Statistical inference in games: stability of pure equilibria

Segismundo S. Izquierdo[a,*] and Luis R. Izquierdo[b]

[a] *Department of Industrial Organization and BioEcoUva, Universidad de Valladolid, Spain.*
[b] *Department of Management Engineering, Universidad de Burgos, Spain.*

Abstract

We consider sampling best response decision protocols with statistical inference in population games. Under these protocols, a revising agent observes the actions of k randomly sampled players in a population, estimates from the sample a probability distribution for the state of the population (using some inference method), and chooses a best response to the estimated distribution. We formulate deterministic approximation dynamics for these protocols. If the inference method is unbiased, strict Nash equilibria are rest points, but they may not be stable. We present tests for stability of pure equilibria under these dynamics. Focusing on maximum-likelihood estimation, we can define an index that measures the strength of each strict Nash equilibrium. In *tacit coordination or weakest-link* games, the stability of equilibria under sampling best response dynamics is consistent with experimental evidence, capturing the effect of *strategic uncertainty* and its sensitivity to the number of players and to the cost/benefit ratio.
JEL classification numbers: C72, C73.

Keywords: Statistical inference; Sampling best response; Stability; Strict Nash; weakest-link games.

1 Introduction

In this paper we study the local stability of pure equilibria under *sampling best response dynamics with statistical inference* (Salant & Cherry, 2020; Sawa & Wu, 2023), which incorporate statistical inference into the *sampling best response* decision protocol introduced by Sandholm (2001). We consider a population of agents who repeatedly make decisions that affect each other (i.e. agents choose an action or pure strategy in a stage game). Specifically, the payoff obtained by an agent who plays strategy a_i depends on

*Correspondence to: Department of Industrial Organization, Universidad de Valladolid, Dr. Mergelina s/n, 47011 Valladolid, Spain. *e-mail*: segismundo.izquierdo@uva.es.
Abbreviations. SBR: Sampling Best Response; BR: Best Response; BEP: Best Experienced Payoff.

the prevalence of each strategy in the population. Under *sampling best response dynamics with statistical inference*, agents follow this procedure to revise their current strategy:

1. Obtain a sample of (the strategies used by) k players, randomly sampled from the population of agents.

2. Use that sample, together with an inference method, to estimate a probability distribution over population states (i.e., an *estimate*).

3. Calculate the expected payoff to each strategy, according to the estimated probability distribution over population states.

4. Adopt one of the strategies with maximum expected payoff.

If the inference method estimates that the prevalence of each strategy in the population is the same as its prevalence in the observed sample, we obtain, as a special case, the Sampling Best Response (SBR) protocol (Osborne & Rubinstein, 2003; Oyama, Sandholm, & Tercieux, 2015; Sandholm, 2001). This case corresponds to the maximum likelihood estimate of the population state from the sample, which is a point estimate or degenerate distribution.

The rest points or equilibria of the SBR process have been named *sampling equilibria* (Osborne & Rubinstein, 2003) and also *action-sampling equilibria* (Arigapudi, Heller, & Milchtaich, 2021; Sethi, 2021).[1] Salant and Cherry (2020) (see also Sawa and Wu (2023)) introduced a generalization of these action-sampling equilibria by considering agent heterogeneity and different estimation methods of the population state from the sample (other than taking the sampled empirical distribution as the estimated population state). For this generalization, they used the term *sampling equilibria with statistical inference (SESI)*.

Regarding the stability of these equilibria, Sandholm (2001) and Oyama et al. (2015) provided conditions for (almost) global asymptotic stability under SBR dynamics.[2] Sawa and Wu (2023) extended some of these results to sampling best response dynamics with statistical inference and heterogeneous preferences, for unbiased inference methods in games with two strategies. These conditions for global stability are all related to γ-dominance. In a single-population game, a strategy is γ-dominant (with $\gamma \in [0, 1]$) if it is the unique best response whenever the fraction of players using it is greater than or equal to γ (Morris, Rob, & Shin, 1995).

Sandholm (2001) showed that, under a SBR protocol with sample size $k \geq 2$, if there exists a $\frac{1}{k}$-dominant strategy,[3] then it corresponds to an almost globally asymptotically

[1] The term *action-sampling equilibrium* has been used to distinguish the equilibria considered here from those corresponding to other decision protocols based on sampling, such as the *best-experienced-payoff* protocol, which selects the strategy that performs best in a (sampled) battery of trials. This protocol gives rise to the so-called *payoff-sampling* equilibria or, more generally, *best-experienced-payoff* equilibria.

[2] Intuitively, if a state is almost globally asymptotically stable, it means that, in large populations, the process will converge to this state with high probability from almost every initial condition.

[3] Note that being $\frac{1}{2}$-dominant is a necessary condition to be $\frac{1}{k}$-dominant for any $k \geq 2$.

stable equilibrium. Being $\frac{1}{k}$-dominant (for any $k \geq 2$) is a refinement of strict Nash equilibrium that may be satisfied by at most one equilibrium in a game (with the condition becoming more stringent as k grows).[4]

Oyama et al. (2015) formulated deterministic dynamics (more precisely, a differential inclusion) that approximate well the population state dynamics under SBR protocols in large populations. Oyama et al. (2015) showed that, in order to find the set of surviving strategies under a SBR dynamic starting from interior initial conditions, some strategies can be eliminated iteratively. If the elimination process leads to one single surviving strategy, it corresponds to an almost globally asymptotically stable strict Nash state. The conditions for a strict Nash state to be almost globally asymptotically stable in Oyama et al. (2015) are more general than in Sandholm (2001), with the elimination process requiring some proper subset of strategies to be $\frac{1}{2}$-dominant (i.e., some proper subset of strategies J such that the best responses to any state x satisfying $\sum_{i:a_i \in J} x_i \geq \frac{1}{2}$ belong to J).

Instead of studying conditions for global stability (of one strict Nash equilibrium) under SBR dynamics, in this paper we analyze the local asymptotic stability of every monomorphic rest point under sampling best response dynamics with statistical inference (and, as a special case, of every strict Nash equilibrium under SBR dynamics). We achieve this by characterizing the Jacobian of the dynamics at such rest points, leading to one test for instability and another test for asymptotic stability.

Focusing on the specific case of SBR dynamics, our local stability analysis can be used to provide an index for each strict Nash equilibrium. By discriminating between different strict Nash equilibria according to their stability under SBR, our approach complements previous approaches based on (almost) global asymptotic stability, which can single out at most one equilibrium from the rest. The following example illustrates some of our contributions.

Example 1.1. Consider the following two-player symmetric game, with strategy set $\{a, b, c\}$ and payoff matrix to the row player:

$U_1 =$

	a	b	c
a	3	8	0
b	0	9	0
c	0	0	4

Figure 1 shows the best response dynamics for this game played in one single population. There are three strict Nash equilibria, each corresponding to one of the three strategies. Most equilibrium selection methods do not distinguish between the three strict Nash equilibria of this game. Also, no strategy or proper subset of strategies is $\frac{1}{2}$-dominant (see figure 1), so the results of Oyama et al. (2015) for SBR dynamics do

[4] In the context of population games, the term "Nash equilibrium" is often used to refer to a Nash state (see section 2.1).

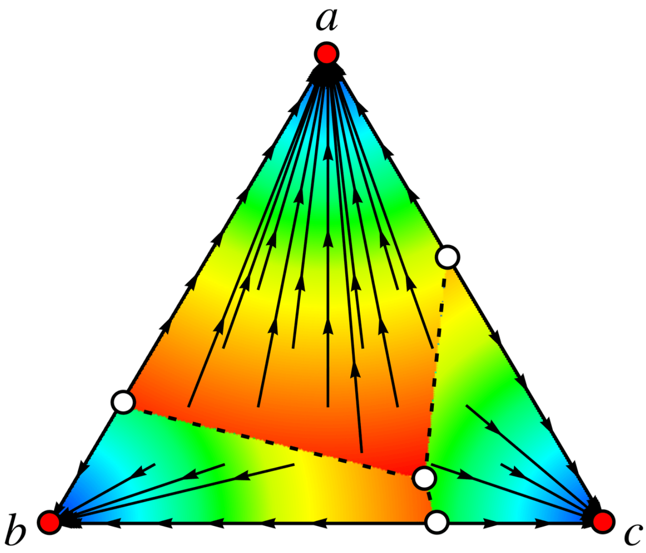

Figure 1: Best response dynamics in the game with payoff matrix U_1. In the figures, background colors represent speed of motion: red is fastest, blue is slowest. Isolated rest points are represented with circles: red if the rest point is asymptotically stable, white if unstable. The dashed black lines in this figure separate the basins of attractions of each asymptotically stable rest point.

not apply here. However, using the methods developed in this paper, the following properties of SBR dynamics can be easily derived (see figure 2):

- No equilibrium is asymptotically stable for $k = 1$.
- The first equilibrium (corresponding to strategy a) is unstable for $k = 2$ and asymptotically stable for $k \geq 3$.
- The second equilibrium (corresponding to strategy b) is unstable for $k \in \{2, 3\}$ and asymptotically stable for $k \geq 5$. Its stability for $k = 4$ depends on how ties between alternative best response strategies are resolved. Figures 2(iii) and 2(iv) show SBR dynamics for $k = 4$ under two different tie breakers.[5]
- The third equilibrium (corresponding to strategy c) is unstable for $k \in \{2, 3\}$ and asymptotically stable for $k \geq 4$.

These results suggest an index for each strict equilibrium, according to the lowest sample size that guarantees asymptotic stability under SBR dynamics, with the index of the first equilibrium being 3, the index of the second equilibrium being 5, and the index of the third equilibrium being 4. The index indicates the sample size needed to

[5] Tie breakers are defined in section 2.2.

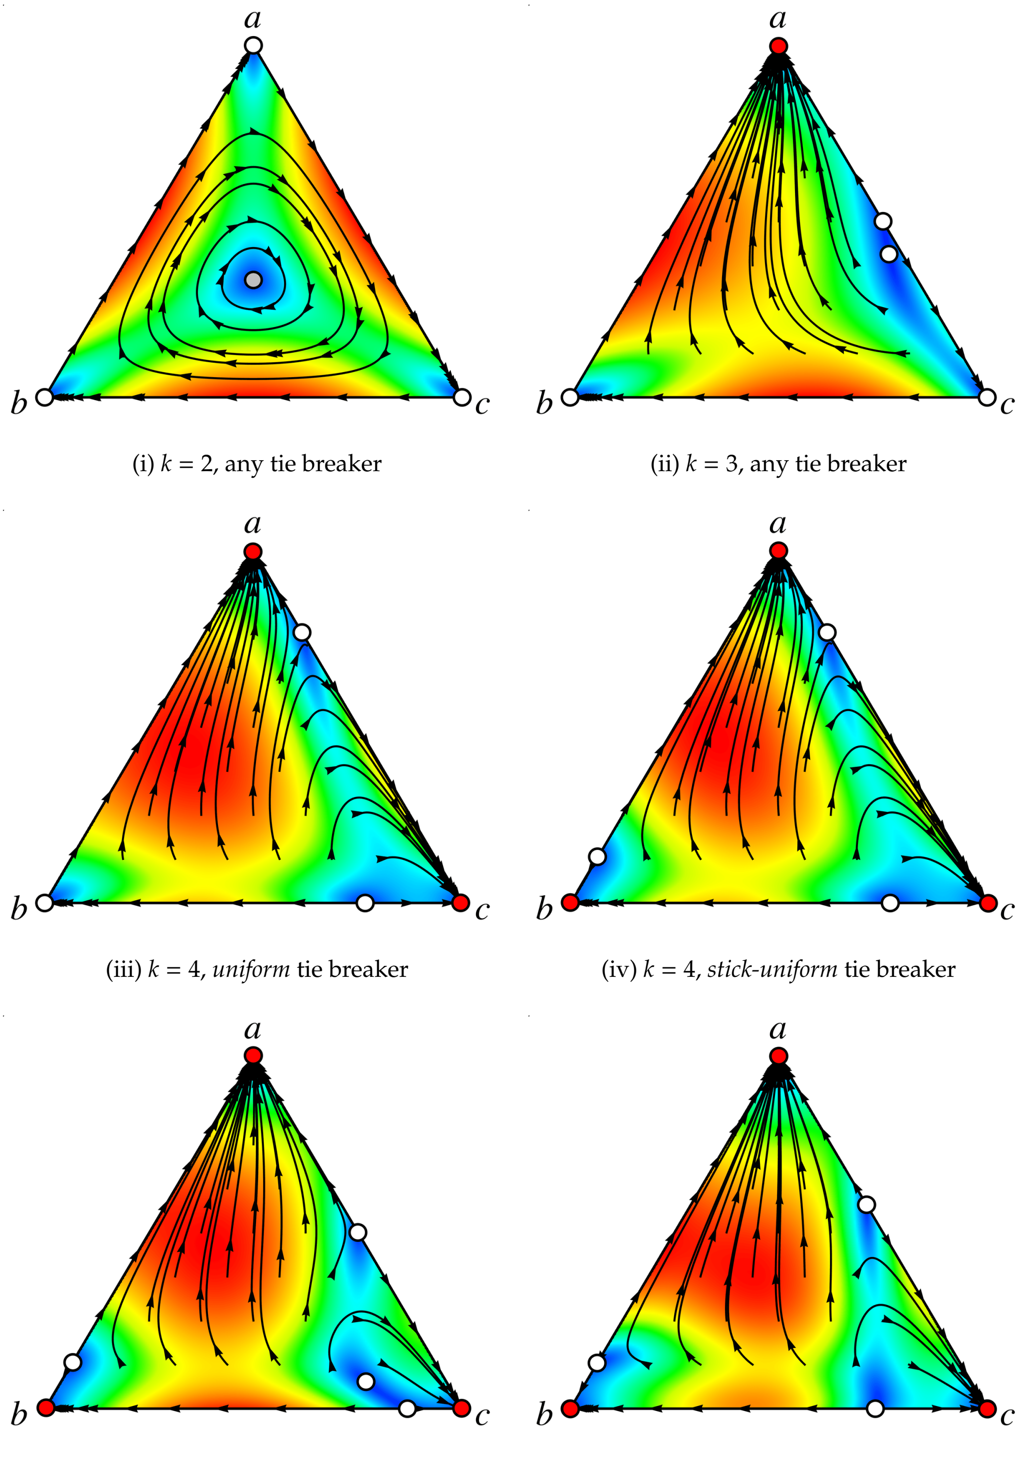

(i) $k = 2$, any tie breaker

(ii) $k = 3$, any tie breaker

(iii) $k = 4$, *uniform* tie breaker

(iv) $k = 4$, *stick-uniform* tie breaker

(v) $k = 5$, *uniform* tie breaker

(vi) $k = 10$, *uniform* tie breaker

Figure 2: Sampling best response dynamics for different values of k and different tie breakers, for the game with payoff matrix U_1.

make a strict equilibrium asymptotically stable under SBR dynamics. The lower the index number, the more robust the equilibrium is.[6]

The type of stability-based discrimination between strict Nash equilibria proposed here cannot be achieved using most evolutionary dynamics that are guided by *exact* (vs. estimated) expected payoffs: strict Nash equilibria are asymptotically stable under most dynamics based on exact expected payoffs because, assuming continuity of the payoff functions, there is a neighborhood of every strict Nash state in which each of the strategies making up the strict equilibrium profile is the corresponding player's unique best response to the population state, so each of those strategies is "locally strictly dominant" in expected payoffs for the player that is using it.

A related protocol that is also based on sampling but which does not involve estimates of the population state, and under which strict Nash equilibria may also be unstable, is the *payoff-sampling* protocol proposed by Osborne and Rubinstein (1998) and Sethi (2000) (more generally, the *best experienced payoff* protocol of Sandholm, Izquierdo, and Izquierdo (2020)). Under this protocol, revising players test (some of) their available strategies against random samples of co-players and choose the strategy that performed best in the test. Many of the proofs in this paper are based on the methods developed by Sandholm et al. (2020) and Izquierdo and Izquierdo (2022) to analyze the stability of strict Nash equilibria under best experienced payoff dynamics (see also Arigapudi et al. (2021)). Differences and similarities between these protocols are discussed with more detail in appendix A.2.

The remaining of the paper is structured as follows. Section 2 introduces *population games* and derives the *mean dynamic* differential equations for sampling best response protocols with statistical inference. In section 3 we derive the Jacobian of the dynamics at monomorphic equilibria and present two tests based on the Jacobian: one for asymptotic stability and one for instability. In this section we also extend our tests of stability to λ-sampling dynamics (section 3.6). Section 4 proposes an index for strict Nash equilibria based on asymptotic stability under SBR dynamics, and studies the implications of the index of an equilibrium. In section 5 we apply our results to a practical case, namely *tacit coordination* (or weakest-link) games, and show that our analysis captures some interesting experimental features of these games. Lastly, in section 6 we present some conclusions. Most of the proofs are relegated to an appendix. Some extensions and special cases are also discussed in appendices.

All figures in this paper can be easily replicated with open-source freely available software which performs exact computations of rest points and exact linearization analyses.[7]

---

[6]In this particular example, there is a perfect correlation between the index of an equilibrium and its efficiency, but this is not necessarily the case in general. To see this, note that we can make the inefficient equilibrium corresponding to strategy a be the most efficient one by adding the same value to each payoff in the first column of U_1, and the SBR dynamics would not change.

[7]*EvoDyn-3s* (Izquierdo, Izquierdo, & Sandholm, 2018) for figures 1-8, and *SBR-TCG* (https://doi.org/10.5281/zenodo.7933941) for figure 10.

2 Sampling best response with statistical inference

2.1 Population games

For notational simplicity, we focus here on symmetric games played in one population, although our results can be extended to multi-population (symmetric or asymmetric) games. The extension to two-population games is discussed in appendix A.5.[8]

Following Sandholm (2010), we consider a unit-mass population of agents who play a game with strategy set $A = \{a_1, ..., a_n\}$, containing n pure strategies or actions. Agents play pure strategies, and aggregate behavior in the population is described by a *population state* $x = (x_i)_{i=1}^n \in \Delta_A$, with $\Delta_A \equiv \{x \in \mathbb{R}^n_+ \colon \sum_{i=1}^n x_i = 1\}$, where component x_i represents the fraction of agents in the population using strategy $a_i \in A$. The standard basis vector $e_i \in \Delta_A$ represents the pure (monomorphic) state at which all agents play strategy a_i.

We assume that there is a continuous expected-payoff function $\pi : \Delta_A \to \mathbb{R}^n$ such that $\pi_i(x)$ is the expected payoff to (a player using) strategy a_i at state x. If e_s is a strict Nash equilibrium state, i.e., if $\pi_s(e_s) > \pi_i(e_s)$ for every $i \neq s$, we say that strategy a_s is a strict Nash strategy.

A special case of this framework corresponds to a large population of agents who are randomly matched to play a p-player symmetric game with payoff function $U \colon A^p \to \mathbb{R}$, where $U(a_i; a_{j_1}, \ldots, a_{j_{p-1}})$ represents the payoff obtained by a strategy-a_i player whose opponents play strategies $a_{j_1}, \ldots, a_{j_{p-1}}$. When $p = 2$, we sometimes write U_{ij} instead of $U(a_i; a_j)$. The symmetry assumption implies that the value of U does not depend on the ordering of the last $p-1$ arguments. In this setting, the expected payoff function to strategy a_i at state x is the usual extension of U to the simplex Δ_A, i.e., $\pi_i(x) = \sum_{\bar{a} \in A^{(p-1)}} \left(\prod_{j=1}^n (x_j)^{\mathbb{I}_j(\bar{a})} \right) U(a_i; \bar{a})$, where $\bar{a}$ is a $(p-1)$-tuple of strategies (one for each co-player), and the exponent $\mathbb{I}_j(\bar{a})$ is the number of occurrences of strategy a_j in $\bar{a}$. The simplest case corresponds to symmetric 2-player normal form games with payoff matrix (U_{ij}), for which $\pi_i(x) = \sum_{j=1}^n U_{ij} x_j$. In the examples based on payoff matrices we will always be assuming this (standard) expected payoff function, but our results are valid for games defined by any other continuous payoff function $\pi : \Delta_A \to \mathbb{R}^n$.

We can establish the following links between the population game approach applied to a p-player game with random (or complete) matching, and traditional game theory. First, the expected payoff to strategy a_i at population state x coincides in the traditional setting with the expected payoff to a player using strategy a_i whose opponents play mixed strategy x. Second, e_s is a strict Nash state in the population game if and only if strategy profile $(a_s, a_s, ..., a_s)$ is a strict Nash equilibrium of the p-player game in the traditional setting. Strict Nash states in a population game are usually also called strict Nash equilibria, with the context indicating whether the term is referring to a population state e_s or to a strategy profile.

[8] The specific case of two-action coordination games played in one or two populations under SBR dynamics is discussed in detail in Arigapudi, Heller, and Schreiber (in press), including the stability of interior rest points and the effect of heterogeneous sample sizes.

2.2 Sampling Best Response dynamics with statistical inference

Under a sampling best response decision protocol with statistical inference, a revising agent observes the actions of k randomly-sampled players in a population, estimates a probability distribution for the state of the population, and chooses a best-response action (one of the best-response actions, if there is more than one) according to the estimated distribution of population states. The choice among optimal strategies is made following some *tie-breaking rule* β.

Consider a sample of size k. Let a *sample vector* (or just a *sample*) $z = (z_1, ..., z_n) \in \mathbb{N}_0^n$, with $\sum_{i=1}^n z_i = k$, be a vector indicating the occurrences of each strategy in a sample of size k, with z_j corresponding to the number of occurrences of strategy a_j in the sample. Let $\mathbb{N}_0^{n,k}$ be the set of all possible sample vectors z of n strategies with sample size k. For instance, if the number of strategies is $n = 3$ and the sample size is $k = 2$, we have that the set of possible sample vectors is

$$\mathbb{N}_0^{n,k} = \mathbb{N}_0^{3,2} = \{(2, 0, 0), (0, 2, 0), (0, 0, 2), (1, 1, 0), (1, 0, 1), (0, 1, 1)\}$$

At state x, the probability of obtaining sample z, $P_x(z)$, is given by the multinomial distribution:

$$P_x(z) = \binom{k}{z_1, \dots, z_n} x_1^{z_1} \dots x_n^{z_n}$$

with the convention that $0^0 = 1$.

For each sample $z \in \mathbb{N}_0^{n,k}$, an inference method G provides a probability distribution over population states, called an *estimate*. For instance, under maximum likelihood estimation, players who obtain sample z estimate the (unique) population state that maximizes the probability $P_x(z)$ of obtaining sample z, which is the state $\frac{z}{k}$.

In general, an estimate is characterized by a generalized[9] probability density function g_z over states with expected state $\bar{x}^G(z) = \int_{\Delta_A} x\, g_z(x)\, dx$ and with expected payoff vector $\pi^G(z) = \int_{\Delta_A} \pi(x)\, g_z(x)\, dx$. For instance, under maximum-likelihood ($G = ML$), we have the expected state $\bar{x}^{ML}(z) = \frac{z}{k}$ and the expected payoff vector $\pi_i^{ML}(z) = \pi_i(\frac{z}{k})$, with the payoff functions $\pi_i(\cdot)$ considered in section 2.1.

An inference method G is unbiased (Salant & Cherry, 2020) if, for every sample z, the expected state given the sample, $\bar{x}^G(z)$, coincides with the empirical proportions $\frac{z}{k}$ observed in sample z, i.e., if $\bar{x}^G(z) = \frac{z}{k}$. Maximum likelihood is an unbiased inference method that, given a sample z, assigns probability 1 to the single state $\frac{z}{k}$. Other inference methods can assign to a sample z a non-degenerate distribution over states. Typically,

[9] In order to consider continuous and discrete (or mixed) distributions under the same notation, as well as distributions with support on subsets of Δ_A such as the faces of Δ_A, we consider generalized density functions that may use the Dirac delta function δ. For instance, the density function corresponding to the maximum-likelihood estimate is given by $g_z(x) = \delta(x - \frac{z}{k})$; if the estimate is discrete and assigns positive probabilities to a finite set of states $\{(x)_1, ..., (x)_m\}$, with corresponding probabilities $\{p_1, ..., p_m\}$, then $g_z(x) = \sum_{i=1}^m p_i \delta(x - (x)_i)$; and a probability distribution characterized by a density function $f_z(x_i, x_j)$ with support on the edge of Δ_A defined by $x_i + x_j = 1$ is extended to Δ_A as $g_z(x) = f_z(x_i, x_j) \prod_{l \in A \setminus \{x_i, x_j\}} \delta(x_l)$.

inference methods that combine the data from a sample z with previous beliefs or information about the population state will not be unbiased.

Note that, for observed samples z with $z_j = 0$, an unbiased inference method must concentrate the estimated probability on states with $x_j = 0$. In particular, if the sample is monomorphic ($\frac{z}{k} = e_i$) then an unbiased estimate method assigns probability 1 to the state $x = e_i$ (the state with $x_i = 1$) and probability 0 to the set of states with $x_i < 1$. Consequently, for unbiased inference methods we have $\pi^G(k\,e_i) = \pi(e_i)$. We will use this result later to show that, for unbiased inference methods, strict Nash states are rest points of the considered dynamics.

To specify which best response action is chosen when several actions obtain the maximum expected payoff $\max_i \pi_i^G(z)$, we consider tie-breaking rules β. A tie-breaking rule β is a set of functions that indicate the probability $\beta_{ij}(\pi)$ with which a revising agent using strategy a_i and having expected payoff vector $\pi \equiv (\pi_h)_{h=1}^n$ switches to strategy a_j. We assume that revising agents select one of the strategies with the maximum expected payoff (according to their inference method), so a tie-breaking rule must satisfy $\beta_{ij}(\pi) = 0$ if $\pi_j < \max_h \pi_h$ (i.e., it only places weight on the estimated maximum-payoff strategies) and, for every i, $\sum_{j=1}^n \beta_{ij}(\pi) = 1$ (i.e., a revising agent selects one of the available strategies).

Well-known results of Benaïm and Weibull (2003) show that the behavior of a large but finite population following the protocol described above is closely approximated by the solution of its associated *mean dynamic*, a differential equation which describes the expected motion of the population at each state. The mean dynamic for a sampling best response protocol with statistical inference method G, sample size k and tie-breaking rule β, denoted by $\text{StatSBR}^{G,k,\beta}$, is[10]

$$\dot{x}_i = \sum_{z \in \mathbb{N}_0^{n,k}} P_x(z) \left[\sum_{h=1}^{n} x_h \, \beta_{hi}\left(\pi^G(z)\right) \right] - x_i \tag{1}$$

where:

- $P_x(z) = \binom{k}{z_1, \ldots, z_n} x_1^{z_1} \ldots x_n^{z_n}$ is the probability of obtaining sample z at state x.
- $\pi^G(z)$ is the vector of expected payoffs corresponding to inference method G when observing sample z.
- $\beta_{hi}(\pi)$ is the probability with which a revising a_h-strategist who obtains the vector of expected payoffs π adopts strategy a_i. This probability is i) 1 if strategy a_i is the only strategy with the maximum payoff in π, ii) 0 if strategy a_i does not have the maximum value in π, and iii) a number between 0 and 1, determined by the

[10]For the specific case of SBR dynamics (i.e., for the maximum likelihood inference method), Oyama et al. (2015) consider a k-sampling best response correspondence (instead of specific functions corresponding to particular tie-breaking rules) and derive a differential inclusion for the dynamics (instead of a differential equation adjusted to each tie-breaking rule). This approach allows to study properties that hold for every tie-breaking rule, but it does not allow to study the effect of different tie-breaking rules on the stability of an equilibrium, which can be relevant in our case (see e.g. figure 6).

tie-breaking rule, if strategy a_i is not the only strategy with the maximum value in π. In the examples, we will use the following tie breakers:

- *uniform*, which randomizes uniformly among all optimal strategies:

$$\beta_{ij}(\pi) = \begin{cases} \frac{1}{\#(\operatorname{argmax}_k \pi_k)} & \text{if } j \in \operatorname{argmax}_k \pi_k, \\ 0 & \text{in any other case.} \end{cases}$$

- *stick-uniform*, which selects the current strategy if it is optimal, and randomizes uniformly among optimal strategies if the current strategy is not optimal:

$$\beta_{ij}(\pi) = \begin{cases} 1 & \text{if } i = j \in \operatorname{argmax}_k \pi_k, \\ \frac{1}{\#(\operatorname{argmax}_k \pi_k)} & \text{if } i \notin \operatorname{argmax}_k \pi_k \text{ and } j \in \operatorname{argmax}_k \pi_k, \\ 0 & \text{in any other case.} \end{cases}$$

- *min*, which selects the optimal strategy with the smallest number:

$$\beta_{ij}(\pi) = \begin{cases} 1 & \text{if } j = \min\left(\operatorname{argmax}_k \pi_k\right), \\ 0 & \text{in any other case.} \end{cases}$$

Note that the differential flows in StatSBR$^{G,k,\beta}$ dynamics (right-hand side in (1)) are polynomials of degree no greater than $k+1$.

For the special case in which the tie-breaking rule is independent of the strategy followed by the revising agent (i.e., if there is a function β_i such that $\beta_{hi}(\pi) = \beta_i(\pi)$ for every $h \in \{1, ..., n\}$), the mean dynamic equations (1) simplify to:

$$\dot{x}_i = \sum_{z \in \mathbb{N}_0^{n,k}} P_x(z)\, \beta_i\left(\pi^G(z)\right) - x_i \tag{2}$$

And for the special case $k = 1$, the mean dynamic equations (1) for any unbiased inference method reduce to

$$\dot{x}_i = \sum_{j=1}^{n} x_j \sum_{h=1}^{n} x_h\, \beta_{hi}(\pi(e_j)) - x_i \tag{3}$$

For the mean dynamics to be a good approximation of the process in large finite populations, it is required that the fraction of agents who revise their strategy at any given time is small, so the state of the population changes slowly between successive revisions, advancing by many small steps. Intuitively, the fact that a large number of small stochastic steps is required before the population state can change significantly is the reason why the stochastic process tends to follow its expected direction of change, which is captured by the mean dynamics (for a technical discussion, see Sandholm (2010)).

3 Stability of monomorphic rest points

3.1 Stationarity. Strict Nash states

At a monomorphic population state $x = e_i$, the only sample that can be obtained is that in which each of the k sampled players uses strategy a_i, i.e., the sample $z = k\, e_i$. This is consequently the only sample that enters dynamics (1) at $x = e_i$. It is then easy to see that a monomorphic state e_i is a rest point of dynamics (1) if and only if, for $z = k\, e_i$:

- strategy a_i obtains the maximum expected payoff under G, and
- if some other strategy also obtains that maximum payoff, strategy a_i is selected.

We can now state the following observation:

Observation 3.1. *Strict Nash states are (monomorphic) rest points under StatSBR$^{G,k,\beta}$ dynamics for any unbiased inference method G.*

Observation 3.1 follows from considering that, if the inference method G is unbiased, then, at a monomorphic state e_s we have expected payoff vector $\pi^G(k\, e_s) = \pi(e_s)$. If e_s is a strict Nash state, the strict Nash strategy a_s is then the only strategy that achieves the maximum payoff in $\pi^G(k\, e_s)$. Consequently, the inflow (positive) and outflow (negative) terms in dynamics (1) at state $x = e_s$ (where $x_s = 1$) are both 1 for strategy a_s, and both 0 for each of the other strategies, so $\dot{x}_i = 0$ for every strategy a_i. If the inference method is not unbiased, it is easy to see that strict Nash states may not be rest points of the dynamics.

We now turn to analyzing the stability of monomorphic rest points under StatSBR$^{G,k,\beta}$ dynamics. The following subsection provides a brief summary of the main mathematical concepts we will use in our analysis.

3.2 Background on stability and linear stability

Consider a C^1 differential equation $\dot{x} = V(x)$ defined on a compact, convex set $X \subset \mathbb{R}^n$ (in our case, $X = \Delta_A$) whose forward solutions $(x(t))_{t\geq 0}$ do not leave X. State x^* is a *rest point* or equilibrium of the dynamics if $V(x^*) = 0$, so that the unique solution starting from x^* is stationary.

A rest point x^* is *Lyapunov stable* if for every neighborhood O of x^*, there exists a neighborhood O' of x^* such that every forward solution that starts in $O' \cap X$ is contained in O. If x^* is not Lyapunov stable it is *unstable*, and it is *repelling* (or a *repellor*) if there is a neighborhood O of x^* such that solutions from all initial conditions in $O \cap X$ leave O.

A rest point x^* is *attracting* if there is a neighborhood O of x^* such that all solutions that start in $O \cap X$ converge to x^*. If a rest point x^* is Lyapunov stable and attracting, it is *asymptotically stable*. In this case, the maximal (relatively) open[11] set of states in X from which solutions converge to x^* is called the *basin* of attraction of x^*. If the basin of

[11] A set is relatively open in X if it is the intersection of X with an open set in $\mathbb{R}^n$.

attraction of x^* is X itself, we call x^* *globally asymptotically stable*. If the basin of attraction of x^* contains the relative interior of X, we call x^* *almost globally asymptotically stable*.

By the definition of the derivative, the value of V in a (relative) neighborhood $O \cap X$ of a rest point x^* can be approximated via

$$V(x) = \mathbf{0} + DV(x^*)(x - x^*) + o(|x - x^*|)$$

where $DV(x^*)$ is the Jacobian matrix of V (more precisely, the Jacobian of a C^1 extension of V to R^n such that the first-order partial derivatives of the component functions of the extension are defined at x^*) evaluated at state x^*. The linear stability of x^* can be analyzed by considering the eigenvalues of $DV(x^*)$ corresponding to those eigenvectors lying in the tangent space $TX = \{z \in \mathbb{R}^n : \sum_i z_i = 0\}$. If all such eigenvalues have negative real parts, then x^* is linearly stable. If any of those eigenvalues has positive real part, then x^* is linearly unstable. A linearly stable rest point is asymptotically stable, and solutions starting near the rest point converge to it at an exponential rate (Perko, 2001; Sandholm, 2010).

3.3 Jacobian of StatSBR$^{G,k,\beta}$ dynamics at monomorphic rest points

Let us now consider the linear stability analysis of monomorphic rest points under StatSBR$^{G,k,\beta}$ dynamics. To do this, we study the Jacobian of the dynamics at such equilibrium states.

At states near a monomorphic state e_s, those samples z in which most players use strategy a_s are most likely to happen. It is then convenient to rewrite the inflow terms in (1),

$$I_i^{G,k,\beta} \equiv \sum_{z \in \mathbb{N}_0^{n,k}} P_x(z) \left[\sum_{h=1}^{n} x_h \beta_{hi} \left(\pi^G(z) \right) \right],$$

singling out the sample where only strategy a_s is observed, i.e., the sample $z = k\, e_s$, and also the samples where strategy a_s has been observed exactly $(k-1)$ times (and some other strategy $a_j \neq a_s$ has been observed exactly once), i.e. the $(n-1)$ samples $z^{s,j} \equiv e_j + (k-1)e_s$, for $j \neq s$. Doing so, and considering that e_s is a rest point, leads to the following expression:

$$I_i^{G,k,\beta}(x) = x_s^k\, \delta_{is} + k\, x_s^{k-1} \sum_{j \neq s} x_j \left[\sum_{h=1}^{n} x_h \beta_{hi} \left(\pi^G(z^{s,j}) \right) \right] + f_i^{G,k,\beta}(x) \tag{4}$$

where δ_{is} is the Kronecker delta and where

$$f_i^{G,k,\beta}(x) \equiv \sum_{z \in \mathbb{N}_0^{n,k} : z_s < k-1} \binom{k}{z_1, \ldots, z_n} x_1^{z_1} \ldots x_n^{z_n} \left[\sum_{h=1}^{n} x_h \beta_{hi} \left(\pi^G(z) \right) \right].$$

Note that $\frac{\partial f_i^{G,k,\beta}}{\partial x_j}(e_s) = 0$, since the partial derivatives of all the monomials $x_1^{z_1} \dots x_n^{z_n}$ in $f_i^{G,k,\beta}$ take the value 0 at e_s (this is so because, in each monomial, the sum of the exponents of the variables that are not x_s is at least 2, and those variables take the value 0 at e_s).

The terms of the Jacobian of the inflow function $I^{G,k,\beta}$ at a rest point $x = e_s$, denoted by $DI^{G,k,\beta}(e_s)$, are consequently $DI_{ij}^{G,k,\beta}(e_s) \equiv \frac{\partial I_i^{G,k,\beta}}{\partial x_j}(e_s) = k\delta_{is} + k\beta_{si}(\pi^G(z^{s,j}))$. This means that the only relevant payoffs for a linear stability analysis of a monomorphic rest point e_s under StatSBR$^{G,k,\beta}$ dynamics are $\pi_i^G(z^{s,j})$, for $i, j \in \{1, \dots, n\}$. Furthermore, if we eliminate coordinate x_s by considering $x_s = 1 - \sum_{i \neq s} x_i$, we have that, for $i, j \neq s$, the terms of the corresponding (reduced) Jacobian of StatSBR$^{G,k,\beta}$ dynamics at the equilibrium e_s are:[12]

$$k\,\beta_{si}\left(\pi^G(z^{s,j})\right) - \delta_{ij}$$

Considering a rest point e_s and a tie-breaking rule β, let us define α_{ij}^k as follows (omitting in the notation the dependence of α_{ij}^k on s, G and β):

$$\alpha_{ij}^k \equiv k\,\beta_{si}\left(\pi^G(z^{s,j})\right)$$

The terms of the reduced Jacobian are then $\alpha_{ij}^k - \delta_{ij}$. A positive value $\alpha_{ij}^k > 0$ indicates that sample $z^{s,j} = (k-1)e_s + e_j$ leads a revising a_s-player who obtains such a sample (with one a_j-player) to adopt strategy a_i with positive probability.

The stability of e_s depends on how the presence of other strategies in states near e_s reinforce the growth of those same strategies. For any non-empty subset of strategies $J \subseteq A \setminus \{a_s\}$ and any strategy $a_i \in J$, let $\alpha_{i\cdot}^k(J) \equiv \sum_{j:a_j \in J} \alpha_{ij}^k$. We can interpret $\alpha_{i\cdot}^k(J)$ as the total growth support that strategy a_i receives near e_s from the strategies in J (including itself, via α_{ii}^k). And let $\alpha_{\cdot i}^k(J) \equiv \sum_{j:a_j \in J} \alpha_{ji}^k$, which we can interpret as the total growth support that strategy a_i provides near e_s to the strategies in J.

Depending on how much growth support near e_s a strategy a_i receives from (or provides to) other strategies in a subset J containing a_i, we have the following definitions, which will be useful to analyze the stability of e_s.

Definition 1. *Consider a StatSBR$^{G,k,\beta}$ dynamic and a rest point e_s, with associated strategy a_s. Let $J \subseteq A \setminus \{a_s\}$ be a non-empty subset of strategies other than a_s. Under the considered StatSBR$^{G,k,\beta}$ dynamic, a strategy $a_i \in J$ is:*

- *e_s-stabilizing in J if either $\alpha_{\cdot i}^k(J) = 0$ or $\alpha_{i\cdot}^k(J) = 0$.*
- *Potentially e_s-stabilizing in J if $\alpha_{\cdot i}^k(J) \leq 1$.*
- *Potentially e_s-stabilized in J if $\alpha_{i\cdot}^k(J) \leq 1$.*

[12]For details, see the proof of proposition 3.2.

3.4 Asymptotic stability

In this section we provide a condition that guarantees asymptotic stability of monomorphic rest points under StatSBR$^{G,k,\beta}$ dynamics. We then use this result to analyze the stability of strict Nash equilibria under unbiased inference methods, and finally we provide specific results for the maximum likelihood inference method (i.e., for SBR dynamics).

Before analyzing local stability, we note that, if the game has a strictly dominant strategy a_s, its associated state is globally asymptotically stable under every StatSBR$^{G,k,\beta}$ dynamic. This follows from a_s being the unique best response at every state.

Proposition 3.2. *Let e_s be a monomorphic rest point under a StatSBR$^{G,k,\beta}$ dynamic. If no strategy survives the iterated elimination[13] of e_s-stabilizing strategies in $A \setminus \{a_s\}$, then e_s is asymptotically stable.*

The proof of proposition 3.2 proceeds by showing that, if no strategy survives iterated elimination, then all the relevant eigenvalues of the Jacobian of dynamics (1) at the rest point are negative. If there is a non-empty set of strategies J_{surv} which survive iterated elimination, the rest point may still be stable. In fact, if either $\alpha^k_{\cdot i}(J_{surv}) < 1$ for every $i \in J_{surv}$ or $\alpha^k_{i\cdot}(J_{surv}) < 1$ for every $i \in J_{surv}$ then the rest point is asymptotically stable.[14]

Example 3.1. Consider a three-player coordination game with available actions C and D. If players coordinate on the action profile (C, C, C) each one obtains payoff 10; if they coordinate on (D, D, D), each one obtains payoff 1. Otherwise, each one obtains payoff 0. The state of the population is the vector $x = (x_C, x_D)$.

Let us study the stability of the inefficient strict Nash equilibrium $e_D = (0, 1)$ under SBR$^{k,\beta}$ dynamics. We only have another strategy, C, so we only need to calculate the term[15] $\alpha^k_{CC} = k\,\beta_{DC}\left(\pi^G(z^{D,C})\right)$ where $z^{D,C}$ represents a sample with one C-player and $(k-1)$ D-players. The estimated state given sample $z^{D,C}$ is $(\bar{x}_C, \bar{x}_D) = (\frac{1}{k}, \frac{k-1}{k})$, and its associated expected payoff vector is $\pi^G(z^{D,C}) = (\pi_C, \pi_D) = (10\frac{1}{k^2}, \frac{(k-1)^2}{k^2})$. For $k \geq 5$ we have $\pi_C < \pi_D$, so $\alpha^k_{CC} = 0$. Proposition 3.2 shows that e_D is asymptotically stable for $k \geq 5$.

Clearly, if all strategies other than a_s are e_s-stabilizing in $A \setminus \{a_s\}$, then the condition in proposition 3.2 holds (we just need to consider one elimination step). Our two next propositions are based on that result. Proposition 3.3 shows that, if the inference method is unbiased, strict Nash states are asymptotically stable for sufficiently large sample size. Proposition 3.4 applies to maximum likelihood estimation (i.e., to SBR$^{k,\beta}$) and establishes a minimal sample size that guarantees asymptotic stability.

[13] For a formal description of the process of iterated elimination see Appendix A.1.

[14] This can be shown by considering that, for a non-negative matrix, the maximum row sum (and the maximum colum sum) is an upper bound for its spectral radius (Horn & Johnson, 1985).

[15] For clarity, to refer to the strategies here, we use their letters C and D instead of their numbers 1 and 2.

Proposition 3.3. *Let e_s be a strict Nash equilibrium of a symmetric game. Under any StatSBR$^{G,k,\beta}$ dynamic with unbiased inference method, e_s is a rest point and there is a finite k_0 such that e_s is asymptotically stable for $k > k_0$.*

The proof of proposition 3.3 considers first that, given that e_s is a strict Nash state, there is a relative neighborhood O around e_s in which a_s is the unique best response to the population state. Next, it shows that, if the inference method G is unbiased, then, given a sample $z^{s,j} = e_j + (k-1)e_s$, the probability that the estimate g_z assigns to the set of states outside O must tend to zero as k grows. Eventually, a_s must be the unique strategy obtaining the maximum payoff in any vector $\pi^G(z^{s,j})$ of estimated payoffs. By proposition 3.2, e_s is then asymptotically stable, as all the other strategies are e_s-stabilizing in $A \setminus \{a_s\}$.

Proposition 3.4. *Let e_s be a strict Nash equilibrium of a symmetric game. If e_s is $(1-\frac{1}{k_0})$-dominant for some $k_0 \in \mathbb{N}$ then e_s is asymptotically stable under every SBR$^{k,\beta}$ dynamic with $k \geq k_0$.*

The proof of proposition 3.4 considers that, if e_s is $(1-\frac{1}{k_0})$-dominant, then a_s is the unique best response at any state $\frac{z^{s,j}}{k}$ with $k \geq k_0$.

In the special case of two-player games with linear payoffs, the following result provides an explicit upper bound for the number of trials beyond which there is asymptotic stability.

Corollary 3.5. *Let e_s be a strict Nash equilibrium of a two-player symmetric game with payoff matrix (U_{ij}) and standard payoff functions $\pi_i(x) = \sum_{j=1}^{n} U_{ij}x_j$. For every $k > 1 + \max_{i\neq s, j\neq s} \frac{U_{ij}-U_{sj}}{U_{ss}-U_{is}}$, e_s is asymptotically stable under SBR$^{k,\beta}$ dynamics.*

If we apply corollary 3.5 to the two-strategy case, we find that a strict Nash state e_1 is asymptotically stable for every $k > \frac{U_{11}-U_{21}+U_{22}-U_{12}}{U_{11}-U_{21}}$. It is also easy to show that, for those values of k, e_1 is $(1-\frac{1}{k})$-dominant. Stability in the special case of equality $k = \frac{U_{11}-U_{21}+U_{22}-U_{12}}{U_{11}-U_{21}}$ depends on the tie-breaking rule.

3.5 Instability

In this section we provide sufficient conditions for monomorphic rest points to be unstable, and even repelling, under StatSBR$^{G,k,\beta}$ dynamics.

Proposition 3.6. *Let e_s be a monomorphic rest point under a StatSBR$^{G,k,\beta}$ dynamic.*

- *If no strategy is potentially e_s-stabilizing in $S \setminus \{a_s\}$, then state e_s is repelling.*
- *If some strategy survives the iterated elimination of potentially e_s-stabilizing strategies in $S \setminus \{a_s\}$, then state e_s is unstable. The same result holds for potentially e_s-stabilized strategies.*

The proof of the first part of proposition 3.6, i.e., the repelling result, shows that, under the indicated conditions, any small deviation from the equilibrium state e_s that ends up in a small relative neighborhood of e_s will be amplified by the dynamics, until the solution trajectory leaves that relative neighborhood. The proof of the second part of proposition 3.6, i.e., the instability result, makes use of lower bounds on the Perron-Frobenius eigenvalue of non-negative matrices to show that the Jacobian of the dynamics at e_s has (at least) one positive eigenvalue, corresponding to an eigenvector in the tangent space of Δ_A that goes from e_s into Δ_A.

Example 3.2. For the coordination game of example 3.1 and the inefficient strict Nash equilibrium e_D under SBR$^{k,\beta}$ dynamics, we have $\alpha^k_{CC} = k$ for $k \leq 4$. Applying proposition 3.6 we find that, for $2 \leq k \leq 4$, e_D is a repellor.

Corollary 3.7. *Let e_s be a monomorphic rest point under a StatSBR$^{G,k,\beta}$ dynamic. If there is some strategy a_i (other than a_s) with $\alpha^k_{ii} > 1$, then e_s is unstable.*

If we apply corollary 3.7 to SBR dynamics in two-player two-strategy games (with standard payoff functions $\pi_i(x) = \sum_{j=1}^{2} U_{ij}x_j$), we find that a strict Nash state e_1 is unstable for $2 \leq k < \frac{U_{11}-U_{21}+U_{22}-U_{12}}{U_{11}-U_{21}}$. It is also easy to show that for those values of k (if there are any), e_1 is not $(1-\frac{1}{k})$-dominant, while it is so for $k > \frac{U_{11}-U_{21}+U_{22}-U_{12}}{U_{11}-U_{21}}$.

Example 3.3. Consider the following game, with strategy set $\{a, b, c\}$ and payoff matrix:

$U_2 =$

	a	b	c
a	2	0	0
b	0	0	3
c	0	3	0

In game U_2, players can coordinate on the strict equilibrium e_a, but action a is not $\frac{1}{2}$-dominant (no proper subset of strategies is). For $k = 2$ and $e_s = e_a$, it is easy to check that, under SBR dynamics, $\alpha^{k=2}_{cb} = \alpha^{k=2}_{bc} = 2$, so strategies b and c are not potentially e_a-stabilizing in $\{b, c\}$. Thus, using proposition 3.6 we can state that e_a is a repellor under SBR dynamics with $k = 2$. These dynamics are illustrated in figure 3 for different tie breakers.

Note that the analysis of this game does not change if we add some fixed amount m to every payoff in the same column of U_2. In fact, by doing so, we can turn any unstable strict equilibrium (such as e_a, for $k = 2$) into the Pareto efficient solution of a game, but this change in efficiency does not change the stability properties of the equilibria under SBR.

3.6 λ-sampling best response dynamics with statistical inference

Instead of considering a fixed sample size k, and following the equivalent generalization in Oyama et al. (2015), we can consider λ-sampling best response dynamics with

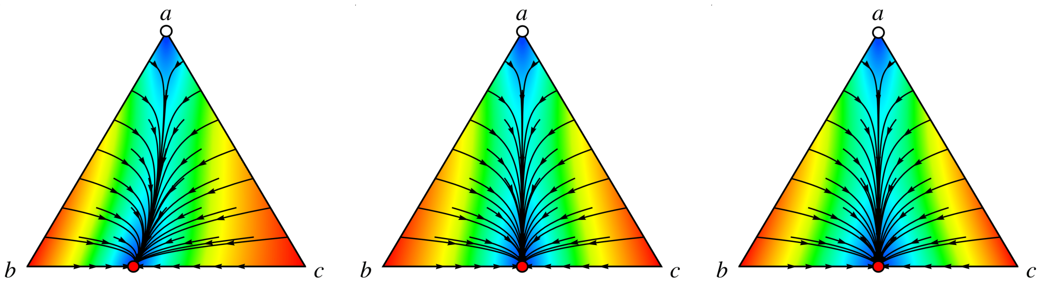

(i) *min* tie breaker (ii) *uniform* tie breaker (iii) *stick-uniform* tie breaker

Figure 3: Sampling best response dynamics with sample size $k = 2$, for the game with payoff matrix U_2, using different tie breakers.

statistical inference, under which the size of an agent's sample is a random draw from a discrete probability distribution λ on the natural numbers. In this case, let λ_k be the probability that a revising agent takes a sample of size k when sampling strategies from the population state, or, equivalently, the fraction of agents that take samples of size k when revising their strategy (this fraction is assumed to be independent of the agent's strategy).

The dynamics in this case are given by

$$\dot{x}_i = I_i^{G,\lambda,\beta}(x) - x_i = \sum_{k \in \mathbb{N}} \lambda_k \, I_i^{G,k,\beta}(x) - x_i,$$

and it follows easily from our previous results that

$$DI_{ij}^{G,\lambda,\beta}(e_s) \equiv \frac{\partial I_i^{G,\lambda,\beta}}{\partial x_j}(e_s) = \sum_{k \in \mathbb{N}} \lambda_k \, k \, \beta_{si}\left(\pi^G(z^{s,j})\right) = \sum_{k \in \mathbb{N}} \lambda_k \, \alpha_{ij}^k.$$

We can consequently define

$$\alpha_{ij}^{\lambda} = \sum_{k \in \mathbb{N}} \lambda_k \, \alpha_{ij}^k, \qquad \alpha_{i\cdot}^{\lambda}(J) = \sum_{j: a_j \in J} \alpha_{ij}^{\lambda} \qquad \text{and} \qquad \alpha_{\cdot i}^{\lambda}(J) = \sum_{j: a_j \in J} \alpha_{ji}^{\lambda}$$

and the adaptations of definition 1, proposition 3.2 and proposition 3.6 are immediate.

Note that, if for a given rest point e_s, inference method G and sample size $k = k_0$ the payoff vectors $\pi^G(z^{s,j})$ are k_0-generic, meaning that (for every strategy $a_j \neq a_s$) the vector $\pi^G(z^{s,j})$ has a unique strategy that obtains the maximum payoff, then $\alpha_{ij}^{k_0} \in \{0, k_0\}$. Consequently, if the payoffs are k_0-generic and proposition 3.6 shows instability for $k = k_0$, then the adaptation of proposition 3.6 for λ-sampling implies instability for any distribution λ satisfying $k_0 \, \lambda_{k_0} > 1$, i.e., if $\lambda_{k_0} > \frac{1}{k_0}$.

For instance, if e_s satisfies the instability conditions in proposition 3.6 for $k = 2$, then (assuming genericity for $k = 2$) e_s is unstable under any distribution λ satisfying $\lambda_2 > \frac{1}{2}$.

Similarly, if it follows from proposition 3.6 that e_s is repelling for $k = 2$ and for $k = 3$, then (assuming generic payoffs $\pi^G(z^{s,j})$ for those sample sizes) e_s is repelling under any distribution λ satisfying $2\lambda_2 + 3\lambda_3 > 1$.

If the tie-breaking rule does not depend on the strategy of the revising agent (see equation (2)), then the formulas in this section are also valid under the alternative assumption that agents have time-invariant yet different sample sizes, with λ_k being the fraction of players in the population using sample size k.

4 An index to rank strict Nash equilibria

We can assign a natural number $index(e_s)$ to each strict Nash state e_s in a game according to the smallest sample size k_0 such that the state is asymptotically stable under SBR$^{k,\beta}$ dynamics, for every $k \geq k_0$ and for every β (i.e., regardless of the tie-breaking rule).

Having index 1 is quite restrictive: in order to have index 1, a strict Nash state needs to be unique (see corollary A.3, appendix A.3). If there is a strictly dominant strategy, then the index of its corresponding (unique) strict Nash state is indeed 1 (proposition 3.4). Proposition 3.3 shows that the index of a strict Nash state is always a finite number. Proposition 3.4 can be used to find an upper bound for the index, and, for two-player games, corollary 3.5 provides an explicit upper bound on the index.

For an intuitive informal interpretation of the index of an equilibrium, note that we can consider SBR$^{k,\beta}$ dynamics as *smoothed* best response dynamics. The approximation is better the higher the value of the sample size k (see figure 4). Mathematically, we can think of SBR$^{k,\beta}$ dynamics as polynomial approximations to best response dynamics, of degree (at most) $k + 1$.[16] Under best response dynamics (BR), every strict Nash state is asymptotically stable and has some basin of attraction. This is also the case under SBR dynamics for sufficiently large values of k. However, as k decreases, the SBR dynamics become increasingly dissimilar from BR (in the sense of corresponding to polynomial approximations of lower degree; see figure 4), and some equilibria may (and typically will) become unstable and lose their basins of attraction (see figure 5). The index k_0 of a strict Nash state is the smallest number that guarantees that it is asymptotically stable for every $k \geq k_0$. While it is often the case that the equilibrium is not asymptotically stable for every $k < k_0$, this is not necessarily so.[17]

The index of an equilibrium is a measure of how robust the local stability of that strict equilibrium is to the noise, limited information or smoothing that is implicit in a SBR decision protocol with small sample size. Lower index indicates more robustness to the perturbations on the exact best response dynamics created by the SBR process.

Observation 4.1 (Index of a strict Nash state in symmetric 2x2 games)**.** *In the special case of two-player two-strategy games with linear payoffs* $\pi_i(x) = \sum_{j=1}^{2} U_{ij} x_j$*, it follows from our previous results that:*

[16]For 2-player 2-strategy games played in one population, SBR$^{k,\beta}$ dynamics are the Berstein polynomial approximation to best response dynamics.

[17]An equilibrium may be stable for some $k = k_1$ but unstable for some $k' > k_1$. For instance, for the game with payoff matrix (rows) $[(3, 1, 0), (0, 0, 5), (2, 0, 2.1)]$, e_1 is stable for $k = 2$ and unstable for $k = 3$.

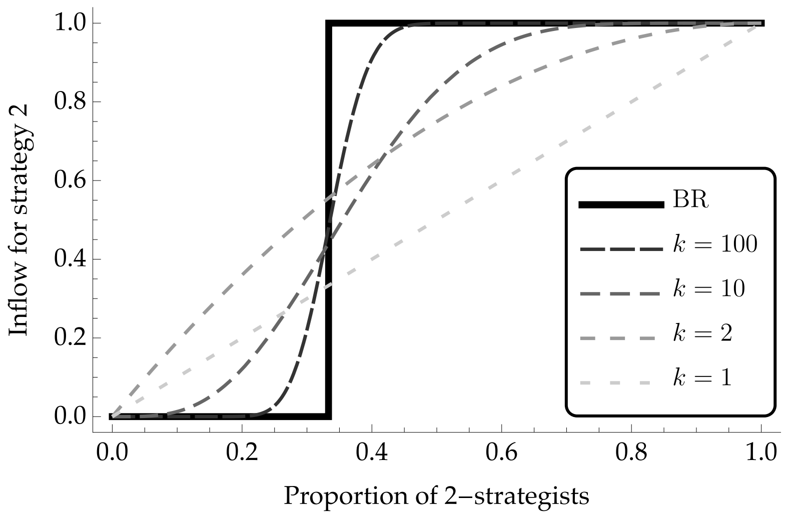

Figure 4: Inflow for strategy 2 as a function of the proportion of 2-strategists, in the pure coordination game $[(1,0),(0,2)]$, for best response dynamics (BR) and for sampling best response dynamics $\mathrm{SBR}^{k,\beta}$ for different values of the sample size k and any tie breaker β. Note that the outflow for strategy 2 in this graph corresponds to the segment that connects points $(0,0)$ and $(1,1)$. This observation allows the reader to easily infer the stability and the size of the basin of attraction of each of the two strict Nash states for the different dynamics.

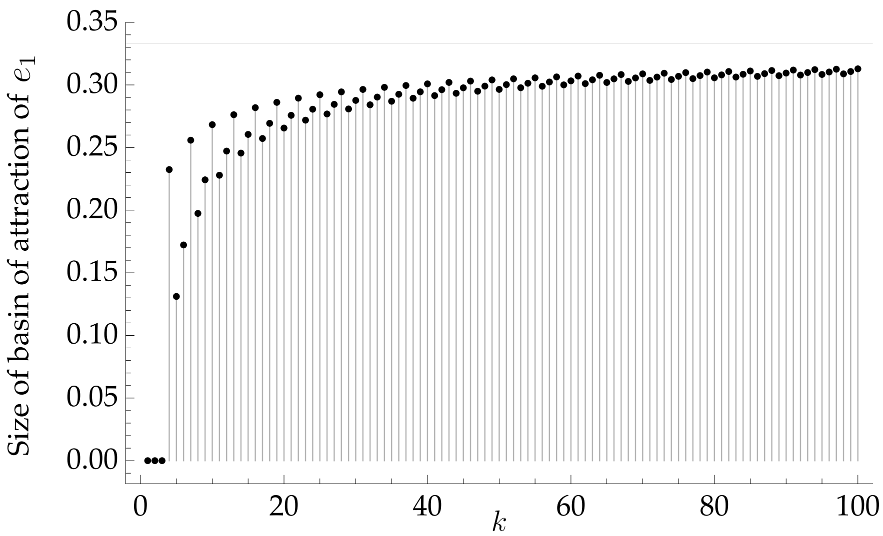

Figure 5: Size of the basin of attraction of strict Nash state e_1 in the pure coordination game $[(1,0),(0,2)]$, for sampling best response dynamics $\mathrm{SBR}^{k,\beta=\mathrm{uniform}}$, for different values of the sample size k.

- *The index of a strict Nash state e_s is 1 if and only if strategy a_s is strictly dominant.*
- *The index of a strict Nash state e_s is the smallest integer k such that a_s is $(1-\frac{1}{k})$-dominant.*
- *In coordination games with two (symmetric) strict Nash equilibria, an equilibrium is risk-dominant if and only if its index is 2, while the index of a non-risk-dominant equilibrium is greater than 2. To be precise, assuming without loss of generality that e_1 is non-risk-dominant, its index is $\lfloor \frac{U_{11}-U_{21}+U_{22}-U_{12}}{\min(U_{11}-U_{21},U_{22}-U_{12})} \rfloor + 1$.*

Example 4.1. Consider the following pure coordination game, with strategy set $\{1,2\}$ and payoff matrix

$$\begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$$

Applying the result for two-player two-strategy coordination games presented in observation 4.1, we can easily state that the index of e_1 is $\lfloor \frac{1+2}{\min(1,2)} \rfloor + 1 = 4$ and the index of e_2 is 2, since it is risk-dominant. Figures 4 and 5 provide several insights for this game.

Example 4.2. Consider the following game, with strategy set $\{1,2\}$ and payoff matrix:

$$\begin{bmatrix} 10 & 8 \\ 0 & U_{22} \end{bmatrix}$$

Applying the results in observation 4.1, we can state:

- If $U_{22} < 8$, a_1 is strictly dominant (0-dominant) and the index of the unique strict Nash equilibrium e_1 is consequently 1.
- If $8 < U_{22} < 18$, there are two strict Nash equilibria and a_1 is risk-dominant ($\frac{1}{2}$-dominant), so the index of e_1 is 2 (this is also its index for $U_{22} = 8$), while the index of e_2 is $\lfloor \frac{U_{22}+2}{U_{22}-8} \rfloor + 1 \geq 3$.
- If $U_{22} > 18$, there are two strict Nash equilibria: a_2 is risk-dominant ($\frac{1}{2}$-dominant), so the index of e_2 is 2, while the index of e_1 is $\lfloor \frac{U_{22}+2}{10} \rfloor + 1 \geq 3$.

Example 4.3. Consider the following game, with strategy set $\{a,b,c\}$ and payoff matrix:

$U_3 =$

	a	b	c
a	1	1	1
b	0	2	2
c	0	0	4

It follows from propositions 3.2 and 3.6 that:

- e_a is asymptotically stable for $k \geq 5$, and unstable for $k = 2,3$ (see figures 6 and 7). Its stability for $k = 4$ depends on the tie-breaking rule (see figure 8), so its index is 5.

- e_b and e_c are asymptotically stable for $k \geq 3$ (see figures 7 and 8). Their stability for $k = 2$ depends on the tie-breaking rule (see figure 6), so their index is 3.

Note that in this example, no proper subset of strategies is $\frac{1}{2}$-dominant, so the results in Oyama et al. (2015) do not apply.

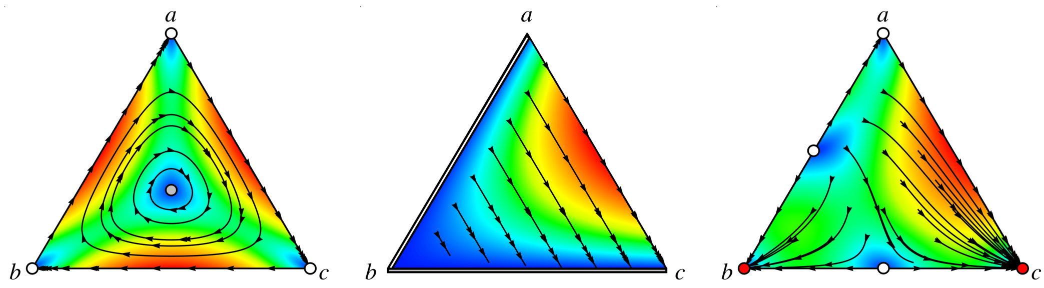

(i) *min* tie breaker (ii) *uniform* tie breaker (iii) *stick-uniform* tie breaker

Figure 6: Sampling best response dynamics with sample size $k = 2$, for the game with payoff matrix U_3, using different tie breakers. Connected components of rest points are represented with solid lines, colored according to their stability in the same way as rest points: red if asymptotically stable, gray if Lyapunov stable (but not asymptotically stable), and white if unstable.

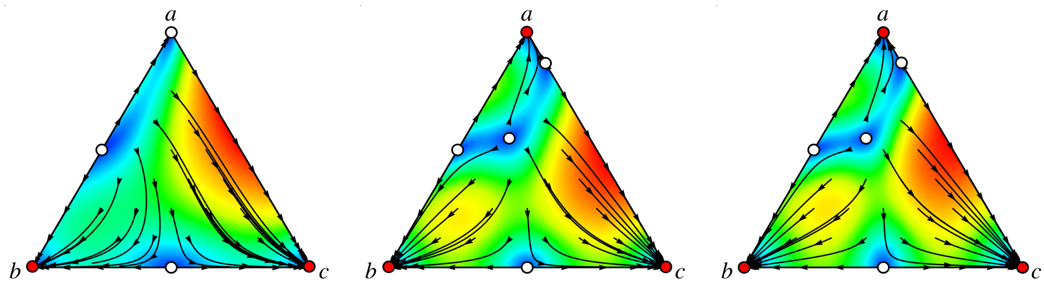

(i) $k = 3$ (ii) $k = 5$ (iii) $k = 10$

Figure 7: Sampling best response dynamics for the game with payoff matrix U_3, for *uniform* tie breaker and various sample sizes k.

5 Application to tacit coordination games

Tacit coordination games, also known as *minimum-effort*, *weak-link* or *weakest-link* games (Camerer, 2003; Engelmann & Normann, 2010; Feri, Gantner, Moffatt, & Erharter, 2022; Van Huyck, Battalio, & Beil, 1990), constitute a paradigmatic case of how individuals can show a tendency to coordinate on some strict Nash equilibria rather than others. In these games, all symmetric strict Nash equilibria satisfy most equilibrium refinements and

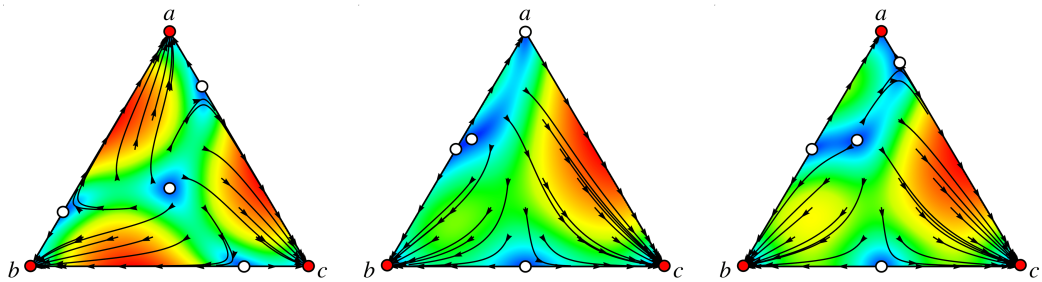

(i) *min* tie breaker (ii) *uniform* tie breaker (iii) *stick-uniform* tie breaker

Figure 8: Sampling best response dynamics with sample size $k = 4$, for the game with payoff matrix U_3, using different tie breakers.

correspond to evolutionarily stable states.[18] However, experimental evidence clearly shows that human subjects do discriminate between different strict equilibria in these games, and their behavior clearly depends on the number of players.

In a tacit coordination game, a group of individuals must decide how much effort to put into a common project. Each unit of effort has an individual cost b, and the output of the project depends solely on the minimum of the individual efforts, i.e. m; specifically, each player gets the minimum individual effort m times the unitary benefit a, with $a > b$. The greater the minimum effort, the greater the profit players will obtain. However, if any individual works less than the rest, the extra effort put by the others goes to waste.

Formally, tacit coordination games are symmetric p-player games with strategy set $A = \{1, ..., n\}$ (denoting the player's effort or contribution) and payoff function

$$U(i; j_1, \ldots, j_{p-1}) = a \min(i, j_1, \ldots, j_{p-1}) - b\, i,$$

where $a > b \geq 0$ are the two parameters controlling the benefit and the cost of effort units, respectively.[19]

Table 1 represents the payoff function for the p-player 3-strategy case ($n = 3$). The row headings on the payoff matrices in table 1 indicate the strategy chosen by the player that receives the payoff. The column headings indicate the minimum value of the strategies chosen by the other $(p - 1)$ players.

Here we focus on the case $b > 0$. Note that every homogeneous pure strategy profile $(i, i, ..., i)$, in which all the p players choose the same effort level i, is a strict Nash equilibrium, and these equilibria are strictly Pareto ranked, with their efficiency growing with the effort level i. However, given any strategy profile, selecting the lowest strategy chosen by the other players is the unique best reply. This means that, at any

[18] All symmetric strict Nash equilibria are evolutionarily stable according to the standard definition of evolutionary stability (Weibull, 1995). Crawford (1991) provides a detailed analysis of these games and shows that the only equilibrium state that satisfies a finite-population definition of evolutionary stability is the secure state e_1.

[19] Often, in experiments, a constant value c is added to every payoff to establish (material or monetary) rewards, so that they fall within a desired range.

Table 1: Payoff matrices for a p-player tacit coordination game with three strategies ($n = 3$). Left: general case. Right: $a = 2$ and $b = 1$.

	min of others' strategies		
	1	2	3
1	$a - b$	$a - b$	$a - b$
2	$a - 2b$	$2a - 2b$	$2a - 2b$
3	$a - 3b$	$2a - 3b$	$3a - 3b$

	min of others' st.		
	1	2	3
1	1	1	1
2	0	2	2
3	−1	1	3

equilibrium $(i, \ldots, i)$, if any player deviates to a lower strategy $j < i$, then following suit and changing to strategy j is the unique best response for the other players. This creates a tension that can induce players to lower their strategy or "effort" as soon as any other player does, or as soon as players *believe* that any other player *may* do it. The only Pareto efficient state is e_n, while the minimax or *secure* profile corresponds to state e_1 (the minimum effort is the strategy that guarantees the greatest worst-case payoff).

Experimental evidence on these games shows that people's behavior depends on the number of players in the game and on the benefit/cost ratio $\frac{a}{b}$. When the game is played in very small groups,[20] experimental results show a clear tendency to choose the efficient highest-effort strategy (Engelmann & Normann, 2010; Van Huyck et al., 1990).[21] In contrast, in groups with several players, the distribution of strategies is initially diverse, and then the vast majority of players approach the lowest effort fairly quickly.[22] This clear pattern of discrimination between strict Nash equilibria, dependent on the number of players and against efficiency in the case of large groups, is difficult to explain along the lines of traditional game theory (Crawford, 1991) [23].

An underlying intuition for this effect of the number of players is that, even if a player were to expect every other player to choose the efficient-equilibrium strategy n with high probability, the larger the number of players, the more likely some player will deviate from the efficient-equilibrium strategy n, and this can make the efficient strategy unattractive. The effect of this uncertainty about which of the possible equilibria other players expect, and consequently about what strategy they will choose, is known as *strategic uncertainty* (Andersson, Argenton, & Weibull, 2014; Van Huyck et al., 1990).

Besides increasing with the number of players, the coordination failure has been shown to decrease with the benefit/cost ratio $\frac{a}{b}$ (Brandts & Cooper, 2006; Feri et al., 2022; Goeree & Holt, 2005), so for low benefit/cost ratios the tendency to low effort can be found already for $p = 2$, while for larger ratios the tendency to high effort persists for

[20] With $\frac{a}{b} = 2$, "small groups" meaning usually $p = 2$ players, sometimes 3 or 4 players per group.

[21] In most cases, the experiment is repeated with the same group of co-players. Van Huyck et al. (1990) also present results on setups where players were randomly paired after every period. In that case, they did not find any stable pattern of behavior.

[22] With $\frac{a}{b} = 2$, this is almost always the case for groups with more than 4 players.

[23] Monderer and Shapley (1996) show that tacit coordination games are potential games, with the potential being maximal at the lowest-effort profile (for $p > \frac{a}{b}$) or at the highest-effort profile (for $p < \frac{a}{b}$). Maximal potential, however, does not imply global convergence or even stability under SBR dynamics, as can be seen in figure 9.

greater numbers of players.

Under sampling best response $SBR^{k,\beta}$ dynamics with $k > 1$, we find (see proofs in the appendix) the following results for the stability of each strict Nash equilibrium.[24] We focus on the effect of the number of players $p \geq 2$, considering the sample size k and the benefit/cost ratio $\frac{a}{b}$ as parameters (see figure 9).

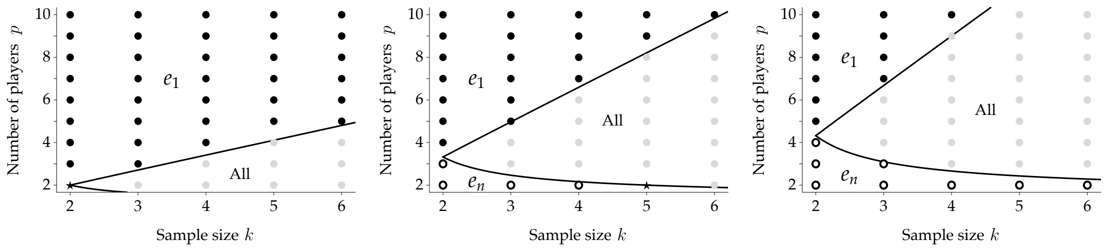

(i) benefit/cost ratio $a/b = 2$ (ii) benefit/cost ratio $a/b = 5$ (iii) benefit/cost ratio $a/b = 10$

Figure 9: Stability of strict Nash states under $SBR^{k,\beta}$ dynamics in tacit coordination games, for different values of sample size k and number of players p, and for different values of the benefit/cost ratio a/b. A black circle located at (k, p) indicates that e_1 is an almost global attractor. Gray circles indicate that all strict Nash states are asymptotically stable, and white circles indicate that e_n is an almost global attractor. Stars indicate borderline cases which depend on the tie breaker.

- If the number of players p is large enough, specifically, if $p > 1 + \frac{\log(a/b)}{\log(k/(k-1))}$, then the lowest-effort state e_1 is almost globally asymptotically stable, because it is an iterated $1/k$-dominant equilibrium (Oyama et al., 2015). This implies that all the other strict Nash states are unstable.
- If the number of players is $p < 1 + \log_k(\frac{a}{b})$, then the highest-effort state e_n is almost globally asymptotically stable, because it is an iterated $1/k$-dominant equilibrium. Note that, if $k < \frac{a}{b}$, this condition is satisfied at least for $p = 2$; while if $k > \frac{a}{b}$, the condition cannot hold, for any $p \geq 2$. As we increase the benefit/cost ratio $\frac{a}{b}$, the number of players for which e_n is almost globally asymptotically stable may increase.
- If the number of players is in the range $\left(1 + \log_k(\frac{a}{b}), 1 + \frac{\log(a/b)}{\log(k/(k-1))}\right)$, then every strict Nash state is asymptotically stable.

In order to observe how an increment in the number of players p favors the stability of the lowest-effort state e_1 and the instability of the highest-effort state e_n under $SBR^{k,\beta}$ dynamics, we can consider a small sample size such as $k = 3$ and a benefit/cost ratio $\frac{a}{b} > 3$ (see figures 9(ii), 9(iii)). In this case:

- For two players ($p = 2$), the Pareto efficient highest-effort state e_n is almost globally asymptotically stable (while all the other strict Nash states are unstable). This can also hold for more players, specifically, it holds while $p < 1 + \log_3(\frac{a}{b})$.

[24]For $k = 1$, we have $\dot{x} = 0$ at every state.

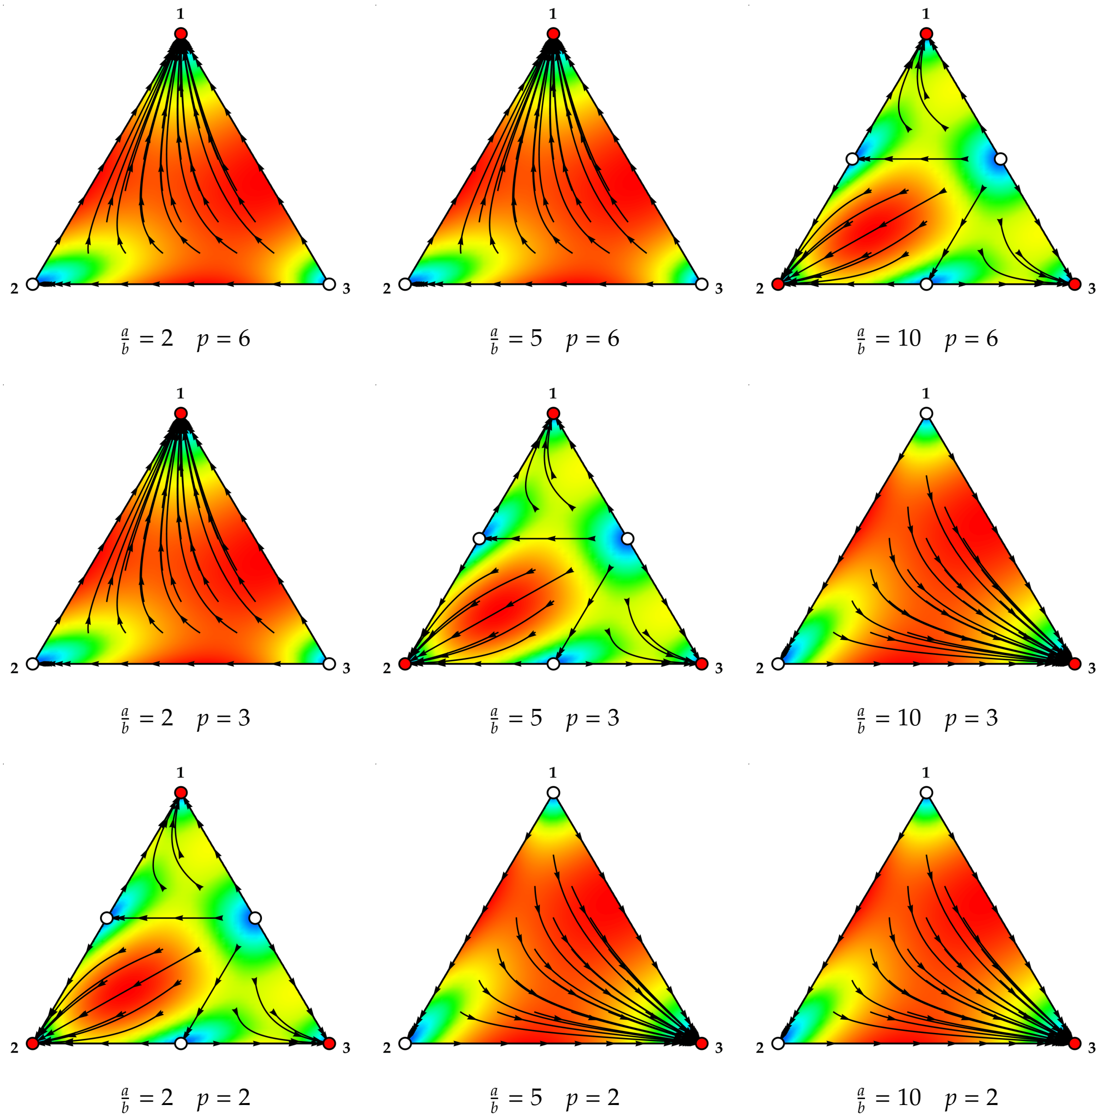

Figure 10: Sampling best response dynamics $\mathrm{SBR}^{k=3,\beta}$ dynamics in tacit coordination games, for different values of the benefit/cost ratio $\frac{a}{b}$ and number of players p. The figures are valid for any tie breaker.

- For number of players in the range $\left(1+\log_3(\frac{a}{b}), 1+\frac{\log_3(\frac{a}{b})}{1-\log_3 2}\right)$, every strict Nash state is asymptotically stable.
- For $p > 1+\frac{\log_3(\frac{a}{b})}{1-\log_3 2}$, the lowest-effort state e_1 is almost globally asymptotically stable (while all the other strict Nash states are unstable).

To conclude, the stability analysis of the strict equilibria of tacit coordination games under SBR dynamics captures two of the more salient experimental findings in these games (see figure 10): the effect of increasing the number of players (which favors the lowest-effort state e_1) and the effect of increasing the benefit/cost ratio (which favors the highest-effort state e_n). This qualitative insight can be quantified in a formal way using the concept of index: the index of the secure state e_1 (i.e. its vulnerability to sampling noise) is a –weakly– increasing function of the benefit/cost ratio and a –weakly– decreasing function of the number of players (see figure 11(i)). By contrast, the index of the efficient state e_n –weakly– increases with the number of players and –weakly– decreases with the benefit/cost ratio (see figure 11(ii)).

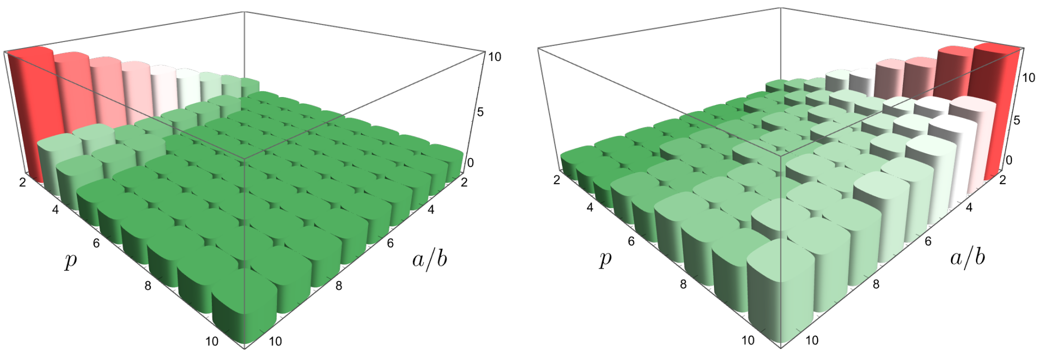

(i) Index of secure state e_1 (ii) Index of efficient state e_n

Figure 11: Index of the secure state e_1 and of the efficient state e_n in tacit coordination games, for different values of the number of players p and of the benefit/cost ratio a/b. For clarity, we only consider integer values of the benefit/cost ratio a/b.

6 Conclusions

In this paper we have derived deterministic approximations (mean dynamics) for sampling best response protocols with statistical inference, characterized the Jacobian of such dynamics at monomorphic rest points, and developed simple tests for asymptotic stability and for instability of such equilibria. The tests can be easily applied to λ-sampling best response dynamics with statistical inference, under which the size of an agent's sample is a random draw from a discrete probability distribution λ on the natural numbers. Considering the specific case of maximum likelihood estimation, or *sampling best response*, the analysis of stability can assign an index to each strict equilibrium in a game, measuring how stable or robust each equilibrium is under small-size sampling (best response) protocols. This index also provides a method to select among the different strict equilibria in a game. In the 2x2 case, this selection method coincides with the *risk dominance* criterion. Tacit coordination or weakest-link games constitute a paradigmatic case study of how players tend to coordinate in some particular strict equilibria instead of others, and, in these games, the analysis of stability of each equilibrium under sampling best response dynamics presents the same qualitative features that are usually observed in experimental studies.

Acknowledgements

We thank several anonymous reviewers for their comments. Financial support from the Spanish State Research Agency (PID2024-159461NB-I00/MICIU and PID2020-118906GB-I00/MCIN, AEI/10.13039/501100011033/EU-FEDER), from the Regional Government of Castilla y León and the EU-FEDER program (CLU-2019-04 - BIOECOUVA Unit of Excellence of the University of Valladolid), and from the Spanish Ministry of Universities (PRX22/00064 and PRX22/00065), is gratefully acknowledged.

A Appendixes

A.1 Iterated elimination of strategies

Here we define the set of survivors of iterated elimination of strategies satisfying condition $\mathbb{C}$ in a finite set Ω. Let $J^0 \equiv \Omega$ and define J^m recursively by

$$J^m = \{i \in J^{m-1} \mid i \text{ does not satisfy condition } \mathbb{C} \text{ in } J^{m-1}\}.$$

The (potentially empty) set $J^{|\Omega|}$ is the set of strategies that survive iterated elimination of strategies satisfying condition $\mathbb{C}$ in set Ω. An algorithm for this procedure is described in algorithm 1.

Algorithm 1 Iterated elimination of strategies satisfying condition $\mathbb{C}$ in set Ω

$J \leftarrow \Omega$
while $\exists j \in J \mid j$ satisfies condition $\mathbb{C}$ in J **do**
  $J \leftarrow J \setminus \{j \in J \mid j$ satisfies condition $\mathbb{C}$ in $J\}$
end while ▷ J at the end is the set of all surviving strategies after iterated elimination

A.2 Sampling Best Response vs. Best Experienced Payoff dynamics

Sampling Best Response with statistical inference (StatSBR) and Best Experienced Payoff (BEP) protocols are both based on sampling (the strategies played by) a group or several groups of players. The decision method that a revising agent is assumed to follow is different in each case. Let us assume that the game is a p-player game played in one population.

Under StatSBR, one single sample (of size k) is obtained and it is used to estimate a population state (more precisely, to obtain an estimate: a probability distribution over the states). The estimate provides one (expected) payoff for each strategy. Players are assumed to be able to make inferences about the population state and to calculate expected payoffs, so they need to know the payoff functions of the game.

Under BEP, each available strategy is tested k times, each time by playing against an independent sample of $(p-1)$ co-players, so, for each strategy, k independent samples of size $(p-1)$ are obtained. Each strategy obtains a total (experienced) payoff, by playing against its k associated samples. Players do not need to make inferences about the population state, or calculate expected payoffs (they may not even know the payoff functions).

When studying the stability of the dynamics at a strict Nash state, the fact that both processes are based on sampling allows for some parallel simplifications in order to calculate the Jacobian. For BEP, those simplifications are discussed by Sandholm et al. (2020), Arigapudi et al. (2021) and Izquierdo and Izquierdo (2022). For StatSBR, those simplifications are discussed here. The stability results can then be quite different depending on the protocol. For instance, strictly dominant strategies are global attractors under StatSBR dynamics, but they can be unstable under BEP dynamics (Osborne & Rubinstein, 1998; Sandholm et al., 2020; Sethi, 2000), and even repelling.[25]

A.3 Dynamics with unbiased inference method and sample size $k = 1$

Let us consider here sampling best response dynamics with an unbiased inference method and sample size $k = 1$. These dynamics are the same for every unbiased inference method, so they are $SBR^{1,\beta}$ dynamics:

$$\dot{x}_i = \sum_{j=1}^{n} x_j \sum_{h=1}^{n} x_h \beta_{hi}(\pi(e_j)) - x_i \tag{5}$$

Observation A.1. *Every strict Nash state is Lyapunov stable under $SBR^{1,\beta}$ dynamics.*

This observation comes from noting that, if e_s is a strict Nash state, then $\beta_{hs}(\pi(e_s)) = 1$ and $SBR^{1,\beta}$ dynamics in (5) guarantee $\dot{x}_s \geq 0$, so every strict Nash state is Lyapunov stable under $SBR^{1,\beta}$ dynamics. Note, however, that strict Nash states may or may not be *asymptotically* stable under $SBR^{1,\beta}$. Example 1.1 showed a game where no strict Nash state is asymptotically stable under $SBR^{1,\beta}$. On the other hand, if there is a strictly dominant strategy, its associated strict Nash state is *globally* asymptotically stable under every $SBR^{k,\beta}$ dynamic.

[25]See, for instance, Mantilla, Sethi, and Cárdenas (2020, section 5.2) and Sandholm et al. (2020, example 5.1), who prove that the dominant-strategy equilibrium in a "Voluntary exchange" game is repelling under a wide range of best experienced payoff protocols.

In order to characterize which strict Nash equilibria are asymptotically stable under $SBR^{1,\beta}$ dynamics, let us consider the *pure best response correspondence*, which assigns to each strategy a_i the set of pure best responses to e_i (i.e., the set of strategies that are best response to the population state in which every player uses strategy a_i).

Proposition A.2. *A necessary condition for a strict Nash state e_s to be asymptotically stable under $SBR^{1,\beta}$ dynamics is that $\{a_s\}$ is the only minimal set of pure strategies closed under the pure best response correspondence.*

The proof of proposition A.2 is based on the fact that, under $SBR^{1,\beta}$ dynamics, if a_s is a strict Nash strategy and the set of strategies J is closed under the pure best response correspondence, with $a_s \notin J$, then the set of states whose support is in $\{a_s\} \cup J$ is forward invariant, and at any of those states we have $\dot{x}_s = 0$, precluding convergence to e_s from some points in any relative neighborhood of e_s.

Considering that every strict Nash strategy is a minimal set of strategies closed under the pure best response correspondence, we have the following corollary of proposition A.2.

Corollary A.3. *A necessary condition for a strict Nash equilibrium e_s to be asymptotically stable under $SBR^{1,\beta}$ dynamics is to be unique, i.e., that there is no other strict Nash equilibrium in the game.*

If every monomorphic state has a unique best response, then $\{a_s\}$ being the only minimal set of pure strategies closed under the pure best response correspondence is equivalent to a_s being the only strategy that survives iterated elimination of strategies that are not best response to any other strategy (in the surviving set). One could then expect the asymptotic disappearance of all such iteratively eliminated strategies under $SBR^{1,\beta}$ dynamics, as our next proposition shows. In this case, being the only minimal set of pure strategies closed under the pure best response correspondence implies global asymptotic stability. This result can also be extended to cases with non unique best response to monomorphic states, if the tie-breaking rule places positive probability on every best response.

Proposition A.4. *Assuming that each monomorphic state has a single best response, state e_s is globally asymptotically stable under $SBR^{1,\beta}$ dynamics if and only if $\{a_s\}$ is the only minimal set of pure strategies closed under the pure best correspondence.*

If monomorphic states present several best-responses (i.e., if there are ties), the result of proposition A.4 still holds for tie-breaking rules that place positive probability on every best response.[26]

[26]For a_s satisfying the condition, it can be shown, using Lemma 1 of Appendix B in Sandholm et al. (2020), that, starting at any initial ($t = 0$) state such that $x_s(t = 0) < 1$, the result $\int_0^T \dot{x}_s(t)dt > 0$ holds for every $T > 0$, leading to global convergence to e_s.

A.4 Proofs

Proof of proposition 3.2. This proof is based on a related result for BEP dynamics in Izquierdo and Izquierdo (2022). Remember that, for $i, j \neq s$, $DI_{ij}^{G,k,\beta}(e_s) = k\,\beta_{si}(\pi^G(z^{s,j})) \equiv \alpha_{ij}^k$. Consider in (1) a change of variables for the population state $(x_1, x_2, ..., x_n)$ that sends the equilibrium e_s to the origin $\mathbf{0}$, by eliminating the coordinate $x_s = 1 - \sum_{i \neq s} x_i$ while keeping the labeling of the other coordinates. In this reduced system W, consider the Jacobian of dynamics (1) at the equilibrium, $DW(\mathbf{0})$, whose components (for $i, j \neq s$) are $DW_{ij}(\mathbf{0}) = DI_{ij}^{G,k,\beta}(e_s) - DI_{is}^{G,k,\beta}(e_s) - \delta_{ij} = \alpha_{ij}^k - \delta_{ij}$, where δ_{ij} is the Kronecker delta, and where it has been considered that, for $i \neq s$, $DI_{is}^{G,k,\beta}(e_s) = 0$ (see section 3.3). Let $J \subseteq A \setminus \{a_s\}$ be a non-empty subset of strategies (not containing strategy a_s) and let $DW_J(\mathbf{0})$ be the square submatrix of $DW(\mathbf{0})$ whose rows and columns correspond to the strategies in J, i.e., the principal submatrix of $DW(\mathbf{0})$ corresponding to J. If a_j is e_s-stabilizing in J, then either the column or the row of $DW_J(\mathbf{0})$ corresponding to strategy a_j is made up by zeros in all non-diagonal positions, with a value -1 at the diagonal position. Let $(a_{j_1}, a_{j_2}, ..., a_{j_{n-1}})$ be an ordering of the $(n-1)$ strategies in $A \setminus \{a_s\}$ that iteratively eliminates e_s-stabilizing strategies. Then either the column or the row of $DW(\mathbf{0})$ corresponding to strategy a_{j_1} is made up by zeros in all non-diagonal positions, with a value -1 at the diagonal position. Considering the cofactor expansion of the determinant of the Jacobian along the column or row corresponding to a_{j_1}, and denoting by $DW_{-\{j_1\}}(\mathbf{0})$ the submatrix of $DW(\mathbf{0})$ obtained by eliminating the column and row corresponding to a_{j_1}, we have that $|DW(\mathbf{0})| = (-1)\,|DW_{-\{j_1\}}(\mathbf{0})|$. Now, either the column or row of $DW_{-\{j_1\}}(\mathbf{0})$ corresponding to strategy a_{j_2} is made up by zeros in all non-diagonal positions, with a value -1 at the diagonal position. Proceeding sequentially with the other strategies we obtain $|DW(\mathbf{0})| = (-1)\,|DW_{-\{j_1\}}(\mathbf{0})| = (-1)^2\,|DW_{-\{j_1,j_2\}}(\mathbf{0})| = ... = (-1)^{n-1}$, i.e., all the eigenvalues of the Jacobian have negative real parts, which implies asymptotic stability of the equilibrium.

□

Proof of proposition 3.3. By continuity of the payoff function π, if a_s is a strict Nash strategy then there is a positive $\epsilon > 0$ and a positive $\delta > 0$ such that $\pi_s(x) - \pi_j(x) > \delta$ for $x \in \{x : x_s > 1 - \epsilon\}$ and for every $j \neq s$. We will show next that, if the estimation method G is unbiased, then the probability of the set $\{x : x_s > 1 - \epsilon\}$ (where $\pi_s(x) - \pi_j(x) > \delta$) under G, given sample $z^{s,j}$, tends to 1 as $k \to \infty$, which (given that the payoffs are bounded) implies that, eventually (i.e., for every k greater than some finite k_0), strategy a_s is the only strategy to obtain the maximum in the payoff vector $\pi^G(x)$, and, by proposition 3.2 (considering that all strategies in $A \setminus \{a_s\}$ are then e_s-stabilizing), e_s is asymptotically stable. To see that the probability of the set $\{x : x_s > 1 - \epsilon\}$ tends to 1 as $k \to \infty$ note that, for unbiased G and sample $z^{s,j}$, the estimate g_z assigns probability 0 to (the set of) states with $x_i > 0$ for $i \notin \{j, s\}$ and, for the probability of states with $x_s + x_j = 1$, there is a generalized probability density function $f_{z^{s,j}} : [0, 1] \to [0, 1]$ such that

$$\frac{1}{k} = \bar{x}_j^G(z^{s,j}) = \int_0^1 x_j f_{z^{s,j}}(x_j)\, dx_j = \int_0^{\epsilon} x_j f_{z^{s,j}}(x_j)\, dx_j + \\ + \int_{\epsilon^+}^1 x_j f_{z^{s,j}}(x_j)\, dx_j \geq \epsilon \int_{\epsilon^+}^1 f_{z^{s,j}}(x_j)\, dx_j$$

where $\int_{\epsilon^+}^1 f_{z^{s,j}}(x_j)\, dx_j$ is the probability (under G, given sample $z^{s,j}$) of the states with $x_s \leq 1-\epsilon$. This probability is consequently bounded above by $\frac{1}{k\epsilon}$ (for a fixed $\epsilon > 0$), and therefore tends to 0 as $k \to \infty$, so the probability of the complement set $\{x : x_s > 1-\epsilon\}$ tends to one. If the estimation method is maximum likelihood (SBR protocol), we find asymptotic stability of e_s for $k > \frac{1}{\epsilon}$. □

Proof of proposition 3.4. At any state $\frac{z^{s,j}}{k}$ the fraction of players using strategy a_s is $\frac{k-1}{k} = 1 - \frac{1}{k}$. For $k \geq k_0$, this fraction is greater or equal than $1 - \frac{1}{k_0}$. Consequently, if e_s is $(1-\frac{1}{k_0})$-dominant, then a_s is the only best response at any state $\frac{z^{s,j}}{k}$ with $k \geq k_0$, so $\alpha_{ij}^k = 0$ for every other strategy $a_i \neq a_s$, i.e., all the other strategies are e_s-stabilizing in $A \setminus \{a_s\}$. □

Proof of corollary 3.5. In two-player games, with payoffs U_{ij}, we have $\pi_i(\frac{z^{s,j}}{k}) = \frac{k-1}{k}U_{is} + \frac{1}{k}U_{ij}$. Consequently, $k(\pi_s(\frac{z^{s,j}}{k}) - \pi_i(\frac{z^{s,j}}{k})) = (k-1)(U_{ss} - U_{is}) + U_{sj} - U_{ij}$. It follows that, if $k - 1 > \frac{U_{ij}-U_{sj}}{U_{ss}-U_{is}}$, then $\pi_s(\frac{z^{s,j}}{k}) > \pi_i(\frac{z^{s,j}}{k})$, i.e., a_s is a better response to $\frac{z^{s,j}}{k}$ than a_i. The bound for k stated in the corollary guarantees that this condition holds for every $i \neq s$ and every $j \neq s$. □

Proof of proposition 3.6. As in the proof of proposition 3.2, consider the Jacobian of the (reduced) dynamics (1) at the equilibrium, $DW(\mathbf{0})$, whose components are $\alpha_{ij}^k - \delta_{ij}$, with $i, j \neq s$.

1. Repulsion.

   Let us first prove the repelling result, based on Sethi (2000) and Sandholm et al. (2020). From (4), the linearization of the reduced dynamics (where x_s has been eliminated) around the equilibrium state $\mathbf{0}$ is $\dot{x}_i = \sum_{j\neq s} \alpha_{ij}^k x_j - x_i$. Consequently,

$$\sum_{i\neq s} \dot{x}_i = \sum_{i\neq s}\sum_{j\neq s} \alpha_{ij}^k x_j - \sum_{i\neq s} x_i = \sum_{j\neq s}\left(\sum_{i\neq s} \alpha_{ij}^k\right) x_j - \sum_{i\neq s} x_i \tag{6}$$

   Let $\alpha \equiv \min_j \sum_{i\neq s} \alpha_{ij}^k$. If no strategy is potentially e_s-stabilizing in $S \setminus \{a_s\}$, then $\sum_{i\neq s} \alpha_{ij}^k > 1$ for every $j \neq s$, so $\alpha > 1$. From (6) we have that the linearized system satisfies

$$\sum_{i\neq s} \dot{x}_i \geq \sum_{j\neq s} \alpha x_j - \sum_{i\neq s} x_i = (\alpha - 1)\sum_{i\neq s} x_i$$

This implies that we can find a positive value ϵ_0 and a positive value $\alpha_1 < \alpha - 1$ such that, in the partial neighborhood of $\mathbf{0}$ where $\sum_{i \neq s} x_i < \epsilon_0$, the non-linearized dynamics satisfy

$$\sum_{i \neq s} \dot{x}_i \geq \alpha_1 \sum_{i \neq s} x_i$$

Consequently, every trajectory starting in the considered partial neighbourhood where $\sum_{i \neq s} x_i < \epsilon_0$, at a point that is not $\mathbf{0}$, will eventually leave the neighborhood (with $\sum_{i \neq s} x_i(t)$ growing at exponential rate), proving that $\mathbf{0}$ is a repellor in the reduced dynamics and, equivalently, that e_s is a repellor in the original dynamics (1).

2. Instability.

   This proof is based on Sandholm et al. (2020) (proof of prop. 5.4). If the iterated elimination of potentially e_s-stabilizing strategies does not eliminate all strategies in $S \setminus \{a_s\}$, then there is some non-empty subset of strategies $J \subseteq S \setminus \{a_s\}$ which does not contain any potentially e_s-stabilizing strategies. This implies that if $a_i \in J$ then $\alpha^k_{\cdot i}(J) > 1$.

   Let $DW^+(\mathbf{0}) \equiv DW(\mathbf{0}) + I_{n-1}$ be the non-negative matrix obtained by adding the identity matrix to $DW(\mathbf{0})$, and whose components are consequently $\alpha^k_{ij} \geq 0$, and let $DW^+_J(\mathbf{0})$ be the principal submatrix of $DW^+(\mathbf{0})$ corresponding to the strategies in J. The addition of the values in each column of $DW^+_J(\mathbf{0})$ is then $\alpha^k_{\cdot i}(J) > 1$. As $DW^+(\mathbf{0})$ is a non-negative real matrix, its Perron-Frobenius eigenvalue ρ is bounded below by the minimum column sum in (any one of) its principal submatrices (Horn & Johnson, 1985, corollary 8.1.20 and theorem 8.1.22), and, consequently, $\rho \geq \min_{i:a_i \in J} \alpha^k_{\cdot i}(J) > 1$. Note that if λ is an eigenvalue of $DW^+(\mathbf{0})$ then $\lambda' = \lambda - 1$ is an eigenvalue of $DW(\mathbf{0})$. Consequently, $DW(\mathbf{0})$ has a positive eigenvalue $(\rho - 1) > 0$ corresponding to a non-negative eigenvector, proving that e_s is unstable under SBR dynamics.

   The proof for potentially e_s-stabilized strategies is equivalent and rests on the fact that the Perron eigenvalue ρ of $DW^+(\mathbf{0})$ is also bounded below by the minimum row sum in (any one of) its principal submatrices (Horn & Johnson, 1985).

□

Proof of proposition A.2. Suppose, by contradiction, that there is some non-empty subset of strategies $J \subseteq A \setminus \{a_s\}$ such that J is closed under the (pure) best response correspondence BR, with $BR(a_i) \equiv \{a_j \in A : \pi_j(e_i) = \max_{h \in \{1,\ldots,n\}} \pi_h(e_i)\}$. Then, considering (3), we have that the face $\Delta_{J \cup \{a_s\}} \equiv \{x \in \Delta_A : \sum_{i:a_i \in J \cup \{a_s\}} x_i = 1\}$ is invariant and, along every trajectory starting at a point in that face, $\dot{x}_s = 0$ (so $x_s(t)$ is constant). Consequently, for every (partial) neighborhood O of e_s there are trajectories starting at states in O (at least those in the intersection of $O \setminus e_s$ with the face $\Delta_{J \cup \{a_s\}}$) that do not converge to e_s, so e_s is not asymptotically stable. Note, however, that this does not rule out that all trajectories starting at the relative interior of Δ_A converge to e_s, even if e_s is not asymptotically stable. □

Proof of proposition A.4. It follows from proposition A.2 that the condition is necessary for asymptotic stability of e_s, and, consequently, it is necessary for global asymptotic stability. Let us show that it is sufficient for global asymptotic stability. Under the conditions in proposition A.4, dynamics (5) simplify to

$$\dot{x}_i = \sum_{j=1}^{n} x_j \beta_i(\pi(e_j)) - x_i$$

or, in matrix form, letting $\mathcal{B}$ be a matrix with components $\mathcal{B}_{ij} \equiv \beta_i(\pi(e_j))$,

$$\dot{x} = (\mathcal{B} - I)x \tag{7}$$

which is a constant coefficient homogeneous linear differential system with solution $x = e^{t(\mathcal{B}-I)}x_0$. If we eliminate the component corresponding to the strict Nash state e_s (considering that, for $i \neq s, \beta_i(\pi(e_s)) = 0$), then, using the subscript $_{-s}$ to indicate that component x_s has been eliminated, we get to the following reduced system:

$$\dot{x}_{-s} = (\mathcal{B}_{-s} - I)\, x_{-s} \tag{8}$$

Note that $\mathcal{B}_{-s}$ is a non-negative matrix, and the addition of the values in each of its columns (i.e. $\sum_{i \neq s} \mathcal{B}_{ij} = \sum_{i \neq s} \beta_i(\pi(e_j))$ for column j) is bounded above by 1, which implies that its Perron-Frobenius eigenvalue is also bounded above by 1 (Horn & Johnson, 1985, theorem 8.1.22). In fact, its Perron-Frobenius eigenvalue must be less than 1, because otherwise there would be (besides a_s) another minimal set of strategies closed under the best response correspondence. This implies that all the eigenvalues of $(\mathcal{B}_{-s} - I)$, which are obtained by substracting 1 from the eigenvalues of $\mathcal{B}_{-s}$, have negative real part, and, consequently, that the linear flow $e^{t(\mathcal{B}_{-s}-I)}$ corresponding to (8) is a contraction (Hirsch & Smale, 1974), so every trajectory tends to the origin as $t \to \infty$, proving global convergence of (7) to e_s. □

Proof of the results for tacit coordination games. We will show that, for $k > 1$:

1. If $\frac{a}{b} < (\frac{k}{k-1})^{p-1}$ then the lowest-effort state e_1 is an iterated $1/k$-dominant equilibrium, which implies that it is almost globally asymptotically stable (Oyama et al., 2015).

2. If $\frac{a}{b} > k^{p-1}$ then e_n is an iterated $1/k$-dominant equilibrium, so it is almost globally asymptotically stable.

3. If $(\frac{k}{k-1})^{p-1} < \frac{a}{b} < k^{p-1}$ then every strict Nash state is asymptotically stable.

Proof of result 1.

When a player is matched with $(p-1)$ co-players at state x, the probability that the minimum of their contributions is n (i.e, the maximum possible contribution) is x_n^{p-1}. The probability that that minimum is less than n is $(1 - x_n^{p-1})$. By considering the two last rows of the payoff table for a tacit coordination game (see table 1), we have

$$\pi_n(x) - \pi_{n-1}(x) = (a-b)x_n^{p-1} - b(1 - x_n^{p-1}) = ax_n^{p-1} - b$$

Consequently, if $x_n^{p-1} < \frac{b}{a}$ then n is not a best response to state x. It follows that if $x_n \leq (1 - \frac{1}{k})$ and $(1 - \frac{1}{k})^{p-1} < \frac{b}{a}$, then strategy n is not a best response. In other words, if $(1 - \frac{1}{k})^{p-1} < \frac{b}{a}$, then $\{1, 2, ..., n-1\}$ is a $1/k$-best response set (Oyama et al., 2015). Now, if we eliminate strategy n and consider the restriction of the game to the strategy set $\{1, 2, ..., n-1\}$, we can apply the same argument to the strategy with the maximum contribution in the set (i.e., strategy $n-1$), and repeat the process until only strategy 1 remains, showing that e_1 is an iterated $1/k$-dominant equilibrium. Condition $(1 - \frac{1}{k})^{p-1} < \frac{b}{a}$ can alternatively be written as $\frac{a}{b} < (\frac{k}{k-1})^{p-1}$, or as $p > 1 + \frac{\log(a/b)}{\log(k/(k-1))}$.

Proof of result 2.

When a player is matched with $(p-1)$ co-players at state x, the probability that the minimum of their contributions is greater than 1 (i.e, greater than the minimum possible contribution) is $(1 - x_1)^{p-1}$. The probability that that minimum equals one is $1 - (1 - x_1)^{p-1}$. By considering the two first rows of the payoff table for a tacit coordination game (see table 1), we have

$$\pi_2(x) - \pi_1(x) = (a-b)\,(1 - x_1)^{p-1} - b\,[1 - (1 - x_1)^{p-1}] = a\,(1 - x_1)^{p-1} - b$$

Consequently, if $(1 - x_1)^{p-1} > \frac{b}{a}$ then contributing 1 is not a best response to state x. It follows that if $x_1 \leq (1 - \frac{1}{k})$ and $(\frac{1}{k})^{p-1} > \frac{b}{a}$, then strategy 1 is not a best response. In other words, if $\frac{a}{b} > k^{p-1}$, then $\{2, ..., n-1\}$ is a $1/k$-best response set. Now, if we eliminate strategy 1 and consider the restriction of the game to the strategy set $\{2, ..., n\}$, we can apply the same argument to the strategy with the minimum contribution in the set (i.e., strategy 2), and repeat the process until only strategy n remains, showing that e_n is an iterated $1/k$-dominant equilibrium. Condition $\frac{a}{b} > k^{p-1}$ can alternatively be written as $p < 1 + \log_k(\frac{a}{b})$.

Proof of result 3.

Let us calculate $\pi_i(\frac{z^{s,j}}{k})$. At state $\frac{z^{s,j}}{k}$: the fraction of a_j-players in the population is $\frac{1}{k}$, and the fraction of a_s-players is $\frac{k-1}{k}$; then, the probability that all (randomly met) $(p-1)$ co-players use strategy a_j is $\frac{1}{k^{p-1}}$, and the probability that they all use strategy a_s

is $(\frac{k-1}{k})^{p-1}$; the expected payoff to strategy a_i is then

$$\pi_i\left(\frac{z^{s,j}}{k}\right) = a\left(\min(i,s)\left(\frac{k-1}{k}\right)^{p-1} + \min(i,j)\frac{1}{k^{p-1}}\right) + a\min(i,j,s)\left(1 - \frac{1}{k^{p-1}} - \left(\frac{k-1}{k}\right)^{p-1}\right) - bi$$

For $j \in \{1, ..., n\}$ and $i \in \{1, ..., n-1\}$, the difference $\pi_{i+1}(\frac{z^{s,j}}{k}) - \pi_i(\frac{z^{s,j}}{k})$ is:

$$(a-b) > 0, \qquad \text{if } i < \min(j,s). \tag{9}$$

$$-b < 0, \qquad \text{if } i \geq \max(j,s). \tag{10}$$

$$\left(\frac{k-1}{k}\right)^{p-1} a - b, \qquad \text{if } j \leq i < s. \tag{11}$$

$$\frac{1}{k^{p-1}} a - b, \qquad \text{if } s \leq i < j. \tag{12}$$

If $\left(\frac{k}{k-1}\right)^{p-1} < \frac{a}{b} < k^{p-1}$, which requires $k > 2$:

- The differences (9) and (11) are positive, and the differences (10) and (12) are negative.
- For any s and j we have that $\pi_i\left(\frac{z^{s,j}}{k}\right)$ increases with i for $i < s$ and decreases with i for $i > s$, which implies that $\alpha_{ij}^k = 0$ for every $i \neq s$ and, by proposition 3.2, e_s is asymptotically stable.

□

A.5 Two-population games

Here we indicate how to extend our results to games played in two populations (X and Y). We focus on the stability and instability conditions for monomorphic rest points (propositions 3.2 and 3.6).

Let A_X be the set of n_X actions or pure strategies available to players in population X, numbered from 1 to n_X, and let A_Y be the set of n_Y actions available to players in population Y, numbered from 1 to n_Y. A population state is defined by a vector $(x|y) \in \mathbb{R}^{n_X+n_Y}$ with the components of x (and the components of y) adding up to 1, i.e., $x \in \Delta_{A_X}$ and $y \in \Delta_{A_Y}$. A pure or monomorphic (in each population) state, at which every player in population X uses strategy (number) s_X and every player in population Y uses strategy (number) s_Y, is represented as $(e_{s_X}|e_{s_Y})$.

To keep the notation manageable, we assume that both populations use the same sample size k, inference method G and family of tie-breaking rules β. This can be relaxed without complications.

The mean dynamic for a sampling best response protocol with statistical inference $\text{StatSBR}^{G,k,\beta}$ in two populations is

$$\dot{x}_i = \sum_{z\in\mathbb{N}_0^{n_Y,k}} P_y(z)\left[\sum_{h=1}^{n_X} x_h\,\beta_{hi}\left(\pi_X^G(z)\right)\right] - x_i$$

$$\dot{y}_j = \sum_{z\in\mathbb{N}_0^{n_X,k}} P_x(z)\left[\sum_{h=1}^{n_Y} y_h\,\beta_{hj}\left(\pi_Y^G(z)\right)\right] - y_j$$

for $i \in \{1, ..., n_X\}$ and $j \in \{1, ..., n_Y\}$, where

- For each possible sample vector $z = (z_1, ..., z_{n_Y})$ of k strategies from population Y, $P_y(z) = \binom{k}{z_1,\dots,z_{n_Y}} y_1^{z_1} \dots y_{n_Y}^{z_{n_Y}}$ is the probability of obtaining sample z at state (x, y). $P_x(z)$ is defined equivalently.
- $\pi_X^G(z)$ is the vector of expected payoffs for the strategies in population X corresponding to inference method G when observing sample z (from population Y). $\pi_Y^G(z)$ is defined equivalently.

At a monomorphic rest point $(e_{s_X}|e_{s_Y})$, we can consider a reduced system that eliminates strategies $s_X \in A_X$ and $s_Y \in A_Y$ by using the equations $x_{s_X} = 1 - \sum_{i=1,i\neq s_X}^{n_X} x_i$ and $y_{s_Y} = 1-\sum_{j=1,j\neq s_Y}^{n_Y} y_j$. The coordinates of the rest point in this reduced system correspond to the origin $(\mathbf{0}|\mathbf{0})$.

Let $z^{s_X,i}$ be a sample from population X where strategy s_X has been observed exactly $(k-1)$ times (and some other strategy $i \neq s_X$ has been observed exactly once), and define $z^{s_Y,j}$ equivalently. For $i \in \{1, ..., n_X\} \setminus \{s_X\}$ and $j \in \{1, ..., n_Y\} \setminus \{s_Y\}$, the Jacobian of the reduced system at the origin has the structure of the matrix shown in (13), which presents the linearization of the reduced system at the origin for the specific case $n_X = 3$, $n_Y = 4$, $s_X = 1$ and $s_Y = 2$ (the structure is easily generalized to other cases), and where $\alpha_{ij}^X \equiv k\beta_{s_X i}\left(\pi_X^G(z^{s_Y,j})\right)$ and $\alpha_{ji}^Y \equiv k\beta_{s_Y j}\left(\pi_Y^G(z^{s_X,i})\right)$.

$$\begin{pmatrix} \dot{x}_2 \\ \dot{x}_3 \\ \hline \dot{y}_1 \\ \dot{y}_3 \\ \dot{y}_4 \end{pmatrix} = \left(\begin{array}{cc|ccc} -1 & 0 & \alpha_{21}^X & \alpha_{23}^X & \alpha_{24}^X \\ 0 & -1 & \alpha_{31}^X & \alpha_{33}^X & \alpha_{34}^X \\ \hline \alpha_{12}^Y & \alpha_{13}^Y & -1 & 0 & 0 \\ \alpha_{32}^Y & \alpha_{33}^Y & 0 & -1 & 0 \\ \alpha_{42}^Y & \alpha_{43}^Y & 0 & 0 & -1 \end{array}\right) \begin{pmatrix} x_2 \\ x_3 \\ \hline y_1 \\ y_3 \\ y_4 \end{pmatrix} \tag{13}$$

Definition 2. *Consider a two-population $\text{StatSBR}^{G,k,\beta}$ dynamic and a rest point $(e_{s_X}|e_{s_Y})$. Let the set of strategies $J = J_X \cup J_Y$ contain a subset of strategies J_X for population X (excluding s_X) and a subset of strategies J_Y for population Y (excluding s_Y). A strategy $i \in J_X$ is:*

- *stabilizing in J if either* $\sum_{j\in J_Y} \alpha_{ij}^X = 0$ *or* $\sum_{j\in J_Y} \alpha_{ji}^Y = 0$.
- *Potentially stabilizing in J if* $\sum_{j\in J_Y} \alpha_{ji}^Y \leq 1$.
- *Potentially stabilized in J if* $\sum_{j\in J_Y} \alpha_{ij}^X \leq 1$.

A strategy $j \in J_Y$ *is:*

- *stabilizing in J if either* $\sum_{i\in J_X} \alpha_{ji}^Y = 0$ *or* $\sum_{j\in J_X} \alpha_{ij}^X = 0$.
- *Potentially stabilizing in J if* $\sum_{i\in J_X} \alpha_{ij}^X \leq 1$.
- *Potentially stabilized in J if* $\sum_{i\in J_X} \alpha_{ji}^Y \leq 1$.

Proposition A.5. *Let* $(e_{s_X}|e_{s_Y})$ *be a rest point under a two-population StatSBR*$^{G,k,\beta}$ *dynamic, and let* $J = (A_X \setminus \{s_X\} \cup A_Y \setminus \{s_Y\})$.

- *If no strategy survives the iterated elimination of stabilizing strategies in J, then the rest point* $(e_{s_X}|e_{s_Y})$ *is asymptotically stable.*
- *If no strategy is potentially stabilizing in J, then the rest point* $(e_{s_X}|e_{s_Y})$ *is repelling.*
- *If some strategy survives the iterated elimination of potentially stabilizing strategies in J, then the rest point* $(e_{s_X}|e_{s_Y})$ *is unstable.*

The last two results hold if we substitute stabilized for stabilizing.

The proof is a direct adaptation of the proofs for the single-population case.

References

Andersson, O., Argenton, C., & Weibull, J. W. (2014). Robustness to strategic uncertainty. *Games and Economic Behavior*, *85*(1), 272–288. doi: 10.1016/J.GEB.2014.01.018

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

Arigapudi, S., Heller, Y., & Schreiber, A. (in press). Heterogeneous noise and stable miscoordination. *American Economic Journal: Microeconomics*. doi: 10.1257/mic.20240109

Benaïm, M., & Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, *71*(3), 873–903. doi: 10.1111/1468-0262.00429

Brandts, J., & Cooper, D. J. (2006). A change would do you good ... An experimental study on how to overcome coordination failure in organizations. *American Economic Review*, *96*(3), 669–693. doi: 10.1257/AER.96.3.669

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton: Princeton University Press.

Crawford, V. P. (1991). An "evolutionary" interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination. *Games and Economic Behavior*, *3*(1), 25–59. doi: 10.1016/0899-8256(91)90004-X

Engelmann, D., & Normann, H. T. (2010). Maximum effort in the minimum-effort game. *Experimental Economics*, *13*(3), 249–259. doi: 10.1007/S10683-010-9239-3

Feri, F., Gantner, A., Moffatt, P. G., & Erharter, D. (2022). Leading to efficient coordination: Individual traits, beliefs and choices in the minimum effort game. *Games and Economic Behavior*, *136*, 403–427. doi: 10.1016/J.GEB.2022.10.003

Goeree, J. K., & Holt, C. A. (2005). An experimental study of costly coordination. *Games and Economic Behavior*, *51*(2 SPEC. ISS.), 349–364. doi: 10.1016/j.geb.2004.08.006

Hirsch, M. W., & Smale, S. (1974). *Differential equations, dynamical systems, and linear algebra*. San Diego: Academic Press.

Horn, R. A., & Johnson, C. R. (1985). *Matrix Analysis*. Cambridge: Cambridge University Press.

Izquierdo, L. R., Izquierdo, S. S., & Sandholm, W. H. (2018). EvoDyn-3s: A Mathematica computable document to analyze evolutionary dynamics in 3-strategy games. *SoftwareX*, *7*, 226–233. doi: 10.1016/j.softx.2018.07.006

Izquierdo, S. S., & Izquierdo, L. R. (2022). Stability of strict equilibria in best experienced payoff dynamics: Simple formulas and applications. *Journal of Economic Theory*, *206*, 105553. doi: 10.1016/J.JET.2022.105553

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

Monderer, D., & Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, *14*(1), 124-143. doi: 10.1006/game.1996.0044

Morris, S., Rob, R., & Shin, H. S. (1995). p-Dominance and Belief Potential. *Econometrica*, *63*(1), 145. doi: 10.2307/2951700

Osborne, M., & Rubinstein, A. (1998). Games with Procedurally Rational Players. *American Economic Review*, *88*(4), 834–847.

Osborne, M., & Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, *45*(2), 434–441. doi: 10.1016/S0899-8256(03)00147-7

Oyama, D., Sandholm, W. H., & Tercieux, O. (2015). Sampling best response dynamics and deterministic equilibrium selection. *Theoretical Economics*, *10*(1), 243–281. doi: 10.3982/TE1405

Perko, L. (2001). *Differential Equations and Dynamical Systems* (Third ed.). New York: Springer.

Salant, Y., & Cherry, J. (2020). Statistical Inference in Games. *Econometrica*, *88*(4), 1725–1752. doi: 10.3982/ECTA17105

Sandholm, W. H. (2001). Almost global convergence to p-dominant equilibrium. *International Journal of Game Theory*, *30*(1), 107–116. doi: 10.1007/S001820100067

Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. The MIT Press.

Sandholm, W. H., Izquierdo, S. S., & Izquierdo, L. R. (2020). Stability for best

experienced payoff dynamics. *Journal of Economic Theory*, *185*, 104957. doi: 10.1016/J.JET.2019.104957
Sawa, R., & Wu, J. (2023). Statistical inference in evolutionary dynamics. *Games and Economic Behavior*, *137*, 294–316. doi: 10.1016/J.GEB.2022.11.008
Sethi, R. (2000). Stability of Equilibria in Games with Procedurally Rational Players. *Games and Economic Behavior*, *32*(1), 85–104. doi: 10.1006/GAME.1999.0753
Sethi, R. (2021, oct). Stable sampling in repeated games. *Journal of Economic Theory*, *197*. doi: 10.1016/j.jet.2021.105343
Van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, *80*(1), 234–248.
Weibull, J. W. (1995). *Evolutionary game theory*. Cambridge: MIT Press.