



Conventions in large games with randomly drawn payoffs[★]

Jonathan Newton ^a, Ryoji Sawa ^b

^a Institute of Economic Research, University of Kyoto, Yoshida-honmachi Sakyo-ku, 606–8501, Kyoto, Japan

^b University of Tsukuba, Tennoudai 1-1-1, Tsukuba, Ibaraki, Japan

ARTICLE INFO

JEL classification:

C71
C73
D01

Keywords:

Evolution
Stability
One-shot
Social choice

ABSTRACT

By showing that a one-shot deviation property holds for almost all large games with i.i.d. randomly drawn payoffs in which a convention (in the sense of Young, 1993a) exists, long-run outcomes under behavioral rules are characterized by social choice rules.

1. Introduction

Lewis (1969) argued that conventions, self-reinforcing regularities in the behavior of individuals in a society, can emerge when individuals follow simple, adaptive behavioral rules. Mathematical foundations for this theory were later presented by Young (1993a) and Kandori et al. (1993), who showed that conventions can be ranked by their stability properties. The stability of conventions under diverse behavioral rules has since been extensively studied (Sandholm, 2010; Newton, 2018).

In general, the stability of conventions depends on non-trivial interaction between the underlying game and individuals' behavioral rules. Stability is a global property. However, it is possible to construct example pairings of games and behavioral rules under which stability is determined by a local property. Specifically, we say that a pair of a game and a behavioral rule satisfy the *one-shot deviation property* if the relative stability of conventions is wholly determined by the probability with which a *deviation* from equilibrium play occurs at each convention. Formally, identify the set of conventions with NE , the set of strict Nash equilibria. Write $OS \subseteq NE$ as the set of equilibria at which the probability of a deviation is lowest. Let SS be the most stable, the *stochastically stable*, equilibria (Foster and Young, 1990). Then, the one-shot deviation property holds if $OS = SS$ (Fig. 1).

In this paper, we leverage recent advances on the structure of large games with randomly drawn payoffs (Johnston et al., 2025)¹ to show that the one-shot deviation property holds for almost all² pairs of games and behavioral rules when payoffs are i.i.d. and the

[★] Newton's work was supported by KAKENHI B funding from the Japan Society for the Promotion of Science (Grant 21H00695), and Sawa's work was supported by KAKENHI funding from the Japan Society for the Promotion of Science (Grant 23K22098, 22H05086). We thank Srinivas Arigapudi, Andrea Collecchio, Bary Pradelski, Bassel Tarbush, Marco Scarsini, Satoru Takahashi, as well as audiences at the Spring 2024 LEG seminar, Econometric Society Asian Meetings, and SWET 2024 at Hokkaido University.

E-mail addresses: newton@kier.kyoto-u.ac.jp (J. Newton), rsawa@sk.tsukuba.ac.jp (R. Sawa).

¹ These games are often known as *random games* but we follow the more explicit nomenclature of *games with random payoffs / utilities* as suggested in Pradelski and Tarbush (2025).

² Throughout, we use *almost all* in the sense used in random graph theory, meaning that a property holds asymptotically almost surely (e.g. Bollobás, 2001; Spencer, 2013) as the number of vertices (here, players) tends to infinity. Given the links between random graphs and best response correspondences of games with random payoffs, this terminology is natural. All conditions in formal results are stated explicitly.

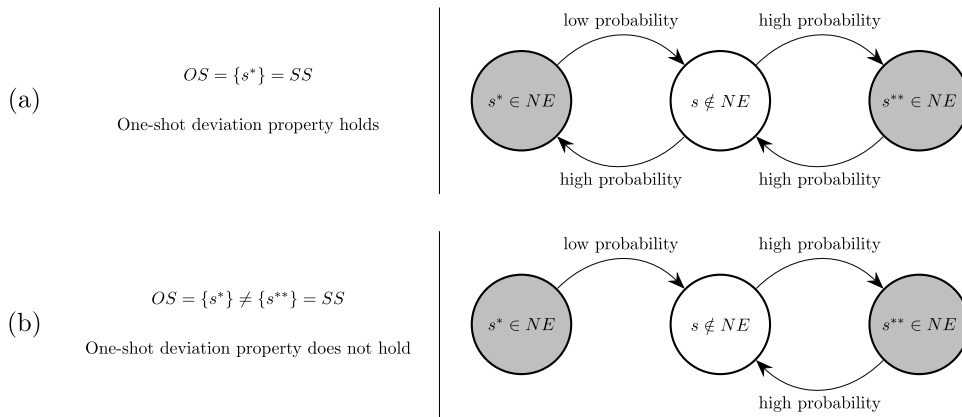


Fig. 1. One-shot deviation property. There are three illustrated strategy profiles, s^*, s', s^{**} . The set of strict Nash equilibria is NE . Probabilities are either ‘low’ or ‘high’. $OS = \{s^*\}$ is the set of profiles from which transitions occur with the lowest probability. In **Panel (a)** we see that upon leaving either profile in NE , the other profile in NE is reached with high probability. Therefore, the most stable equilibria SS are those for which the initial deviation has the lowest probability. $SS = \{s^*\} = OS$ so the one-shot deviation property holds. In **Panel (b)**, although transitions from s^* have low probability, after they occur s^* is never again reached. Consequently, $SS = \{s^{**}\} \neq OS$ so the one-shot deviation property does not hold.

behavioral rule is based on best response. Specifically, for any *universal behavioral rule* that specifies behavior in any game, for any finite bound on the size of strategy sets, the one-shot deviation property holds asymptotically almost surely in the number of players [Theorem 1]. That is, for large games with randomly drawn payoffs, it is very likely that the one-shot deviation property holds. The problem of determining SS reduces to the problem of determining OS , which only depends on local properties of equilibria.

Our result allows us to definitively characterize the outcomes of given behavioral rules across almost all games. In particular, as strategy choice depends on payoffs, the properties of conventions in OS can depend on payoff comparisons both within and between players, giving results of the form:

For almost all games, X is stochastically stable under behavioral rule Y if and only if X is chosen by the social choice rule Z .

By *social choice rule* we mean a mapping from players’ preferences to outcomes. In contrast to normative analysis, which considers properties (moral, desirable etc.) of social choice rules, our approach is closer to positive analysis, focusing on how social choice rules emerge as long run equilibria under different behavioral rules.

Interestingly, popular behavioral rules lead to social choice rules that are related to established concepts in noncooperative and cooperative game theory. Best response with uniform deviations selects NE , the set of strict Nash equilibria [Proposition 1]. Payoff-difference based rules, including logit and probit choice, select the least- NE , a non-cooperative analogy of the least-core in cooperative game theory [Proposition 3]. Condition dependent rules, under which deviations are more likely when players have low payoffs, select Rawlsian- NE , the subset of strict Nash equilibria that maximize the payoff of the least well-off player [Proposition 4].

Analytically, the one-shot deviation property brings further benefits. First, the speed of convergence to SS is governed by the most likely deviation available at each state outside SS , specifically by the least likely of these across all such states [Theorem 2]. Second, it becomes easy to make statements about medium term behavioral dynamics: trajectories followed by the process before it hits SS . This makes it simple to calculate the relative stability of all conventions [Theorem 3]. Third, all equilibria in SS will be robustly stochastically stable (Alós-Ferrer and Netzer, 2015) in the sense that selection does not depend on synchronous versus asynchronous strategy updating.

1.1. Alignment and misalignment with existing literature

Many studies in the existing literature consider games in which players’ payoffs are positively or negatively correlated to some extent, and our results for games with random payoffs do not directly apply to such settings. However, by proving results for almost all games with i.i.d. randomly drawn payoffs, we provide a benchmark against which existing or future results can be interpreted. For a given behavioral rule based on best response, the baseline expectation is that results align with the current paper. When results do not align, the divergence must arise from some aspect of the model structure other than the basic behavioral rule. Here we consider some examples of alignment and non-alignment.

A famous result in the stochastic stability literature is the emergence of risk dominant conventions in populations playing 2×2 coordination games. By definition, the risk dominant equilibrium maximizes the payoff loss from playing a non-best response. This makes it the least- NE outcome and our prediction under logit deviations [Proposition 3]. This is indeed the outcome that typically arises under logit (Newton, 2021; Blume, 2003; Alós-Ferrer and Netzer, 2010). In contrast, our prediction under uniform deviations is the set of strict Nash equilibria [Proposition 1] rather than the risk dominant outcomes found in the literature (Young, 1993a; Kandori et al., 1993; Newton, 2021). This misalignment arises because the cited papers effectively mimic the payoff dependence of logit by the game being played against a population rather than just a given individual.

In one-to-one matching problems with non-transferable utility, a one-shot deviation principle holds (Newton and Sawa, 2015), and consequently we have alignment. Uniform deviations select all stable matchings (Jackson and Watts, 2002a), logit gives a form of the least-core (Newton and Sawa, 2015), and condition dependence gives Rawlsian stable matchings (Bilancini et al., 2020). Similarly, in matching with transferable utility, uniform deviations give almost all stable matchings (Klaus and Newton, 2016) and logit gives the least-core (Nax and Pradelski, 2015). However, when Klaus and Newton (2016) study logit, they obtain almost no selection. This misalignment arises because the model allows local deviations that do not generate enough disorder to generate a one-shot result.

There is partial alignment in the housing market model of Ben-Shoham et al. (2004). One-shot results applicable to matching models are not applicable, since houses have no preferences or agency. Under uniform deviations, they select efficient (but not all) stable allocations. When serious mistakes are less likely, their deviations are an ordinal discretization of logit (Kandori et al., 2008). In fact, due to the process satisfying detailed balance conditions, the effects turn out to be aggregative, with logit selecting allocations that maximize total utility (in the ordinal environment, equivalent to minimizing aggregate envy).

In two player bargaining problems, logit selects the egalitarian solution (Hwang et al., 2018), effectively the least-core and thus aligned with Proposition 3. In contrast, under uniform deviations, Young (1993b) does not find no selection as predicted by Proposition 1, but rather the emergence of the Kalai-Smorodinsky bargaining solution. This is due to a sampling/population process giving rise to patterns of behavior that are similar to condition dependence, but with payoffs for any given player measured relative to their best possible payoff.

Experimental literature supports payoff and condition dependent perturbations. In coordination games, Mäs and Nax (2016) and Lim and Neary (2016) observe behavior that is primarily myopic best response, with deviations that are payoff dependent. Bilancini et al. (2021) show that deviations depend on both payoff differences and condition, with higher non-best response rates for subjects who are currently doing poorly. In bargaining experiments, Hwang et al. (2018) find that non-best-response play is payoff dependent and exhibits intentional, distributional bias, again rejecting uniform perturbations. Taken together, these results indicate that empirically relevant behavioral rules are payoff-difference based and condition dependent, while uniform perturbations are a theoretical simplification.

1.2. Related literature

Our paper contributes to a suite of methods for the study of the evolution of conventions. Some of these methods are fully general, such as the tree characterization (Young, 1993a; Kandori et al., 1993) and cyclic decompositions (Cui and Zhai, 2010; Levine and Modica, 2016; Newton and Sandholm, 2022). Other methods gain simplicity at the cost of generality. These include radius-coradius methods (Ellison, 2000) and asymmetry (Peski, 2010; Newton, 2021). When the one-shot deviation property holds, we are clearly in the domain of simplicity, so some generality must be sacrificed. For example, Newton and Sawa (2015) prove a one-shot result ($SS \subseteq OS$) in the domain of matching problems. The current paper deals with normal form games and shows that, somewhat remarkably, in this domain the one-shot deviation property ($SS = OS$) is asymptotically general.

The Evolutionary Nash Program is the study of connections between evolutionary game theory and cooperative game theory (Newton, 2018). Examples include the study of recontracting and Nash demand games (Feldman, 1974; Green, 1974; Young, 1993b; Agastya, 1997, 1999; Newton, 2012; Serrano and Volij, 2008; Sawa, 2021; Khan, 2022; Montero and Possajennikov, 2023), matching with transferable utility (Nax and Pradelski, 2015; Pradelski, 2015; Klaus and Newton, 2016; Nax et al., 2013), matching with non-transferable utility (Newton and Sawa, 2015; Bilancini et al., 2020; Jackson and Watts, 2002b; Roth and Vande Vate, 1990; Diamantoudi et al., 2004; Kojima and Ünver, 2008), the study of bargaining solutions as long term outcomes of coordination games (Young, 1998; Naidu et al., 2010; Hwang et al., 2018). The current paper contributes to this literature by giving the most general characterizations to date of the implications of behavioral rules in terms of social choice functions.

Finally, this paper contributes to the study of games with randomly drawn payoffs. This literature includes results on the number of (Nash) equilibria (Dresher, 1970; Stanford, 1996; Rinott and Scarsini, 2000; McLennan and Berg, 2005; Takahashi, 2008; Arieli and Babichenko, 2016; Pei and Takahashi, 2019) and convergence of adaptive dynamics, typically best response dynamics, to equilibrium (Amiet et al., 2021; Wiese and Heinrich, 2022; Heinrich et al., 2023; Johnston et al., 2025; Mimun et al., 2024). We build on this literature by combining it with stochastic evolutionary game theory and showing clean equilibrium selection under all regular perturbed best response dynamics in almost all normal form games that have at least one Nash equilibrium. Our main result [Theorem 1] uses continuous payoff distributions, but we also show that the result extends to distributions containing small atoms [Theorem 4].

2. Model

2.1. Games

A normal form game $G = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ consists of a finite set of players $N = \{1, \dots, n\}$, and for each player i , a finite set of strategies $S_i = \{1, \dots, |S_i|\}$ and payoff function $u_i : S \rightarrow \mathbb{R}$, where $S = \times_{i \in N} S_i$.^{3,4}

³ Under our dynamic process of strategy updating (Section 2.2), players can choose their strategies randomly, but the strategies that they choose are pure. This keeps the state space finite and analysis tractable. For issues that arise on infinite state spaces, see Newton (2015).

⁴ Note this is a single game played by n players. This differs from models in which, for example, n players are randomly matched to play a 2-player game. There is no random matching in our model. The randomness in our games enters via their payoffs (Section 2.4).

The set of best responses by player i to strategy profile $s = (s_i, s_{-i})$ is

$$B_i(s) := \left\{ \tilde{s}_i \in S_i : u_i(\tilde{s}_i, s_{-i}) = \max_{s'_i \in S_i} u_i(s'_i, s_{-i}) \right\}. \tag{1}$$

If, for all $i \in N$, s_i is a best response to s , then s is a *pure Nash equilibrium*. If all best responses are unique, the equilibrium is *strict*. The set of strict Nash equilibria of G is

$$NE := \left\{ s \in S : \forall i \in N, B_i(s) = \{s_i\} \right\}. \tag{2}$$

2.2. Dynamics and behavioral rules

Given a game G , play evolves according to a discrete time Markov process on S . Define a family of Markov processes $P = \{P^\epsilon\}_\epsilon$ indexed by $\epsilon \in [0, 1)$, where higher values of ϵ correspond to a greater frequency of perturbations from the *unperturbed process* P^0 . Let the state at time t be $s^t \in S$. Let P^ϵ be described by the following steps. At time $t + 1$, select an updating player $i \in N$ according to a probability measure π with full support on N . Then, let s^{t+1} be randomly determined according to a probability measure $P_i^\epsilon(s^t, \cdot)$ satisfying $P_i^\epsilon(s^t, s) = 0$ if $s_{-i} \neq s_{-i}^t$.

Note that P_i^ϵ is also a Markov process on S . We shall refer to the family $\{P_i^\epsilon\}_\epsilon$ as a *behavioral rule* for i . In summary, the two step strategy updating process selects an updating player before (possibly) updating his strategy, leaving the strategies of other players unchanged. The relationship between P^ϵ and $\{P_i^\epsilon\}_{i \in N}$ is given by

$$P^\epsilon(s, \cdot) = \sum_{i \in N} \pi(i) P_i^\epsilon(s, \cdot). \tag{3}$$

We consider *regular* behavioral rules (Young, 1993a), the class of rules that satisfy the following conditions. Let P_i^ϵ be continuous in ϵ . If $P_i^0(s, s') = 0$ and $P_i^{\hat{\epsilon}}(s, s') > 0$ for some $\hat{\epsilon} > 0$, let $\{P_i^\epsilon(s, s')\}_\epsilon$ satisfy

$$P_i^\epsilon(s, s') = (a + o(1)) \epsilon^k \quad \text{for some } a > 0, k > 0, \tag{4}$$

where a, k may depend on s, s', i , but not on ϵ ; and $o(1)$ represents a term that vanishes as $\epsilon \rightarrow 0$.⁵ This class of rules includes popular rules such as the logit choice rule and best response with uniform deviations.⁶

Assume that the unperturbed process follows best response,

$$P_i^0(s, (s'_i, s_{-i})) > 0 \iff s'_i \in B_i(s), \tag{5}$$

and that the perturbed process is non-deterministic,

$$\text{For all } \epsilon > 0, s \in S, s'_i \in S_i, P_i^\epsilon(s, (s'_i, s_{-i})) < 1. \tag{6}$$

We do not assume that P^ϵ has a unique recurrent class.⁷ However, when it does, we let μ^ϵ denote the unique invariant probability measure on the state space S . If there exists $\hat{\epsilon}$ such that P^ϵ has a unique recurrent class for $\epsilon \in (0, \hat{\epsilon})$, then standard arguments imply that the limit of μ^ϵ as $\epsilon \rightarrow 0$ exists. For small ϵ , the process will spend most of the time at states which have positive probability under this limiting measure. These are known as *stochastically stable* states (Foster and Young, 1990).

$$SS := \begin{cases} \{s \in S : \lim_{\epsilon \rightarrow 0} \mu^\epsilon(s) > 0\}, & \text{if there exists } \hat{\epsilon} \text{ such that } P^\epsilon \text{ has a unique} \\ \emptyset, & \text{recurrent class for all } \epsilon \in (0, \hat{\epsilon}), \\ & \text{otherwise.} \end{cases} \tag{7}$$

2.3. One-shot deviation property

Define the cost $c(s, s')$ of a transition from s to s' as the exponential rate of decay of the probability of such a transition as $\epsilon \rightarrow 0$.

$$c(s, s') := \begin{cases} \lim_{\epsilon \rightarrow 0} \frac{\log P_i^\epsilon(s, s')}{\log \epsilon} & \text{if } P_i^{\hat{\epsilon}}(s, s') > 0 \text{ for some } i \in N, \hat{\epsilon} > 0, \\ \infty & \text{otherwise.} \end{cases} \tag{8}$$

Cost functions measure the order of magnitude of transition probabilities for low values of ϵ . Transitions with a high cost are less likely than transitions with a low cost. From (8), we see that if $P_i^0(s, s') > 0$, then $c(s, s') = 0$. That is, transitions that can occur under the unperturbed process have zero cost. In contrast, if a transition is only possible for $\epsilon > 0$, then $c(s, s') = k$, where k is the k from expression (4).

⁵ Note that, as our state space is finite, it is possible to uniformly bound a and k across all possible transitions. This fact is later used in our proofs.

⁶ For example, consider logit choice when $S_i = \{s_i, s'_i\}$ and $u_i(s_i, s_{-i}) = u > u' = u_i(s'_i, s_{-i})$,

$$P_i^\epsilon(s, (s'_i, s_{-i})) = \frac{\epsilon^{-u_i(s'_i, s_{-i})}}{\epsilon^{-u_i(s, s_{-i})} + \epsilon^{-u_i(s'_i, s_{-i})}} = \left(\underbrace{1}_{a=1} - \frac{\epsilon^{u-u'}}{1 + \epsilon^{u-u'}} \right) \underbrace{\epsilon^{u-u'}}_{k=u-u'}.$$

⁷ In the literature, irreducibility is often assumed, but a unique recurrent class will usually suffice, as the process is then effectively irreducible on the state space restricted to the recurrent class. In our case, a unique recurrent class arises ‘for free’ under the conditions of our theorems.

		Player 2	
		A	B
Player 1	A	3, 3, 3	0, 0, 1
	B	0, 1, 0	1, 0, 0
		Player 3: A	

		Player 2	
		A	B
Player 1	A	1, 0, 0	0, 1, 0
	B	0, 0, 1	4, 4, 2
		Player 3: B	

Fig. 2. A three player coordination game, in which Player 3 chooses a matrix. This game has the one-shot deviation property for any regular best response rule.

Define the minimum cost of leaving strategy profile $s \in S$ as

$$c(s) := \min_{s' \in S \setminus \{s\}} c(s, s'). \tag{9}$$

Roughly speaking, $c(s)$ measures how unlikely is the most probable departure from strategy profile s . Denote the set of states most resilient to a single deviation by

$$OS := \left\{ s \in S : c(s) = \max_{s' \in S} c(s') \right\}. \tag{10}$$

Our definitions imply that $c(s) > 0$ if and only if $s \in NE$. Therefore, if $NE \neq \emptyset$, it must be that $OS \subseteq NE$. Importantly, OS is determined by the local properties of the dynamics at each state. This contrasts with SS , which is a globally determined property. Thus, if the sets coincide, we simplify the problem of determining SS .

Definition 1. Given $G, \{P^\epsilon\}_\epsilon$, the one-shot deviation property holds if $SS = OS$.

It will be shown in Section 3 that the one-shot deviation property holds for almost all large games that have a Nash equilibrium, played under any behavioral rules that satisfy our assumptions. This will allow us to give simple characterizations of SS for the most popular behavioral rules in the literature.

Example 1 (Coordination game for which the one-shot deviation property holds).

Consider the coordination game in Fig. 2. Best responses of players 1, 2, 3 are to play the same strategy as players 2, 3, 1 respectively. There are two pure Nash equilibria, $NE = \{(A, A, A), (B, B, B)\}$.

After specifying a game G , it would be natural to specify dynamics $\{P^\epsilon\}_\epsilon$. In this example, we do not do this. Consequently, we cannot determine SS and OS . We can, however, show that $SS = OS$ for any $\{P^\epsilon\}_\epsilon$. That is, for this game, the one-shot deviation property holds for any process satisfying our model assumptions.

Observe that any deviation from either equilibrium can lead to the alternative equilibrium by best response alone, without further costly deviations. For example,

$$(A, A, A) \xrightarrow{\text{deviation}} (B, A, A) \xrightarrow{\text{best resp.}} (B, A, B) \xrightarrow{\text{best resp.}} (B, B, B).$$

Consequently, the order of magnitude (in ϵ) of the probability of a transition from (A, A, A) to (B, B, B) equals the order of magnitude of the probability of the least cost deviation from (A, A, A) . The same is true for transitions from (B, B, B) to (A, A, A) . Consequently, for small ϵ , the process will spend most of the time at the equilibrium most resilient to a single deviation. The one-shot deviation property holds for any behavioral rule satisfying our model assumptions.

2.4. Universal behavioral rules

It can be useful to define rules of behavior before a game is known. This could be as simple as specifying that logit choice be used in every possible game. In contrast, one might specify qualitative variation in behavioral rules depending on the game. For our purposes, we do not require any restrictions in this respect.

Definition 2. A universal behavioral rule \mathcal{P} specifies, for every game G , a process $\mathcal{P}(G) = \{P^\epsilon\}_\epsilon$ satisfying conditions (3) to (6).

The next step is to consider distributions over games with given players and strategy sets. Let $\mathcal{G}(N, S)$ be the set of all games with players N and strategies S . Let Γ map every possible pair (N, S) to a probability measure $\Gamma(N, S)$ on $\mathcal{G}(N, S)$. That is, $\Gamma(N, S)$ is a probability measure on $\{u_i\}_{i \in N}, u_i : S \rightarrow \mathbb{R}$.

A natural set of $\Gamma(N, S)$ to consider are those that, for a given player i , have no bias towards one strategy profile over another. That is, ex-ante, for any $s, s' \in S$, the inequalities $u_i(s) > u_i(s')$ and $u_i(s') > u_i(s)$ are equally likely. Under such $\Gamma(N, S)$, every strict ordering of payoffs that a player might have is equally likely.

Definition 3. Γ is unbiased if, for all (N, S) , for all $i \in N$, $\Gamma(N, S)$ is such that $\{u_i(s)\}_{s \in S}$ are independent identically distributed continuous random variables.

Unbiasedness is not without loss of generality. In particular, it excludes the possibility of payoffs being correlated across players. The study of correlated payoffs is outside the scope of the current paper, but is a potentially interesting avenue for future work.

3. Main result

Our main result is that the one-shot deviation property holds for almost all games with a pure Nash equilibrium. Specifically, for all unbiased distributions over games, universal behavioral rules and size of strategy sets, the property holds asymptotically almost surely as the number of players becomes large.

Theorem 1. *For all unbiased Γ , universal behavioral rules \mathcal{P} , constants $\delta > 0$, integers $m \geq 2$, there exists \bar{n} such that for all (N, S) satisfying $|N| \geq \bar{n}$, $\max_{i \in N} |S_i| \leq m$, the probability that the one-shot deviation property holds conditional on the game having a pure Nash equilibrium is greater than $1 - \delta$.*

Hence, for unbiased Γ , $SS = OS$ for almost all games with a pure Nash equilibrium. For given behavioral rules, this allows us to give simple characterizations of SS in terms of social choice functions. In Section 4, we do this for some popular behavioral rules. Before that, we make some remarks on Theorem 1.

Remark 1. As $|N|$ becomes large, it follows from Rinott and Scarsini (2000, Proposition 2.4 dealing with uncorrelated payoffs) that $\Gamma(N, S)$ places limiting probability $1 - e^{-1} \approx 0.63$ on existence of at least one pure Nash equilibrium.⁸ That is, Theorem 1 gives equilibrium selection $SS = OS \subseteq NE$ in approximately 63% of games.

Remark 2. The proof of Theorem 1 has the following intuition. Starting from $s \in NE$, we use Theorem 1 of Johnston et al. (2025) to show that best response paths exist from profiles close to s to target profile $s' \in NE$. Given this, the aggregate cost of a transition from s to s' is given by $c(s)$. We then show that the Freidlin and Wentzell (1984) tree characterization of SS reduces to considering such transitions and that minimal cost trees are rooted at s that maximize $c(s)$, giving $SS = OS$.

Remark 3. Cost $c(s)$ is a local property of s . The set OS is determined by $c(s)$, $s \in S$. That is, OS is determined by a local property of each strategy profile. This contrasts with SS , which is determined by the tree characterization, a global property. When the one-shot deviation property holds, the two coincide.

Remark 4. Existence of the best response paths in Remark 2 implies that the set of recurrent classes of the unperturbed process P^0 equals NE . Furthermore, considering $OS \neq NE$ and defining the cost of a path of transitions to be the sum of costs along that path, (i) the least cost path from $s \in OS$ to any state in $NE \setminus OS$ has cost $\max_{s \in S} c(s)$; (ii) the highest cost (across $s \in NE \setminus OS$) least cost path to any state in OS has cost $\max_{s \in S \setminus OS} c(s)$. The cost in (i) is greater than the cost in (ii). In the terminology of Ellison (2000), the radius of OS is greater than the coradius of OS . This is the sufficient condition for $SS \subseteq OS$ from Theorem 1 of Ellison (2000). That is, our Theorem 1 implies that Theorem 1 of Ellison (2000) applies to almost all games. To understand why, consider that one-shot methods are simpler but less general than radius-coradius methods, which are in turn simpler but less general than the tree characterization. Thus, when one-shot methods asymptotically approach full generality, radius-coradius methods do too.

Remark 5. Newton and Sawa (2015) show that a one-shot deviation property holds in matching problems. In their model, it is possible to move via pairwise better-response from states close to stable matching m to some stable matching \bar{m} that is closer to a target stable matching m' . Using this, they exploit the structure of stable matchings to prove $SS \subseteq OS$. The current paper deals with normal form games and gives asymptotic results. Notably, intermediate equilibria turn out to be unnecessary in our constructions (see Remark 2), so we have the stronger result $SS = OS$.

4. Stability under different best response rules

In this section, we characterize stochastic stability in large games with randomly drawn payoffs under different behavioral rules. Throughout, the assumption that games are drawn from an unbiased Γ is maintained.

4.1. Uniform deviations

A player i who follows best response with uniform deviations (Kandori et al., 1993; Young, 1993a) chooses a best response with probability $1 - \epsilon$ and chooses a strategy at random with probability ϵ .

$$P_i^\epsilon(s, (s'_i, s_{-i})) = \begin{cases} \frac{1-\epsilon}{|B_i(s)|} + \frac{\epsilon}{|S_i|} & \text{if } s'_i \in B_i(s), \\ \frac{\epsilon}{|S_i|} & \text{otherwise.} \end{cases}$$

The cost function is such that, for all $s, s' = (s'_i, s_{-i})$, $i \in N$,

$$c(s, s') = \begin{cases} 0, & \text{if } s'_i \in B_i(s), \\ 1, & \text{otherwise.} \end{cases} \tag{11}$$

It follows that, from any $s \in NE$, the cost $c(s, s')$ of a deviation to $s' = (s'_i, s_{-i})$, $s'_i \notin B_i(s)$, $i \in N$, is equal to one. That is, for all $s \in NE$, the least cost deviation $c(s) = 1$. For any $s \notin NE$, we have $c(s) = 0$, as there exists a zero cost best response to some profile other than s . Therefore, s is in OS if and only if it is in NE .

⁸ This limiting probability was previously established by Arratia et al. (1989).

Proposition 1. Under uniform deviations, if $NE \neq \emptyset$, then $OS = NE$.

Together with Theorem 1, Proposition 1 implies that, under the universal behavioral rule that plays best response with uniform deviations in every game, for almost all large games that have a pure Nash equilibrium, the set of stochastically stable strategy profiles equals the set of pure Nash equilibria. That is, for most large games, uniform deviations have no power to select between Nash equilibria.

Remark 6. Extending the comparison with the matching literature (see Remark 5), Proposition 1 together with Theorem 1 is analogous to the result of Jackson and Watts (2002b, Theorem 2) that SS equals the set of all stable matchings.

4.2. Payoff-difference based deviations

A player i follows the *logit rule* (Blume, 1993; Alós-Ferrer and Netzer, 2010) if the difference in log probabilities of choosing any two strategies is linear in their utilities. Writing $\eta = -(\log \epsilon)^{-1}$,

$$P_i^\epsilon(s, (s'_i, s_{-i})) = \frac{e^{\frac{1}{\eta} u_i(s'_i, s_{-i})}}{\sum_{s''_i \in S_i} e^{\frac{1}{\eta} u_i(s''_i, s_{-i})}}.$$

The cost function is such that, for all $s, s' = (s'_i, s_{-i}), i \in N$,

$$c(s, s') = \max \left\{ 0, \max_{\bar{s}_i \in \bar{S}_i} u_i(\bar{s}_i, s_{-i}) - u_i(s') \right\}. \tag{12}$$

Logit is a special case of the class of (strictly) *payoff-difference based rules*⁹ that satisfy, for all $s, s' = (s'_i, s_{-i}), i \in N$,

$$c(s, s') = \begin{cases} 0, & \text{if } s'_i \in B_i(s), \\ f \left(\max_{\bar{s}_i \in \bar{S}_i} u_i(\bar{s}_i, s_{-i}) - u_i(s') \right), & \text{otherwise.} \end{cases} \tag{13}$$

for strictly increasing function $f : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$. Logit is the special case $f(x) = x$.

Again, for any $s \notin NE$, we have $c(s) = 0$, as there exists a zero cost best response to some profile other than s . For $s^* \in NE$,

$$c(s^*) = \min_{\substack{i \in N \\ s_i \neq s_i^*}} f \left(\underbrace{\max_{\bar{s}_i \in \bar{S}_i} u_i(\bar{s}_i, s_{-i}^*) - u_i(s_i, s_{-i}^*)}_{=u_i(s^*) \text{ as } s \in NE} \right) \tag{14}$$

$$= \min_{\substack{i \in N \\ s_i \neq s_i^*}} f(u_i(s^*) - u_i(s_i, s_{-i}^*)). \tag{15}$$

Substituting into $OS = \operatorname{argmax}_{s \in S} c(s) = \operatorname{argmax}_{s^* \in NE} c(s^*)$ and using the fact that f is strictly increasing, we obtain the following proposition.

Proposition 2. Under payoff-difference based deviations, if $NE \neq \emptyset$, then

$$OS = \operatorname{arg} \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} u_i(s^*) - u_i(s_i, s_{-i}^*). \tag{16}$$

To better understand this result, we consider *approximate-equilibria*.¹⁰ The set of ρ -best responses by player i to strategy profile s is

$$B_i^\rho(s) := \left\{ \bar{s}_i \in S_i : u_i(\bar{s}_i, s_{-i}) \geq \max_{s'_i \in S_i \setminus \bar{s}_i} u_i(s'_i, s_{-i}) - \rho \right\}, \tag{17}$$

It is standard to have $\rho \geq 0$, in which case $B_i^\rho(s)$ is nonempty. However, for our purposes, we allow the possibility of $\rho < 0$, in which case $B_i^\rho(s)$ may be empty.

Define the set of ρ -NE of G as

$$\rho\text{-NE} := \{s \in S : \forall i \in N, s_i \in B_i^\rho(s)\}, \tag{18}$$

Taking inspiration from the least-core of Maschler et al. (1979), which is defined using the approximate-core, we define the least-NE of G as

$$\text{least-NE} := \bigcap_{\substack{\rho \in \mathbb{R} \\ \rho\text{-NE} \neq \emptyset}} \rho\text{-NE}. \tag{19}$$

⁹ These behavioral rules roughly correspond to *skew-symmetric rules* (Blume, 2003), *payoff-based rules* (Peski, 2010), and *self-regarding payoff-difference-based rules* (Newton, 2021).

¹⁰ These are sometimes referred to as *epsilon-equilibria* in the literature, but as we use ϵ for our perturbation parameter, we use ρ instead.

The definitions of ρ -NE and $B_i^\rho(\cdot)$ imply that ρ -NE is nonempty for large enough ρ . Furthermore, if $\rho' < \rho''$ then ρ' -NE \subseteq ρ'' -NE. Taken together, these two facts imply that least-NE is nonempty.

Note that least-NE is nonempty even if NE is empty. That is, if $NE = \emptyset$, then least-NE is a relaxation of NE. In contrast, if $NE \neq \emptyset$, then the definitions of ρ -NE and NE imply $\emptyset \neq \rho$ -NE \subseteq NE \subseteq 0-NE for some $\rho < 0$, therefore least-NE \subseteq NE. That is, if NE is nonempty, then least-NE is a refinement of NE.

We can now restate Proposition 2.

Proposition 3. *Under payoff-difference based deviations, if $NE \neq \emptyset$, then*

$$OS = \text{least-NE}. \tag{20}$$

Together with Theorem 1, Proposition 3 implies that, under a universal behavioral rule that plays the same payoff-difference based rule in every game, for almost all large games that have a pure Nash equilibrium, the set of stochastically stable strategy profiles is the set of least-Nash equilibria.

Remark 7. Extending the comparison with the matching literature (see Remark 5), Proposition 3 is analogous to the result of Newton and Sawa (2015, Proposition 4.3) that OS equals the least-core. However, Theorem 1 of the current paper implies $SS = OS$, whereas Theorem 3.4 of the cited paper has $SS \subseteq OS$.

4.3. Condition dependent deviations

In this section, non-best response behavior is more likely when a player is currently experiencing low payoffs. An interpretation is that players who earn high payoffs are less inclined to experimentally change their behavior in search of higher payoffs. Deviations are condition dependent (Bilancini and Boncinelli, 2020) if, for all $s, s' = (s'_i, s_{-i}), i \in N$,

$$c(s, s') = \begin{cases} 0, & \text{if } s'_i \in B_i(s), \\ g(u_i(s)), & \text{otherwise.} \end{cases} \tag{21}$$

for strictly increasing function $g : \mathbb{R} \rightarrow \mathbb{R}_{++}$.

Again, for any $s \notin NE$, we have $c(s) = 0$, as there exists a zero cost best response to some profile other than s . For $s^* \in NE$,

$$c(s^*) = \min_{i \in N} g(u_i(s^*)). \tag{22}$$

Define the set of Rawlsian-NE, the strict Nash equilibria that maximize the payoff of the least well off player.

$$\text{Rawlsian-NE} := \arg \max_{s^* \in NE} \min_{i \in N} u_i(s^*). \tag{23}$$

Substituting (22) into $OS = \arg \max_{s \in S} c(s) = \arg \max_{s^* \in NE} c(s^*)$ and using the fact that g is strictly increasing, we obtain the following proposition.

Proposition 4. *Under condition dependent deviations, if $NE \neq \emptyset$, then*

$$OS = \text{Rawlsian-NE}. \tag{24}$$

Together with Theorem 1, Proposition 4 implies that, under a universal behavioral rule that plays the same condition dependent rule in every game, for almost all large games that have a pure Nash equilibrium, the set of stochastically stable strategy profiles is the set of Rawlsian-Nash equilibria.

Remark 8. Extending the comparison with the matching literature (see Remark 5), Bilancini et al. (2020) show that $SS \subseteq OS \subseteq \text{Rawlsian Stable Matchings}$. In contrast, Proposition 4 and Theorem 1 of the current paper imply the equality $SS = OS = \text{Rawlsian-NE}$. The cited paper also shows that if a behavioral rule is not condition dependent, then we can find payoffs that give $SS \not\subseteq \text{Rawlsian Stable Matchings}$. In the current setting, we can find such payoffs even for condition dependent rules, but the probability of such payoffs disappears asymptotically.

5. Discussion

5.1. Speed of convergence

Since the early days of the evolution of conventions literature in economics, there has been interest in the time required to reach SS from states outside of SS.¹¹ Here, we show that for almost all large games, expected convergence time to SS is of the same order of magnitude as the waiting time to exit the states in $S \setminus OS$ that are most resilient to a single deviation.

Denote the highest cost least cost deviation from a strategy outside of OS by

$$\bar{c} := \max_{s \in S \setminus OS} c(s). \tag{25}$$

Under the same conditions and probabilities as Theorem 1, the expected waiting time to enter OS will be of order of magnitude $\varepsilon^{-\bar{c}}$.

¹¹ See, for example, Ellison (2000, 1993), Arieli and Young (2016), Kreindler and Young (2013, 2014), Arieli et al. (2020), Newton and Angus (2015).

Theorem 2. For all unbiased Γ , universal behavioral rules \mathcal{P} , constants $\delta > 0$, integers $m \geq 2$, there exists \bar{n} such that for all (N, S) satisfying $|N| = n \geq \bar{n}$, $\max_{i \in N} |S_i| \leq m$, the probability, conditional on the game having a pure Nash equilibrium, that

- There exists $\kappa(n)$, $\hat{\epsilon}$ such that for all $\epsilon < \hat{\epsilon}$, starting from any strategy profile, the expected waiting time for the process to enter OS is less than $\kappa(n)\epsilon^{-\bar{c}}$,

is greater than $1 - \delta$.

Observe that for dynamics such as the ones discussed in Section 4, \bar{c} will not increase unboundedly in the number of players.¹² This does not imply that waiting time is also bounded, as $\kappa(n)$ may increase in $n = |N|$. Instead, it suggests that factors specific to stochastic stability (as opposed to best response dynamics alone) will not lead to exponential increase in convergence times.

5.2. Medium term behavior and the invariant measure

Aside from NE and SS , we can ask what are the likely trajectories that the process will follow across NE . This information can be used to decompose the process and determine further information about the invariant measure.¹³

Define $r(s)$ such that the invariant measure on $s \in S$, $\mu^\epsilon(s)$, is of order $\epsilon^{r(s)}$,

$$r(s) := \lim_{\epsilon \rightarrow 0} \frac{\log \mu^\epsilon(s)}{\log \epsilon}. \tag{26}$$

We deduce a simple expression for $r(\cdot)$ that applies to most large games.

Let P^{NE} be the Markov process on state space NE such that $P^{NE}(s, s')$, $s, s' \in NE$, is the probability under P that, from profile s at $t = 0$, the first strategy profile reached in the set NE at time $t > 0$ is s' . Define an accompanying cost function $c^{NE}(\cdot, \cdot)$ and let $\sigma(\cdot)$ be the least cost correspondence,

$$\sigma(s) := \arg \min_{s' \in NE \setminus s} c^{NE}(s, s'). \tag{27}$$

A cycle is a subset $\Gamma \subseteq NE$ such that $\Gamma = \bigcup_{m=0}^{\bar{m}} \{s^m\}$ for some sequence $s^0, \dots, s^{\bar{m}}$ satisfying $s^1 \in \sigma(s^0), s^2 \in \sigma(s^1), \dots, s^0 \in \sigma(s^{\bar{m}})$. A cycle Γ is closed if it is closed under σ in that $\sigma(\Gamma) \subseteq \Gamma$. From states within a closed cycle, transitions within the cycle happen with higher probability than transitions to states outside of the cycle. Thus, cycles can be considered to model behavior in the medium term, in contrast to the short term (the unperturbed process) and the long term (stochastic stability).

Under the conditions of Theorem 1, with high probability there exists only one closed cycle and it comprises all profiles in NE , including those in $OS = SS$ (see proof of Theorem 3). That is, from any starting point, P^{NE} is as likely to reach OS (long term stability under the one-shot deviation property) as it is to move around a subset of NE that doesn't include OS (possible medium term dynamics). Another way of putting this is that the medium term and the long term operate on the same timescale.

An implication of $\Gamma = NE$ being a cycle under P^{NE} is that, for $s \in NE$, the function $r(\cdot)$ that gives the order of magnitudes of our invariant probabilities under P has a particularly simple expression.

Theorem 3. For all unbiased Γ , universal behavioral rules \mathcal{P} , constants $\delta > 0$, integers $m \geq 2$, there exists \bar{n} such that for all (N, S) satisfying $|N| \geq \bar{n}$, $\max_{i \in N} |S_i| \leq m$, the probability, conditional on the game having a pure Nash equilibrium, that

- For all $s \in NE$, we have $r(s) = \max_{s^* \in NE} c(s^*) - c(s)$,

is greater than $1 - \delta$.

5.3. Robust stochastic stability

Under the processes we have considered, a single player updates his strategy each period. This assumption aids clarity of exposition, but is not necessary to our results. Recall our assumption that, each period, player i is selected to update with probability $\pi(i) > 0$. As long as π has full support on N , the cost function $c(\cdot, \cdot)$ and the results of the paper are independent of the exact probabilities given by π .

Consider the following alternative. Each period, a set of players $J \subseteq N$ is selected with probability $\pi(J)$ to update their strategies. That is, π is a probability measure on subsets of N . Assume that $\pi(\{i\}) > 0$ for all $i \in N$, so there is still positive probability that player i is the only updating player. Let Π be the set of all such distributions. Players in J update their strategies independently of one another.

$$\text{For } s, s', J = \{i \in N : s_i \neq s'_i\}, \quad P^\epsilon(s, s') = \pi(J) \prod_{i \in J} P_i^\epsilon(s, (s'_i, s_{-i})), \tag{28}$$

¹² Of course, we could specify a universal behavioral rule that explicitly specifies probabilities dependent on $|N|$ so as to give, for example, $c(s) \approx e^{|N|}$ for $s \in NE$, but this would be unusual.

¹³ These methods, based on Freidlin and Wentzell (1984), have recently elicited interest within economics (Cui and Zhai, 2010; Levine and Modica, 2016; Newton and Sandholm, 2022).

which gives the cost function

$$c^\pi(s, s') = \sum_{i \in J} c(s, (s'_i, s_{-i})), \tag{29}$$

Consequently, stochastic stability can depend on π . For $\pi \in \Pi$, let SS^π denote the set of stochastically stable states. A profile s is *robustly stochastically stable* (Alós-Ferrer and Netzer, 2015) if it is stochastically stable under all π .

$$\text{robust-}SS = \{s \in S : \text{for all } \pi \in \Pi, s \in SS^\pi\}. \tag{30}$$

The probabilities of transition paths used in the proof of Theorem 1 depend on $c(s)$ for $s \in NE$. It follows from (29) that lower cost transitions cannot be found.

$$c^\pi(s) = \min_{\text{by defn}}_{s' \in S \setminus s} c^\pi(s, s') = \min_{(29) \text{ s.t. } i \in N} c^\pi(s, s') = \min_{(29) \text{ s.t. } i \in N} c(s, s') = \min_{\text{by defn}} c(s) \tag{31}$$

As a consequence, the most likely transition paths between profiles in NE are unaffected by π and we have the following proposition.¹⁴

Proposition 5. *Theorem 1 continues to hold if we replace SS with robust- SS in the definition of the one-shot deviation property.*

5.4. Weakly regular rules and probit choice

Instead of regular behavioral rules, consider *weakly regular* behavioral rules, under which, rather than satisfy (4), $\{P_i^\epsilon(s, s')\}_\epsilon$ satisfy

$$P_i^\epsilon(s, s') = \epsilon^{k+o(1)} \quad \text{for some } k > 0, \tag{32}$$

where k may depend on s, s', i , but not on ϵ ; and $o(1)$ represents a term that vanishes as $\epsilon \rightarrow 0$. The tree characterization of SS under regular rules (Young, 1993a; Kandori et al., 1993) characterizes a weak version of stochastic stability under weakly regular rules (Sandholm, 2010),¹⁵

$$\text{weak-}SS := \begin{cases} \left\{ s \in S : \lim_{\epsilon \rightarrow 0} \frac{\log \mu^\epsilon(s)}{\log \epsilon} = 0 \right\}, & \text{if there exists } \hat{\epsilon} \text{ such that } P^\epsilon \text{ has a unique} \\ \emptyset, & \text{recurrent class for all } \epsilon \in (0, \hat{\epsilon}), \\ & \text{otherwise.} \end{cases} \tag{33}$$

It follows that, if we define a *weak one-shot deviation property* by substituting weak- SS for SS in Definition 1, then a version of Theorem 1 will hold for this weak one-shot deviation property under weakly regular rules. The proof will be identical to that of Theorem 1 except that weak- SS will replace SS .

An example of a weakly regular but not regular rule is multinomial probit (Myatt and Wallace, 2003; Dokumaci and Sandholm, 2011),

$$P_i^\epsilon(s, (s'_i, s_{-i})) = \mathbb{P} \left(\bigcap_{\tilde{s}_i \in S_i} \left\{ u_i(s'_i, s_{-i}) + Z_{s'_i}^\epsilon \geq u_i(\tilde{s}_i, s_{-i}) + Z_{\tilde{s}_i}^\epsilon \right\} \right),$$

where, for all $s_i \in S_i$, random variable $Z_{s_i}^\epsilon$ is normal with mean 0 and variance $\sigma^2 / (-\log \epsilon)$. Writing $m = -\log \epsilon$, we can regard $Z_{s_i}^\epsilon$ as an average of m signals, each normal with mean 0 and variance σ^2 . Thus $\epsilon \rightarrow 0$ corresponds to $m \rightarrow \infty$, with more signals associated with lower uncertainty about payoffs.

The cost function $c(\cdot, \cdot)$ for multinomial probit is rather complicated (see Appendix A.5). Qualitatively, Dokumaci and Sandholm (2011) write

...the multinomial probit model introduces a novel feature: the rate of decay in the probability of choosing a suboptimal strategy is neither independent of payoffs [...], nor dependent only on the gap between its payoff and the optimal strategy's payoff [...], but can depend on the gaps between its payoff and those of all better performing strategies.

Consider $s \in NE$. The least cost deviation for player i is to whichever s'_i gives the second highest payoff. The only better performing strategy is s_i , the equilibrium strategy itself. Therefore, when the (weak) one-shot deviation property holds, (weak) stochastic stability is determined by payoff-differences. The additional complexity introduced by multinomial choice disappears when we consider one-shot deviation.

Proposition 6. *Under multinomial probit, if $NE \neq \emptyset$, then*

$$OS = \text{least-}NE \tag{34}$$

¹⁴ Interestingly, if we were to require that players update in a deterministic order, as in the clockwork sequence of Heinrich et al. (2023), then certain transition paths would be ruled out and our result would no longer hold. However, even in this case, our result can be recovered by adding an arbitrary inertia probability $\alpha \in (0, 1)$ with which any updating player maintains his current strategy.

¹⁵ By definition, weak- $SS \supseteq SS$. For regular behavioral rules, weak- $SS = SS$.

5.5. Payoff distributions with atoms

So far, we have considered generic payoff distributions in which there are no atoms. In this section, we discuss non-generic payoff distributions and the effect of atoms on the analysis. This can be thought of as a first step towards extending our analysis beyond unbiased, continuous Γ . To consider atoms in a concise way, we flatten some range of possible payoffs to create an atom. Specifically, map realized payoffs as follows,

$$\bar{u}_i(s) := \begin{cases} u_i(s) & \text{if } u_i(s) < 0, \\ 0 & \text{if } u_i(s) \in [0, \bar{u}], \\ u_i(s) - \bar{u} & \text{if } u_i(s) > \bar{u}, \end{cases} \tag{35}$$

where $\bar{u} > 0$ determines the range of payoffs flattened to create an atom. To ensure flattened payoff distributions include atoms, assume a continuous distribution F with a support that includes zero in its interior, such that for all (N, S) , $\Gamma(N, S)$ is such that $\{u_i(s)\}_{i \in N, s \in S}$ are independent draws from F .

We shall use subscripts to indicate concepts taken under each set of payoffs. For example, NE_u and $B_{i,u}(\cdot)$ are the strict Nash equilibria and best responses under payoffs $u = \{u_i(s)\}_{i \in N, s \in S}$, whereas $NE_{\bar{u}}$ and $B_{i,\bar{u}}(\cdot)$ are the equivalent objects under payoffs $\bar{u} = \{\bar{u}_i(s)\}_{i \in N, s \in S}$. Observe that the flattening of payoffs in (35) weakly increases the set of best responses from any given strategy profile.

Remark 9. For all $i \in N, s \in S$, we have $B_{i,u}(s) \subseteq B_{i,\bar{u}}(s)$.

To simplify exposition, assume that $|S_i| = m$ for all $i \in N$. It turns out that with atoms, there are asymptotically almost surely no strict Nash equilibria. The number of strategy profiles grows at rate $m^{|N|}$, whereas the ex-ante probability of any given profile being a strict Nash equilibrium decreases faster than $m^{-|N|}$. Thus, the expected number of strict Nash equilibria approaches zero as $|N| \rightarrow \infty$.¹⁶

Proposition 7. Fixing $m \geq 2$, we have $Prob_{\Gamma(N,S)}[NE_{\bar{u}} = \emptyset] \rightarrow 1$ as $|N| \rightarrow \infty$.

However, despite the set of strict Nash equilibria $NE_{\bar{u}}$ being generically empty under \bar{u} , it transpires that we can still prove an analogy of our main result conditional on $NE_{\bar{u}} \neq \emptyset$. The proof of [Theorem 1](#) (see [Remark 2](#)) relies on paths of best responses from $s' \notin NE_u$ to $s^* \in NE_u$ under payoffs u . [Remark 9](#) implies that such a path is also a path of best responses under payoffs \bar{u} . By showing that, for small enough \bar{u} , the probability of such paths not existing vanishes faster than $Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset]$, we extend our main result to games with atoms.^{17,18}

Theorem 4. For all unbiased Γ , all universal behavioral rules \mathcal{P} , constants $\delta > 0$, integers $m \geq 2$, there exist \bar{n} and $\bar{U} > 0$ such that for all (N, S) satisfying $|N| \geq \bar{n}, |S_i| = m$ for all $i \in N, \bar{u} \leq \bar{U}$, the probability that, under payoffs \bar{u} , the one-shot deviation property holds conditional on $NE_{\bar{u}} \neq \emptyset$ is greater than $1 - \delta$.

Flattened best responses also allow us to comment on the generic situation in which $NE_u \neq \emptyset$ but $NE_{\bar{u}} = \emptyset$ in accordance with [Proposition 7](#). $NE_{\bar{u}} = \emptyset$ implies $c_{\bar{u}}(s) = 0$ for all $s \in S$, so $OS_{\bar{u}} = S$. That is, $OS_{\bar{u}}$ has no selective power.

Consider $s^*, s^{**} \in NE_u$. As s^{**} is a strict Nash equilibrium under u , [Remark 9](#) implies that it is a Nash equilibrium under \bar{u} . However, as $NE_{\bar{u}} = \emptyset$, we have $s^{**} \notin NE_{\bar{u}}$. That is, s^{**} is a Nash equilibrium, but not a strict Nash equilibrium under \bar{u} . As $s^{**} \notin NE_{\bar{u}}$, there exist $i \in N, s'_i \in S_i, s'_i \neq s^{**}_i$ such that $s'_i \in B_{i,\bar{u}}(s^{**})$.

Now note that as $s^{**} \in NE_u, s'_i \neq s^{**}_i$ implies that $s'_i \notin B_{i,u}(s^{**}) = B_{i,u}(s')$. Therefore, $s' \notin NE_u$. Thus, in large games there is almost always a path of best responses s', \dots, s^* under both u and \bar{u} . Combining, we have a path of best responses s^{**}, s', \dots, s^* under \bar{u} . In summary, from any $s \in S$, there exists a path of best responses to s^* under \bar{u} .

Therefore, the unperturbed process P^0 for payoffs \bar{u} has a unique recurrent class of which NE_u is a strict subset. Every state in this recurrent class is stochastically stable for all families of perturbed best response dynamics $P = \{P^\epsilon\}_\epsilon$.

This concludes the main body of the paper.

Data availability

No data was used for the research described in the article.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT in order to improve readability. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

¹⁶ This is shown in [Amiet et al. \(2021\)](#) for $m = 2$.

¹⁷ Bounds on the former quantity are from [Johnston et al. \(2025\)](#). Bounds on the latter quantity are derived by adapting results of [Arratia et al. \(1989\)](#) that were previously applied to games without atoms by [Rinott and Scarsini \(2000\)](#).

¹⁸ Consistent with the rest of this section, the proof of [Theorem 4](#) assumes Γ has every payoff in $\{u_i(s)\}_{i \in N, s \in S}$ drawn from the same distribution F . Nothing important is lost. If each player had a distinct distribution F_i , then the quantity $\beta = Prob_{\Gamma(N,S)}[\{s_i\} = B_{i,\bar{u}}(s)] \leq 1/m$ would differ across players, but would still approach $1/m$ as $\bar{u} \rightarrow 0$, which is what matters for the proof.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Jonathan Newton reports financial support was provided by JSPS KAKENHI Grants-in-Aid for Scientific Research. Ryoji Sawa reports financial support was provided by JSPS KAKENHI Grants-in-Aid for Scientific Research. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Proofs

A.1. Additional definitions

Definition 4. An ordinal game $H = (N, \{S_i\}_{i \in N}, \{\succeq_i\}_{i \in N})$ consists of a finite set of players $N = \{1, \dots, n\}$, and for each player i , a finite set of strategies $S_i = \{1, \dots, |S_i|\}$ and a total order \succeq_i on S_i , where $S = \times_{i \in N} S_i$. Let \succ_i denote the asymmetric part of \succeq_i . Let $\overline{H}(N, S)$ be the (finite) set of all ordinal games with player set N and strategies S .

Definition 5. The ordinal game associated with game $G = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ is the unique ordinal game $H(G) = (N, \{S_i\}_{i \in N}, \{\succeq_i\}_{i \in N})$ such that $s \succeq_i s'$ if and only if $u_i(s) \geq u_i(s')$.

Definition 6. The ordinal game $H = (N, \{S_i\}_{i \in N}, \{\succeq_i\}_{i \in N})$ is generic if, for all $i \in N$, $(s_i, s_{-i}), (s'_i, s_{-i}) \in S$, either $(s_i, s_{-i}) \succ_i (s'_i, s_{-i})$ or $(s'_i, s_{-i}) \succ_i (s_i, s_{-i})$. Let $\mathcal{H}(N, S)$ be the set of all generic ordinal games with player set N and strategies S .

Definition 7. For given H , s_i is a best response to s_{-i} if $(s_i, s_{-i}) \succeq_i (s'_i, s_{-i})$ for all $s'_i \in S_i$.

Definition 8. s is a pure Nash equilibrium of $H = (N, \{S_i\}_{i \in N}, \{\succeq_i\}_{i \in N})$ if, for all $i \in N$, s_i is a best response to s_{-i} .

Remark 10. s is a pure Nash equilibrium of G if and only if s is a pure Nash equilibrium of the associated ordinal game $H(G)$.

Definition 9. The best response graph of H is the directed graph (S, \rightarrow) with vertex set S and edge set \rightarrow such that $s \rightarrow s'$ if and only if there exists $i \in N$ such that

- $s_{-i} = s'_{-i}$,
- s_i is not a best response to s_{-i} ,
- s'_i is a best response to s_{-i} .

Note that a vertex $s \in S$ is a Nash equilibrium if and only if it has no outgoing edges. H is connected if it has at least one Nash equilibrium and from every non-Nash equilibrium strategy profile, to every Nash equilibrium, there exists a path on the best response graph.

A.2. Lemmas

The following lemma builds upon [Theorem 1 of Johnston et al. \(2025\)](#), which bounds the proportion of ordinal games with a pure Nash equilibrium that are connected. The distribution Γ over games in our setting implies a distribution over the associated ordinal games. When Γ is unbiased, every generic ordinal game occurs with equal probability, so the bound in [Johnston et al. \(2025\)](#) also bounds the probability that a game selected according to Γ is connected.

Lemma 1. Fix unbiased Γ , constants $\delta > 0$, integers $m \geq 2$. There exists \bar{n} such that for all (N, S) such that $|N| \geq \bar{n}$, $\max_{i \in N} |S_i| \leq m$, the probability under $\Gamma(N, S)$ that $H(G)$ is generic and connected conditional on having at least one pure Nash equilibrium is greater than $1 - \delta$.

Proof.

Step 1

Consider a given (N, S) . By definition, $\Gamma(N, S)$ is a probability measure over games $\mathcal{G}(N, S)$. As H maps each $G \in \mathcal{G}(N, S)$ to $H(G) \in \overline{H}(N, S)$, $\Gamma(N, S)$ also gives a probability measure over ordinal games $\overline{H}(N, S)$.

Step 2

As Γ is unbiased, $\Gamma(N, S)$ is such that, for any given $i \in N$, $\{u_i(s)\}_{s \in S}$ are independent identically distributed continuous random variables.

It follows that, for given $i \in N$, $s, s' \in S$, $s \neq s'$, we have that $u_i(s)$ and $u_i(s')$ are iid and continuous, therefore $u_i(s) = u_i(s')$ with probability zero.

As N, S are finite, this implies that with probability one, G is such that for all $i \in N$, $s, s' \in S$, $s \neq s'$, we have $u_i(s) \neq u_i(s')$.

Combined with the definition of $H(G)$, this implies that, with probability one, $H(G)$ is such that for all $i \in N$, $s, s' \in S$, $s \neq s'$, we have either $s \succ_i s'$ or $s' \succ_i s$. In other words, with probability one, $H(G)$ is generic, $H(G) \in \mathcal{H}(N, S)$.

Therefore, the probability that $H(G)$ satisfies any given property X equals the probability that $H(G)$ is generic and satisfies property X .

Step 3

The definition of connectedness implies that any connected H has at least one pure Nash equilibrium. Therefore, the probability under $\Gamma(N, S)$ that $H(G)$ is connected, conditional on having at least one Nash equilibrium equals

$$\frac{Prob_{\Gamma(N,S)}[H(G) \text{ is connected}]}{Prob_{\Gamma(N,S)}[H(G) \text{ has at least one pure Nash eq}^m]}, \tag{A.1}$$

which by Step 2 equals

$$\frac{Prob_{\Gamma(N,S)}[H(G) \text{ is generic and connected}]}{Prob_{\Gamma(N,S)}[H(G) \text{ is generic and has at least one pure Nash eq}^m]}. \tag{A.2}$$

Step 4

Unbiasedness of Γ implies that every generic ordinal game in $\mathcal{H}(N, S)$ occurs with equal probability under $\Gamma(N, S)$. Thus, for sets $\mathbf{H}, \mathbf{H}' \subseteq \mathcal{H}(N, S)$,

$$\frac{Prob_{\Gamma(N,S)}[H(G) \in \mathbf{H}]}{Prob_{\Gamma(N,S)}[H(G) \in \mathbf{H}']} = \frac{|\mathbf{H}|}{|\mathbf{H}'|}. \tag{A.3}$$

Therefore, (A.2) equals

$$\frac{|\{H \in \mathcal{H}(N, S) : H \text{ is connected}\}|}{|\{H \in \mathcal{H}(N, S) : H \text{ has at least one pure Nash eq}^m\}|}. \tag{A.4}$$

Step 5

Theorem 1 of Johnston et al. (2025) states that there exists $c > 0$ such that if $n = |N|$ is large enough relative to $\max_{i \in N} |S_i|$, then (A.4) is greater than or equal to $1 - e^{-cn}$. In statement of the lemma, we bound $\max_{i \in N} |S_i| \leq m$, therefore there exists \bar{n} such that, for $|N| \geq \bar{n}$, (A.4) is greater than or equal to $1 - e^{-c|N|}$.

Furthermore, for any $\delta > 0$, we can let \bar{n} be large enough that $1 - e^{-c\bar{n}} > 1 - \delta$. This completes the proof. \square

The next step is to prove that, for any of the dynamics that we consider, the one-shot deviation property holds for any G such that $H(G)$ is generic and connected.

Lemma 2. *If $H(G)$ is generic and connected, then the OS deviation property holds for G and any $\{P^\epsilon\}_\epsilon$ satisfying our model assumptions.*

Proof.

Step 1

Consider the weighted, directed graph $(S, c(\cdot, \cdot))$ with vertex set S and the weight of the edge from s to s' given by $c(s, s')$.

A tree rooted at s is a subgraph of $(S, c(\cdot, \cdot))$ such that there exists a unique path from every $s' \neq s$ to s and the graph has no cycles. The stochastic potential of s is the minimum sum of edge weights across all trees rooted at s .

The assumption that $\{P^\epsilon\}_\epsilon$ is non-deterministic implies that, from any $s \in NE$, with positive probability the process moves to $s' = (s'_i, s_{-i})$ for some $s'_i \neq s_i$. Furthermore, as $s \in NE$, $B_i(s') = B_i(s) = \{s_i\}$, therefore $s' \notin NE$. Connectedness implies that there exists a finite path of best responses from any $s' \notin NE$ to any given target $s^* \in NE$. Therefore, P^ϵ has a unique recurrent class for $\epsilon > 0$.

If $P = \{P^\epsilon\}_\epsilon$ is regular and has a unique recurrent class for $\epsilon > 0$, then $s^* \in SS$ if and only if s^* minimizes stochastic potential across all strategy profiles in $s \in S$ (see, e.g. Sandholm, 2010, Theorem 12.A.2).

Step 2

Let s^* minimize stochastic potential across $s \in S$. Take a tree \mathcal{T}^0 rooted at s^* with sum of edge weights equal to the stochastic potential of s^* .

Note that genericity implies that the set of pure Nash equilibria equals the set of strict Nash equilibria NE .

Case 1: $s^* \in NE$.

Let $s^1 \in \operatorname{argmin}_{s \in S \setminus s^*} c(s^*, s)$. Thus, $c(s^*, s^1) = c(s^*)$. Add the edge $s^* \rightarrow s^1$ to \mathcal{T}^0 to obtain the graph \mathcal{T}^1 .

For some $i \in N$, $s_{-i}^* = s_{-i}^1$ and $s_i^* \neq s_i^1$. If this were not the case, then we would have $c(s^*, s^1) = \infty$. However, by the assumption that the perturbed process is non-deterministic, there exists $s \neq s^*$ such that $c(s^*, s) < \infty$.

It follows that s^1 is not a NE. Indeed, by genericity, there is a unique best response for i to s_{-i}^* , so $s^1 = (s_i^1, s_{-i}^*)$ being an NE would contradict $s^* = (s_i^*, s_{-i}^*)$ being a NE.

Case 2: $s^* \notin NE$.

Set $s^1 = s^*$ and let $\mathcal{T}^1 = \mathcal{T}^0$.

Considering both cases, we see that the sum of edge weights in \mathcal{T}^1 equals the sum of edge weights in \mathcal{T}^0 plus $c(s^*)$.

Step 3

Connectedness of $H(G)$ implies that, for all $s' \in NE$, there exists a finite sequence $s^1, \dots, s^T = s'$ such that for each $t = 1, \dots, T - 1$, there exists $i \in N$ such that $s_i^{t+1} \neq s_i^t$, $s_{-i}^{t+1} = s_{-i}^t$, and s_i^{t+1} is a best response to s_{-i}^t .

Adjust \mathcal{T}^1 by first removing any edges leaving s^1, \dots, s^T , then, for $t = 1, \dots, T - 1$, adding edges $s^t \rightarrow s^{t+1}$. The resulting graph \mathcal{T}^T is a tree rooted at s^T .

As our unperturbed process follows best response, we have that $c(s^t, s^{t+1}) = 0$ for $t = 1, \dots, T - 1$. It follows that the sum of edge weights in \mathcal{T}^T is less than or equal to the sum of edge weights in \mathcal{T}^1 minus $c(s^T)$.

Combining with Step 2, this implies that the sum of edge weights in \mathcal{T}^T is less than or equal to the sum of edge weights in \mathcal{T}^0 plus $c(s^*)$ minus $c(s^T)$.

Consider $s^* \notin OS$. The definition of OS and $NE \neq \emptyset$ together imply $\emptyset \neq OS \subseteq NE$, therefore we can choose s^T such that $s^T \in OS$. By definition of OS , this implies that $c(s^*) - c(s^T) < 0$. Therefore, the sum of edge weights of \mathcal{T}^T is strictly less than the sum of edge weights of \mathcal{T}^0 . That is, s^T has strictly lower stochastic potential than s^* , a contradiction.

Therefore, it must be that $s^* \in OS$ and any $s^T \in OS$ has the same stochastic potential as s^* . Therefore, $SS = OS$.

□

A.3. Proof of Theorem 1

Proof of Theorem 1.

Fix an unbiased Γ , universal behavioral rule \mathcal{P} , constant $\delta > 0$, integer $m \geq 2$.

Combining Lemmas 1 and 2, we have that there exists \bar{n} such that for all (N, S) such that $|N| \geq \bar{n}$, $\max_{i \in N} |S_i| \leq m$, the probability under $\Gamma(N, S)$ that the OS deviation property holds for G and any $\{P^\epsilon\}_\epsilon$ satisfying our model assumptions, conditional on G having at least one pure Nash equilibrium, is greater than $1 - \delta$.

As the above statement is true for any process $P = \{P^\epsilon\}_\epsilon$ satisfying our model assumptions, it must be true for the process $\mathcal{P}(G)$ specified by the universal behavioral rule \mathcal{P} . This completes the proof.

A.4. Proofs of Propositions 1, 2, 3, 4

First, recall some facts common to all results proved in this section. These facts are common to all behavioral rules satisfying our model assumptions.

(i) By definition of $s^* \notin NE$, there exists $i \in N$, $s_i \neq s_i^*$ such that $s_i \in B_i(s^*)$, $c(s^*, (s_i, s_{-i}^*)) = 0$. Therefore, $c(s^*) = 0$.

(ii) By definition of $s^* \in NE$, for all $i \in N$, $B_i(s^*) = \{s_i^*\}$. Therefore for all $s \neq s^*$, $c(s^*, s) > 0$. Therefore, $c(s^*) > 0$.

(iii) Consider s^* , s such that for $i, j \in N$, $i \neq j$, $s_i \neq s_i^*$, $s_j \neq s_j^*$. Then, $c(s^*, s) = \infty$.

Facts (i) and (ii) imply that, if $NE \neq \emptyset$, then $OS \subseteq NE$. Fact (iii) implies that, to determine OS , it suffices to know $c(s^*, s)$ for all $s^* \in NE$, $s = (s_i, s_{-i}^*)$, $i \in N$. Specifically,

$$OS = \arg \max_{s^* \in NE} c(s^*) = \arg \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} c(s^*, (s_i, s_{-i}^*)). \tag{A.5}$$

Proof of Proposition 1.

Consider $s^* \in NE$. For all $i \in N$, $s_i \neq s_i^*$, (11) implies that $c(s^*, (s_i, s_{-i}^*)) = 1$. Substituting into (A.5), we obtain $OS = NE$.

Proof of Proposition 2.

Consider $s^* \in NE$. For all $i \in N$, $s_i \neq s_i^*$, (13) implies that

$$\begin{aligned} c(s^*, (s_i, s_{-i}^*)) &= f \left(\underbrace{\max_{\tilde{s}_i \in S_i} u_i(\tilde{s}_i, s_{-i}^*)}_{=u_i(s^*) \text{ as } s^* \in NE} - u_i(s_i, s_{-i}^*) \right) \\ &= f \left(u_i(s^*) - u_i(s_i, s_{-i}^*) \right). \end{aligned} \tag{A.6}$$

Substituting into (A.5),

$$\begin{aligned} OS &= \arg \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} f \left(u_i(s^*) - u_i(s_i, s_{-i}^*) \right) \\ &\stackrel{\substack{= \\ \text{as } f \\ \text{strictly} \\ \text{increasing}}}{=} \arg \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} u_i(s^*) - u_i(s_i, s_{-i}^*). \end{aligned} \tag{A.7}$$

Proof of Proposition 3.

From the definition of $B_i^\rho(\cdot)$, s_i' is a ρ -best response by player i to s' if and only if,

$$u_i(s_i') \geq \max_{s_i \neq s_i'} u_i(s_i, s_{-i}') - \rho, \tag{A.8}$$

which rearranges to

$$-\rho \leq u_i(s_i') - \max_{s_i \neq s_i'} u_i(s_i, s_{-i}') = \min_{s_i \neq s_i'} [u_i(s_i') - u_i(s_i, s_{-i}')]. \tag{A.9}$$

Therefore, from the definition of ρ -NE, s' is an ρ -NE if and only if (A.9) holds for all $i \in N$. That is,

$$-\rho \leq \min_{\substack{i \in N \\ s_i \neq s'_i}} [u_i(s') - u_i(s_i, s'_{-i})]. \tag{A.10}$$

Consider

$$\rho^* := - \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} [u_i(s^*) - u_i(s_i, s_{-i}^*)]. \tag{A.11}$$

Then, substituting into (A.10), s' is a ρ^* -NE if and only if

$$\max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} [u_i(s^*) - u_i(s_i, s_{-i}^*)] \leq \min_{\substack{i \in N \\ s_i \neq s'_i}} [u_i(s') - u_i(s_i, s'_{-i})]. \tag{A.12}$$

Inequality (A.12) holds if and only if

$$s' \in \arg \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} [u_i(s^*) - u_i(s_i, s_{-i}^*)] = OS. \tag{A.13}$$

For $\rho > \rho^*$, (A.12), (A.13) and the definition of ρ -NE imply that $OS \subseteq \rho$ -NE.

For $\rho < \rho^*$, (A.12) cannot be satisfied, therefore ρ -NE = \emptyset .

Therefore,

$$\text{least-NE} = \bigcap_{\substack{\rho \in \mathbb{R} \\ \rho\text{-NE} \neq \emptyset}} \rho\text{-NE} = (\rho^*\text{-NE}) \cap \left(\bigcap_{\rho > \rho^*} \rho\text{-NE} \right) = OS. \tag{A.14}$$

Proof of Proposition 4.

Consider $s^* \in NE$. For all $i \in N$, $s_i \neq s_i^*$, (21) implies that

$$c(s^*, (s_i, s_{-i}^*)) = g(u_i(s^*)). \tag{A.15}$$

Substituting into (A.5),

$$\begin{aligned} OS &= \arg \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} g(u_i(s^*)) = \arg \max_{s^* \in NE} \min_{i \in N} g(u_i(s^*)) \\ &\stackrel{\substack{\text{as } g \\ \text{strictly} \\ \text{increasing}}}{=} \arg \max_{s^* \in NE} \min_{i \in N} u_i(s^*). \end{aligned} \tag{A.16}$$

A.5. Proofs of results in Section 5

Proof of Theorem 2.

By Lemma 1, with probability greater than $1 - \delta$, $H(G)$ is generic and connected. Henceforth, assume this to be the case.

Fix some $s' \in OS \subseteq NE$. Consider the sequences in Steps 2 and 3 of the proof of Lemma 2. Specifically, for any $s^* \notin NE$, pick one such sequence $s^* = s^1, \dots, s^T$, and for any $s^* \in NE \setminus OS$, pick one such sequence s^*, s^1, \dots, s^T . Let \bar{T} be the maximum T amongst the sequences we have chosen.

Recall that, in each period, the probability that the updating player is player i is given by $\pi(i)$. Write $\underline{\pi} = \min_{i \in N} \pi(i)$. Full support of π implies $\underline{\pi} > 0$.

Consider any $s^* \notin OS$ and associated s^1, \dots, s^T , regardless of whether $s^* = s^1$. Starting from s^1 , the probability of transition sequence s^1, \dots, s^T is

$$\prod_{t=1}^{T-1} P^\epsilon(s^t, s^{t+1}) \underset{\substack{\text{for small} \\ \text{enough } \epsilon}}{>} \left(\frac{1}{2}\underline{\pi}\right)^{T-1} \underset{\substack{\text{by defn} \\ \text{of } T}}{\geq} \left(\frac{1}{2}\underline{\pi}\right)^{\bar{T}-1} =: \kappa_1, \tag{A.17}$$

where the first inequality follows because, for each t , the probability of choosing the correct player (to get from s^t to s^{t+1}) to best respond is at least $\underline{\pi}$ and the probability of this player best responding is at least $1/2$ for small enough ϵ .¹⁹ Genericity implies that best responses are unique.

Next, consider $s^* \in NE \setminus OS$, so that $s^* \neq s^1$. Let i be such that $s_i^* \neq s_i^1$. Then,

$$\begin{aligned} P^\epsilon(s^*, s^1) &= \pi(i) P_i^\epsilon(s^*, s^1) \underset{\substack{\text{for } a, k > 0 \\ \text{by (4)}}}{=} \pi(i) (a + o(1))\epsilon^k \\ &\underset{\substack{\text{as } \epsilon \rightarrow 0 \\ \text{and } o(1) \rightarrow 0}}{>} \pi(i) \frac{a}{2} \epsilon^k \underset{\substack{\text{by defn} \\ \text{of } \underline{\pi}}}{\geq} \frac{\underline{\pi}}{2} \epsilon^k \underset{\substack{\text{by defn} \\ \text{of } \bar{c}}}{\geq} \frac{\underline{\pi}}{2} \epsilon^{\bar{c}} =: \kappa_2 \epsilon^{\bar{c}}. \end{aligned} \tag{A.18}$$

¹⁹ This follows from the fact that the probability of each possible non-best response is of order ϵ^k for some strictly positive k , and there are a finite number of non-best responses, so k is uniformly bounded above 0. Thus the probability of a non-best response occurring approaches 0 as $\epsilon \rightarrow 0$.

Combining (A.17) and (A.18), starting from any $s^* \notin OS$, the probability of following the chosen sequence is, for small enough ϵ , greater than $\kappa_1 \kappa_2 \epsilon^{\bar{c}}$. Thus, from any $s^* \notin OS$, the probability of hitting OS within \bar{T} periods is greater than $\kappa_1 \kappa_2 \epsilon^{\bar{c}}$.

It follows that, from any initial state $s^* \notin OS$, the expected hitting time of OS is less than \bar{T} multiplied by the expected number of trials until the first success of a geometric distribution with success probability $\kappa_1 \kappa_2 \epsilon^{\bar{c}}$. That is the hitting time is less than $\bar{T} \frac{1}{\kappa_1 \kappa_2} \epsilon^{-\bar{c}}$.

Setting $\kappa := \bar{T} \frac{1}{\kappa_1 \kappa_2}$, we have our result.

Proof of Theorem 3.

By Lemma 1, with probability greater than $1 - \delta$, $H(G)$ is generic and connected. Henceforth, assume this to be the case.

Consider $P^{NE}(s^*, s')$ and $c^{NE}(s^*, s')$, $s^*, s' \in NE$, $s^* \neq s'$. Recall that $P^{NE}(s^*, s')$ is the probability under P that, from profile s^* at $t = 0$, the first strategy profile reached in the set NE at time $t > 0$ is s' .

Step 1: Sequences defining $P^{NE}(s^*, s')$

If $P^{NE}(s^*, s') \neq 0$, then there exists a sequence $s^*, s^1, \dots, s^T = s'$ such that transitions under P between consecutive states in the sequence happen with positive probability. Indeed, genericity and connectedness implies that such a sequence must exist, as following any deviation from s^* , a sequence of unique best responses leading to s' exists.

Step 2: Upper bound on $P^{NE}(s^*, s')$

For any given $s^1, s_i^* \neq s_i^1, s_{-i}^* = s_{-i}^1$, we have

$$P^\epsilon(s^*, s^1) \underset{\text{as } s_i^* \neq s_i^1}{=} \pi(i) P_i^\epsilon(s^*, s^1) \underset{\text{for } a, k > 0 \text{ by (4)}}{=} \pi(i) (a + o(1)) \epsilon^k \tag{A.19}$$

and therefore, for small enough ϵ ,

$$\frac{\pi}{2} \epsilon^k < P^\epsilon(s^*, s^1) < 2a \epsilon^k \leq 2a \epsilon^{c(s^*)}, \tag{A.20}$$

where the final inequality follows from $k \geq c(s^*)$ by definition of $c(\cdot)$.

As there a finite set of such s^1 , the probability of any of them occurring can thus be bounded above by $A \epsilon^{c(s^*)}$ for some $A > 0$.

Step 3: Lower bound on $P^{NE}(s^*, s')$

The probability of following a sequence of best responses $s^1, \dots, s^T = s'$ is bounded below by $(\frac{1}{2}\pi)^{T-1}$ (as in the proof of Theorem 2). Therefore, choosing s' such that $c(s^*, s') = c(s^*)$, the probability of following the sequence $s^*, s^1, \dots, s^T = s'$ is bounded below by $(\frac{1}{2}\pi)^{T-1} \frac{\pi}{2} \epsilon^{c(s^*)}$.

Step 4: NE is a cycle

Given the bounds on $P^{NE}(s^*, s')$, it follows from the definition of cost function that $c^{NE}(s^*, s') = c(s^*)$. Importantly, this is independent of s' . Consequently, $\sigma(s^*) = NE \setminus s^*$. This implies that $\Gamma = NE$ is a closed cycle of P^{NE} .

Step 5: Cyclic decomposition implying result

Cycles can be consolidated to give a *Freidlin-Wentzell decomposition* of P^{NE} . In the case under consideration, we have a single closed cycle $\Gamma = NE$ that can be consolidated into a single element. Thus we have a decomposition consisting of only two partitions of NE , $\mathcal{P}_0^{FW} = \bigcup_{s \in NE} \{\{s\}\}$, $\mathcal{P}_1^{FW} = \{NE\}$.

Given this decomposition, it follows immediately from Theorem 3.5 of Newton and Sandholm (2022) that $r(s) = \max_{s^* \in NE} c(s^*) - c(s)$.

Proof of Proposition 5.

Consider the proof of Theorem 1. The part of the proof showing the OS deviation property is Lemma 2.

Fix behavioral rules $\{P_i^\epsilon\}_\epsilon$. Observe that (31) implies that OS is constant with respect to π . Consider the proof of Lemma 2, but under c^π rather than c .

Step 1 of the proof remains the same.

Consider Step 2. In Case 1, $s^1 \in \operatorname{argmin}_{s \in S \setminus s^*} c(s^*, s)$. From (29) and (31), it follows that we also have $s^1 \in \operatorname{argmin}_{s \in S \setminus s^*} c^\pi(s^*, s)$. The remainder of Case 1, as well as Case 2, remain the same.

Consider Step 3. (31) implies that $c^\pi(s^*) = c(s^*)$ and $c^\pi(s') = c(s')$. Therefore, Step 3 remains the same.

Therefore, $\mathcal{R}S^\pi = OS$. As this holds for all π and OS does not depend on π , by definition of Robust- SS we have Robust- $SS = OS$.

Proof of Proposition 6.

Step 1: Cost function for multinomial probit

From Dokumaci and Sandholm (2011, Proposition 2.5), for $s', s'' = (s'_i, s'_{-i})$,

$$c(s', s'') = \sum_{\tilde{s}_i \in S_i} \frac{(z_{\tilde{s}_i}^*)^2}{2\sigma^2}, \tag{A.21}$$

where

$$z_{s_i}^* := \begin{cases} \left(\frac{1}{|\hat{S}_i \cup \{s_i''\}|} \sum_{\bar{s}_i \in \hat{S}_i \cup \{s_i''\}} u_i(\bar{s}_i, s_{-i}^l) \right) - u_i(s_i, s_{-i}^l), & \text{if } s_i \in \hat{S}_i \cup \{s_i''\}, \\ 0, & \text{otherwise,} \end{cases} \tag{A.22}$$

=average payoff from strategies in $\hat{S}_i \cup \{s_i''\}$

with the set $\hat{S}_i \subset S \setminus \{s_i''\}$ uniquely determined by the requirement that

$$s_i \in \hat{S}_i \iff u_i(s_i, s_{-i}^l) > \frac{1}{|\hat{S}_i \cup \{s_i''\}|} \sum_{\bar{s}_i \in \hat{S}_i \cup \{s_i''\}} u_i(\bar{s}_i, s_{-i}^l). \tag{A.23}$$

That is, \hat{S}_i is the maximal set of strategies that give higher payoffs than the average payoff from strategies in $\hat{S}_i \cup \{s_i''\}$. \hat{S}_i can be constructed by adding strategies in descending order of payoffs $u_i(\cdot, s_{-i}^l)$, stopping when an addition will no longer increase the average payoff from strategies in $\hat{S}_i \cup \{s_i''\}$.

Step 2a: Best responses have zero cost

Consider $s_i'' \in B_i(s^l)$, so that $u_i(s_i'', s_{-i}^l) \geq u_i(s_i, s_{-i}^l)$ for all $s_i \in S_i$. Together with (A.23), this implies $\hat{S}_i = \emptyset$. Together with (A.22) and (A.21), this implies $z_{\bar{s}_i}^* = 0$ for all \bar{s}_i and hence $c(s^l, s'') = 0$.

Step 2b: Second best responses have payoff-difference based cost

Now consider $B_i(s^l) = \{s^0\}$ and $u_i(s_i'', s_{-i}^l) \geq u_i(s_i, s_{-i}^l)$ for all $s_i \in S_i \setminus s^0$. Together with (A.23), this implies $\hat{S}_i = \{s^0\}$. Together with (A.22), this implies

$$z_{s_i^0}^* = \left(\frac{u_i(s_i^0, s_{-i}^l) + u_i(s_i'', s_{-i}^l)}{2} \right) - u_i(s_i^0, s_{-i}^l) = \frac{u_i(s_i'', s_{-i}^l) - u_i(s_i^0, s_{-i}^l)}{2}, \tag{A.24}$$

$$z_{s_i''}^* = \left(\frac{u_i(s_i^0, s_{-i}^l) + u_i(s_i'', s_{-i}^l)}{2} \right) - u_i(s_i'', s_{-i}^l) = \frac{u_i(s_i^0, s_{-i}^l) - u_i(s_i'', s_{-i}^l)}{2}.$$

Substituting into (A.21),

$$c(s^l, s'') = 2 \frac{(u_i(s_i^0, s_{-i}^l)/2 - u_i(s_i'', s_{-i}^l)/2)^2}{2\sigma^2} = \frac{(u_i(s_i^0, s_{-i}^l) - u_i(s_i'', s_{-i}^l))^2}{4\sigma^2}. \tag{A.25}$$

Note that the right hand side of (A.25) is a strictly increasing function $f : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$, with $f(x) = x^2/4\sigma^2$ taking argument

$$\underbrace{u_i(s_i^0, s_{-i}^l)}_{=\max_{\bar{s}_i \in S_i} u_i(\bar{s}_i, s_{-i}^l) \text{ by } s_i^0 \in B_i(s^l)} - u_i(s_i'', s_{-i}^l) = \max_{\bar{s}_i \in S_i} u_i(\bar{s}_i, s_{-i}^l) - u_i(s_i'', s_{-i}^l). \tag{A.26}$$

That is, second best responses have a cost that is a strictly increasing function of payoff differences.

Step 2c: Monotonicity of costs in payoffs

Additionally, it can be checked that for any $s_i^1, s_i^2 \in S_i$ such that $u_i(s_i^1, s_{-i}^l) \geq u_i(s_i^2, s_{-i}^l)$, we have $c(s^l, (s_i^1, s_{-i}^l)) \leq c(s^l, (s_i^2, s_{-i}^l))$.

Step 3: OS = least-NE

Consider $s^* \in NE$. By definition of NE, $B_i(s^*)$ is a singleton, say s^0 . By Step 2c, the least cost deviation by player i is to some s_i'' such that $u_i(s_i'', s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$ for all $s_i \in S_i \setminus s^0$. By Step 2b, the cost of this deviation is given by (A.25). Using (A.26) and taking the minimum over all $i \in N$,

$$c(s^*) = \min_{\substack{i \in N \\ s_i \neq s_i^*}} f \left(\underbrace{\max_{\bar{s}_i \in S_i} u_i(\bar{s}_i, s_{-i}^*)}_{=u_i(s^*) \text{ as } s^* \in NE} - u_i(s_i, s_{-i}^*) \right) \tag{A.27}$$

$$= \min_{\substack{i \in N \\ s_i \neq s_i^*}} f \left(u_i(s^*) - u_i(s_i, s_{-i}^*) \right).$$

Substituting into (A.5),

$$OS = \arg \max_{s^* \in NE} \min_{\substack{i \in N \\ s_i \neq s_i^*}} f \left(u_i(s^*) - u_i(s_i, s_{-i}^*) \right) \tag{A.28}$$

$$\begin{aligned}
 &= \underset{\substack{\text{as } f \\ \text{strictly} \\ \text{increasing}}}{\text{arg max}} \underset{\substack{s^* \in NE \\ i \in N \\ s_i \neq s_i^*}}{\text{min}} [u_i(s^*) - u_i(s_i, s_{-i}^*)] \\
 &= \text{least-}NE. \\
 &\text{by Prop.3}
 \end{aligned}$$

Proof of Proposition 7.

$$\begin{aligned}
 \mathbb{E}_{\Gamma(N,S)}[|NE_{\bar{u}}|] &= \sum_{s \in S} Prob_{\Gamma(N,S)}[s \in NE_{\bar{u}}] \\
 &= \sum_{s \in S} \prod_{i \in N} \underbrace{Prob_{\Gamma(N,S)}[\{s_i\} = B_{i,\bar{u}}(s)]}_{=: \beta < \frac{1}{m}} = |S| \beta^{|N|} = m^{|N|} \beta^{|N|} \xrightarrow{|N| \rightarrow \infty} 0,
 \end{aligned} \tag{A.29}$$

which together with

$$\mathbb{E}_{\Gamma(N,S)}[|NE_{\bar{u}}|] = \sum_{k \in \mathbb{Z}_+} k Prob_{\Gamma(N,S)}[|NE_{\bar{u}}| = k], \tag{A.30}$$

implies

$$Prob_{\Gamma(N,S)}[NE_{\bar{u}} = \emptyset] = Prob_{\Gamma(N,S)}[|NE_{\bar{u}}| = 0] \xrightarrow{N \rightarrow \infty} 1. \tag{A.31}$$

Proof of Theorem 4.

Step 1: Poisson approximation and lower bound on $Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset]$

Fix $\Gamma, \mathcal{P}, \delta, m \geq 2$. Let $\beta = Prob_{\Gamma(N,S)}[\{s_i\} = B_{i,\bar{u}}(s)]$ for arbitrary $i \in N$. Due to the existence of atoms, we have $\beta < \frac{1}{m}$. As in the proof of Proposition 7, recalling that $n = |N|$,

$$Prob_{\Gamma(N,S)}[s \in NE_{\bar{u}}] = Prob_{\Gamma(N,S)}[\mathbb{1}_{NE_{\bar{u}}}(s) = 1] = \beta^n, \tag{A.32}$$

$$\mathbb{E}_{\Gamma(N,S)}[|NE_{\bar{u}}|] = \mathbb{E}_{\Gamma(N,S)}\left[\sum_{s \in S} \mathbb{1}_{NE_{\bar{u}}}(s)\right] = \sum_{s \in S} \beta^n = m^n \beta^n. \tag{A.33}$$

Note that

- $\mathbb{1}_{NE_{\bar{u}}}(s)$ is not independent of $\mathbb{1}_{NE_{\bar{u}}}(s')$ for $s' = (s'_i, s_{-i})$,
- $\mathbb{1}_{NE_{\bar{u}}}(s)$ is independent of $\mathbb{1}_{NE_{\bar{u}}}(s'')$ for $s''_i \neq s_i, s''_j \neq s_j, i \neq j$.

This allows us to use Theorem 1 of Arratia et al. (1989) to derive the following lemma that approximates $|NE_{\bar{u}}|$ by a Poisson distribution with mean $\mathbb{E}_{\Gamma(N,S)}[|NE_{\bar{u}}|]$.²⁰

Lemma 3. Let Y be a Poisson random variable with mean $\mathbb{E}_{\Gamma(N,S)}[|NE_{\bar{u}}|]$. Then,

$$\sup_{Z \subseteq \mathbb{Z}_+} \left| Prob_{\Gamma(N,S)}[|NE_{\bar{u}}| \in Z] - Prob[Y \in Z] \right| \leq m^n(n(m-1) + 1)\beta^{2n}.$$

Lemma 3 implies that

$$\begin{aligned}
 Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset] &\geq Prob[Y > 0] - m^n(n(m-1) + 1)\beta^{2n} \\
 &= 1 - e^{-\beta^n m^n} - m^n(n(m-1) + 1)\beta^{2n}.
 \end{aligned} \tag{A.34}$$

Step 2: Lower bound on $Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u \cup NE_u = \emptyset]$

Consider the event that either $H(G)$ is connected_u or $NE_u = \emptyset$. We have

$$\begin{aligned}
 Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u \cup NE_u = \emptyset] &= \\
 &Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_u \neq \emptyset] Prob_{\Gamma(N,S)}[NE_u \neq \emptyset] \\
 &+ Prob_{\Gamma(N,S)}[NE_u = \emptyset] \\
 &\geq Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_u \neq \emptyset] \underset{\substack{\text{for large enough } n, \\ \text{for some } c > 0}}{\geq} 1 - e^{-cn},
 \end{aligned} \tag{A.35}$$

where the last inequality follows from Johnston et al. (2025, Theorem 1). for the remainder of the proof, assume n is large enough that this holds.

²⁰ Rinott and Scarsini (2000) apply Arratia et al. (1989) to games without atoms. With reference to Theorem 3.1 of Rinott and Scarsini (2000), the difference is that the bound in our lemma is $m^n(n(m-1) + 1)\beta^{2n}$, whereas without atoms β is replaced by $1/m$ to give $m^n(n(m-1) + 1)(1/m)^{2n}$.

Step 3: Lower bound on $Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_{\bar{u}} \neq \emptyset]$

Note that $NE_u = \emptyset$ implies $NE_{\bar{u}} = \emptyset$. Therefore,

$$\begin{aligned} Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u \cup NE_u = \emptyset | NE_{\bar{u}} \neq \emptyset] \\ = Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_{\bar{u}} \neq \emptyset]. \end{aligned} \tag{A.36}$$

Let A denote the event “ $H(G)$ is connected $_u \cup NE_u = \emptyset$ ”. Then,

$$\begin{aligned} Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_{\bar{u}} \neq \emptyset] \\ \stackrel{\text{by (71)}}{=} Prob_{\Gamma(N,S)}[A | NE_{\bar{u}} \neq \emptyset] = \frac{Prob_{\Gamma(N,S)}[A \cap NE_{\bar{u}} \neq \emptyset]}{Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset]} \\ \geq \frac{Prob_{\Gamma(N,S)}[A] - (1 - Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset])}{Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset]} = 1 - \frac{1 - Prob_{\Gamma(N,S)}[A]}{Prob_{\Gamma(N,S)}[NE_{\bar{u}} \neq \emptyset]} \\ \stackrel{\text{by (70),(69)}}{\geq} 1 - \frac{e^{-cn}}{1 - e^{-\beta^n m^n} - m^n(n(m-1) + 1)\beta^{2n}} \end{aligned} \tag{A.37}$$

Step 4: For small \bar{u} , $Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_{\bar{u}} \neq \emptyset] \rightarrow 1$ as $n \rightarrow \infty$

We shall show that, for small enough \bar{u} , the final expression of (A.37) approaches one as $n \rightarrow \infty$. Proceed by dividing each term in the numerator and denominator of the fraction in the expression by $\beta^n m^n$ and considering the limit $n \rightarrow \infty$.

- (i) The Maclaurin series of $1 - e^{-\beta^n m^n}$ implies that $\lim_{n \rightarrow \infty} \frac{1 - e^{-\beta^n m^n}}{\beta^n m^n} = 1$.
- (ii) In contrast, $\lim_{n \rightarrow \infty} \frac{m^n(n(m-1)+1)\beta^{2n}}{\beta^n m^n} = \lim_{n \rightarrow \infty} (n(m-1) + 1)\beta^n = 0$.
- (iii) Finally, consider $\lim_{n \rightarrow \infty} \frac{e^{-cn}}{\beta^n m^n} = \lim_{n \rightarrow \infty} e^{-n(c + \log \beta m)}$. This equals zero if $c + \log \beta m > 0$.

Considering the definition of β (for arbitrary $i \in N$) together with the fact that payoff distribution F is continuous,

$$\beta = Prob_{\Gamma(N,S)}[\{s_i\} = B_{i,\bar{u}}(s)] \xrightarrow{\bar{u} \rightarrow 0} \frac{1}{m}, \tag{A.38}$$

we have that $\beta m \rightarrow 1$ and $\log \beta m \rightarrow 0$. Therefore, for \bar{u} small enough that $c + \log \beta m > 0$, the inequality in (iii) holds and $\lim_{n \rightarrow \infty} \frac{e^{-cn}}{\beta^n m^n} \rightarrow 0$.

Considering (i), (ii), (iii), (A.37) together, for \bar{u} small enough that $c + \log \beta m > 0$,

$$Prob_{\Gamma(N,S)}[H(G) \text{ is connected}_u | NE_{\bar{u}} \neq \emptyset] \xrightarrow{n \rightarrow \infty} 1 \tag{A.39}$$

Step 5: If $H(G)$ is connected $_u$ and $NE_{\bar{u}} \neq \emptyset$, then one-shot deviation holds

Assume $H(G)$ is connected $_u$ and $NE_{\bar{u}} \neq \emptyset$. Consider $s^* \in NE_{\bar{u}}$.

- (i) For $s' \notin NE_u$, as $H(G)$ is connected under u , there is a path of best responses s', \dots, s^* under u . Remark 9 implies that this is also a path of best responses under \bar{u} .
- (ii) Consider $s^{**} \in NE_u \setminus NE_{\bar{u}}$, $s^{**} \notin NE_{\bar{u}}$, so there exist $i \in N$, $s'_i \in S_i$, $s'_i \neq s^{**}_i$ such that $s'_i \in B_i(s^{**})$. $s^{**} \in NE_u$ implies that $s' = (s'_i, s^{**}_{-i}) \notin NE_{\bar{u}}$. Combining with (i), we have a path of best responses s^{**}, s', \dots, s^* under \bar{u} .
- (iii) Consider $s^{**} \in NE_{\bar{u}}$. By definition of \bar{u} , this implies $s^{**} \in NE_u$. Let $s' = (s'_i, s^{**}_{-i})$ be such that $c_{\bar{u}}(s^{**}, s') = c_{\bar{u}}(s^{**})$. $s^{**} \in NE_u$ implies that $s' \notin NE_{\bar{u}}$. Combining with (i), a transition to s^* can be completed via the path of best responses s', \dots, s^* under \bar{u} .

The above points imply that we can apply the arguments of Lemma 2, only for $NE_{\bar{u}}$ rather than NE_u . For all $s^* \in NE_{\bar{u}}$, stochastic potential is $\sum_{s \in NE_{\bar{u}} \setminus s^*} c_{\bar{u}}(s)$, therefore $SS_{\bar{u}}$ comprises the maximizers of $c_{\bar{u}}(s^*)$. That is, $SS_{\bar{u}} = OS_{\bar{u}}$.

References

Agastya, M., 1997. Adaptive play in multiplayer bargaining situations. *Rev. Econ. Stud.* 64 (3), 411–26. <http://ideas.repec.org/a/bla/restud/v64y1997i3p411-26.html>.
 Agastya, M., 1999. Perturbed adaptive dynamics in coalition form games. *J. Econ. Theory* 89 (2), 207–233. <http://ideas.repec.org/a/eee/jetheo/v89y1999i2p207-233.html>.
 Alós-Ferrer, C., Netzer, N., 2010. The logit-response dynamics. *Games Econ. Behav.* 68, 413–427.
 Alós-Ferrer, C., Netzer, N., 2015. Robust stochastic stability. *Econ. Theory* 58 (1), 31–57.
 Amiet, B., Collecchio, A., Scarsini, M., Zhong, Z., 2021. Pure Nash equilibria and best-response dynamics in random games. *Math. Oper. Res.* 46 (4), 1552–1572.
 Arieli, I., Babichenko, Y., 2016. Random extensive form games. *J. Econ. Theory* 166, 517–535. <https://doi.org/10.1016/j.jet.2016.09.010>
 Arieli, I., Babichenko, Y., Peretz, R., Young, H.P., 2020. The speed of innovation diffusion in social networks. *Econometrica* 88 (2), 569–594.
 Arieli, I., Young, H.P., 2016. Stochastic learning dynamics and speed of convergence in population games. *Econometrica* 84 (2), 627–676.
 Arratia, R., Goldstein, L., Gordon, L., 1989. Two moments suffice for poisson approximations: the chen-stein method. *Annal. Probab.* 17 (1), 9–25. <http://www.jstor.org/stable/2244193>.
 Ben-Shoham, A., Serrano, R., Volij, O., 2004. The evolution of exchange. *J. Econ. Theory* 114 (2), 310–328.
 Bilancini, E., Boncinelli, L., 2020. The evolution of conventions under condition-dependent mistakes. *Econ. Theory* 69 (2), 497–521.
 Bilancini, E., Boncinelli, L., Nax, H.H., 2021. What noise matters? experimental evidence for stochastic deviations in social norms. *J. Behav. Exp. Econ.* 90, 101626.

- Bilancini, E., Boncinelli, L., Newton, J., 2020. Evolution and Rawlsian social choice in matching. *Games Econ. Behav.* 123, 68–80.
- Blume, L.E., 1993. The statistical mechanics of strategic interaction. *Games Econ. Behav.* 5 (3), 387–424. <https://doi.org/10.1006/game.1993.1023>
- Blume, L.E., 2003. How noise matters. *Games Econ. Behav.* 44 (2), 251–271.
- Bollobás, B., 2001. *Random Graphs*. Cambridge Studies in Advanced Mathematics, Cambridge University Press, Cambridge, UK. 2 edition. <https://doi.org/10.1017/CBO9780511814068>
- Cui, Z., Zhai, J., 2010. Escape dynamics and equilibria selection by iterative cycle decomposition. *J. Math. Econ.* 46 (6), 1015–1029. <https://doi.org/10.1016/j.jmateco.2009.11.014>
- Diamantoudi, E., Xue, L., Miyagawa, E., 2004. Random paths to stability in the roommate problem. *Games Econ. Behav.* 48 (1), 18–28.
- Dokumaci, E., Sandholm, W.H., 2011. Large deviations and multinomial probit choice. *J. Econ. Theory* 146 (5), 2151–2158. <http://ideas.repec.org/a/eee/jetheo/v146y2011i5p2151-2158.HTML>.
- Dresher, M., 1970. Probability of a pure equilibrium point in n-person games. *J. Combinator. Theory* 8 (1), 134–145.
- Ellison, G., 1993. Learning, local interaction, and coordination. *Econometrica* 61 (5), 1047–71. <http://ideas.repec.org/a/econ/emetrp/v61y1993i5p1047-71.HTML>.
- Ellison, G., 2000. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Rev. Econ. Stud.* 67 (1), 17–45. <http://ideas.repec.org/a/bla/restud/v67y2000i1p17-45.HTML>.
- Feldman, A.M., 1974. Reconstructing stability. *Econometrica* 42 (1), 35–44. <http://www.jstor.org/stable/1913683>.
- Foster, D., Young, H.P., 1990. Stochastic evolutionary game dynamics. *Theor. Popul. Biol.* 38, 219–232.
- Freidlin, M.I., Wentzell, A.D., 1984. *Random Perturbations of Dynamical Systems*, 2nd ed.(1998). Springer, Berlin, Germany.
- Green, J.R., 1974. The stability of Edgeworth's reconstructing process. *Econometrica* 42 (1), 21–34. <http://www.jstor.org/stable/1913682>.
- Heinrich, T., Jang, Y., Mungo, L., Pangallo, M., Scott, A., Tarbush, B., Wiese, S., 2023. Best-response dynamics, playing sequences, and convergence to equilibrium in random games. *Int. J. Game Theory* 52 (3), 703–735.
- Hwang, S.-H., Lim, W., Neary, P., Newton, J., 2018. Conventional contracts, intentional behavior and logit choice: equality without symmetry. *Games Econ. Behav.* 110, 273–294. <https://doi.org/10.1016/j.geb.2018.05.002>
- Jackson, M.O., Watts, A., 2002a. The evolution of social and economic networks. *J. Econ. Theory* 106 (2), 265–295.
- Jackson, M.O., Watts, A., 2002b. On the formation of interaction networks in social coordination games. *Games Econ. Behav.* 41 (2), 265–291.
- Johnston, T., Savery, M., Scott, A., Tarbush, B., 2025. Game connectivity and adaptive dynamics. [arXiv:2309.10609](https://arxiv.org/abs/2309.10609).
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61 (1), 29–56. <http://ideas.repec.org/a/econ/emetrp/v61y1993i1p29-56.HTML>.
- Kandori, M., Serrano, R., Volij, O., 2008. Decentralized trade, random utility and the evolution of social welfare. *J. Econ. Theory* 140 (1), 328–338.
- Khan, A., 2022. Expected utility versus cumulative prospect theory in an evolutionary model of bargaining. *J. Econ. Dyn. Control* 137, 104332. <https://doi.org/10.1016/j.jedc.2022.104332>
- Klaus, B., Newton, J., 2016. Stochastic stability in assignment problems. *J. Math. Econ.* 62, 62–74. <https://doi.org/10.1016/j.jmateco.2015.11.002>
- Kojima, F., Ünver, M.U., 2008. Random paths to pairwise stability in many-to-many matching problems: a study on market equilibration. *Int. J. Game Theory* 36 (3–4), 473–488. <https://doi.org/10.1007/s00182-006-0037-2>
- Kreindler, G.E., Young, H.P., 2013. Fast convergence in evolutionary equilibrium selection. *Games Econ. Behav.* 80 (0), 39–67. <https://doi.org/10.1016/j.geb.2013.02.004>
- Kreindler, G.E., Young, H.P., 2014. Rapid innovation diffusion in social networks. *Proc. Natl. Acad. Sci.* 111 (Supplement 3), 10881–10888.
- Levine, D.K., Modica, S., 2016. Dynamics in stochastic evolutionary models. *Theor. Econ.* 11 (1), 89–131.
- Lewis, D., 1969. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA, USA.
- Lim, W., Neary, P.R., 2016. An experimental investigation of stochastic adjustment dynamics. *Games Econ. Behav.* 100, 208–219.
- Mäs, M., Nax, H.H., 2016. A behavioral study of “noise” in coordination games. *J. Econ. Theory* 162, 195–208.
- Maschler, M., Peleg, B., Shapley, L.S., 1979. Geometric properties of the kernel, nucleolus, and related solution concepts. *Math. Oper. Res.* 4 (4), 303–338. <http://www.jstor.org/stable/3689220>.
- McLennan, A., Berg, J., 2005. Asymptotic expected number of Nash equilibria of two-player normal form games. *Games Econ. Behav.* 51 (2), 264–295. Special Issue in Honor of Richard D. McKelvey. <https://doi.org/10.1016/j.geb.2004.10.008>
- Mimun, H.A., Quattropiani, M., Scarsini, M., 2024. Best-response dynamics in two-person random games with correlated payoffs. *Games Econ. Behav.* 145, 239–262. <https://doi.org/10.1016/j.geb.2024.03.011>
- Montero, M., Possajennikov, A., 2023. “Greedy” demand adjustment in cooperative games. *Ann. Oper. Res.* 10.1007/s10479-023-05179-8. <https://doi.org/10.1007/s10479-023-05179-8>
- Myatt, D.P., Wallace, C., 2003. A multinomial probit model of stochastic evolution. *J. Econ. Theory* 113 (2), 286–301. <http://ideas.repec.org/a/eee/jetheo/v113y2003i2p286-301.HTML>.
- Naidu, S., Hwang, S.-H., Bowles, S., 2010. Evolutionary bargaining with intentional idiosyncratic play. *Econ. Lett.* 109 (1), 31–33. <https://doi.org/10.1016/j.econlet.2010.07.005>
- Nax, H.H., Pradelski, B. S.R., 2015. Evolutionary dynamics and equitable core selection in assignment games. *Int. J. Game Theory* 44 (4), 903–932. <https://doi.org/10.1007/s00182-014-0459-1>
- Nax, H.H., Pradelski, B. S.R., Young, H.P., 2013. Decentralized dynamics to optimal and stable states in the assignment game. *Proc. 52nd IEEE Conf. Decis. Control*, 2391–2397.
- Newton, J., 2012. Coalitional stochastic stability. *Games Econ. Behav.* 75 (2), 842–84. <https://doi.org/10.1016/j.geb.2012.02.014>
- Newton, J., 2015. Stochastic stability on general state spaces. *J. Math. Econ.* 58, 46–60. <https://doi.org/10.1016/j.jmateco.2015.03.007>
- Newton, J., 2018. Evolutionary game theory: a renaissance. *Games* 9 (2), 31. <https://doi.org/10.3390/g9020031>
- Newton, J., 2021. Conventions under heterogeneous behavioural rules. *Rev. Econ. Stud.* 88 (4), 2094–2118.
- Newton, J., Angus, S.D., 2015. Coalitions, tipping points and the speed of evolution. *J. Econ. Theory* 157 (0), 172–187. <https://doi.org/10.1016/j.jet.2015.01.003>
- Newton, J., Sandholm, W.H., 2022. Stochastic dynamics and Edmonds' algorithm. *J. Dyn. Games* 9 (4), 529–546.
- Newton, J., Sawa, R., 2015. A one-shot deviation principle for stability in matching problems. *J. Econ. Theory* 157 (0), 1–27. <https://doi.org/10.1016/j.jet.2014.11.015>
- Pei, T., Takahashi, S., 2019. Rationalizable strategies in random games. *Games Econ. Behav.* 118, 110–125. <https://doi.org/10.1016/j.geb.2019.08.011>
- Peski, M., 2010. Generalized risk-dominance and asymmetric dynamics. *J. Econ. Theory* 145 (1), 216–248. <https://doi.org/10.1016/j.jet.2009.05.007>
- Pradelski, B. S.R., 2015. Decentralized dynamics and fast convergence in the assignment game. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM.
- Pradelski, B. S.R., Tarbush, B., 2025. On the existence of pure epsilon-equilibrium. Online available at <https://arxiv.org/abs/2502.07585>.
- Rinott, Y., Scarsini, M., 2000. On the number of pure strategy Nash equilibria in random games. *Games Econ. Behav.* 33 (2), 274–293. <https://doi.org/10.1006/game.1999.0775>
- Roth, A.E., Vande Vate, J.H., 1990. Random paths to stability in two-sided matching. *Econometrica* 58 (6), 1475–80. <http://ideas.repec.org/a/econ/emetrp/v58y1990i6p1475-80.HTML>.
- Sandholm, W.H., 2010. *Population games and evolutionary dynamics*. Economic learning and social evolution, Cambridge, Mass. MIT Press. <http://opac.inria.fr/record=b1132817>.
- Sawa, R., 2021. A prospect theory Nash bargaining solution and its stochastic stability. *J. Econ. Behav. Org.* 184, 692–711. <https://doi.org/10.1016/j.jebo.2020.11.009>
- Serrano, R., Volij, O., 2008. Mistakes in cooperation: the stochastic stability of Edgeworth's reconstructing. *Econ. J.* 118 (532), 1719–1741. <http://ideas.repec.org/a/ecj/econjl/v118y2008i532p1719-1741.HTML>.
- Spencer, J., 2013. *The strange logic of random graphs*. Vol. 22. Springer Science & Business Media.

- Stanford, W., 1996. The limit distribution of pure strategy Nash equilibria in symmetric bimatrix games. *Math. Oper. Res.* 21 (3), 726–733. <http://www.jstor.org/stable/3690306>.
- Takahashi, S., 2008. The number of pure Nash equilibria in a random game with nondecreasing best responses. *Games Econ. Behav.* 63 (1), 328–340. <https://doi.org/10.1016/j.geb.2007.10.003>
- Wiese, S.C., Heinrich, T., 2022. The frequency of convergent games under best-response dynamics. *Dyn. Games Appl.* 12 (2), 689–700. 10.1007/s13235-021-00401-3. <https://doi.org/10.1007/s13235-021-00401-3>
- Young, H.P., 1993a. The evolution of conventions. *Econometrica* 61 (1), 57–84. <http://ideas.repec.org/a/ecm/emetrp/v61y1993i1p57-84.HTML>.
- Young, H.P., 1993b. An evolutionary model of bargaining. *J. Econ. Theory* 59 (1), 145–168. <https://doi.org/10.1006/jeth.1993.1009>
- Young, H.P., 1998. Conventional contracts. *Rev. Econ. Stud.* 65 (4), 773–92. <http://ideas.repec.org/a/bla/restud/v65y1998i4p773-92.HTML>.