.$$

When $z = 0.9$,

$$\tilde{z} = \frac{0.9^{0.6}}{0.9^{0.6} + 0.1^{0.6}} \approx 0.789 < 0.9.$$

Example 3 (Truncated Normal (TN))**.** The agents observing z believe that α is distributed according to a normal distribution truncated symmetrically around the mean z with the variance being $\frac{z(1-z)}{k}$.

Note that both MLE and OLP result in point estimators, while TN results in a continuous cumulative distribution function.

Let us start by considering static solution concepts for this game. It is useful to first examine the one proposed by Salant and Cherry [2020]:

Definition 1 (SESI)**.** *A* sampling equilibrium with statistical inference (SESI) *is a number $\alpha_{k,\mathcal{G}} \in [0,1]$ such that an $\alpha_{k,\mathcal{G}}$ proportion of agents chooses action A when each agent obtains k independent observations from the aggregate population state $\alpha = \alpha_{k,G}$, forms an estimate according to inference procedure $\mathcal{G}$, and best-responds to this estimate in choosing an action. We refer to $\alpha_{k,\mathcal{G}}$ that constitutes a SESI as a* SESI proportion *(of degree k with respect to the inference procedure $\mathcal{G}$).*

Next, we characterize the SESI (proportion) $\alpha_{k,\mathcal{G}}$. The pure best response correspondence of agents with preference θ, inference procedure $\mathcal{G}$, and observation (k, z) is denoted by $b_{\mathcal{G}}^{\theta}$. It is described as

$$b^{\theta}_{\mathcal{G}}(k,z) = \operatorname*{argmax}_{s\in\mathcal{S}} \int_{\alpha\in\mathcal{A}} F^{\theta}_{s}(\alpha) d\mathcal{G}_{k,z}(\alpha) \qquad \forall z \in \left\{0, \frac{1}{k} \ldots, 1\right\}.$$

Recall that the agents choose action B when there is a tie. This means that $b^{\theta}_{\mathcal{G}}(k,z) \in \{A, B\}$. Let

$$B^{\theta}_{\mathcal{G}}(z) = \begin{cases} 1 & \text{if } b^{\theta}_{\mathcal{G}}(k,z) = A, \\ 0 & \text{if } b^{\theta}_{\mathcal{G}}(k,z) = B. \end{cases}$$

In words, $B^{\theta}_{\mathcal{G}}(z)$ is an index function that becomes one if action A is the best response to the estimate for the agents with preference θ and sample mean z.

Using the above expressions, a SESI proportion $\alpha_{k,\mathcal{G}}$ can be written as follows.

$$\alpha_{k,\mathcal{G}} = \int_{\theta\in\Theta} \sum_{j=0}^{k} \binom{k}{j} \alpha^{j}_{k,\mathcal{G}} (1-\alpha_{k,\mathcal{G}})^{k-j} B^{\theta}_{\mathcal{G}} \left(\frac{j}{k}\right) d\lambda. \tag{2}$$

Let $F_{k,z}$ denote the expected cost f under $\mathcal{G}_{k,z}$, that is,

$$F_{k,z} = \int_{\alpha\in\mathcal{A}} f(\alpha) d\mathcal{G}_{k,z}(\alpha).$$

Observe that the integral of the index function $B^{\theta}_{\mathcal{G}}(z)$ over Θ can be rewritten as

$$\begin{aligned} \int_{\theta\in\Theta} B^{\theta}_{\mathcal{G}}(z) d\lambda &= \int_{\theta\in\Theta} \mathbb{1} \left[\int_{\alpha\in\mathcal{A}} F^{\theta}_{A}(\alpha) d\mathcal{G}_{k,z}(\alpha) > 0 \right] d\lambda = \int_{\theta\in\Theta} \mathbb{1} \left[\theta > F_{k,z}\right] d\lambda \\ &= 1 - \Lambda(F_{k,z}), \end{aligned}$$

where $\Lambda(x) = \int_0^x d\lambda$. Then, Eq.(2) is written as follows.

$$\alpha_{k,\mathcal{G}} = \sum_{j=0}^{k} \binom{k}{j} \alpha^{j}_{k,\mathcal{G}} (1-\alpha_{k,\mathcal{G}})^{k-j} (1 - \Lambda(F_{k,j/k})).$$

Define the kth-order Bernstein polynomial of function v as[4]

$$Bern_k(\alpha; v) = \sum_{j=0}^{k} \binom{k}{j} \alpha^j (1-\alpha)^{k-j} v(j/k) \qquad \forall \alpha \in [0,1]. \tag{3}$$

The SESI proportion $\alpha_{k,\mathcal{G}}$ can be rewritten as follows.

$$1 - \alpha_{k,\mathcal{G}} = Bern_k(\alpha_{k,\mathcal{G}}; \Lambda_{F,k}), \tag{4}$$

where $\Lambda_{F,k}$ is a function such that $\Lambda_{F,k}(z) = \Lambda(F_{k,z})$.[5] We have the following observation. This comes from the fact that Eq.(3) is a polynomial function of α with $\Lambda(F_{k,z})$ bounded for all $z = \{0, 1/k, \ldots, 1\}$.

Observation 1. *For any finite $k > 0$, the Bernstein polynomial $Bern_k(\alpha_{k,\mathcal{G}}; \Lambda_{F,k})$ is Lipschitz continuous in $\alpha_{k,\mathcal{G}}$.*

The SESI proportion $\alpha_{k,\mathcal{G}}$ is a concise expression for an equilibrium. However, it is uncertain that a scalar value is sufficient to describe an equilibrium when it comes to large population games where heterogeneous agents may occasionally update.

To examine such dynamics, we use the notion of Bayesian strategy (Ely and Sandholm 2005), which describes the action distribution for each preference.

A *Bayesian strategy* is a mapping $\sigma : \Theta \to [0,1]$, where $\sigma(\theta)$ is the fraction of the agents with preference $\theta \in \Theta$ choosing action A. Let $\Sigma = \{\sigma : \Theta \to [0,1]\}$ denote the set of all Bayesian strategies that are measurable (with respect to $(\Theta, L(\Theta), \lambda)$, where $L(\Theta)$ is the set of all Lebesgue measurable sets on Θ). Recall that the distribution of preferences is described by the probability measure λ on Θ. Using the Bayesian strategy σ, the aggregate population state can be expressed as $\alpha(\sigma) = \int_\Theta \sigma(\theta) d\lambda$.

We define the k-sampling statistical inference (k-SI) best response correspon-

[4]See Chapter 7 of Phillips [2003] for the Bernstein polynomial. Salant and Cherry [2020] list some of its properties.

[5]When the preference distribution is uniform, we have $\Lambda(F_{k,z}) = F_{k,z}$, which gives the identical characterization of SESI in Salant and Cherry [2020].

dence as

$$B^{k,\theta}(\alpha) = \sum_{j=0}^{k} \binom{k}{j} \alpha^j (1-\alpha)^{k-j} B_{\mathcal{G}}^{\theta}\left(\frac{j}{k}\right).$$

$B^{k,\theta}(\alpha)$ is the probability of action A being the best response to the estimate when an agent with preference θ samples k agents from the aggregate populate state α.[6] Using $B^{k,\theta}(\cdot)$, we define a Bayesian SESI. A Bayesian SESI σ^* is the same as a SESI except that it offers a mathematical expression that is amenable to formal evolutionary analysis of individual strategy adjustment.

Definition 2. *A* Bayesian SESI *is a Bayesian strategy* $\sigma^* : \Theta \to [0,1]$ *such that*

$$\sigma^* = B^k(\alpha(\sigma^*)) \equiv \left\{B^{k,\theta}(\alpha(\sigma^*))\right\}_{\theta\in\Theta}$$

where $\alpha(\sigma^*) = \int_{\Theta} \sigma^*(\theta) d\lambda$. *In the above expression,* $\sigma^* = \{B^{k,\theta}(\alpha(\sigma^*))\}_{\theta\in\Theta}$ *is interpreted as* $\sigma^*(\theta) = B^{k,\theta}(\alpha(\sigma^*))$ *for all* $\theta \in \Theta$.

In words, σ^* is a strategy profile where every agent with inference procedure $\mathcal{G}$ best-responds to the estimate. Obviously, $\alpha(\sigma^*)$ is a SESI proportion if σ^* is a Bayesian SESI. The converse, however, does not necessarily hold. To see this, consider the following example.

Example 4. Assume that $f(\cdot)$ is continuous, convex, and increasing on $[0,1]$, and λ is the uniform distribution on $[0,1]$. Suppose σ^* is the unique Bayesian SESI (see Proposition 1 in Section 4.1 for the existence and the uniqueness) and $\alpha(\sigma^*)$ is the corresponding SESI proportion. Suppose the Bayesian strategy σ satisfies that for all agents with $\theta < \alpha(\sigma^*)$, they choose action A and for all agents with $\theta \geq \alpha(\sigma^*)$, they choose action B. σ is not a Bayesian SESI, but $\alpha(\sigma) = \alpha(\sigma^*)$.

Consider a Bayesian strategy with a SESI proportion as in Example 4. If all agents simultaneously update their strategy, then the Bayesian strategy will become a Bayesian SESI in the next period. It is, however, not obvious whether the Bayesian strategy converges to a Bayesian SESI if not all agents simultaneously update. Thus, Example 4 raises the following question. Would there be any difference in terms of convergence between the dynamic defined on the individual level

[6]The k-SI best response correspondence coincides with the sampling best response correspondence of Oyama et al. [2015] when the inference procedure is MLE in Example 1.

and the one defined on the aggregate level? We will answer this question in the next section.

3 Evolutionary dynamics

Now we turn to the dynamic version of the game. Every agent receives a revision opportunity with probability ε in each discrete period $t \in \{0, 1, \ldots\}$. When an agent receives a revision opportunity, she randomly samples k agents from the population, uses statistical inference to estimate the distribution of actions, and chooses a best response to the estimate.

Recall that $B^k(\alpha) = \{B^{k,\theta}(\alpha)\}_{\theta \in \Theta}$, and that every agent revises strategy with probability ε in each discrete time. The *Bayesian k-sampling statistical inference (k-SI) best response dynamic* (Bayesian dynamic, for short) is described by the following difference equation:

$$\sigma_t = (1-\varepsilon)\sigma_{t-1} + \varepsilon B^k(\alpha_{t-1}), \tag{5}$$

where $\alpha_{t-1} = \int_\Theta \sigma_{t-1}(\theta) d\lambda$.

The *aggregate dynamic* of the population, i.e. the dynamic of the fraction of action-A agents, is described as,

$$\alpha_t = \int_\Theta \sigma_t(\theta) d\lambda = (1-\varepsilon)\alpha_{t-1} + \varepsilon \int_\Theta B^{k,\theta}(\alpha_{t-1}) d\lambda. \tag{6}$$

Using the Bernstein polynomial in Eq.(3), the last term of Eq.(6) is written as follows.

$$\int_\Theta B^{k,\theta}(\alpha_{t-1}) d\lambda = \sum_{j=0}^{k} \binom{k}{j} \alpha_{t-1}^{j} (1-\alpha_{t-1})^{k-j} \int_\Theta B_{\mathcal{G}}^{\theta}\left(\frac{j}{k}\right) d\lambda = 1 - Bern_k(\alpha_{t-1}; \Lambda_{F,k}),$$

where recall that $\Lambda_{F,k}$ is such that $\Lambda_{F,k}(z) = \Lambda(F_{k,z})$. The aggregate dynamic (6) is rewritten as

$$\alpha_t = (1-\varepsilon)\alpha_{t-1} + \varepsilon(1 - Bern_k(\alpha_{t-1}; \Lambda_{F,k})). \tag{7}$$

This is equivalent to the dynamic presented in Salant and Cherry [2020].

Recall that $\alpha(\sigma) = \int_\Theta \sigma(\theta) d\lambda$. Define $\|\sigma\| = \int_\Theta |\sigma(\theta)| d\lambda$ for all $\sigma \in \mathbb{R}^\Theta$. Then $\|\sigma - \hat{\sigma}\| > 0$ for all $\sigma, \hat{\sigma} \in \Sigma$ with $\sigma \neq \hat{\sigma}$, and $\|\sigma - \hat{\sigma}\| = 0$ if and only if $\sigma = \hat{\sigma}$. In other words, $\|\cdot\|$ is the L^1 norm for a pair of Bayesian strategies. We say that σ_t converges to σ^* if $\lim_{t\to\infty} \|\sigma_t - \sigma^*\| = 0$; α_t converges to α^* if $\lim_{t\to\infty} |\alpha_t - \alpha^*| = 0$.

A Bayesian strategy that is not a Bayesian SESI can induce the SESI proportion as shown in Example 4. Hence, the convergence of the aggregate dynamic (6) to the SESI proportion does not guarantee the convergence of the Bayesian dynamic (5). The following lemma shows that if the aggregate dynamic is at a SESI proportion α^*, then the Bayesian strategy converges to $B^k(\alpha^*)$. Note that we relegate all the proofs to the Appendix.

Lemma 1. *A SESI proportion α^* is a stationary aggregate population state under the aggregate dynamic (6), that is, $\alpha_{t+1} = \alpha^*$ if $\alpha_t = \alpha^*$. If the aggregate dynamic is at α^*, then any Bayesian strategy σ with $\alpha(\sigma) = \alpha^*$ converges to $B^k(\alpha^*)$ under the Bayesian dynamic (5).*

More importantly, our first main result shows that the Bayesian strategy converges if and only if the aggregate population state converges.

Theorem 1. *The Bayesian strategy σ converges to σ^* under the Bayesian dynamic (5) if and only if the aggregate population state α converges to $\alpha^* = \alpha(\sigma^*)$ under the aggregate dynamic (6).*

Thus, the introduction of the Bayesian SESI and the Bayesian k-sampling statistical inference (k-SI) best response dynamic provides further supports for the analysis of the aggregate dynamic in Salant and Cherry [2020].

Next, we consider asymptotic stability of SESI. Though Theorem 1 implies that the convergence to a Bayesian SESI must coincide with the convergence to a SESI proportion, it does not tell whether the Bayesian dynamic (or the aggregate dynamic) ever converges. We will answer the question in the affirmative. For what follows, recall Observation 1 and let $K_{B(k)}$ be a Lipschitz constant for $Bern_k(\alpha; \Lambda_{F,k})$ with sample size k. To simplify the analysis, we assume that $(1-\alpha)$ and $Bern_k(\alpha; \Lambda_{F,k})$ are not tangent to each other and the intersection of the two curves contains at most finite points.

Lemma 2. *Let α^* be a SESI proportion. Fix an open interval $(\underline{\alpha}, \overline{\alpha})$ that satisfies the three conditions: i) $\alpha^* \in (\underline{\alpha}, \overline{\alpha})$, ii) $1-\alpha > Bern_k(\alpha; \Lambda_{F,k})$ for all $\alpha \in (\underline{\alpha}, \alpha^*)$, and iii) $1-\alpha <$*

$Bern_k(\alpha; \Lambda_{F,k})$ *for all* $\alpha \in (\alpha^*, \overline{\alpha})$. *The aggregate population state starting with some* $\alpha \in (\underline{\alpha}, \overline{\alpha})$ *converges to* α^* *for any* $\varepsilon \in (0, \overline{\varepsilon})$, *where* $\overline{\varepsilon} = \min\{1/K_{B(k)}, \alpha^* - \underline{\alpha}, \overline{\alpha} - \alpha^*\}$.

Recall Theorem 1, which implies that two different Bayesian strategies converge to the same strategy if their aggregate populations states converge to the same aggregate state. An implicaton of Lemma 2 is that all Bayesian strategies that have the same aggregate population state converge to the same Bayesian SESI regardless of any difference they may have. With Lemma 2 in hand, we can show global asymptotic stability of the set of Bayesian SESIs.

Theorem 2. *For any inference procedure, any sample size k, any preference distribution, and all sufficiently small* $\varepsilon > 0$, *the Bayesian strategy converges to some Bayesian SESI under the Bayesian dynamic (5).*

In the proof of Lemma 2 and Theorem 2, we use a discrete version of the Lyapunov stability theorem (see Theorem 7 in the Appendix) to show convergence. The implications of Lemma 2 and Theorem 2 include not only the global convergence to SESIs but also a characterization of stability of SESIs. The next example provides an illustration.

Example 5 (Convergence and (in)stability of SESIs)**.** Consider the uniformly distributed λ on $[0, 1]$ and the two continuous cost functions f_1 and f_2 below:

$$f_1(\alpha) = \alpha^2,$$
$$f_2(\alpha) = \max\{0.9 - 1.3\alpha^2, 0\}.$$

f_1 is a typical convex cost function, while f_2 is a decreasing cost function. The latter represents settings where action A has positive externality. The initial adoption cost is high, but the cost decreases as more agents adopt A. The left of Figure 1 shows the graphs of $1 - \alpha$, $Bern_k(\alpha; \Lambda_{F,k})$ with MLE, and $Bern_k(\alpha; \Lambda_{F,k})$ with OLP, with the cost function f_1. The right of Figure 1 shows the graphs of $1 - \alpha$, $Bern_k(\alpha; \Lambda_{F,k})$ with MLE, with the cost function f_2. Each intersection of $1 - \alpha$ and $Bern_k(\alpha; \Lambda_{F,k})$ represents a SESI proportion. The sample size is 10 for both settings.

Theorem 2 implies that the Bayesian strategy converges to some Bayesian SESI from any initial state. For the setting with f_1, this means that the aggregate population state globally converges to the unique SESI proportion for each statistical inference. For the setting with f_2, a SESI proportion to which the aggregate

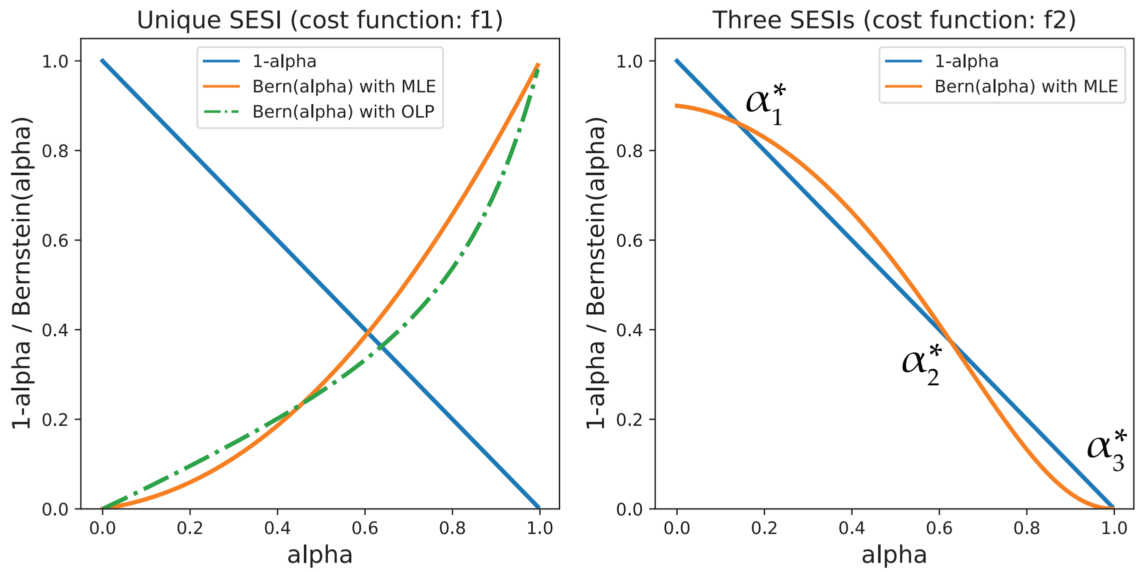

Figure 1: $(1-\alpha)$ and $Bern_k(\alpha; \Lambda_{F,k})$ with function f_1 (left) and with f_2 (right)

population state converges depends on the initial state. Lemma 2, however, tells that some SESIs may be unstable. To see this, let α_1^*, α_2^*, and α_3^* denote the SESIs ($\alpha_1^* < \alpha_2^* < \alpha_3^* = 1$). According to Lemma 2, the aggregate population state converges to α_1^* if the initial state is in $[0, \alpha_2^*)$, while it converges to α_3^* if the initial state is in $(\alpha_2^*, 1]$. The middle SESI α_2^* is unstable, that is, the aggregate population state starting from any close neighborhood of α_2^* moves away from α_2^*. Any small perturbation to α_2^* will move the state to a different SESI.

Note also that, in the setting with f_2, the total utility of agents is maximized in α_3^* because the adoption cost is zero in α_3^* and thus all agents receive positive utility. However, if the society starts with a low adoption rate, it will be trapped in an inefficient SESI α_1^*.

In Section 4, we discuss a sufficient condition on the cost function which guarantees the uniqueness of SESI, and an implication of our results on social welfare. We also discuss a tax scheme improving social welfare. In Section 5, we characterize a class of games where (almost) global convergence to a SESI is obtained even if there are more than one SESI. The analysis of global convergence is extended to determine diffusion in network games in Section 6.1.

4 Anti-coordination games

4.1 Equilibrium existence and uniqueness

We consider *anti-coordination games*, which are games where $f(\cdot)$ is weakly increasing on $[0,1]$ and $0 \leq f(0) < f(1) \leq 1$. The game of this type is studied in Salant and Cherry [2020]. It captures the nature of anti-coordination — agents have an incentive to differ from others. Evolutionary dynamics in such games are studied by Kojima and Takahashi [2007], Zhang [2016], and Zusai [2017], among others. We first extend some results of Salant and Cherry [2020] for any preference distribution, and then discuss a tax scheme that improves social welfare.

Roughly speaking, the Bernstein polynomial $Bern_k(\alpha; \Lambda_{F,k})$ is the fraction of agents whose best response to the estimate is action B when the fraction of agents choosing A is α. This implies that the Bernstein polynomial is weakly increasing in α in anti-coordination games, and further implies the next proposition.

Proposition 1. *In an anti-coordination game, for any inference procedure, any sample size, and any preference distribution, there exists a unique SESI.*

The next corollary is immediate from Lemma 2 and Theorem 2. Its implication on the Bayesian strategy and social welfare is illustrated in Example 6.

Corollary 1. *Let α^* be the unique SESI proportion in the anti-coordination game for a given tuple of the inference procedure, sample size k, and preference distribution. Then, under the Bayesian dynamic (5) for any $\varepsilon \in (0, 1/K_{B(k)})$, the aggregate population state converges to α^*, and the Bayesian strategy converges to $B^k(\alpha^*)$.*

Example 6. Consider the game with the cost function f_1 in Example 5. The utility from action B is 0, the utility from action A is $u(\theta, \alpha) = \theta - \alpha^2$, and the preference θ is uniformly distributed on $[0,1]$. Corollary 1 guarantees the convergence of the Bayesian strategy to $B^k(\alpha^*)$. Furthermore, Lemma 1 implies that the aggregate population state never deviates from α^* once it is reached. That is, if α_t reaches α^*, then the Bayesian strategy σ_t will approach $B^k(\alpha^*)$ without changing the aggregate population state thereafter.

Figure 2 illustrates the dynamics of the aggregate population state and Bayesian strategy for agents using MLE. Initially, the fraction of agents choosing A is the SESI proportion (≈ 0.6) for all preferences. The left figure shows the dynamics of

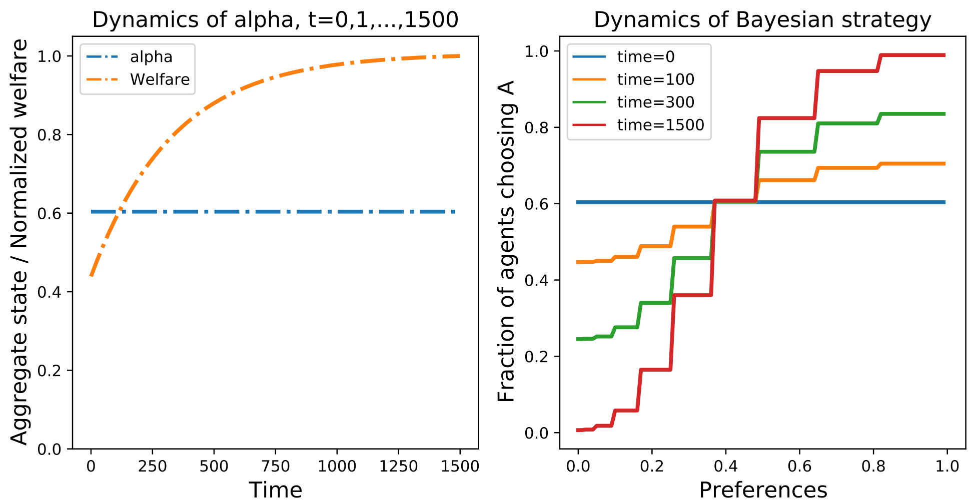

Figure 2: Evolution of the aggregate state (left) and Bayesian strategy (right) with MLE. The sample size is 10, and the revision probability is $\varepsilon = 3 \times 10^{-3}$.

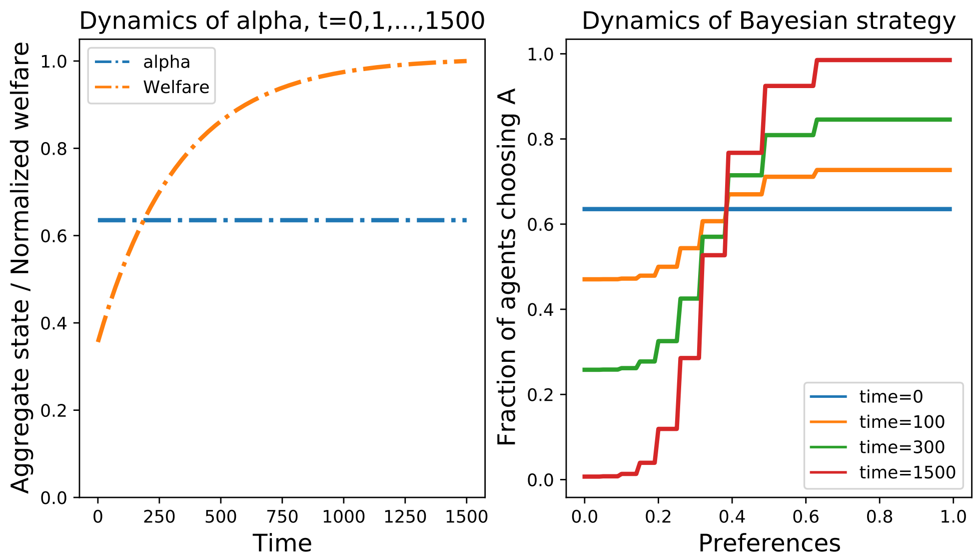

Figure 3: Evolution of the aggregate state (left) and Bayesian strategy (right) with OLP ($\delta = 0.6$). The sample size is 10, and the revision probability is $\varepsilon = 3 \times 10^{-3}$.

the aggregate population state α and the sum of agents' payoffs normalized to that sum in the Bayesian SESI, which measures the social welfare of a state relative to the Bayesian SESI. The right figure shows the dynamics of the Bayesian strategies, where agents' preferences are on the horizontal axis and the fraction of action-

A agents for each preference is on the vertical axis. As illustrated in Figure 2, the aggregate state does not fluctuate. The agents, however, frequently switch their actions until they reach the Bayesian SESI, which is roughly illustrated by the graph at time $t = 1500$ in the right figure. The action distribution is becoming more diverse across preferences over time. This suggests that the aggregate population state may not be a good indicator for welfare. Even if it is in the steady state, the social welfare may be far from the one in equilibrium as shown in Figure 2.

Figure 3 illustrates the dynamics for the agents using OLP with $\delta = .6$. The parameters are the same as the ones with MLE, except that the initial fraction of agents choosing A is the SESI proportion under OLP (≈ 0.63) for all preferences. We have an observation similar to MLE. Recall that agents with OLP moderately evaluate objective probabilities, that is, their subjective probabilities are closer to 50% than the objective probabilities are. The action distribution with OLP in the Bayesian SESI is more polarized than the action distribution with MLE. For example, with OLP, the agents with $\theta > .6$ choose A with probability close to one. In contrast, with MLE, only the agents with $\theta > .8$ do so. The reason is that, for example, if the agents with OLP draw a sample with only one agent choosing A, they think that the fraction of agents choosing A in the population is around 0.211 instead of 0.1 (see Example 2). This makes more agents with higher θ choose A.

4.2 On improving social welfare

We discuss that some intervention by the social planner may improve the social welfare. We restrict our attention to simpler settings — anti-coordination games with a fixed fee on the choice of action A. The revenue from the fee is equally distributed to all agents, and thus is not considered as a welfare loss. The intervention we consider is akin to an evolutionary implementation, e.g. Lahkar and Mukherjee [2019, 2021], Sandholm [2002, 2005, 2007], in the sense that a simple price scheme in evolutionary dynamics may lead to welfare improvements.[7]

We assume that agents pay a tax $c \in \mathbb{R}_+$ to choose action A, and that the cost function f is continuously differentiable. Observe that if an agent has to pay c to

[7] Differences from the above mentioned papers are that agents do not perfectly observe the aggregate population state and that there is an infinite set of types of agents.

choose A, the best response function $b^{\theta}_{\mathcal{G}}(k,z)$ can be written as

$$b^{\theta}_{\mathcal{G}}(k,z|c) = \underset{s\in\mathcal{S}}{\operatorname{argmax}} \int_{\alpha\in\mathcal{A}} F^{\theta}_{s}(\alpha) d\mathcal{G}_{k,z}(\alpha) - c\mathbb{1}_{[s=A]} \qquad \forall z \in \left\{0, \frac{1}{k} \dots, 1\right\},$$

where $\mathbb{1}_{[s=A]}$ is 1 when $s = A$ and 0 otherwise, and $b^{\theta}_{\mathcal{G}}(k,z|c) = B$ when there is a tie. The index function $B^{\theta}_{\mathcal{G}}(z|c)$ is defined the same way as $B^{\theta}_{\mathcal{G}}(z)$. Recall that $B^{k,\theta}(\alpha)$ is the probability of action A being the best response to the estimate when an agent with preference θ samples k agents from the aggregate populate state α. We define a similar probability with tax c as

$$B^{k,\theta}(\alpha|c) = \sum_{j=0}^{k} \binom{k}{j} \alpha^{j}(1-\alpha)^{k-j} B^{\theta}_{\mathcal{G}}\left(\frac{j}{k}\middle| c\right).$$

Definition 3. *A* Bayesian SESI *with tax c is a Bayesian strategy* $\sigma^{*}_{c} : \Theta \to [0,1]$ *such that* $\sigma^{*}_{c} \equiv \{B^{k,\theta}(\alpha(\sigma^{*}_{c})|c)\}_{\theta\in\Theta}$*, where recall that* $\alpha(\sigma^{*}_{c}) = \int_{\Theta} \sigma^{*}_{c}(\theta) d\lambda$.

Let σ^{*}_{c} be a Bayesian SESI with tax c, and σ^{*}_{0} be the one without a tax. We define the social welfare of the game and the set of welfare maximizers as below.

$$SW(\sigma) = \int_{\theta\in\Theta} [\theta - f(\alpha(\sigma))]\, \sigma(\theta) d\lambda,$$
$$C^{*} = \underset{c}{\operatorname{argmax}}\, SW(\sigma^{*}_{c}).$$

The social welfare $SW(\cdot)$ is simply the total utility of agents. C^{*} is the set of taxes that maximize the social welfare in the SESI. C^{*} always exists (Lemma 4 in the Appendix). The next theorem shows (i) the uniqueness and convergence result, and (ii) a condition under which an intervention can improve the social welfare. Claim (i) shows that a tax c dynamically implements σ^{*}_{c}. Thus, if a tax in C^{*} is imposed on A, the behavior will converge to a state that maximizes the social welfare. Claim (ii) gives a simple condition for when the social planner should intervene. This condition only depends on the aggregate population state and does not require the social planner to inspect the strategy distribution over preferences.

Theorem 3. *There is a unique Bayesian SESI* σ^{*}_{c} *for every* $c \in \mathbb{R}_{+}$*. The Bayesian strategy converges to* σ^{*}_{c} *under the Bayesian dynamic (5) for all sufficiently small* ε*. If* $1 - f(\alpha(\sigma^{*}_{0})) - \alpha(\sigma^{*}_{0}) f'(\alpha(\sigma^{*}_{0})) < 0$*, then* $c^{*} > 0$ *for all* $c^{*} \in C^{*}$.

5 Coordination games

In this section, we consider a class of games that resemble coordination games with two strict Nash equilibria (NEs). The equilibrium selection problem in coordination games lies at the heart of evolutionary game theory (see Foster and Young [1990], Young [1993] and Kandori et al. [1993], for example). We investigate if our Bayesian dynamic has any equilibrium selection property.

Salant and Cherry [2020] only focus on games with interior NEs because in these games SESIs differ from NEs. When there exists a corner NE, where every agent is choosing the same action, the corner NE must be a SESI if the inference procedure $\mathcal{G}$ is *unbiased* as defined in Salant and Cherry [2020]:

Definition 4. *An inference procedure $\mathcal{G}$ is unbiased if for any sample (k, z), the expected value of the estimate $\mathcal{G}_{k,z}$ is equal to the sample mean, that is,*

$$\int_0^1 \alpha d\mathcal{G}_{k,z}(\alpha) = z, \text{ for any sample } (k, z).$$

The rationale is that when $\mathcal{G}$ is unbiased, if an agent's sample includes only failures $(z = 0)$ or successes $(z = 1)$, the agent concentrates his estimate on the sample mean.[8]

We assume that f is weakly decreasing and we relax the range of f by assuming that $f(0) > 1$ and $f(1) < 0$. Given these assumptions, both $\alpha^* = 0$, $\alpha^{**} = 1$ are strict NEs, and they are also SESIs given unbiased $\mathcal{G}$. Nevertheless, we will show that the evolutionary dynamic has a sharp equilibrium selection result. To begin, we define the *p-dominance* in the spirit of Morris et al. [1995]. To avoid repetition, we focus our analysis on action B being p-dominant. The analysis for action A being p-dominant is similar.[9]

Definition 5. *Action $B \in S$ is p-dominant in F if $b^\theta(\alpha) = \{B\}$ for all $\theta \in \Theta$ and all $\alpha \leq 1 - p$. A Bayesian strategy where all agents play the same p-dominant action is called a p-dominant equilibrium.*

[8]Note that the MLE inference procedure in Example 1 and the TN inference procedure in Example 3 are both unbiased, while the OLP inference procedure in Example 2 is not.

[9] A symmetric result holds for action A. If A is $1/k$-dominant, the sample size ranges from 2 to k, $\mathcal{G}$ is unbiased, and f is concave, then the aggregate dynamic (6) converges to $\alpha^{**} = 1$ for all $\alpha_0 \neq 0$.

One can observe that it is an extension of the one in Morris et al. [1995] to population games, which is also used in Sandholm [2001] and Oyama et al. [2015]. A similar extended definition on compact action sets is used in Tercieux [2006]. Given the game structure, we have an immediate observation:

Observation 2. *Action B is p-dominant if and only if $f(1-p) > 1$.*

We have the following selection result:

Theorem 4. *Assume that B is $1/k$-dominant for some integer $k > 1$ and the sample size $l \in \{2, ..., k\}$. If f is convex and $\mathcal{G}$ is unbiased, for any initial proportion $\alpha_0 \neq 1$, the aggregate population state converges to $\alpha^* = 0$ under the aggregate dynamic (6).*[10]

Theorem 4 shows that the evolutionary dynamic almost globally converges to the $1/k$-dominant equilibrium, which generalizes the selection result of Sandholm [2001] and Oyama et al. [2015] to games with heterogeneous preferences and a broad class of statistical inferences. In contrast to the equilibrium selection results in stochastic stability analysis which generally rely on long waiting times, our convergence result happens within finite periods. The result demonstrates that in coordination games beyond the consideration of Salant and Cherry [2020], even though both strict NEs are SESIs in the static setting, statistical decision making still has important dynamic implications.

6 Discussion

6.1 Application: statistical inference in network games

Our evolutionary dynamics can be applied to diffusion processes in network games with agents using statistical inference. Following Jackson and Yariv [2007], we model a large social network of a unit-mass population of agents through the distribution of the number of neighbors. The fraction of agents with k neighbors is given by the degree distribution γ_k, where $\sum_{k \in \mathbf{K}} \gamma_k = 1$ for $\mathbf{K} = \{1, \ldots, K\}$. We call an agent with k neighbors a degree-k agent. Let λ_k denote the probability measure

[10] Theorem 1 guarantees the convergence of the Bayesian dynamic (5) to the corresponding Bayesian SESI for $\alpha^* = 0$.

on the preference set $\Theta = [0, 1]$ for degree-k agents. That is, $\gamma_k = \int_\Theta d\lambda_k$.[11] The game is similar to the one in Section 2. Agents choose an action from $\mathcal{S} = \{A, B\}$. The utility from B is 0. The utility from A is $u(\theta, \rho) = \theta - f(\rho)$, where θ is the agent's preference, and $f(\rho)$ is the cost of choosing action A when the probability that a neighbor chooses A is ρ.[12] f is continuous on $[0, 1]$.

In each discrete period $t \in \{0, 1, \dots\}$, a social network is randomly formed according to the degree distribution $\{\gamma_k\}_{k \in \mathbf{K}}$, and then the agents receive a revision opportunity with probability ε.[13] Each revising agent observes their neighbors' actions, estimates the probability ρ using statistical inference, and best-responds to the estimate. For a revising degree-k agent, let z denote the mean of action A of their neighbors. An *inference procedure* $\mathcal{G} = \{\mathcal{G}_{k,z}\}$ assigns a cumulative distribution function $\mathcal{G}_{k,z}$ to every observation (k, z) such that $\mathcal{G}_{k,\hat{z}}$ strictly first-order stochastically dominates $\mathcal{G}_{k,z}$ when $\hat{z} > z$.

The pure best response correspondence of agents with preference θ, inference procedure $\mathcal{G}$, and observation (k, z) is defined as

$$b^\theta_{\mathcal{G}}(k, z) = \operatorname*{argmax}_{s \in \{A,B\}} \int_{\rho \in [0,1]} F^\theta_s(\rho) d\mathcal{G}_{k,z}(\rho) \qquad \forall z \in \{0, 1/k, \dots, 1\},$$

where $F^\theta_A(\rho) = \theta - f(\rho)$, and $F^\theta_B(\rho) = 0$. Agents choose action B when there is a tie. Let $B^\theta_{\mathcal{G}}(z)$ be an index function that becomes one if action A is the best response to the estimate for the agents with preference θ and mean z.

Let $\sigma^k_t(\theta)$ be the fraction of degree-k agents with preference θ choosing action A in period t. Define the weighted mean of action A as

$$\rho_t = \sum_k \frac{k}{\bar{k}} \int_\Theta \sigma^k_t(\theta) d\lambda_k, \quad \text{where} \quad \bar{k} = \sum_k \gamma_k k.$$

ρ_t is the probability that an agent chooses action A if we randomly choose a link and pick the agent at either end of the link in period t.

[11] We assume that the degree distribution does not differ across preferences, but the analysis should be extended to more general degree distributions using the extended model in Section 6.2.

[12] The utility function follows one of the utility functions considered in Jackson and Yariv [2007]. The agents gather information from close friends/relatives (to make an estimate), and their payoff depends on the average play of their neighbors.

[13] The set of k-degree agents does not differ over time for all $k \in \mathbf{K}$, that is, agents have k neighbors in period t if they have k neighbors in period $t - 1$.

Following Jackson and Yariv [2007], we write the mean-field dynamics for $\sigma_t^k(\theta)$ as

$$\sigma_t^k(\theta) = (1-\varepsilon)\sigma_{t-1}^k(\theta) + \varepsilon \sum_{j=0}^{k} \binom{k}{j} (\rho_{t-1})^j (1-\rho_{t-1})^{k-j} B_{\mathcal{G}}^{\theta}\left(\frac{j}{k}\right) \quad \forall \theta \in \Theta.$$

See also Section 5.1.3 of Jackson and Zenou [2015] for the mean-field dynamics of diffusion on networks. In network games, the probability that a neighbor chooses action A in period $t-1$ is ρ_{t-1} instead of the fraction of agents choosing A in period $t-1$. The equation can be viewed as the dynamics of Eq.(5) by replacing α_{t-1} with ρ_{t-1}. Our evolutionary dynamics can be applied with this slight modification.

Recall that for all $z \in [0,1]$,

$$\int_{\Theta} B_{\mathcal{G}}^{\theta}(z) d\lambda_k = \gamma_k - \int_0^{F_{k,z}} d\lambda_k,$$

where $F_{k,z} = \int_{\alpha \in \mathcal{A}} f(\alpha) d\mathcal{G}_{k,z}(\alpha)$, the expected cost of action A when observing a sample (k,z). Note that the first term in the right-hand side is γ_k since the fraction of k-degree agents is γ_k. Then the dynamics of ρ_t can be written as

$$\begin{aligned}
\rho_t &= \sum_{k \in \mathbf{K}} \frac{k}{\bar{k}} \int_{\Theta} \sigma_t^k(\theta) d\lambda_k \\
&= (1-\varepsilon)\rho_{t-1} + \varepsilon \sum_{k \in \mathbf{K}} \frac{k}{\bar{k}} \sum_{j=0}^{k} \binom{k}{j} (\rho_{t-1})^j (1-\rho_{t-1})^{k-j} \int_{\Theta} B_{\mathcal{G}}^{\theta}\left(\frac{j}{k}\right) d\lambda_k \\
&= (1-\varepsilon)\rho_{t-1} + \varepsilon \sum_{k \in \mathbf{K}} \frac{k}{\bar{k}} (\gamma_k - Bern_k(\rho_{t-1}; \Lambda_{F,k})), \qquad (8)
\end{aligned}$$

where $\Lambda_{F,k} = \Lambda_k(F_{k,z}) \equiv \int_0^{F_{k,z}} d\lambda_k$. Eq.(8) shows that the diffusion process can be expressed by a dynamic process with one state variable ρ_t.

Example 7 (k-regular networks)**.** Consider that the network is a k-regular graph, that is, $\gamma_k = 1$ for some $k \geq 2$. Then, $\bar{k} = k$, and Eq.(8) is reduced to

$$\rho_t = (1-\varepsilon)\rho_{t-1} + \varepsilon(1 - Bern_k(\rho_{t-1}; \Lambda_{F,k})).$$

This is the same as Eq.(7). For k-regular networks, the probability that a neighbor chooses action A is the same as the fraction of agents choosing A, i.e. $\rho_t = \alpha_t$.

The next proposition characterizes a sufficient condition under which action B spreads in the entire network with agents using statistical inference. In this setting, action B is said to be *p-dominant* if $F_B^\theta(\rho) > F_A^\theta(\rho)$ for all $\theta \in \Theta$ and all $\rho \leq 1 - p$. Note that by definition, $\rho_t = 0$ implies that $\sigma_t^k(\theta) = 0$ for almost all $\theta \in \Theta$.

Proposition 2. *Assume that B is $1/K$-dominant for some integer $K > 1$. If f is convex, $\mathcal{G}$ is unbiased, and the set of degrees is $\mathbf{K} \subseteq \{2, \ldots, K\}$, then for any initial state $\rho_0 \neq 1$ and any degree distribution $\{\gamma_k\}_{k \in \mathbf{K}}$, the state converges to $\rho^* = 0$ under the diffusion process expressed by Eq.(8).*

The intuition is similar to that of Theorem 4, and thus we omit the formal proof. When B is $1/K$-dominant and all conditions in the proposition are met, $Bern_k(\rho_{t-1}; \Lambda_{F,k}) = \gamma_k(1 - \rho_{t-1}^k)$. Then, $(\gamma_k - Bern_k(\rho_{t-1}; \Lambda_{F,k}))$ is reduced to $\gamma_k \rho_{t-1}^k$ for all $k \in \mathbf{K}$. Eq.(8) guarantees that $\rho_t < \rho_{t-1}$ for any $\rho_{t-1} \in (0, 1)$.

6.2 Heterogeneity

So far, we have assumed that all agents adopt the same statistical inference procedure. We relax the assumption and consider heterogeneous inference procedures and samples sizes in this section. Let $\mathbf{G} = \{\mathcal{G}^1, \ldots, \mathcal{G}^M\}$ denote the set of inference procedures, and $\mathbf{K} = \{1, \ldots, K\}$ denote the set of sample sizes. Let λ_{ik} denote the probability measure on the preference set $\Theta = [0, 1]$ for agents who make an estimate by sampling k agents and using inference procedure $\mathcal{G}^i$, and let γ_{ik} denote the mass of such agents. That is, $\gamma_{ik} = \int_\Theta d\lambda_{ik}$ and $\sum_{(\mathcal{G}^i, k) \in \mathbf{G} \times \mathbf{K}} \gamma_{ik} = 1$.

We extend the notations in Sections 2 and 3 to this setting. $\mathcal{G}_{k,z}^i$ denotes the cumulative distribution assigned by $\mathcal{G}^i$ for sample (k, z), where recall that k is a sample size and z is a sample mean. The pure best response correspondence and the best response index function are written as

$$b_{\mathcal{G}^i}^\theta(k, z) = \operatorname*{argmax}_{s \in \mathcal{S}} \int_{\alpha \in \mathcal{A}} F_s^\theta(\alpha) d\mathcal{G}_{k,z}^i(\alpha) \qquad \forall z \in \left\{0, \frac{1}{k} \ldots, 1\right\},$$

$$B_{\mathcal{G}^i}^{\theta,k}(z) = \begin{cases} 1 & \text{if } b_{\mathcal{G}^i}^\theta(k, z) = A, \\ 0 & \text{if } b_{\mathcal{G}^i}^\theta(k, z) = B. \end{cases}$$

$B_{\mathcal{G}^i}^{\theta,k}(z)$ is an index function that becomes one if A is the best response to the estimate for the agents with preference θ, inference procedure $\mathcal{G}^i$, sample size k, and sample mean z. A SESI proportion $\alpha_{\mathbf{GK}}$ for this setting can be characterized as follows.

$$\alpha_{\mathbf{GK}} = \sum_{(\mathcal{G}^i,k)\in\mathbf{G}\times\mathbf{K}} \int_{\theta\in\Theta} \sum_{j=0}^{k} \binom{k}{j} \alpha_{\mathbf{GK}}^{j} (1-\alpha_{\mathbf{GK}})^{k-j} B_{\mathcal{G}^i}^{\theta,k}\left(\frac{j}{k}\right) d\lambda_{ik}. \tag{9}$$

Let $F_{k,z}^{i}$ denote the expected cost f under $\mathcal{G}_{k,z}^{i}$, that is, $F_{k,z}^{i} = \int_{\alpha\in\mathcal{A}} f(\alpha) d\mathcal{G}_{k,z}^{i}(\alpha)$. Observe that

$$\int_{\theta\in\Theta} B_{\mathcal{G}^i}^{\theta,k}(z) d\lambda_{ik} = \int_{\theta\in\Theta} \mathbb{1}\left[\int_{\alpha\in\mathcal{A}} F_{A}^{\theta}(\alpha) d\mathcal{G}_{k,z}^{i}(\alpha) > 0\right] d\lambda_{ik} = \gamma_{ik} - \Lambda_{ik}(F_{k,z}^{i}),$$

where $\Lambda_{ik}(x) = \int_0^x d\lambda_{ik}$. Eq.(9) is written as follows.

$$\alpha_{\mathbf{GK}} = \sum_{(\mathcal{G}^i,k)\in\mathbf{G}\times\mathbf{K}} \sum_{j=0}^{k} \binom{k}{j} \alpha_{\mathbf{GK}}^{j} (1-\alpha_{\mathbf{GK}})^{k-j} (\gamma_{ik} - \Lambda_{ik}(F_{k,j/k}^{i})).$$

Recall the Bernstein polynomial in Eq.(3). $\alpha_{\mathbf{GK}}$ can be rewritten as follows.[14]

$$1 - \alpha_{\mathbf{GK}} = \sum_{(\mathcal{G}^i,k)\in\mathbf{G}\times\mathbf{K}} Bern_k(\alpha_{\mathbf{GK}}; \Lambda_{F,k}^{ik}), \tag{10}$$

where $\Lambda_{F,k}^{ik}$ is a function such that $\Lambda_{F,k}^{ik}(z) = \Lambda_{ik}(F_{k,z}^{i})$. A Bayesian SESI for this setting is defined as follows.

Definition 6. *A* Bayesian SESI *is a Bayesian strategy* $\sigma^* : \Theta \to [0,1]$ *such that*

$$\sigma^* = B^{\mathbf{GK}}(\alpha(\sigma^*)) \equiv \left\{ \sum_{j=0}^{k} \binom{k}{j} \alpha(\sigma^*)^j (1-\alpha(\sigma^*))^{k-j} B_{\mathcal{G}^i}^{\theta,k}\left(\frac{j}{k}\right) \right\}_{\theta\in\Theta, \mathcal{G}^i\in\mathbf{G}, k\in\mathbf{K}},$$

where $\alpha(\sigma^*) = \sum_{(\mathcal{G}^i,k)\in\mathbf{G}\times\mathbf{K}} \int_{\Theta} \sigma^*(\theta) d\lambda_{ik}$*. Note that* $\alpha(\sigma^*)$ *must be a SESI proportion.*

[14] Eq.(10) is slightly different from Eq.(8) in Salant and Cherry [2020] since we allow the preference distribution to be different among agents who adopt different inference procedures or sample sizes. If we assume the common preference distribution λ, i.e. $\lambda(\theta) = \lambda_{ik}(\theta)/\gamma_{ik}$ for all $(\mathcal{G}^i, k) \in \mathbf{G} \times \mathbf{K}$, then we will have a similar expression to theirs.

The aggregate and Bayesian dynamics can be written as

$$\alpha_t = (1-\varepsilon)\alpha_{t-1} + \varepsilon \left(1 - \sum_{\mathcal{G}^i, k \in \mathbf{G}\times\mathbf{K}} Bern_k(\alpha_{t-1}; \Lambda^{ik}_{F,k}) \right), \tag{11}$$

$$\sigma_t = (1-\varepsilon)\sigma_{t-1} + \varepsilon B^{\mathbf{GK}}(\alpha_{t-1}). \tag{12}$$

Theorems 1 and 2 can be extended to this setting. Our convergence results are robust with respect to the heterogeneity of inference procedures, sample sizes, and preference distributions (associated with each pair of inference procedure and sample size).

Theorem 5. *For any set of inference procedures* $\mathbf{G}$*, any set of sample sizes* $\mathbf{K}$*, and any set of preference distributions* $\{\lambda_{ik}\}$*, the Bayesian strategy* σ *converges to* σ^* *under the Bayesian dynamic (12) if and only if the aggregate population state* α *converges to* $\alpha^* = \alpha(\sigma^*)$ *under the aggregate dynamic (11).*

Theorem 6. *For any set of inference procedures* $\mathbf{G}$*, any set of sample sizes* $\mathbf{K}$*, any set of preference distributions* $\{\lambda_{ik}\}$*, and all sufficiently small* $\varepsilon > 0$*, the Bayesian strategy converges to some Bayesian SESI under the Bayesian dynamic (12).*

7 Conclusion

We propose a novel evolutionary model incorporating statistical inference, show global convergence to SESIs, and apply the analysis to several important settings. We conclude by listing several future research directions. A major limitation of this study is the assumption of two actions. It will be interesting to consider multi-action games. For network games, it has been reported that the message-passing method describes diffusion on networks more accurately (Gleeson and Porter 2018, Kobayashi and Onaga 2021). Adopting the message passing method with agents using statistical inference may produce some new insights to the diffusion of behavior on networks.

Appendix: Proofs

Proof of Lemma 1. Let α^* denote a SESI proportion. Note that $\int_\Theta B^{k,\theta}(\alpha^*)d\lambda = \alpha^*$ by definition of the SESI proportion. Suppose $\sigma_0 \in \Sigma$ with $\alpha(\sigma_0) = \alpha^*$. Eq.(7) implies that the aggregate population state in period 1 is that $\alpha_1 = (1-\varepsilon)\alpha^* + \varepsilon(1 - Bern_k(\alpha^*; \Lambda_{F,k})) = \alpha^*$. Thus, it is stationary. The Bayesian strategy in period 1 is computed using Eq.(5),

$$\sigma_1 = (1-\varepsilon)\sigma_0 + \varepsilon B^k(\alpha^*).$$

This implies that $\sigma_t = (1-\varepsilon)^t\sigma_0 + (1-(1-\varepsilon)^t)B^k(\alpha^*)$ for all $t \in \{1, 2, \ldots\}$. The claim follows. □

Observation 3. $\|\sigma - \hat{\sigma}\| \geq |\alpha(\sigma) - \alpha(\hat{\sigma})|$, *for all* $\sigma, \hat{\sigma} \in \mathbb{R}^\Theta$.

Proof. Recall that $\alpha(\sigma) = \int_\Theta \sigma(\theta)d\lambda$, and that $\|\sigma\| = \int_\Theta |\sigma(\theta)|d\lambda$ for all $\sigma \in \mathbb{R}^\Theta$. Because $|w-x| + |y-z| \geq |w+y-(x+z)|$ for all $w, x, y, z \geq 0$, we have

$$\|\sigma - \hat{\sigma}\| = \int_\Theta |\sigma(\theta) - \hat{\sigma}(\theta)|\, d\lambda \geq \left| \int_\Theta \sigma(\theta)d\lambda - \int_\Theta \hat{\sigma}(\theta)d\lambda \right| = |\alpha(\sigma) - \alpha(\hat{\sigma})|.$$

□

The next lemma shows that the Bayesian best response correspondence $B^k(\cdot)$ is Lipschitz continuous with respect to the L^1 norm. This lemma together with Observation 3 is used in the proof of Theorem 1.

Lemma 3. $B^k(\cdot)$ *is Lipschitz continuous on* $\mathcal{A}$, *that is, there exists* $K > 0$ *such that* $\|B^k(\alpha) - B^k(\beta)\| < K|\alpha - \beta|$ *for all* $\alpha, \beta \in \mathcal{A}$.

Proof of Lemma 3. Fix $\alpha, \beta \in \mathcal{A}$. Observe that

$$\begin{aligned}\|B^k(\alpha) - B^k(\beta)\| &= \int_\Theta |B^{k,\theta}(\alpha) - B^{k,\theta}(\beta)|d\lambda \\ &= \int_\Theta \left| \sum_{j=0}^{k} \binom{k}{j} (\alpha^j(1-\alpha)^{k-j} - \beta^j(1-\beta)^{k-j}) B_G^\theta\left(\frac{j}{k}\right) \right| d\lambda.\end{aligned}$$

Recall that $B_G^\theta(j/k) \in \{0, 1\}$, i.e. bounded by 1. Note also that there exists some

$\hat{K} > 0$ such that

$$|\alpha^j(1-\alpha)^{k-j} - \beta^j(1-\beta)^{k-j})| \leq \hat{K}|\alpha - \beta| \qquad \forall \alpha, \beta \in \mathcal{A} = [0,1].$$

This is because a function $g(x) = x^j(1-x)^{k-j}$ is a polynomial function on the closed interval $[0,1]$ and thus is Lipschitz continuous. We can bound the norm as follows.

$$\|B^k(\alpha) - B^k(\beta)\| \leq \int_\Theta \sum_{j=0}^{k} \binom{k}{j} \hat{K}|\alpha - \beta| d\lambda < K|\alpha - \beta|,$$

where $K > \hat{K} \sum_{j=0}^{k} \binom{k}{j}$. □

Proof of Theorem 1. We first prove the 'only if' part. Fix $\eta > 0$. Suppose that the Bayesian strategy converges to σ^* under the Bayesian best response dynamic (5). Let σ_0 be the Bayesian strategy in period 0. There exists some $T > 0$ such that $\| \sigma_t - \sigma^* \| < \eta$ for all $t > T$. Then, $\eta > \| \sigma_t - \sigma^* \| \geq |\alpha(\sigma_t) - \alpha^*|$ for all $t > T$, by Observation 3.

Next, we prove the 'if' part. Suppose that the aggregate population state converges to α^*. Let $\sigma^* = B^k(\alpha^*)$. We show that the Bayesian strategy converges to σ^*. Fix $\eta > 0$. It suffices to show that there exists $T > 0$ such that $\| \sigma_t - \sigma^* \| < \eta$ for all $t > T$.

By Lemma 3, B^k is Lipschitz continuous. Let K denote the Lipschitz constant for B^k. Choose τ_1 such that $\|B^k(\alpha_t) - B^k(\alpha^*)\| < K|\alpha_t - \alpha^*| < \eta/2$ for all $t > \tau_1$. Since the aggregate population state converges to α^*, such τ_1 exists. Choose τ_2 such that $(1-\varepsilon)^t \| \sigma - \hat{\sigma} \| < \eta/2$ for all $\sigma, \hat{\sigma} \in \Sigma$ and all $t > \tau_2$. Since $\|\sigma - \hat{\sigma}\|$ is bounded, such τ_2 exists.

Fix $T > \tau_1 + \tau_2$. Recall Eq.(5), that is, $\sigma_t = (1-\varepsilon)\sigma_{t-1} + \varepsilon B^k(\alpha_{t-1})$. Observe that for all $t > T$,

$$\begin{aligned}
\|\sigma_t - \sigma^*\| &= \|(1-\varepsilon)\sigma_{t-1} + \varepsilon B^k(\alpha_{t-1}) - \sigma^*\| \\
&= \|(1-\varepsilon)^2\sigma_{t-2} + \varepsilon(1-\varepsilon)B^k(\alpha_{t-2}) + \varepsilon B^k(\alpha_{t-1}) - \sigma^*\| \\
&= \left\|(1-\varepsilon)^{\tau_2}\sigma_{t-\tau_2} + \sum_{j=1}^{\tau_2} \varepsilon(1-\varepsilon)^{j-1} B^k(\alpha_{t-j}) - \sigma^*\right\|
\end{aligned}$$

Let $j^* = \text{argmax}_{1\le j\le \tau_2} \|B^k(\alpha_{t-j}) - \sigma^*\|$, and define $\underline{B}^k \equiv B^k_{t-j^*}$. Since $t - \tau_2 > \tau_1$, $\|\underline{B}^k - \sigma^*\| < \eta/2$. Then,

$$\begin{aligned}\|\sigma_t - \sigma^*\| &\le (1-\varepsilon)^{\tau_2}\|\sigma_{t-\tau_2} - \sigma^*\| + (1-(1-\varepsilon)^{\tau_2})\|\underline{B}^k - \sigma^*\| \\ &< \frac{\eta}{2} + \frac{\eta}{2} = \eta.\end{aligned}$$

This proves the 'if' part. □

The proofs of Lemma 2 and Theorem 2 rely on the discrete time version of Lyapunov stability theorem below. A discrete-time dynamic process is *autonomous* if for all $s, t \in \{1, 2, \ldots\}$, $x_{s+1} = x_{t+1}$ if $x_s = x_t$. In words, an autonomous process has no time-dependent variable. The aggregate dynamic characterized by Eq.(7) is autonomous.

Theorem 7 (Lyapunov Stability Theorem)**.** *Consider an autonomous process $\alpha_t = h(\alpha_{t-1})$, where h is Lipschitz continuous in $D \in \mathbb{R}$ and $0 \in D$. Suppose $h(0) = 0$, i.e., 0 is a steady state (all this can be extended for an steady state different from 0). Let $V : \mathbb{R} \to \mathbb{R}$ be a continuous function such that $V(0) = 0$; $V(\alpha) > 0$, for $\alpha \in D - \{0\}$; $\lim_{|\alpha|\to\infty} V(\alpha) = \infty$; and $V(h(\alpha)) - V(\alpha) < 0$ for any $\alpha \in D\backslash\{0\}$, then $\lim_{t\to\infty} \alpha_t = 0$ for any initial point $\alpha_0 \in D$.*[15]

Proof. See Kalman and Bertram [1960], Khalil [1996] and Bof et al. [2017]. The above statement follows the version of Bof et al. [2017]. □

Proof of Lemma 2. Let $D = (\underline{\alpha}, \overline{\alpha})$. Define $V(\alpha) = |\alpha - \alpha^*|$. We show that $V(\alpha)$ serves as the Lyapunov function on D. Assume that $\alpha_t \in D$ with $\alpha_t < \alpha^*$. Observe that

$$\begin{aligned}|\alpha_{t+1} - \alpha^*| &= |(1-\varepsilon)\alpha_t + \varepsilon(1 - Bern_k(\alpha_t; \Lambda_{F,k})) - \alpha^*| \\ &= |(1-\varepsilon)(\alpha_t - \alpha^*) + \varepsilon(Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k}))|.\end{aligned}$$

For the second equality, we use that $\alpha^* = 1 - Bern_k(\alpha^*; \Lambda_{F,k})$.

[15] The condition that $\lim_{|\alpha|\to\infty} V(\alpha) = \infty$ is irrelevant to our analysis. This ensures that for any $c > 0$, there exists $B_r = \{\alpha \in D : |\alpha| \le r\}$ such that $B_r \supset V_c \equiv \{\alpha \in D : V(\alpha) \le c\}$. In our model, $B_r = \{\alpha : \alpha \le 1\}$ (or $B_r = \{\alpha : \alpha \le c\}$ for Theorem 3) always includes such V_c. We leave this condition in the statement to precisely describe the theorem.

First, consider that $Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k}) \geq 0$. Recall that $K_{B(k)}$ is a Lipschitz constant for $Bern_k(\alpha; \Lambda_{F,k})$. Fix $\varepsilon < \bar{\varepsilon} \leq 1/K_{B(k)}$. If $(1-\varepsilon)(\alpha^* - \alpha_t) \geq \varepsilon(Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k}))$, then it is obvious that $|\alpha_{t+1} - \alpha^*| < |\alpha_t - \alpha^*|$. If $(1-\varepsilon)(\alpha^* - \alpha_t) < \varepsilon(Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k}))$, then observe that

$$\begin{aligned}
|\alpha_{t+1} - \alpha^*| &= \; |(1-\varepsilon)(\alpha_t - \alpha^*) + \varepsilon(Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k}))| \\
&= \; \varepsilon(Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k})) - (1-\varepsilon)(\alpha^* - \alpha_t) \\
&< \; \frac{1}{K_{B(k)}}(K_{B(k)}(\alpha^* - \alpha_t)) - (1-\varepsilon)(\alpha^* - \alpha_t) \\
&< \; |\alpha_t - \alpha^*|.
\end{aligned}$$

Note that $\varepsilon < \bar{\varepsilon} \leq \bar{\alpha} - \alpha^*$. This upper-bound guarantees that $\alpha_{t+1} \in (\underline{\alpha}, \bar{\alpha})$ even if $\alpha_{t+1} > \alpha^*$. The same applies to the discussion to follow.

Second, consider that $Bern_k(\alpha^*; \Lambda_{F,k}) - Bern_k(\alpha_t; \Lambda_{F,k}) < 0$. We can rewrite $|\alpha_{t+1} - \alpha^*|$ and bound it as follows.

$$\begin{aligned}
|\alpha_{t+1} - \alpha^*| &= (1-\varepsilon)(\alpha^* - \alpha_t) + \varepsilon\left(Bern_k(\alpha_t; \Lambda_{F,k}) - Bern_k(\alpha^*; \Lambda_{F,k})\right) \\
&< \; |\alpha_t - \alpha^*|.
\end{aligned}$$

In the inequality, we use that $Bern_k(\alpha_t; \Lambda_{F,k}) < 1 - \alpha_t$ and $Bern_k(\alpha^*; \Lambda_{F,k}) = 1 - \alpha^*$.

We can prove the case that $\alpha_t > \alpha^*$ similarly. This implies that $V(\alpha_{t+1}) - V(\alpha_t) < 0$ for all $\alpha_t \neq \alpha^*$. Observe also that $V(\alpha) = 0$ if $\alpha = \alpha^*$, $V(\alpha) > 0$ otherwise. Then, by the Lyapunov stability theorem, we can conclude that the aggregate population state (starting with some point in D) converges to α^*. □

Proof of Theorem 2. If the initial state is a SESI proportion, then Lemma 1 implies the claim. Choose α_t that is not a SESI proportion. If $1 - \alpha_t > Bern_k(\alpha_t; \Lambda_{F,k})$, then there must exist some SESI $\alpha^* > \alpha_t$. To see this, observe that $(1 - \alpha)$ is continuously decreasing and reaches 0 at $\alpha = 1$, and $Bern_k(\alpha; \Lambda_{F,k})$ is continuous and bounded by $[0, 1]$. The two curves must intersect at some point to the right of α_t (by an appeal to the Brouwer's fixed point theorem). Choose the smallest SESI $\alpha^* > \alpha_t$. Then, $1 - \alpha > Bern_k(\alpha; \Lambda_{F,k})$ for all $\alpha \in (\alpha_t, \alpha^*)$. Since $Bern_k(\alpha; \Lambda_{F,k})$ crosses $1 - \alpha$ from below (at α^*), there exists some $\bar{\alpha} > \alpha^*$ such that $1 - \alpha < Bern_k(\alpha; \Lambda_{F,k})$ for all $\alpha \in (\alpha^*, \bar{\alpha})$. Lemma 2 together with Theorem 1 proves the claim. We can prove similarly for the case that $1 - \alpha_t < Bern_k(\alpha_t; \Lambda_{F,k})$.

In the above discussion, we implicitly assume an interior SESI proportion, that is, $\alpha^* \in (0,1)$. We can prove the case that $\alpha^* \in \{0,1\}$ in a similar way. For example, suppose that $\alpha^* = 1$ and $1-\alpha > Bern_k(\alpha;\Lambda_{F,k})$ for all $\alpha \in (\underline{\alpha},1)$ for some $\underline{\alpha}$. Let $V(\alpha) = 1-\alpha$. Then, $V(\alpha_{t+1}) - V(\alpha_t) < 0$ for all $\alpha_t \in (\underline{\alpha},1)$. To see this, recall that Eq.(7) implies that $\alpha_t > \alpha_{t-1}$ if $1-\alpha_{t-1} > Bern_k(\alpha_{t-1};\Lambda_{F,k})$. The Lyapunov stability theorem implies that the aggregate dynamic starting from $(\underline{\alpha},1]$ converges to $\alpha^* = 1$ for any $\varepsilon \in (0,1)$. □

Proof of Proposition 1. Fix an inference procedure $\mathcal{G}$, a sample size k, and a preference distribution. Rewrite Eq.(4) as $\alpha_{k,\mathcal{G}} = 1 - Bern_k(\alpha_{k,\mathcal{G}};\Lambda_{F,k})$. Define $g(\alpha) \equiv 1 - Bern_k(\alpha;\Lambda_{F,k})$. Function $g(\cdot)$ is continuous on $[0,1]$ and maps $[0,1]$ to itself. Brouwer's fixed point theorem implies that there exists $\alpha^* = 1 - Bern_k(\alpha^*;\Lambda_{F,k})$. Thus a SESI exists. It is unique since the Bernstein polynomial is weakly increasing in α. That is, $\alpha < 1 - Bern_k(\alpha;\Lambda_{F,k})$ if $\alpha < \alpha^*$, and $\alpha > 1 - Bern_k(\alpha;\Lambda_{F,k})$ if $\alpha > \alpha^*$. □

Lemma 4 below shows the existence of C^* defined in Section 4.2.

Lemma 4. *$C^* \neq \emptyset$.*

Proof of Lemma 4. For all $c \geq 1$, $b^{\theta}_{\mathcal{G}}(k,z|c) = B$ for all θ, k, z. Thus, $SW(\sigma^*_c) = SW(\sigma^*_1)$ for all $c > 1$. Without loss of generality, we restrict attention to $c \in [0,1]$.

We show that $SW(\sigma^*_c)$ is continuous on $[0,1]$. Suppose some $c \in [0,1)$ and a small $\Delta > 0$ with $c + \Delta \leq 1$. Let θ_j be the preference for which

$$\int_{\alpha \in \mathcal{A}} F^{\theta_j}_A(\alpha) d\mathcal{G}_{k,j/k}(\alpha) - c = 0.$$

That is, agents with θ_j are indifferent between A and B when they sample j of action A. If c increases by Δ, then the best response to the estimate for agents with preference $\theta \in (\theta_j, \theta_j + \Delta]$ will change from A to B. This is because $F^{\theta}_A(\alpha)$ is linear in θ. The absolute continuity of λ implies that

$$\lim_{\Delta \to 0} \lambda([\theta_j, \theta_j + \Delta]) = 0 \qquad\qquad \forall j \in \{0,\ldots,k\}.$$

In words, the set of agents whose best response to the estimate is affected by the Δ increase in c has measure zero in the limit of small Δ. This implies that

$\lim_{\Delta \to 0} |\sigma^*_{c+\Delta} - \sigma^*_c| = 0$, which further implies that $\lim_{\Delta \to 0} \alpha(\sigma^*_{c+\Delta}) = \alpha(\sigma^*_c)$. Thus, $\lim_{\Delta \to 0} SW(\sigma^*_{c+\Delta}) = \sigma^*_c$.

Similarly, we can show that $\lim_{\Delta \to 0} SW(\sigma^*_{c-\Delta}) = \sigma^*_c$ for $c \in (0, 1]$. Thus, $SW(\sigma^*_c)$ is continuous on $[0, 1]$, and the claim follows. □

Proof of Theorem 3. Observe that a game with tax c is equivalent to a game in Section 2 where the payoff function is defined as $F^\theta_A(\alpha) = \theta - f(\alpha) - c$ and $F^\theta_B(\alpha) = 0$. Then, the uniqueness and convergence results are immediate from Corollary 1.

For the last claim, let Δ_c denote the decrease in α when a tax c is imposed, or $\Delta_c = \alpha(\sigma^*_0) - \alpha(\sigma^*_c)$. Note that Δ_c is non-negative due to that $B^\theta_G(j/k|c)$ is weakly decreasing in c for all $j \in \{0, \ldots, k\}$ and $\theta \in \Theta$. Since payoffs are continuous in c and the preference measure λ is absolutely continuous, we can make Δ_c arbitrarily small by choosing an appropriately small $c > 0$. Observe that

$$SW(\sigma^*_c) - SW(\sigma^*_0) = \int_{\theta \in \Theta} \theta(B^{k,\theta}(\alpha(\sigma^*_c)|c) - B^{k,\theta}(\alpha(\sigma^*_0)))d\lambda \\ - \left[(\alpha(\sigma^*_0) - \Delta_c)f(\alpha(\sigma^*_0) - \Delta_c) - \alpha(\sigma^*_0)f(\alpha(\sigma^*_0))\right].$$

Note that $\int_{\theta \in \Theta}(B^{k,\theta}(\alpha(\sigma^*_c)|c) - B^{k,\theta}(\alpha(\sigma^*_0)))d\lambda = -\Delta_c$. Since $\theta \leq 1$, the difference in the social welfare can be bounded as follows.

$$SW(\sigma^*_c) - SW(\sigma^*_0) > -\Delta_c - \left[(\alpha(\sigma^*_0) - \Delta_c)f(\alpha(\sigma^*_0) - \Delta_c) - \alpha(\sigma^*_0)f(\alpha(\sigma^*_0))\right].$$

Let $H(\alpha) = \alpha f(\alpha)$. For all sufficiently small Δ_c, we can approximate the lower bound as

$$SW(\sigma^*_c) - SW(\sigma^*_0) > -\Delta_c \left[1 - \frac{H(\alpha(\sigma^*_0)) - H(\alpha(\sigma^*_0) - \Delta_c)}{\Delta_c}\right] \\ \approx -\Delta_c \left[1 - f(\alpha(\sigma^*_0)) - \alpha(\sigma^*_0)f'(\alpha(\sigma^*_0))\right].$$

Thus, if $1 - f(\alpha(\sigma^*_0)) - \alpha(\sigma^*_0)f'(\alpha(\sigma^*_0)) < 0$, then at least some small positive tax can increase the welfare. □

Proof of Observation 2. If $f(1-p) > 1$, for all $\alpha \leq 1-p$, $f(\alpha) \geq f(1-p) > 1$. Hence, for all $\alpha \leq 1-p$, $0 > \theta - f(\alpha)$ for all $\theta \in \Theta$. If B is p-dominant, then for all $\theta \in \Theta$ and all $\alpha \leq 1-p$, $\theta - f(\alpha) < 0$, which implies that $f(\alpha) > 1$ for all $\alpha \leq 1-p$. Hence, $f(1-p) > 1$. □

Proof of Theorem 4. We prove the theorem in three steps:
Step 1: Since $\mathcal{G}$ is unbiased, for any sample (l, z), $\mathcal{G}_{l,z}$ is a mean preserving spread of $\mathcal{G}_{l,z}^{MLE}$, where $\mathcal{G}_{l,z}^{MLE}$ denotes the estimate derived from the maximum likelihood estimation method. Given that $-f$ is nondecreasing and concave,

$$F_{l,z} = -\int_{\alpha \in \mathcal{A}} -f(\alpha) d\mathcal{G}_{l,z}(\alpha) \geq -\int_{\alpha \in \mathcal{A}} -f(\alpha) d\mathcal{G}_{l,z}^{MLE}(\alpha) = F_{l,z}^{MLE}.$$

Step 2: Since $\mathcal{G}$ is unbiased, $F_{l,1} = F_{l,1}^{MLE} = f(1) < 0$, implying that $\Lambda(F_{l,1}) = 0$. By Step 1 and Observation 2, since B is $1/k$-dominant, we have $F_{l,j/l} \geq F_{l,j/l}^{MLE} = f(j/l) \geq f(1 - 1/l) \geq f(1 - 1/k) > 1$ for $j \leq l - 1$, implying that $\Lambda(F_{l,j/l}) = 1$. Hence, $Bern_l(\alpha; \Lambda_{F,l}) = \sum_{j=0}^{l-1} \binom{l}{j} \alpha^j (1-\alpha)^{l-j} = 1 - \alpha^l$. This implies that $1 - \alpha < 1 - \alpha^l = Bern_l(\alpha; \Lambda_{F,l})$ for all $\alpha \in (0, 1)$.
Step 3: There is no SESI proportion other than $\alpha \in \{0, 1\}$ since $1 - \alpha < Bern_l(\alpha; \Lambda_{F,l})$ for all $\alpha \in (0, 1)$. The observation in Step 2 also implies that $\alpha_t < \alpha_{t-1}$ for all $\alpha_{t-1} \in (0, 1)$. Theorem 2 guarantees that α_t converges to $\alpha^* = 0$. □

We omit the formal proofs of Theorems 5 and 6 to avoid redundancy. They are similar to those of Theorems 1 and 2. We only offer a sketch of proofs instead. For Theorem 5, we can show that $B^{\mathbf{GK}}(\cdot)$ is Lipschitz continuous in a similar way we do for $B^k(\cdot)$ in Lemma 3. Then we can prove the theorem by replacing $B^k(\cdot)$ with $B^{\mathbf{GK}}(\cdot)$ in the proof of Theorem 1.

For Theorem 6, we briefly prove Lemma 5 below, which is analogous to Lemma 2. Then, Theorem 6 can be proved similarly to Theorem 2. Note that Observation 1 applies to $Bern_k(\alpha; \Lambda_{F,k}^{ik})$ for all inference procedures $\mathcal{G}^i \in \mathbf{G}$ and sample sizes $k \in \mathbf{K}$. Let $K_{B(k,i)}$ be a Lipschitz constant for $Bern_k(\alpha; \Lambda_{F,k}^{ik})$, and let $\overline{K} = \max_{\mathcal{G}^i, k \in \mathbf{G} \times \mathbf{K}} K_{B(k,i)}$.

Lemma 5. *Let α^* be a SESI proportion. Fix an open interval $(\underline{\alpha}, \overline{\alpha})$ that satisfies the three conditions: i) $\alpha^* \in (\underline{\alpha}, \overline{\alpha})$, ii) $1 - \alpha > \sum_{(\mathcal{G}^i,k) \in \mathbf{G} \times \mathbf{K}} Bern_k(\alpha; \Lambda_{F,k}^{ik})$ for all $\alpha \in (\underline{\alpha}, \alpha^*)$, and iii) $1 - \alpha < \sum_{(\mathcal{G}^i,k) \in \mathbf{G} \times \mathbf{K}} Bern_k(\alpha; \Lambda_{F,k}^{ik})$ for all $\alpha \in (\alpha^*, \overline{\alpha})$. The aggregate population state starting with some $\alpha \in (\underline{\alpha}, \overline{\alpha})$ converges to α^* for any $\varepsilon \in (0, \overline{\varepsilon})$, where $\overline{\varepsilon} = \min\{1/\overline{K}, \alpha^* - \underline{\alpha}, \overline{\alpha} - \alpha^*\}$.*

Proof of Lemma 5. Let $D = (\underline{\alpha}, \overline{\alpha})$, and $V(\alpha) = |\alpha - \alpha^*|$. It suffices to show that

$|\alpha_{t+1} - \alpha^*| < |\alpha_t - \alpha^*|$. Assume that $\alpha_t \in D$ with $\alpha_t < \alpha^*$. Observe that

$$|\alpha_{t+1} - \alpha^*| = \left|(1-\varepsilon)(\alpha_t - \alpha^*) + \varepsilon \sum_{\mathcal{G}^i, k \in \mathbf{G} \times \mathbf{K}} \left[Bern_k(\alpha^*; \Lambda_{F,k}^{ik}) - Bern_k(\alpha_t; \Lambda_{F,k}^{ik})\right]\right|.$$

We can prove the claim similarly to Lemma 2. To see this, replace $Bern_k(\alpha^*; \Lambda_{F,k})$ with $\sum_{\mathcal{G}^i, k \in \mathbf{G} \times \mathbf{K}} Bern_k(\alpha^*; \Lambda_{F,k}^{ik})$, and $Bern_k(\alpha_t; \Lambda_{F,k})$ with $\sum_{\mathcal{G}^i, k \in \mathbf{G} \times \mathbf{K}} Bern_k(\alpha_t; \Lambda_{F,k}^{ik})$ in the proof of Lemma 2. □

References

I. Arieli, Y. Babichenko, R. Peretz, and H. P. Young. The speed of innovation diffusion in social networks. 88(2):569–594, 2020. doi: 10.3982/ecta17007.

S. Arigapudi, Y. Heller, and I. Milchtaich. Instability of defection in the prisoner's dilemma under best experienced payoff dynamics. *Journal of Economic Theory*, 197:105174, 2021. doi: 10.1016/j.jet.2020.105174.

B. D. Bernheim and C. Sprenger. On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting. *Econometrica*, 88(4):1363–1409, 2020. doi: 10.3982/ecta16646.

N. Bof, R. Carli, and L. Schenato. Lyapunov theory for discrete time systems. Techical Report, 2017. URL `http://automatica.dei.unipd.it/tl files/utenti2/bof/Papers/NoteDiscreteLyapunov.pdf`.

J. C. Ely and W. H. Sandholm. Evolution in bayesian games i: Theory. *Games and Economic Behavior*, 53(1):83–109, 2005. doi: 10.1016/j.geb.2004.09.003.

D. P. Foster and P. H. Young. Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232, 1990.

A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. Network games. *Review of Economic Studies*, 77(1):218–244, 2010. doi: 10.1111/j.1467-937x.2009.00570.x.

J. P. Gleeson and M. A. Porter. Message-passing methods for complex contagions. In *Computational Social Sciences*, pages 81–95. Springer International Publishing, 2018.

Y. Heller and E. Mohlin. Social learning and the shadow of the past. *Journal of Economic Theory*, 177:426–460, 2018. doi: 10.1016/j.jet.2018.07.003.

M. O. Jackson and L. Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, 2007. doi: 10.1257/aer.97.2.92.

M. O. Jackson and Y. Zenou. Games on networks. In *Handbook of Game Theory with Economic Applications*, pages 95–163. Elsevier, 2015. doi: 10.1016/b978-0-444-53766-9.00003-3.

D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979. doi: 10.2307/1914185.

R. E. Kalman and J. E. Bertram. Control system analysis and design via the "second method" of lyapunov ii: Discrete time systems. *Journal of Basic Engineering*, 82 (2):394–400, 1960. doi: 10.1115/1.3662605.

M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.

H. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 1996.

T. Kobayashi and T. Onaga. Dynamics of diffusion on monoplex and multiplex networks: A message-passing approach. *arXiv:2103.09417*, 2021. https://arxiv.org/abs/2103.09417.

F. Kojima and S. Takahashi. Anti-coordination games and dynamic stability. *International Game Theory Review*, 09(04):667–688, 2007. doi: 10.1142/s0219198907001655.

M. Kosfeld, E. Droste, and M. Voorneveld. A myopic adjustment process leading to best-reply matching. *Games and Economic Behavior*, 40(2):270–298, 2002. doi: 10.1016/s0899-8256(02)00007-6.

G. E. Kreindler and H. P. Young. Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67, 2013. doi: 10.1016/j.geb.2013.02.004.

R. Lahkar and S. Mukherjee. Evolutionary implementation in a public goods game. *Journal of Economic Theory*, 181:423–460, 2019. doi: 10.1016/j.jet.2019.03.009.

R. Lahkar and S. Mukherjee. Evolutionary implementation in aggregative games. *Mathematical Social Sciences*, 109:137–151, 2021. doi: 10.1016/j.mathsocsci.2020.11.004.

C. Mantilla, R. Sethi, and J. C. Cárdenas. Efficiency and stability of sampling equilibrium in public goods games. *Journal of Public Economic Theory*, 22(2):355–370, 2020. doi: 10.1111/jpet.12351.

J. Miekisz and M. Ramsza. Sampling dynamics of a symmetric ultimatum game. *Dynamic Games and Applications*, 3(3):374–386, 2012. doi: 10.1007/s13235-012-0064-5.

S. Morris. Contagion. *Review of Economic Studies*, 67:57–78, 2000.

S. Morris, R. Rob, and H. S. Shin. p-dominance and belief potential. *Econometrica*, 63(1):145–157, 1995. doi: 10.2307/2951700.

J. Newton and S. D. Angus. Coalitions, tipping points and the speed of evolution. *Journal of Economic Theory*, 157:172–187, 2015.

M. J. Osborne and A. Rubinstein. Games with procedurally rational players. *American Economic Review*, 88:834–874, 1998.

M. J. Osborne and A. Rubinstein. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45:434–411, 2003.

D. Oyama and S. Takahashi. Contagion and uninvadability in local interaction games: The bilingual game and general supermodular games. *Journal of Economic Theory*, 157:100–127, may 2015. doi: 10.1016/j.jet.2014.12.012.

D. Oyama, W. H. Sandholm, and O. Tercieux. Sampling best response dynamics and deterministic equilibrium selection. *Theoretical Economics*, 10(1):243–281, 2015. doi: 10.3982/te1405.

G. M. Phillips. *Interpolation and approximation by polynomials*. Springer, New York, 2003. ISBN 9780387216829.

Y. Salant and J. Cherry. Statistical inference in games. *Econometrica*, 88(4):1725–1752, 2020. doi: 10.3982/ecta17105.

W. H. Sandholm. Almost global convergence to p-dominant equilibrium. *International Journal of Game Theory*, 30:107–116, 2001. doi: 10.1007/s001820100067.

W. H. Sandholm. Evolutionary implementation and congestion pricing. *Review of Economic Studies*, 69(3):667–689, 2002. doi: 10.1111/1467-937x.t01-1-00026.

W. H. Sandholm. Negative externalities and evolutionary implementation. *Review of Economic Studies*, 72(3):885–915, 2005. doi: 10.1111/j.1467-937x.2005.00355.x.

W. H. Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007. doi: 10.1016/j.jet.2005.09.005.

W. H. Sandholm, S. S. Izquierdo, and L. R. Izquierdo. Best experienced payoff dynamics and cooperation in the centipede game. *Theoretical Economics*, 14(4): 1347–1385, 2019. doi: 10.3982/te3565.

W. H. Sandholm, S. S. Izquierdo, and L. R. Izquierdo. Stability for best experienced payoff dynamics. *Journal of Economic Theory*, 185:104957, 2020. doi: 10.1016/j.jet.2019.104957.

R. Sethi. Stability of equilibria in games with procedurally rational players. *Games and Economic Behavior*, 32:85–104, 2000.

R. Sethi. Stable sampling in repeated games. *Journal of Economic Theory*, 197:105343, 2021. doi: 10.1016/j.jet.2021.105343.

R. Spiegler. Competition over agents with boundedly rational expectations. *Theoretical Economics*, 1(2):207–231, 2006a.

R. Spiegler. The market for quacks. *Review of Economic Studies*, 73(4):1113–1131, 2006b. doi: 10.1111/j.1467-937x.2006.00410.x.

O. Tercieux. p-best response set. *Journal of Economic Theory*, 131(1):45–70, 2006. doi: 10.1016/j.jet.2005.06.001.

A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992. doi: 10.1007/bf00122574.

G. Wu and R. Gonzalez. Curvature of the probability weighting function. *Management Science*, 42(12):1676–1690, 1996. doi: 10.1287/mnsc.42.12.1676.

P. H. Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.

P. H. Young. *Individual Strategy and Social Structure*. Princeton University Press, 1998.

B. Zhang. Quantal response methods for equilibrium selection in normal form games. *Journal of Mathematical Economics*, 64:113–123, may 2016. doi: 10.1016/j.jmateco.2016.04.003.

D. Zusai. Tempered best response dynamics. *International Journal of Game Theory*, 47(1):1–34, apr 2017. doi: 10.1007/s00182-017-0575-9.

D. Zusai. Heterogeneity and aggregation in evolutionary dynamics. mimeo, Temple University, 2018. URL `https://arxiv.org/pdf/1805.04897v1.pdf`.

D. Zusai. Distributional stability and deterministic equilibrium selection under heterogeneous evolutionary dynamics. mimeo, Temple University, 2019a. URL `https://arxiv.org/pdf/1805.04895.pdf`.

D. Zusai. Evolutionary dynamics in heterogeneous populations: a general framework for an arbitrary type distribution. mimeo, Temple University, 2019b. URL `https://arxiv.org/abs/1805.04897`.