

Ferrer-Alós, Carlos; Netzer, Nick

Working Paper

Robust stochastic stability

Working Paper, No. 63 [rev.]

Provided in Cooperation with:

Department of Economics, University of Zurich

Suggested Citation: Ferrer-Alós, Carlos; Netzer, Nick (2014) : Robust stochastic stability, Working Paper, No. 63 [rev.], University of Zurich, Department of Economics, Zurich, <https://doi.org/10.5167/uzh-59571>

This Version is available at:

<https://hdl.handle.net/10419/111185>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series
ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 63

Robust Stochastic Stability

Carlos Alós-Ferrer and Nick Netzer

Revised version, January 2014

Robust Stochastic Stability*

Carlos Alós-Ferrer[†] and Nick Netzer[‡]

Revised Version
January 2014

Abstract

A strategy profile of a game is called robustly stochastically stable if it is stochastically stable for a given behavioral model independently of the specification of revision opportunities and tie-breaking assumptions in the dynamics. We provide a simple radius-coradius result for robust stochastic stability and examine several applications. For the logit-response dynamics, the selection of potential maximizers is robust for the subclass of supermodular symmetric binary-action games. For the mistakes model, the weaker property of strategic complementarity suffices for robustness in this class of games. We also investigate the robustness of the selection of risk-dominant strategies in coordination games under best-reply and the selection of Walrasian strategies in aggregative games under imitation.

Keywords: learning in games, stochastic stability, radius-coradius theorems, logit-response dynamics, mutations, imitation.

JEL Classification Numbers: C72, D83.

*We are grateful to Larry Blume, Josef Hofbauer, and Jörgen Weibull for helpful discussions. We also thank the participants in the SAET 2011 conference in Faro and the World Congress of the Game Theory Society 2012 for comments on earlier versions of this paper.

[†]Corresponding author. Department of Economics, University of Cologne, Albertus-Magnus Platz, D-50923 Cologne, Germany. E-mail: carlos.alos-ferrer@uni-koeln.de

[‡]Department of Economics, University of Zurich, Bluemlisalpstrasse 10, CH-8006 Zurich, Switzerland. E-mail: nick.netzer@econ.uzh.ch

1 Introduction

The concept of stochastic stability was introduced in game theory in a series of seminal papers by Blume (1993), Kandori, Mailath, and Rob (1993), Young (1993), and Ellison (1993). A large number of applications have since been developed, a series of theoretical improvements have ensued, and several striking results have been proven relying on this concept. Among the best known results, we single out three which have had a lasting impact in the literature. First, the selection of risk-dominant equilibria (even in the presence of alternative, Pareto-efficient ones) in coordination games under best-reply or imitation dynamics (Kandori, Mailath, and Rob, 1993; Kandori and Rob, 1995; Sandholm, 1998). Second, the selection of potential maximizers in exact potential games in logit-response dynamics (Blume, 1993, 1997). Third, the selection of Walrasian equilibria in oligopolies with imitating firms (Vega-Redondo, 1997), which has been shown to generalize to the class of aggregative games (Alós-Ferrer and Ania, 2005). These are all important insights which have shaped our understanding of equilibrium (and non-equilibrium) selection and stability.

The literature has also made a number of weaknesses apparent, some of which have the status of unwritten “folk wisdom”. The main weakness of the stochastic stability literature as a whole is probably the fact that many results might depend, or might be perceived to depend on modeling details, thereby casting doubt on the main insights gained from this approach. A number of failed robustness checks have demonstrated this issue. We would like to argue that, while some of these checks are substantial and have further sharpened our intuition, others have in fact arisen due to a fundamental lack of robustness in the very concept of stochastic stability.

Among the substantial results we count the analysis in Robson and Vega-Redondo (1996), which showed that the selection of risk-dominant equilibria under the imitation dynamics of Kandori, Mailath, and Rob (1993) depends on the postulated interaction structure, with round-robin interaction leading to risk-dominant equilibria but true random matching favoring Pareto-efficient ones (this distinction would not exist if myopic best-reply is assumed). In our opinion, this result does not correspond to a weakness in the general approach. On the contrary, it is a substantial contribution that points at the importance of the interaction structure, and it should not be considered a robustness check. Indeed, the importance of both the interaction structure and the behavioral rule for equilibrium selection has been made apparent in the closely related literature on games in networks (see Weidenholzer, 2010 for a review). For instance, Morris (2000) shows that best-reply dynamics lead to risk-dominant equilibria in quite general networks, while Alós-Ferrer and Weidenholzer (2008) show that imitation favors Pareto-efficient outcomes under comparatively mild conditions on the network.

Among the worrying failed robustness tests we count the fact that changing the specification of either revision opportunities or tie-breaking assumptions might sometimes influence the long-run outcomes in a given dynamics. This affects, for instance, the well-known result that the original logit dynamics of Blume (1993, 1997) selects potential maximizers in exact potential games. Alós-Ferrer and Netzer (2010) have shown that this result depends crucially on the assumption of asynchronous learning, that is, a dynamic specification in which every period one and only one agent is selected and allowed to revise his or her strategy, while all other players are required to stay put. The result vanishes if more general revision processes are allowed for.

Tie-breaking assumptions are also not always harmless. Suppose that a behavioral rule fully specifies the set of strategies that a player might choose from, e.g. the set of payoff maximizers given other players' strategies (as in the case of myopic best-reply) or the set of strategies leading to currently maximal, observed payoffs (as in the case of imitation). Even abstracting from revision opportunities, this still does not fully specify the dynamics. One might for instance require that all optimal strategies be chosen with positive (maybe equal) probability; it might, however, be equally reasonable to postulate that players who are already employing one of the optimal strategies do not switch away. These are all reasonable choices, which sometimes have consequences for the dynamic analysis (contrast e.g. Oechssler, 1997 and Alós-Ferrer, 2003; see also Sandholm, 1998).

The possible dependence of long-run predictions on the specification of revision opportunities and tie-breaking assumptions is an important consideration. A result which depends on such modeling details, which can be argued to be orthogonal to the behavioral rule and the interaction structure, should not be considered to be on equal grounds with a result which is immune to the specification thereof. In the present research, we aim to provide and apply a simple result which helps establish when a long-run prediction is robust to the specification of revision opportunities and/or tie-breaking assumptions.

We concentrate on these modeling specification issues because they are the ones which “should not matter”, and hence robustness failures are particularly worrying. The issue of robustness, however, is more general. One of the early criticisms on the literature, due to Bergin and Lipman (1996), was that results might change if the perturbations (mistakes or mutations) which are used to define stochastic stability can be defined in a state-dependent way. By specifying arbitrary asymmetric perturbations it is possible to stabilize any outcome. The question is then rather whether results are robust to the specification of perturbations within a reasonable class. An excellent example of such a robustness analysis is van Damme and Weibull (2002), which showed the robustness of the results of Kandori, Mailath, and Rob (1993) if mutation rates are endogenously determined, assuming that players can effortfully control the probability of implementing their intended strategies. Further, for specific dynamics, additional

dimensions not considered here might be natural candidates for a robustness analysis. Consider the case of imitation with memory (e.g. Alós-Ferrer, 2004, 2008; Bergin and Bernhardt, 2009; Alós-Ferrer and Shi, 2012). A nice example of robustness analysis within this category is Josephson and Matros (2004) which considers imitation dynamics for multiple populations analogous to Young (1993) and derives selection results which are robust to the specification of memory length and information sample size.

The paper is structured as follows. Section 2 describes the general framework for the analysis. Section 3 introduces the robustness concept and presents the main result. Section 4 analyzes symmetric binary-action games both for the logit-response dynamics and for the popular best-reply mistakes model. Section 5 investigates evolutionary stability and aggregate-taking behavior for perturbed imitation dynamics. Section 6 concludes.

2 Learning in Games: A General Framework

2.1 Stage Model

Consider a finite population of N agents who repeatedly interact in discrete time $t = 1, 2, \dots$ according to a pre-specified stage model, formalized as a finite, normal-form game $\Gamma = (I, (S_i, u_i)_{i \in I})$, where $I = \{1, 2, \dots, N\}$ is the set of players, S_i are the pure strategy sets, and $u_i : S \rightarrow \mathbb{R}$ are payoff functions, where $S = \prod_{i \in I} S_i$. We let $S_{-i} = \prod_{j \neq i} S_j$ be the set of strategy profiles of all players except i , and write $s = (s_i, s_{-i})$ and $u_i(s_i, s_{-i})$. The strategies chosen and the stage model determine the payoffs agents receive at the end of the period t . The stage model can simply be taken to be an arbitrary, asymmetric N -player game, as in Blume (1993) or Alós-Ferrer and Netzer (2010), or it can incorporate additional structure. For example, it might specify that agents play a bilateral finite game sequentially against each other agent in the population (round robin tournament), as in Kandori, Mailath, and Rob (1993) (hereafter KMR), where the bilateral game is a symmetric 2×2 coordination game.

2.2 Behavioral Rules and Correspondences, and Tie-Breaking

The game is played by boundedly rational players whose behavior is summarized by *behavioral rules*. At the beginning of each period, a certain subset of agents is chosen to update their actions. Each of them then chooses a pure strategy according to a pre-specified behavioral rule $B_i : S \mapsto \Delta S_i$. That is, $B_i(s)(s'_i)$ is the probability with which player i will choose strategy $s'_i \in S_i$ after the profile $s \in S$ has been played. A simple behavioral rule which has been extensively studied in the literature of learning in games is the *myopic best-reply* rule, where players are assumed to be able to compute best-replies to the current profile of strategies of their opponents, and choose one of them. In games with alternative best-replies, the need for tie-breaking gives rise to a

family of rules. That is, a rule B_i^{BR} is a best-reply rule if

$$B_i^{BR}(s)(s'_i) > 0 \implies u_i(s'_i, s_{-i}) \geq u_i(s''_i, s_{-i}) \forall s''_i \in S_i. \quad (1)$$

Let us call Γ a symmetric game if $S_i = S_j = S_0$ for all $i, j \in I$ and payoffs are given by a symmetric mapping, i.e., the payoff of a player choosing strategy s_i against the profile of strategies s_{-i} is $u_i(s_i, s_{-i}) = u(s_i | s_{-i})$, where the latter is invariant to permutations of s_{-i} . For symmetric games, a second prominent example of behavioral rule (or, rather, family thereof) is given by *imitate-the-best* rules as in KMR, Vega-Redondo (1997), or Alós-Ferrer and Ania (2005), where players just adopt one of the strategies leading to the highest, currently observed payoff. That is, again taking into account the need for tie-breaking assumptions, a rule B_i^{IB} is an imitate-the-best rule if

$$B_i^{IB}(s)(s'_i) > 0 \implies \begin{aligned} & s'_i = s_j \text{ for some } j \in I \text{ with} \\ & u_j(s) \geq u_k(s) \forall k \in I. \end{aligned} \quad (2)$$

The description of both the best-reply and the imitation rules allows for different tie-breaking assumptions. We will now provide a formal approach to their specification.¹ A *behavioral correspondence* for player i is a correspondence $\widehat{B}_i : S \rightrightarrows S_i$. That is, $\widehat{B}_i(s)$ is the set of strategies $s'_i \in S_i$ which player i *might* choose after the profile $s \in S$ has been played.² A behavioral rule B_i is said to *agree* with correspondence \widehat{B}_i if

$$B_i(s)(s'_i) > 0 \implies s'_i \in \widehat{B}_i(s) \quad (3)$$

for all $s'_i \in S_i$ and all $s \in S$. For instance, myopic best-reply rules as in (1) are those agreeing with the best-reply correspondence

$$\widehat{B}_i^{BR}(s) = \{s'_i \in S_i \mid u_i(s'_i, s_{-i}) \geq u_i(s''_i, s_{-i}) \forall s''_i \in S_i\}.$$

Imitate-the-best rules as in (2) are those agreeing with the imitate-the-best correspondence

$$\widehat{B}_i^{IB}(s) = \{s'_i \in S_i \mid s'_i = s_j \text{ for some } j \in I \text{ with } u_j(s) \geq u_k(s) \forall k \in I\}.$$

Given a behavioral correspondence \widehat{B}_i , we say that a behavioral rule B_i agreeing with

¹Note that a rule's definition might formally rely on the payoff functions, but the interpretation of the actual knowledge of the game that players have might be very different. Under best-reply, players do know the payoff function and can use it to (myopically) optimize their behavior. For imitation rules, the use of the payoff function is just a modeling device capturing the informational assumption that players observe realized payoffs, but do not necessarily know the game or are able to perform optimizing computations.

²Ritzberger and Weibull (1995) introduce the term *behavior correspondence* in a different context. Apart from technical issues (as e.g. whether they are defined on profiles of pure or mixed strategies), the main difference is that we do not require behavioral correspondences to be extensions of the best-reply correspondence.

\widehat{B}_i is *balanced* if $B_i(s)(s_i) > 0$ whenever $s_i \in \widehat{B}_i(s)$, where s_i is player i 's strategy in the profile s , and $B_i(s)(s'_i) > 0$ for all $s'_i \in \widehat{B}_i(s)$ whenever $s_i \notin \widehat{B}_i(s)$. With a balanced behavioral rule, players who find their current behavior optimal according to the behavioral correspondence will not abandon it for sure (although they might also not stick to it for sure). If the current behavior is not optimal, then balanced rules respect anonymity of the strategies in the sense that they place positive (although not necessarily identical) probability on all optimal strategies.³ Let \mathcal{T}_i denote the set of all balanced behavioral rules that agree with a given correspondence \widehat{B}_i . (We suppress dependency of \mathcal{T}_i on \widehat{B}_i for notational convenience.)

Now consider two balanced behavioral rules B_i and B'_i from \mathcal{T}_i . We say that B'_i is (weakly) *swifter* than B_i , written $B_i \preceq B'_i$, if $B_i(s)(s'_i) > 0$ implies $B'_i(s)(s'_i) > 0$, for all $s'_i \in S_i$ and all $s \in S$. That is, the support of B'_i is always weakly larger than the support of B_i . We say that the two rules are *equally swift*, written $B_i \simeq B'_i$, if $B_i \preceq B'_i$ and $B'_i \preceq B_i$, so that the sets $\{s'_i \in S_i \mid B_i(s)(s'_i) > 0\}$ and $\{s'_i \in S_i \mid B'_i(s)(s'_i) > 0\}$ always coincide. By construction, the relation \simeq is a binary equivalence relation on \mathcal{T}_i . In the following, we will informally identify two behavioral rules if they are equally swift, i.e., if they differ in specific probabilities assigned to strategies, but not in their support. Formally, we work in the quotient set \mathcal{T}_i / \simeq , on which the swiftness relation \preceq becomes a partial order.

Among all rules in \mathcal{T}_i , we consider two distinguished rules (modulo equal swiftness). The *cautious tie-breaking rule* B_i^0 is the balanced rule specifying $B_i^0(s)(s_i) = 1$ whenever $s_i \in \widehat{B}_i(s)$. That is, under the cautious rule a player always sticks to his or her current action if it is one of the optimal ones according to the behavioral correspondence \widehat{B}_i . The *random tie-breaking rule* B_i^X is the rule given by $B_i^X(s)(s'_i) > 0$ for all $s'_i \in \widehat{B}_i(s)$, that is, all strategies that are optimal according to \widehat{B}_i are always chosen with strictly positive probability. The following observation is now immediate.

Lemma 1. *Any balanced behavioral rule B_i satisfies $B_i^0 \preceq B_i \preceq B_i^X$.*

That is, the poset \mathcal{T}_i / \simeq has a top and a bottom element. It is straightforward to show that it is actually a complete lattice. We denote profiles of balanced behavioral rules for all players by $B = (B_i)_{i \in I} \in \mathcal{T} := \prod_{i \in I} \mathcal{T}_i$. Consider the product order on \mathcal{T} , i.e., $B \preceq B'$ if and only if B'_i is weakly swifter than B_i for all $i \in I$. Then we also obtain $B^0 \preceq B \preceq B^X$ for any balanced profile B and the two extreme profiles $B^0 = (B_i^0)_{i \in I}$ and $B^X = (B_i^X)_{i \in I}$.

³Lexicographic tie-breaking conditions as e.g. choosing the most popular action in case of ties would violate balancedness. They could, however, also be built into the behavioral correspondence, and the modified rules would then be balanced by virtue of a “tie” meaning then both equal payoffs and identical popularity.

2.3 Revision Opportunities

A learning dynamics for a game Γ is made of a behavioral rule for each player, which includes tie-breaking assumptions, and a specification of revision opportunities, i.e., a way of determining which players receive the opportunity to update their actions in a given period. Intuitively, revision opportunities are closely related to the speed of the dynamics. A dynamics where only one agent is allowed to revise per period is more gradual than one where the whole population might switch away simultaneously, enabling abrupt transition phenomena. Following Alós-Ferrer and Netzer (2010), a *revision process* is a probability measure q on the set of all subsets of I , with the property that

$$\forall i \in I, \exists J \subseteq I \text{ such that } i \in J \text{ and } q(J) > 0, \quad (4)$$

where $q(J)$ (also written q_J) is interpreted as the probability that exactly the players in J receive a revision opportunity (independently across periods). Condition (4) implies that each player gets the opportunity to revise with strictly positive probability. A revision process is called *regular* if $q_i = q_{\{i\}} > 0$ for all $i \in I$, so that for each player there is a strictly positive probability of being the only player who is allowed to revise. Let \mathcal{Q} denote the set of all regular revision processes.

Analogously to the previous subsection, we can define a binary relation \preceq on \mathcal{Q} as follows.⁴ For any $q, q' \in \mathcal{Q}$ we say that q' is (weakly) *quicker* than q , written $q \preceq q'$, if $q_J > 0$ implies $q'_J > 0$, for all $J \subseteq I$. That is, the revision process q' includes more possibilities than q . We say that q and q' have *the same speed*, written $q \simeq q'$, if $q \preceq q'$ and $q' \preceq q$. By construction, the relation \simeq is a binary equivalence relation. Consider again the quotient set \mathcal{Q}/\simeq , where two revision processes belong to the same class if and only if they have the same speed, i.e., they differ in specific probabilities assigned to player subsets but not in their support. We will again identify two processes that have the same speed and treat \preceq as a partial order.

Among all processes in \mathcal{Q} , we again consider two distinguished elements (modulo equal speed). The *asynchronous learning* process q^{AL} satisfies $q_J^{AL} = 0$ whenever $|J| \geq 2$. The *independent learning* process q^{IL} satisfies $q_J^{IL} > 0$ for all $J \subseteq I$.⁵ The following observation is now again immediate.

Lemma 2. *Any regular revision process q satisfies $q^{AL} \preceq q \preceq q^{IL}$.*

Therefore the poset \mathcal{Q}/\simeq has a top and a bottom element as well. It is again a simple exercise to show that it is actually a complete lattice.

⁴We use the same symbol for the binary relations on \mathcal{T} and on \mathcal{Q} for convenience.

⁵These concepts are again taken from Alós-Ferrer and Netzer (2010). The model of Blume (1993) postulates $q_i = 1/N$ and is therefore an instance of asynchronous learning. Independent inertia as in Sandholm (1998), where $q_J = p^{|J|} (1-p)^{N-|J|}$ for some $0 < p < 1$, is an instance of independent learning. The simultaneous learning process, where $q_I = 1$, is the simplest example of a process which is not regular.

2.4 Stochastic Stability

Given a profile of behavioral rules B and a revision process q , the *unperturbed dynamics* (B, q) induces a Markov chain on the state space S . An *absorbing set* of (B, q) is a minimal set of states with respect to the property that it cannot be left with positive probability under (B, q) . An *absorbing state* is a state $s \in S$ such that the singleton set $\{s\}$ is absorbing. The fact that unperturbed learning dynamics will typically have multiple absorbing sets or states motivates the introduction of noise to select between them.

We can apply a *noise process* to B to derive associated profiles of behavioral rules with noise $B^\varepsilon = (B_i^\varepsilon)_{i \in I}$, where $\varepsilon \in (0, 1)$ measures how strongly the players' behavior is perturbed from B .⁶ For the first noise process that we consider, the *mistakes model*, we fix a noise rule $E_i : S \mapsto \Delta S_i$ for every player $i \in I$, where $E_i(s)(s'_i)$ is independent of s and satisfies $E_i(s)(s'_i) > 0$ for all $s'_i \in S_i$. Then each player's behavioral rule B_i is perturbed to $B_i^{M, \varepsilon}$ by

$$B_i^{M, \varepsilon}(s)(s'_i) = (1 - \varepsilon)B_i(s)(s'_i) + \varepsilon E_i(s)(s'_i). \quad (5)$$

For instance, the best-reply version of the well-known KMR model, first studied in Kandori and Rob (1995), proceeds exactly like this to derive the best-reply with mistakes $B_i^{BR, M, \varepsilon}$ from an unperturbed best-reply rule B_i^{BR} . As $\varepsilon \rightarrow 0$, behavior converges to the best-reply rule. The noisy version $B_i^{IB, M, \varepsilon}$ of an imitate-the-best rule B_i^{IB} can be constructed analogously.⁷ Importantly, the tie-breaking assumptions implicit in B_i^{BR} or B_i^{IB} carry over to the noisy rules when the mistakes approach is used. When we start from a behavioral correspondence such as \hat{B}_i^{BR} or \hat{B}_i^{IB} , for instance, the mistakes model associates to every behavioral rule $B_i \in \mathcal{T}_i$ a distinct behavioral rule with noise B_i^ε , which converges to B_i as $\varepsilon \rightarrow 0$. We say that noise processes with this property *respect tie-breaking*. The second noise process that we will consider is the *logit choice function*, which has been used in the literature to obtain noisy versions of the best-reply dynamics (see e.g. Blume, 1993 or Alós-Ferrer and Netzer, 2010). Formally, the probability of player i choosing s'_i is given by

$$B_i^{BR, L, \varepsilon}(s)(s'_i) = \frac{e^{(1/\varepsilon)u_i(s'_i, s_{-i})}}{\sum_{s''_i \in S_i} e^{(1/\varepsilon)u_i(s''_i, s_{-i})}}. \quad (6)$$

Again, all actions are chosen with strictly positive probability whenever $\varepsilon > 0$ and choice concentrates on myopic best-replies as $\varepsilon \rightarrow 0$. The logit perturbation, however, leaves no freedom in tie-breaking assumptions. As $\varepsilon \rightarrow 0$, the behavioral rule $B_i^{BR, L, \varepsilon}$

⁶See Bergin and Lipman (1996) for a general treatment of noise processes.

⁷The original KMR model can be readily interpreted as a model of imitation (see KMR p. 31, Rhode and Stegeman, 1996, and Sandholm, 1998) where agents mimic the actions which led to the highest payoffs in the last period.

converges to the specific best-reply rule that breaks ties with equal probabilities. The logit approach is therefore not suited to associate a distinct noisy rule to every $B_i \in \mathcal{T}_i$ for a given behavioral correspondence \hat{B}_i . By using logit choice, we rather select a specific $B_i \in \mathcal{T}_i$, the one with equal tie-breaking.⁸ We say that noise processes with this property *impose tie-breaking*. Other particular examples of noise processes could also be considered. For instance, Myatt and Wallace (2003) and Dokumaci and Sandholm (2008) consider dynamics based on probit choice functions, which, as in the case of logit, impose tie-breaking. Sandholm (2010) considers general “noisy revision protocols” (where the word revision is used in a different sense than in this paper) including the mistakes model and logit and probit choice. The two prominent examples presented above are those for which we develop specific applications later.

Now consider any *perturbed dynamics* (B^ε, q) derived from an unperturbed dynamics (B, q) according to some noise process. Suppose that all B_i^ε have full support whenever $0 < \varepsilon < 1$, as in the examples above. Then, the perturbed dynamics induces an irreducible and aperiodic Markov chain on the state space S with transition probabilities

$$P_{s,s'}^\varepsilon = \sum_{J \subseteq I | s_k = s'_k \forall k \in I \setminus J} q_J \prod_{i \in J} B_i^\varepsilon(s)(s'_i), \quad (7)$$

and it has a unique invariant distribution, denoted μ^ε . A strategy profile or state $s \in S$ is *stochastically stable for* (B^ε, q) if $\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(s) > 0$. Stochastic stability for the mistakes model can be characterized along the lines introduced in KMR or Young (1993), with a useful “radius-coradius” sufficient condition developed by Ellison (2000). Alós-Ferrer and Netzer (2010) provide an analogous general characterization for the logit-response dynamics, and a similar radius-coradius result.⁹

3 Robustness

3.1 Definitions

We are interested in the following two concepts of robustness. First, suppose we consider a given profile of behavioral rules with noise B^ε , based on some underlying profile of unperturbed behavioral rules B . Hence we treat as fixed a specification of tie-breaking assumptions. This is always the case when the noise process imposes tie-breaking, as with the logit-response dynamics, but it can be done for any behavioral rule and noise process as detailed above. Robustness now refers to the specification of revision opportunities alone.

⁸One could also study the logit perturbation $B_i^{IB,L,\varepsilon}$ of an imitate-the-best dynamics B_i^{IB} , which would converge to the equal tie-breaking imitation rule as noise vanishes.

⁹Several earlier contributions have studied logit behavior for special classes of games or dynamics (e.g. Blume, 1993, 1997; Maruta, 2002; Myatt and Wallace, 2008a,b). See Beggs (2005) for radius-coradius results with a focus on waiting times.

Definition 1. Fix a profile of behavioral rules with noise B^ε . A state $s \in S$ is *robustly stochastically stable with respect to revision processes* (RSS-R) if it is stochastically stable for (B^ε, q) for any regular revision process $q \in \mathcal{Q}$.

Second, suppose we consider a profile of behavioral correspondences $\hat{B} = (\hat{B}_i)_{i \in I}$, with \mathcal{T} being the set of profiles of balanced behavioral rules that agree player-wise with \hat{B} . For each $B \in \mathcal{T}$ we then apply a noise process that respects tie-breaking to associate a profile with noise B^ε . Robustness then refers to the specification of both tie-breaking assumptions and revision opportunities.

Definition 2. Fix a profile of behavioral correspondences \hat{B} and a noise process that respects tie-breaking. A state $s \in S$ is *robustly stochastically stable with respect to revision processes and tie-breaking* (RSS-RT) if it is stochastically stable for (B^ε, q) for any $q \in \mathcal{Q}$ and any B^ε that the noise process associates to some element of \mathcal{T} .

We aim to provide a method allowing us to identify RSS-R and RSS-RT states based on the concepts of radius and coradius. We first introduce the following auxiliary concept.

Definition 3. An *operator on revision processes* is a mapping

$$\begin{aligned} T : \mathcal{Q} \times S &\mapsto \mathbb{R} \\ (q, s) &\mapsto T^q(s). \end{aligned}$$

Given some $s \in S$, the operator T is *monotone for s* if $T^q(s) \geq T^{q'}(s)$ whenever $q \preceq q'$. It is *monotone* if it is monotone for all $s \in S$.

In contrast to the usual approach, we will not define radius and coradius from a primitive such as cost (Ellison, 2000) or waste (Alós-Ferrer and Netzer, 2010). The only property of the different radius and coradius concepts that we need in the following is that they are monotone operators (for the candidate prediction) that yield sufficient conditions for stochastic stability.

Definition 4. Fix a profile of behavioral rules with noise B^ε . A *radius-coradius pair* (R, CR) for B^ε is a pair of operators on revision processes such that, whenever $R^q(s) > CR^q(s)$ for some absorbing state $s \in S$ of the unperturbed dynamics (B, q) , it follows that s is the unique stochastically stable state for (B^ε, q) . Given some $s \in S$, a radius-coradius pair is *monotone for s* if both operators are monotone for s .

As we will show below, the existing radius and coradius concepts of Ellison (2000) and Alós-Ferrer and Netzer (2010) are monotone operators under minimal conditions. Once monotonicity is established, the fact that the property embodied in the definition above is fulfilled follows from the radius-coradius theorems in Ellison (2000) and Alós-Ferrer and Netzer (2010).

Analogously, we can define (monotone) operators and radius-coradius pairs with respect to both revision opportunities and tie-breaking rules as follows. Fix a profile of behavioral correspondences \widehat{B} , which induces the set \mathcal{T} , and consider the product order on $\mathcal{T} \times \mathcal{Q}$, i.e., $(B, q) \preceq (B', q')$ if and only if $B \preceq B'$ and $q \preceq q'$.

Definition 5. An operator on revision processes and tie-breaking rules is a mapping

$$\begin{aligned} T : \mathcal{T} \times \mathcal{Q} \times S &\mapsto \mathbb{R} \\ (B, q, s) &\mapsto T^{B,q}(s). \end{aligned}$$

Given some $s \in S$, the operator T is *monotone for s* if $T^{B,q}(s) \geq T^{B',q'}(s)$ whenever $(B, q) \preceq (B', q')$. It is *monotone* if it is monotone for all $s \in S$.

Definition 6. Fix a profile of behavioral correspondences \widehat{B} and a noise process that respects tie-breaking. A *radius-coradius pair* (R, CR) for \widehat{B} and the noise process is a pair of operators on revision processes and tie-breaking rules such that, whenever $R^{B,q}(s) > CR^{B,q}(s)$ for some absorbing state $s \in S$ of the unperturbed dynamics (B, q) , it follows that s is the unique stochastically stable state for (B^ϵ, q) . Given some $s \in S$, a radius-coradius pair is *monotone for s* if both operators are monotone for s .

3.2 Results

The following proposition embodies the main idea behind our results.

Proposition 1. (i) Fix a profile of behavioral rules with noise B^ϵ . Let (R, CR) be a radius-coradius pair for B^ϵ . Let $q^1, q^2 \in \mathcal{Q}$ with $q^1 \preceq q^2$. If there exists an absorbing state $s \in S$ of (B, q^2) such that (R, CR) is monotone for s and

$$R^{q^2}(s) > CR^{q^1}(s),$$

then s is the unique stochastically stable state for any (B^ϵ, q) with $q^1 \preceq q \preceq q^2$.

(ii) Fix a profile of behavioral correspondences \widehat{B} and a noise process that respects tie-breaking. Let (R, CR) be a radius-coradius pair for \widehat{B} and the noise process. Let $(B^1, q^1), (B^2, q^2) \in \mathcal{T} \times \mathcal{Q}$ with $(B^1, q^1) \preceq (B^2, q^2)$. If there exists an absorbing state $s \in S$ of (B^2, q^2) such that (R, CR) is monotone for s and

$$R^{B^2, q^2}(s) > CR^{B^1, q^1}(s),$$

then s is the unique stochastically stable state for any (B^ϵ, q) with $(B^1, q^1) \preceq (B, q) \preceq (B^2, q^2)$.

Proof. We prove statement (ii). Statement (i) is proven analogously. Consider an arbitrary $(B, q) \in \mathcal{T} \times \mathcal{Q}$ with $(B^1, q^1) \preceq (B, q) \preceq (B^2, q^2)$. Since s is an absorbing state of (B^2, q^2) , it is also an absorbing state of (B, q) , because the latter (unperturbed)

dynamics enables fewer positive-probability transitions than (B^2, q^2) . It then suffices to notice that, by monotonicity,

$$R^{B,q}(s) \geq R^{B^2,q^2}(s) > CR^{B^1,q^1}(s) \geq CR^{B,q}(s),$$

which implies the statement by definition of radius-coradius pair. \square

Using part (i) of this result and Lemma 2, we obtain an immediate corollary which delivers a simple condition for RSS-R.

Corollary 1. *Fix a profile of behavioral rules with noise B^ε . Let (R, CR) be a radius-coradius pair for B^ε . If there exists an absorbing state $s \in S$ of (B, q^{IL}) such that (R, CR) is monotone for s and*

$$R^{q^{IL}}(s) > CR^{q^{AL}}(s),$$

then s is the unique RSS-R state.

This corollary applies directly to the logit-response dynamics or to any mistakes model for pre-specified tie-breaking assumptions. The result states that establishing robust stochastic stability is just as simple (or just as complex) as establishing stochastic stability with the help of a radius-coradius result. The only difference is that one must focus on asynchronous learning for computing the coradius and on independent learning for computing the radius.

Using part (ii) of Proposition 1 and Lemmata 1 and 2, we also obtain an immediate corollary about robustness with respect to tie-breaking rules in addition to revision processes.

Corollary 2. *Fix a profile of behavioral correspondences \hat{B} and a noise process that respects tie-breaking. Let (R, CR) be a radius-coradius pair for \hat{B} and the noise process. If there exists an absorbing state $s \in S$ of (B^X, q^{IL}) such that (R, CR) is monotone for s and*

$$R^{B^X, q^{IL}}(s) > CR^{B^0, q^{AL}}(s),$$

then s is the unique RSS-RT state.

Hence, even when we require robustness to cover both revision processes and tie-breaking assumptions, a radius-coradius result applies. Again we need to focus on two different, focal dynamic specifications only: independent learning with random tie-breaking, and asynchronous learning with cautious tie-breaking.

3.3 Applying the Results

In order to apply the results, one needs to identify a radius-coradius pair and guarantee monotonicity. This is easy to do for the existing concepts of Ellison (2000) for the

mistakes model and Alós-Ferrer and Netzer (2010) for the logit-response dynamics.

The coradius of a state s due to Ellison (2000) is the maximum across states $s' \neq s$ of the minimum number of mistakes required to enable a positive-probability transition from s' to s in the unperturbed dynamics (B, q) .¹⁰ A mistake here is a strategy choice of agent i that has zero probability under rule B_i . If an unperturbed dynamics includes more positive-probability transitions, then the minima can only weakly decrease and the maximum among all the minima can only be weakly smaller than before. The reasoning for logit-response is analogous, with the number of mistakes replaced by the utility differences between the chosen actions and the myopically optimal ones. This shows the following.

Proposition 2. *The coradius concept of Ellison (2000) is a monotone operator on revision processes and tie-breaking rules. The coradius concept of Alós-Ferrer and Netzer (2010) is a monotone operator on revision processes.*

The radius of a state s due to Ellison (2000) is the minimal number of mistakes needed to enable a positive-probability transition from s to some absorbing set not containing s . Again, the concept for logit-response is analogous, with the number of mistakes replaced by the appropriate utility differences. In general, the radius will also be a monotone operator. The intuition is again that, if one considers an unperturbed dynamics that includes more positive-probability transitions, then the path which previously realized the radius is still feasible and associated with a weakly smaller number of mistakes (or utility losses). Hence the minimum can only become weakly smaller. The only problem with this argument is that the absorbing set reached under the initial dynamics might cease to be absorbing under a dynamics enabling more transitions. In some cases, this might create difficulties. Hence, we provide a result delivering a sufficient condition for monotonicity. We formulate it only for the radius of Ellison (2000). The analogous result also holds for the radius of the logit-response dynamics (where only revision opportunities can be varied).

Proposition 3. *Fix a profile of behavioral correspondences \hat{B} and let $(B^1, q^1), (B^2, q^2) \in \mathcal{T} \times \mathcal{Q}$ with $(B^1, q^1) \preceq (B^2, q^2)$. Suppose there exists an absorbing state $s \in S$ of (B^2, q^2) such that*

- (α) *in every absorbing set $S' \neq \{s\}$ of (B^1, q^1) there exists a state s' such that the probability of reaching s from s' under (B^2, q^2) is zero.*

Then the radius concept of Ellison (2000) is a monotone operator for s on the set of all (B, q) with $(B^1, q^1) \preceq (B, q) \preceq (B^2, q^2)$.

Proof. Say that $s'' \in S$ is (B, q) -reachable from $s' \in S$ if there exists a positive-probability path from s' to s'' in dynamics (B, q) . Note that if $(B, q) \preceq (B', q')$ and s''

¹⁰Radius and coradius are often defined only for (the elements of) absorbing sets. The difference is inconsequential for our purposes.

is (B, q) -reachable from s' , then s'' is also (B', q') -reachable from s' , because the latter dynamics allows for more positive-probability paths. As an immediate consequence, s is an absorbing state of any dynamics (B, q) with $(B, q) \preceq (B^2, q^2)$.

Fix any $(B, q), (B', q') \in \mathcal{T} \times \mathcal{Q}$ with $(B^1, q^1) \preceq (B, q) \preceq (B', q') \preceq (B^2, q^2)$. Let $R^{B, q}(s) = m$. We aim to prove that $R^{B', q'}(s) \leq m$. Consider a path for (B, q) that realizes the radius, starting at s and reaching some s' in an absorbing set S' of (B, q) with $s \notin S'$. For (B', q') , this path is still feasible and associated with a number of mistakes weakly smaller than m . The set S' , however, is not necessarily an absorbing set of (B', q') . However, since states outside of S' are not (B, q) -reachable from states in S' , they are also not (B^1, q^1) -reachable from states in S' . It follows that there exists an absorbing set S'' of (B^1, q^1) with $S'' \subseteq S'$ and an $s'' \in S''$ which is (B^1, q^1) -reachable, hence (B', q') -reachable, from s' . By (α) , there is an $s''' \in S''$ such that s is not (B^2, q^2) -reachable, and hence not (B', q') -reachable, from s''' . Since $s'', s''' \in S''$, which is an absorbing set of (B^1, q^1) , s''' is (B^1, q^1) -reachable, hence (B', q') -reachable, from s'' . Since s is not (B', q') -reachable from s''' (and the state space is finite) it follows that there exists some absorbing set S^* of (B', q') with $s \notin S^*$ which is (B', q') -reachable from s''' . Connecting the paths from s to s' , from the latter to s'' , from the latter to s''' , and from the latter to S^* , we obtain a path from s to an absorbing set of (B', q') not containing s , with a number of mistakes weakly smaller than m for (B', q') . Hence, $R^{B', q'}(s) \leq m$. \square

Condition (α) in the proposition merely captures the idea that absorbing sets should not become fully embodied in the “basin of attraction” of s as one moves from the less quick/swift dynamics to the quickest/swiftest one. A sufficient condition fulfilled in many applications is simply that absorbing sets do not change across dynamics.

Remark 1. Monotonicity of the radius and coradius operators, either in general or just for the appropriate state, is of course just a sufficient condition for our results to hold. One could replace them by even weaker conditions. For instance, as can be seen from its proof, Proposition 1 also holds if we replace the requirement that the radius be a monotone operator by the condition that the radius of s under the upper bound dynamics should be weakly smaller than the radius of s under any of the other dynamics. For Corollary 1, this translates into $R^q(s) \geq R^{q^{IL}}(s)$ for all q with $q^{AL} \preceq q \preceq q^{IL}$. For Corollary 2, the condition becomes $R^{B, q}(s) \geq R^{B^X, q^{IL}}(s)$ for all (B, q) with $(B^0, q^{AL}) \preceq (B, q) \preceq (B^X, q^{IL})$. These weaker conditions are harder to verify in general. As we will see, however, for some particular cases they are still applicable.

4 Symmetric Binary Action Games

4.1 Notation and Definitions

Let Γ be a symmetric binary action game (see e.g. Kim, 1996; Maruta, 2002; Blume, 2003) with strategy set $S_0 = \{A, B\}$. Symmetry implies that each player's payoff depends only on the own action and on the number of opponents choosing each action.¹¹ Given a strategy profile $s \in S$, denote by $m(s)$ the number of players choosing A in s . Let $\pi^A(n)$ be the payoff of an A -player if n players choose action A (including the respective player herself) and let $\pi^B(n)$ be the payoff of a B -player if n players choose A . We can then write the payoff functions as $u_i(s_i, s_{-i}) = \pi^{s_i}(m(s_i, s_{-i}))$. Furthermore, we define $\Delta(n) = \pi^A(n) - \pi^B(n-1)$ for $1 \leq n \leq N$ as the payoff change of a player who switches from action B to action A , given that $n-1$ of the opponents choose action A so that the overall number of A -players is n after the switch. We consider two examples.

Example 1. Consider a unanimity game (e.g. Young, 1998a, Section 9) where $\pi^B(0) > 0$ and $\pi^A(N) > 0$, but $\pi^A(n) = 0$ if $n < N$ and $\pi^B(n) = 0$ if $n > 0$. The game has two strict Nash equilibria, $\mathbf{A} = (A, \dots, A)$ and $\mathbf{B} = (B, \dots, B)$. In addition, every profile $s \in S$ with $2 \leq m(s) \leq N-2$ is a non-strict Nash equilibrium. The difference function $\Delta(n)$ of the unanimity game is given by $\Delta(1) = -\pi^B(0)$, $\Delta(n) = 0$ for all $2 \leq n \leq N-1$, and $\Delta(N) = \pi^A(N)$.

Example 2. The unanimity game can be generalized in different ways. As a particularly interesting example for our purpose, consider a team project game with two projects, A and B , where each of the N players must participate in exactly one of the projects. Participation is costless, but the success of project A requires the participation of at least $1 \leq n_A \leq N$ players, while project B is successful if at least $1 \leq n_B \leq N$ players participate. Assume further that $n_A + n_B > N + 1$, which implies that the two projects cannot be realized jointly and that there is the possibility that none of them is successful. If project A (B) is successful, it generates an overall benefit of size $a > 0$ ($b > 0$), which is distributed equally among all participating players. Players who do not participate in a successful project obtain a payoff of zero. Hence payoffs are

$$\pi^A(n) = \begin{cases} a/n & \text{if } n \geq n_A, \\ 0 & \text{if } n < n_A, \end{cases} \quad \pi^B(n) = \begin{cases} 0 & \text{if } n > N - n_B, \\ b/(N - n) & \text{if } n \leq N - n_B. \end{cases}$$

The two monomorphic profiles \mathbf{A} and \mathbf{B} are again strict Nash equilibria, and profiles $s \in S$ with $N - n_B + 2 \leq m(s) \leq n_A - 2$ are non-strict Nash equilibria. We obtain the

¹¹Sandholm (2010) also considers symmetric binary action games, concentrating on the asymptotics as noise vanishes and population size goes to infinity. Staudigl (2012) follows the same approach for asymmetric binary action games.

difference function

$$\Delta(n) = \begin{cases} a/n & \text{if } n_A \leq n \leq N, \\ 0 & \text{if } N - n_B + 2 \leq n \leq n_A - 1, \\ -b/(N - n + 1) & \text{if } 1 \leq n \leq N - n_B + 1. \end{cases}$$

The team project game becomes the unanimity game if $n_A = n_B = N$.¹²

As shown by Hofbauer and Sorger (2002, p. 27), every symmetric binary action game is an exact potential game in the sense of Monderer and Shapley (1996). It is straightforward to verify that

$$\rho(s) = \begin{cases} \sum_{j=1}^{m(s)} \Delta(j) & \text{if } m(s) \geq 1, \\ 0 & \text{if } m(s) = 0. \end{cases} \quad (8)$$

is a potential function for these games.¹³ Potential games are relevant for the logit-response dynamics, which, as mentioned in the introduction, selects the potential-maximizing states as stochastically stable under certain assumptions.

We now introduce two additional properties that are crucial for robust stochastic stability, *supermodularity* and *strategic complementarity* (Topkis, 1998; Vives, 2005). For symmetric binary action games, the usual definition of supermodularity based on increasing differences of the payoff function is equivalent to the requirement that the difference function $\Delta(n)$ is weakly increasing in n .¹⁴ Figure 1 depicts an exemplary difference function of a supermodular game. The unanimity game is supermodular. Another example would be a population game where N players are matched pairwise in a round-robin tournament to play a symmetric 2×2 coordination game.

The team project game is not supermodular (except if $n_A = n_B = N$) because $\Delta(n)$ is decreasing from 1 to $N - n_B + 1$ and from n_A to N . Still, the game has monotonic best-responses and hence satisfies the weaker condition of strategic complementarity. For symmetric binary action games, strategic complementarity is equivalent to the existence of two values \underline{n} and \bar{n} with the property that $\Delta(n) < 0$ if and only if $n \leq \underline{n}$ and $\Delta(n) > 0$ if and only if $n > \bar{n}$. Thus any supermodular game exhibits

¹²Maruta (2002) and Maruta and Okada (2009) generalize unanimity games to the different class of “binary coordination games” (see our discussion below). Our team project game is also related to the collective-action games studied by Myatt and Wallace (2008a,b) for general quantal response dynamics under asynchronous learning (and also simultaneous learning in Myatt and Wallace, 2008b). The games in Myatt and Wallace (2008a,b) are not necessarily symmetric, they exhibit a single project only, and all players obtain a positive payoff if the project is successful.

¹³Maruta (2002) shows that symmetric binary coordination games are exact potential games, with a potential function as given in (8). Myatt and Wallace (2008b) show that their collective-action games are potential games under a symmetry condition, again with a potential function similar to (8).

¹⁴Such games are called binary coordination games by Maruta (2002) and Maruta and Okada (2009). Technically speaking, Maruta (2002) requires the difference function Δ to be strictly increasing. Maruta and Okada (2009) allow for games that are not necessarily symmetric.




Figure 1: Supermodularity

strategic complementarity but the converse is not true. Figure 2 illustrates the case of a game that exhibits strategic complementarity but is not supermodular. We assume throughout that $\underline{n} > 0$ and $\bar{n} < N$, to render the problem of equilibrium selection meaningful. If Γ exhibits strategic complementarity, only **A** or **B** can maximize the potential function ρ given in (8), which adds the values $\Delta(n)$ up to $m(s)$ for any s .




Figure 2: Strategic Complementarity

4.2 Logit-Response

We first study the logit-response dynamics based on myopic best-response. An earlier result by Blume (1993, 1997) for the class of exact potential games implies that the potential maximizing strategy profile will be stochastically stable under asynchronous learning. With the potential function (8), the difference in potential between two states s and s' corresponds to the accumulated utility changes of moving asynchronously from s to s' . Moving towards a profile with larger potential is thus always easier under logit response if only one player can update at a time. Consider the unanimity game, for instance. We only need to compare the value of the potential between **A** and **B**. Straightforward calculations reveal that $\rho(\mathbf{B}) = 0$ and $\rho(\mathbf{A}) = (a - b)/N$, so that a project is stochastically stable with asynchronous logit-response if and only if it is

Pareto efficient. We now want to examine under which conditions the selection of the potential maximizers by the logit-response dynamics is robust with respect to revision processes.

Theorem 1. *Let Γ be a supermodular symmetric binary action game. Consider the logit-response dynamics. Then, the potential maximizers are RSS-R.*

Proof. Let Γ be a supermodular symmetric binary action game, with critical values $\underline{n} > 0$ and $\bar{n} < N$ as defined above. We will rely on Corollary 1 and the radius-coradius from Alós-Ferrer and Netzer (2010). The coradius is always monotone, by Proposition 2. For any process $q \in \mathcal{Q}$, the unperturbed myopic best-reply dynamics with random (equal probability) tie-breaking has two absorbing sets, the singletons $\{\mathbf{A}\}$ and $\{\mathbf{B}\}$. Hence the radius is also monotone for both absorbing states, by Proposition 3.

Consider asynchronous learning. The waste caused by a single player switching from B to A in the presence of $n - 1$ other A -players is $\max\{-\Delta(n), 0\}$ and hence non-zero if and only if $n \leq \underline{n}$. Analogously, the waste that is generated if one of n A -players switches to B is $\max\{\Delta(n), 0\}$ and non-zero if and only if $n > \bar{n}$. Fix any $s' \neq \mathbf{B}$. Construct a minimal waste path $P = (s', \dots, \mathbf{B})$ by letting A -players switch to B sequentially. We obtain the waste

$$W(P) = \sum_{j=1}^{m(s')} \max\{\Delta(j), 0\}.$$

It is maximal if $s' = \mathbf{A}$ so that $m(s') = N$, which yields

$$CR^{q^{AL}}(\mathbf{B}) = \sum_{j=1}^N \max\{\Delta(j), 0\} = \sum_{j=\bar{n}+1}^N \Delta(j),$$

and, from analogous arguments,

$$CR^{q^{AL}}(\mathbf{A}) = \sum_{j=1}^N \max\{-\Delta(j), 0\} = -\sum_{j=1}^{\underline{n}} \Delta(j).$$

Consider independent learning. Since $\Delta(n)$ is increasing in n by supermodularity, the waste caused by a B -player switching to A in the presence of $n - 1$ A -players, $\max\{-\Delta(n), 0\}$, is decreasing in n . Analogously, the waste of an A -player switching to B , $\max\{\Delta(n), 0\}$, is increasing in n . Hence the waste caused by several players switching simultaneously is weakly larger than the waste caused by sequential switching, so that among minimal waste paths between \mathbf{A} and \mathbf{B} there are always paths that make use of sequential revisions only. This implies

$$R^{q^{IL}}(\mathbf{B}) = -\sum_{j=1}^{\underline{n}} \Delta(j), \quad R^{q^{IL}}(\mathbf{A}) = \sum_{j=\bar{n}+1}^N \Delta(j).$$

Now suppose \mathbf{B} is the unique potential maximizer, i.e., $\sum_{j=1}^N \Delta(j) < 0$. This can be rearranged to

$$-\sum_{j=1}^{\underline{n}} \Delta(j) > \sum_{j=\bar{n}+1}^N \Delta(j),$$

because $\Delta(n) = 0$ for $\underline{n} < n \leq \bar{n}$. This is equivalent to $R^{q^{IL}}(\mathbf{B}) > CR^{q^{AL}}(\mathbf{B})$ and implies that \mathbf{B} is the unique RSS-R state by Corollary 1. The argument for \mathbf{A} is analogous. If both \mathbf{A} and \mathbf{B} maximize the potential, i.e., $\sum_{j=1}^N \Delta(j) = 0$, Corollary 1 is not applicable. From the above arguments about supermodularity it is still true that minimal waste revision trees (Alós-Ferrer and Netzer, 2010) can be constructed using singleton revising sets only, so the stochastically stable states for any regular revision process must be the potential maximizers as for asynchronous learning. \square

The selection of potential maximizers in exact potential games (Blume, 1993, 1997) has been shown to be knife-edge by Alós-Ferrer and Netzer (2010), in the sense that it neither holds for general revision processes beyond asynchronous learning even for exact potential games, nor for generalized potential games even for asynchronous learning.¹⁵ Theorem 1 shows that potential maximizers do become a robust prediction for the particular class of supermodular symmetric binary action games. Hence, the relevance of potential maximizers does extend beyond asynchronous learning, at the price of considering a smaller class of games. The above mentioned selection of the Pareto efficient equilibrium in the unanimity game is robust by its supermodularity property. The result that the risk-dominant equilibrium of a symmetric 2×2 coordination game played in a round-robin tournament or on a (weighted) network (Young, 1998b) will be selected by the logit dynamics is also robust due to supermodularity. The same is true for the results that Maruta (2002) obtains for binary coordination and hence supermodular games under asynchronous logit response.

Theorem 1 has a straightforward intuition. With a logit choice rule, the likelihood of a player choosing a non-best-reply is decreasing in the size of the associated payoff loss. In supermodular games, this payoff loss becomes smaller the more players have already switched to the non-optimal action. Hence a mistake becomes more likely the more players have already made that mistake before. The minimal resistance paths between absorbing states that are relevant for stochastic stability are therefore constructed by letting players switch sequentially as under asynchronous learning, so that the stochastically stable states under asynchronous learning are stochastically stable

¹⁵Interestingly, however, Okada and Tercieux (2008) show that, under supermodularity, the asynchronous version of the logit-response dynamics selects local potential maximizers, a generalization of potential maximizers. Marden and Shamma (2012) describe different ways to restore the selection of potential maximizers for more general revision processes, including updating probabilities that depend on the level of noise and restrictions on actions sets or information available to the players. Marden and Shamma (2012) and Candogan, Ozdaglar, and Parrilo (2013) also investigate games that are approximately potential games.

for any regular revision process. The team project game is not supermodular. The mistake of leaving a successful project becomes (myopically) more expensive and hence less likely under logit choice the more players have already made that mistake before, because the project benefit is then shared among a smaller number of people. With independent learning, for instance, minimal resistance paths between the absorbing states are constructed by letting players switch simultaneously. Hence the selection result of logit-response in the team project game will depend on the revision process, and an RSS-R state does not generally exist. This illustrates that Theorem 1 cannot be generalized from supermodular games to the broader class of games with strategic complements.

4.3 Mistakes Model

In this section, we apply the myopic best-response based mistakes model. We proceed in parallel to the previous section and first investigate robustness with respect to revision processes, given random tie-breaking.

Theorem 2. *Let Γ be a symmetric binary action game with strategic complementarity. Consider the profile of random tie-breaking best-reply rules and the mistakes model. Then, the following states are the unique RSS-R states:*

$$\mathbf{A} \text{ if and only if } \underline{n} + \bar{n} \leq N, \quad \mathbf{B} \text{ if and only if } \underline{n} + \bar{n} \geq N.$$

Proof. Let Γ be a symmetric binary action game with strategic complementarity, with critical values $\underline{n} > 0$ and $\bar{n} < N$. We will rely on Corollary 1 and radius-coradius from Ellison (2000). The coradius is monotone by Proposition 2. For any $q \in \mathcal{Q}$, the unperturbed myopic best-reply dynamics with random tie-breaking has two absorbing sets, the singletons $\{\mathbf{A}\}$ and $\{\mathbf{B}\}$. Hence the radius is monotone for both absorbing states, by Proposition 3.

Consider asynchronous learning. Since $\Delta(n) < 0$ if and only if $n \leq \underline{n}$ and $\Delta(n) > 0$ if and only if $n > \bar{n}$, by strategic complementarity, switching from B to A in the presence of $n - 1$ other A -players is a mistake if and only if $n \leq \underline{n}$. Analogously, switching to B is a mistake for any of the n A -players if and only if $n > \bar{n}$. Fix any $s' \neq \mathbf{B}$. Construct a minimal cost path $P = (s', \dots, \mathbf{B})$ by letting A -players switch to B sequentially. The cost of this path is $C(P) = \max\{m(s') - \bar{n}, 0\}$. It is maximal if $s' = \mathbf{A}$ so that $m(s') = N$, which yields

$$CR^{q^{AL}}(\mathbf{B}) = N - \bar{n},$$

and, from analogous arguments,

$$CR^{q^{AL}}(\mathbf{A}) = \underline{n}.$$

Consider independent learning. By the previous arguments, at least \underline{n} mistakes are necessary to enable a transition from \mathbf{B} to \mathbf{A} , for any regular revision process. With random tie-breaking, states with $m(s) \geq \underline{n}$ can be connected to \mathbf{A} without mistakes (by letting indifferent B -players switch to A until A becomes a best response). Together with the analogous arguments for \mathbf{A} , this implies

$$R^{q^{IL}}(\mathbf{B}) = \underline{n}, \quad R^{q^{IL}}(\mathbf{A}) = N - \bar{n}.$$

Now suppose $\underline{n} + \bar{n} > N$. This is identical to $R^{q^{IL}}(\mathbf{B}) > CR^{q^{AL}}(\mathbf{B})$ and implies that \mathbf{B} is the unique RSS-R state by Corollary 1. The same applies to \mathbf{A} if $\underline{n} + \bar{n} < N$. If $\underline{n} + \bar{n} = N$, Corollary 1 is not applicable. The above arguments, however, imply that in this case both (and only) \mathbf{A} and \mathbf{B} have minimal stochastic potential (Kandori and Rob, 1995) for any regular revision process, so that they are the unique RSS-R states. \square

Compared to the logit-response dynamics, the mistakes model requires only the weaker property of strategic complementarity for robustness of its selection result in symmetric binary action games. In this sense, the mistakes model is more robust than the logit-response dynamics. The reason is, of course, that it makes use of the payoff structure of the game to a lesser extent. Strategic complementarity implies that the basin of attraction of each monomorphic state contains in its interior no area where the unperturbed dynamics would lead away from the monomorphic state, and thus it suffices to compare the size of the basins, irrespective of the specific regular revision process. If a basin encompassed an area from which the unperturbed dynamics no longer gravitates back to the absorbing state, then a simultaneous strategy change by several players might allow for cost-saving jumps away from the absorbing state, generating a dependence of the selection result on the revision process. This illustrates the role of strategic complementarity for Theorem 2.

For the team project game, where $\underline{n} = N - n_B + 1$ and $\bar{n} = n_A - 1$, Theorem 2 implies that the project with smaller participation requirement is RSS-R, i.e., \mathbf{A} if and only if $n_A \leq n_B$ and \mathbf{B} if and only if $n_B \leq n_A$.¹⁶ As another immediate application, we again obtain the robustness of the selection of the risk-dominant equilibrium of a symmetric 2×2 coordination game played in a round-robin tournament. Similarly, the comparable results of Maruta and Okada (2009) for the symmetric case are robust due to strategic complementarity of their binary coordination games.

Since the mistakes model is a noise process that respects tie-breaking, we can additionally investigate robustness with respect to tie-breaking assumptions.

¹⁶Specifically, both projects are stochastically stable in the unanimity game, as already pointed out by Young (1998a). See Maruta and Okada (2009) for a treatment of generalized, asymmetric unanimity games under perturbed adaptive play as in Young (1993).

Proposition 4. *Let Γ be a symmetric binary action game with strategic complementarity. Consider the profile of best-reply correspondences and the mistakes model. Then, the following states are the unique RSS-RT states:*

$$\mathbf{A} \text{ if } \bar{n} < N/2, \quad \mathbf{B} \text{ if } \underline{n} > N/2.$$

Proof. We will rely on Corollary 2 together with Remark 1 (as a substitute for monotonicity of the radius). Observe that $\{\mathbf{A}\}$ and $\{\mathbf{B}\}$ are absorbing sets of any unperturbed myopic best-reply dynamics $(B, q) \in \mathcal{T} \times \mathcal{Q}$ (albeit not necessarily the only ones). From the arguments in the proof of Theorem 2, we have $R^{B,q}(\mathbf{B}) \geq \underline{n}$ for any $(B, q) \in \mathcal{T} \times \mathcal{Q}$. With random tie-breaking and independent learning, states with $m(s) \geq \underline{n}$ can again be connected to \mathbf{A} without mistakes. Hence we have

$$R^{B,q}(\mathbf{B}) \geq R^{B^X, q^{IL}}(\mathbf{B}) = \underline{n}, \quad \forall (B, q) \in \mathcal{T} \times \mathcal{Q},$$

and, from analogous arguments,

$$R^{B,q}(\mathbf{A}) \geq R^{B^X, q^{IL}}(\mathbf{A}) = N - \bar{n}, \quad \forall (B, q) \in \mathcal{T} \times \mathcal{Q}.$$

Consider cautious tie-breaking and asynchronous learning. Fix any $s' \neq \mathbf{B}$. Construct a minimal cost path $P = (s', \dots, \mathbf{B})$ by letting A -players switch to B sequentially. The cost of this path is $C(P) = \max\{m(s') - \underline{n}, 0\}$, because the switch of an indifferent A -player to B is a mistake with cautious tie-breaking. It is maximal for $s' = \mathbf{A}$, which yields

$$CR^{B^0, q^{AL}}(\mathbf{B}) = N - \underline{n},$$

and, from analogous arguments,

$$CR^{B^0, q^{AL}}(\mathbf{A}) = \bar{n}.$$

The statement of the proposition now follows immediately from Corollary 2. \square

Observe that our approach delivers only sufficient conditions for RSS-RT, hence we cannot conclude that there are no RSS-RT states if the conditions in the proposition are not satisfied. It shares this limitation with the older radius-coradius results.¹⁷ The relatively strict conditions in Proposition 4 arise because the absorbing sets of the unperturbed dynamics depend on tie-breaking assumptions. With random tie-breaking, only the monomorphic states are absorbing, as both are reached with positive probability from any state s with $\underline{n} \leq m(s) \leq \bar{n}$ after indifferent players change their strategy. With cautious tie-breaking, states with $\underline{n} < m(s) < \bar{n}$ become absorbing as

¹⁷We were able to state necessary and sufficient conditions for RSS-R in Theorem 2, because there our radius-coradius approach was always able to identify the unique stable state (except for the case where both monomorphic states are stable, for which a direct argument was made in the proof).

well, as no indifferent player will switch. More complicated non-singleton absorbing sets emerge for intermediate tie-breaking assumptions, such as when some players break ties cautiously and others randomly. Our approach requires to calculate the radius for the quickest/swiftest and the coradius for the least quick/swift dynamics, giving rise to a relatively small radius and relatively large coradius, and hence the sufficient conditions in Proposition 4.

5 Symmetric Games and Globally Stable ESS

Consider a symmetric game Γ as defined earlier. Relevant examples include Cournot oligopolies, rent-seeking games, and other classes of games (see Alós-Ferrer and Ania, 2005 for details). Following Schaffer (1988, 1989), a strategy $s^* \in S_0$ is a strict globally stable ESS (where ESS stands for evolutionarily stable strategy) if for all $s' \in S_0$, $s' \neq s^*$,

$$u(s^*|s', \dots, s', s^*, \dots, s^*) > u(s'|s', \dots, s', s^*, \dots, s^*) \quad (9)$$

for all $m \in \{1, \dots, N-1\}$. In words, s^* earns larger payoffs than any alternative strategy in any profile where only those two strategies are present.

In Cournot oligopolies, the Walrasian quantity has been shown by Vega-Redondo (1997) to be a strict globally stable ESS, and stochastically stable in imitation-based dynamics with mistakes. Alós-Ferrer and Ania (2005) extended this result, showing that any strict globally stable ESS in any symmetric game is stochastically stable in a mistakes dynamics with imitate-the-best and independent inertia (see also Alós-Ferrer and Schlag, 2009). Bergin and Bernhardt (2004) provide further results for convergence of imitative dynamics to “relative equilibria” (a closely related concept) in general games. Alós-Ferrer and Ania (2005) also generalize Vega-Redondo (1997) to show that strict global stability includes a family of outcomes of special interest (see also Bergin and Bernhardt, 2004 for related results). Informally, an aggregative game is a symmetric game such that players’ payoffs depend only on the own strategy and an aggregate of all strategies. If the own strategy and the aggregate exhibit an ordinal substitutability (quasisubmodularity), as is the case e.g. in Cournot oligopolies and rent-seeking games, strict global stability follows from a more economic concept, aggregate-taking-strategy, i.e., a generalization of Walrasian equilibrium where each player maximizes payoffs taking the aggregate of all strategies as given.

For any unperturbed imitate-the-best dynamics, the singletons of monomorphic states (s, \dots, s) are absorbing. Consider the dynamics with random tie-breaking and independent learning. If one mutant appears at state (s^*, \dots, s^*) , inequality (9) with $m = 1$ indicates that s^* -players earn strictly more than the mutant. Hence, one mutation is not enough to move from (s^*, \dots, s^*) to a different absorbing set. We conclude that $R^{B^X, q^{LL}}(s^*, \dots, s^*) > 1$. Consider now the dynamics with asynchronous learn-

ing and cautious tie-breaking. For any monomorphic state (s, \dots, s) with $s \neq s^*$, if a single mutation to s^* occurs, by inequality (9) with $m = N - 1$, we obtain that the mutant earns strictly more than the incumbents, and hence eventually an incumbent will be selected and switch to s^* . By (9) with $m = N - 2$, the next incumbent to be selected will also switch to s^* . Iterating, the dynamics will reach the state (s^*, \dots, s^*) . We might thus be tempted to conclude that $CR^{B^0, q^{AL}}(s^*, \dots, s^*) = 1$, and, by virtue of Corollary 2, that (s^*, \dots, s^*) is an RSS-RT state. This tempting conclusion, however, is incorrect. The problem is that, with the specified behavioral rule, both asynchronous learning and cautious tie-breaking create additional absorbing sets which do not correspond to monomorphic states, and the above coradius computation fails. The following example shows a failure of robustness with respect to tie-breaking assumptions, showing that a strict globally stable ESS needs not be RSS-RT.

Example 3. Let N be even, $N \geq 6$, and consider a symmetric game with strategies $S_0 = \{A, B, C\}$. Let $m^A(s)$ be the number of players choosing A and $\pi^A(s)$ the payoff of an A -player in profile s , and analogously for B and C . Hence we can write $u_i(s_i, s_{-i}) = \pi^{s_i}(s_i, s_{-i})$. Let $C(s) \subseteq S_0$ be the set of strategies actually played in s . Payoffs are either 1 (win) or 0 (lose) and given as follows. First, $\pi^A(\mathbf{A}) = \pi^B(\mathbf{B}) = \pi^C(\mathbf{C}) = 1$. Second, $\pi^A(s) = 1$ and $\pi^B(s) = 0$ whenever $C(s) = \{A, B\}$, and $\pi^A(s) = 1$ and $\pi^C(s) = 0$ whenever $C(s) = \{A, C\}$. Third, $\pi^A(s) = 0$ and $\pi^B(s) = \pi^C(s) = 1$ whenever $C(s) = \{A, B, C\}$. Last, whenever $C(s) = \{B, C\}$, $\pi^B(s) = 1$ if $m^B(s) \leq m^C(s)$ and $\pi^B(s) = 0$ otherwise, and symmetrically for C , i.e., in this case the weak minority wins. Note that A is a strict globally stable ESS of this game.

Now consider imitate-the-best with cautious tie-breaking and asynchronous learning. Analyze the dynamics on the reduced state space where states with the same number of players for each strategy are identified. There are four absorbing sets, all of them singletons. They correspond to the three monomorphic states plus the state \hat{s} with equal number of B - and C -players, and no A -player. A single mutation to B suffices for a transition from \mathbf{C} to \hat{s} and symmetrically for \mathbf{B} , and two mutations (one to B , one to C) suffice for a transition from \mathbf{A} to \hat{s} . Hence $CR^{B^0, q^{AL}}(\hat{s}) = 2$. However, at least $N/2 \geq 3$ mutations are needed to leave \hat{s} for a different absorbing set, so $R^{B^0, q^{AL}}(\hat{s}) > CR^{B^0, q^{AL}}(\hat{s})$ and it follows that \hat{s} is the only stochastically stable state.

The last example relies crucially on the combination of cautious tie-breaking (which renders \hat{s} absorbing) and asynchronous learning (which implies that more than one mutation is necessary to destabilize \hat{s}). It raises the natural question of whether a strict globally stable ESS would at least be RSS-R if one maintains the assumption of random tie-breaking. The answer is negative. The next example shows a failure of robustness with respect to revision opportunities.

Example 4. Let N be even, $N \geq 8$, and consider a symmetric game with strategies $S_0 = \{A, B, C\}$. Let the notation be as in the previous example. Payoffs are as follows.

First, $\pi^A(\mathbf{A}) = \pi^B(\mathbf{B}) = \pi^C(\mathbf{C}) = 1$. Second, $\pi^A(s) = 1$ and $\pi^B(s) = 0$ whenever $C(s) = \{A, B\}$, and $\pi^A(s) = 1$ and $\pi^C(s) = 0$ whenever $C(s) = \{A, C\}$. Third, whenever $C(s) = \{A, B, C\}$ or $C(s) = \{B, C\}$, $\pi^A(s) = 0$ (if $A \in C(s)$), $\pi^B(s) = 1$ if $m^B(s) \leq m^C(s)$ and $\pi^B(s) = 0$ otherwise, and symmetrically for C , i.e., the weak minority between B and C wins even if A is present. Again, A is a strict globally stable ESS of this game.

Now consider imitate-the-best with random tie-breaking and asynchronous learning. It is easy to see that there is a non-singleton absorbing set $S' = \{s_-, s_+, s_+\}$, with $m^B(s_-) = m^C(s_-) = N/2$, $m^B(s_+) = m^C(s_-) = (N/2) + 1$, and $m^B(s_-) = m^C(s_+) = (N/2) - 1$. The unperturbed dynamics transitions into S' from all states with $C(s) = \{A, B, C\}$ or $C(s) = \{B, C\}$. A single mutation suffices to reach one of these states from \mathbf{B} , \mathbf{C} , and any state with $C(s) = \{A, B\}$ or $C(s) = \{A, C\}$. It takes two mutations to reach one of these states from \mathbf{A} . To leave S' for a different absorbing set, at least $(N/2) - 1 \geq 3$ mutations are necessary. It follows that the states in S' are the only stochastically stable states.

The examples above are, admittedly, rather extreme. They give, however, an example of a well-established stochastic stability result which fails the robustness test(s). This robustness failure arises from a conceptual problem with the definition of finite-population ESS. That definition is based on payoff differences between mutant and incumbents when only two strategies are present, which allows wild specifications of payoffs when more than two strategies are part of the strategy profile. A concept based on pairwise comparisons, while yielding stochastic stability for a reasonable set of dynamics, can be “tricked” by exploiting the freedom of specification left by the concept when more than two strategies are present. This kind of problem is precisely what our robustness test was meant to identify.

6 Conclusion

Stochastic stability is and remains an important concept in game theory. In our view, it is well suited to analyze questions of outcome selection in noisy environments, as long as the different ingredients of the model are clearly differentiated. Ideally, a strong, clear-cut result is one linking a particular behavioral assumption (captured by a behavioral rule or correspondence) under a particular interaction structure (as a proxy for the socioeconomic setting, e.g. the network structure) to the selection of a particular outcome. Failing that, it is still important to understand how more subtle elements of the model have an influence on long-run outcomes. Our concept of robust stochastic stability aims to differentiate clear-cut predictions from more subtle ones. For noisy behavioral rules where specific tie-breaking assumptions are built into the rule, as in the case of logit choice, robust stochastic stability requires robustness with respect to the specification of revision opportunities. For noisy rules which remain

silent (or are less vocal) on the issue of tie-breaking, as in the case of the mistakes model, robustness also includes the latter.

We have provided an easy-to-use sufficient condition for robust stochastic stability, and have illustrated its application for different games and dynamics. The condition makes use of an order structure of the space of dynamics, by observing that the radius and coradius concepts introduced in the literature are typically monotone operators in this space. Our result itself reduces to a radius-coradius approach, with the difference that the radius is taken with respect to the “quickest” dynamics (independent learning and random tie-breaking) and the coradius is taken with respect to the “slowest” one (asynchronous learning and cautious tie-breaking). Hence in the quest to obtain results which are independent of certain parts of the specification of the dynamics, we are led to concentrate on two particular, extreme dynamics.

In our illustrations, we have focused on two important classes of behavioral rules, myopic best reply and imitate-the-best. Our results can of course be applied to other rules. For instance, suppose one wishes to consider robustness within the class of rules that focus on strategies delivering better payoffs than the current one, either in a forward-looking sense (“better-reply dynamics”, as in e.g. Friedman and Mezzetti, 2001) or in an observational sense (“imitate if better”, as in e.g. Ellison and Fudenberg, 1995). If one defines the behavioral correspondence in the natural way, the question reduces to robustness with respect to tie-breaking rules within the class of behavioral rules agreeing with that correspondence.

We conclude by emphasizing that working with radius and coradius of two different dynamics is not the same as establishing stochastic stability for these two dynamics, i.e., as checking conventional radius-coradius conditions for the two. Robustness is stronger, as it requires stochastic stability also for all intermediate (balanced and regular) dynamics. The following logit-response based example shows that a state can be stochastically stable for both independent and asynchronous learning, but fail to be stochastically stable for all regular revision processes. Consider a symmetric binary action game with $N = 7$ and a difference function $\Delta(n)$ as depicted in Figure 3. If learning is asynchronous and $-\Delta(1) + \Delta(2) + \Delta(3) < \Delta(5) + \Delta(6) + \Delta(7)$, then **A** is stochastically stable because the utility losses of moving from **A** to **B** are strictly larger than those of moving from **B** to **A**. If learning is independent and $-\Delta(1) < \Delta(7)$, then the same is true because optimal transitions between **A** and **B** involve simultaneous strategy changes by 3 players, generating minimal utility losses of $3\Delta(7)$ and $-3\Delta(1)$, respectively. Now consider a regular revision process $q \in \mathcal{Q}$ where $q_{\{1,2\}} > 0$ and $q_J = 0$ for all other $J \subseteq I$ with $|J| \geq 2$. Under the assumption that $-\Delta(2) < -[\Delta(1) + \Delta(3)]/2$, a minimal utility loss transition from **B** to **A** involves a single player switching to *A* first, followed by a simultaneous switch of players 1 and 2. The associated utility losses are $-\Delta(1) + 2\Delta(2)$. An analogous argument shows that moving from **A** to **B** generates minimal utility losses of $[\Delta(7) + 2\Delta(6)]$ if




Figure 3: Non-Robustness

$\Delta(6) < [\Delta(5) + \Delta(7)]/2$. Now consider a game where

$$\Delta(1) = -1, \Delta(2) = -10, \Delta(3) = -20, \Delta(5) = 20, \Delta(6) = 7, \Delta(7) = 5.$$

These parameters satisfy all above assumptions, and they imply that **B** is stochastically stable under q because $-\Delta(1) + 2\Delta(2) > \Delta(7) + 2\Delta(6)$.

References

- ALÓS-FERRER, C. (2003): “Finite Population Dynamics and Mixed Equilibria,” *International Game Theory Review*, 5, 263–290.
- (2004): “Cournot vs. Walras in Oligopoly Models with Memory,” *International Journal of Industrial Organization*, 22, 193–217.
- (2008): “Learning, Bounded Memory, and Inertia,” *Economics Letters*, 101(2), 134–136.
- ALÓS-FERRER, C., AND A. B. ANIA (2005): “The Evolutionary Stability of Perfectly Competitive Behavior,” *Economic Theory*, 26, 497–516.
- ALÓS-FERRER, C., AND N. NETZER (2010): “The Logit-Response Dynamics,” *Games and Economic Behavior*, 68, 413–427.
- ALÓS-FERRER, C., AND K. SCHLAG (2009): “Imitation and Learning,” in *The Handbook of Rational and Social Choice*, ed. by P. Anand, P. Pattanaik, and C. Puppe. Oxford University Press.
- ALÓS-FERRER, C., AND F. SHI (2012): “Imitation with Asymmetric Memory,” *Economic Theory*, 49(1), 193–215.
- ALÓS-FERRER, C., AND S. WEIDENHOLZER (2008): “Contagion and Efficiency,” *Journal of Economic Theory*, 143, 251–274.

- BEGGS, A. (2005): “Waiting Times and Equilibrium Selection,” *Economic Theory*, 25, 599–628.
- BERGIN, J., AND D. BERNHARDT (2004): “Comparative Learning Dynamics,” *International Economic Review*, 45, 431–465.
- (2009): “Cooperation through Imitation,” *Games and Economic Behavior*, 67, 376–388.
- BERGIN, J., AND B. L. LIPMAN (1996): “Evolution with State-Dependent Mutations,” *Econometrica*, 64, 943–956.
- BLUME, L. (1993): “The Statistical Mechanics of Strategic Interaction,” *Games and Economic Behavior*, 5, 387–424.
- (1997): “Population Games,” in *The Economy as an Evolving Complex System II*, ed. by B. Arthur, S. Durlauf, and D. Lane, pp. 425–460. Addison-Wesley, Reading, MA.
- (2003): “How noise matters,” *Games and Economic Behavior*, 44, 251–271.
- CANDOGAN, O., A. OZDAGLAR, AND P. PARRILO (2013): “Dynamics in Near-Potential Games,” *Games and Economic Behavior*, forthcoming.
- DOKUMACI, E., AND W. H. SANDHOLM (2008): “Stochastic Evolution with Perturbed Payoffs and Rapid Play,” Mimeo, University of Wisconsin-Madison.
- ELLISON, G. (1993): “Learning, Local Interaction, and Coordination,” *Econometrica*, 61, 1047–1071.
- (2000): “Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution,” *Review of Economic Studies*, 67, 17–45.
- ELLISON, G., AND D. FUDENBERG (1995): “Word-of Mouth Communication and Social Learning,” *Quarterly Journal of Economics*, 110, 95–126.
- FRIEDMAN, J. W., AND C. MEZZETTI (2001): “Learning in Games by Random Sampling,” *Journal of Economic Theory*, 98, 55–84.
- HOFBAUER, J., AND G. SORGER (2002): “A Differential Game Approach To Evolutionary Equilibrium Selection,” *International Game Theory Review*, 4, 17–31.
- JOSEPHSON, J., AND A. MATROS (2004): “Stochastic Imitation,” *Games and Economic Behavior*, 49, 244–259.
- KANDORI, M., G. J. MAILATH, AND R. ROB (1993): “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 61, 29–56.

- KANDORI, M., AND R. ROB (1995): “Evolution of Equilibria in the Long Run: A General Theory and Applications,” *Journal of Economic Theory*, 65, 383–414.
- KIM, Y. (1996): “Equilibrium Selection in n-Person Coordination Games,” *Games and Economic Behavior*, 15, 203–227.
- MARDEN, J. R., AND J. S. SHAMMA (2012): “Revisiting Log-Linear Learning: Asynchrony, Completeness and Payoff-Based Implementation,” *Games and Economic Behavior*, 75, 788–808.
- MARUTA, T. (2002): “Binary Games with State Dependent Stochastic Choice,” *Journal of Economic Theory*, 103(2), 351–376.
- MARUTA, T., AND A. OKADA (2009): “Stochastically Stable Equilibria in Coordination Games with Multiple Populations,” mimeo.
- MONDERER, D., AND L. SHAPLEY (1996): “Potential Games,” *Games and Economic Behavior*, 14, 124–143.
- MORRIS, S. (2000): “Contagion,” *Review of Economic Studies*, 67, 57–78.
- MYATT, D. P., AND C. WALLACE (2003): “A Multinomial Probit Model of Stochastic Evolution,” *Journal of Economic Theory*, 113, 286–301.
- (2008a): “An Evolutionary Analysis of the Volunteer’s Dilemma,” *Games and Economic Behavior*, 62, 67–76.
- (2008b): “When Does One Bad Apple Spoil the Barrel? An Evolutionary Analysis of Collective Action,” *Review of Economic Studies*, 75, 499–527.
- OECHSSLER, J. (1997): “An Evolutionary Interpretation of Mixed-Strategy Equilibria,” *Games and Economic Behavior*, 21, 203–237.
- OKADA, D., AND O. TERCIEUX (2008): “Log-linear Dynamics and Local Potential,” Mimeo.
- RHODE, P., AND M. STEGEMAN (1996): “A Comment on “Learning, Mutation, and Long-Run Equilibria in Games”,” *Econometrica*, 64, 443–449.
- RITZBERGER, K., AND J. WEIBULL (1995): “Evolutionary Selection in Normal-Form Games,” *Econometrica*, 63, 1371–1399.
- ROBSON, A. J., AND F. VEGA-REDONDO (1996): “Efficient Equilibrium Selection in Evolutionary Games with Random Matching,” *Journal of Economic Theory*, 70, 65–92.

- SANDHOLM, W. H. (1998): “Simple and Clever Decision Rules for a Model of Evolution,” *Economics Letters*, 61, 165–170.
- (2010): “Orders of Limits for Stationary Distributions, Stochastic Dominance, and Stochastic Stability,” *Theoretical Economics*, 5, 1–26.
- SCHAFFER, M. (1988): “Evolutionarily Stable strategies for a Finite Population and a Variable Contest Size,” *Journal of Theoretical Biology*, 132, 469–478.
- (1989): “Are Profit-Maximisers the Best Survivors?,” *Journal of Economic Behavior and Organization*, 12, 29–45.
- STAUDIGL, M. (2012): “Stochastic Stability in Binary Choice Coordination Games,” *Games and Economic Behavior*, 75, 372–401.
- TOPKIS, D. M. (1998): *Supermodularity and Complementarity*. Princeton University Press, Princeton, New Jersey.
- VAN DAMME, E., AND J. W. WEIBULL (2002): “Evolution in Games with Endogenous Mistake Probabilities,” *Journal of Economic Theory*, 106, 296–315.
- VEGA-REDONDO, F. (1997): “The Evolution of Walrasian Behavior,” *Econometrica*, 65, 375–384.
- VIVES, X. (2005): “Complementarities and Games: New Developments,” *Journal of Economic Literature*, 43, 437–479.
- WEIDENHOLZER, S. (2010): “Coordination Games and Local Interactions: A Survey of the Game Theoretic Literature,” *Games*, 1(4), 551–585.
- YOUNG, P. (1993): “The Evolution of Conventions,” *Econometrica*, 61, 57–84.
- (1998a): “Conventional Contracts,” *Review of Economic Studies*, 65, 773–792.
- (1998b): *Individual Strategy and Social Structure*. Princeton University Press, Princeton, New Jersey.