

Sampling Best Response Dynamics and Deterministic Equilibrium Selection*

Daisuke Oyama[†], William H. Sandholm[‡] and Olivier Tercieux[§]

March 28, 2014

Abstract

We consider a model of evolution in games in which a revising agent observes the actions of a random number of randomly sampled opponents and then chooses a best response to the distribution of actions in the sample. We provide a condition on the distribution of sample sizes under which an iterated p -dominant equilibrium is almost globally asymptotically stable under these dynamics. We show under an additional condition on the sample size distribution that in supermodular games, an almost globally asymptotically stable state must be an iterated p -dominant equilibrium. Since our selection results are for deterministic dynamics, any selected equilibrium is reached quickly; the long waiting times associated with equilibrium selection in stochastic stability models are absent.

1. Introduction

Evolutionary game theory concerns the dynamics of behavior in large populations of strategically interacting agents whose decisions are based on simple myopic rules. Most analyses in this field follow one of two general approaches. One approach, with origins in the work of Foster and Young (1990), Kandori et al. (1993), and Young (1993), focuses on the very long run behavior of stochastic evolutionary processes when players' choice rules include the possibility of rare "mistakes". These stochastic stability models

*We thank Drew Fudenberg, Mario Pagliero, Satoru Takahashi, Jörgen Weibull, Dai Zusai, Martin Osborne, three anonymous referees, many seminar audiences, and especially Sergiu Hart and Josef Hofbauer for helpful discussions. Financial support from JSPS Grant-in-Aid for Young Scientists (B), the Seimeikai Foundation, and NSF Grants SES-0851580 and SES-1155135 is gratefully acknowledged.

[†]Faculty of Economics, University of Tokyo, Hongo, Bunkyo-ku, Tokyo 113-0033, Japan. e-mail: oyama@e.u-tokyo.ac.jp; website: www.oyama.e.u-tokyo.ac.jp.

[‡]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu; website: www.ssc.wisc.edu/~whs.

[§]Paris School of Economics and CNRS, 48 Boulevard Jourdan, 75014 Paris, France. e-mail: tercieux@pse.ens.fr; website: www.pse.ens.fr/tercieux.

are prized for their ability to provide unique predictions of play in games with multiple strict equilibria. At the same time, these equilibrium selection results are subject to criticism. Since stochastic stability analyses capture a process's infinite horizon behavior, the predictions they provide are not appropriate in all applications. If the population size is large and "mistakes" are rare, the equilibrium that the population reaches first is likely to remain in place for a very long time, regardless of whether it is stochastically stable.¹ This qualification constrains the set of environments in which the predictions of stochastic stability theory are relevant.

An earlier approach to evolutionary game theory, building on models of biological natural selection (Taylor and Jonker (1978)), focuses on deterministic dynamics, modeling the evolution of aggregate behavior using ordinary differential equations and related methods. When the population of interest is large, these deterministic models accurately approximate the behavior of the underlying stochastic model over medium to long time spans. Deterministic dynamics provide clear predictions about the likely course of play starting from any given initial condition, dismiss certain equilibria as unstable, and capture the possibility of stable cycles or more complex recurrent behavior.

A central aim of the literature on deterministic game dynamics has been to provide an interpretation of Nash equilibrium as describing stationary aggregate behavior in populations of agents employing simple decision rules. Attention has therefore focused on dynamics exhibiting strong agreement between stationary states and Nash equilibrium. These dynamics regard each strict equilibrium as locally stable, attracting all solutions starting in a surrounding region.² Thus most deterministic dynamics studied to date have not provided equilibrium selection results.³

A stochastic evolutionary model of Sandholm (2001) provides a starting point for obtaining such results. In this model, a revising agent in a finite-population game observes the actions of $k \geq 2$ randomly sampled opponents. The agent views the empirical distribution of actions in his sample as an estimate of the distribution of actions in the population, and chooses an action that is optimal against this empirical distribution. Using birth-death chain methods, Sandholm (2001) shows that if a game has a $\frac{1}{k}$ -dominant action⁴ that

¹This waiting time critique is formalized by Ellison (1993, 2000) and Binmore et al. (1995), among others. Settings where this critique holds less force are discussed below.

²More generally, any regular evolutionarily stable state (ESS) (Maynard Smith and Price (1973), Taylor and Jonker (1978)) is locally stable under the replicator dynamic and other imitative dynamics, the best response dynamic, the BNN dynamic and other excess payoff dynamics, and the Smith dynamic and other pairwise comparison dynamics. See Taylor and Jonker (1978), Cressman (1997), Hofbauer (2000), and Sandholm (2010a).

³Exceptions include Sethi (2000), discussed in Section 1.2, as well as Hofbauer (1999) and Matsui (1991).

⁴In a single-population game, a pure action is *p-dominant* (Morris et al. (1995)) if it is the unique best response whenever it is played by at least fraction p of the population.

is initially played by a positive fraction of the population, and if the population size is sufficiently large, then all agents ultimately play this action with probability close to one.

Sandholm (2001) studies a stochastic process describing the evolution of aggregate behavior in a finite population. Taking the limit of this process as the population size grows large causes its stochastic component to be averaged away.⁵ The resulting deterministic process, which we call the *k-sampling best response dynamic*, is an inexact-information analogue of the best response dynamic of Gilboa and Matsui (1991). While under the latter dynamic all strict equilibria are locally stable, a simple observation from Sandholm (2001) implies that under the *k-sampling best response dynamic*, any $\frac{1}{k}$ -dominant equilibrium attracts solutions from almost all initial conditions.

In this paper, we develop a number of new results for sampling best response dynamics in single-population games. First, we introduce *λ -sampling best response dynamics*, under which the size of an agent's sample is a random draw from a distribution λ on the positive integers. Theorem 3.1 shows that if this distribution satisfies *k-goodness*, a lower bound on the probabilities of sample sizes of k or less, then any $\frac{1}{k}$ -dominant equilibrium is almost globally asymptotically stable: solution trajectories starting from almost all population states converge to the equilibrium, and solutions that start close enough to the equilibrium always remain close by. The *k-goodness* condition is surprisingly weak, requiring only moderate probabilities of small sample sizes even for low values of k . In other words, equilibrium selection only requires that there be some core probability of agents being poorly informed. Even if the complementary probability is assigned to agents having very large samples, almost global stability of the $\frac{1}{k}$ -dominant equilibrium is still assured.

The heart of the paper extends this selection result to solution concepts based on iterated application of *p*-dominance. The kernel of these concepts is the notion of a *p-best response set*: by definition, when actions in such a set are played by at least fraction p of the population, only actions in the set are best responses. Iterated application of this idea leads to the notion of iterated *p*-best response sets and, when one action remains, of iterated *p*-dominant equilibrium (Tercieux (2006), Oyama and Tercieux (2009)). Our main result, Theorem 4.2, shows that under the *k-goodness* condition described above, any iterated $\frac{1}{k}$ -best response set or iterated $\frac{1}{k}$ -dominant equilibrium is almost globally asymptotically stable.

A key argument used in the proof of this result is a transitivity theorem for nested asymptotically stable sets of deterministic dynamics (cf. Conley (1978)). Ignoring technicalities, the transitivity theorem says that if $C \subseteq B \subseteq A$ are sets of population states, if B is asymptotically stable as a subset of A , and if C is asymptotically stable as a subset

⁵See Benaïm and Weibull (2003), Roth and Sandholm (2013), and Section 2.2 below.

of B , then C is asymptotically stable as a subset of A . While we use this theorem in the context of iterated p -dominance, various other traditional solution concepts for games are also based on the iterated application of a basic principle. If this principle is respected by a game dynamic, then as our analysis illustrates, one can use the transitivity theorem to transfer the logic of the iteration from the analysis of the underlying game to the analysis of the dynamic itself. We are hopeful that this approach will find application in future work on game dynamics.

Finally, we turn our attention to the behavior of sampling best response dynamics in supermodular games. Applying a comparison theorem from the theory of cooperative differential equations (Walter (1970)), we obtain conditions on the distribution of sample sizes under which iterated p -dominant equilibrium is not only a sufficient condition for almost global asymptotic stability, but also a necessary condition (Theorem 5.1). This analysis reveals the surprisingly tight links among supermodularity, the notion of iterated p -dominance, and the sampling best response dynamics introduced here.

In summary, sampling best response dynamics combine positive features of the two approaches to evolutionary game theory described at the outset. In certain games with multiple strict equilibria, sampling best response dynamics provide unique predictions of play, but without the long waiting times common in the stochastic stability approach. While deterministic equilibrium selection results are not possible for every game, in the games where they are available they are relevant quickly.

1.1 Examples

Before introducing our dynamics formally, we present examples that illustrate our results. To do so, we first define our strategic environment: population games played by a single unit-mass population of agents. We denote by $S = \{1, \dots, n\}$ the finite set of actions available to each agent. Then $X = \{x \in \mathbb{R}_+^n : \sum_i x_i = 1\}$, the simplex in \mathbb{R}^n , is the set of *population states*; for each $x \in X$, x_i is the fraction of agents playing action i . (Later, when describing agents' choice rules, we use $y \in X$ to denote a mixed action; in this case, y_i is the probability with which action i is chosen.) For each $i \in S$, e_i denotes the i th standard basis vector in \mathbb{R}^n , here representing the monomorphic population state at which all agents play action i . The continuous function $F: X \rightarrow \mathbb{R}^n$ describes a game's payoffs, with $F_i(x)$ being the payoff obtained at state x by agents playing action i . When the set of actions S is understood, we can identify a population game with its payoff function F .

The simplest population games are obtained by matching a population of agents to play a symmetric two-player normal form game $A \in \mathbb{R}^{n \times n}$, where A_{ij} is the payoff obtained by an agent playing action i when matched against an opponent playing action j . The

induced population game F is defined by the linear function $F_i(x) = \sum_{j \in S} A_{ij}x_j = (Ax)_i$, or equivalently, by $F(x) = Ax$. The payoff $F_i(x)$ can represent either the expected payoff to an agent playing action i from a random match, or the realized payoff of an agent playing action i if all pairs of agents are matched.

Population games of this linear form are the easiest to define, and we use them in the examples to follow. But all of our results apply to population games in which payoffs may be nonlinear functions of the population state, including congestion games (Beckmann et al. (1956)) and other models of multilateral externalities. See Sandholm (2010b) for accounts of these and other examples.

The presentation of the dynamics in the following two examples is informal, with precise definitions deferred until subsequent sections. A third example illustrating our results for supermodular games can be found in Section 5.1.

Example 1.1. Consider the population game

$$F(x) = Ax = \begin{pmatrix} 4 & 0 \\ 3 & 3 + \varepsilon \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix},$$

where $\varepsilon > 0$. While both actions in this game correspond to strict equilibria, action 2 is $\frac{1}{4}$ -dominant: if at least fraction $\frac{1}{4}$ of the population plays action 2, then the expected payoff to this action is at least $3 + \frac{\varepsilon}{4}$, while the expected payoff to action 1 is at most 3.

Suppose that the current population state is x , and that a randomly chosen agent is given the opportunity to revise his action. During this opportunity, the agent takes a sample of size 4 from the population; he then plays a best response to the empirical distribution of the actions chosen by those in his sample.

If the population size is very large, then the evolution of the population state will be determined by its expected motion. Thus, following Sandholm (2001), we compute the expected effect of the next revision opportunity on the population state. It is simpler to focus on the fraction of agents playing action 1. Because this revision opportunity is randomly assigned, it goes to such an agent with probability x_1 . Since action 2 is $\frac{1}{4}$ -dominant, a revising agent will play action 1 only if all four members of his sample play action 1; such a sample is drawn with probability $(x_1)^4$. Since $(x_1)^4 < x_1$ whenever $x_1 \in (0, 1)$, the expected effect of the next revision opportunity is to reduce the fraction of action 1 players. One can therefore show that if the fraction of agents playing action 2 is initially greater than zero, then it eventually converges to one.

While this computation relies on revising agents always viewing samples of size 4, the conclusion does not. We show in Section 3.3 that as long as a sample size of 4 has a probability of at least .37, then convergence is assured under any distribution of the

remaining sample sizes. Alternatively, if the sample size follows a positive-truncated Poisson(μ) distribution, convergence to the $\frac{1}{4}$ -dominant equilibrium is guaranteed whenever $\mu < 4.97$. Thus, equilibrium selection is assured so long as small sample sizes are not too uncommon. ♦

Example 1.2. Consider the following three-action coordination game of Young (1993):

$$(1.1) \quad F(x) = Ax = \begin{pmatrix} 6 & 0 & 0 \\ 5 & 7 & 5 \\ 0 & 5 & 8 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}.$$

Figure 1(i) presents a phase diagram for this game under the best response dynamic of Gilboa and Matsui (1991), which describes the evolution of play when revising agents always switch to a current best response. The phase diagram divides the set of population states into the best response regions for the three actions. Evidently, each pure equilibrium has a nonnegligible basin of attraction under the best response dynamic.

Now consider the evolution of play if agents play best responses to samples of size 2 (or more generally, to samples whose sizes are drawn from a 2-good distribution; see condition (3.1)). Again examining Figure 1(i), we see that there is no action that is $\frac{1}{2}$ -dominant: while action 2 has the biggest best response region, it is not optimal when the population is split equally between actions 2 and 3.

Nevertheless, a selection result can still be obtained. To start, notice that if at least half of the population plays action 2 or 3, then only these actions can be best responses. In other words, $\{2, 3\}$ is a $\frac{1}{2}$ -best response set. Thus by the logic of Example 1.1, the set $X_{\{2,3\}} \subseteq X$ of states at which only actions 2 and 3 are used is asymptotically stable under the 2-sampling best response dynamic, attracting solutions from every initial condition in the interior of X .

Moreover, in the restricted game in which only actions 2 and 3 are available, action 3 is $\frac{1}{2}$ -dominant. Here the logic of Example 1.1 implies that from any state in the relative interior of $X_{\{2,3\}}$, the 2-sampling best response dynamic converges to state e_3 .

Thus the set $X_{\{2,3\}}$ is an almost global attractor with respect to the full state space X , and the state e_3 is an almost global attractor with respect to $X_{\{2,3\}}$. Arguing from these premises, we would like to conclude that under the 2-sampling best response dynamic, state e_3 is an almost global attractor with respect to X . This conclusion is established in Theorem 4.2, the paper's main result. The analysis relies on precise estimates of the behavior of solution trajectories near faces of the simplex corresponding to actions eliminated during a given iteration (Section 4.3.3), and on the transitivity theorem described in the introduction

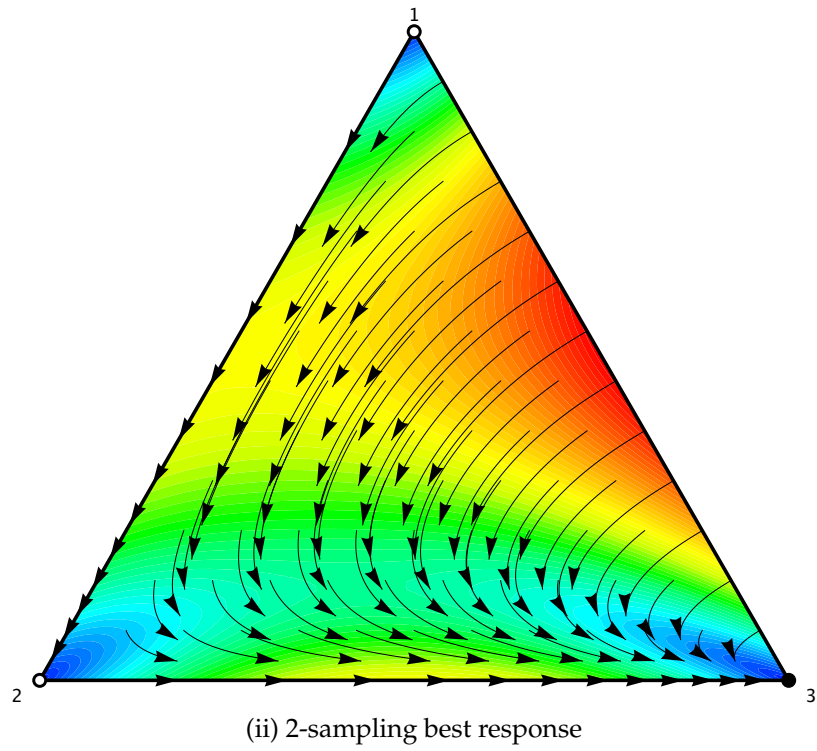
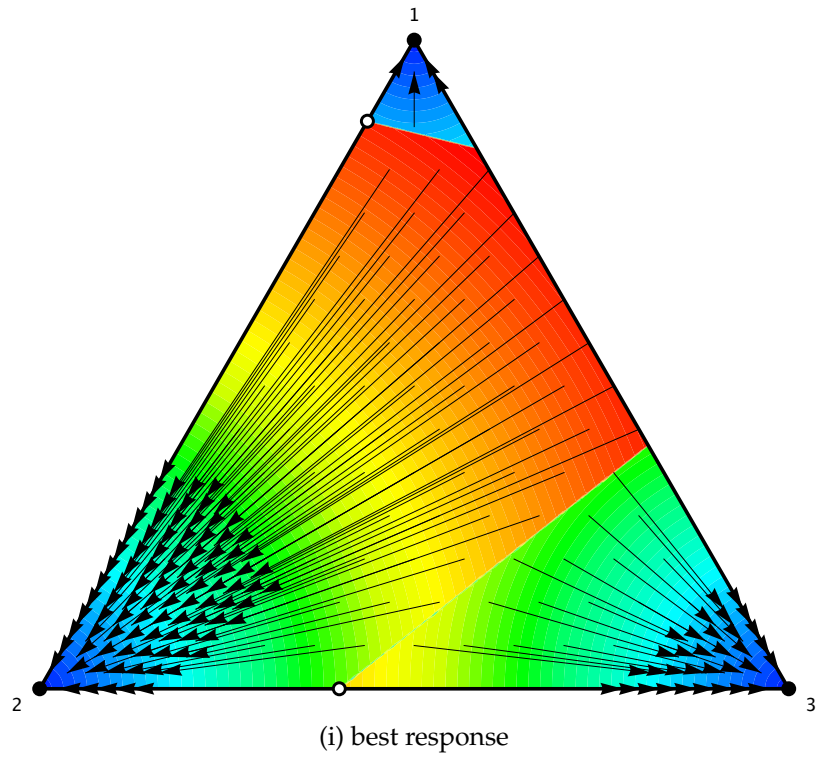


Figure 1: Phase diagrams for dynamics in Young's game. Colors represent speeds: red is fast, blue is slow.

(Theorem 4.7). Figure 1(ii) confirms this stability result.⁶ ♦

1.2 Related literature

A basic goal of this paper is to obtain evolutionary equilibrium selection results that are meaningful over time spans appropriate for typical economic applications. Our work is thus in the spirit of Ellison (1993, 2000), who uses stochastic stability techniques to obtain fast equilibrium selection results for settings in which agents only interact with neighbors. In recent work that complements ours, Kreindler and Young (2013) study waiting times in certain stochastic, two-action evolutionary models with global interaction. Using techniques from stochastic approximation theory, they show that when the deterministic mean dynamic selects among strict equilibria, as in the environments considered here, the expected time until that equilibrium is reached is bounded independently of the population size, and can be estimated as a function of the noise level and of the payoff parameters of the underlying game. They thus provide global-interaction analogues of Ellison’s results.

That agents optimize against samples is a basic feature of Young’s (1993, 1998) influential models of stochastic learning in games. In these models, a group of N agents is repeatedly matched to play an N -player normal form game. In each period, players sample a subset of the realized action profiles from the recent history of play. In the baseline model, players always choose a best response to the distribution of behavior in their samples. Young (1993) shows that this learning procedure always leads to equilibrium play in weakly acyclic games, while Young (1998) (see also Hurkens (1995)) proves that in generic games, play converges to a minimal curb set (Basu and Weibull (1991)). If players are assumed to experiment occasionally with suboptimal actions, then stochastic stability analysis yields a selection result: over the infinite horizon, most periods of play are spent within a single minimal curb set regardless of the players’ initial behavior.

While both Young’s model and ours are based on optimization against random samples of play, the models differ in other important respects. First, while in Young’s model, samples are taken from the action profiles played in the recent past, here the samples are of the current choices of the members of a large population.⁷ Second, while Young’s analyses sometimes rely on unusual (though inevitable) sequences of samples that lead

⁶It is worth noting that the the stochastic evolutionary dynamics of Kandori et al. (1993) and Young (1993) do not respect iterated p -dominance in this game, as it can be shown that action 2 is stochastically stable.

⁷Ellison et al. (2009) study a related evolutionary model in which revising agents face a variable number of random matches between revision opportunities. Their analysis provides conditions under which the effect of matching noise on infinite horizon behavior becomes negligible as the number of matches per period approaches infinity and the probability of “mistakes” approaches zero.

the population into irreducible sets of outcomes, our deterministic analysis focuses on the most likely course of play. Finally, Young’s equilibrium selection results are driven by the relative probabilities of rare transitions between equilibria, and so are subject to the waiting time critique mentioned at the onset. While our analysis based on deterministic dynamics does not always yield unique predictions, these predictions are relevant quickly in the cases where they obtain.

Sethi (2000) introduces action sampling dynamics for recurrent play of normal form games, providing a dynamic foundation for Osborne and Rubinstein’s (1998) notion of $S(1)$ equilibrium.⁸ Under these dynamics, a revising agent plays each available action once against a randomly drawn opponent, and switches to the action that yielded the highest realized payoff. While also being based on sampling, these dynamics differ from those studied here in several important respects: for instance, they are designed for a normal form context rather than for population games, and they only use ordinal properties of payoffs, while in our dynamics, the agents know their payoff functions. Action sampling dynamics therefore have quite different properties than the dynamics studied here. For instance, Sethi (2000) provides an example in which the only stable rest point of his dynamic has a large fraction of the population choosing strictly dominated actions. Under the dynamics studied here, agents only switch to actions that are optimal at some population state; thus strictly dominated actions are eliminated.

2. Best Response and Sampling Best Response Dynamics

2.1 The best response dynamic

Evolutionary game theory considers models of behavior dynamics for large populations of strategically interacting agents. The *best response dynamic* (Gilboa and Matsui (1991), Hofbauer (1995)) provides a natural point of connection between evolutionary and traditional game theoretic modeling, describing the aggregate behavior of populations of myopic but perfectly informed and optimizing agents.

To introduce this dynamic, we denote the pure and mixed best response correspondences for the population game F by $b: X \Rightarrow S$ and $B: X \Rightarrow X$, which are defined by

$$b(x) = \operatorname{argmax}_{i \in S} F_i(x) \text{ and}$$

⁸Osborne and Rubinstein (2003) introduce a notion of equilibrium for contexts in which information about opponents’ play takes the form of random samples. For experimental analyses of various notions of sampling equilibrium, see Selten and Chmura (2008).

$$B(x) = \operatorname{argmax}_{y \in X} y'F(x) = \{y \in X: \operatorname{support}(y) \subseteq b(x)\}.$$

State x is a *Nash equilibrium* of F if all agents are playing a pure action that is optimal given the others' behavior. Thus x is a Nash equilibrium if $x_i > 0$ implies that $i \in b(x)$, or equivalently, if $x \in B(x)$.

The best response dynamic is defined by the differential inclusion⁹

$$(B) \quad \dot{x} \in B(x) - x.$$

To interpret this dynamic, suppose that members of the population receive revision opportunities according to independent, rate 1 Poisson processes, and that an agent who receives a revision opportunity switches to a best response to the current population state. Then at the aggregate level, the rate of “outflow” of agents from action i is governed by the mass x_i of action i players, while only actions that are currently optimal receive “inflows” of agents. Together, these two forces yield expression (B).

While quite appealing from a mathematical point of view, the best response dynamic is founded on the implicit assumption that agents either can perfectly observe the current population state, or can observe each action's exact payoff directly. As such precise information about the strategic environment seems exceptional in many environments in which evolutionary models are appropriate, it is natural to consider alternate models that impose milder informational requirements.

2.2 k -sampling best response dynamics

We now introduce analogues of the best response dynamic that describe behavior when agents' information about their opponents' choices is based on finite and possibly small samples.

To introduce these dynamics, we suppose again that agents receive revision opportunities according to independent, rate 1 Poisson processes. Now when an agent receives a revision opportunity, he observes the actions chosen by k randomly sampled opponents, and plays a best response to the distribution of actions in his sample.

To describe the consequences of this revision procedure formally, we let $\mathbb{Z}_+^{n,k} = \{z \in \mathbb{Z}_+^n : \sum_{i \in S} z_i = k\}$ be the set of possible outcomes of samples of size k . We then define the

⁹Background on differential inclusions is provided in Appendix A.1.

k -sampling best response correspondence $B^k: X \Rightarrow X$ as follows: $y \in B^k(x)$ if and only if

$$y = \sum_{z \in \mathbb{Z}_+^{n,k}} \binom{k}{z_1 \dots z_n} (x_1^{z_1} \dots x_n^{z_n}) \alpha(z),$$

where for each $z \in \mathbb{Z}_+^{n,k}$, $\alpha(z)$ is some element of $B(\frac{1}{k}z)$, and where we adopt the convention that $0^0 = 1$. This correspondence can be described more concisely as

$$B^k(x) = \sum_{z \in \mathbb{Z}_+^{n,k}} \binom{k}{z_1 \dots z_n} (x_1^{z_1} \dots x_n^{z_n}) B\left(\frac{1}{k}z\right).$$

If F is k -generic, meaning that each feasible sample $z \in \mathbb{Z}_+^{n,k}$ induces a unique best response, then B^k is single-valued, and can be expressed as

$$(2.1) \quad B_i^k(x) = \sum_{z \in \mathbb{Z}_+^{n,k} : b(\frac{1}{k}z) = \{i\}} \binom{k}{z_1 \dots z_n} x_1^{z_1} \dots x_n^{z_n}.$$

Evidently, $B_i^k(x)$ is the probability at state x of drawing a random sample of size k to which the best response is i . In this generic case, if all agents follow the revision procedure described above, the expected change in the number of agents playing action i during the next dt time units is

$$(2.2) \quad B_i^k(x) dt - x_i dt.$$

The first term is an inflow term, representing the expected number of revising agents who switch to action i , while the second is an outflow term, representing the expected number of revising agents who switch away from action i . Thus (2.2) describes the expected net change in the use of action i . By taking appropriate limits, we obtain the k -sampling best response dynamic:

$$(2.3) \quad \dot{x} = B^k(x) - x.$$

If F is k -generic, the map B^k is a polynomial of order at most k ; thus (2.3) is a smooth differential equation, and so admits a unique solution trajectory from every initial state in X .¹⁰

¹⁰Results from stochastic approximation theory can be used to establish formal links between the behavior of the stochastic processes described above and the corresponding deterministic dynamics. Benaïm (1998) and Benaïm and Weibull (2003) develop this theory for cases like (2.3) in which the deterministic system

In nongeneric cases, there may be samples that admit multiple best responses. Then the sampling best response set $B^k(x) \subseteq X$ includes every mixed action generated by some assignment of mixed best responses to samples in $\mathbb{Z}_+^{n,k}$. Since in general B^k is a multivalued map, our general formulation of the sampling best response dynamic is as a differential inclusion:

$$(S_k) \quad \dot{x} \in B^k(x) - x.$$

Because B^k is compact- and convex-valued and upper semicontinuous, the dynamic (S_k) admits at least one solution trajectory from every initial condition in X , although these solution trajectories need not be unique. In Appendix A.1, we present some facts about solutions to differential inclusions that are needed for our analysis.

2.3 λ -sampling best response dynamics

While the k -sampling best response dynamics provide an important baseline case, it seems more realistic to allow agents' samples to vary in size. To introduce this possibility, we let $\lambda = \{\lambda_k\}_{k=1}^\infty$ be a probability distribution on the natural numbers, and consider the following revision protocol: When an agent receives an opportunity to switch actions, he draws a random sample size according to distribution λ . If his realized sample size is k , he then draws a sample of size k from the population and plays a best response to the population state defined by the empirical distribution of his sample.

For a given game F with k -sampling best response correspondences $\{B^k\}_{k=1}^\infty$, this procedure generates deterministic dynamics of the form

$$(S_\lambda) \quad \dot{x} \in \sum_k \lambda_k B^k(x) - x.$$

We call such dynamics *λ -sampling best response dynamics*. If F is *λ -generic*, meaning that it is k -generic for every sample size k in the support of λ , and if the support of λ is bounded, then (S_λ) too is a polynomial. For the general case, we show in Appendix A.2 that (S_λ) is again a compact- and convex-valued, upper semicontinuous differential inclusion, so that solutions exist starting from every initial condition in X . This conclusion remains true if we allow a positive probability of an infinite sample size, enabling agents who receive such samples to play exact best responses.

Before beginning the analysis, we note one basic property of sampling best response

is Lipschitz continuous. Roth and Sandholm (2013) extend the theory to allow for upper semicontinuous differential inclusions, including the best response dynamic (B) and the dynamics (S_k) and (S_λ) below.

dynamics.

Proposition 2.1. *Suppose that action i is a strict equilibrium of F : that is, $F_i(e_i) > F_j(e_i)$ for all $j \neq i$. Then under any λ -sampling best response dynamic, the unique solution starting at state e_i is the stationary solution.*

Proposition 2.1 is an immediate consequence of Lemma 4.3 below.

Proposition 2.1 shows one sense in which sampling best response dynamics agree with traditional game-theoretic predictions: strict equilibria correspond to rest points of the dynamic. The main idea behind the proposition is simple. At state e_i , all samples consist entirely of i players. Since i is a strict equilibrium, the unique best response to such a sample is to continue to play action i .¹¹ In this sense, sampling best response dynamics agree with the dynamics commonly considered in the literature. The rest of the paper shows that sampling best response dynamics differ markedly from the standard dynamics in other respects.

3. p -Dominance and Monotone Convergence

3.1 p -dominance

To begin, we recall the following refinement of Nash equilibrium due to Morris et al. (1995).

Definition 3.1. *Action $a^* \in S$ is a p -dominant equilibrium of F if $b(x) = \{a^*\}$ for all $x \in X$ with $x_{a^*} \geq p$.*

Lowering the value of p makes this concept more demanding: a 1-dominant equilibrium is a strict equilibrium, while a 0-dominant equilibrium is a strictly dominant equilibrium.

3.2 k -sampling best response dynamics

To begin our analysis, we argue that if a game has a $\frac{1}{k}$ -dominant equilibrium, where $k \geq 2$, then this equilibrium attracts solutions of the k -sampling best response dynamic from almost all initial conditions. In the remainder of this section, we minimize notation by focusing on the case of a k -generic, two-action coordination game in which action 2 is $\frac{1}{k}$ -dominant, but not $\frac{1}{1+k}$ -dominant. (cf. Example 1.1). The result for arbitrary $\frac{1}{k}$ -dominant equilibria follows almost immediately from the analysis of this case; see Lemma 4.4.

¹¹If we knew that the solution starting from e_i was unique, this argument would be enough to prove the proposition; Lemma 4.3 establishes this uniqueness.

Let $\mathcal{B}_{\ell,q}$ denote a binomial random variable with parameters $\ell \in \mathbb{N}$ and $q \in [0,1]$, defined on an arbitrary probability space. At population state x an agent drawing a sample of size ℓ will choose action 2 with probability

$$B_2^\ell(x) = \mathbb{P}\left(\frac{1}{\ell}\mathcal{B}_{\ell,x_2} \geq \frac{1}{k}\right) = \mathbb{P}\left(\mathcal{B}_{\ell,x_2} \geq \frac{\ell}{k}\right) = \sum_{j=\lceil \ell/k \rceil}^{\ell} \binom{\ell}{j} (x_2)^j (1-x_2)^{\ell-j}.$$

If the sample is of size k , then an agent will choose action 1 only if all k members of his sample choose action 1. He thus chooses action 2 with probability

$$B_2^k(x) = \mathbb{P}(\mathcal{B}_{k,x_2} \geq 1) = 1 - (1-x_2)^k.$$

When $x_2 \in (0,1)$, this probability, representing the inflow into action 2, exceeds x_2 itself, which represents the outflow from action 2:

$$\dot{x}_2 = B_2^k(x) - x_2 = \left(1 - (1-x_2)^k\right) - x_2 = (1-x_2) - (1-x_2)^k > 0.$$

Thus state $x_2 = 1$ attracts all interior initial conditions.

This conclusion still obtains if sample sizes smaller than k are also allowed. If an agent observes a sample of size $\ell \in \{2, \dots, k\}$, it is still the case that he will choose action 1 only if all members of his sample choose this action. Since

$$\mathbb{P}\left(\mathcal{B}_{\ell,x_2} \geq \frac{\ell}{k}\right) = \mathbb{P}(\mathcal{B}_{\ell,x_2} \geq 1) = 1 - (1-x_2)^\ell \text{ when } \ell \in \{2, \dots, k\},$$

smaller sample sizes still lead action 2 to be chosen more quickly than it is abandoned, so that global stability is preserved—see Figure 2.

3.3 λ -sampling best response dynamics

We now turn to λ -sampling best response dynamics. If the distribution λ is such that only sample sizes between 1 and k receive positive probability, with sample size 1 receiving probability less than 1, then the relevant “inflow” is described by an appropriate weighted average of the curves in Figure 2, and the same conclusion obtains.¹²

If agents may obtain samples of size $\ell > k$, then matters become more complicated, as agents with such sample sizes who observe small numbers of action 2 players will still select action 1. (See Figure 3.) Thus, if only samples of size larger than k receive positive

¹²When the sample size is 1, then the inflow to and outflow from action 2 are both x_2 . As long as this sample size is used with probability less than 1, selection of the $\frac{1}{k}$ -dominant equilibrium is still assured.

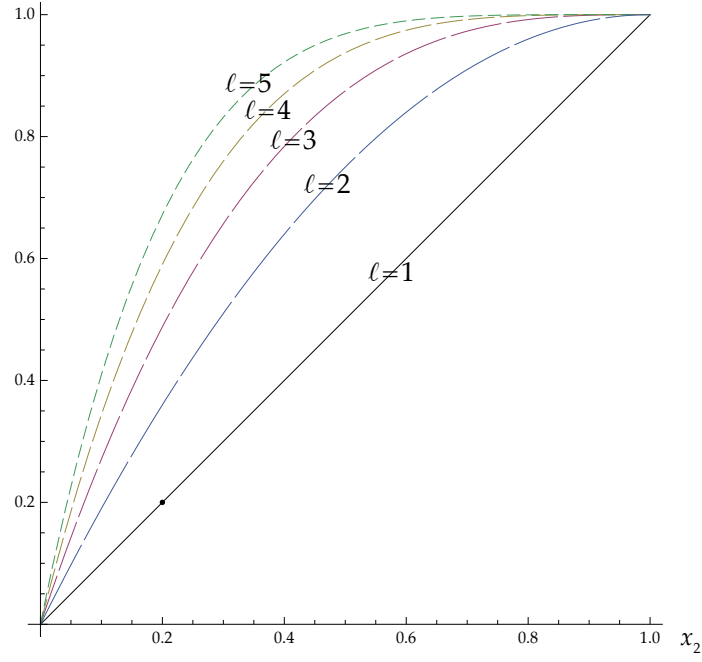


Figure 2: The function $\mathbb{P}(\mathcal{B}_{\ell, x_2} \geq \frac{\ell}{5})$ for $\ell = 1, 2, 3, 4$, and 5.

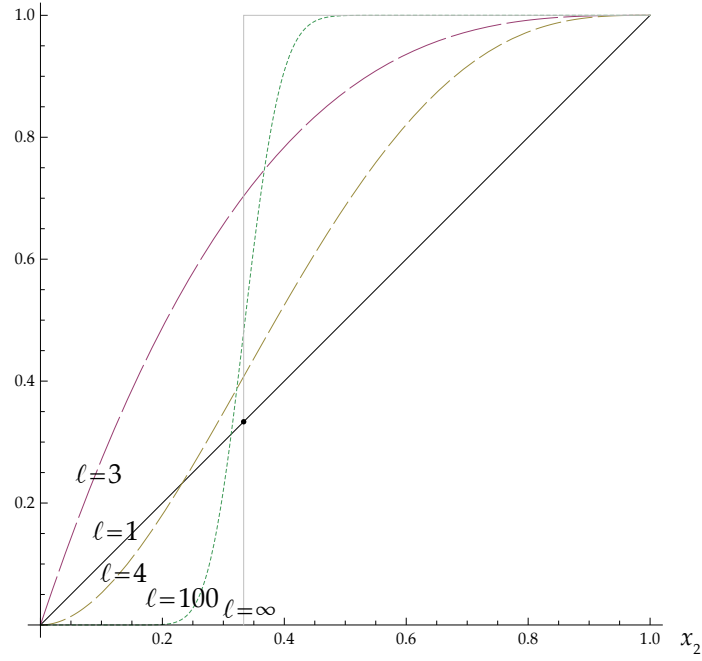


Figure 3: The function $\mathbb{P}(\mathcal{B}_{\ell, x_2} \geq \frac{\ell}{3})$ for $\ell = 1, 3, 4$, and 100, along with the exact best response function for a coordination game with mixed equilibrium $x_2^* = \frac{1}{3}$.

probability, the outflow of action 2 players will exceed the inflow at small positive values of x_2 , so that coordination on action 1 will be locally stable.

What if sample sizes of at most k and larger than k both have positive probability? In this case, it is possible that the “good” behavior under the smaller sample sizes might compensate for the “bad” behavior under the larger ones, so that selection of $\frac{1}{k}$ -dominant equilibrium is preserved.

For this to occur, the λ -weighted average of the functions $\mathbb{P}(\mathcal{B}_{\ell, x_2} \geq \frac{\ell}{k})$ must always lie above the 45° line. Since $\mathbb{P}(\mathcal{B}_{\ell, x_2} \geq \frac{\ell}{k})$ is close to 0 when x_2 is close to 0 for any sample size ℓ , one might expect the prospects for global convergence to be determined by behavior in this region. But since

$$\left. \frac{d}{dx_2} \mathbb{P}(\mathcal{B}_{\ell, x_2} \geq \frac{\ell}{k}) \right|_{x_2=0} = \begin{cases} \ell & \text{if } \ell \leq k, \\ 0 & \text{if } \ell > k, \end{cases}$$

the “good” cases easily compensate for the bad ones here under appropriate conditions on the sample size distribution (again, see Figure 3). Thus, whether $\frac{1}{k}$ -dominant equilibrium is selected depends largely on whether the required compensation occurs at larger values of x_2 . We obtain sufficient conditions for this to occur by relying on some less obvious properties of the function $\mathbb{P}(\mathcal{B}_{\ell, x_2} \geq \frac{\ell}{k})$; see Lemma 3.4 and especially Lemma 3.5.

Theorem 3.1 and the observations that follow show that the total probability that must be assigned to small sample sizes to ensure selection of $\frac{1}{k}$ -dominant equilibrium can be surprisingly small. The theorem is stated for two-action games in which action 2 is $\frac{1}{k}$ -dominant, but the assumption that this action is not $\frac{1}{k+1}$ -dominant is dropped.

Theorem 3.1. *Consider a two-action game in which action 2 is $\frac{1}{k}$ -dominant, where $k \geq 2$. Suppose that the sample size distribution $\lambda = \{\lambda_\ell\}_{\ell \geq 1}$ is k -good:*

$$(3.1) \quad \sum_{\ell=1}^k \lambda_\ell \left(1 - \left(\frac{k-1}{k}\right)^\ell\right) > \frac{1}{k}.$$

Then under the λ -sampling best response dynamic, $\dot{x}_2 > 0$ whenever $x_2 \in (0, 1)$. Thus if play begins at a state at which a positive mass of agents choose action 2, then this mass increases monotonically over time and converges to 1.

The proof of Theorem 3.1 is presented below.

Clearly, if λ places all mass on sample sizes 1 through k and some mass on sample sizes above 1 (that is, if $\lambda_1 + \dots + \lambda_k = 1$ and $\lambda_1 < 1$), then λ is k -good. But to provide a sense of the weakness of the k -goodness condition, we consider three other simple cases.

For the first, we introduce a condition that will play an important role in our analysis of supermodular games in Section 5. We say that distribution λ is *k-minimal* if it only places mass on sample sizes of k or higher: that is, if $\lambda_1 = \dots = \lambda_{k-1} = 0$. Under this condition, all of the mass placed on smaller sample sizes is concentrated on the one that is most favorable for selection of $\frac{1}{k}$ -dominant equilibrium.

Observation 3.2. *A sufficient condition for distribution λ to be k-good is that*

$$\lambda_k > \underline{\lambda}_k \equiv \frac{1}{k \left(1 - \left(\frac{k-1}{k}\right)^k\right)}.$$

If λ is k-minimal this condition is also necessary.

In Table 1, we present some values of the lower bound $\underline{\lambda}_k$ corresponding to small values of k . This lower bound is quite far from 1, even for small values of k , as shown in Table 1. Evidently, $\underline{\lambda}_k \approx 1/(k(1 - e^{-1})) \approx 1.5820/k$ for larger values of k .¹³

To consider a more typical case, we next suppose that all sample sizes between 1 and k are equally likely to be employed.

Observation 3.3. *Suppose that $\lambda_1 = \dots = \lambda_k$, and let $\tau_k = \lambda_1 + \dots + \lambda_k$. Then λ is k-good if and only if*

$$\tau_k > \underline{\tau}_k \equiv \frac{1}{1 + (k-1)\left(\frac{k-1}{k}\right)^k}.$$

Some values of $\underline{\tau}_k$ for small k are presented in Table 1. As before, the probability that must be placed on small sample sizes to ensure equilibrium selection is quite far from 1 even for small values of k . Indeed, allowing all sample sizes up to k does not make the bound obtained here dramatically more demanding than the one from Observation 3.2. This is clear from the table when k is small, and it remains so when k is large: in the latter case, we have that $\underline{\tau}_k \approx 1/(1 + (k-1)e^{-1})$, implying that the ratio of the bounds $\underline{\tau}_k/\underline{\lambda}_k \approx (k(1 - e^{-1}))/((1 + (k-1)e^{-1}))$ is approximately $e - 1 \approx 1.7183$.

Finally, a natural way to model the sample size distribution is to suppose that an agent has a small probability of observing the action of each opponent, with the occurrence of each observation being independent of the others. Since the number of opponents is large, the distribution of sample sizes is well approximated by a Poisson(μ) distribution for some $\mu > 0$.¹⁴ If we assume that only agents with nonempty samples revise, then

¹³This comes from a Poisson approximation of a binomial probability: $\left(\frac{k-1}{k}\right)^k$ is the probability that a sample of size k from a population with $x_2 = \frac{1}{k}$ contains no action 2 players; its limit is the probability that a Poisson(1) random variable equals 0, namely e^{-1} .

¹⁴See, for instance, Durrett (2005, Sec. 2.6).

k	$\underline{\lambda}_k$	$\underline{\tau}_k$
2	.6667	.8000
3	.4736	.6279
4	.3657	.5130
5	.2974	.4328
6	.2506	.3739
10	.1535	.2417

Table 1: Values of the lower bounds $\underline{\lambda}_k$ and $\underline{\tau}_k$ for some small values of k . For large k , $\underline{\lambda}_k \approx 1/(k(1 - e^{-1}))$ and $\underline{\tau}_k \approx 1/(1 + (k - 1)e^{-1})$, implying that $\underline{\tau}_k/\underline{\lambda}_k \approx e - 1$.

k	$\bar{\mu}_k$	$\mathbb{E}\mathcal{N}_{\bar{\mu}_k}$
2	1.109	1.655
3	3.162	3.302
4	4.978	5.012
5	6.647	6.656
6	8.230	8.232
10	14.126	14.126

Table 2: Values of the Poisson parameter upper bounds $\bar{\mu}_k$ and mean sample size upper bounds $\mathbb{E}\mathcal{N}_{\bar{\mu}_k} = \bar{\mu}_k/(1 - e^{-\bar{\mu}_k})$ for some small values of k .

the relevant distribution is $\text{Poisson}(\mu)$, but truncated to exclude realizations of 0; letting \mathcal{N}_μ denote a random variable with this truncated distribution, we have that $\mathbb{P}(\mathcal{N}_\mu = \ell) = (\mu^\ell e^{-\mu})/((1 - e^{-\mu})\ell!)$ and $\mathbb{E}\mathcal{N}_\mu = \mu/(1 - e^{-\mu})$. In Table 2, we present numerical upper bounds $\bar{\mu}_k$ on the values of the parameter μ for which the truncated $\text{Poisson}(\mu)$ distribution is k -good, as well as the corresponding values of the means $\mathbb{E}\mathcal{N}_{\bar{\mu}_k}$. The table suggests that once k is larger than 2, the k -goodness condition is satisfied whenever $\mu \leq k$.

This analysis shows that equilibrium selection is quite robust to the possibility that agents are well informed. As long as a certain core probability is placed on small sample sizes in the sense required by k -goodness, selection of $\frac{1}{k}$ -dominant equilibrium obtains regardless of how well informed the remaining agents are. This robustness is equally true of the more general equilibrium selection results developed in Section 4.

While the k -goodness condition is sufficient for almost global convergence, it is not necessary. By following the graphical procedure described above, it is easy to check whether almost global convergence obtains for any given combination of game and sample size distribution.

Proof of Theorem 3.1. As before, let $\mathcal{B}_{\ell,p}$ denote a binomial random variable with parameters

$\ell \in \mathbb{N}$ and $p \in [0, 1]$. Observe that

$$(3.2) \quad \dot{x}_2 \geq \left(\sum_{j \leq k} \lambda_j \mathbb{P}(\mathcal{B}_{j,x_2} \geq 1) + \sum_{\ell > k} \lambda_\ell \mathbb{P}(\mathcal{B}_{\ell,x_2} \geq \frac{\ell}{k}) \right) - x_2.$$

Now for $j \leq k$,

$$(3.3) \quad \mathbb{P}(\mathcal{B}_{j,x_2} \geq 1) = 1 - (1 - x_2)^j$$

is strictly concave in x_2 , and its fixed points occur at 0 and 1. Moreover, the following properties of $\mathbb{P}(\mathcal{B}_{\ell,q} \geq \frac{\ell}{k})$ when $\ell > k$ are well known (see, e.g., Kaniovski and Young (1995)).

Lemma 3.4. *If $\ell > k$, $\mathbb{P}(\mathcal{B}_{\ell,q} \geq \frac{\ell}{k})$ is convex-concave in q , with fixed points occurring at 0, 1, and exactly one additional point in $[0, 1]$.*

The key to our analysis is the following estimate.

Lemma 3.5. $\mathbb{P}(\mathcal{B}_{\ell, \frac{1}{k}} \geq \frac{\ell}{k}) \geq \frac{1}{k}$ for all ℓ and k .

In words, Lemma 3.5 says that if a coin with bias $\frac{1}{k}$ is tossed any number of times, then the probability that the proportion of heads is at least $\frac{1}{k}$ is itself at least $\frac{1}{k}$.

*Proof.*¹⁵ Let \mathcal{M} be a k -dimensional multinomial random vector with parameters ℓ and $(\frac{1}{k}, \dots, \frac{1}{k})$. Clearly, $\mathbb{P}(\mathcal{B}_{\ell, \frac{1}{k}} \geq \frac{\ell}{k}) = \mathbb{P}(\mathcal{M}_i \geq \frac{\ell}{k})$ for all $i \in \{1, \dots, k\}$. Call this common probability p . Since the sum of the components of \mathcal{M} is always ℓ , there is always at least one component of \mathcal{M} whose realization is at least $\frac{\ell}{k}$. Thus

$$1 = \mathbb{P}\left(\bigcup_{i=1}^k \{\mathcal{M}_i \geq \frac{\ell}{k}\}\right) \leq \sum_{i=1}^k \mathbb{P}(\mathcal{M}_i \geq \frac{\ell}{k}) = kp,$$

and so $p \geq \frac{1}{k}$. ■

To prove the theorem, it is enough to show that $\dot{x}_2 > 0$ whenever $x_2 \in (0, 1)$. For such x_2 , we have that $\mathbb{P}(\mathcal{B}_{1,x_2} \geq 1) = x_2$, and also that $\mathbb{P}(\mathcal{B}_{j,x_2} \geq 1) > x_2$ for all $j \in \{2, \dots, k\}$. Moreover, Lemmas 3.4 and 3.5 together imply that for each $\ell > k$, the interior fixed point of $\mathbb{P}(\mathcal{B}_{\ell,x_2} \geq \frac{\ell}{k})$ is at most $\frac{1}{k}$, and hence that $\mathbb{P}(\mathcal{B}_{\ell,x_2} \geq \frac{\ell}{k}) > x_2$ whenever $x_2 \in (\frac{1}{k}, 1)$ (see Figure 3). Thus, since the k -goodness of λ ensures that $\lambda_1 < 1$, inequality (3.2) implies that $\dot{x}_2 > 0$ whenever $x_2 \in (\frac{1}{k}, 1)$.

¹⁵We thank Sergiu Hart for suggesting the following proof.

To verify that $\dot{x}_2 > 0$ whenever $x_2 \in (0, \frac{1}{k}]$, it follows from expressions (3.2) and (3.3) that it is enough to show that

$$(3.4) \quad \sum_{j \leq k} \lambda_j (1 - (1 - x_2)^j) > x_2.$$

In fact, since the left hand side of (3.4) is strictly concave and equals 0 when $x_2 = 0$, it is enough to check that (3.4) holds at $x_2 = \frac{1}{k}$. This reduces to the k -goodness condition (3.1), so the proof of the theorem is complete. ■

4. Iterated p -Dominance and Almost Global Asymptotic Stability

4.1 p -best response sets and iterated p -dominance

We now argue that the k -goodness criterion is relevant more broadly. To do so, we follow Tercieux (2006) and Oyama and Tercieux (2009) by defining p -best response sets, which are set-valued extensions of p -dominant equilibrium, and their iterated analogues.

We start with p -best response sets. To ease this and future definitions, for any nonempty set $R \subseteq S$ and any $x \in X$, we let $x_R = \sum_{i \in R} x_i$ denote the total mass placed on actions in R at population state x .

Definition 4.1. *A nonempty set of actions $R \subseteq S$ is a p -best response set of F if $b(x) \subseteq R$ holds for all $x \in X$ with $x_R \geq p$.*

In words, R is a p -best response set if whenever a player assigns a probability of at least p to his opponents playing actions in R , all of the player's best responses are themselves in R .

To define iterated p -best response sets and iterated p -dominant equilibria, we require notation for restricted games. Let $R \subseteq S$ be a nonempty set containing $n' \leq n$ actions. We let $X_R = \{x \in X : x_R = 1\}$ denote the set of population states assigning mass 1 to R , and let $F|_R : X_R \rightarrow \mathbb{R}^{n'}$ be the restricted version of F in which agents may only choose actions in R .¹⁶ The restricted best response correspondences $b|_R : X_R \Rightarrow R$ and $B|_R : X_R \Rightarrow X_R$ are defined in the obvious way.

Definition 4.2. *A nonempty set of actions $S^* \subseteq S$ is an iterated p -best response set of F if for some $m \geq 1$ there exists a sequence S^0, S^1, \dots, S^m with $S = S^0 \supset S^1 \supseteq \dots \supseteq S^m = S^*$ such that S^ℓ is a p -best response set in $F|_{S^{\ell-1}}$ for each $\ell = 1, \dots, m$.*

¹⁶There is a minor abuse of notation here, in that the domain of game $F|_R$ is $X_R \subseteq X \subseteq \mathbb{R}^n$ rather than the simplex in $\mathbb{R}^{n'}$.

An action $a^* \in S$ is an iterated p -dominant equilibrium of F if $\{a^*\}$ is an iterated p -best response set of F .

For an iterated p -best response set S^* , the sequence S^0, S^1, \dots, S^m in the definition will be called an *associated sequence* of S^* .

Tercieux (2006) establishes that for each $p \leq \frac{1}{2}$, there exists a smallest p -best response set: that is, a p -best response set that is contained in all other p -best response sets. Proposition 4.1, whose proof can be found in Appendix A.3, extends this result to iterated p -best response sets.

Proposition 4.1. *For any $p \leq \frac{1}{2}$, any game has a smallest iterated p -best response set.*

Of course, any game that is dominance solvable admits an iterated p -dominant equilibrium. For an economic example outside this class, iterated p -dominant equilibrium also exists in a model of technology adoption introduced in Oyama and Tercieux (2009).¹⁷

4.2 The main result

In order to state our main result, we need some definitions concerning stability in dynamical systems, which we now state in the context of solutions to the differential inclusion (S_λ) . Further facts on stability under differential inclusions can be found in Appendix A.1; for a general treatment, see Benaïm et al. (2005).

The differential inclusion (S_λ) induces a set-valued dynamical system $\Phi: [0, \infty) \times X \Rightarrow X$ defined by

$$\Phi_t(\xi) = \{x(t) : x \text{ is a solution to } (S_\lambda) \text{ with } x(0) = \xi\}.$$

For $I \subseteq [0, \infty)$ and $Y \subseteq X$, we write $\Phi_I(\xi) = \bigcup_{t \in I} \Phi_t(\xi)$ and $\Phi_I(Y) = \bigcup_{t \in I} \bigcup_{\xi \in Y} \Phi_t(\xi)$. We call the set $Y \subseteq X$ *strongly positively invariant* if $\Phi_{[0, \infty)}(Y) \subseteq Y$, so that no solution starting in Y ever leaves Y .

Let $\omega(\xi)$ be the ω -limit set of $\xi \in X$: that is, $\omega(\xi) = \bigcap_{t \geq 0} \overline{\Phi_{[t, \infty)}(\xi)}$. The set $Y \subseteq X$ is *attractive* if there exists a neighborhood U of Y such that

$$(4.1) \quad \bigcup_{\xi \in U} \omega(\xi) \subseteq Y.$$

¹⁷Specifically, Propositions 2.15 and the proof of Proposition 2.9 of Oyama and Tercieux (2009) show that when the technology adoption games considered there are supermodular, symmetric, and satisfy a diminishing marginal cost condition, they admit iterated $\frac{1}{2}$ -dominant equilibria.

The set Y is *Lyapunov stable* if for any neighborhood U of Y , there exists a neighborhood V of Y such that $\Phi_{[0,\infty)}(V) \subseteq U$. A set is *asymptotically stable* if it is both attractive and Lyapunov stable.

Finally, we introduce stronger notions of stability for faces of the simplex X . Let $R \subseteq S$ be nonempty, and let $O_R = \{x \in X : x_R > 0\}$ be the set of states that put positive mass on actions in R . If condition (4.1) holds with $Y = X_R$ and $U = O_R$, we call X_R *almost globally attractive*. If in addition X_R is Lyapunov stable, we call X_R *almost globally asymptotically stable*.

Theorem 4.2. *Let S^* be the smallest iterated $\frac{1}{k}$ -best response set of game F , and suppose that λ is k -good. Then the set X_{S^*} is almost globally asymptotically stable under the λ -sampling best response dynamics for F . In particular, if a^* is an iterated $\frac{1}{k}$ -dominant equilibrium, then e_{a^*} is almost globally asymptotically stable.*

For game (1.1) from Example 1.2, Figure 1(i) shows that the set $\{2, 3\}$ is a $\frac{1}{2}$ -best response set, and action 3 is an iterated $\frac{1}{2}$ -dominant equilibrium. Theorem 4.2 thus implies that if λ is 2-good, then equilibrium e_3 is almost globally asymptotically stable under the λ -sampling best response dynamic. Figure 1(ii) confirms this conclusion for the case of the 2-sampling best response dynamic.

It is worth comparing Theorem 4.2 to stability results for the sets $X_{S_0^*}$ and X_D , where S_0^* is the smallest iterated 0-best response set (that is, the set of actions that survive iterated removal of actions that are never a best response), and $D \supseteq S_0^*$ is the set of actions that survive iterated removal of actions that are strictly dominated by pure actions. In the case of the best response dynamic (B), it is well known that $X_{S_0^*}$ is globally attractive,¹⁸ and one can show that $X_{S_0^*}$ is asymptotically stable.¹⁹ In the case of imitative dynamics, Nachbar (1990) and Samuelson and Zhang (1992) show that the set X_D attracts solutions from all interior initial conditions. However, one can construct simple examples in which X_D is not even Lyapunov stable under any imitative dynamics.²⁰ Hofbauer and Sandholm (2011) show that for all dynamics from a large class that excludes these two cases, there are games in which strictly dominated actions survive in perpetuity from most initial conditions.

Oyama and Tercieux (2009) show that iterated p -dominance is respected by the perfect foresight dynamics of Matsui and Matsuyama (1995). By contrast, the stochastic evolutionary dynamics of Kandori et al. (1993) and Young (1993) do not respect iterated p -dominance: in game (1.1) from Example 1.2, action 2 is stochastically stable, whereas action 3 is an iterated $\frac{1}{2}$ -dominant equilibrium.

¹⁸See Sandholm (2010b, Theorem 7.4.2).

¹⁹This is easily deduced from the transitivity theorem below (Section 4.3.2), or from the more elementary Lemma A.2 (Appendix A.1).

²⁰See Sandholm (2013, Example 8.2).

4.3 Analysis

4.3.1 Preliminaries

We start by presenting a few preliminary results. We consider a λ -sampling best response dynamic (S_λ) in which the distribution λ is k -good.

We begin with a basic invariance property that we alluded to in Section 3.3, and that we prove in Appendix A.4.²¹

Lemma 4.3. *If $R \subseteq S$ is a 1-best response set, then X_R is strongly positively invariant.*

Next, we extend the conclusion of Theorem 3.1 from $\frac{1}{k}$ -dominant equilibria to $\frac{1}{k}$ -best response sets.

Lemma 4.4. *Let $R \subseteq S$ be a $\frac{1}{k}$ -best response set. If $x_R \in (0, 1)$, then $\dot{x}_R > 0$.*

The proof of Lemma 4.4 is identical to that Theorem 3.1 if all instances of x_2 and \dot{x}_2 are replaced with x_R and \dot{x}_R , respectively. Lemma 4.4 implies that the continuous function $L(x) = x_R$ serves as a Lyapunov function for X_R on the set $O_R = \{x \in X: x_R > 0\}$, and so establishes that X_R is almost globally asymptotically stable.²²

4.3.2 Asymptotic stability

In this section, we establish that any iterated $\frac{1}{k}$ -best response set is asymptotically stable. Let $S = S^0 \supseteq S^1 \supseteq \dots \supseteq S^m = S^*$ be an associated sequence of an iterated $\frac{1}{k}$ -best response set S^* . The idea of the argument is to develop an analogue of the iterative deletion of actions that are not in the $\frac{1}{k}$ -best-response sets for the context of the dynamical system (S_λ).

The first step in our argument is analogous to a single step of the iterative deletion: we show that the set of states X_{S^ℓ} is asymptotically stable in $X_{S^{\ell-1}}$. For $\ell = 1, \dots, m$, let $\bar{S}^\ell = S^{\ell-1} \setminus S^\ell$ be the set of actions removed during the ℓ th round of removing actions that are not in the $\frac{1}{k}$ -best response sets. Let

$$X_{S^{\ell-1}}^\circ = X_{S^{\ell-1}} \setminus X_{\bar{S}^\ell} = \{x \in X_{S^{\ell-1}}: x_{\bar{S}^\ell} > 0\}$$

be the set of population states that put all of their mass on actions surviving $\ell - 1$ rounds of removal and at least some mass on actions surviving the ℓ th round of removal.

²¹In the course of proving Lemma 4.3, we also show that a 1-best response set is asymptotically stable under the k -sampling best response dynamic for all sufficiently large k .

²²See Benaïm et al. (2005, Proposition 3.25).

Lemma 4.5 shows that if only actions in $S^{\ell-1}$ are played and if some actions in S^ℓ are played, then the mass on actions in S^ℓ increases over time.

Lemma 4.5. *For all $\ell \in \{1, \dots, m\}$, if $x \in X_{S^{\ell-1}}^\circ$ and $x_{S^\ell} < 1$, then $\dot{x}_{S^\ell} > 0$.*

Since S^ℓ is a $\frac{1}{k}$ -best response set in $F|_{S^{\ell-1}}$, this lemma is a direct consequence of Lemma 4.4.

For $Z \subseteq Y \subseteq X$, we say that Z is *asymptotically stable in Y* if Y is a strongly positively invariant set and Z is asymptotically stable with respect to the dynamical system restricted to Y . Lemma 4.5 implies that the continuous function $L(x) = x_{S^\ell}$ is a Lyapunov function for X_{S^ℓ} on $X_{S^{\ell-1}}^\circ$, and so with Lemma 4.3 establishes

Corollary 4.6. *X_{S^ℓ} is asymptotically stable in $X_{S^{\ell-1}}$.*

The proof of asymptotic stability of S^* is completed by an appeal to an intuitive but powerful theorem from the theory of dynamical systems. The theorem establishes a basic transitivity property of nested sequences of asymptotically stable sets. The statement of the theorem requires one further definition. We call $Y \subseteq X$ *strongly negatively invariant* if for any solution $y: [0, \infty) \rightarrow X$ and any $t \geq 0$, we have $y(0) \in Y$ whenever $y(t) \in Y$.

Theorem 4.7 (Transitivity theorem). *Let $C \subseteq B \subseteq A$ be compact sets. If B is strongly negatively invariant and asymptotically stable in A , and if C is asymptotically stable in B , then C is asymptotically stable in A .*

The transitivity theorem is proved by Conley (1978) for single-valued dynamical systems defined at all times $t \in (-\infty, \infty)$. In Appendix A.5, we extend his analysis to set-valued systems, and, more challengingly, to systems that are only defined in forward time.

For any set of actions $R \subseteq S$, the set X_R is strongly negatively invariant, as solutions to (S_λ) with initial conditions in $X \setminus X_R$ cannot reach face X_R in finite time.²³ Therefore, Corollary 4.6 and Theorem 4.7 together yield

Proposition 4.8. *For all $\ell \in \{1, \dots, m\}$, X_{S^ℓ} is asymptotically stable (in X).*

4.3.3 Almost global attraction

It remains to show that X_{S^*} is almost globally attractive, i.e., for any $\xi \in O_{S^*}$, the ω -limit set $\omega(\xi)$ is contained in X_{S^*} . To do so, we show that for each $\ell = 1, \dots, m$, the solution trajectory from ξ never reaches a neighborhood of the set X_{S^ℓ} , the face of X whose vertices correspond to actions eliminated during the ℓ th round of removal of actions that are not in the $\frac{1}{k}$ -best response sets. The proof combines estimates of the behavior of the dynamic near X_{S^ℓ} based on k -goodness with the Lyapunov stability of each set X_{S^ℓ} established in Proposition 4.8. The details of this argument are presented in Appendix A.6.

²³This follows from the fact that $\dot{x}_R(t) \geq -x_R(t)$, which implies that $x_R(t) \geq e^{-t}x_R(0)$; cf. equation (A.3).

5. Supermodular Games

5.1 Characterization of almost global asymptotic stability

In this section we study evolution under sampling best response dynamics in supermodular games. To introduce these games, we define the stochastic dominance matrix $\Sigma \in \mathbb{R}^{(n-1) \times n}$ and the difference matrix $\tilde{\Sigma} \in \mathbb{R}^{n \times (n-1)}$ by

$$\Sigma = \begin{pmatrix} 0 & 1 & \cdots & 1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 \end{pmatrix} \quad \text{and} \quad \tilde{\Sigma} = \begin{pmatrix} -1 & 0 & \cdots & 0 \\ 1 & -1 & \ddots & \vdots \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & -1 \\ 0 & \ddots & 0 & 1 \end{pmatrix}.$$

Observe that $(\Sigma x)_i = \sum_{j=i+1}^n x_j$ equals the total mass on actions greater than i at population state x , so that $\Sigma y \geq \Sigma x$ if and only if y stochastically dominates x .²⁴ Moreover, if we let $X = \Sigma X \subseteq \mathbb{R}^{n-1}$ denote the image of X under Σ , then the map from X to X defined by $\chi \mapsto \tilde{\Sigma}\chi + e_1$ is the inverse of Σ .

We call F a *supermodular game* if

$$(5.1) \quad \Sigma y \geq \Sigma x \text{ implies that } F_{i+1}(y) - F_i(y) \geq F_{i+1}(x) - F_i(x) \text{ for all } i < n.$$

In words: if y stochastically dominates x , then for any action $i < n$, the payoff advantage of $i + 1$ over i is greater at y than at x .

The main result of this section shows that, under certain conditions on the distribution λ , being an iterated $\frac{1}{k}$ -dominant equilibrium is not only sufficient for almost global asymptotic stability under λ -sampling best response dynamic, but also necessary.

To state the theorem, we recall that the distribution λ is k -minimal if only sample sizes of k or higher receive mass under λ , and we let $\mu_\lambda = \sum_k \lambda_k k$ denote its mean.

Theorem 5.1. *Suppose that λ is k -good and k -minimal, that μ_λ is finite, and that F is supermodular.*

- (i) *Let S^* be the smallest iterated $\frac{1}{k}$ -best response set of F . Suppose that F is λ -generic. Then states $e_{\min S^*}$ and $e_{\max S^*}$ are asymptotically stable under the λ -sampling best response dynamic.*
- (ii) *State e_{a^*} is almost globally asymptotically stable under the λ -sampling best response dynamic if and only if a^* is an iterated $\frac{1}{k}$ -dominant equilibrium.*

²⁴Here $y \geq x$ means that $y_i \geq x_i$ for all $i \in S$, and $y \gg x$ means that $y_i > x_i$ for all $i \in S$.

Since the distribution with $\lambda_k = 1$ is both k -good and k -minimal, Theorem 5.1 applies immediately to the k -sampling best response dynamic.

We observed in Section 4.2 that the set $X_{S_0^*}$, where S_0^* is the smallest iterated 0-best response set, is globally asymptotically stable under the best response dynamic (B). Now if F is supermodular, then $\min S_0^*$ and $\max S_0^*$ are Nash equilibria; if they are strict equilibria, their associated states are asymptotically stable under (B). Thus by the argument following Proposition 5.3 below, we can conclude that in supermodular games whose pure Nash equilibria are strict, state e_{a^*} is globally asymptotically stable under (B) if and only if a^* is an iterated 0-dominant equilibrium.

Although this result for the best response dynamic parallels Theorem 5.1, the proof of the latter requires more subtle arguments. While global asymptotic stability of $X_{S_0^*}$ under (B) follows from the simple Lemma A.2 (Appendix A.1), the proof of *almost global* asymptotic stability of X_{S^*} under (S_λ) requires the transitivity theorem and some additional involved arguments as in Appendix A.6.²⁵ And while the asymptotic stability of strict equilibria under (B) is immediate, proving the asymptotic stability of $e_{\min S^*}$ and $e_{\max S^*}$ under (S_λ) requires tools from the theory of cooperative differential equations, as we explain below.

Example 5.2. In a game with a bilingual option, two basic actions correspond to strict equilibria. In choosing the a third action, the *bilingual option*, one pays a cost to be able to coordinate with agents playing either basic action.²⁶

Consider the following class of games with a bilingual option:

$$(5.2) \quad F(x) = Ax = \begin{pmatrix} k & k & 0 \\ k-c & k-c & k-1+\varepsilon-c \\ k-1 & k-1+\varepsilon & k-1+\varepsilon \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix},$$

where $k \geq 2$, $c > 0$, and $0 < \varepsilon < 1$. It is easy to verify that game (5.2) is supermodular. Actions 1 and 3 correspond to strict equilibria, with the former being Pareto dominant. If action 2 is omitted, then action 3 is $\frac{1}{k}$ -dominant (and not $\frac{1}{k+1}$ -dominant) in the resulting two-action game. But by choosing action 2, a player is able to coordinate with players of both action 1 and action 3, but must pay a cost of c in order to do so.

²⁵The complication arises because iterating p -dominance typically eliminates pure equilibrium states, which are rest points under (S_λ) .

²⁶Games with a bilingual option are studied by Galesloot and Goyal (1997), Goyal and Janssen (1997), and Oyama and Takahashi (2014). Galesloot and Goyal (1997) show that in a global interaction framework, if the cost of being bilingual is small enough, the Pareto dominant action is stochastically stable. Goyal and Janssen (1997) and Oyama and Takahashi (2014) prove that when such games are played by players situated in a network, the Pareto dominant action can spread contagiously in certain network structures.

If the distribution λ is k -good and k -minimal, and if it makes the game F above λ -generic, then the nature of equilibrium selection under the λ -sampling best response dynamic depends on the value of the cost parameter c . If $c > \frac{k-1}{k}$, then action 3 is a $\frac{1}{k}$ -dominant equilibrium, so state e_3 is almost globally asymptotically stable. If $c \in [\frac{1-\varepsilon}{k}, \frac{k-1}{k}]$, then the whole action set $S = \{1, 2, 3\}$ is the smallest iterated $\frac{1}{k}$ -best response set, so if $F(x) = Ax$ is λ -generic, states e_1 and e_3 are both locally asymptotically stable. Finally, if $c < \frac{1-\varepsilon}{k}$, then action 1 is an iterated $\frac{1}{k}$ -dominant equilibrium, so state e_1 is almost globally asymptotically stable. ♦

5.2 Analysis

To begin the proof of Theorem 5.1, we let $\underline{b}(x) = \min b(x)$ and $\bar{b}(x) = \max b(x)$ denote the smallest and largest pure best responses at population state x . It is well known (see, e.g., Sandholm (2010b, Section 3.4)) that in supermodular games, $\underline{b}(x)$ and $\bar{b}(x)$ are nondecreasing in the stochastic dominance order:

$$\text{If } \Sigma y \geq \Sigma x, \text{ then } \underline{b}(y) \geq \underline{b}(x) \text{ and } \bar{b}(y) \geq \bar{b}(x).$$

To proceed, we introduce two selections from the dynamic (S_λ) obtained by assuming that revising agents with multiple best responses always choose the smallest one, or, alternatively, the largest one. We define the functions $\underline{B}^k: X \rightarrow X$ and $\bar{B}^k: X \rightarrow X$ by

$$\underline{B}^k(x) = \min B^k(x) \text{ and } \bar{B}^k(x) = \max B^k(x),$$

where the min and max are taken with respect to the stochastic dominance order. Equivalently, $\underline{B}^k(x)$ and $\bar{B}^k(x)$ can be obtained by replacing $b(\frac{1}{k}z)$ in equation (2.1) with $\underline{b}(\frac{1}{k}z)$ and $\bar{b}(\frac{1}{k}z)$, respectively. We then define our two dynamics as follows:

$$\begin{aligned} (S_\lambda^-) \quad \dot{x} &= \sum_k \lambda_k \underline{B}^k(x) - x, \text{ and} \\ (S_\lambda^+) \quad \dot{x} &= \sum_k \lambda_k \bar{B}^k(x) - x. \end{aligned}$$

Since for each k , \underline{B}^k and \bar{B}^k are polynomials with values in the compact set X , equations (S_λ^-) and (S_λ^+) are continuous. Thus (S_λ^-) and (S_λ^+) admit solutions from each initial condition, which are also solutions to (S_λ) . If F is λ -generic, then (S_λ^-) , (S_λ^+) , and (S_λ) are identical.

To begin the analysis, we consider games in which the only $\frac{1}{k}$ -best response set is the full action set S . Proposition 5.3 shows that if λ is k -minimal, so that agents' samples are sufficiently informative, and if its mean μ_λ is finite, then the states corresponding to the

smallest and largest actions in S are asymptotically stable under the dynamics given by (S_λ^-) and (S_λ^+) , respectively.

Proposition 5.3. *Suppose that λ is k -minimal and that μ_λ is finite. Suppose that F is supermodular, and that $S = \{1, \dots, n\}$ itself is its smallest $\frac{1}{k}$ -best response set. Then state e_1 is asymptotically stable under (S_λ^-) and state e_n is asymptotically stable under (S_λ^+) .*

Theorem 5.1 is an easy consequence of this proposition and our earlier results. Part (i) of the theorem follows directly from Theorem 4.2, the transitivity theorem, this proposition, and the fact that (S_λ^-) , (S_λ^+) , and (S_λ) are identical when F is λ -generic. The “if” direction of part (ii) of the theorem follows directly from Theorem 4.2. For the “only if” direction, note first that since solutions to (S_λ^-) are solutions to (S_λ) , Theorem 4.2, the transitivity theorem, and the proposition imply that $e_{\min S^*}$ is asymptotically stable under (S_λ^-) even when S is not the smallest $\frac{1}{k}$ -best response set. It follows that from every state in some neighborhood of $e_{\min S^*}$, there is a solution to (S_λ) that converges to $e_{\min S^*}$. Analogous reasoning leads to the same conclusion about (S_λ) and $e_{\max S^*}$. Therefore, in order for (S_λ) to have an almost globally attractive state e_{a^*} , it must be that $e_{\min S^*} = e_{a^*} = e_{\max S^*}$, implying that a^* is an iterated $\frac{1}{k}$ -dominant equilibrium.

It remains to prove Proposition 5.3; by symmetry, it is enough to show that state e_1 is asymptotically stable under (S_λ^-) . To begin, we obtain a simple upper bound on the function \underline{b} as an implication of the assumptions that the game F is supermodular and that S is the smallest $\frac{1}{k}$ -best response set. We use this bound to show that after a change of variable, the dynamic (S_λ^-) can be bounded above by a new differential equation, (A.11). The asymptotic stability of state e_1 ’s counterpart under (A.11) is established via direct evaluation of the eigenvalues of the derivative matrix of (A.11). A variant of this calculation also shows that (A.11) is a *cooperative differential equation*: increasing any component of the state variable increases the growth rates of all other components. The comparison theorem for cooperative differential equations, Theorem A.11, then allows us to bound solution trajectories of (S_λ^-) by those of (A.11), and this in turn allows us to transfer the conclusion of asymptotic stability from (A.11) to (S_λ^-) . The details of this argument are provided in Appendix A.7.²⁷

²⁷Tools from the theory of monotone dynamical systems have been used in analyses of other game dynamics. Hofbauer and Sandholm (2002, 2007) show that stochastically perturbed best response dynamics for supermodular games are monotone dynamical systems, and use this fact to prove convergence results for stochastic fictitious play (Fudenberg and Kreps (1993)). Oyama et al. (2008) establish a comparison theorem for the perfect foresight dynamics (Matsui and Matsuyama (1995)). Their result establishes an order relation between perfect foresight equilibrium paths for two games, at least one of which is supermodular, and can be used to transfer equilibrium stability properties from one game to the other.

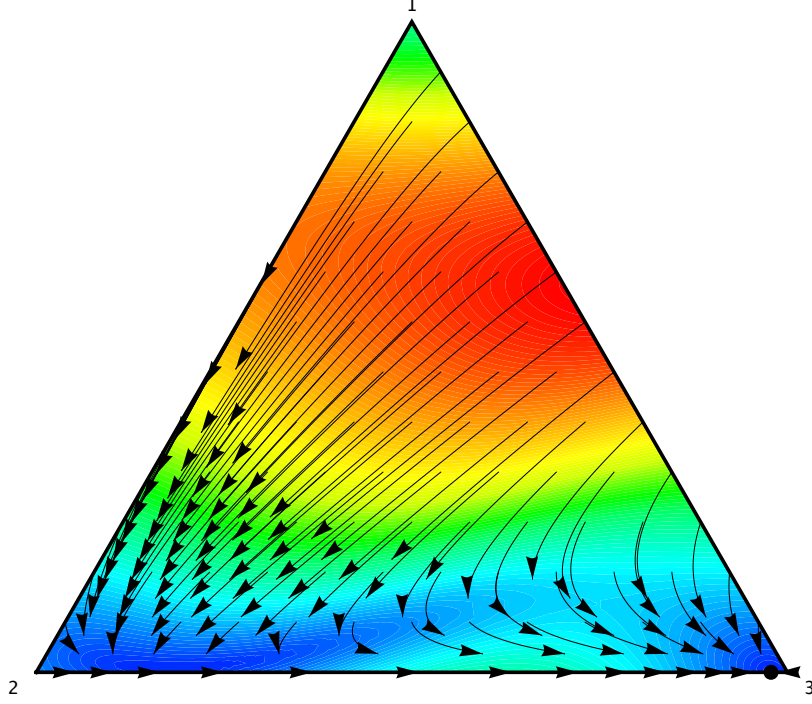


Figure 4: Phase diagram for the $\text{logit}(.76)$ dynamic in Young's game .

6. Discussion: Deterministic Equilibrium Selection

Our main result establishes a deterministic equilibrium selection result for games with iterated $\frac{1}{k}$ -dominant equilibria. If a game has such an equilibrium, then if λ is k -good, this equilibrium is almost globally asymptotically stable under the λ -sampling best response dynamic. This selection result does not apply in every game, but in the games where it can be used, the prediction it provides becomes relevant after little time has elapsed.

An important question for future research is to understand when deterministic equilibrium selection results can be obtained for dynamics based on other choice principles. As we noted in the introduction, the absence of such results may be attributed in part to the literature's focus on dynamics whose stability properties agree with those suggested by traditional solution concepts; by definition, such dynamics are unlikely to lead to selection results. In contrast, this paper shows that there are dynamics generated by natural choice rules under which equilibrium selection occurs.

We expect selection results to be available for dynamics based on other choice rules. For an example that is not so distant from the analysis here, consider the *logit dynamic*

with *noise level* η (Fudenberg and Levine (1998)):

$$(6.1) \quad \dot{x} = \tilde{M}^\eta(F(x)) - x, \text{ where } \tilde{M}_i^\eta(\pi) = \frac{\exp(\eta^{-1}\pi_i)}{\sum_{j \in S} \exp(\eta^{-1}\pi_j)}.$$

At low values of η , the logit dynamic closely approximates the exact best response dynamic (B). But at moderate noise levels, the connection with the exact best response dynamic weakens, allowing the possibility of deterministic equilibrium selection. For instance, in the case of Young's game (1.1), we have seen that the best response dynamic divides the simplex into three basins of attraction, one for each strict equilibrium (Figure 1(i)), while the 2-sampling best response dynamic generates almost global convergence to the exact Nash equilibrium e_3 (Figure 1(ii)). Figure 4 presents a phase diagram for this game under the logit dynamic with noise level $\eta = .76$. Evidently, this dynamic too engenders equilibrium selection, generating *global* convergence to the *approximate* Nash equilibrium $\tilde{x} \approx (.0000, .0218, .9782)$.

Of course, one can define deterministic evolutionary game dynamics based on many principles other than noisy optimization. The resulting criteria for equilibrium selection would likely differ from those considered here.

A. Appendix

A.1 Facts from the theory of differential inclusions

We now state some properties of differential inclusions to be used in our analysis. For further background, see Aubin and Cellina (1984), Smirnov (2002), and Benaïm et al. (2005, henceforth BHS).

Let $X \subseteq \mathbb{R}^n$ be a nonempty, compact, convex set, endowed with the max norm. Let $L = \max\{|a - b| : a, b \in X\}$, and let C^L be the set of Lipschitz continuous functions from the interval $[0, \infty)$ to X with Lipschitz constant L , endowed with the topology of uniform convergence on compact intervals, which is metrizable by, for example, the metric $D(x, y) = \sup_{t \geq 0} e^{-t}|x(t) - y(t)|$, where $x, y \in C^L$.

We focus on differential inclusions of the form

$$(DI) \quad \dot{x} \in G(x) - x,$$

where $G: X \rightrightarrows X$ is a nonempty-, compact-, and convex-valued upper semicontinuous correspondence. Our k - and λ -sampling best response dynamics belong to this class, as

does the best response dynamic of Gilboa and Matsui (1991) and Hofbauer (1995).

A (Carathéodory) solution to (DI) with initial condition $\xi \in X$ is a trajectory $x \in C^L$ such that $x(0) = \xi$ and $\dot{x}(t) + x(t) \in G(x(t))$ for almost all $t \geq 0$. It is known that (DI) admits (typically nonunique) solutions from every initial condition (cf. BHS, Remark 1.2), and that the set of all solutions to (DI) forms a compact subset of C^L .

Let $\Phi: [0, \infty) \times X \Rightarrow X$ be the set-valued dynamical system on the compact set X induced by (DI):

$$\Phi_t(\xi) = \{x(t) : x \text{ is a solution to (DI) with } x(0) = \xi\}.$$

One can verify that Φ is an upper semicontinuous map that satisfies $\Phi_0(\xi) = \xi$ and $\Phi_t(\Phi_s(\xi)) = \Phi_{t+s}(\xi)$ for all $t, s \geq 0$. Moreover, the upper semicontinuity of Φ implies that for any compact interval $I \subseteq [0, \infty)$, the correspondence $\Phi_I: X \Rightarrow X$ defined by $\Phi_I(\xi) = \bigcup_{t \in I} \Phi_t(\xi)$ is itself upper semicontinuous.

For a nonempty set $A \subseteq X$, define the ω -limit set of A by

$$\omega(A) = \bigcap_{t \geq 0} \overline{\Phi_{[t, \infty)}(A)}.$$

That is, $z \in \omega(A)$ if and only if $z = \lim_{n \rightarrow \infty} x_n(t_n)$ for some sequence (x_n) of solutions to (DI) with $x_n(0) \in A$ and some sequence (t_n) that approaches infinity. Note that $\bigcup_{\xi \in A} \omega(\xi) \subseteq \omega(A)$, and that in general this inclusion is strict, since the definition of $\omega(A)$ allows one to move between different initial conditions in A as one considers later times t_n . By compactness of X , $\omega(A)$ is nonempty and compact. It is immediate from the definition that if $A \subseteq B$, then $\omega(A) \subseteq \omega(B)$. We also have the following.

Lemma A.1. *For $A \subseteq X$, for any $z \in \omega(A)$ and $t \geq 0$, there exists a solution x to (DI) such that $x(0) \in \omega(A)$ and $x(t) = z$.*

Proof. Suppose that $z \in \omega(A)$, so that there are sequences of times (t_n) and solutions (x_n) such that $\lim_{n \rightarrow \infty} t_n = \infty$, $x_n(0) \in A$, and $\lim_{n \rightarrow \infty} x_n(t_n) = z$. Fix any $t \geq 0$, and define a sequence of solutions (y_n) by $y_n(s) = x_n(t_n - t + s)$ (which makes sense for values of n such that $t_n \geq t$). By compactness of the set of solutions, there is a subsequence (y_{n_k}) of (y_n) converging to some solution y satisfying $y(0) = \lim_{n_k \rightarrow \infty} x_{n_k}(t_{n_k} - t) \in \omega(A)$ and $y(t) = z$. ■

Lemma A.2. *For $A \subseteq X$, we have $\omega(A) \subseteq \omega(\omega(A))$. Thus, if $\omega(A) \subseteq B$, then $\omega(A) \subseteq \omega(B)$.*

Proof. Let $z \in \omega(A)$. Then Lemma A.1 implies that $z \in \Phi_t(\omega(A))$ for all $t \geq 0$. It follows that $z \in \omega(\omega(A))$. Thus, if $\omega(A) \subseteq B$, then $\omega(A) \subseteq \omega(\omega(A)) \subseteq \omega(B)$. ■

Asymptotic stability as defined in Section 4.2 has a characterization in terms of so-called attracting sets. We call $A \subseteq X$ an *attracting set* if it is compact and if there exists a neighborhood U of A such that for any $\varepsilon > 0$, there exists a time $t_\varepsilon > 0$ such that

$$(A.1) \quad \Phi_{[t_\varepsilon, \infty)}(U) \subseteq N^\varepsilon(A),$$

where $N^\varepsilon(A) = \{x \in X: \inf_{y \in A} |x - y| < \varepsilon\}$. Such a set U is called a *fundamental neighborhood* of A . Lemma A.3 (BHS Proposition 3.10(i)) provides an alternative characterization of attracting sets and their fundamental neighborhoods in terms of ω -limits, while Lemma A.4 (BHS Corollary 3.17) provides a characterization of asymptotic stability. (Note that unlike that in BHS, the definition of asymptotically stable sets used here does not require them to be invariant.)

Lemma A.3. *Let A be a nonempty compact subset of X , and let U be a neighborhood of A . Then A is an attracting set with fundamental neighborhood U if and only if $\omega(U) \subseteq A$.*

Lemma A.4. *A compact set A is attracting and strongly positively invariant if and only if it is asymptotically stable.*

A.2 Upper semicontinuity of λ -sampling best response dynamics

The following result was noted in Section 2.3.

Proposition A.5. *Fix any probability distribution on the natural numbers $\lambda = \{\lambda_k\}_{k=1}^\infty$. The correspondence $\sum_k \lambda_k B^k: X \Rightarrow X$ is compact- and convex-valued and upper semicontinuous.*

Proof. Convex valuedness is immediate. To prove that $\sum_k \lambda_k B^k$ is compact-valued and upper semicontinuous, we introduce the countable product of compact sets $X^\infty = \prod_{k=1}^\infty X$, endowed with the topology of pointwise convergence. This topology is metrizable by, for example, the metric $d(\mathbf{x}, \mathbf{y}) = \sup_k \frac{1}{k} |x^k - y^k|$, where $\mathbf{x} = (x^1, x^2, \dots)$, $\mathbf{y} = (y^1, y^2, \dots) \in X^\infty$, and one can readily show that X^∞ is compact with respect to this metric (see Munkres (2000, p. 280)). Define the correspondence $G: X \Rightarrow X^\infty$ by $G(x) = \prod_{k=1}^\infty B^k(x)$ and the function $g: X^\infty \rightarrow X$ by $g(\mathbf{x}) = \sum_k \lambda_k x^k$. That G has a closed graph follows easily from the fact that B^k has a closed graph for each $k \geq 1$. Since the codomain X^∞ of G is a compact metric space, it follows that G is compact-valued and upper semicontinuous. Since g is clearly continuous, $\sum_k \lambda_k B^k = g \circ G$ must be compact-valued and upper semicontinuous as well. ■

A.3 Proof of Proposition 4.1

We begin with a lemma.

- Lemma A.6.** (i) Let $T \subseteq S$, and suppose that T is a p -best response set in $F|_S$. If $T \subseteq S' \subseteq S$, then T is a p -best response set in $F|_{S'}$.
- (ii) Let $p \leq \frac{1}{2}$. Let $R, T \subseteq S$, and suppose that R and T are p -best response sets in $F|_S$. Then $R \cap T$ is nonempty and it is a p -best response set in $F|_R$ and $F|_T$.

Proof. The proof of part (i) is immediate. To prove part (ii), note that since $F|_S$ has a smallest p -best response set, it must be that $R \cap T \neq \emptyset$. It follows immediately that $R \cap T$ is a p -best response set in $F|_S$, so the claim follows from part (i). ■

We proceed with the proof of Proposition 4.1. Let $p \leq \frac{1}{2}$. By the existence results of Tercieux (2006), we can define $\{S^\ell\}_{\ell=0}^\infty$ as the sequence such that $S^0 = S$, and S^ℓ is the smallest p -best response set in $F|_{S^{\ell-1}}$ for all $\ell = 1, 2, \dots$. Let $S^* = \bigcap_{\ell=0}^\infty S^\ell$. By the finiteness of S , there is a finite m such that $S^* = S^m$. We show that S^* is the smallest iterated p -best response set.

Let T be any iterated p -best response set, and $S = T^0 \supseteq \dots \supseteq T^{m'} = T$ its associated sequence. We assume without loss of generality that $m = m'$. We show that

$$(I_\ell) \quad S^\ell \text{ is an iterated } p\text{-best response set in } F|_{T^\ell}$$

for all $\ell = 0, \dots, m$.

First, (I_0) is trivially true. So assume (I_ℓ) , and let $\{R^h\}_{h=0}^{\bar{h}}$ with $T^\ell = R^0 \supseteq R^1 \supseteq \dots \supseteq R^{\bar{h}} = S^\ell$ be an associated sequence. Set $R^{\bar{h}+1} = S^{\ell+1}$. We claim that $S^{\ell+1}$ is an iterated p -best response set in $F|_{T^{\ell+1}}$ such that the sequence $\{\tilde{R}^h\}_{h=0}^{\bar{h}+1}$ defined by $\tilde{R}^0 = T^{\ell+1}$ and $\tilde{R}^h = \tilde{R}^{h-1} \cap R^h$ for $h = 1, \dots, \bar{h} + 1$ is an associated sequence.

Indeed, since $\tilde{R}^0 (= T^{\ell+1})$ and R^1 are p -best response sets in $F|_{T^\ell}$, $\tilde{R}^1 (= \tilde{R}^0 \cap R^1)$ is a p -best response set in $F|_{\tilde{R}^0}$ and $F|_{R^1}$ by Lemma A.6(ii); and if \tilde{R}^h is a p -best response set in $F|_{\tilde{R}^{h-1}}$ and $F|_{R^h}$, then, since R^{h+1} is a p -best response set in $F|_{R^h}$, $\tilde{R}^{h+1} (= \tilde{R}^h \cap R^{h+1})$ is a p -best response set in $F|_{\tilde{R}^h}$ and $F|_{R^{h+1}}$ by Lemma A.6(ii). Finally, since $\tilde{R}^{\bar{h}+1} (= S^{\ell+1})$ is the smallest p -best response set in $F|_{\tilde{R}^{\bar{h}}} (= F|_{S^\ell})$ and $\tilde{R}^{\bar{h}}$ is a p -best response set in $F|_{R^{\bar{h}}}$, we have $\tilde{R}^{\bar{h}+1} \subseteq \tilde{R}^{\bar{h}}$, so that $\tilde{R}^{\bar{h}+1} = \tilde{R}^{\bar{h}+1} = S^{\ell+1}$. ■

A.4 Proof of Lemma 4.3

Since R is a 1-best response set, all best responses to samples in which all those sampled play actions in R are themselves in R . It follows that $\dot{x}_{S \setminus R} = 0$ whenever $x_R = 1$. If (S_λ) is Lipschitz continuous, the result easily follows.

To prove the general case, let $Q = S \setminus R$. Since R is a 1-best response set, there is a $p < 1$ for which it is a p -best response set. Thus for a revising agent to choose an action in Q , it is necessary that at least fraction $1 - p$ of his sample choose an action in Q . Hence for the ℓ -sampling best response dynamic, we have the bound

$$(A.2) \quad \dot{x}_Q \leq \mathbb{P}(\mathcal{B}_{\ell, x_Q} \geq (1 - p)\ell) - x_Q.$$

We show next that for all sufficiently large ℓ , the right hand side of (A.2) is negative when $x_Q \in (0, \frac{1}{2}(1 - p)]$. (It follows that for such ℓ , $\Lambda(x) = x_R$ is a local Lyapunov function for X_R under the ℓ -sampling best response dynamic, and hence that X_R is asymptotically stable under this dynamic.)

Write $f^\ell(q) = \mathbb{P}(\mathcal{B}_{\ell, q} \geq (1 - p)\ell)$. The proof of Lemma 2 of Kaniovski and Young (1995) shows that f^ℓ is convex on $[0, \lfloor (1 - p)\ell \rfloor / (\ell - 1)]$; in particular, f^ℓ is convex on $[0, \frac{1}{2}(1 - p)]$ for all ℓ large enough (see Figure 2). And clearly, $f^\ell(0) = 0$ for all ℓ , and $\lim_{\ell \rightarrow \infty} f^\ell(q) = 0$ for all $q < 1 - p$. Thus there exists an $L < \infty$ such that $f^\ell(\frac{1}{2}(1 - p)) < \frac{1}{2}(1 - p)$ for all $\ell > L$, and thus, by convexity, $f^\ell(q) < q$ for all $q \in (0, \frac{1}{2}(1 - p)]$ and $\ell > L$.

Returning to the λ -sampling best response dynamic, take any solution $x: [0, \infty) \rightarrow X$ of (S_λ) with $x_Q(0) = 0$, and let $T = \sup\{t \geq 0: x_Q(s) \leq \frac{1}{2}(1 - p) \text{ for all } s \in [0, t]\}$, which is positive by the continuity of x . Then the previous paragraph implies that for all $t \in [0, T]$,

$$\begin{aligned} \dot{x}_Q(t) &\leq \sum_{\ell=1}^L \lambda_\ell \left(\mathbb{P}(\mathcal{B}_{\ell, x_Q(t)} \geq 1) - x_Q(t) \right) + \sum_{\ell=L+1}^{\infty} \lambda_\ell \left(\mathbb{P}(\mathcal{B}_{\ell, x_Q(t)} \geq (1 - p)\ell) - x_Q(t) \right) \\ &\leq \sum_{\ell=1}^L \lambda_\ell \left(1 - (1 - x_Q(t))^\ell - x_Q(t) \right) \\ &\leq M x_Q(t), \end{aligned}$$

where $M = \sum_{\ell=1}^L \lambda_\ell(\ell - 1)$, and where the third inequality follows from taking tangent approximations of the concave functions $1 - (1 - x_Q)^j$ at $x_Q = 0$. Thus

$$(A.3) \quad \frac{d}{dt}(e^{-Mt} x_Q(t)) = -M e^{-Mt} x_Q(t) + e^{-Mt} \dot{x}_Q(t) \leq 0,$$

so integration yields $x_Q(t) \leq e^{Mt} x_Q(0)$. We conclude that $T = \infty$, and hence that $x_Q(t) = 0$ for all $t \in [0, \infty)$. In other words, no solution to (S_λ) can escape X_R . ■

A.5 Proof of Theorem 4.7

The proof of Theorem 4.7 is a modification of that of Theorem 5.3.D in Conley (1978). While Conley (1978) considers single-valued dynamical systems defined at all times $t \in (-\infty, \infty)$, we allow for set-valued systems, and for systems that are only defined in forward time. In particular, while Conley's (1978) analysis makes use of the notion of a repeller, we do not make use of that notion here.

Recall that B is asymptotically stable in A if A is strongly positively invariant and B is asymptotically stable with respect to $\Phi|_A$, the restriction of Φ to A . Let U_B be a fundamental neighborhood of B in A (so that $\omega(U_B) \subseteq B$), and U_C a fundamental neighborhood of C in B (so that $\omega(U_C) \subseteq C$). Let U be a compact neighborhood of C in A such that $U \subseteq U_B$ and $U \cap B \subseteq U_C$.

Claim 1. *For all $z \in \partial U$, there exists $t_z > 0$ such that, for any solution x and for all $t \geq t_z$, if $x(t) = z$, then $x(t - t_z) \notin U$.*

Proof. Suppose not. Then, there exist $z \in \partial U$, a sequence of solutions (x_n) , and a sequence (t_n) approaching infinity such that $x_n(0) \in U$ and $x_n(t_n) = z$. If $z \in \partial U \setminus B$, then, since $(x_n(0)) \subseteq U \subseteq U_B$, by definition we have $z \in \omega(U_B)$, which contradicts $\omega(U_B) \subseteq B$. If $z \in \partial U \cap B$, then, since $(x_n(0)) \subseteq B$ by the strong negative invariance of B and hence $(x_n(0)) \subseteq U \cap B \subseteq U_C$, we have $z \in \omega(U_C)$, which contradicts $\omega(U_C) \subseteq C$. ■

By the compactness of the set of solutions (see Appendix A.1), the correspondence $\xi \mapsto \{x(0) : x \text{ is a solution to } (S_\lambda) \text{ with } x(t) = \xi\}$ is upper semicontinuous for any t . Thus for each $z \in \partial U$, there exists an open neighborhood V_z of z such that, for any solution x and for all $t \geq t_z$, if $x(t) \in V_z$, then $x(t - t_z) \notin U$, where t_z is specified in Claim 1. Since ∂U is compact, there are states $z^1, \dots, z^N \in \partial U$ such that $\bigcup_i V_{z^i} \supseteq \partial U$. Thus, if $z \in \partial U$, then for some z^j , $z \in V_{z^j}$ and, for any $t \geq t_{z^j}$, $x(t - t_{z^j}) \notin U$ whenever $x(t) = z$. Now define $T = \max_i t_{z^i}$.

Claim 2. *For all $\xi \in U$, if $\Phi_{[0,T]}(\xi) \subseteq U$, then $\Phi_{[0,\infty)}(\xi) \subseteq U$.*

Proof. Suppose that $\Phi_{[0,T]}(\xi) \subseteq U$. Let $t' = \sup\{t : \Phi_{[0,t]}(\xi) \subseteq U\}$ ($\geq T$). Assume that $t' < \infty$. Then we have $\Phi_{t'}(\xi) \subseteq U$ (by the closedness of U and the continuity of solutions) and $\Phi_{[t'-T, t']}(\xi) \subseteq U$. If $\Phi_{t'}(\xi) \cap \partial U \neq \emptyset$, then for any $z \in \Phi_{t'}(\xi) \cap \partial U$ and for any solution x such that $x(t') = z$, we have $x(t' - t_{z^j}) \notin U$ by construction of the sets V_{z^j} (where z^j is such that $z \in V_{z^j}$), which is a contradiction. Therefore $\Phi_{t'}(\xi) \cap \partial U = \emptyset$, which implies that $\Phi_{t'+\varepsilon}(\xi) \subseteq U$ for some $\varepsilon > 0$ by the upper semicontinuity of Φ , but then t' could not be the supremum as defined. Hence, $t' = \infty$. ■

Let U^0 be the upper inverse image of U under the correspondence $\Phi_{[0,T]}(\cdot)$, i.e., $U^0 = \{\xi : \Phi_{[0,T]}(\xi) \subseteq U\} (\subseteq U)$. Since C is strongly positively invariant, we have $\Phi_{[0,T]}(C) \subseteq C \subseteq U$. Therefore, by the upper semicontinuity of $\Phi_{[0,T]}(\cdot)$, U^0 is a neighborhood of C in A .

Now, by Claim 2, we have $\Phi_{[0,\infty)}(U^0) \subseteq U$, and hence, $\omega(U^0) \subseteq U$ (by the closedness of U). Since $U^0 \subseteq U_B$, $\omega(U^0) \subseteq \omega(U_B) \subseteq B$ by the attraction of B in A . Thus, $\omega(U^0) \subseteq U \cap B \subseteq U_C$. Therefore, we have $\omega(U^0) \subseteq \omega(U_C) \subseteq C$ by Lemma A.2 and the attraction of C in B , which implies, by Lemma A.3, that C is an attracting set in A . Since C is strongly positively invariant, it follows from Lemma A.4 that C is asymptotically stable in A . This completes the proof of Theorem 4.7. ■

A.6 Completion of the proof of Theorem 4.2

In this appendix, we establish that X_{S^*} is almost globally attractive under the dynamic (S_λ) .

Since λ satisfies the k -goodness condition (3.1), it immediately follows that

$$\sum_{j \leq k} \lambda_j j > 1, \text{ and hence } \alpha \equiv \left(\sum_{j \leq k} \lambda_j j \right)^{\frac{1}{1-k}} \in (0, 1).$$

We therefore let

$$Y^\ell = \{x \in O_{S^*} : x_{\bar{S}^\ell} \geq \alpha\}$$

be the set of states that give the set S^* (and hence S^ℓ) positive mass and give actions in \bar{S}^ℓ a total mass of at least α , as shown in Figure 5. In terms of that figure, the following lemma states that within the set Y^ℓ , solution trajectories cross lines parallel to the dashed line from left to right.

Lemma A.7. *For all $\ell \in \{1, \dots, m\}$ and for all $\tau_0 < \tau_1$, if $x(t) \in Y^\ell$ for all $t \in [\tau_0, \tau_1]$, then there exists $c > 0$ such that $\dot{x}_{S^\ell}(t) \geq c$ for almost all $t \in [\tau_0, \tau_1]$.*

Proof. Let $z \in \mathbb{Z}_+^{n,j}$ be the result of a sample of size $j \leq k$. By construction, if z satisfies $z_{S^{\ell-1}} = j$ and $z_{S^\ell} \geq 1$, then $b(\frac{1}{j}z) \subseteq S^\ell$. This observation, the definition of \bar{S}^ℓ , and the binomial theorem imply that

$$\begin{aligned} \dot{x}_{S^\ell} &\geq \sum_{j \leq k} \lambda_j \left((x_{S^{\ell-1}})^j - (x_{\bar{S}^\ell})^j \right) - x_{S^\ell} \\ &= \sum_{j \leq k} \lambda_j \left((x_{S^\ell} + x_{\bar{S}^\ell})^j - (x_{\bar{S}^\ell})^j \right) - x_{S^\ell} \end{aligned}$$

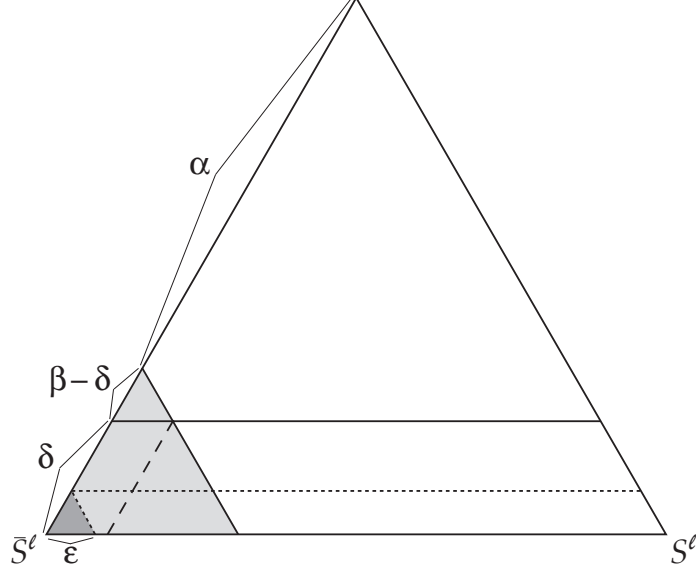


Figure 5: The proofs of Lemmas A.7 and A.8. Y^ℓ is the union of the shaded regions, excluding the northwest boundary. The southern boundary of the simplex represents $X_{S^{\ell-1}}$. Points southeast of the dashed line satisfy $x_{S^\ell} > \beta - \delta > \varepsilon$.

$$\begin{aligned}
&= \sum_{j \leq k} \lambda_j \sum_{h=1}^j \binom{j}{h} (x_{S^\ell})^h (x_{\bar{S}^\ell})^{j-h} - x_{S^\ell} \\
&= \sum_{j \leq k} \lambda_j \sum_{h=2}^j \binom{j}{h} (x_{S^\ell})^h (x_{\bar{S}^\ell})^{j-h} + \left(\sum_{j \leq k} \lambda_j j (x_{\bar{S}^\ell})^{j-1} - 1 \right) x_{S^\ell} \\
&\geq \sum_{j \leq k} \lambda_j \sum_{h=2}^j \binom{j}{h} (x_{S^\ell})^h (x_{\bar{S}^\ell})^{j-h} + \left((x_{\bar{S}^\ell})^{k-1} \sum_{j \leq k} \lambda_j j - 1 \right) x_{S^\ell}.
\end{aligned}$$

The first term in the last line is positive on Y^ℓ , and since $x_{\bar{S}^\ell} \geq \alpha$ on Y^ℓ , the expression in parentheses is nonnegative on Y^ℓ . Therefore, \dot{x}_{S^ℓ} has a positive lower bound on the set $\{y \in Y^\ell : y_{S^\ell} \geq x_{S^\ell}(\tau_0)\}$, proving the lemma. ■

Lemma A.8 shows that the ω -limit set of any $\xi \in O_{S^*}$ cannot intersect $X_{\bar{S}^\ell}$. The proof of this lemma relies on the Lyapunov stability of the set $X_{S^{\ell-1}}$ in X established in Proposition 4.8, and so makes indirect use of the transitivity theorem.

Lemma A.8. *For any $\xi \in O_{S^*}$, $\omega(\xi) \cap X_{\bar{S}^\ell} = \emptyset$ for all $\ell \in \{1, \dots, m\}$.*

Proof. Let $\beta = 1 - \alpha \in (0, 1)$, and choose $\delta \in (0, \beta)$. Since $X_{S^{\ell-1}}$ is Lyapunov stable by Proposition 4.8, we can choose an $\varepsilon \in (0, \beta - \delta)$ such that any solution to (S_λ) whose initial condition $\chi \in X$ satisfies $\chi_{S^{\ell-1}} > 1 - \varepsilon$ never reaches a state $\hat{\chi}$ with $\hat{\chi}_{S^{\ell-1}} \leq 1 - \delta$ (see Figure 5).

Let $\xi \in O_{S^*}$, and suppose that $\omega(\xi) \cap X_{\bar{S}^\ell} \neq \emptyset$. By definition, there is a solution x of (S_λ) with initial condition $x(0) = \xi$ for which there is a time $t_0 > 0$ such that $x_{\bar{S}^\ell}(t_0) > 1 - \varepsilon$ ($> \alpha$); since the form of (S_λ) implies that for all $t \geq 0$,

$$(A.4) \quad x_{S^*}(t) \geq x_{S^*}(0) e^{-t} > 0,$$

and hence $x(t) \in O_{S^*}$, it follows that $x(t_0) \in Y^\ell$. By Lemma A.7, there must be a time $t_1 > t_0$ such that $x(t_1) \notin Y^\ell$.

To complete the proof of the lemma, we now argue that the solution x cannot satisfy $x_{\bar{S}^\ell}(t) > 1 - \varepsilon$ at any time $t > t_1$, contradicting our supposition that $\omega(\xi) \cap X_{\bar{S}^\ell} \neq \emptyset$. Indeed, if $x_{\bar{S}^\ell}(t_3) > 1 - \varepsilon$ at some time $t_3 > t_1$, then accounting for equation (A.4), there must be a time $t_2 \in (t_1, t_3)$ such that $x_{\bar{S}^\ell}(t_2) = 1 - \beta$ and such that $x(t) \in Y^\ell$ ($\subseteq O_{S^*}$) for all $t \in [t_2, t_3]$. But since $x_{S^{\ell-1}}(t_0) > x_{\bar{S}^\ell}(t_0) > 1 - \varepsilon$, the definitions of ε and δ imply that $x_{S^{\ell-1}}(t_2) > 1 - \delta$ (by Lyapunov stability); in terms of Figure 5, the trajectory crosses, at time t_2 , the northeast boundary of Y^ℓ below the solid horizontal line. Combining these facts with the definition of \bar{S}^ℓ yields

$$x_{S^\ell}(t_2) = x_{S^{\ell-1}}(t_2) - x_{\bar{S}^\ell}(t_2) > (1 - \delta) - (1 - \beta) = \beta - \delta > \varepsilon.$$

Thus, Lemma A.7 implies that $x_{S^\ell}(t) \geq x_{S^\ell}(t_2) > \varepsilon$ for all $t \in [t_2, t_3]$, which contradicts the existence of a time t_3 with the required properties. This completes the proof of the lemma. ■

We now use Lemma A.8 to establish almost global attraction.

Proposition A.9. *For any $\xi \in O_{S^*}$, $\omega(\xi) \subseteq X_{S^\ell}$ for all $\ell \in \{1, \dots, m\}$.*

Proof. Let $\xi \in O_{S^*}$. Our proof is by induction. That $\omega(\xi) \subseteq X_{S^1}$ follows from Lemma 4.4. For the inductive step, for $\ell \in \{2, \dots, m\}$ suppose that $\omega(\xi) \subseteq X_{S^{\ell-1}}$. Lemma A.8 tells us that $\omega(\xi) \subseteq X_{S^{\ell-1}}^\circ (= X_{S^{\ell-1}} \setminus X_{\bar{S}^\ell})$.

Now let $z \in \operatorname{argmin}_{y \in \omega(\xi)} y_{S^\ell}$, which is well defined since $\omega(\xi)$ is compact. By Lemma A.1 (with $A = \{\xi\}$), there exist a solution x to (S_λ) (restricted to $X_{S^{\ell-1}}$) and a $t > 0$ such that $x(0) \in \omega(\xi)$ and $x(t) = z$. If $z_{S^\ell} < 1$, then Lemma 4.5 would imply that $z_{S^\ell} > x_{S^\ell}(0)$, which contradicts the definition of z . We thus conclude that $z_{S^\ell} = 1$, and so that $\omega(\xi) \subseteq X_{S^\ell}$. ■

Together, Propositions 4.8 and A.9 imply Theorem 4.2.

A.7 Proof of Proposition 5.3

We begin with a simple lemma that bounds the minimal and maximal best responses to extreme samples of size $\ell \geq k$. As an aside, we note that together with the fact that an iterated p -best response set is a 1-best response set, this lemma implies that if S^* is the smallest iterated p -best response set of a supermodular game, then $\min S^*$ and $\max S^*$ are Nash equilibria; let $i = j = 1$ in (A.5) and $i = j = n$ in (A.6). In addition, if the underlying game is k -generic for some $k \geq 1$, then $\min S^*$ and $\max S^*$ are strict equilibria.

Lemma A.10. *Suppose that F is supermodular, that S itself is the smallest $\frac{1}{k}$ -best response set, and that $\ell \geq k$. Then*

$$(A.5) \quad \underline{b}(\frac{\ell-1}{\ell}e_1 + \frac{1}{\ell}e_j) \leq i \text{ for all } i \leq n-1 \text{ and } j \leq i+1, \text{ and}$$

$$(A.6) \quad \bar{b}(\frac{1}{\ell}e_j + \frac{\ell-1}{\ell}e_n) \geq i \text{ for all } i \geq 2 \text{ and } j \geq i-1.$$

Proof. We prove (A.5). Because $\{i+1, \dots, n\}$ is not a $\frac{1}{k}$ -best response set, there is a state $x \in X$ with $x_{\{i+1, \dots, n\}} \geq \frac{1}{k}$ such that $\underline{b}(x) \leq i$. But $\Sigma(\frac{\ell-1}{\ell}e_1 + \frac{1}{\ell}e_j) \leq \Sigma x$, so the claim follows from the monotonicity of \underline{b} . The proof of (A.6) is similar. ■

Fix a sample size $\ell \geq k$. Condition (A.5) tells us that for sample outcomes $z \in \mathbb{Z}_+^{n, \ell}$, $\underline{b}(\frac{1}{\ell}z) \leq i$ whenever (i) $z_1 = \ell$, or (ii) $z_1 = \ell - 1$ and $z_j = 1$ for some $j \leq i+1$. It follows that

$$(A.7) \quad \sum_{h \geq i+1} \underline{B}_h^\ell(x) \leq 1 - (x_1)^\ell - \ell(x_1)^{\ell-1}x_{\{2, \dots, i+1\}} \quad \text{for } i \in \{1, \dots, n-1\}.$$

It will be useful to study the dynamic (S_λ^-) after employing the linear change of variable defined by the stochastic dominance operator Σ . Doing so transports the dynamic from the simplex $X \subseteq \mathbb{R}^n$ to the convex set $\mathcal{X} \subseteq \mathbb{R}^{n-1}$. The transported dynamic can be expressed as

$$(A.8) \quad \dot{\chi} = \mathcal{V}(\chi) \equiv \sum_\ell \lambda_\ell \underline{\mathcal{B}}^\ell(\chi) - \chi,$$

where $\underline{\mathcal{B}}^\ell(\chi) = \Sigma \underline{B}^\ell(\tilde{\Sigma}\chi + e_1)$. Since dynamics (S_λ^-) and (A.8) are linearly conjugate, to prove the proposition it is sufficient to show that $\Sigma e_1 = \mathbf{0} \in \mathcal{X}$ is asymptotically stable under (A.8).

After the change of variable, (A.7) says that for sample sizes $\ell \geq k$, and thus for all sample sizes with $\lambda_\ell > 0$,

$$(A.9) \quad \underline{\mathcal{B}}_i^\ell(\chi) \leq 1 - (1 - \chi_1)^\ell - \ell(1 - \chi_1)^{\ell-1}(\chi_1 - \chi_{i+1}) \quad \text{for } i \in \{1, \dots, n-1\},$$

where we set $\chi_n \equiv 0$. Define the function $\tilde{\mathcal{B}}^\ell: \mathcal{X} \rightarrow \mathcal{X}$ by the right-hand side of (A.9). Then direct calculation shows that for $\ell \geq k$,

$$(A.10) \quad \frac{\partial \tilde{\mathcal{B}}_i^\ell}{\partial \chi_j}(\chi) = \begin{cases} \ell(\ell-1)(1-\chi_1)^{\ell-2}(\chi_1 - \chi_{i+1}) & \text{if } j = 1, \\ \ell(1-\chi_1)^{\ell-1} & \text{if } j = i+1, \\ 0 & \text{otherwise.} \end{cases}$$

One can verify that for each $\ell \geq k$, $0 \leq \frac{\partial \tilde{\mathcal{B}}_i^\ell}{\partial \chi_j}(\chi) \leq \ell$ for all $\chi \in \mathcal{X}$. Since λ has finite mean μ_λ , the function $\sum_\ell \lambda_\ell \tilde{\mathcal{B}}^\ell$ is differentiable by the dominated convergence theorem, and is Lipschitz continuous with Lipschitz constant μ_λ .

Now consider the differential equation

$$(A.11) \quad \dot{\chi} = \tilde{\mathcal{V}}(\chi) \equiv \sum_\ell \lambda_\ell \tilde{\mathcal{B}}^\ell(\chi) - \chi$$

on \mathcal{X} . Equation (A.10) implies that the derivative matrix of $\tilde{\mathcal{V}}$ at $\mathbf{0}$ is

$$D\tilde{\mathcal{V}}(\mathbf{0}) = \begin{pmatrix} -1 & \mu_\lambda & 0 & \cdots & 0 \\ 0 & -1 & \mu_\lambda & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ \vdots & & \ddots & -1 & \mu_\lambda \\ 0 & \cdots & \cdots & 0 & -1 \end{pmatrix},$$

where $\mu_\lambda < \infty$ is the mean of distribution λ . As this matrix is triangular, it is evident that all $n-1$ of its eigenvalues equal -1 . Thus, $\mathbf{0} \in \mathcal{X}$ is asymptotically stable under (A.11).

Equation (A.10) also implies that the dynamic (A.11) is a *cooperative differential equation*, meaning that all cross-partial derivatives of $\tilde{\mathcal{V}}$ are nonnegative throughout \mathcal{X} . Condition (A.9) implies that for all $\chi \in \mathcal{X}$, $\mathcal{V}(\chi) \leq \tilde{\mathcal{V}}(\chi)$. Proposition 5.3 will thus follow from the following comparison theorem for cooperative differential equations, which is Theorem I.6.VII of Walter (1970) applied to our context.

Theorem A.11 (Comparison theorem). *Let $f: \mathcal{X} \rightarrow \mathbb{R}^{n-1}$ be a Lipschitz continuous function satisfying $\frac{\partial f_i}{\partial \chi_j} \geq 0$ for all $j \neq i$. If absolutely continuous functions $y, z: [0, \infty) \rightarrow \mathcal{X}$ satisfy $y(0) \ll z(0)$, $\dot{y}(t) \leq f(y(t))$, and $\dot{z}(t) \geq f(z(t))$ for almost all $t \geq 0$, then $y(t) \ll z(t)$ for all $t \geq 0$.*

Proof. Let $T = \sup\{t \geq 0: y(s) \ll z(s) \text{ for all } s \in [0, t)\} > 0$. Suppose by way of contradiction that $T < \infty$. Then $z(t) - y(t) \geq 0$ for all $t \in [0, T]$ and $z_i(T) - y_i(T) = 0$ for some i . For any

such i ,

$$\begin{aligned}
(A.12) \quad \frac{d}{dt}(z_i(t) - y_i(t)) &\geq f_i(z(t)) - f_i(y(t)) \\
&\geq f_i(y_1(t), \dots, y_{i-1}(t), z_i(t), y_{i+1}(t), \dots, y_{n-1}(t)) - f_i(y(t)) \\
&\geq -K(z_i(t) - y_i(t))
\end{aligned}$$

for almost all $t \in [0, T]$, where the second inequality follows from the fact that $\frac{\partial f_i}{\partial \chi_j} \geq 0$ for all $j \neq i$, and the third from the Lipschitz continuity of f with Lipschitz constant $K > 0$. Mimicking (A.3), we use (A.12) to show that $\frac{d}{dt}e^{Kt}(z_i(t) - y_i(t)) \geq 0$, and then integrate to obtain $z_i(T) - y_i(T) \geq e^{-KT}(z_i(0) - y_i(0))$. But since the right hand side of this inequality is positive, we have a contradiction of the definition of T . ■

We can now complete the proof of Proposition 5.3. Theorem A.11 has the following implication: Let y be the solution to differential equation (A.8), and let z be the solution to differential equation (A.11) from an initial condition $z(0)$ satisfying $y(0) \ll z(0)$. As shown above, we have $\dot{y} \leq \tilde{\mathcal{V}}(y)$, $\dot{z} = \tilde{\mathcal{V}}(z)$, and $\frac{\partial \tilde{\mathcal{V}}_i}{\partial \chi_j} \geq 0$ for all $j \neq i$. It therefore follows from Theorem A.11 that $y(t) \ll z(t)$ for all $t \geq 0$.

It remains to show that the asymptotic stability of $\mathbf{0} \in \mathcal{X}$ under (A.11) implies the same under (A.8). Since $\mathbf{0}$ is attractive under (A.11), there is $c > 0$ such that the solution trajectory z of (A.11) with initial condition $z(0) = (c, \dots, c)$ converges to $\mathbf{0}$. Thus, Theorem A.11 implies that all solution trajectories y of (A.8) starting in the neighborhood $O_c = \{\chi \in \mathcal{X} : \max_i |\chi_i| < c\}$ of $\mathbf{0}$ satisfy $y(t) \ll z(t)$ for all $t \geq 0$, and so we have $\bigcup_{\chi \in O_c} \omega(\chi) = \{\mathbf{0}\}$. This establishes the attractiveness of $\mathbf{0}$ under (A.8). Lyapunov stability of $\mathbf{0}$ under (A.8) is established in a similar fashion.

This completes the proof of Proposition 5.3.

References

- Aubin, J.-P. and Cellina, A. (1984). *Differential Inclusions*. Springer, Berlin.
- Basu, K. and Weibull, J. W. (1991). Strategy sets closed under rational behavior. *Economics Letters*, 36:141–146.
- Beckmann, M., McGuire, C. B., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven.
- Benaïm, M. (1998). Recursive algorithms, urn processes, and the chaining number of chain recurrent sets. *Ergodic Theory and Dynamical Systems*, 18:53–87.

- Benaïm, M., Hofbauer, J., and Sorin, S. (2005). Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44:328–348.
- Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.
- Binmore, K., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35. Erratum, 21 (1997), 325.
- Conley, C. (1978). *Isolated Invariant Sets and the Morse Index*. American Mathematical Society, Providence, RI.
- Cressman, R. (1997). Local stability of smooth selection dynamics for normal form games. *Mathematical Social Sciences*, 34:1–19.
- Durrett, R. (2005). *Probability: Theory and Examples*. Brooks-Cole, Belmont, CA, third edition.
- Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61:1047–1071.
- Ellison, G. (2000). Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45.
- Ellison, G., Fudenberg, D., and Imhof, L. A. (2009). Random matching in adaptive dynamics. *Games and Economic Behavior*, 66:98–114.
- Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232. Corrigendum, 51 (1997), 77–78.
- Fudenberg, D. and Kreps, D. M. (1993). Learning mixed equilibria. *Games and Economic Behavior*, 5:320–367.
- Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge.
- Galeslout, B. M. and Goyal, S. (1997). Costs of flexibility and equilibrium selection. *Journal of Mathematical Economics*, 28:249–264.
- Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.
- Goyal, S. and Janssen, M. C. W. (1997). Non-exclusive conventions and social coordination. *Journal of Economic Theory*, 77:34–57.
- Hofbauer, J. (1995). Stability for the best response dynamics. Unpublished manuscript, University of Vienna.
- Hofbauer, J. (1999). The spatially dominant equilibrium of a game. *Annals of Operations Research*, 89:233–251.

- Hofbauer, J. (2000). From Nash and Brown to Maynard Smith: Equilibria, dynamics, and ESS. *Selection*, 1:81–88.
- Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294.
- Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69.
- Hofbauer, J. and Sandholm, W. H. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6:341–377.
- Hurkens, S. (1995). Learning by forgetful players. *Games and Economic Behavior*, 11:304–329.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.
- Kaniovski, Y. M. and Young, H. P. (1995). Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior*, 11:330–363.
- Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67.
- Matsui, A. (1991). Cheap-talk and cooperation in a society. *Journal of Economic Theory*, 54:245–258.
- Matsui, A. and Matsuyama, K. (1995). An approach to equilibrium selection. *Journal of Economic Theory*, 65:415–434.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.
- Morris, S., Rob, R., and Shin, H. S. (1995). p -dominance and belief potential. *Econometrica*, 63:145–157.
- Munkres, J. R. (2000). *Topology*. Pearson, New York, second edition.
- Nachbar, J. H. (1990). ‘Evolutionary’ selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19:59–89.
- Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, 88:834–847.
- Osborne, M. J. and Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45:434–441.
- Oyama, D. and Takahashi, S. (2014). Contagion and uninvadability in local interaction games with strategic complementarities. Unpublished manuscript, University of Tokyo and National University of Singapore.

- Oyama, D., Takahashi, S., and Hofbauer, J. (2008). Monotone methods for equilibrium selection under perfect foresight dynamics. *Theoretical Economics*, 3:155–192.
- Oyama, D. and Tercieux, O. (2009). Iterated potential and robustness of equilibria. *Journal of Economic Theory*, 144:1726–1769.
- Roth, G. and Sandholm, W. H. (2013). Stochastic approximations with constant step size and differential inclusions. *SIAM Journal on Control and Optimization*, 51:525–555.
- Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391.
- Sandholm, W. H. (2001). Almost global convergence to p -dominant equilibrium. *International Journal of Game Theory*, 30:107–116.
- Sandholm, W. H. (2010a). Local stability under evolutionary game dynamics. *Theoretical Economics*, 5:27–50.
- Sandholm, W. H. (2010b). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.
- Sandholm, W. H. (2013). Population games and deterministic evolutionary dynamics. Forthcoming in the *Handbook of Game Theory with Economic Applications*, volume 4.
- Selten, R. and Chmura, T. (2008). Stationary concepts for experimental 2×2 -games. *American Economic Review*, 98:938–966.
- Sethi, R. (2000). Stability of equilibria in games with procedurally rational players. *Games and Economic Behavior*, 32:85–104.
- Smirnov, G. V. (2002). *Introduction to the Theory of Differential Inclusions*. American Mathematical Society, Providence, RI.
- Taylor, P. D. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.
- Tercieux, O. (2006). p -best response set. *Journal of Economic Theory*, 131:45–70.
- Walter, W. (1970). *Differential and Integral Inequalities*. Springer-Verlag, Berlin.
- Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.
- Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.