Sampling Logit Equilibrium and Endogenous Payoff Distortion

Minoru Osawa

Institute of Economic Research, Kyoto University

osawa.minoru@gmail.com

March 10, 2026

We introduce the *sampling logit equilibrium* (SLE), a stationary concept for population games in which agents evaluate actions using a finite sample of opponents' plays and respond according to a logit choice rule. This framework combines informational frictions from finite sampling with stochastic choice. When the sample size is large, SLE is well approximated by a logit equilibrium of a *virtual game* whose payoffs incorporate explicit distortion terms generated by sampling noise. Examples illustrate how finite sampling can systematically shift equilibrium behavior and generate equilibrium selection effects.

Keywords: Evolutionary game dynamics, bounded rationality, quantal response equilibria, sampling equilibrium, equilibrium selection.

JEL Classification: C72, C73.

Website: mino-osawa.github.io. I thank Srinivas Arigapudi, Chiaki Hara, Jonathan Newton, Daisuke Oyama, Yannick Viossat, and Dai Zusai for their comments. I am grateful to the late Professor Sandholm for his beautiful book, "*Population Games and Evolutionary Dynamics*," which has been a constant source of inspiration. Financial support from JSPS KAKENHI Grant Number JP22K04353 is warmly acknowledged.

1 Introduction

Decision making in strategic environments often departs from perfect rationality through two distinct channels. One is *stochastic choice*: even when agents observe payoffs correctly, their responses may be probabilistic due to idiosyncratic shocks or cognitive noise. The other is *informational constraint*: agents frequently evaluate actions using only limited observations of the environment, forming payoff assessments from a small sample of interactions. These mechanisms introduce noise in different ways. Stochastic choice perturbs the decision rule, whereas finite sampling distorts the evaluation of payoffs themselves.

Existing approaches in game theory typically treat these channels separately. Payoff-perturbation models such as the *quantal response equilibrium* (QRE) capture stochastic choice through random utility models while assuming that agents evaluate expected payoffs under full information (McKelvey and Palfrey, 1995, 1998; Goeree et al., 2005). Conversely, *sampling* models assume that agents observe only a limited subset of the environment, but typically retain deterministic best-response behavior to the sampled observations (e.g., Osborne and Rubinstein, 1998, 2003; Salant and Cherry, 2020). This separation leaves open a natural question: how does behavior change when agents both observe only a finite sample of the environment and respond stochastically to the resulting payoff signals?

This study analyzes that interaction in a simple framework combining finite sampling with logit stochastic choice. Under the (k, η)*-sampling logit choice rule*, each agent draws k independent samples of opponents from the population, evaluates the resulting sample-based payoffs, and then selects an action according to a logit rule with noise level $\eta > 0$. A *sampling logit equilibrium* (SLE) is defined as a fixed point of this process.

The framework connects naturally to existing models at two limiting regimes. When the sample size k is large, agents effectively observe the true population state and SLE approaches the standard logit equilibrium. When the decision noise η vanishes while k remains finite, SLE converges to the *sampling equilibrium* of Osborne and Rubinstein (2003), equivalently the stationary states of the sampling best-response dynamics studied by Oyama et al. (2015). By varying (k, η), the model therefore spans a continuum between fully informed stochastic choice and limited-information deterministic response.

Our main result shows that the interaction between finite sampling and stochastic choice generates systematic distortions in incentives that admit a simple deterministic representation. When the sample size k is sufficiently large, the sampling logit equilibrium of a game can be approximated by the logit equilibrium of a *virtual game* with modified payoffs. The modification consists of two components. A *variance premium* favors actions whose payoffs fluctuate more across samples, because payoff variability increases the chance that an action appears temporarily attractive under the logit rule. A *curvature premium* arises from nonlinearities in payoff functions: sampling noise interacts with payoff curvature to produce Jensen-type shifts in effective payoffs. Finite sampling therefore does not merely add randomness to behavior but systematically alters the incentives agents face.

This representation provides a tractable way to analyze the joint effects of noisy observation and noisy choice. Agents behave *as if* they were responding to distorted payoffs, allowing the equilibrium implications of the model to be studied using familiar tools for logit equilibrium applied to a game with endogenously modified incentives.

Beyond this approximation, we establish several exact results in benchmark environments. When each agent observes only one or two opponents, the equilibrium is unique and globally attracting under the sampling logit dynamics. Moreover, as the logit noise vanishes, the equilibrium converges to the risk-dominant Nash equilibrium in coordination games. These results illustrate how finite sampling can sharpen equilibrium selection when agents operate under limited information.

Related literature

The present study relates to four strands of literature: extensions of logit QRE, models of finite sampling and limited observation, stochastic evolutionary dynamics, and the virtual-payoff approach.

First, the framework extends the *quantal response equilibrium* (QRE) paradigm, which incorporates stochastic choice into equilibrium analysis through random utility models (McKelvey and Palfrey, 1995, 1998; Goeree et al., 2005). In QRE, agents evaluate expected payoffs under full information and randomness arises from idiosyncratic payoff shocks. The present study instead introduces stochasticity through finite sampling of the environment. Payoff evaluations are therefore based on a small number of observations, so randomness enters through distorted payoff estimates rather than purely idiosyncratic payoff noise. Experimental comparisons of stationary concepts have shown that models incorporating sampling often fit observed play better than QRE (Selten and Chmura, 2008), suggesting that sampling may represent an important source of behavioral noise.

Second, the analysis builds on the literature studying limited observation in games. Early contributions include Osborne and Rubinstein (1998, 2003), who introduced the concept of sampling equilibrium. Related work examines strategic behavior when agents observe only partial information about the environment or other players' actions (e.g., Spiegler, 2006a,b; Salant and Cherry, 2020). Learning dynamics under such informational constraints have also been studied extensively (e.g., Sethi, 2000; Sandholm, 2001; Oyama et al., 2015; Heller and Mohlin, 2018; Mantilla et al., 2020; Sawa and Wu, 2023; Arigapudi et al., 2024; Danenberg and Spiegler, 2025). The framework here is particularly close to Oyama et al. (2015), who analyze dynamics in which agents occasionally sample opponents and best respond to the resulting empirical distribution. The present study replaces deterministic best responses with logit responses, yielding a smoother dynamic and enabling a tractable approximation in terms of a virtual game.

Third, the model connects to the literature on *stochastic evolutionary dynamics* and stochastic stability in games (Foster and Young, 1990; Young, 1993; Kandori et al., 1993; Ellison, 1993). In these models, random experimentation or errors drive transitions between states, and equilibrium selection is studied through the long-run behavior of the induced Markov process. Logit choice has been widely used as a modeling device in this context (e.g., Blume, 1993, 1995; Alós-Ferrer and Netzer, 2010; Marden and Shamma, 2012). The study most closely related to the present analysis is Kreindler and Young (2013), who analyze a finite-population model in which agents observe a random sample of opponents before making a logit choice. Their focus is the speed of convergence in two-action coordination games. In contrast, the present study introduces a static equilibrium concept for large-population games and characterizes its properties in general n-action environments.

Finally, the analysis relates to the *virtual payoff* approach for representing stochastic choice models through deterministic payoff perturbations. In large-population games, Hofbauer and Sandholm (2007) show that stationary behavior under stochastic choice can often be represented as if agents were responding to appropriately modified payoffs. Related ideas appear in earlier discussions of logit-based models (Anderson et al., 1992). The present study derives an explicit virtual payoff representation that captures the distortions generated by finite sampling under stochastic choice. In equilibrium, agents behave *as if* they were responding to these modified incentives.

2 Model

2.1 Population games

We focus on a large-population game played by a single homogeneous population. There is a unit mass of anonymous agents each of whom chooses their pure action, where $S \equiv \{1, 2, \dots, n\}$ denotes the common, finite set of available actions. The n-simplex $X \equiv \{x \in \mathbb{R}^n_{\geq 0} : \sum_{i \in S} x_i = 1\}$ represents the set of *population states* where $x_i \in [0, 1]$ is the fraction of agents playing action i. A state at which all agents play action i is called a *pure population state* and denoted by e_i. The function $F : X \to \mathbb{R}^n$ describes a game's payoffs, with $F_i(x)$ being the payoff obtained at state x by nonatomic and anonymous agents playing action $i \in S$. Where S is understood, F identifies a *population game*.

The most basic instances of population games are those generated by random matching in normal form games. Suppose that, given the population state, agents are randomly matched to play a normal form game with payoff matrix $A = [a_{ij}] \in \mathbb{R}^{n \times n}$, where a_{ij} is the payoff for an agent playing action $i \in S$ matched to another agent playing action $j \in S$. Then, the expected payoffs are given by $F(x) = Ax$, which we may call a *linear population game*. We can identify a linear population game with its base payoff matrix A.

Below, we introduce four choice rules in population games relevant for our discussion. For each rule, the associated evolutionary game dynamic is defined as the expected motion of population state (see Sandholm, 2010, Ch.4 "Revision Protocols and Evolutionary Dynamics").

2.2 Best response

Given a population game F, the pure and mixed best response correspondences $\mathrm{br} : X \rightrightarrows S$ and $\mathrm{BR} : X \rightrightarrows X$ are respectively defined as

$$\mathrm{br}(x) \equiv \arg\max_{i \in S} F_i(x) \quad \text{and} \tag{1}$$

$$\mathrm{BR}(x) \equiv \arg\max_{y \in X} \langle y, F(x) \rangle = \{ y \in X : \mathrm{support}(y) \subseteq \mathrm{br}(x) \}, \tag{2}$$

where we define $\langle u, v \rangle \equiv \sum_i u_i v_i$ for vectors. A population state $x \in X$ is a *Nash equilibrium* of F if every agent is playing a pure action that is optimal given the others' behavior, for which case $x \in \mathrm{BR}(x)$. The *best response dynamic* (Gilboa and Matsui, 1991; Hofbauer, 1995) is defined as the following differential inclusion:[1]

$$\dot{x} \in \mathrm{BR}(x) - x. \tag{BRD}$$

The dynamic assumes each agent has perfect information of the current population state and is able to choose an optimal action, which can be demanding depending on the context.

2.3 Best response under finite sampling

Oyama et al. (2015) considers an alternative model that impose milder informational requirements. When an agent receives a revision opportunity, the agent first observes $k \geq 1$ independent samples of opponents' play from the population. The set of possible outcomes of samples of size k is $Z^k \equiv \{ z \in \mathbb{Z}^n_{\geq 0} : \sum_{i \in S} z_i = k \}$. The *empirical population state* according to a sample $z \in Z^k$ is $w = \frac{1}{k} z \in X$. The agent then evaluates the payoffs based on $w \in X$, and plays a best response. The probability that $z \in Z^k$ is drawn at $x \in X$ follows the multinomial distribution $\mathrm{Mult}(k \,|\, x)$. That is, if $M^k(z \,|\, x) \in [0, 1]$ denotes the probability mass of drawing $z \in Z^k$ at $x \in X$, we have

$$M^k(z \,|\, x) = \binom{k}{z_1, z_2, \cdots, z_n} \cdot x_1^{z_1} \cdot x_2^{z_2} \cdot \cdots \cdot x_n^{z_n}. \tag{3}$$

The empirical population state $w = \frac{1}{k} z$ is the maximum likelihood estimator of the population state under the multinomial sampling model. For each k, the *k-sampling best response correspondence* $\mathrm{BR}^k : X \rightrightarrows X$ is

$$\mathrm{BR}^k(x) \equiv \mathbb{E}\left[\mathrm{BR}(\tfrac{1}{k} z)\right] = \sum_{z \in Z^k} M^k(z \,|\, x) \cdot \mathrm{BR}(\tfrac{1}{k} z). \tag{4}$$

More precisely, $y \in \mathrm{BR}^k(x)$ if and only if $y = \sum_{z \in Z^k} M^k(z \mid x) \cdot \alpha(z)$ where $\alpha(z) \in \mathrm{BR}(\frac{1}{k} z)$ for each $z \in Z^k$. A (k-) *sampling equilibrium* (Osborne and

[1] See Oyama et al. (2015, Appendix A.1) for a concise summary of differential inclusions.

Rubinstein, 2003) is a fixed point of BR^k satisfying $x \in \mathrm{BR}^k(x)$.[2] The *k-sampling best response dynamic* is defined as the following differential inclusion:

$$\dot{x} \in \mathrm{BR}^k(x) - x. \tag{SBRD}$$

While demanding less information about the population state, this model assumes agents are rational responders.

2.4 Logit choice

We next recall the logit choice and the logit dynamic. At each $x \in X$, the *η-logit choice rule $P^\eta : X \to X$* with *noise level $\eta > 0$* is defined as the following mixed strategy given the current payoff $F(x)$:

$$P_i^\eta(x) \equiv \frac{\exp\left(\eta^{-1}F_i(x)\right)}{\sum_{j\in S}\exp\left(\eta^{-1}F_j(x)\right)}. \tag{5}$$

The logit choice rule can be interpreted either as a random utility model or as the solution to a deterministically perturbed optimization problem.[3]

A population state x is a *η-logit equilibrium* if it is consistent with the η-logit choice rule: $x = P^\eta(x)$. Logit equilibrium is by far the most common formulation of QRE. The large-population *η-logit dynamic* (Fudenberg and Levine, 1998, Ch.4) is defined by the following differential equation:[4]

$$\dot{x} = P^\eta(x) - x. \tag{LD}$$

The η-logit equilibria for game F coincide with the Nash equilibria for the distorted game $\widetilde{F}$ where $\widetilde{F}_i(x) \equiv F_i(x) - \eta\log(x_i)$. In this sense, $\widetilde{F}$ is a "virtual" payoff for the η-logit equilibrium problem (Hofbauer and Sandholm, 2007, Appendix).[5]

[2]Sampling equilibrium is a special case of *sampling equilibrium with statistical inference* (SESI) (Salant and Cherry, 2020). Sawa and Wu (2023) extensively discusses Bayesian large-population dynamics corresponding to SESI.

[3]Our analysis is agnostic to the precise microfoundation and focuses on the interaction between finite sampling and this widely used stochastic choice rule.

[4]For finite-population settings, logit choice yields Markov chain/process rather than differential equation, as considered in, e.g., Blume (1993, 1995, 2003); Hofbauer and Sandholm (2007); Alós-Ferrer and Netzer (2010); Marden and Shamma (2012), among others. Another foundation for logit dynamic can be found in stochastic fictitious play (Hofbauer and Sandholm, 2002).

[5]Requiring $x^* \in X$ and $\lambda^* = \widetilde{F}_i(x^*)$ for all $i \in S$ implies $\lambda^* = \eta\log\sum_{j\in S}\exp(\eta^{-1}F_i(x_i^*))$ and $x^* = P^\eta(x^*)$. See Behrens and Murata (2021) for an application in spatial economics, where

2.5 Logit choice under finite sampling

Our framework is a natural synthesis of the sampling best response and the logit choice. Assume that, after observing a sample $z \in Z^k$, each agent follows the η-logit choice rule P^η instead of the mixed best response BR. The induced aggregate choice rule $L^{k,\eta} : X \to X$ defines the (k,η)*-sampling logit choice rule*:

$$L^{k,\eta}(x) \equiv \mathbb{E}\left[P^\eta(\tfrac{1}{k}z)\right] = \sum_{z \in Z^k} M^k(z \,|\, x) \cdot P^\eta(\tfrac{1}{k}z). \tag{6}$$

A (k,η)*-sampling logit equilibrium* (SLE) is a fixed point of $L^{k,\eta}$ satisfying $x = L^{k,\eta}(x)$. The (k,η)*-sampling logit dynamic* can be defined as

$$\dot{x} = L^{k,\eta}(x) - x. \tag{SLD}$$

Brouwer's fixed point theorem implies that SLE exist for any $1 \leq k < \infty$ and $\eta > 0$. The dynamic (SLD) admits unique global solution for every initial population state in X as $L^{k,\eta}$ is continuous and globally Lipschitz on X.

All SLE are necessarily positive because $L^{k,\eta}$ is strictly positive. For example, a pure population state cannot be a fixed point of $L^{k,\eta}$ because $L^{k,\eta}(e_i) = P^\eta(e_i) > \mathbf{0}$ for any $i \in S$ and $k \geq 1$. We have a uniqueness result for the case $k = 1$.

Proposition 1. *For any $\eta > 0$, if $k = 1$, there exists a unique SLE, and it is globally asymptotically stable under the sampling logit dynamics* (SLD).

3 Examples

Selected examples serve to build intuitions. Proofs are in Appendix A.

3.1 A two-action coordination game

This section presents benchmark environments where exact properties of SLE can be derived and where its qualitative behavior differs from both logit equilibrium and sampling equilibrium. For two-action population games, we have a strong characterization similar to Proposition 1 for the case $k = 2$.

Proposition 2. *Consider a two-action population game. For any $\eta > 0$, if $k = 2$, there exists a unique SLE, and it is globally asymptotically stable under the sampling logit dynamics* (SLD).

$-\eta \log(x_i)$ represents *congestion externalities* incurred by the households residing in *location* $i \in S$.

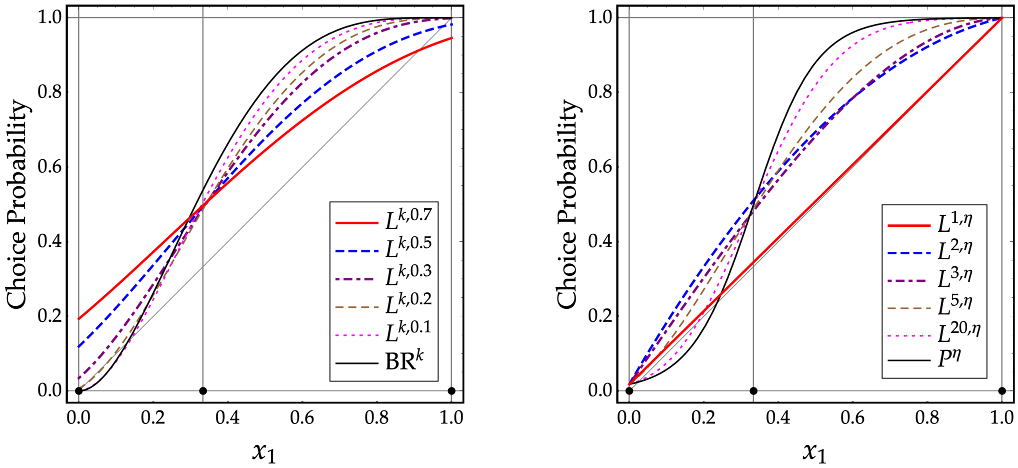

(A) BR_1^k and $L_1^{k,\eta}$ for different η ($k = 5$) (B) P_1^{η} and $L_1^{k,\eta}$ for different k ($\eta = 0.25$)

Figure 1: Choice probability of action 1 in the game (7) under different rules.

As a concrete example, consider a linear population game

$$A = \begin{bmatrix} s & 0 \\ 0 & t \end{bmatrix} \qquad (s > t > 0). \tag{7}$$

This is a coordination game, and the Nash equilibria satisfy $x_1 \in \{0, \frac{t}{s+t}, 1\}$, and $x = (1, 0) = e_1$ is *risk dominant*.

Assuming $s = 2$ and $t = 1$, Figure 1 compares the choice probability of action 1 under different choice rules. We see $L^{k,\eta}$ tends toward BR^k as η decreases, and toward P^{η} as k increases.

For $k = 1$, the unique SLE for the game is given by

$$x_1^* = \frac{1 + \mathrm{e}^{-s/\eta}}{1 + \mathrm{e}^{-(s-t)/\eta} + 2\mathrm{e}^{-s/\eta}} \in (\tfrac{1}{2}, 1). \tag{8}$$

From $s > t > 0$, we see $x_1^* \to 1$ as $\eta \to 0$, which yields a method of equilibrium refinement. In fact, the unique $(1, \eta)$-SLE can be alternatively identified as the stationary distribution of a logit-perturbed Markov chain, and selection under the $\eta \to 0$ limit yields analogous predictions as the *stochastic stability* approach under log-linear learning rules (Young, 1993; Blume, 1993, 1995, 2003).[6]

[6]The selection in our context can be "fast" because convergence is described by a deterministic ordinary differential equation, rather than requiring long-run sampling of a stochastic evolutionary dynamics. This aspect of sampling dynamics is stressed in the literature (Kreindler and Young, 2013; Oyama et al., 2015).

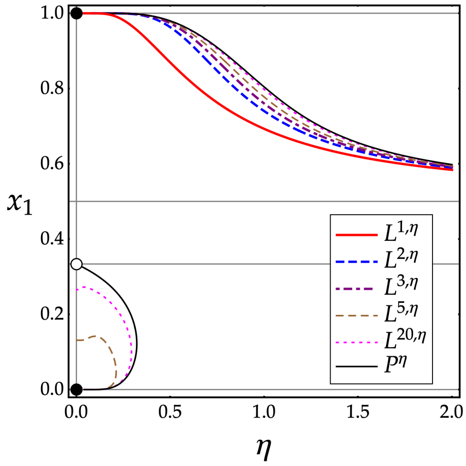

Figure 2: Sampling logit equilibria for $k \in \{1, 2, 3, 5, 20\}$ in the game (7).

If $k = 2$, there is a unique SLE x_1^{**} such that $x_1^{**} \to 1$ as $\eta \to 0$ (see the proof of Proposition 1). These limiting results can be seen as a sampling version of the convergence result for the principal branch of quantal response functions in 2×2 games (Turocy, 2005, Theorem 7), which in particular shows this branch converges to risk-dominant equilibria in coordination games. This behavior also resembles *almost global asymptotic stability* under (SBRD): all its trajectories starting from $x_1 \in X \setminus \{0\} = (0, 1]$ converges to the risk dominant equilibrium if $k = 2$ (Oyama et al., 2015, Theorem 1). The origin x_1 is the exception because it is another sampling equilibrium, albeit locally unstable. Under (SLD), in contrast, $x_1 = 0$ is not even a fixed point, and the unique SLE is globally attracting.

Figure 2 considers the same setting as Fig. 1 to illustrate how small k leads to equilibrium selection. The logit equilibrium curves are shown for reference, for which we note multiplicity of equilibria for small η and the convergence to either of the three Nash equilibria as η approaches zero. On the other hand, for each $k \in \{1, 2, 3\}$, the SLE is unique for all $\eta > 0$, allowing for selection in the limit $\eta \to 0$. Naturally, SLE approximate logit equilibria as we increase k, and multiple equilibria emerge.

3.2 Young (1993)'s game

Consider the linear population game based on the 3×3 game of Young (1993):

$$A = \begin{bmatrix} 6 & 0 & 0 \\ 5 & 7 & 5 \\ 0 & 5 & 8 \end{bmatrix}. \tag{9}$$

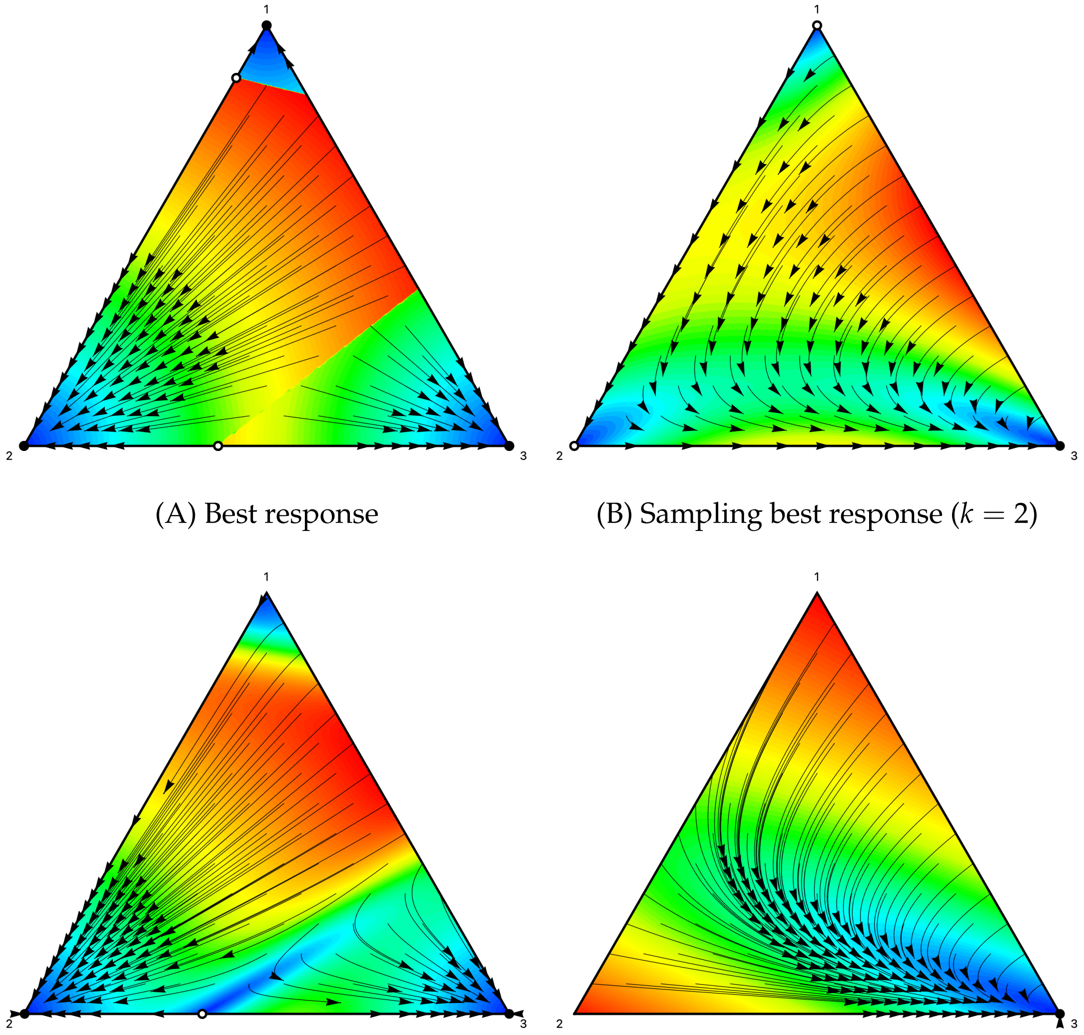

(A) Best response

(B) Sampling best response ($k = 2$)

(C) Logit ($\eta = 0.3$)

(D) Sampling logit ($k = 2, \eta = 0.3$)

Figure 3: Phase diagrams of the four dynamics in Young's game (9). Arrows show sample trajectories, and background contours depict the speed of adjustment: warmer colors indicate faster adjustment, whereas cooler colors indicate slower adjustment. This figure and the next were generated with the `Dynamo` software (Franchetti and Sandholm, 2013).

The game possesses three strict equilibria $\{e_1, e_2, e_3\}$ and two mixed Nash equilibria on the boundary of X. Figure 3 depicts the phase portraits for the four dynamics introduced in Section 2.

The best response dynamic (BRD) provides a basic refinement, where the two boundary equilibria are deemed unstable (Fig. 3A). However, it does not yield a unique prediction, as all three corners are locally stable. Figure 3B demonstrates that (SBRD) with $k = 2$ selects $x = e_3$ because e_3 attracts all trajectories starting from $X \setminus \{e_1, e_2\}$ (Oyama et al., 2015, Example 2).

Figure 3C considers the logit dynamic (LD). As the dynamic approximates (BRD) for small η, multiple stationary states are locally stable (Fig. 3C). Figure 3D illustrates that, with sampling noise, a unique prediction can be obtained. The correspondence between Figs. 3C and 3D is analogous to that between Figs. 3A and 3B. With two types of noises, the SLE nearby e_3 is globally attracting with no exceptions at the other corners.[7]

4 Approximation and virtual payoffs

In the remainder of the paper, we employ an approximation approach to study the qualitative implications of the sampling logit choice. The results in Section 3 provide exact characterizations of SLE in several benchmark environments. In more general settings, however, the noisy nature of logit choice makes direct characterization difficult.[8]

For noisy evolutionary dynamics, it is often useful to analyze deterministic approximations of stochastic revision processes. For instance, Benaïm and Weibull (2003) used stochastic approximation theory to derive deterministic game dynamics as the expected behavior of stochastic revision processes in finite populations. Following this approach, the analysis below proceeds by approximating the expected choice rule. The equilibrium distortions identified in the following sections are derived from this approximated rule, which provides an accurate representation of the true choice rule provided that k is sufficiently large.

We suppress k and η from notation whenever they are understood from the context. For example, $L = L^{k,\eta}$, $P = P^{\eta}$, and so on.

[7]While the comparison of Figs. 3A to 3D may give the impression that the sampling dynamics (SBRD) and (SLD) exhibit broadly similar qualitative behavior, this conclusion is game-dependent. Appendix B presents a numerical illustration based on the *bilingual game*. In that example, the introduction of sampling noise changes the geometry of the phase portrait in a more substantial way than in Young's game. Such behavior suggests that the interaction between payoff noise and sampling noise can generate nontrivial dynamical phenomena that are not visible in simpler examples.

[8]For example, Oyama et al. (2015) derived conditions under which an "iterated p-dominant equilibrium" is almost globally asymptotically stable under their sampling dynamics.

4.1 Approximation by the delta method

It is noted that the empirical population state $w = \frac{1}{k}z$ admits an asymptotic normal approximation. Provided that k is sufficiently large and the population state x is not too close to the boundary of X, we can assume that w approximately follows the multivariate normal distribution with mean $\mathbb{E}[w] = x$ and covariance matrix $\mathrm{Var}[w] = \frac{1}{k}\Sigma$. Here, $\Sigma \equiv \mathrm{diag}[x] - xx^\top$ is the covariance matrix for $\mathrm{Mult}(k \,|\, x)$, with $\top$ denoting transpose.

The asymptotic normality of w allows us to employ the *delta method* (e.g., van der Vaart, 2000, Ch.3) to approximate the expectation over samples. Specifically, the second-order Taylor approximation of $P_i(w)$ is

$$P_i^\star(w) \equiv P_i(x) + \langle P_i'(x), w - x\rangle + \frac{1}{2}(w - x)^\top P_i''(x)(w - x), \tag{10}$$

where P_i' and P_i'' denotes the gradient and the Hessian matrix of P_i, respectively.[9] Then, the second-order delta method approximates $L_i(x) = \mathbb{E}[P_i(w)]$ by

$$\widetilde{L}_i(x) \equiv \mathbb{E}\left[P_i^\star(w)\right] = P_i(x) + \frac{1}{2k}\langle P_i''(x), \Sigma(x)\rangle. \tag{11}$$

Here, given square matrices A, B, we define $\langle A, B\rangle \equiv \mathrm{trace}[AB] = \sum_{k,l} a_{kl} b_{kl}$.[10]

To represent $\widetilde{L}$ in terms of F, it is useful to introduce some notation. For any collection $\{y_i\}_{i \in S}$ of scalars or vectors, let $\overline{y}(x)$ be the *logit-weighted* average, and let $\widehat{y}_i(x)$ be the relative value with respect to $\overline{y}(x)$:

$$\overline{y}(x) \equiv \sum_{l \in S} P_l(x) y_l \quad \text{and} \quad \widehat{y}_i(x) \equiv y_i - \overline{y}(x). \tag{12}$$

By definition, $\sum_{i \in S} P_i(x)\ \widehat{y}_i(x) = \sum_{i \in S} P_i(x)(y_i - \overline{y}(x)) = 0$. With these notations, the delta method yields the uniform approximation of L:

Theorem 1 (Uniform approximation). *Assume that F admits a C^4 extension to a neighborhood of X. Define the η-logit choice rule P with multiplicative corrections as follows:*

$$\widetilde{L}_i(x) = (1 + \widehat{v}_i(x) + \widehat{q}_i(x))\, P_i(x) \qquad \forall i \in S,\ x \in X, \tag{13}$$

[9] We assume that F is differentiable as desired. For simplicity, we interpret differentiability of functions defined on X via extensions of the functions to an open neighborhood of X.

[10] This approximation is impossible for the sampling best response correspondence $\mathrm{BR}^k(x) = \mathbb{E}[\mathrm{BR}(w)]$, simply because $\mathrm{BR}(\cdot)$ is generally not differentiable.

where the functions $v : X \to \mathbb{R}^n_{\geq 0}$ and $q : X \to \mathbb{R}^n$ are defined by

$$v_i(x) = \frac{1}{2k\eta^2} \, \widehat{F_i'(x)}^{\top} \, \Sigma(x) \, \widehat{F_i'(x)}, \qquad q_i(x) = \frac{1}{2k\eta} \langle F_i''(x), \Sigma(x) \rangle \tag{14}$$

with $\Sigma(x) = \operatorname{diag}[x] - xx^\top$. For $\epsilon > 0$, let $X_\epsilon = \{x \in X : x_i \geq \epsilon \; \forall i \in S\}$. Then, for any $\epsilon > 0$, $\eta > 0$, and $\zeta > 0$, there exists $k^* < \infty$ such that for all $k \geq k^*$,

$$\sup_{x \in X_\epsilon} \|L(x) - \widetilde{L}(x)\| < \begin{cases} \zeta \, \eta^{-2} & \text{if } \eta \leq 1, \\ \zeta \, \eta^{-1} & \text{if } \eta \geq 1. \end{cases} \tag{15}$$

Thus, for each fixed η, the approximated choice rule $\widetilde{L}$ can be made arbitrarily close to the true expected choice rule L as k becomes large. Conversely, Equation (15) shows that for fixed k the approximation deteriorates as η decreases. In particular, the approximation error is of smaller order than the leading correction term whenever $k\eta^2$ is sufficiently large. Accordingly, the analysis and interpretation below based on the approximated rule $\widetilde{L}$ should be understood as describing the behavior of L in regimes where k is large relative to η^{-2}.

4.2 Virtual payoff premiums and their origins

The approximation (13) yields an interpretation of SLE based on *virtual payoffs* capturing agents' effective decision biases in aggregate, in the spirit of Hofbauer and Sandholm (2002, Appendix).

Theorem 2 (Virtual payoff representations)**.** *Equilibria under the approximated choice rule $\widetilde{L}$ are equivalently represented as:*

(a) the η-logit equilibria for the "virtual" population game $\widetilde{F} = F + G$, where

$$G_i(x) = \eta \log\left(1 + \widehat{v}_i(x) + \widehat{q}_i(x)\right) \qquad \forall i \in S, x \in X \tag{16}$$

is the deterministic payoff distortion that encapsulates the role of sampling noise, provided that $1 + \widehat{v}_i(x) + \widehat{q}_i(x) > 0$ for all $i \in S, x \in X$; or

(b) the Nash equilibria for the "virtual" population game $F + G + H$, where G is the same as above and $H_i(x) \equiv -\eta \log(x_i)$.

To see Theorem 2 (a), one can simply confirm that

$$\widetilde{L}_i(x) = \frac{\exp\left(\eta^{-1}\widetilde{F}_i(x)\right)}{\sum_{j\in S}\exp\left(\eta^{-1}\widetilde{F}_j(x)\right)} \qquad \forall i \in S, x \in X. \tag{17}$$

Thus, an equilibrium under the approximated choice rule, namely $x \in X$ satisfying $x = \widetilde{L}(x)$, is nothing but an η-logit equilibrium for the virtual population game $\widetilde{F}$. Theorem 2 (b) then follows from a known result for the logit choice discussed in Section 2.4.[11]

Theorem 1 shows that actions with relatively higher values of v and/or q become exaggerated in aggregate behavior. In the language of Theorem 2, they are "virtually" preferred by agents beyond the true payoffs would imply. Then, to obtain insights into SLE, the biases introduced by the deterministic payoff distortion G, or equivalently, the correction terms $\hat{v}$ and $\hat{q}$ need to be elucidated. What are these correction terms?

Naturally, v and q originate from finite-sample variability of inferred payoffs. The payoffs inferred from a sampled distribution $w = \frac{1}{k}z$ is expanded as

$$F_i(w) \approx F_i(x) + \underbrace{\langle F_i'(x), w - x\rangle}_{\text{First-order error: } \zeta_1} + \underbrace{\tfrac{1}{2}(w-x)^\top F_i''(x)(w-x)}_{\text{Second-order error: } \zeta_2}. \tag{18}$$

Since w is a random variable, both the first-order error ζ_1 and the second-order error ζ_2 are also random variables. We have $\mathbb{E}[\zeta_1] = 0$ because $\mathbb{E}[w] = x$. On the other hand, as we note $\mathrm{Var}[w] = \frac{1}{k}\Sigma$, basic statistical identities imply

$$v_i(x) = \frac{1}{2\eta^2}\mathrm{Var}\left[\langle \widehat{F_i'(x)}, w - x\rangle\right] \quad \text{and} \tag{19}$$

$$q_i(x) = \frac{1}{2\eta}\mathbb{E}\left[(w-x)^\top F_i''(x)(w-x)\right]. \tag{20}$$

From Eqs. (18) to (20), $v_i(x)$ corresponds to the variance of the (relative) first-order error, and $q_i(x)$ corresponds to the expected second-order error.

We can thus designate $v(x)$ as the *variance premium* and $q(x)$ as the *curvature premium*. The former reflects the tendency to overweight strategies with higher variability of inferred relative payoffs, while the latter captures systematic dis-

[11]The condition $1 + \hat{v}_i(x) + \hat{q}_i(x) > 0$ in Theorem 2 (a) ensures that the multiplicative representation is well defined. In the regimes where the approximation is accurate, these correction terms are small and the condition is automatically satisfied.

tortions due to the local curvature of the payoff function. We discuss these two distortions separately below.

5 How variance matters

We first focus on variance premium $\widehat{v}$. To this end, it is useful to consider linear population games for which F_i'' vanishes and hence the curvature premium is absent. Then, $\widehat{v}_i(x)$ and $G_i(x) = \eta \log(1 + \widehat{v}_i(x))$ have the same signs.

Since $F(x) = Ax$, we have $F_i'(x) = a_i \equiv [a_{i1}, a_{i2}, \ldots, a_{in}]^\top$, where a_i the ith row of A as a column vector. Then, $\widehat{a}_i(x) = a_i - \overline{a}(x)$ represents the *relative marginal payoffs* at x under the hypothesis that all other agents apply the η-logit choice rule, since we see $\langle \widehat{a}_i(x), x \rangle = F_i(x) - \overline{F}(x)$. In linear population games, variance premium is evaluated as $v_i(x) = \frac{1}{2k\eta^2}\sigma_i(x)$ where

$$\sigma_i(x) \equiv \widehat{a}_i(x)^\top \Sigma(x) \widehat{a}_i(x) = \sum_{l \in S} (\widehat{a}_{il}(x))^2 x_l - \Big(\sum_{l \in S} \widehat{a}_{il}(x) x_l\Big)^2 \tag{21}$$

is the variance of relative marginal payoffs. Thus, we have the following observation from Theorem 1 in the context of linear population game, as expected from Eq. (19).

Observation 1. In linear population games, the approximated choice rule $\widetilde{L}$ assigns higher probability than the plain η-logit choice rule P on actions that have higher variance $\sigma_i(x)$ of relative marginal payoffs. ◀

5.1 Intrinsic bias in logit choice and variance premium

The variance premium stems from strict convexity of $\exp(\cdot)$ in the logit choice formula. As discussed, finite-sample payoff evaluation is subject to errors. Then, an upward error tend to increase the choice probability of action i more than an equally sized downward error decreases it. Specifically, since $p(\mu) = \exp(\eta^{-1}\mu)$ is increasing and convex ($p' > 0, p'' > 0$), for any symmetric perturbation $\pm\zeta$,

$$\underbrace{p(\mu + \zeta) - p(\mu)}_{\text{Upside increment}} = \int_{\mu}^{\mu+\zeta} p'(t)\mathrm{d}t \geq \int_{\mu-\zeta}^{\mu} p'(t)\mathrm{d}t = \underbrace{p(\mu) - p(\mu - \zeta)}_{\text{Downside decrement}}. \tag{22}$$

Or, in terms of Jensen's inequality, for any zero-mean noise ζ,

$$\mathbb{E}[p(\mu + \zeta)] \geq p\left(\mathbb{E}[\mu + \zeta]\right) = p(\mu). \tag{23}$$

Thus, the expected value is systematically pushed upward under zero-mean noise. The magnitude of this amplification can be explicitly estimated as[12]

$$\frac{\mathbb{E}[p(\mu+\zeta)] - p(\mu)}{p(\mu)} \approx \frac{1}{2}\frac{p''(x)}{p(x)}\operatorname{Var}[\zeta] = \frac{1}{2\eta^2}\operatorname{Var}[\zeta], \tag{24}$$

again confirming connection to the identity (19).

While the full logit choice probability $P_i(w) = \frac{p(F_i(w))}{\sum_{j\in S} p(F_j(w))}$ is not strictly convex nor concave, all $\exp(\cdot)$ terms in $P_i(w)$ are subject to this "lucky-draw" effect. Hence, *relative* magnitudes of these biases determine the net corrections, as captured by Eq. (19).

The logit choice rule is known to approximate the best response and agents' choice behavior becomes more deterministic and sensitive as η diminishes. In light of this, Equation (24) reveals that the variance premium, or the sensitivity of choice behavior to (first-order) finite-sample errors in payoff evaluation, becomes more pronounced in such environments.

5.2 Two-action games

As a concrete example, we consider a two-action linear population game

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{with} \quad \beta \equiv (a-c) - (b-d) \neq 0. \tag{25}$$

The game is a coordination game if $\beta > 0$, and anti-coordination game if $\beta < 0$.

Proposition 3. *Consider a linear population game* (25)*. Let $i, j \in S = \{1, 2\}$ with the convention being $i \neq j$. Then, $\sigma_i(x) = P_j(x)^2 \cdot \sigma_A(x)$, where $\sigma_A(x) = \beta^2 x_1 x_2 \geq 0$ with $\beta = (a-c) - (b-d) \neq 0$. In turn, variance premium satisfies*

$$\widehat{v}_i(x) = \frac{1}{2k\eta^2}\,\widehat{\sigma}_i(x) = \frac{1}{2k\eta^2} P_j(x)\,(1 - 2P_i(x))\,\sigma_A(x) \qquad \forall x \in X. \tag{26}$$

Observation 2. Variance premium vanishes ($\widehat{v}(x) = \mathbf{0}$) for $x_1 \in \{0, 1\}$ since $\sigma_A(x) = 0$ at the boundaries. This is because any sample drawn at these extreme states happen to allow agents to infer the population state correctly, and sample-dependent payoff evaluation errors cannot occur. If the population

[12] For small ζ, from $p(\mu+\zeta) \approx p(\mu) + p'(\mu)\zeta + \frac{1}{2}p''(\mu)\zeta^2$, up to the second order, $\mathbb{E}[p(\mu + \zeta)] \approx p(\mu) + p'(\mu)\mathbb{E}[\zeta] + \frac{1}{2}p''(\mu)\mathbb{E}[\zeta^2] \approx (1 + \frac{1}{\eta}\mathbb{E}[\zeta] + \frac{1}{2\eta^2}\operatorname{Var}[\zeta]) \cdot p(\mu)$. In linear population games, $\mathbb{E}[\zeta] = 0$ because $\mathbb{E}[Aw - Ax] = A\mathbb{E}[w - x] = \mathbf{0}$.

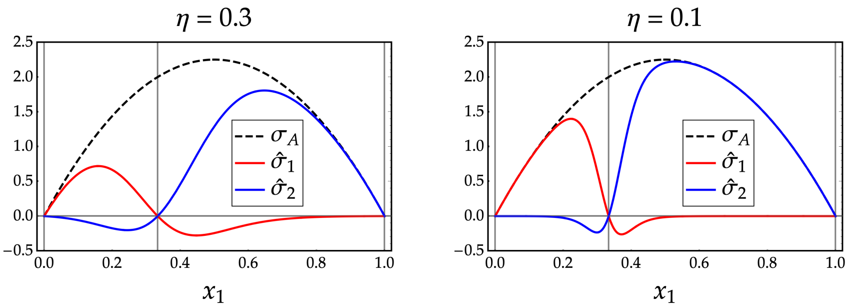

Figure 4: Illustration of Corollary 1 by the plots of $\widehat{\sigma}_i(x) = 2k\eta^2\ \widehat{v}_i(x)$ for the case $A = \left[\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right]$ with different η. The interior Nash equilibrium of the game is $x_1^* \equiv \frac{1}{3}$, and $\mathrm{br}(x) = \{2\}$ if $x < x_1^*$ and $\mathrm{br}(x) = \{1\}$ if $x_1 > x_1^*$.

state is more balanced, inference based on samples fluctuate more, and the resulting bias $\widehat{v}$ tend to become large. This is reflected by $\sigma_A(x)$ attaining its maximum at $x_1 = \frac{1}{2}$. ◀

A more interesting observation from Proposition 3 is summarized as follows:

Corollary 1 (Virtual preference for the suboptimal)**.** *Assume $\sigma_A(x) \neq 0$. Then,*

$$\widehat{v}_i(x) \gtreqless 0 \quad \Leftrightarrow \quad \frac{1}{2} \gtreqless P_i(x) \quad \Leftrightarrow \quad F_j(x) \gtreqless F_i(x), \qquad \forall i, j \in S, i \neq j, \tag{27}$$

with the same sign for the inequalities. In particular, for $x_1 \in (0, 1)$, $\widehat{v}(x) = \mathbf{0}$ if and only if x is a Nash equilibrium. Also,

$$\lim_{\eta \downarrow 0} \widehat{\sigma}_i(x) = \begin{cases} 0 & \text{if } i \in \mathrm{br}(x) \\ \sigma_A(x) & \text{if } i \notin \mathrm{br}(x) \end{cases} \qquad \forall i \in S. \tag{28}$$

Corollary 1 indicates that, in 2×2 games, the population behaves *as if* agents prefer the suboptimal option. Such a property already exists in the η-logit choice rule because it assigns a positive choice probability to the suboptimal alternative. Equation (27) indicates that sampling errors amplify this bias.

Figure 4 illustrates Corollary 1 for the case $A = \left[\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right]$ as considered in Figs. 1 and 2. For example, we can confirm that $\widehat{\sigma}_1(x) > 0$ only if $x_1 < \frac{1}{3}$ where $\mathrm{br}(x) = \{2\}$, and for this range $\widehat{\sigma}_1(x)$ approaches $\sigma_A(x)$ as $\eta \to 0$.

5.3 The loci of sampling logit equilibria

To the extent that our approximation is valid, interior SLE in two-action game can be analytically characterized. Continue with the general game (25), and assume that there is an interior Nash equilibrium satisfying

$$x_1 = x^* \equiv \frac{d-b}{\beta} \in (0,1) \quad \text{with} \quad \beta = (a-c)-(b-d) \neq 0. \tag{29}$$

If $\beta > 0$, then combined with the assumption that $x^* \in (0,1)$, the game is a coordination game with two strict Nash equilibria such that $x_1 \in \{0,1\}$. If $\beta < 0$, then the game is a *stable game* (or *contractive game*) (Hofbauer and Sandholm, 2009) and x^* is the unique Nash equilibrium. The following result characterizes the shift of the interior equilibrium induced by dual noises.[13]

Proposition 4 (Analytical approximation for the interior equilibrium)**.** *Consider a linear population game $F(x) = Ax$ where A is given by Eq. (25). Consider a generic case where $x^* > 0$ is smaller than $\frac{1}{2}$. If both $\theta \equiv \frac{1}{2k\eta^2}$ and η are sufficiently small, the interior SLE $\tilde{x} \in (0,1)$ corresponding to x^* is approximated by*

$$\tilde{x} = x^* - \frac{\eta}{\beta}\log\frac{1-x^*}{x^*} - \frac{\eta}{\beta}\log\left(1+v^*\right), \tag{30}$$

where $\beta = (a-c)-(b-d) \neq 0$ and $v^* = \theta\sigma_A(x^*) = \theta\beta^2 x^*(1-x^*)$. In particular, $\tilde{x} < x^*$ if $\beta > 0$ and $\tilde{x} > x^*$ if $\beta < 0$.

The third term in Eq. (30) captures sampling effects. If v^* is sufficiently small, the first-order approximation yields

$$-\frac{\eta}{\beta}\log\left(1+v^*\right) \approx -\frac{\eta}{\beta}\cdot v^* = -\frac{\eta}{\beta}\cdot\theta\sigma_A(x^*) = -\frac{\beta}{2k\eta}x^*(1-x^*). \tag{31}$$

This term vanishes in the "logit limit" $\theta \to 0$, that is, when $k \to \infty$ with fixed $\eta > 0$, or more generally when k grows sufficiently faster than η decreases. In this plain logit limit, only the term $-\frac{\eta}{\beta}\log\frac{1-x^*}{x^*}$ remains. Thus, this represents the basic bias introduced by the logit choice. For $\beta > 0$, the sampling effect shifts the interior fixed point toward $x_1 = 0$, consistent with the intuition in Section 5.2.

Direct computation based on Eq. (30) shows that $\tilde{x}$ approaches x^* as k increases and/or η decreases, which is natural:

[13] As discussed in the proof, we assume x^* is sufficiently smaller than $\frac{1}{2}$ only to exclude nearly symmetric cases. A symmetric result holds true for the case $x^* > \frac{1}{2}$.

Corollary 2. *The comparative statics for decreasing η and increasing k satisfy*

$$-\operatorname{sign}\frac{\partial}{\partial\eta}|\tilde{x}-x^*|<0 \quad \textit{for small } \theta=\frac{1}{2k\eta^2}, \textit{ and} \tag{32}$$

$$\operatorname{sign}\frac{\partial}{\partial k}|\tilde{x}-x^*|<0 \tag{33}$$

where we treat k as a continuous variable assuming large k.

6 How curvature matters

It remains to understand the curvature premium q. As a parsimonious example, consider a *separable game* in which each F_i depends only on x_i. For simplicity, we write $F_i(x)=F_i(x_i)$. In this class of games, we compute

$$q_i(x)=\frac{1}{2k\eta}\langle F_i''(x),\Sigma(x)\rangle=\frac{1}{2k\eta}F_i''(x_i)\cdot\Sigma_{ii}(x)=\frac{1}{2k\eta}F_i''(x_i)\cdot x_i(1-x_i), \tag{34}$$

since $[F_i''(x)]_{ij}=0$ unless $i=j$, implying the following observation.

Observation 3. Other things being equal, at each state, agents behave *as if* they prefer actions with (i) higher payoff curvature because of $F_i''(x_i)$ and/or (ii) relatively high but not too high popularity because of $x_i(1-x_i)$. ◀

This is another form of Jensen-type biases due to convexity or concavity as we have discussed in Section 5.1. Specifically, actions with larger payoff curvature F_i'' are preferred because convex payoff functions exaggerate the apparent benefits due to upside evaluation errors. Under sampling noise, convexity makes the expected payoff appear higher than the payoff at the mean, while concavity has the opposite effect. Thus, agents behave as if actions whose payoff functions are locally more convex are more attractive, even when expected payoffs are the same.[14]

Also, $x_i(1-x_i)$ is largest in the interior of the simplex and vanishes at the boundary, reflecting sampling fluctuations (as discussed in Observation 2). Sampling noise plays an important role in shaping behavior at interior population states, where each individual observation carries relatively less information about the state.

[14] The importance of payoff curvature for evolutionary incentives also appears in Hofbauer and Weibull (1996); Viossat (2015), where convex or concave transformations of payoffs generate qualitatively different evolutionary behavior. Our curvature premium can be viewed as an analogous mechanism arising from sampling noise.

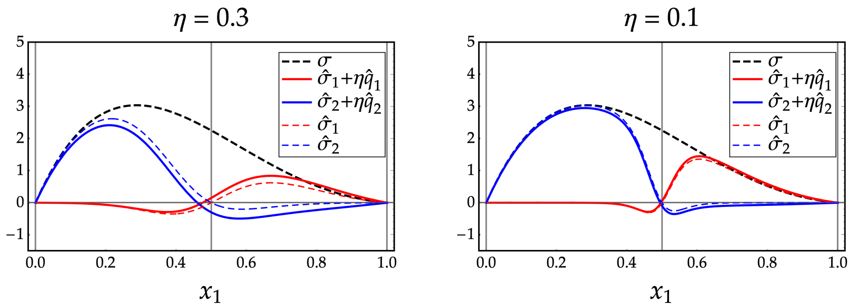

Figure 5: Illustration of Proposition 5 based on graphs of $2k\eta^2(\widehat{v}+\widehat{q}) = \widehat{\sigma}+\eta\widehat{q}$ under different values of η in a separable two-action congestion game. $\mathrm{br}(x) = \{1\}$ if $x_1 < \frac{1}{2}$ and $\mathrm{br}(x) = \{2\}$ if $x > \frac{1}{2}$. For comparison, the dashed curves show only $\widehat{\sigma}$.

For small η, the curvature premium q becomes smaller in magnitude compared to the variance premium v. This difference reflects the source of each correction in Eq. (18). The variance premium arises from the leading effect of erroneous payoff evaluation, and the curvature premium is the second-order correction.

A concrete example is provided below for separable two-action games:

Proposition 5. *For a separable two-action game,*

$$\widehat{v}_1(x) = \frac{1}{2k\eta^2} P_2(x)\left(1-2P_1(x)\right)\sigma(x) \quad \textit{and} \tag{35}$$

$$\widehat{q}_1(x) = \frac{1}{2k\eta} P_2(x)(F_1''(x_1) - F_2''(x_2))x_1x_2, \tag{36}$$

where $\sigma(x) = \left(F_1'(x_1)+F_2'(x_2)\right)^2 x_1x_2$.

The variance premium biases toward suboptimal action as discussed in Section 5.2. The curvature premiums clearly favors the action with a greater curvature at each x. Importantly, the sign of $\widehat{v}$ and $\widehat{q}$ can be different, reflecting their distinct origins.

Figure 5 considers a *congestion game* $F(x) = (-x_1, -2(x_2)^2)$ in which the unique Nash equilibrium is $x_1^* = \frac{1}{2}$. The curvature premium biases the aggregate choice toward action 1 as we see $F_1'' = 0 > -1 = F_2''$. Notably, $\sigma(x) = (5-4x_1)^2 x_1x_2$ is not symmetric about the midpoint $\frac{1}{2}$, which contrasts to linear population games (Fig. 4). This is because marginal payoffs are state dependent. While $\widehat{v}$ vanishes at x_1^* as $P_1(x^*) = \frac{1}{2}$, nonzero distortion remains at x_1^* due to $\widehat{q}$. Thus, in this congestion game with a non-zero curvature term, the population at an interior Nash equilibrium behave as if it prefer less risky option (i.e., action 1) for which expected cost is lower under sampling noise.

7 Concluding remarks

We introduced a tractable framework for analyzing strategic behavior when agents face both stochastic choice and informational constraints. Finite sampling does not simply add noise to behavior but systematically alters incentives by distorting effective payoffs. The virtual payoff representation makes these distortions explicit and yields a useful approximation for analyzing equilibrium behavior when k is large. The results suggest several directions for future work, including extensions to dynamic learning models and applications to environments where payoff observations are intrinsically limited.

One of the key feature of Oyama et al. (2015), the immediate basis for this study, is the heterogeneity of sample sizes whereby possible number of observations $k \geq 1$ is also a random variable. This possibility is abstracted away in this study because our emphasis is on the connection between the number of observations k and the accuracy of decision represented by η, as well as the resulting aggregate biases under positive η. Since the heterogeneity is the key for their equilibrium selection results, understanding the role of randomness of k in our context is an interesting extension. The examples in Section 3 suggest similar conclusions on equilibrium selection may be drawn. It is also important to explore how other predictions based on the sampling best response dynamic (e.g., Sawa and Wu, 2023; Arigapudi et al., 2024; Arigapudi and Heller, 2025) are robust against logit noise.

Finally, several apparent limitations of the present study should be noted. First, our discussion in Sections 5 and 6 concerning systematic biases rely on an approximation under large sample size k. Second, both the sample size k and the noise level η are treated as exogenous. A natural extension as a model of learning would be to endogenize these parameters by incorporating costs of information acquisition or cognitive effort. Finally, while our framework provides sharp predictions in two-action games, its analytical tractability in general n-action games, possibly in potential games (Sandholm, 2001, 2009) or stable games (Hofbauer and Sandholm, 2009), warrants further investigation.

References

Alós-Ferrer, C. and Netzer, N. (2010). The logit-response dynamics. *Games and Economic Behavior*, 68(2):413–427.

Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.

Arigapudi, S. (2020). Transitions between equilibria in bilingual games under logit choice. *Journal of Mathematical Economics*, 86:24–34.

Arigapudi, S. (2024a). Evolutionary dynamics in bilingual games. *Journal of Economic Dynamics and Control*, 165:104898.

Arigapudi, S. (2024b). Transitions between equilibria in bilingual games under probit choice. *Journal of Mathematical Economics*, 111:102956.

Arigapudi, S. and Heller, Y. (2025). Stable mixing in hawk-dove games under best experienced payoff dynamics. *Games and Economic Behavior*, 151:148–161.

Arigapudi, S., Heller, Y., and Schreiber, A. (2024). Heterogeneous noise and stable miscoordination. *American Economic Journal: Microeconomics*. Forthcoming.

Behrens, K. and Murata, Y. (2021). On quantitative spatial economic models. *Journal of Urban Economics*, 123:103348.

Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424.

Blume, L. E. (1995). The statistical mechanics of best-response strategy revision. *Games and Economic Behavior*, 11(2):111–145.

Blume, L. E. (2003). How noise matters. *Games and Economic Behavior*, 44(2):251–271.

Danenberg, T. and Spiegler, R. (2025). A representative-sampling model of stochastic choice. *Journal of Political Economy: Microeconomics*. Forthcoming.

Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071.

Foster, D. and Young, P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38(2):219–232.

Franchetti, F. and Sandholm, W. H. (2013). An introduction to *Dynamo*: Diagrams for evolutionary game dynamics. *Biological Theory*, 8(2):167–178.

Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press.

Galesloot, B. M. and Goyal, S. (1997). Costs of flexibility and equilibrium selection. *Journal of Mathematical Economics*, 28(3):249–264.

Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59(3):859–867.

Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2005). Regular quantal response equilibrium. *Experimental Economics*, 8(4):347–367.

Goyal, S. and Janssen, M. C. (1997). Non-exclusive conventions and social coordination. *Journal of Economic Theory*, 77(1):34–57.

Heller, Y. and Mohlin, E. (2018). Observations on cooperation. *The Review of Economic Studies*, 85(4):2253–2282.

Hofbauer, J. (1995). Stability for the best response dynamics. Mimeo, University of Vienna.

Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294.

Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132(1):47–69.

Hofbauer, J. and Sandholm, W. H. (2009). Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665–1693.

Hofbauer, J. and Weibull, J. W. (1996). Evolutionary selection against dominated strategies. *Journal of economic theory*, 71(2):558–573.

Immorlica, N., Kleinberg, J., Mahdian, M., and Wexler, T. (2007). The role of compatibility in the diffusion of technologies through social networks. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 75–83.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67.

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

Marden, J. R. and Shamma, J. S. (2012). Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation. *Games and Economic Behavior*, 75(2):788–808.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41.

Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, 88(4):834–847.

Osborne, M. J. and Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45(2):434–441.

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

Oyama, D. and Takahashi, S. (2015). Contagion and uninvadability in local interaction games: The bilingual game and general supermodular games. *Journal of Economic Theory*, 157:100–127.

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

Sandholm, W. H. (2001). Potential games with continuous player sets. *Journal of Economic Theory*, 97(1):81–108.

Sandholm, W. H. (2009). Large population potential games. *Journal of Economic Theory*, 144(4):1710–1725.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press.

Sawa, R. and Wu, J. (2023). Statistical inference in evolutionary dynamics. *Games and Economic Behavior*, 137:294–316.

Selten, R. and Chmura, T. (2008). Stationary concepts for experimental 2x2-games. *American Economic Review*, 98(3):938–966.

Sethi, R. (2000). Stability of equilibria in games with procedurally rational players. *Games and Economic Behavior*, 32(1):85–104.

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

Spiegler, R. (2006b). The market for quacks. *The Review of Economic Studies*, 73(4):1113–1131.

Turocy, T. L. (2005). A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51(2):243–263.

van der Vaart, A. W. (2000). *Asymptotic Statistics*. Cambridge University Press.

Viossat, Y. (2015). Evolutionary dynamics and dominated strategies. *Economic Theory Bulletin*, 3(1):91–113.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.

A Proofs

Proof of Proposition 1. We have $L^{1,\eta}(x) = \sum_{j\in S} M^1_x(e_j)\cdot P^\eta(e_j) = \sum_{j\in S}\Pi_{ij}x_j$ where $\Pi_{ij} \equiv P^\eta_i(e_j) \in (0,1)$ is the η-logit choice probability of action $i \in S$ at e_j. The equilibrium condition $x = L^{1,\eta}(x)$ then reads $x_i = \sum_{j\in S}\Pi_{ij}x_j$ or $x = \Pi x$. Because $\Pi \in \mathbb{R}^{n\times n}_{>0}$ can be seen as a transition probability matrix of an irreducible Markov chain, there exists a unique "stationary distribution" $x^* \in X$ satisfying $x^* = \Pi x^*$. Since Π is a positive matrix, x^* is proportional to the positive eigenvector of Π associated with the (Perron–Frobenius) eigenvalue 1. Furthermore, (SLD) reduces to a linear dynamical system $\dot{x} = \Pi x - x = (\Pi - I)x$. It follows via standard arguments that X is forward invariant and x^* is globally asymptotically stable in X. ■

Proof of Proposition 2. Assume $k = 2$ and $n = 2$. For brevity, let $y = x_1$. Let $z \in \{0,1,2\}$ denote how many times action-1 player is drawn in a sample of size $k = 2$. Then, $\Pr(z = 2) = x_1^2 = y^2$, $\Pr(z = 1) = 2x_1x_2 = 2y(1-y)$, and $\Pr(z = 0) = x_2^2 = (1-y)^2$. Let $q_z \in (0,1)$ be the choice probability of action 1 for each realization of z. That is, $q_0 = P^\eta_1(e_2)$, $q_1 = P^\eta_1(\frac{1}{2}(e_1+e_2))$, and $q_2 = P^\eta_1(e_1)$. Define

$$f(y) \equiv L^{k,\eta}_1(y, 1-y) - y \tag{37}$$

$$= q_2 \times y^2 + 2q_1 \times y(1-y) + q_0 \times (1-y)^2 - y \tag{38}$$

$$= (q_2 - 2q_1 + q_0)y^2 + (2(q_1 - q_0) - 1)\,y + q_0. \tag{39}$$

Then, an SLE is a solution to $f(y) = 0$ satisfying $y \in (0,1)$. Observe that $f(0) = q_0 > 0$ and $f(1) = q_2 - 1 < 0$. Because f is a quadratic function, there is a unique $y^* \in (0,1)$ that solves $f(y^*) = 0$. Furthermore, y^* is globally asymptotically stable because $\dot{y} = f(y) > 0$ for $x \in [0, y^*)$ and $\dot{y} = f(y) < 0$ for $x \in (y^*, 1]$. In fact, under $k = 2$, the sampling logit choice rule is a quadratic function for general $n \geq 2$.

Remark (Selection in a coordination game). If $F(x) = Ax$ with $A = \begin{bmatrix} s & 0 \\ 0 & t \end{bmatrix}$ ($s > t > 0$), $q_0 = p(-t)$, $q_1 = p(\frac{1}{2}(s-t))$, and $q_2 = p(s)$ with $p(\Delta) \equiv \left(1 + \exp(-\eta^{-1}\Delta)\right)^{-1}$. Since $q_2 - 2q_1 + q_0 < 0$,

$$y^* = \frac{-b - \sqrt{b^2 - 4ac}}{2a} \in (0,1) \tag{40}$$

solves $f(y^*) = 0$, with $a \equiv q_2 - 2q_1 + q_0$, $b \equiv 2(q_1 - q_0) - 1$, and $c \equiv q_0$. As $\eta \downarrow 0$, $(q_0, q_1, q_2) \to (0, 1, 1)$ and $(a, b, c) \to (-1, 1, 0)$, and hence $y^* \to 1$. ∎

Proof of Theorem 1. Let $\theta \equiv \eta^{-1} > 0$. The gradient of P_i is given by

$$P_i'(x) = \theta P_i(x)\big(F_i'(x) - F'(x)^\top P(x)\big) = \theta P_i(x)\widehat{F_i'(x)}. \tag{41}$$

Let $R_i \equiv \widehat{F_i'(x)}$ for brevity, so that $P_i' = \theta P_i R_i$. The Hessian matrix of P_i is

$$P_i''(x) = \theta R_i P_i'^\top + \theta P_i R_i' = \theta^2 P_i R_i R_i^\top + \theta P_i R_i'. \tag{42}$$

The Jacobian matrix of R_i is computed as

$$\begin{aligned}
R_i' = \big(F_i' - \big(\sum_l P_l F_l'\big)\big)' &= -\sum_l \big(F_l'(P_l')^\top + P_l F_l''\big) + F_i'' &(43)\\
&= -\theta \sum_l P_l F_l' R_l^\top + F_i'' - \sum_l P_l F_l'' &(44)\\
&= -\theta \sum_l P_l (F_l' - \overline{F'}) R_l^\top - \theta \sum_l P_l \overline{F'} R_l^\top + \widehat{F_i''} &(45)\\
&= -\theta \sum_l P_l R_l R_l^\top + \widehat{F_i''}, &(46)
\end{aligned}$$

where we note that $\sum_l P_l \overline{F'} R_l^\top = O$. Together with Eq. (42), we see

$$P_i''(x) = \theta^2 P_i\Big(R_i R_i^\top - \sum_l P_l R_l R_l^\top\Big) + \theta P_i \widehat{F_i''} = \theta^2 P_i \cdot \widehat{R_i R_i^\top} + \theta P_i \widehat{F_i''}. \tag{47}$$

Finally, from the identity $\langle bb^\top, A\rangle = \sum_{i,j} b_i b_j A_{ij} = b^\top A b$ for any $b \in \mathbb{R}^n$,

$$\frac{1}{2k}\langle P_i''(x), \Sigma(x)\rangle = \frac{1}{2k\eta^2} P_i \cdot \widehat{R_i^\top \Sigma R_i} + \frac{1}{2k\eta} P_i \cdot \widehat{\langle F_i'', \Sigma\rangle} \tag{48}$$

where we recall $\Sigma = \mathrm{diag}[x] - xx^\top$. Collecting terms, we obtain Theorem 1 (a).

The accuracy of the approximation should be discussed in relation to (k, η). The required scaling of k and η is characterized as follows. Below, $\|\cdot\|$ denotes any fixed norm on the relevant finite-dimensional spaces, and multilinear maps are equipped with the corresponding induced operator norms. The constants below may depend on this choice of norm, but only up to norm-equivalence factors. In particular, the asymptotic orders in k and η are unaffected.

Proposition A (Error estimate)**.** *Fix $\epsilon > 0$, and let $X_\epsilon \equiv \{x \in X : \min_{i \in S} x_i \geq \epsilon\}$.*

Assume that F admits a C^4 extension to an open neighborhood of X_ϵ. Then there exists a constant $C_\epsilon > 0$ such that

$$\sup_{x \in X_\epsilon} \|L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x)\| \leq \frac{C_\epsilon}{k^2} \sum_{m=1}^{4} \eta^{-m} \qquad \forall k \geq 1,\ \forall \eta > 0. \tag{49}$$

In particular,

$$\sup_{x \in X_\epsilon} \|L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x)\| \leq C_\epsilon k^{-2} \eta^{-4}, \qquad 0 < \eta \leq 1, \tag{50}$$

$$\sup_{x \in X_\epsilon} \|L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x)\| \leq C_\epsilon k^{-2} \eta^{-1}, \qquad \eta \geq 1. \tag{51}$$

Hence, the approximation error is negligible relative to $k^{-1}\eta^{-2}$ in the regime $k\eta^2 \to \infty$ for $\eta \leq 1$, and relative to $k^{-1}\eta^{-1}$ whenever $k \to \infty$ for $\eta \geq 1$.

Proof. Let $\Delta \equiv w - x$. By construction, $\mathbb{E}[\Delta] = 0$ and $\mathbb{E}[\Delta\Delta^\top] = \frac{1}{k}\Sigma(x)$. Also, standard moment bounds yield

$$\sup_{x \in X} \left\| \mathbb{E}[\Delta^{\otimes 3}] \right\| = O(k^{-2}) \quad \text{and} \quad \sup_{x \in X} \mathbb{E}[\|\Delta\|^4] = O(k^{-2}). \tag{52}$$

Next, we recall $P = \Lambda^\eta \circ F$, where Λ^η is the η-logit map. Since F has a C^4 extension to a neighborhood of X_ϵ, and since the mth derivatives of Λ^η are of order η^{-m}, there is a constant $C'_\epsilon > 0$ satisfying

$$\sup_{x \in X_\epsilon} \|D^m P(x)\| \leq C'_\epsilon \sum_{r=1}^{m} \eta^{-r}, \qquad m = 3, 4. \tag{53}$$

Applying Taylor's theorem of order three to each coordinate of $P(x + \Delta)$ around x and taking expectations, we obtain

$$L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x) = \frac{1}{6} \langle D^3 P(x), \mathbb{E}[\Delta^{\otimes 3}] \rangle + R(x), \tag{54}$$

where the remainder R satisfies

$$\|R(x)\| \leq C''_\epsilon \sup_{y \in X_\epsilon} \|D^4 P(y)\| \, \mathbb{E}[\|\Delta\|^4] \tag{55}$$

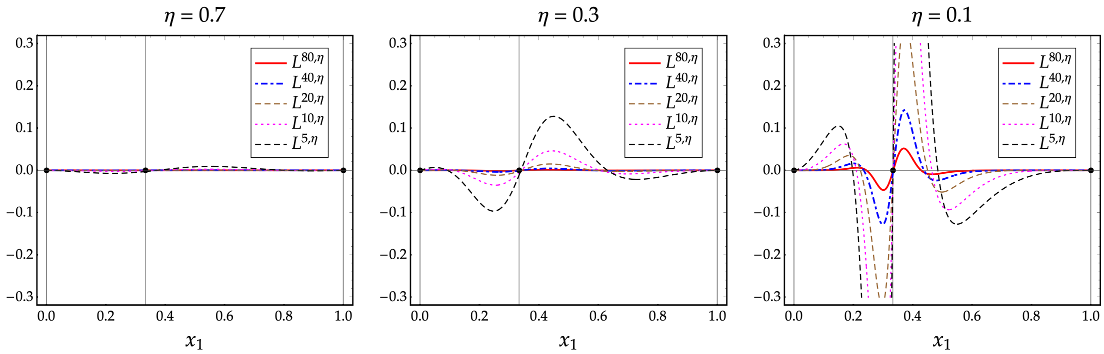

Figure 6: Approximation errors for the choice probability of action 1 ($L_1^{k,\eta} - \widetilde{L}_1$) under different η and k. The coordination game (7) with $(s,t) = (2,1)$ as considered in Figs. 1 and 2 is assumed. Black markers show the Nash equilibria.

for some constant $C''_\epsilon > 0$. Combining this with (52) and (53) gives

$$\sup_{x \in X_\epsilon} \left\| L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x) \right\| \leq \frac{C_\epsilon}{k^2} \sum_{m=1}^{4} \eta^{-m} \tag{56}$$

for some constant $C_\epsilon > 0$, proving the first claim (49). For Eqs. (50) and (51), if $\eta \in (0,1]$, then $\sum_{m=1}^{4} \eta^{-m} \leq 4\eta^{-4}$, so

$$\sup_{x \in X_\epsilon} \left\| L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x) \right\| \leq \frac{4C_\epsilon}{k^2 \eta^4}. \tag{57}$$

Similarly, if $\eta \geq 1$, then $\sum_{m=1}^{4} \eta^{-m} \leq 4\eta^{-1}$, and therefore

$$\sup_{x \in X_\epsilon} \| L^{k,\eta}(x) - \widetilde{L}^{k,\eta}(x) \| \leq \frac{4C_\epsilon}{k^2 \eta}. \tag{58}$$

This completes the proof. ■

Theorem 1 is a slightly informal version of Proposition A. Figure 6 illustrates how the accuracy of the approximation deteriorates when k is small, especially for relatively small η. ■

Proof of Theorem 2. Let $Z(x) \equiv \sum_{j \in S} \exp(\eta^{-1} F_j(x))$. Let $\delta \equiv \widehat{v} + \widehat{q}$. Then,

$$\begin{aligned} \sum_{j \in S} (1 + \delta_j(x)) \exp\left(\eta^{-1} F_j(x)\right) &= Z(x) \sum_{j \in S} (1 + \delta_j(x)) \frac{\exp\left(\eta^{-1} F_j(x)\right)}{Z(x)} && (59) \\ &= Z(x) \sum_{j \in S} (1 + \delta_j(x)) P_j(x) = Z(x) && (60) \end{aligned}$$

since $\sum_{j\in S}\delta_j(x)P_j(x)=0$. Define $\widetilde{F}_i(x)\equiv F_i(x)+\eta\log(1+\delta_i(x))$. Then,

$$\frac{\exp\left(\eta^{-1}\widetilde{F}_i(x)\right)}{\sum_{j\in S}\exp\left(\eta^{-1}\widetilde{F}_j(x)\right)}=\frac{(1+\delta_i(x))\exp\left(\eta^{-1}F_i(x)\right)}{\sum_{j\in S}(1+\delta_j(x))\exp\left(\eta^{-1}F_j(x)\right)} \tag{61}$$

$$=\frac{(1+\delta_i(x))\exp\left(\eta^{-1}F_i(x)\right)}{Z(x)} \tag{62}$$

$$=(1+\delta_i(x))P_i(x)=\widetilde{L}_i(x). \tag{63}$$

Thus, the condition $x=\widetilde{L}(x)$ is nothing but the η-logit equilibrium condition for the virtual payoff $\widetilde{F}$, showing part (a). A direct application of known results for the logit choice yields (b) (e.g., Hofbauer and Sandholm, 2007, Appendix). The requirement that $1+\delta>0$ hols true if k is sufficiently large. ■

Proof of Proposition 3. Let $A_1=(a,b)^\top$ and $A_2=(c,d)^\top$. Define $A_0\equiv A_1-A_2=(a-c,b-d)^\top$, and denote $p_1=P_1(x)$ and $p_2=P_2(x)$. Let $R_i\equiv\widehat{F_i'(x)}$,

$$R_1=A_1-(p_1A_1+p_2A_2)=p_2(A_1-A_2)=p_2A_0, \tag{64}$$

$$R_2=A_2-(p_1A_1+p_2A_2)=-p_1(A_1-A_2)=-p_1A_0. \tag{65}$$

Let $\sigma_A(x)\equiv A_0^\top\Sigma A_0=((a-c)-(b-d))^2\cdot x_1x_2\geq 0$, with equality iff $a-c=b-d$ or $x_1x_2=0$. Then, we confirm

$$\sigma_1=R_1^\top\Sigma R_1=p_2^2A_0^\top\Sigma A_0=p_2^2\sigma_A, \tag{66}$$

$$\sigma_2=R_2^\top\Sigma R_2=p_1^2A_0^\top\Sigma A_0=p_1^2\sigma_A, \tag{67}$$

$$p_1\sigma_1+p_2\sigma_2=p_1p_2^2\sigma_A+p_2p_1^2\sigma_A=(p_1+p_2)p_1p_2\sigma_A=p_1p_2\sigma_A. \tag{68}$$

Thus, $\widehat{\sigma}_1(x)=p_2(p_2-p_1)\sigma_A$ and $\widehat{\sigma}_2(x)=p_1(p_1-p_2)\sigma_A$. ■

Proof of Proposition 4. For simplicity, we use x_1 as the state variable. Note that $F_1(x_1)-F_2(x_1)=\beta(x_1-x^*)$. To marginally economize on notation, let $\delta_i\equiv\widehat{v}_i$.

For now, assume that $\beta>0$. Then, $\mathrm{br}(x_1)=\{1\}$ if $x_1>x^*$ and $\mathrm{br}(x_1)=\{2\}$ if $x_1<x^*$. Corollary 1 implies

$$(\delta_1(x_1),\delta_2(x_1))=\begin{cases}(0,\delta(x_1)) & \text{if } x_1>x^*\\ (0,0) & \text{if } x_1=x^*\\ (\delta(x_1),0) & \text{if } x_1<x^*,\end{cases} \tag{69}$$

where we set $\delta(x_1) \equiv \theta\sigma_A = \theta\beta^2 x_1(1-x_1)$. Let $y \in (0,1)$ denote the interior fixed point of $\widetilde{L}(x_1)$. Then,

$$\frac{y}{1-y} = \frac{\widetilde{L}_1(y)}{\widetilde{L}_2(y)} = \exp\left(\frac{\beta(y-x^*)}{\eta}\right) \cdot \frac{1+\delta_1(y)}{1+\delta_2(y)}. \tag{70}$$

Taking log of both sides,

$$\log\frac{y}{1-y} = \frac{\beta}{\eta}(y-x^*) + \rho(y), \quad \rho(y) = \begin{cases} -\log(1+\delta(y)) & \text{if } y > x^* \\ 0 & \text{if } y = x^* \\ \log(1+\delta(y)) & \text{if } y < x^*. \end{cases} \tag{71}$$

Set $\epsilon \equiv y - x^*$ and assume $\epsilon = O(\eta)$. From the first-order expansion of Eq. (71),

$$\log\frac{x^*}{1-x^*} + \frac{\epsilon}{x^*(1-x^*)} = \frac{\beta}{\eta}\epsilon + \rho(y). \tag{72}$$

It is noted that the special case $y = x^* \in (0,1)$ occurs if and only if $x^* = \frac{1}{2}$ because Eq. (70) implies $\frac{y}{1-y} = 1$. Then, Equation (72) yields $\epsilon = 0$.

For $y \neq x^*$, for sufficiently small η, up to the first order of η,

$$\begin{aligned}
\epsilon &= \left(\log\frac{x^*}{1-x^*} - \rho(y)\right)\left(\frac{\beta}{\eta} - \frac{1}{x^*(1-x^*)}\right)^{-1} && (73)\\
&= \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} - \rho(y)\right)\left(1 - \frac{\eta}{\beta}\frac{1}{x^*(1-x^*)}\right)^{-1} && (74)\\
&= \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} - \rho(y)\right)(1+O(\eta)) \quad (\because (1-c)^{-1} = 1 + c + O(c^2)) && (75)\\
&= \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} - \rho(y)\right). && (76)
\end{aligned}$$

To examine the dependence on y, define

$$\epsilon_\pm = \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} \pm \log(1+\delta^*)\right) \tag{77}$$

where $\delta^* \equiv \delta(x^*) = \theta\beta^2 x^*(1-x^*) > 0$. Note that $\log(1+\delta^*) > 0$, so that $\epsilon_+ > \epsilon_-$ because we assume $\beta > 0$. It is also remarked that we did not expand $\delta(y)$ around x^* explicitly in Eq. (72) because it amounts to considering a second-order term of η and thus irrelevant for the first-order approximation. From

Eq. (71) and $\beta > 0$, we have ϵ_+ if $y > x^*$, and ϵ_- if $y < x^*$. Because $\epsilon = y - x^*$, we need to check the consistency conditions $\epsilon_+ > 0$ and $\epsilon_- < 0$. We see

$$\epsilon_+ > 0 \quad \Leftrightarrow \quad \log \frac{x^*}{1-x^*} + \log(1+\delta^*) > 0 \quad \Leftrightarrow \quad x^* > x_+ \equiv \frac{1}{2+\delta^*}, \tag{78}$$

$$\epsilon_- < 0 \quad \Leftrightarrow \quad \log \frac{x^*}{1-x^*} - \log(1+\delta^*) < 0 \quad \Leftrightarrow \quad x^* < x_- \equiv \frac{1+\delta^*}{2+\delta^*}. \tag{79}$$

If $x_- < x^* < x_+$, both possibilities remain valid. However, for small δ^*, this requires x^* to be sufficiently close to $\frac{1}{2}$, which is not satisfied for generic games. Assuming that θ (and thus δ^*) is sufficiently small, only ϵ_+ is valid if x^* is sufficiently smaller than $\frac{1}{2}$ and hence than x_+.

If $\beta < 0$, $\mathrm{br}(x_1) = \{2\}$ if $x_1 > x^*$ and $\mathrm{br}(x_1) = \{1\}$ if $x_1 < x^*$. Repeating the same line of logic, again, only ϵ_- is valid if x^* is sufficiently smaller than $\frac{1}{2}$. Thus, ϵ_- is the only possibility irrespective of the sign of β. Symmetric arguments show that only ϵ_+ is valid if instead x^* is sufficiently larger than $\frac{1}{2}$. ■

Proof of Proposition 5. For the variance term, Equation (26) is applicable as we replace $\sigma_A(x)$ with $\sigma_{F'(x)}(x) = (F_1'(x_1) + F_2'(x_2))^2 x_1 x_2$. For the curvature term,

$$\begin{aligned} F_1''(x_1)x_1(1-x_1) - p_1 F_1''(x_1)x_1(1-x_1) - p_2 F_2''(x_2)x_2(1-x_2) \\ &= (1-p_1)F_1''(x_1)x_1(1-x_1) - p_2 F_2''(x_2)x_2(1-x_2) \\ &= (1-p_1)(F_1''(x_1) - F_2''(x_2))x_1(1-x_1), \end{aligned}$$

as $p_2 = 1 - p_1$ and $x_1(1-x_1) = x_2(1-x_2)$. ■

Supplementary Materials for

"Sampling Bias in Logit Choice"

March 10, 2026

Minoru Osawa

B Bilingual game

The *bilingual game* (Galesloot and Goyal, 1997; Goyal and Janssen, 1997) models language coordination with a bilingual option, and has been an important subject of study in the literature on learning and contagion in games (e.g., Immorlica et al., 2007; Oyama and Takahashi, 2015; Arigapudi, 2020, 2024b,a).

Consider the linear population game with the base payoff matrix being

$$A = \begin{bmatrix} 1+g & 0 & 1+g \\ 1 & 1 & 1 \\ 1+g-c & 1-c & 1+g-c \end{bmatrix} \qquad \left(0 < g < 1,\ 0 < c < \tfrac{g}{1+g}\right). \tag{80}$$

The game (80) is a variant of the bilingual game in which agents choose among three actions: adopting technology 1, adopting technology 2, or using a compatible interface that allows interaction with both technologies. Technology 1 yields a higher coordination payoff than technology 2: when two agents using technology 1 meet they obtain $1+g$, whereas coordination on technology 2 yields payoff 1, with $0 < g < 1$. The compatibility option enables interaction with either technology but incurs a cost $c > 0$. Under the parameter restriction $0 < c < \frac{g}{1+g}$, the compatibility option is viable but strictly costly. It allows agents to hedge against coordination risk while remaining dominated by successful coordination on the superior technology 1.

Figure 7 depicts the phase portraits of the four dynamics. Under the best response dynamic (BRD), trajectories eventually reach the superior pure state e_1 (Fig. 7A), except for those starting close to e_2. In particular, states with a substantial share of the compatibility action tend to move first toward higher use of action 3, and subsequently toward action 1, reflecting the transitional role of the bridge technology. When sampling noise is introduced, the qualitative geometry of the flow changes, yielding a sharp selection (Fig. 7B), just as in Young's game.

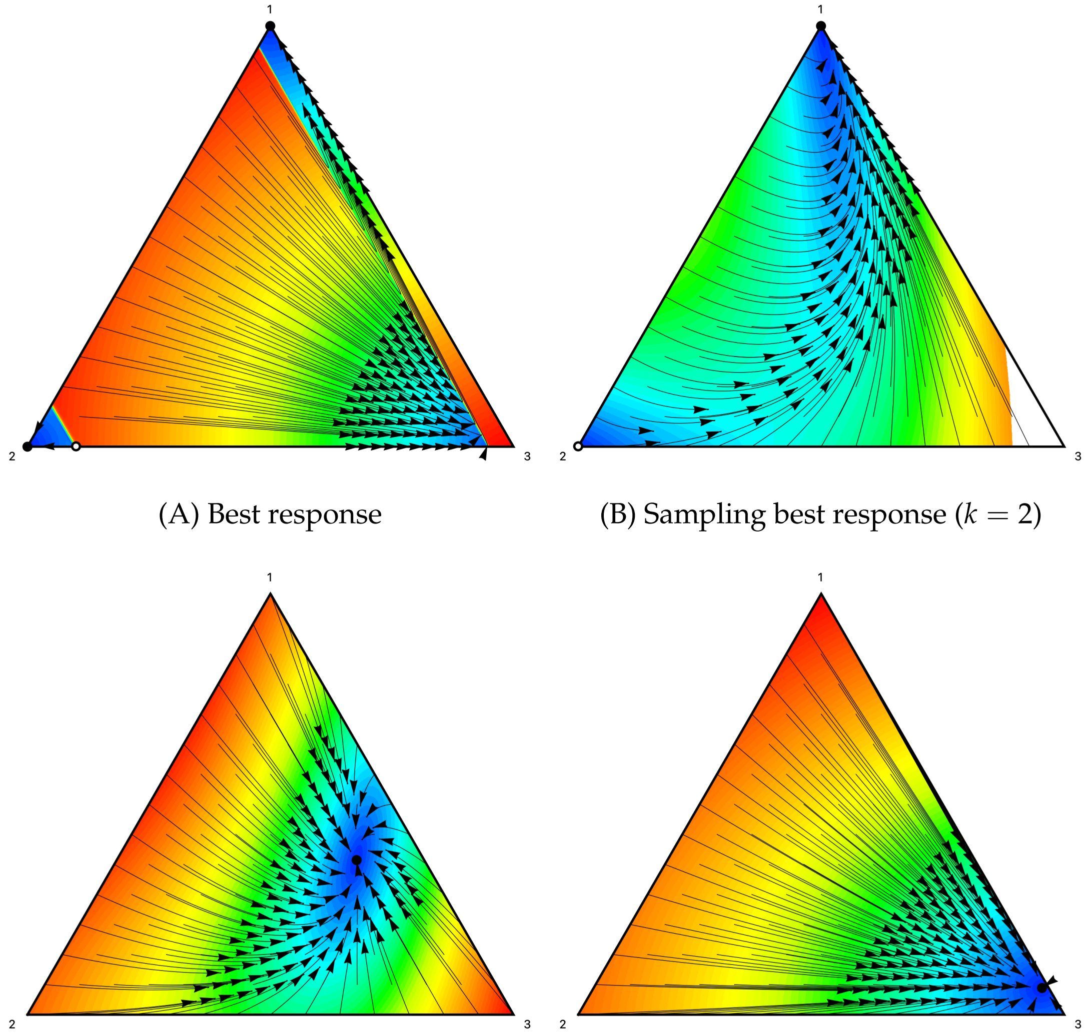

(A) Best response

(B) Sampling best response ($k = 2$)

(C) Logit ($\eta = 0.3$)

(D) Sampling logit ($k = 2, \eta = 0.3$)

Figure 7: Phase diagrams of the four dynamics in the bilingual game (80) for the case $g = 0.5$ and $c = 0.05$.

Both the logit dynamic and the sampling logit dynamic admit interior stationary states (Figs. 7C and 7D). Notably, although e_3 is nevern an equilibrium of the game, the sampling logit dynamic (SLD) converges to a neighborhood of e_3. This feature contrasts with the best response dynamics, where trajectories ultimately leave that region and converge to e_1. A systematic characterization of such approximate stationary states and their dependence on the parameters (k, η) for specific games is left for future work.

C Perturbed potential

The connection between *large-population potential games* (Sandholm, 2001, 2009) is of interest. Equilibrium problems in this class of games are represented as maximization problems of scalar-valued functions over $x \in X$. The question is that whether there is an appropriately defined scalar-valued function representing sampling effects. The answer is yes, albeit in an approximate sense.

We focus on general two-action population games. Unfortunately, generalizations beyond this class of games appears to be difficult. In the two-action case, a *potential function* $f : X \to \mathbb{R}$ is a function that satisfies

$$\frac{\partial f(x)}{\partial x_1} - \frac{\partial f(x)}{\partial x_2} = F_1(x) - F_2(x) \qquad \forall x \in X. \tag{81}$$

In fact, with $y(t) \equiv (t, 1-t)$, the following function satisfies Eq. (81):

$$f(x) \equiv \int_0^{x_1} (F_1(y(t)) - F_2(y(t)))\, \mathrm{d}t. \tag{82}$$

The Nash equilibria of the game F are known to coincide with the *stationary points* of the maximization problem $\max_{x \in X} f(x)$, denoted by $\mathrm{SP}(f)$, satisfying the first-order necessary conditions for optimality (including not only local maxima but also saddle points and local minima).

Likewise, it is known that η-logit equilibria of F correspond to the stationary points of the maximization problem of the following *perturbed potential function* (Hofbauer and Sandholm, 2002, 2007):

$$f^\eta(x) \equiv f(x) + h(x) \quad \text{where} \quad h(x) \equiv -\eta \sum_{i \in S} x_i \log x_i. \tag{83}$$

Here, $h : X \to \mathbb{R}$ is the entropy function with convention $0 \log 0 \equiv 0$.

Combining these facts and Theorem 2, it is straightforward to see that SLE in a game can be represented by an optimization problem:

Proposition 7. *Consider a two-action population game F. Define the perturbed potential function $f^{k,\eta} : X \to \mathbb{R}$ by*

$$f^{k,\eta}(x) \equiv \int_0^{x_1} \left(\widetilde{F}_1(y(t)) - \widetilde{F}_2(y(t))\right) \mathrm{d}t + h(x) \tag{84}$$

with $y(t) \equiv (t, 1-t) \in X$ and the virtual payoff function $\widetilde{F}$ defined in Theorem 2.

Then, the set of fixed points of the approximated choice rule $\widetilde{L}$ in Eq. (13) is $\mathrm{SP}(f^{k,\eta})$.

Proof of Proposition 7. The slope of $f^{k,\eta}$ along the tangent space of X satisfy

$$\mathrm{d}f \equiv \frac{\partial f^{k,\eta}(x)}{\partial x_1} - \frac{\partial f^{k,\eta}(x)}{\partial x_2} = \widetilde{F}_1(x) - \widetilde{F}_2(x) - \eta \log x_1 - \eta + \eta \log x_2 + \eta. \quad (85)$$

Thus, the corner solutions $x_1 = 0$ or $x_2 = 0$ cannot be a stationary point of the maximization problem $\max_{x \in X} f^{k,\eta}(x)$, because $\mathrm{d}f \to \infty$ as $x_1 \downarrow 0$ and $-\mathrm{d}f \to \infty$ as $x_2 \downarrow 0$. Thus, every stationary point must be positive and satisfies $\mathrm{d}f = 0$. Solving for $\mathrm{d}f = 0$, we have $\eta \log \frac{x_1}{x_2} = \widetilde{F}_1(x) - \widetilde{F}_2(x)$, which reduces to the condition $x_i = \widetilde{L}_i(x)$, showing that x is must be a fixed point of $\widetilde{L}$. ■

It should be reiterated that the connection between SLE and $\mathrm{SP}(f^{k,\eta})$ is of an approximate sense. That is, $\mathrm{SP}(f^{k,\eta})$ is close to SLE only if the approximation $L^{k,\eta} \approx \widetilde{L}$ is sufficiently good. Nonetheless, under this hypothesis, Proposition 7 yields a simple dynamic characterization of SLE by a direct application of Theorem 3.2 in Hofbauer and Sandholm (2007):

Proposition 8. *Consider a two-action population game F. For $\eta > 0$ and $k \gg \eta$, let the approximated (k, η)-sampling logit dynamic be defined by $\dot{x} = \widetilde{L}(x) - x$. Then, every solution trajectory of the dynamic converges to connected subsets of* $\mathrm{SP}(f^{k,\eta})$. *If* $\mathrm{SP}(f^{k,\eta})$ *is a singleton, then it is globally asymptotically stable.*

Proof of Proposition 8. From Theorem 2, the approximated dynamic can be seen as the η-logit dynamic in the modified population game $\widetilde{F}$. Since $f^{k,\eta}$ is the perturbed potential function associated to this setting, it is a strictly increasing Lyapunov function for the dynamic. From this, Theorem 3.2 of Hofbauer and Sandholm (2007) applies. ■

We can rearrange $f^{k,\eta}$ such that $f^{k,\eta}(x) = f(x) + g(x) + h(x)$ where

$$g(x) \equiv \int_0^{x_1} (G_1(y(t)) - G_2(y(t)))\mathrm{d}t. \quad (86)$$

Since $f^{k,\eta}(x) = f^{\eta}(x) + g(x)$, it is observed that g is the new perturbation that encapsulates the aggregate impacts of sampling under logit choice.

The shape of g is of interest because it represents how the extrema of the potential function f are shifted in $f^{k,\eta}$. For example, the entropy function h is maximized at $\bar{x} \equiv (\frac{1}{2}, \frac{1}{2})$, so that extrema of f^{η} must be shifted toward this $\bar{x}$, and in the extreme case $\eta \to \infty$, f^{η} is maximized at $\bar{x}$.

For illustration, we again consider the general 2×2 linear population game (25). With $\beta = (a-c) - (b-d)$, the potential in terms of x_1 can be written as $f(x) = \frac{\beta}{2}(x_1 - x_1^*)^2$ where $x_1^* = \frac{d-b}{\beta}$ and we assume $x_1^* \in (0,1)$. Depending on the sign of β, the potential function f is globally minimized or maximized at the interior Nash equilibrium. We have the following characterization:

Proposition 9. *Consider the general linear two-action game (25) and assume that there is an interior Nash equilibrium x^* such that $x_1^*, x_2^* \in (0,1)$. If $\beta > 0$, g as a function on X is a strictly quasiconcave and maximized at x^*. If $\beta = 0$, $g(x) = 0$ for all $x \in X$. If $\beta < 0$, g on X is q strictly quasiconvex and minimized at x^*.*

Proof of Proposition 9. We note that g is continuous and differentiable.

If $\beta > 0$, the game is a coordination game. Corollary 1 implies that, for $0 < x_1 < 1$,

$$\frac{\partial g(x)}{\partial x_1} - \frac{\partial g(x)}{\partial x_2} = G_1(x) - G_2(x) \begin{cases} > 0 & \text{if } x_1 < x^* \ (\mathrm{br}(x) = \{2\}) \\ = 0 & \text{if } x_1 = x^* \ (\mathrm{br}(x) = \{1,2\}) \\ < 0 & \text{if } x_1 > x^* \ (\mathrm{br}(x) = \{1\}). \end{cases} \tag{87}$$

Thus, g is a unimodal function whose maximum is attained at $x_1 = x^*$.

Similarly, if $\beta < 0$, $-g$ becomes unimodal. That is, g is locally maximized at the boundaries $x_1 = 0$ and $x_1 = 1$, and globally minimized at $x_1 = x^*$.

Also, Observation 2 implies that $\widehat{v}(x) = \mathbf{0}$ if $x_1 \in \{0,1\}$, and hence $G_1(x) - G_2(x) = 0$ at $x_1 = 0$ and $G_2(x) - G_1(x)$ at $x_1 = 1$. That is, the directional derivatives of g vanishes at the boundaries of X.

Summing up, g is not concave nor convex, except for the degenerate case $g \equiv 0$ where $\beta = 0$; g is strictly quasiconcave (quasiconvex) if $\beta > 0$ ($\beta < 0$). ■

Figure 8 illustrates the perturbed potential $f^{k,\eta}$ as well as the perturbations g and h, from which we confirm quasiconvexity or quasiconcavity of g for respective cases, as well as nonconvexity. The slope of g vanishes also at the boundaries because Proposition 3 implies $G_1(x) = G_2(x) = 0$ if $x_1 x_2 = 0$.

The sampling-induced perturbation g inherits the structure of the underlying payoff environment. This contrasts to the entropy h, which is always maximized at the uniform state $(\frac{1}{2}, \frac{1}{2})$ independently of the game. This reflects the fundamentally different sources of perturbation. The entropy h captures idiosyncratic, independent noise that does not depend on the game, whereas the sampling-induced perturbation g reflects systematic, state-dependent distortions arising

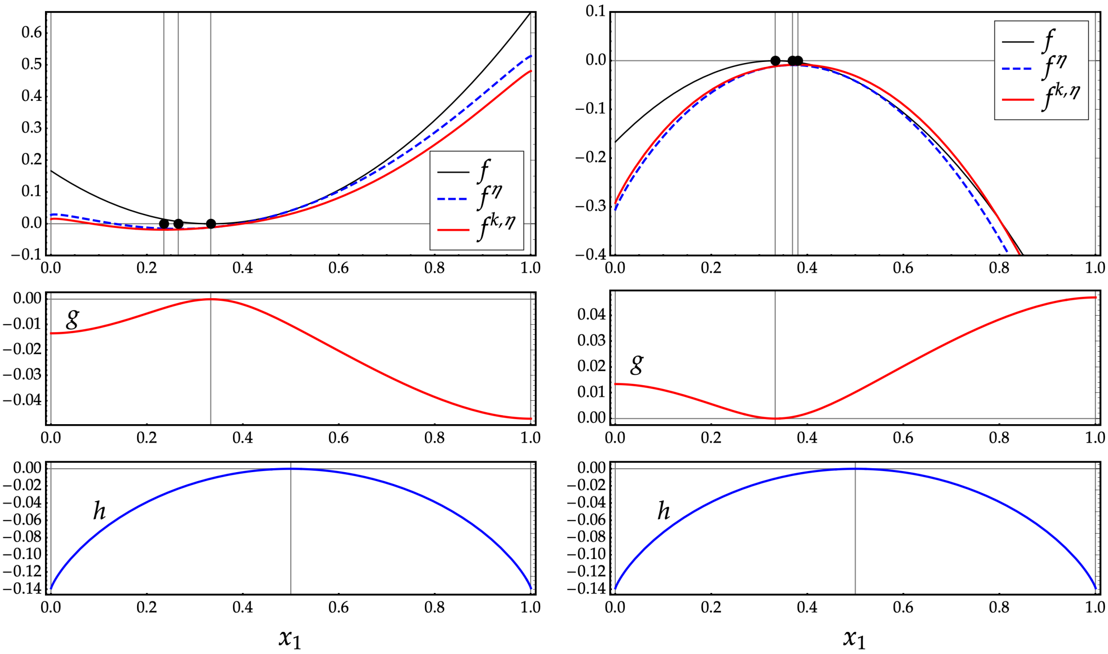

(A) Coordination game ($\beta > 0$)   (B) Anti-coordination game ($\beta < 0$)

Figure 8: Potential f, perturbed potentials f^{η} and $f^{k,\eta}$, and perturbations g and h for linear population game with $A = \left[\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right]$ (left panel) and $-A$ (right panel). We set $(k, \eta) = (40, 0.2)$. The unique mixed Nash equilibrium is $x_1^* = \frac{1}{3}$. The corresponding interior stationary points of f, f^{η}, and $f^{k,\eta}$ are indicated by black markers. For ease of comparison, g and h are vertically shifted so that each equals zero at its extremum.

from the finite-sample evaluation of payoffs.