Equilibrium Distortion with Dual Noise: The Sampling Logit Approach

Minoru Osawa

Institute of Economic Research, Kyoto University

osawa.minoru@gmail.com

October 1, 2025

Preliminary Draft

Decisions are rarely based on perfect information or flawless evaluation. This study introduces the sampling logit choice, a unified large-population framework that combines finite sampling with idiosyncratic decision errors. A central finding is that equilibria in this setting can be represented as logit equilibria with endogenously distorted payoffs. Two systematic distortions arise: a *variance premium*, which favors actions with higher payoff variability across samples, and a *curvature premium*, which reflects convexity or concavity of payoff functions. Together, these forces bias aggregate play in predictable directions, offering a structural account of equilibrium distortion under dual noise. Examples demonstrate how this framework sharpens equilibrium selection and delivers testable comparative statics.

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

JEL Classification: C72, C73.

I thank Srinivas Arigapudi and Jonathan Newton for their comments. I am grateful to the late William H. 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 often involves bounded rationality from two distinct channels of noise. The first is *stochastic choice*: even when payoffs are perfectly observed, agents may make mistakes or respond probabilistically reflecting idiosyncratic factors. The second is *informational constraint*: agents often have access to only a small or noisy sample of the environment, which induces systematic distortions in perceived payoffs.

Extant approaches in game theory tend to address these factors in isolation. On the one hand, payoff-perturbation models such as the *quantal response equilibrium* (QRE) (McKelvey and Palfrey, 1995, 1998; Goeree et al., 2005) introduce idiosyncratic randomness in choice as in random utility models (McFadden, 1974), but assume agents evaluate all available alternatives with full information. On the other hand, *sampling* models assume agents have limited information in that they observe only a subset of the environment, but usually posit deterministic best response behavior to those observations (e.g., Osborne and Rubinstein, 1998, 2003; Salant and Cherry, 2020). This dichotomy leaves a gap in our understanding of boundedly rational behavior, since in realistic environments decision makers often face both imprecise choice and imperfect information at the same time.

This study proposes the *sampling logit choice*, a unified framework that combines finite sampling with logit-style stochastic choice. Formally, under the (k, η)-sampling logit choice rule, each agent draws k independent samples of opponents' actions from the population, evaluates the resulting sample-based payoffs, and then selects an action according to the logit choice rule with noise level $\eta > 0$. A *sampling logit equilibrium* (SLE) is then defined as a fixed point of this process: the population state must coincide with the aggregate distribution of actions generated by agents who, given their informational limitation k and decisional imprecision η, choose according to this rule.

This framework naturally converges to extant frameworks at the extremes. As the sample size k grows large while η is fixed, agents effectively observe the true population state and SLE reduces to the standard logit QRE. Conversely, as decision noise η vanishes under a fixed k, SLE approaches *sampling equilibrium* (Osborne and Rubinstein, 2003), or equivalently, stationary points of the *sampling best response dynamic* (Oyama et al., 2015). By varying (k, η), the model spans a continuum between fully informed noisy optimization and limited-information best response.

In fact, the combination of parameters (k, η) provides a natural interpretation of different behavioral regimes. A higher η combined with small k appears to be a natural assumption in a complex or unfamiliar setting, so that players make more errors or idiosyncratic choice relying only on a few observations or analogies. In contrast, in a well-understood environment or after sufficient learning, η would be low and k effectively large, pushing the equilibrium outcome closer to rational prediction. The sampling logit equilibrium can thus offer a flexible mapping from environmental conditions to observable play, unifying how we think about bounded rationality across different contexts.

Our main contribution is to uncover systematic distortions induced by finite sampling under logit choice. Equilibria can be represented as if players were optimizing distorted "virtual" payoffs, where two distinct bias forces emerge. The *variance premium* favors strategies whose payoffs fluctuate more strongly across samples, while the *curvature premium* arises from convexity or concavity in payoff functions shifting expected payoffs upward or downward. Together, they bias aggregate play in predictable directions, sharpen equilibrium refinement, and yield empirically testable comparative statics. Importantly, these distortions are derived through an application of the *delta method* (e.g., van der Vaart, 2000, Ch.3), a classical tool from statistics that has rarely been exploited in game theory, which enables a tractable approximation of stochastic sampling effects.

Beyond these qualitative insights, the paper establishes several rigorous properties of SLE for special cases that suggest its role in equilibrium refinement. If each agent observes just one or two opponents, the SLE is unique and globally attracting under the *sampling logit dynamic* naturally associated to our choice rule. As logit noise diminishes, the unique SLE in these cases converges to the risk-dominant Nash equilibrium of the game, which corroborates with the literature and provides another foundation for why certain equilibria might emerge when players have limited information about others.

Related literature

Our framework connects to four strands of literature: (i) extensions of logit QRE, (ii) models of finite sampling, (iii) stochastic stability and finite-population dynamics, and (iv) the "virtual payoff" and rational inattention methods.

We extend the QRE framework by introducing finite sampling. The QRE framework (McKelvey and Palfrey, 1995, 1998; Goeree et al., 2005) introduces stochastic choice into equilibrium analysis in games. In our framework, the introduction of sampling noise yields new insights not captured by logit QRE, such as a systematic bias toward higher-variance strategies. Experimental comparisons of different stationary concepts including QRE have shown that models incorporating sampling often fit observed play better than QRE (Selten and Chmura, 2008), highlighting the importance of sampling as a source of noisy behavior. Our model is suited to settings where players have limited experience or observational data.

Our framework directly builds on the literature studying limited observation in games, especially the *sampling best response dynamic* (Oyama et al., 2015). In their evolutionary model, players occasionally update by taking a sample of k opponents' actions and best-responding to the payoffs inferred by the sampled action distribution. We incorporate logit errors into this setting. This merger yields a smoother dynamic that facilitates analysis and interpretations through tractable approximations. Other notable models that incorporate deterministic best response behavior under some forms of sampling include Osborne and Rubinstein (1998, 2003), Spiegler (2006a,b), and Salant and Cherry (2020), to note a few. In evolutionary context, we should mention Sethi (2000), Sandholm (2001), Mantilla et al. (2020), Arigapudi et al. (2024), and Sawa and Wu (2023), among others.

Erroneous choices or sampling in a finite population has been one of the main motivations to consider *stochastic evolutionary dynamics*, which yields the *stochastic stability* method for equilibrium selection in games (Foster and Young, 1990; Young, 1993; Kandori et al., 1993; Ellison, 1993). In this context, logit response has been one of the main approaches (e.g., Blume, 1993, 1995; Alós-Ferrer and Netzer, 2010; Marden and Shamma, 2012). In this literature, the model by Kreindler and Young (2013) can be seen as a finite-population analogue of our approach. They consider a logit stochastic evolution model in a finite population where the agents observe a random finite sample of other agents' play upon decision. Their focus is the speed of convergence in a two-action coordination game. Our approach is complementary, as we introduce a static equilibrium concept in large-population game and explore its properties.

The concept of deterministic *virtual payoffs* representing stochastic choice in large-population games is due to Hofbauer and Sandholm (2007). A related early discussion can be found in Anderson et al. (1992). In this literature, it is known that one can recast the expected behavior under a stochastic choice model as a deterministic behavior under an appropriately perturbed model. We adopt a similar approach and derive a virtual payoff function encapsulating the effect of dual noises, and in aggregate agents behave as if they were optimizing this virtual payoff. In two-strategy games, this link allows us to associate a *perturbed potential function* whose extrema approximate SLE. It is noted that Matějka and McKay (2015) and Fosgerau et al. (2020) established the connection between discrete choice models and the rational inattention framework (Sims, 2003) in terms of equivalent optimization problems. In light of this, if we recast our large-population result to the decision problem of a single agent, our perturbation function might be interpreted as representing a new type of information cost, or endogenous control cost in the sense of van Damme and Weibull (2002).

2 Model

2.1 Population games

We focus on 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 $S \equiv \{1, 2, \ldots, n\}$ 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*. For each $x \in X$, x_i represents 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, $F(x) = Ax$, which we may call a *linear population game*, and 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. Given a choice 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

$$\mathrm{br}(x) \equiv \operatorname*{arg\,max}_{i \in S} F_i(x) \quad \text{and} \tag{1}$$

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

where we define $\langle u, v \rangle \equiv \sum_i u_i v_i$ for same-sized 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 essentially assumes each agent has perfect information about the current population state and is able to choose an optimal action, which can be demanding depending on the context.

2.3 Sampling best response

Oyama et al. (2015) considers an alternative model that impose milder informational requirements. Assume that 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 $\mathrm{Mult}(k \mid x)$. That is, if $M^k(z \mid x) \in [0, 1]$ denotes the probability mass of drawing $z \in Z^k$ at $x \in X$, we have

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

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

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

k-sampling best response correspondence $\mathrm{BR}^k : X \rightrightarrows X$ is

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

In the element-wise manner, $y \in \mathrm{BR}^k(x)$ if and only if $y = \sum_{z\in Z^k} M^k(z\,|\,x)\cdot \alpha(z)$ where $\alpha(z) \in \mathrm{BR}(\frac{1}{k}z)$ for each $z \in Z^k$. A (k-) *sampling equilibrium* (Osborne and 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 $F(x)$:

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

A population state x is a *η-logit equilibrium* if it is consistent with the η-logit choice rule, that is, if $x = P^\eta(x)$. 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:[3]

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

It is recalled that the η-logit equilibria for game F coincide with the Nash equilibria for the distorted game $\tilde{F}$ where $\tilde{F}_i(x) \equiv F_i(x) - \eta\log(x_i)$. In fact, requiring

[2] Sampling equilibrium in this context is a special case of *sampling equilibrium with statistical inference* (Salant and Cherry, 2020) in which agents apply the maximum likelihood estimation to infer the population state from the finite sample. Sawa and Wu (2023) extensively discusses Bayesian large-population dynamics corresponding to this equilibrium concept.

[3] 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).

$x^* \in X$ and $\lambda^* = \tilde{F}_i(x^*)$ for all $i \in S$ implies $\lambda^* = \eta \log \sum_{j \in S} \exp(\eta^{-1} F_i(x_i^*))$ and $x^* = P^\eta(x^*)$. In this sense, $\tilde{F}$ is a "virtual" payoff for the η-logit equilibrium problem (Hofbauer and Sandholm, 2007, Appendix).[4]

2.5 Sampling logit choice

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

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

The existence of SLE for any fixed $1 \le k < \infty$ and $\eta > 0$ follows from Brouwer's fixed point theorem. All SLE are necessarily positive as $L^{k,\eta}$ is strictly positive. In particular, it is noted that $L^{k,\eta}(e_i) = P^\eta(e_i)$ for any $i \in S$. Also, the dynamic admits unique global solution for every initial population state in X as $L^{k,\eta}$ is continuous and globally Lipschitz on X.

3 Examples

Selected examples serve to build intuitions. Equilibrium selection under the sampling logit choice is briefly discussed.

3.1 Two-action coordination game

A formal result for special cases follow. All proofs are relegated to Appendix A.

Proposition 1. *Consider a two-action population game and $\eta > 0$. For any $\eta > 0$, there is a unique SLE that is globally asymptotically stable under* (SLD) *either (a) if $k = 1$ or (b) if $k = 2$.*

[4] See Behrens and Murata (2021) for an application in spatial economics, where $-\eta \log(x_i)$ represents *congestion externalities* incurred by the households residing in *location* $i \in S$.

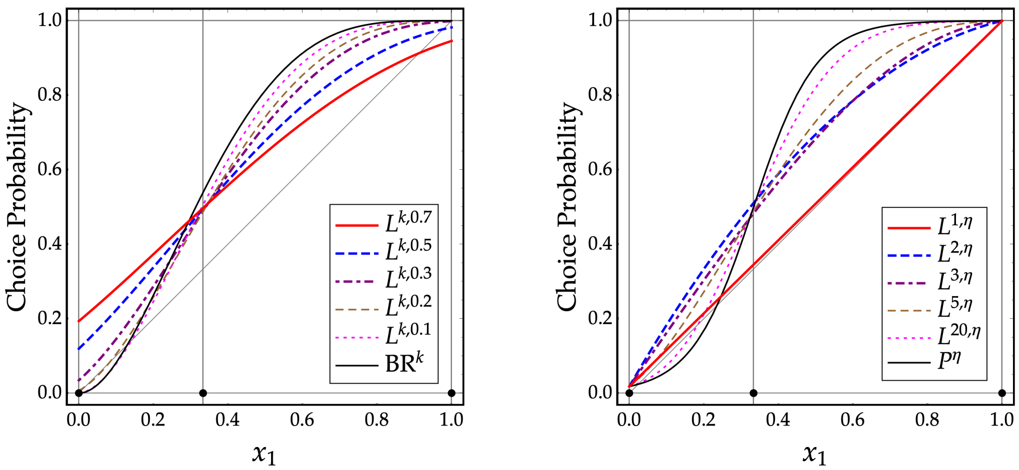

(A) BR_1^k and $L_1^{k,\eta}$ for different η ($k = 5$) (B) P_1^{η} and $L_1^{k,\eta}$ for different k ($\eta = 0.25$)

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

As a concrete example, consider a linear population game with

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

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).[5]

If $k = 2$, there is a unique SLE x_1^{**} such that $x_1^{**} \to 1$ as $\eta \to 0$ (see the proof

[5]While the proof techniques are similar to the stochastic stability approach, the selection in our context can be relatively "fast" in the sense that 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 by Oyama et al. (2015).

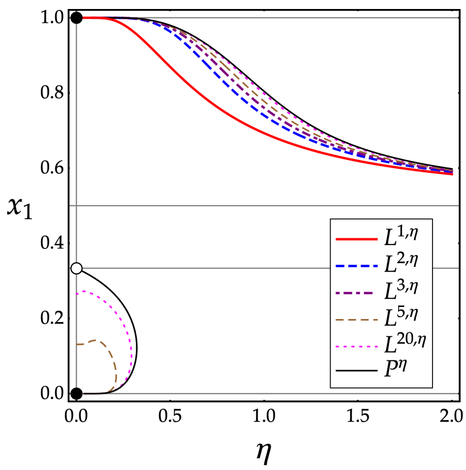

Figure 2: Sampling logit equilibria for $k \in \{1, 2, 3, 5, 20\}$ in the game (7).

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

To further illustrate the role of sampling noise in equilibrium selection, we consider the linear population game based on the 3×3 game of Young (1993):

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

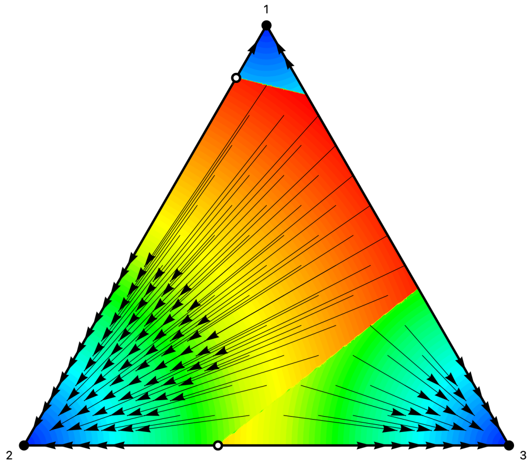

(A) Best response

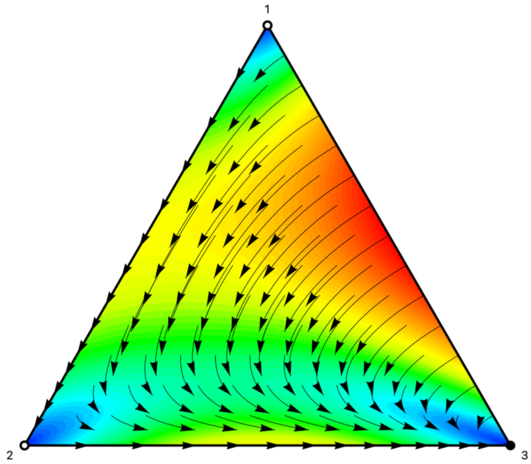

(B) Sampling best response ($k = 2$)

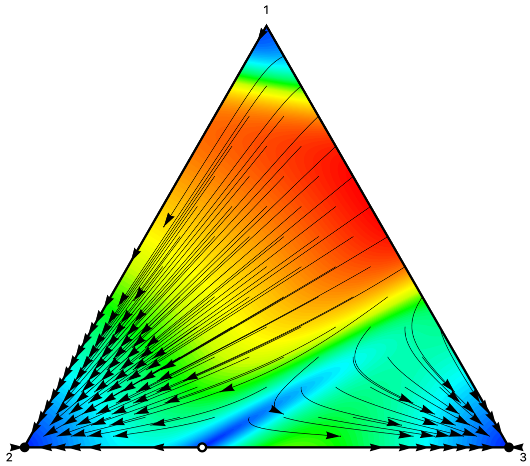

(C) Logit ($\eta = 0.3$)

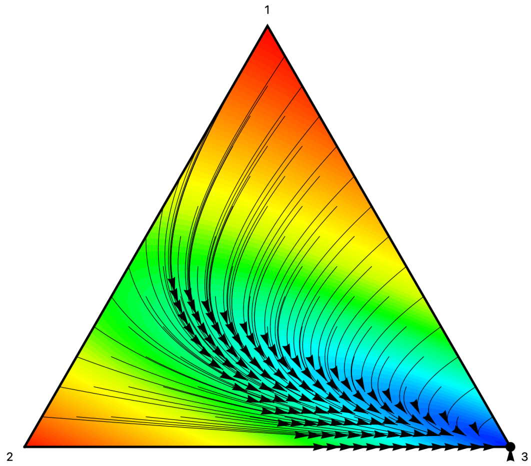

(D) Sampling logit ($k = 2, \eta = 0.3$)

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

Figure 3A illustrates that the best response dynamic (BRD) provides a basic refinement, where the two boundary equilibria are deemed unstable. However, it does not yield a unique prediction, as all three corners are locally stable.

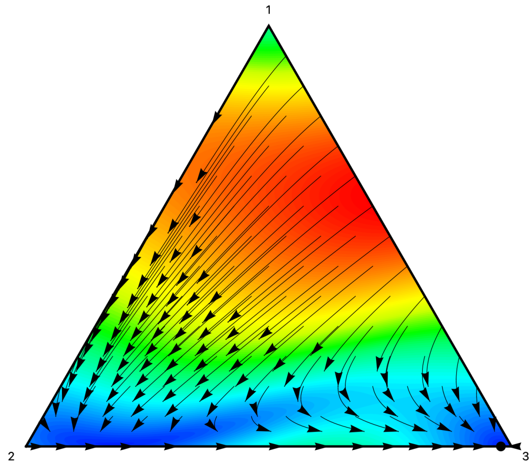

Figure 4: Unique prediction under the logit dynamic with $\eta = 0.76$ in Young's game (9). This example is due to Oyama et al. (2015, p.268)

Figure 3B, on the other hand, demonstrates that (SBRD) with $k = 2$ essentially selects $x = e_3$ because it attracts all trajectories starting from $X \setminus \{e_1, e_2\}$ (Oyama et al., 2015, Example 2).

Figure 3C considers the logit dynamic (LD). The dynamic is known to approximate (BRD) if η is sufficiently small. Reflecting this, multiple stable equilibria can be seen in 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 sampling logit dynamic (SLD) yields a sharper refinement than (SBRD): the SLE nearby e_3 is globally attracting with no exceptions at the other corners.

On equilibrium selection, logit equilibrium is known to yield a unique prediction when noise level η is sufficiently large. If $\eta = 0.76$ for example, there is a unique logit equilibrium nearby e_3, and this yields an approximate selection among Nash equilibria (Figure 4). However, because the choice rule is relatively noisy, the logit equilibrium is close to but noticeably different from e_3. With sampling noise, Figure 3D demonstrates that selection can occur even with a much smaller η, and the SLE is closer to e_3 than the logit equilibrium in Fig. 4.

Another interesting aspect of the sampling logit dynamic (SLD) is that, in Fig. 3, its trajectories (Fig. 3D) appear to head more directly toward the bottom-right vertex of the simplex, compared to those of the sampling best response dynamic (SBRD) in Fig. 3B or the logit dynamic (LD) in Fig. 3C.

4 Approximation and virtual payoffs

While the special cases discussed in Section 3 provide intuitive results on equilibrium selection, the noisy nature of logit choice does not allow for clean and rigorous characterizations based on standard game-theoretic concepts.[6] To study the expected behavior of a stochastic model, it is often useful to consider approximation approaches. For instance, Benaïm and Weibull (2003) employed stochastic approximation theory to obtain deterministic evolutionary game dynamics as the expected behavior of stochastic evolutionary process in finite populations. Below, we pursue approximation arguments to explore qualitative implications of the sampling logit choice. Below, for brevity, we suppress k and η from notations whenever they are understood from the context. For example, $L = L^{k,\eta}$, $P = P^{\eta}$, and so on.

4.1 Approximation by the delta method

It is noted that, as a random variable, 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 classical *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

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

where P_i' and P_i'' denotes the gradient and the Hessian matrix of P_i, respectively.[7] Then, the second-order delta method approximates $L_i(x) = \mathbb{E}[P_i(w)]$ by

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

[6] For example, a core result of Oyama et al. (2015) is a sufficient condition under which an "iterated p-dominant equilibrium" is almost globally asymptotically stable under a generalized version of (SBRD).

[7] 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.

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

To represent $\tilde{L}$ 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, $\sum_{i\in S} P_i(x)\, \widehat{y}_i(x) = \sum_{i\in S} P_i(x)(y_i - \overline{y}(x)) = 0$.

Direct computation by the delta method gives the following result:

Theorem 1. *Assume that F is twice differentiable. Then:*

(a) Given $\eta > 0$, for sufficiently large k, the sampling logit choice rule $L^{k,\eta}$ can be approximated as the η-logit choice rule P with multiplicative corrections:

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

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

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

with $\Sigma(x) = \operatorname{diag}[x] - xx^\top$.

(b) The approximation error $\|L^{k,\eta} - \tilde{L}\|$ vanishes uniformly as $k\eta \to \infty$ when F is linear, and as $k\eta \to \infty$ together with $k\eta^2 \to \infty$ when F is nonlinear.

4.2 Virtual payoff premiums and their origins

The representation (13) yields an interpretation of SLE based on *virtual payoffs* capturing agents' effective decision biases in aggregate, in the spirit of Hofbauer and Sandholm (2007, Appendix).

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

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

$$G_i(x) = \eta \log(1 + \widehat{v}_i(x) + \widehat{q}_i(x)) \qquad \forall i \in S, x \in X \tag{15}$$

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 $\check{F} = F + G + H$*, where* G *is the same as above and* $H_i(x) \equiv -\eta \log(x_i)$*.*

To see Proposition 2 (a), one can simply confirm

$$\tilde{L}_i(x) = \frac{\exp\left(\eta^{-1}\tilde{F}_i(x)\right)}{\sum_{j\in S}\exp\left(\eta^{-1}\tilde{F}_j(x)\right)} \qquad \forall i \in S, x \in X. \tag{16}$$

Thus, the equilibrium condition $x = \tilde{L}(x)$ is nothing but the η-logit equilibrium condition for the virtual population game $\tilde{F} = F + G$. Proposition 2 (b) then follows from a known result for the logit choice discussed in Section 2.4.

Theorem 1 shows that actions with relatively higher values of v and/or q become exaggerated in aggregate behavior. In the language of Proposition 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

$$F_i(w) \approx F_i(x) + \underbrace{\langle F_i'(x), w - x\rangle}_{\text{First-order error: } \zeta_1} + \underbrace{\tfrac{1}{2}(w-x)^\top F_i''(x)(w-x)}_{\text{Second-order error: } \zeta_2}. \tag{17}$$

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 $\text{Var}[w] = \frac{1}{k}\Sigma$, basic statistical identities imply

$$v_i(x) = \frac{1}{2\eta^2}\text{Var}\left[\langle \widehat{F_i'(x)}, w - x\rangle\right] \quad \text{and} \tag{18}$$

$$q_i(x) = \frac{1}{2\eta}\mathbb{E}\left[(w-x)^\top F_i''(x)(w-x)\right]. \tag{19}$$

From Eqs. (17) to (19), $v_i(x)$ 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* reflecting finite-sample noise. The former reflects the tendency to overweight strategies with higher variability of inferred payoffs, while the latter captures systematic distortions due to the local curvature of the payoff function.

5 How variance matters

We first focus on how variance premium $\widehat{v}$. To this end, it is useful to consider linear population games for which $F_i'' = O$ 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

$$\sigma_i(x) \equiv \widehat{a}_i(x)^\top \Sigma(x) \widehat{a}_i(x) = \sum_{l \in S} (\widehat{a}_{il}(x))^2 x_l - \Big(\sum_{l \in S} \widehat{a}_{il}(x) x_l \Big)^2 \tag{20}$$

is the variance of relative marginal payoffs. Thus, we have the following observation from Theorem 1 in the context of linear population game.

Observation 1. In linear population games, the approximated choice rule $\tilde{L}$ 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 $\pm\zeta$,

$$\underbrace{p(\mu+\zeta) - p(\mu)}_{\text{Upside increment}} = \int_{\mu}^{\mu+\zeta} p'(t)\mathrm{d}t \geq \int_{\mu-\zeta}^{\mu} p'(t)\mathrm{d}t = \underbrace{p(\mu) - p(\mu-\zeta)}_{\text{Downside decrement}}. \tag{21}$$

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

$$\mathbb{E}[p(\mu+\zeta)] \geq p\left(\mathbb{E}[\mu+\zeta]\right) = p(\mu). \tag{22}$$

Thus, the expected value is systematically pushed upward under zero-mean noise. The magnitude of this amplification can be explicitly estimated as[8]

$$\frac{\mathbb{E}[p(\mu+\zeta)] - p(\mu)}{p(\mu)} \approx \frac{1}{2}\frac{p''(x)}{p(x)}\operatorname{Var}[\zeta] = \frac{1}{2\eta^2}\operatorname{Var}[\zeta], \tag{23}$$

again confirming connection to the identity (18).

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 or 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 precisely captured by Eq. (18).

We recall the logit choice rule approximates the best response as $\eta \to 0$ and agents' choice behavior becomes more deterministic and sensitive as η decreases. In light of this, Equation (23) intuitively reveals that the variance premium, or the sensitivity of choice behavior to (first-order) finite-sample errors in payoff evaluation, becomes more pronounced as η decreases.

5.2 Two-action games

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

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{with} \quad \beta \equiv (a-c) - (b-d) \neq 0. \tag{24}$$

Proposition 3. *Consider a linear population game (24). 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*

$$\widehat{v}_i(x) = \frac{1}{2k\eta^2}\,\widehat{\sigma}_i(x) = \frac{1}{2k\eta^2} P_j(x)\,(1 - 2P_i(x))\,\sigma_A(x) \qquad \forall x \in X. \tag{25}$$

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 state is more balanced, inference based on samples fluctuate more, and the

[8] From $p(\mu+\zeta) \approx p(\mu) + p'(\mu)\zeta + \frac{1}{2}p''(\mu)\zeta^2$, for small ζ for which $\mathbb{E}[\zeta]^2 \approx 0$, 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, we can assume that $\mathbb{E}[\zeta] = 0$ because $\mathbb{E}[Aw - Ax] = A\mathbb{E}[w-x] = \mathbf{0}$.

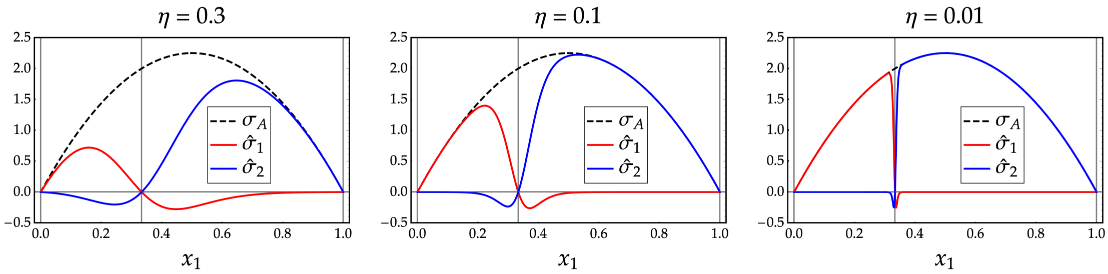

Figure 5: Illustration of Corollary 1 by the plots of $\widehat{\sigma}_i(x) = 2k\eta^2\, \widehat{v}_i(x)$ for the case $A = \left[\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right]$ with different η. The interior Nash equilibrium of the game is $x_1^* \equiv \frac{1}{3}$, and $\mathrm{br}(x) = \{2\}$ if $x < x_1^*$ and $\mathrm{br}(x) = \{1\}$ if $x_1 > x_1^*$.

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,*

$$\widehat{v}_i(x) \gtreqless 0 \quad \Leftrightarrow \quad \frac{1}{2} \gtreqless P_i(x) \quad \Leftrightarrow \quad F_j(x) \gtreqless F_i(x), \qquad \forall i, j \in S, i \neq j, \tag{26}$$

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,

$$\lim_{\eta \downarrow 0} \widehat{\sigma}_i(x) = \begin{cases} 0 & \text{if } i \in \mathrm{br}(x) \\ \sigma_A(x) & \text{if } i \notin \mathrm{br}(x) \end{cases} \qquad \forall i \in S. \tag{27}$$

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 plain η-logit choice rule because it assigns a positive choice probability to the suboptimal alternative. Equation (26) indicates that sampling errors amplify this bias.

Figure 5 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 (24), and assume that there is an interior Nash equilibrium satisfying

$$x_1 = x^* \equiv \frac{d-b}{\beta} \in (0,1) \quad \text{with} \quad \beta = (a-c)-(b-d) \neq 0. \tag{28}$$

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.[9]

Proposition 4 (Analytical approximation for the interior equilibrium)**.** *Consider a linear population game $F(x) = Ax$ where A is given by Eq. (24). Suppose $x^* > 0$ is sufficiently 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*

$$\tilde{x} = x^* - \frac{\eta}{\beta}\log\frac{1-x^*}{x^*} - \frac{\eta}{\beta}\log\left(1+v^*\right), \tag{29}$$

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. (29) captures sampling effects. If v^* is sufficiently small, the first-order approximation yields

$$-\frac{\eta}{\beta}\log\left(1+v^*\right) \approx -\frac{\eta}{\beta}\cdot v^* = -\frac{\eta}{\beta}\cdot\theta\sigma_A(x^*) = -\frac{\beta}{2k\eta}x^*(1-x^*). \tag{30}$$

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. (29) shows that $\tilde{x}$ approaches x^* as k increases and/or η decreases, which is natural:

[9] 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*

$$-\operatorname{sign}\frac{\partial}{\partial\eta}|\tilde{x}-x^*|<0 \quad \text{for small } \theta=\frac{1}{2k\eta^2}, \text{ and} \tag{31}$$

$$\operatorname{sign}\frac{\partial}{\partial k}|\tilde{x}-x^*|<0 \tag{32}$$

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

$$q_i(x)=\frac{1}{2k\eta}\langle F_i''(x),\Sigma(x)\rangle=\frac{1}{2k\eta}F_i''(x_i)\cdot\Sigma_{ii}(x)=\frac{1}{2k\eta}F_i''(x_i)\cdot x_i(1-x_i), \tag{33}$$

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 with "more convex" payoff functions are more attractive, even when expected payoffs are the same.

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.

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. (17). The variance premium arises from the leading effect of

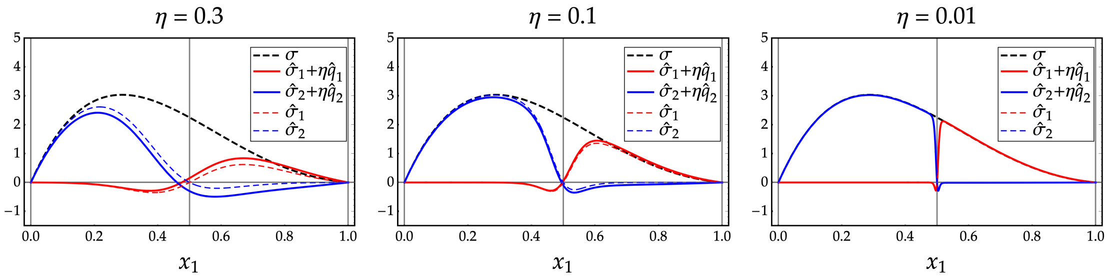

Figure 6: Illustration of Proposition 5 based on graphs of $2k\eta^2(\hat{v}+\hat{q}) = \hat{\sigma}+\eta\hat{q}$ under different values of η in a separable two-action congestion game. $\mathrm{br}(x) = \{1\}$ if $x_1 < \frac{1}{2}$ and $\mathrm{br}(x) = \{2\}$ if $x > \frac{1}{2}$. For comparison, the dashed curves show only $\hat{\sigma}$.

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,*

$$\hat{v}_1(x) = \frac{1}{2k\eta^2} P_2(x)\,(1-2P_1(x))\,\sigma(x) \quad \textit{and} \tag{34}$$

$$\hat{q}_1(x) = \frac{1}{2k\eta} P_2(x)(F_1''(x_1) - F_2''(x_2))x_1x_2, \tag{35}$$

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

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 $\hat{v}$ and $\hat{q}$ can be different, reflecting their distinct origins.

Figure 6 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. 5). This is because marginal payoffs are state dependent. While $\hat{v}$ vanishes at x_1^* as $P_1(x^*) = \frac{1}{2}$, nonzero distortion remains at x_1^* due to $\hat{q}$. 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 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

$$\frac{\partial f(x)}{\partial x_1} - \frac{\partial f(x)}{\partial x_2} = F_1(x) - F_2(x) \qquad \forall x \in X. \tag{36}$$

In fact, with $y(t) \equiv (t, 1-t)$, the following function satisfies Eq. (36):

$$f(x) \equiv \int_0^{x_1} \left(F_1(y(t)) - F_2(y(t))\right) \mathrm{d}t. \tag{37}$$

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):

$$f^{\eta}(x) \equiv f(x) + h(x) \quad \text{where} \quad h(x) \equiv -\eta \sum_{i \in S} x_i \log x_i. \tag{38}$$

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

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

Proposition 6. *Consider a two-action population game F. Define the perturbed potential function $f^{k,\eta} : X \to \mathbb{R}$ by*

$$f^{k,\eta}(x) \equiv \int_0^{x_1} \left(\tilde{F}_1(y(t)) - \tilde{F}_2(y(t))\right) \mathrm{d}t + h(x) \tag{39}$$

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

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

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 $L^{k,\eta} \approx \tilde{L}$ is sufficiently good. Nonetheless, under this hypothesis, Proposition 6 yields a simple dynamic characterization of SLE by a direct application of Theorem 3.2 in Hofbauer and Sandholm (2007):

Proposition 7. *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 $\dot{x} = \tilde{L}(x) - x$. 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.*

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

$$g(x) \equiv \int_0^{x_1} (G_1(y(t)) - G_2(y(t)))\mathrm{d}t. \tag{40}$$

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 (24). 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 we23 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 8. *Consider the general linear two-action game (24) 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^*.*

Figure 7 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 $x_1 x_2 = 0$.

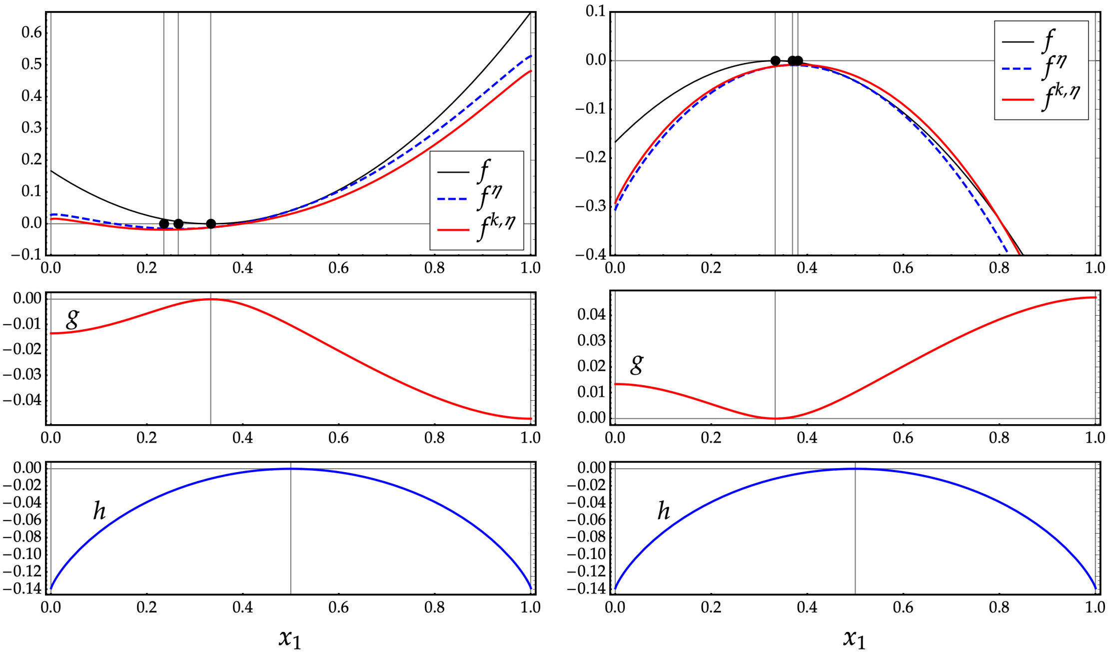

(A) Coordination game ($\beta > 0$) (B) Anti-coordination game ($\beta < 0$)

Figure 7: 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.

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 from the finite-sample evaluation of payoffs.

8 Discussions

This study contributes a unifying equilibrium concept that connects and extends both the QRE/logit models of noisy choice and the sampling-based models of limited information. Combined with the virtual payoff representation, this leads to insights into the interaction of dual noises. We identify which equilibrium outcomes are biased and in what direction due to finite sampling, among which we highlight the *variance premium*, bias for actions with higher variability of perceived payoffs. The mechanism here is not a primitive risk preference, but rather a structural bias induced by the convexity properties of the logit choice mapping under sampling noise. The effect parallels convexity biases well discussed in statistics and can be viewed as a form of virtual risk loving emerging from noisy observation.

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.

On the empirical side, the comparative statics of the sampling logit equilibrium might be explored through laboratory experiments. The two parameters play distinct roles: the sample size k reflects how much information agents can observe, while the noise level η reflects how precisely they act on it. In principle, k could be varied by adjusting the amount of feedback or observations, and η by altering time pressure or cognitive load. We would then expect larger k to weaken the bias toward high-variance strategies, while smaller η mainly sharpens responsiveness to payoff differences. Such experiments could provide an empirical benchmark for distinguishing informational from decisional sources of bounded rationality.

Finally, several apparent limitations of the present study should be noted. First, our results concerning systematic biases rely on an approximation that assumes a sufficiently large sample size k. Figure 8 illustrates how the accuracy of the approximation deteriorates when k is small, especially for relatively small η. The precise dynamics for small k remains an open question beyond the specific cases examined. 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

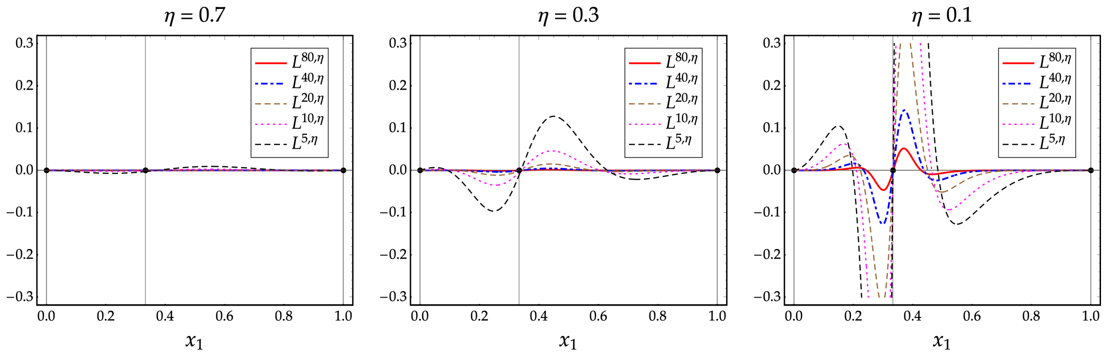

Figure 8: Approximation errors for the choice probability of action 1 ($L_1^{k,\eta} - \tilde{L}_1$) under different η and k. The coordination game (7) with $(s, t) = (2, 1)$ as considered in Figs. 1 and 2 is assumed. Black markers show the Nash equilibria.

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 warrants further investigation. Stable games (Hofbauer and Sandholm, 2009) would be a natural starting point for such analysis.

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. 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.

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

Fosgerau, M., Melo, E., de Palma, A., and Shum, M. (2020). Discrete choice and rational inattention: A general equivalence result. *International Economic Review*, 61(4):1569–1589.

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.

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.

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.

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.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York.

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.

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.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

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 Damme, E. and Weibull, J. W. (2002). Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory*, 106(2):296–315.

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

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

A Proofs

Proof of Proposition 1. *(a) The case $k = 1$.* We show uniqueness and global asymptotic stability for general $n \geq 2$. Note that $L^{1,\eta}(x) = \sum_{z \in \{e_j\}_{j \in S}} M_x^1(z) \cdot P^\eta(z) = \sum_{j \in S} \Pi_{ij} x_j$ where $\Pi_{ij} \equiv P_i^\eta(e_j) \in (0,1)$ 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 $\Pi \in \mathbb{R}_{>0}^{n \times n}$ 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.

(b) The case $k = 1$ 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, $q_0 = P_1^\eta(e_2)$, $q_1 = P_1^\eta(\frac{1}{2}(e_1 + e_2))$, and $q_2 = P_1^\eta(e_1)$. Define

$$f(y) \equiv L_1^{k,\eta}(y, 1-y) - y \tag{41}$$

$$= q_2 \times y^2 + 2q_1 \times y(1-y) + q_0 \times (1-y)^2 - y \tag{42}$$

$$= (q_2 - 2q_1 + q_0)y^2 + (2(q_1 - q_0) - 1)\, y + q_0. \tag{43}$$

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 $n \geq 2$.

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

$$y^* = \frac{-b - \sqrt{b^2 - 4ac}}{2a} \in (0,1) \tag{44}$$

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

$$P_i'(x) = \theta P_i(x)\big(F_i'(x) - F'(x)^\top P(x)\big) = \theta P_i(x)\widehat{F_i'(x)}. \tag{45}$$

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

$$P_i''(x) = \theta R_i P_i'^\top + \theta P_i R_i' = \theta^2 P_i R_i R_i^\top + \theta P_i R_i'. \tag{46}$$

The Jacobian matrix of R_i is computed as

$$\begin{aligned}
R_i' = \big(F_i' - (\sum_l P_l F_l')\big)' &= -\sum_l \big(F_l'(P_l')^\top + P_l F_l''\big) + F_i'' && (47)\\
&= -\theta \sum_l P_l F_l' R_l^\top + F_i'' - \sum_l P_l F_l'' && (48)\\
&= -\theta \sum_l P_l (F_l' - \overline{F'}) R_l^\top - \theta \sum_l P_l \overline{F'} R_l^\top + \widehat{F_i''} && (49)\\
&= -\theta \sum_l P_l R_l R_l^\top + \widehat{F_i''}, && (50)
\end{aligned}$$

where we note that $\sum_l P_l \overline{F'} R_l^\top = O$. Together with Eq. (46), we see

$$P_i''(x) = \theta^2 P_i\Big(R_i R_i^\top - \sum_l P_l R_l R_l^\top\Big) + \theta P_i \widehat{F_i''} = \theta^2 P_i \cdot \widehat{R_i R_i^\top} + \theta P_i \widehat{F_i''}. \tag{51}$$

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

$$\frac{1}{2k}\langle P_i''(x), \Sigma(x)\rangle = \frac{1}{2k\eta^2} P_i \cdot \widehat{R_i^\top \Sigma R_i} + \frac{1}{2k\eta} P_i \cdot \widehat{\langle F_i'', \Sigma\rangle}. \tag{52}$$

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.

Proposition A. *For any $\epsilon > 0$, $\|L^{k,\eta}(x) - \tilde{L}(x)\| \to 0$ uniformly on the set $X_\epsilon \equiv \{x \in X : \min_{i \in S} x_i \geq \epsilon\}$ as $k\eta \to \infty$ and $k\eta^2 \to \infty$, that is, if $k \gg \max\{\eta^{-1}, \eta^{-2}\}$. If F is linear, then $k \gg \eta^{-1}$ suffices.*

Proof. Let the expected residual of the proposed approximation be

$$R(x) \equiv \mathbb{E}[P(w)] - \tilde{L}(x) = L^{k,\eta}(x) - \Big(P(x) + \frac{1}{2k}\langle P''(x), \Sigma(x)\rangle\Big), \tag{53}$$

where $\langle P''(x), \Sigma(x)\rangle$ is interpreted in the element-wise manner, i.e., its ith element is $\langle P_i''(x), \Sigma(x)\rangle$. The leading term of R has an order proportional to the third-order derivative of the logit choice map P''' and third-order moment of the random variable w.

Since P' diverges at the boundary of X, we must focus on X_ϵ. Derivatives of the logit choice map scale as $\|P'\| = O(\eta^{-1})$, $\|P''\| = O(\eta^{-2})$, and $\|P'''\| = O(\eta^{-3})$ as $\eta \downarrow 0$ under standard norms.

Also, w is the sum of i.i.d. random variables: $w = \frac{1}{k}(Y_1 + Y_2 + \cdots + Y_k)$ where $Y_l \in \{e_i\}_{i\in S}$ and $\Pr(Y_l = e_i) = x_i$. Let $\kappa(Z) \in \mathbb{R}^{n\times n\times n}$ denote the third-order cumulant of a $\mathbb{R}^n$-valued random variable Z. From standard properties of cumulants, $\kappa(aZ) = a^3\kappa(Z)$, and $\kappa(\sum_l Z_l) = \sum_l \kappa(Z_l)$ for independent $\{Z_l\}$. Thus, $\kappa(w) = \frac{1}{k^3}\kappa(\sum_{l=1}^k Y_l) = \frac{1}{k^3}\sum_{l=1}^k \kappa(Y_l) = \frac{1}{k^3}\cdot k \cdot \kappa(Y_1) = \frac{1}{k^2}\kappa(Y_1)$. Since $\kappa(Y_1)$ is constant with respect to k and η, $\|\kappa(w)\| = O(k^{-2})$ under any fixed operator norm on tensors. To sum up, there exist appropriately chosen constants C^* and C^{**} such that

$$\|R(x)\| \leq C^* \cdot \|P'''(x)\| \cdot \|\kappa(W)\| \leq \frac{C^{**}}{k^2\eta^3} \tag{54}$$

uniformly for all $x \in X_\epsilon$. The constants C^* and C^{**} depend on ϵ, F, and the chosen norm for P''' and κ, but independent of k and η. That is, $\|R(x)\| = O(k^{-2}\eta^{-3})$. It is noted that the correction term $\delta \equiv \widehat{v} + \widehat{q}$ in Theorem 1 is $O(k^{-1}\eta^{-2})$ for $\eta < 1$ and $O(k^{-1}\eta^{-1})$ for $\eta \geq 1$. From Equation (54), this implies

$$\text{(bound from } \hat{v}\text{)} \quad \sup_{x\in X_\epsilon} \frac{\|R(x)\|}{k^{-1}\eta^{-2}} \leq \frac{C^{**}}{k\eta} \to 0 \quad \text{as} \quad k\eta \to \infty \quad \text{and} \tag{55}$$

$$\text{(bound from } \hat{q}\text{)} \quad \sup_{x\in X_\epsilon} \frac{\|R(x)\|}{k^{-1}\eta^{-1}} \leq \frac{C^{**}}{k\eta^2} \to 0 \quad \text{as} \quad k\eta^2 \to \infty. \tag{56}$$

Thus, $k \gg \eta^{-1}$ and $k \gg \eta^{-2}$. If F is linear, $\hat{q}$ is absent and $\delta = O(k^{-1}\eta^{-2})$, for which case $k \gg \eta^{-1}$ suffices. For nonlinear cases, the former condition is active if $\eta \geq 1$ and the latter is active for $\eta < 1$. □

Theorem 1 (b) is a slightly informal version of Proposition A. ■

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

$$\sum_{j\in S}(1 + \delta_j(x)) \exp\left(\eta^{-1}F_j(x)\right) = Z(x)\sum_{j\in S}(1 + \delta_j(x))\frac{\exp\left(\eta^{-1}F_j(x)\right)}{Z(x)} \tag{57}$$

$$= Z(x) \sum_{j \in S} (1 + \delta_j(x)) P_j(x) = Z(x) \tag{58}$$

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

$$\frac{\exp\left(\eta^{-1} \tilde{F}_i(x)\right)}{\sum_{j \in S} \exp\left(\eta^{-1} \tilde{F}_j(x)\right)} = \frac{(1 + \delta_i(x)) \exp\left(\eta^{-1} F_i(x)\right)}{\sum_{j \in S} (1 + \delta_j(x)) \exp\left(\eta^{-1} F_j(x)\right)} \tag{59}$$

$$= \frac{(1 + \delta_i(x)) \exp\left(\eta^{-1} F_i(x)\right)}{Z(x)} \tag{60}$$

$$= (1 + \delta_i(x)) P_i(x) = \tilde{L}_i(x). \tag{61}$$

Thus, the condition $x = \tilde{L}(x)$ is nothing but the η-logit equilibrium condition for the virtual payoff $\tilde{F}$, showing part (a). A direct application of known results for the logit choice yields (b) (e.g., Hofbauer and Sandholm, 2007, Appendix). ■

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 $R_i \equiv \widehat{F_i'(x)}$,

$$R_1 = A_1 - (p_1 A_1 + p_2 A_2) = p_2 (A_1 - A_2) = p_2 A_0, \tag{62}$$

$$R_2 = A_2 - (p_1 A_1 + p_2 A_2) = -p_1 (A_1 - A_2) = -p_1 A_0. \tag{63}$$

Let $\sigma_A(x) \equiv A_0^\top \Sigma A_0 = ((a - c) - (b - d))^2 \cdot x_1 x_2 \geq 0$, with equality iff $a - c = b - d$ or $x_1 x_2 = 0$. Then, we confirm

$$\sigma_1 = R_1^\top \Sigma R_1 = p_2^2 A_0^\top \Sigma A_0 = p_2^2 \sigma_A, \tag{64}$$

$$\sigma_2 = R_2^\top \Sigma R_2 = p_1^2 A_0^\top \Sigma A_0 = p_1^2 \sigma_A, \tag{65}$$

$$p_1 \sigma_1 + p_2 \sigma_2 = p_1 p_2^2 \sigma_A + p_2 p_1^2 \sigma_A = (p_1 + p_2) p_1 p_2 \sigma_A = p_1 p_2 \sigma_A. \tag{66}$$

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

$$(\delta_1(x_1), \delta_2(x_1)) = \begin{cases} (0, \delta(x_1)) & \text{if } x_1 > x^* \\ (0, 0) & \text{if } x_1 = x^* \\ (\delta(x_1), 0) & \text{if } x_1 < x^*, \end{cases} \tag{67}$$

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 $\tilde{L}(x_1)$. Then,

$$\frac{y}{1-y} = \frac{\tilde{L}_1(y)}{\tilde{L}_2(y)} = \exp\left(\frac{\beta(y-x^*)}{\eta}\right) \cdot \frac{1+\delta_1(y)}{1+\delta_2(y)}. \tag{68}$$

Taking log of both sides,

$$\log\frac{y}{1-y} = \frac{\beta}{\eta}(y-x^*) + \rho(y), \quad \rho(y) = \begin{cases} -\log(1+\delta(y)) & \text{if } y > x^* \\ 0 & \text{if } y = x^* \\ \log(1+\delta(y)) & \text{if } y < x^*. \end{cases} \tag{69}$$

Set $\epsilon \equiv y - x^*$ and assume $\epsilon = O(\eta)$. From the first-order expansion of Eq. (69),

$$\log\frac{x^*}{1-x^*} + \frac{\epsilon}{x^*(1-x^*)} = \frac{\beta}{\eta}\epsilon + \rho(y). \tag{70}$$

It is noted that the special case $y = x^* \in (0,1)$ occurs if and only if $x^* = \frac{1}{2}$ because Eq. (68) implies $\frac{y}{1-y} = 1$. Then, Equation (70) yields $\epsilon = 0$.

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

$$\begin{aligned}
\epsilon &= \left(\log\frac{x^*}{1-x^*} - \rho(y)\right)\left(\frac{\beta}{\eta} - \frac{1}{x^*(1-x^*)}\right)^{-1} && (71)\\
&= \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} - \rho(y)\right)\left(1 - \frac{\eta}{\beta}\frac{1}{x^*(1-x^*)}\right)^{-1} && (72)\\
&= \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} - \rho(y)\right)(1+O(\eta)) \quad (\because (1-c)^{-1} = 1 + c + O(c^2)) && (73)\\
&= \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} - \rho(y)\right). && (74)
\end{aligned}$$

To examine the dependence on y, define

$$\epsilon_\pm = \frac{\eta}{\beta}\left(\log\frac{x^*}{1-x^*} \pm \log(1+\delta^*)\right) \tag{75}$$

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. (70) because it amounts to considering a second-order term of η and thus irrelevant for the first-order approximation. From

Eq. (69) 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

$$\epsilon_+ > 0 \quad \Leftrightarrow \quad \log \frac{x^*}{1-x^*} + \log(1+\delta^*) > 0 \quad \Leftrightarrow \quad x^* > x_+ \equiv \frac{1}{2+\delta^*}, \tag{76}$$

$$\epsilon_- < 0 \quad \Leftrightarrow \quad \log \frac{x^*}{1-x^*} - \log(1+\delta^*) < 0 \quad \Leftrightarrow \quad x^* < x_- \equiv \frac{1+\delta^*}{2+\delta^*}. \tag{77}$$

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 (25) 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 $x_1(1-x_1) = x_2(1-x_2)$. ■

Proof of Proposition 6. The slope of $f^{k,\eta}$ along the tangent space of X satisfy

$$\mathrm{d}f \equiv \frac{\partial f^{k,\eta}(x)}{\partial x_1} - \frac{\partial f^{k,\eta}(x)}{\partial x_2} = \tilde{F}_1(x) - \tilde{F}_2(x) - \eta \log x_1 - \eta + \eta \log x_2 + \eta. \tag{78}$$

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 $\eta \log \frac{x_1}{x_2} = \tilde{F}_1(x) - \tilde{F}_2(x)$, which reduces to the condition $x_i = \tilde{L}_i(x)$, showing that x is must be a fixed point of $\tilde{L}$. ■

Proof of Proposition 7. From Proposition 2, the approximated dynamic can be seen as the η-logit dynamic in the modified population game $\tilde{F}$. 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. ■

Proof of Proposition 8. We note that g is continuous and differentiable.

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

$$\frac{\partial g(x)}{\partial x_1} - \frac{\partial g(x)}{\partial x_2} = G_1(x) - G_2(x) \begin{cases} > 0 & \text{if } x_1 < x^* \; (\mathrm{br}(x) = \{2\}) \\ = 0 & \text{if } x_1 = x^* \; (\mathrm{br}(x) = \{1, 2\}) \\ < 0 & \text{if } x_1 > x^* \; (\mathrm{br}(x) = \{1\}). \end{cases} \tag{79}$$

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 $\beta > 0$ ($\beta < 0$). ■