Econ Theory
DOI 10.1007/s00199-016-0988-x

RESEARCH ARTICLE

Payoff-dependent dynamics and coordination games

Sung-Ha Hwang[1] · **Jonathan Newton**[2]

Received: 12 January 2015 / Accepted: 15 June 2016

Abstract This paper considers populations of agents whose behavior when playing some underlying game is governed by perturbed best (or better) response dynamics with perturbation probabilities that depend log-linearly on payoffs, a class that includes the logit choice rule. A convention is a state at which every agent plays a strategy that corresponds to the same strict Nash equilibrium of the underlying game. For coordination games with zero payoff off-diagonal, it is shown that the difficulty of leaving the basin of attraction of a convention can be well approximated by only considering paths of transitions on which an identical perturbation repeatedly affects one of the populations.

Keywords Evolution · Coordination · Logit · Payoff dependence

JEL Classification C73 · C78

First Version: January 31, 2014, This version: June 27, 2016. S.-H.Hwang was supported by the National Research Foundation of Korea Grant funded by the Korean Government (NRF-2014S1A5A8019513). The work was completed while J. Newton was supported by a Discovery Early Career Researcher Award funded by the Australian Research Council (Grant Number: DE130101768).

✉ Sung-Ha Hwang
sunghah@sogang.ac.kr

Jonathan Newton
jonathan.newton@sydney.edu.au

1 Sogang University, Seoul, Korea

2 University of Sydney, Sydney, Australia

Published online: 04 July 2016

1 Introduction

Consider two populations corresponding to positions in a two-player normal form game. The behavior of agents within the populations is governed by a perturbed individualistic best response dynamic (Young 1993; Kandori et al. 1993). Such dynamics involve *individuals* who will usually play a *best response* to the distribution of play of the opposing population, but whose behavior is *perturbed* in that they occasionally play something other than a best response. One popular class of perturbations involves non-best response strategies being chosen with probabilities that decrease log-linearly in the payoff lost by choosing them rather than a best response (Blume 1993; Sandholm 2010). The logit choice rule falls into this class. For potential games (Monderer and Shapley 1996), when agents update their strategies one at a time and perturbations are small, such dynamics spend most of their time close to strategy profiles which globally maximize the potential function.[1] However, the dependency of perturbations on payoffs that makes such dynamics amenable to the analysis of potential games creates obstacles to the analysis of games which are not potential games.[2]

In this paper, we analyze coordination games with zero payoff off-diagonal, which may not be potential games. A *convention* is a state at which every agent plays a strategy corresponding to the same strict Nash equilibrium of the game. In the absence of perturbations, conventions would be rest points of the dynamic process. An important quantity for the study of adaptive dynamics under small perturbations is the *cost* of transiting from a convention to outside of its basin of attraction. The cost corresponds to the exponential rate of decay of the probability of a transition as perturbations become rare. When perturbations are rare, so that players almost always choose a best response, high-cost transitions occur much less often than low-cost transitions. For the class of games considered in this paper, we show that such costs can be well approximated even if we restrict attention to paths on which the only non-best response behavior involves members of a single population playing the same action. Our proof technique is elementary and relies on the construction of an explicit lower bound function for the cost of escaping a basin of attraction. We use our result to characterize stochastically stable states in coordination games for which the best coordination outcome for a player corresponds to the worst outcome for his opponent.

The results of the current paper hold for coordination games with an arbitrary number of strategies. Potential games are non-generic within this class. Furthermore, most such games do not have a "local potential maximizer" (Morris and Ui 2005) and so the results of Okada and Tercieux (2012) are of limited applicability. Complementary results by other authors hold for either two-strategy, two-population models (Staudigl 2012) or for three-strategy, one-population models (Sandholm and Staudigl 2016).

[1] This does not hold when more than a single agent can update their strategy at the same time (Alós-Ferrer and Netzer 2010). Furthermore, if the dynamic allows for agency at a level greater than the individual (Newton 2012), then potential maximizing strategy profiles are not necessarily even rest points of the unperturbed dynamic (Newton and Angus 2015, 2013).

[2] Similarly, the analysis of potential games is much easier with payoff-dependent perturbations than with payoff-independent perturbations. To see this, compare the difficulty of proving stochastic stability of risk-dominant actions in symmetric two-by-two coordination games on networks for payoff-dependent perturbations (Young 1998b) and for payoff-independent perturbations (Peski 2010).

The cited works show that as population size approaches infinity, transition costs can be found by solving a continuous optimal control problem. In contrast, the current paper uses an explicit construction. This has three benefits. Firstly, our result is as easily proven for n strategies as it is for 3 strategies. Secondly, our method allows us to obtain bounds for fixed, finite population size, which is useful for empirical work. Thirdly, our proof is short, elementary, and easily understood without specialist knowledge.

Finally, we note that the result of the current paper applies to a broad class of exponential revision rules (Sandholm 2010) and not just logit. In particular, the Theorem applies to behavioral rules in which agents do not play best responses, but rather better responses, whereby an agent compares his current action to a single alternative strategy at a time rather than seeking the optimal strategy from the entire set of strategies (Friedman and Mezzetti 2001; Josephson 2008; Dindoš and Mezzetti 2006).

Possible applications of our results arise whenever there are multiple ways in which two parties can coordinate to generate surplus. One such example is the literature on intra-household bargaining (e.g., Manser and Brown 1980; McElroy and Horney 1981; Lundberg and Pollak 1993) in which it is assumed that the distribution of gains from marriage corresponds to some solution from cooperative game theory (e.g., Nash 1950; Kalai and Smorodinsky 1975; Kalai 1977). Similarly, in the literature on search and matching, the division of surplus among matched players is often assumed to satisfy some cooperative solution.[3] See Rogerson et al. (2005) and Lagos et al. (2016) for surveys of this work as it relates to the labor market and monetary economics, respectively. Our work raises the prospect of endogenizing surplus allocation in these models, so that sharing norms emerge from behavioral rules applied to the context under consideration. Steps in this direction have already been taken in Hwang and Newton (2014) and Hwang et al. (2016) who consider bargaining sets characterized by decreasing, concave efficient frontiers. Using the Theorem of the current paper, these works link behavioral rules to bargaining solutions, using experimental data to identify the Egalitarian solution (Kalai 1977) as the most plausible long-run behavioral norm in this context.

The paper is organized as follows. Section 2 gives the model. Section 3 presents and discusses the approximation theorem which is the main result of the paper. Section 4 uses the Theorem to prove a result on stochastic stability in coordination games. Section 5 contains proofs.

2 Model

Consider two populations of agents — α- and β-populations — of size N. Two agents, one from each population, are matched to play a coordination game. The common strategy set is $S := \{0, 1, 2, \ldots, n\}$. A strategy profile or *state* is described by $x := (x_\alpha, x_\beta)$, where x_α and x_β are vectors giving the number of agents in each population who are using each strategy. Thus, the state space Ξ is

[3] We have in mind what is known as the Diamond–Mortensen–Pissarides framework of search and matching. See, for example, Diamond (1982); Mortensen and Pissarides (1994, 1999).

$$\Xi := \left\{ (x_\alpha, x_\beta) : \sum_{l \in S} x_\alpha(l) = N,\ x_\alpha(l) \in \mathbb{N}_0,\ \sum_{l \in S} x_\beta(l) = N, x_\beta(l) \in \mathbb{N}_0 \right\}$$

More explicitly, we have $(x_\alpha, x_\beta) = ((x_\alpha(0), x_\alpha(1), \ldots, x_\alpha(n)), (x_\beta(0), x_\beta(1), \ldots, x_\beta(n)))$, where $x_\beta(2)$, for example, denotes the number of β-agents playing strategy 2.

We consider coordination games with payoffs given by

$$(\pi_\alpha(i, j),\ \pi_\beta(j, i)) = \begin{cases} (\pi_\alpha(i), \pi_\beta(i)), \pi_\alpha(i), \pi_\beta(i) \geq 0 & \text{if } i = j \\ (0, 0) & \text{otherwise} \end{cases}.$$

We assume that agents from each population are randomly matched to play the game, and thus, the expected payoff of an α-agent who plays strategy i is $\pi_\alpha(i, x) := \sum_{l \in S} \pi_\alpha(i, l)\, x_\beta(l)/N$, given that the fraction of the β-population using strategy l is $x_\beta(l)/N$. Similarly, the expected payoff of a β-agent who plays strategy i is $\pi_\beta(i, x) := \sum_{l \in S} \pi_\beta(i, l)\, x_\alpha(l)/N$.

We consider a discrete time strategy updating process defined as follows. At each period, an agent from either the α-population or the β-population is uniformly chosen at random. The chosen agent selects a strategy based on his evaluation of the expected payoffs of the different strategies. The agent may idiosyncratically experiment with non-optimal strategies, or simply make mistakes. The probability of such mistakes will be parameterized by a parameter η, and larger values of η will correspond to higher mistake probabilities.

Specifically, we consider the following revision rules from the class of exponential revision protocols (Sandholm 2010). From state x, when an agent from population $\gamma \in \{\alpha, \beta\}$ who is currently playing strategy i is chosen to update his strategy, he switches to strategy $j \neq i$ with probability

$$p^\eta_\gamma(j|i, x) \propto \frac{\exp(\eta^{-1}\pi_\gamma(j, x))}{\sum_{l \in C_{i,j}} \exp(\eta^{-1}\pi_\gamma(l, x))}, \tag{1}$$

where $C_{i,j} \subseteq S$ and $i, j \in C_{i,j}$. This set $C_{i,j}$ can be understood as the set of alternatives with which the current strategy i and prospective strategy j are compared. Examples of rules satisfying (1) are the logit choice rule, for which $C_{i,j} = S$, and the exponential better reply rule, for which $C_{i,j} = \{i, j\}$.

Note that as η approaches zero, the probability of choosing a strategy from some comparison group $C_{i,j}$ that is inferior to some other strategy in $C_{i,j}$ approaches zero. Taking the limit of the $p^\eta_\gamma(j|i, .)$ as $\eta \to 0$ gives the *unperturbed process*. Unlike the process with $\eta > 0$, the unperturbed process need not be irreducible and may have multiple recurrent classes. In fact, the recurrent classes of the unperturbed process are the absorbing states in which all α- and β-agents coordinate on the same strategy, and each agent type receives nonzero payoff (Young 1998a). We shall denote by E_i, $i \in \{0, \ldots, n\}$, the state in which all agents play strategy i, $x_\alpha(i) = N, x_\beta(i) = N$. The absorbing states of the process are precisely those in the set $\Lambda := \{E_i\ :\ \pi_\alpha(i), \pi_\beta(i) >$

$0\}$. Following Young (1993), we refer to these states as *conventions*. Assume that there are at least two conventions, $|\Lambda| \geq 2$.

We shall write as $x^{\gamma,i,j}$ the state induced from x by a γ-population ($\gamma = \alpha$ or β) agent's strategy change from i to j. An important quantity in the study of perturbed adaptive processes is the *resistance* of transitions between states x and y, which measures the rarity of transitions from x to y. Specific to the current context, the resistance $V(x, x^{\gamma,i,j}) := \lim_{\eta \to 0} -\eta \log p_{\gamma}^{\eta}(j|i, x)$ measures the rarity of switches by agents from strategy i to strategy j at state x. Under revision rules satisfying (1), we have

$$V(x, x^{\gamma,i,j}) = \max\{\pi_{\gamma}(l, x) : l \in C_{i,j}\} - \pi_{\gamma}(j, x). \tag{2}$$

Note that $V(.,.)$ is nonnegative but may be equal to zero for some transitions. Transitions for which $V(.,.) > 0$ become rare as $\eta \to 0$. For a convention $x = E_i \in \Lambda$, $V(x, x^{\gamma,i,j}) > 0$ for all $\gamma \in \{\alpha, \beta\}$, $j \neq i$. That is, perturbations are required to move away from a convention.

In a similar way that $V(\cdot, \cdot)$ measures the rarity of single steps in the dynamic, we will use a concept, overall cost, that measures the rarity of a transition between any two states over any number of periods. Let $\mathcal{P}(x, x')$ be the set of finite sequences of states $(x_1, x_2, \ldots, x_T)$ such that $x_1 = x$, $x_T = x'$ and for $\tau = 1, \ldots, T-1$, $x_{\tau+1} = (x_{\tau})^{\gamma,k,l} > 0$ for some γ, k, l. The *overall cost* of a transition between $x, x' \in \Xi$ is:

$$c(x, x') := \min_{\{x_1,\ldots,x_T\} \in \mathcal{P}(x,x')} \sum_{\tau=1}^{T-1} V(x_{\tau}, x_{\tau+1}).$$

Define the basin of attraction $D(E_i)$ of a convention E_i as the set of states from which the unperturbed process converges to E_i and not to any other convention.

$$D(E_i) := \{x \in \Xi : \ c(x, E_i) = 0, \ c(x, E_j) > 0 \ \text{for all} \ j \neq i\}$$

From a given convention, E_i, we seek to determine the lowest cost transition path to some state outside of the convention's basin of attraction, $D(E_i)$. This is known as the first exit problem and is absolutely crucial to the study of invariant distributions and the stochastic stability concept of Young (1993); Kandori et al. (1993).

3 Approximating the costs of escaping a basin of attraction

The need for the approximation result of the current paper arises from the fact that, in contrast to the case of state-independent error probabilities (Young 1998a), the most probable escape path from the basin of attraction of a rest point of the unperturbed dynamic can involve errors being made by both populations.[4] Errors in one population

[4] For state-independent errors, any least cost transition from a given convention to outside its basin of attraction is driven by errors within a single population: errors occur in one of the populations, following

can facilitate errors in the other population by reducing the payoff loss when they are made. The possibility of such escape paths being the most probable ones is caused by population effects and, in contrast to examples given in Alós-Ferrer and Netzer (2010), occurs even for potential games of coordination under asynchronous strategy updating. Previous work does not explicitly study the implications of such transitions (e.g., Belloc and Bowles 2013), deals with the limit as population sizes grow to infinity (Staudigl 2012), or restricts analysis to potential games (Blume 1993, 1997). As precise invariant distributions can be calculated for potential games played under asynchronous dynamics, the latter strand of research has not required study of the first exit problem. Whether or not the first exit problem for such games is worth studying despite this, it remains the case that potential games are non-generic in the class of games considered in the current paper.[5,6]

From a convention $E_i \in \Lambda$, define $\bar{\theta}_{ij}$ as the maximum number of α-agents who can switch to strategy j such that β-agents still receive a higher expected payoff from i than from j, that is, the state remains in the basin of attraction of E_i. Similarly, define $\bar{\zeta}_{ij}$ as the maximum number of β-agents who can switch to strategy j such that the state remains in the basin of attraction of E_i.

$$\bar{\theta}_{ij} = \left\lceil N \frac{\pi_\beta(i)}{\pi_\beta(i) + \pi_\beta(j)} \right\rceil - 1, \qquad \bar{\zeta}_{ij} = \left\lceil N \frac{\pi_\alpha(i)}{\pi_\alpha(i) + \pi_\alpha(j)} \right\rceil - 1,$$

where $\lceil . \rceil$ is a ceiling function, denoting the smallest integer greater than or equal to its argument. These quantities are well defined, as for $E_i \in \Lambda$, $\pi_\alpha(i)$ and $\pi_\beta(i)$ are strictly positive. Note that, by definition, starting from E_i, exactly $\bar{\theta}_{ij} + 1$ (respectively, $\bar{\zeta}_{ij} + 1$) instances of α (respectively, β) agents erroneously choosing j will suffice for the process to exit the basin of attraction of E_i. However, as remarked, this does not provide a lower bound on exit cost, which we now illustrate with an example.

Example 1 Consider the logit dynamics and the following (potential) game:

		β-agent **0**	β-agent **1**
α-agent	**0**	5, 4	0, 0
	1	0, 0	7, 8

We suppose that $N = 5$. Starting from E_0, we compute the minimum cost of escaping $D(E_0)$ via transitions by only one population. As $\bar{\theta}_{01} = 1$, $\bar{\zeta}_{01} = 2$, we have

$$(\bar{\theta}_{01} + 1)\pi_\alpha(0) = 10, \quad (\bar{\zeta}_{01} + 1)\pi_\beta(0) = 12.$$

Footnote 4 continued
which agents from the other population can best/better respond in a way which differs from the initial convention. Errors in the population which best responds differently would be superfluous.

5 If payoffs on the diagonal are independently drawn from continuous distributions, then the resulting game will almost surely not be a potential game.

6 Likewise, by considering the conditions given in Example 1 of Okada and Tercieux (2012), it becomes clear that there does not always exist a "local potential maximizer" (Morris and Ui 2005) in the games we consider. Moreover, even when there does exist a local potential maximizer, the games are not necessarily supermodular, so the results of Okada and Tercieux (2012) have limited applicability in this context.

Next, consider the path of transitions whereby $\bar{\zeta}_{01}$ of the β-agents switch from strategy 0 to strategy 1, following which $\bar{\theta}_{01}+1$ of the α-agents switch from strategy 0 to strategy 1. This gives a path from E_0 to $D(E_1)$, and the cost of the path is given by

$$\begin{aligned}&\bar{\zeta}_{01}\pi_\beta(0)+(\bar{\theta}_{01}+1)\times\left[\frac{1}{N}\left(N-\bar{\zeta}_{01}\right)\pi_\alpha(0)-\frac{1}{N}\bar{\zeta}_{01}\pi_\alpha(1)\right]\\&\quad=2\,\pi_\beta(0)+2\left(\frac{3}{5}\pi_\alpha(0)-\frac{2}{5}\pi_\alpha(1)\right)=8.4\end{aligned}$$

which is smaller than the minimum costs of transitions driven by a single population.

So we see that in Example 1, the least cost transition from E_0 to E_1 requires errors to be made by agents in both populations. This is due to the behavior of the process close to the boundary of the basin of attraction of E_0. After β-agents make errors, the cost of errors by α-agents is reduced. A single error by a β-agent has a lower cost than the consequent reduction in the cost of two errors by α-agents. Two errors by β-agents reduce the cost further still. However, after two errors have been made by β-agents, subsequent errors by β-agents no longer have a linear effect on the cost of an error by an α-agent due to the zero lower bound on $V(\cdot,\cdot)$. Following two errors by β-agents, the cost of a third error by a β-agent is higher than the cost of two errors by α-agents. A moment's consideration leads one to see that, for any given population size, examples can be constructed for which least cost transitions involve errors by both populations. Note that this effect is driven by the finiteness of the populations, that is, by the ceiling functions $\lceil . \rceil$ in the calculations.[7]

Fortunately, when the population size is large, we can exploit linearity to show that, starting from a convention E_i, the cost of the least cost transition path out of the basin of attraction of E_i can be approximated by the lowest cost transitions which involve errors being made by agents in only one of the populations, and those agents making only one type of error. We construct precise bounds for *fixed finite population size* and do not rely on limiting arguments. However, these bounds become more accurate, and hence more useful, as the population size grows.

Theorem 1 *Let $E_i \in \Lambda$ be given. Let $\underline{\theta}_i$ and $\underline{\zeta}_i$ be the minima of $\bar{\theta}_{ij}$, $\bar{\zeta}_{ij}$ over all possible alternative strategies j.*

$$\underline{\theta}_i := \min\left\{\bar{\theta}_{ij} : j\in S,\ j\neq i\right\},\qquad \underline{\zeta}_i := \min\left\{\bar{\zeta}_{ij} : j\in S,\ j\neq i\right\}.$$

Then,

$$\begin{aligned}\min\left\{\pi_\alpha(i)\underline{\theta}_i\,,\ \pi_\beta(i)\underline{\zeta}_i\right\} &\le \min_{j\neq i} c(E_i,E_j)\\ &\le \min\left\{\pi_\alpha(i)(\underline{\theta}_i+1)\,,\ \pi_\beta(i)(\underline{\zeta}_i+1)\right\}.\end{aligned}$$

7 This can be contrasted with Example 3 of Alós-Ferrer and Netzer (2010), whose model is not a population model, and where least cost transitions can involve different players due to the game considered not being a coordination game with zero payoffs off-diagonal.

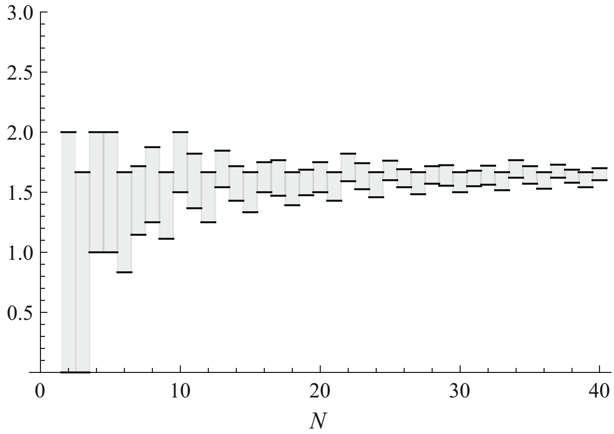

Fig. 1 For the game of Example 1, varying N, here we illustrate the bounds for $c(E_0, E_1)$ given by Theorem 1, normalized by a factor $1/N$

From the definitions of $\bar{\theta}_{ij}$, $\bar{\zeta}_{ij}$, observe that although both bounds in the Theorem increase (asymptotically proportionally) in N, the ratio of the upper to lower bounds approaches unity as $N \to \infty$. This implies that when comparing the costs of exiting the basins of attraction of differing conventions, there is vanishing loss of accuracy (as $N \to \infty$) in considering paths which are driven by errors in only a single population. In the limit, the Theorem extends the result of Staudigl (2012), who uses a different methodology based on optimal control problems to derive such an implication for 2 by 2 games. In Fig. 1, we illustrate Theorem 1 for the game in Example 1. As population size N increases, the lower and upper bounds in the figure converge.

Note that proving Theorem 1 for games with an arbitrary number of strategies is more difficult than proving it for the two-strategy case. Consider paths of length t from convention E_i to outside of its basin of attraction. The number of such paths grows polynomially in the number of alternative strategies (n^t). This is because of the possibility that on a path exiting $D(E_i)$, agents switch from i to j and subsequently to k. In contrast, when there are only two strategies, switches on a path exiting $D(E_i)$ must be from i to j as there are no other alternatives.

The idea of the proof of Theorem 1 is as follows. To estimate the minimum bound for the lowest cost transitions, we study the minimization problem of the cost function over all possible paths escaping $D(E_i)$. Estimation of such minima is complicated when the cost function of a given path loses linearity at the boundary of the basin of attraction, as is illustrated by Example 1. To overcome this problem, we explicitly estimate the size of the basin of attraction of E_i (Lemma 1), showing that to exit $D(E_i)$ requires strictly more than $\underline{\theta}_i$ transitions away from i by α-agents or $\underline{\zeta}_i$ transitions away from i by β-agents. We then restrict attention to the problem of the lowest cost path from E_i to close to the boundary of $D(E_i)$ in the sense of there having been exactly $\underline{\theta}_i$ transitions away from i by α-agents or $\underline{\zeta}_i$ transitions away from i by β-agents. The cost of such a path is a lower bound on the cost of a path that exits $D(E_i)$, and moreover, linearity is retained, giving us the corner solutions that characterize the bounds of the Theorem.

4 Stochastic stability

Consider the (unique) invariant probability measures of the process for fixed η.[8] As $\eta \to 0$, the probability given by these measures to states outside of some set of states $SS \subseteq \Lambda$ approaches zero. States in SS are known as *stochastically stable* states (Young 1993). The transition costs $c(.,.)$ are important quantities for determining SS.

For $E_i \in \Lambda$, let $G(E_i)$ be the set of all directed graphs on Λ satisfying (i) each $E_j \in \Lambda$, $j \neq i$, has outdegree 1, and (ii) the graph has no cycles. For $g \in G(E_i)$, let $E_j \to E_k \in g$ denote an edge from E_j to E_k in graph g. The *stochastic potential* of E_i is defined as

$$SP(E_i) := \min_{g \in G(E_i)} \sum_{E_j \to E_k \in g} c(E_j, E_k).$$

Young (1993); Kandori et al. (1993) show that the stochastically stable states are those that minimize stochastic potential, $SS = \arg\min_{E_i \in \Lambda} SP(E_i)$.[9]

We shall use our Theorem to prove a stochastic stability result under the assumption that there is some coincidence between the best payoff for α-agents and the worst possible payoff for β-agents, and vice versa. This could be the case if agents had to coordinate over some allocation of surplus, with the possibility that one party gets everything while his counterparty gets nothing.

Assumption 1 For $\gamma \in \{\alpha, \beta\}$, let $\pi^*_\gamma := \max_i \pi_\gamma(i)$. Then, there exist $0 \leq j^*_\alpha, j^*_\beta \leq n$ such that $(\pi_\alpha(j^*_\alpha), \pi_\beta(j^*_\alpha)) = (\pi^*_\alpha, 0)$ and $(\pi_\alpha(j^*_\beta), \pi_\beta(j^*_\beta)) = (0, \pi^*_\beta)$.

Consider the following sequence of transitions. Starting from E_i, $\underline{\theta}_i + 1$ α-agents switch to j^*_β. As Assumption 1 implies $\bar{\theta}_{ij^*_\beta} = \underline{\theta}_i$, this makes j^*_β a best response for β-agents (and hence a zero resistance transition). Let all of the β-agents switch to j^*_β. Following this, α-agents have an expected payoff of zero from all possible strategies, so that any j, $0 \leq j \leq n$, is a best response. Let all α-agents switch to some arbitrary j. Following this, j is a best response for β-agents, and we let all β-agents also switch to j, thus attaining E_j. The only non-best response behavior in this sequence of transitions was the initial $\underline{\theta}_i + 1$ switches by α-agents. Therefore, for any $j \neq i$, $c(E_i, E_j) \leq \pi_\alpha(i)(\underline{\theta}_i + 1)$. A similar argument shows that $c(E_i, E_j) \leq \pi_\beta(i)(\underline{\zeta}_i + 1)$. These inequalities allow us to strengthen the inequalities of Theorem 1 by omitting the minimization from the central term. For all $E_j, E_k \in \Lambda$,

$$\begin{aligned} \min\left\{\pi_\alpha(j)\underline{\theta}_j\,,\ \pi_\beta(j)\underline{\zeta}_j\right\} &\leq c(E_j, E_k) \\ &\leq \min\left\{\pi_\alpha(j)(\underline{\theta}_j + 1)\,,\ \pi_\beta(j)(\underline{\zeta}_j + 1)\right\}, \end{aligned}$$

[8] For a proof that these invariant measures are unique, see Hwang and Newton (2014).

[9] See Sandholm (2010) for intricacies that arise for some dynamics regarding the “if and only if” part of this statement.

		β-agent			
		0	**1**	**2**	**3**
α-agent	**0**	2, 0	0, 0	0, 0	0, 0
	1	0, 0	2, 1	0, 0	0, 0
	2	0, 0	0, 0	1, 2	0, 0
	3	0, 0	0, 0	0, 0	0, 3

Fig. 2 Game considered in Example 2

which, substituting the definitions of $\underline{\theta}_j$, $\underline{\zeta}_j$ and taking limits, implies

$$\frac{c(E_j, E_k)}{N} \xrightarrow{N\to\infty} \min\left\{\frac{\pi_\alpha(j)\pi_\beta(j)}{\pi_\beta(j)+\pi^*_\beta}, \frac{\pi_\alpha(j)\pi_\beta(j)}{\pi_\alpha(j)+\pi^*_\alpha}\right\}. \tag{3}$$

Note that the RHS of (3) is independent of k. As any $g \in G(E_i)$ must contain exactly one edge from each E_j, $j \neq i$, the definition of $SP(E_i)$ and (3) implies

$$\frac{SP(E_i)}{N} \xrightarrow{N\to\infty} \sum_{j\neq i} \min\left\{\frac{\pi_\alpha(j)\pi_\beta(j)}{\pi_\beta(j)+\pi^*_\beta}, \frac{\pi_\alpha(j)\pi_\beta(j)}{\pi_\alpha(j)+\pi^*_\alpha}\right\}. \tag{4}$$

As the summation in the RHS of expression (4) omits the summand for $j = i$, it must be that higher values of $\min\left\{\frac{\pi_\alpha(i)\pi_\beta(i)}{\pi_\beta(i)+\pi^*_\beta}, \frac{\pi_\alpha(i)\pi_\beta(i)}{\pi_\alpha(i)+\pi^*_\alpha}\right\}$ correspond to lower values of $SP(E_i)$. We have proved the following corollary to Theorem 1.

Corollary 1 *Under Assumption 1, if strategy $i \in S$ uniquely maximizes $\min\left\{\frac{\pi_\alpha(i)\pi_\beta(i)}{\pi_\beta(i)+\pi^*_\beta}, \frac{\pi_\alpha(i)\pi_\beta(i)}{\pi_\alpha(i)+\pi^*_\alpha}\right\}$, then there exists $\underline{N}$ such that for all $N > \underline{N}$, E_i is the unique stochastically stable state.*

Example 2 Consider the logit dynamics and the game in Fig. 2. $\Lambda = \{E_1, E_2\}$. Calculating the values of $\min\left\{\frac{\pi_\alpha(i)\pi_\beta(i)}{\pi_\beta(i)+\pi^*_\beta}, \frac{\pi_\alpha(i)\pi_\beta(i)}{\pi_\alpha(i)+\pi^*_\alpha}\right\}$, we obtain $1/2$ for $i = 1$ and $2/5$ for $i = 2$. Therefore, by Corollary 1, for large enough N, $SS = \{E_1\}$. Note the dependence on N being large enough. If $N = 1$, then a single mistake by the β-agent suffices to move from E_1 to E_2 at cost $c(E_1, E_2) = 1$. Similarly, a single mistake by the α-agent suffices to move from E_2 to E_1 at cost $c(E_2, E_1) = 1$. Therefore, if $N = 1$, then $SS = \{E_1, E_2\}$.

Under Assumption 1, from any convention E_i, the least cost escape path can be approximated by escape paths to states at which every agent in one position (α or β) plays a strategy that corresponds to an extremal Nash equilibrium that would give them their maximum payoff and would give agents in the other position zero payoff. From these states, agents in the other position obtain zero expected payoff from any strategy, so any strategy is a best response and *any* convention can be reached with zero cost. Hence, Assumption 1 implies that, for large N, the transition cost from a given convention E_i to *any* other convention is approximately equal to the cost of escaping the basin of attraction of E_i. Therefore, if it is harder to escape the basin of attraction of E_i than it is to escape the basin of attraction of any other convention, then E_i will have lower stochastic potential than any other convention and thus be uniquely stochastically stable.

One setup that can satisfy Assumption 1 is when the payoffs on the diagonal of a coordination game represent the efficient frontier of a bargaining set as in Young (1998a). The cited paper found a link between coordination games with many strategies and the Kalai and Smorodinsky (1975) bargaining solution. Naidu et al. (2010) subsequently found a link between coordination games and the Nash (1950) bargaining solution. Both of the above results were for uniform error processes. For logit errors, such problems were unsolved until Hwang and Newton (2014), which uses the results of the current paper.

Of course, Assumption 1 is quite restrictive. The further characterization of stochastic stability for coordination games that are neither potential games nor satisfy Assumption 1 is an open problem, although Theorem 1 should prove helpful in this respect.

5 Proof of theorem

To express transitions by agents from one strategy to another more succinctly, we write

$$e_i^{\alpha} := ((0, \ldots, N, \ldots, 0), (0, \ldots, 0, \ldots, 0)), \quad N \text{ in } i\text{th position}$$
$$e_j^{\beta} := ((0, \ldots, 0, \ldots, 0), (0, \ldots, N, \ldots, 0)), \quad N \text{ in } j\text{th position.}$$

Define $v^{\alpha,i,j}$, $v^{\beta,i,j}$ functions, noting that they are linear in x,

$$\begin{aligned} v^{\alpha,i,j}(x_\beta) &= \frac{x_\beta(i)}{N}\pi_\alpha(i) - \frac{x_\beta(j)}{N}\pi_\alpha(j), \\ v^{\beta,i,j}(x_\alpha) &= \frac{x_\alpha(i)}{N}\pi_\beta(i) - \frac{x_\alpha(j)}{N}\pi_\beta(j). \end{aligned} \tag{5}$$

Note that if $x \in D(E_i)$, then $V(x, x^{\gamma,i,j}) = v^{\gamma,i,j}(x_{\gamma^-})$. In particular, if $x_\alpha(i) \geq N - \bar{\theta}_{ij}$, $x_\alpha(i) + x_\alpha(j) = N$ and $x_\beta(i) \geq N - \bar{\zeta}_{ij}$, $x_\beta(i) + x_\beta(j) = N$, then $x \in D(E_i)$.

For a path $\Gamma = (x_1, x_2, \ldots, x_t) \in \mathcal{P}(x_1, x_t)$, we write $V(\Gamma) := \sum_{l=1}^{t-1} V(x_l, x_{l+1})$. Let

$$\bar{D}(E_j) := \{x \in \Xi : \text{there exists a path } \Gamma \text{ from } x \text{ to } E_j \text{ such that } V(\Gamma) = 0\}.$$

We shall use the notation $(\alpha, k, l; \theta)$ to denote a number θ of α-agents switching, in succession, from action k to action l. Similarly, let $(\beta, k', l'; \zeta)$ denote a number ζ of β-agents switching from action k to action l. Consider a path escaping $D(E_i)$, $\Gamma = (x_1, x_2, \ldots, x_t)$, $x_1 = E_i$, $x_1, \ldots, x_{t-1} \in D(E_i)$ and $x_t \in \bar{D}(E_j)$ for some $j \neq i$. Suppose Γ consists of the following transitions:

$$(\alpha, k_1, l_1; \theta_1) \to (\alpha, k_2, l_2; \theta_2) \to (\beta, k_1', l_1'; \zeta_1) \to (\alpha, k_3, l_3; \theta_3) \to \cdots \to (\beta, k_{M'}', l_{M'}'; \zeta_{M'}) \to (\alpha, k_M, l_M; \zeta_M), \tag{6}$$

where θ_m denotes the number of consecutive transitions in which α-agents switch from k_m to l_m, and ζ_q denotes the number of consecutive transitions in which β-agents switch from k_q' to l_q'. The following lemma states that if the resistance of a transition from i to some other strategy is zero, then it must be the case that the total number of previous transitions from i exceeds either $\underline{\theta}_i$ or $\underline{\zeta}_i$.

Lemma 1 *(Estimation of basin of attraction) For a path of transitions described as in (6), the following statements hold.*
(1) Let y be a state in Γ immediately after the transitions $(\alpha, k_{\tilde{m}}, l_{\tilde{m}}; \theta_{\tilde{m}})$. If $v^{\beta,i,k}(y_\alpha) \leq 0$ for some k, then $\sum_{\{m:k_m=i,\, m\leq\tilde{m}\}} \theta_m > \underline{\theta}_i$
(2) Let y be a state in Γ immediately after the transitions $(\beta, k_{\tilde{q}}', l_{\tilde{q}}'; \zeta_{\tilde{q}})$. If $v^{\alpha,i,k}(y_\beta) \leq 0$ for some k, then $\sum_{\{q:k_q'=i,\, q\leq\tilde{q}\}} \zeta_q > \underline{\zeta}_i$

Proof We first show that (1) holds. First, we establish that if $y_\alpha(i) \geq N - \underline{\theta}_i$, then $v^{\beta,i,k}(y_\alpha) > 0$ for all k. Let $y \in \Xi$ be such that $y_\alpha(i) \geq N - \underline{\theta}_i$. If $y = (e_i^\alpha, y_\beta)$, then $v^{\beta,i,k}(y_\alpha) > 0$ for all k, and we are done. Thus, suppose that $y_\alpha(i) \neq N$. We define

$$c_j := \frac{y_\alpha(j)}{N - y_\alpha(i)}.$$

for $j = 1, 2, \ldots, i-1, i+1, \ldots, n$. Then, $\sum_{j\neq i} c_j = 1$, and

$$\begin{aligned} y_\alpha &= c_1(N - y_\alpha(i), 0, \ldots, 0, y_\alpha(i), 0, \ldots, 0) \\ &\quad + c_2(0, N - y_\alpha(i), \ldots, 0, y_\alpha(i), 0, \cdots, 0) \\ &\quad + c_n(0, \ldots, 0, y_\alpha(i), 0, \ldots, N - y_\alpha(i)) \\ &= \sum_{j\neq i} c_j \left(\frac{y_\alpha(i)}{N} e_i^\alpha + \frac{N - y_\alpha(i)}{N} e_j^\alpha \right). \end{aligned}$$

Therefore,

$$\begin{aligned} v^{\beta,i,k}(y_\alpha) &= v^{\beta,i,k} \left(\sum_{j\neq i} c_j \left(\frac{y_\alpha(i)}{N} e_i^\alpha + \frac{N - y_\alpha(i)}{N} e_j^\alpha \right) \right) \\ &= \sum_{j\neq i} c_j v^{\beta,i,k} \left(\frac{y_\alpha(i)}{N} e_i^\alpha + \frac{N - y_\alpha(i)}{N} e_j^\alpha \right) \\ &\geq v^{\beta,i,k} \left(\frac{y_\alpha(i)}{N} e_i^\alpha + \frac{N - y_\alpha(i)}{N} e_k^\alpha \right) > 0 \end{aligned}$$

where the first inequality follows from the fact that $v^{\beta,i,k}(x)$ is decreasing in $x_\alpha(k)$, and the second inequality follows from $y_\alpha(i) \geq N - \underline{\theta}_i \geq N - \bar{\theta}_{ik}$. This shows that if $y_\alpha(i) \geq N - \underline{\theta}_i$, $v^{\beta,i,k}(y_\alpha) > 0$ for all k. Thus, if $v^{\beta,i,k}(y_\alpha) \leq 0$ for some k, $y_\alpha(i) < N - \underline{\theta}_i$. Let y be the state in Γ immediately after $(\alpha, k_{\tilde{m}}, l_{\tilde{m}}; \theta_{\tilde{m}})$. We have

$$y_\alpha(i) = N - \sum_{\substack{\{m:k_m=i,\\ m\leq\tilde{m}\}}} \theta_m + \sum_{\substack{\{m:l_m=i,\\ m\leq\tilde{m}\}}} \theta_m$$

So

$$\sum_{\substack{\{m:k_m=i,\\ m\leq\tilde{m}\}}} \theta_m = N - y_\alpha(i) + \sum_{\substack{\{m:l_m=i,\\ m\leq\tilde{m}\}}} \theta_m \geq N - y_\alpha(i) > N - (N - \underline{\theta}_i) = \underline{\theta}_i,$$

and Statement (1) of the Lemma is proven. Statement (2) follows similarly. □

Consider again a path Γ from E_i to $\bar{D}(E_j)$ described as in (6). We seek a lower bound for $V(\Gamma)$. As Γ exits $D(E_i)$, and for $x \notin D(E_i)$ we have $v^{\alpha,i,k}(x_\beta) \leq 0$ or $v^{\alpha,i,k}(x_\alpha) \leq 0$, by Lemma 1, it must be that $\sum_{\{m:k_m=i\}} \theta_m > \underline{\theta}_i$ or $\sum_{\{q:k_q=i\}} \zeta_q > \underline{\zeta}_i$. Let $\tilde{\Gamma}$ be the path that is identical to Γ up until the point at which either exactly $\underline{\theta}_i$ transitions away from i have been made by α-agents or exactly $\underline{\zeta}_i$ transitions away from i have been made by β-agents, whichever occurs first, at which point $\tilde{\Gamma}$ terminates. By the definition of $V(.)$, it must be that $V(\Gamma)$ is bounded below by $V(\tilde{\Gamma})$.

Without loss of generality, assume that on Γ the $\underline{\theta}_i$th transition away from i by an α-agent occurs before the $\underline{\zeta}_i$th transition away from i (if it occurs) by a β-agent. Then, letting $\tilde{M}$ index the final transitions of $\tilde{\Gamma}$, we have $\sum_{\substack{\{m:k_m=i,\\ m\leq\tilde{M}\}}} \theta_m = \underline{\theta}_i$.

If x is the state in $\tilde{\Gamma}$ immediately after $(\alpha, k_{\tilde{m}}, l_{\tilde{m}}; \theta_{\tilde{m}})$, then

$$V(x, x^{\beta,i,l}) = v^{\beta,i,l}(x_\alpha) \geq v^{\beta,i,l}(y_\alpha^{\tilde{m}}), \text{ where}$$
$$y_\alpha^{\tilde{m}}(i) = N - \sum_{\substack{\{m:k_m=i,\\ m\leq\tilde{m}\}}} \theta_m, \; y_\alpha^{\tilde{m}}(l) = \sum_{\substack{\{m:k_m=i,\\ m\leq\tilde{m}\}}} \theta_m.$$

That is, the resistance $V(.,.)$, of a transition by a β-agent from i to l is bounded below by what the resistance would be if every previous transition away from i were to l and there were no other transitions. A similar inequality applies for $V(x, x^{\alpha,i,l})$.

Let q_m denote the index of the latest transition by β-agents prior to the mth transition by α-agents. Similarly, let m_q denote the index of the latest transition by α-agents prior to the qth transition by α-agents. Define

$$r_m(\zeta_1, \ldots, \zeta_{q_m}) := v^{\alpha,i,l_m}(y_\beta^{q_m}), \;\; u_q(\theta_1, \ldots, \theta_{m_q}) := v^{\beta,i,l'_q}(y_\alpha^{m_q}).$$

Note that r_m, u_q are affine functions. Then, omitting any terms in $V(\tilde{\Gamma})$ related to transitions other than those from i ($k_m, k'_q \neq i$), $V(\tilde{\Gamma})$, and hence $V(\Gamma)$, is bounded below by

$$\varphi(\theta,\zeta) := r_1\theta_1 + u_1(\theta_1)\zeta_1 + \cdots + r_{m_{\tilde{q}}}(\zeta_1,\ldots,\zeta_{\tilde{q}-1})\theta_{m_{\tilde{q}}} + u_{\tilde{q}}(\theta_1,\ldots,\theta_{m_{\tilde{q}}})\zeta_{\tilde{q}} + \cdots + r_{\tilde{M}}(\zeta_1,\ldots,\zeta_{\tilde{M}'})\theta_{\tilde{M}}. \quad (7)$$

This implies that $V(\Gamma)$ is bounded below by the solution to the following minimization problem:

$$\min\left\{\varphi(\theta,\zeta) : 0 \le \theta_m \le \underline{\theta}_i,\ 0 \le \zeta_q \le \underline{\zeta}_i,\ 1 \le m \le \tilde{M},\ 1 \le q \le \tilde{M}',\ \sum_{m=1}^{\tilde{M}}\theta_m = \underline{\theta}_i,\ \sum_{q=1}^{\tilde{M}'}\zeta_q \le \underline{\zeta}_i\right\}. \quad (8)$$

Similar problems can be defined for φ functions whose last term has a $u_{\tilde{M}'}(.)$ rather than a $r_{\tilde{M}}(.)$. Note that by definition of $\underline{\theta}_i$, $\underline{\zeta}_i$, we have that $v^{\alpha,i,l_m}(y_\beta^{q_m})$, $v^{\beta,i,l'_q}(y_\alpha^{m_q})$ and hence $r_m(\zeta_1,\ldots,\zeta_{q_m})$, $u_q(\theta_1,\ldots,\theta_{m_q})$ are strictly positive as long as $\sum_{m=1}^{\tilde{M}}\theta_m \le \underline{\theta}_i$ and $\sum_{q=1}^{\tilde{M}'}\zeta_q \le \underline{\zeta}_i$.

Proof of Theorem 1 Consider again a path Γ from E_i to $\bar{D}(E_j)$ described as in (6). By Lemma 1, such a path must include either at least $\underline{\theta}_i$ transitions away from i by α-agents, or at least $\underline{\zeta}_i$ transitions away from i by β-agents. Assume that on the path Γ, the $\underline{\theta}_i$th transition away from i by an α-agent occurs before the $\underline{\zeta}_i$th transition away from i (if it occurs) by a β-agent (the alternative case follows similarly). Consider a φ function based on Γ and the minimization problem given by (7) and (8). Let (θ^*,ζ^*) be the optimal choices. As $\sum_{m=1}^{\tilde{M}}\theta_m = \underline{\theta}_i^*$, $\sum_{q=1}^{\tilde{M}'}\zeta_q^* \le \underline{\zeta}_i$, we have $r_m(\zeta^*) > 0$, $u_q(\theta^*) > 0$ for all m, q.

We first show that $\theta_{\bar{m}}^* = \underline{\theta}_i$ for some $\bar{m}$ and $\theta_m^* = 0$ for all $m \ne \bar{m}$. Suppose $\theta_{m_1}^*$ and $\theta_{m_2}^*$ such that $m_1 < m_2$ and $0 < \theta_{m_1}^*, \theta_{m_2}^* < \underline{\theta}_i$. Now, by linearity, replacing $(\theta_{m_1}^*, \theta_{m_2}^*)$ by either $(\theta_{m_1}^*+1, \theta_{m_2}^*-1)$ or $(\theta_{m_1}^*-1, \theta_{m_2}^*+1)$ must lead to a weakly lower value of $\varphi(.,.)$. A strictly lower value contradicts optimality of θ^*. If the new value equals $\varphi(\theta^*,\zeta^*)$, then repeat the argument until $\theta_{m_1}^* = 0$ or $\theta_{m_2}^* = 0$. Repeat with other pairs until $\theta_{\bar{m}}^* = \underline{\theta}_i$ for some $\bar{m}$ and $\theta_m^* = 0$ for all $m \ne \bar{m}$.

For $q > q_{\bar{m}}$, setting $\zeta_q^* > 0$ cannot help as $\theta_m = 0$ for $m > \bar{m}$, so any change to $r_m(.)$ will have no effect. So $\zeta_q^* = 0$ for $q > q_{\bar{m}}$. For $q \le q_{\bar{m}}$, as $u_q(\theta^*) > 0$ for all q and $\theta_m = 0$ for $m < \bar{m}$, by a similar linearity argument to that used above, it must be that $\zeta_q^* \in \{0, \underline{\zeta}_i\}$. If $\zeta_q^* = 0$ for all $q \le q_{\bar{m}}$, then $r_{\bar{m}}(\zeta^*) = v^{\alpha,i,l_{\bar{m}}}(e_i^\beta) = \pi_\alpha(i)$. If $\zeta_{\bar{q}}^* = \underline{\zeta}_i$ for some $\bar{q} \le q_{\bar{m}}$, then $u_{\bar{q}}(\theta^*) = v^{\beta,i,l_{\bar{q}}}(e_i^\alpha) = \pi_\beta(i)$. Taken together, these values for (θ^*,ζ^*) imply that $\varphi(\theta^*,\zeta^*) \ge \min\{\pi_\alpha(i)\underline{\theta}_i, \pi_\beta(i)\underline{\zeta}_i\}$.

Similarly, the alternative case of the $\underline{\zeta}_i$th transition away from i by an β-agent occurring before the $\underline{\theta}_i$th transition away from i (if it occurs) by an α-agent leads to the same lower bound.

Concerning the upper bound, let j_α^* and j_β^* be the states to which the direct escaping costs are minimal, that is, j_α^* solves $\min_{j\in S}(\bar{\theta}_{ij}+1)$, and j_β^* solves $\min_{j\in S}(\bar{\zeta}_{ij}+1)$. The upper bound follows by either choosing a path consisting solely of $\bar{\theta}_{ij^*_\alpha} + 1$ transitions by α-agents from i to j^*_α, or a path comprising $\bar{\zeta}_{ij^*_\beta} + 1$ transitions by β-agents from i to j^*_β. □

References

Alós-Ferrer, C., Netzer, N.: The logit-response dynamics. Games Econ. Behav. **68**, 413–427 (2010)

Belloc, M., Bowles, S.: The persistence of inferior cultural-institutional conventions. Am. Econ. Rev. **103**, 93–98 (2013)

Blume, L.E.: The statistical mechanics of strategic interaction. Games Econ. Behav. **5**, 387–424 (1993)

Blume, L.E.: Population games. In: Arthur, W.B., Durlauf, S.N., Lane, D.A. (eds.) The Economy as an Evolving Complex System II, pp. 425–460. Addison-Wesley, Reading, Boston (1997)

Diamond, P.A.: Wage determination and efficiency in search equilibrium. Rev. Econ. Stud. **49**, 217–227 (1982)

Dindoš, M., Mezzetti, C.: Better-reply dynamics and global convergence to Nash equilibrium in aggregative games. Games Econ. Behav. **54**, 261–292 (2006)

Friedman, J., Mezzetti, C.: Learning in games by random sampling. J. Econ. Theory **98**, 55–84 (2001)

Hwang, S.H., Lim, W., Neary, P., Newton, J.: Conventional contracts, intentional behavior and logit choice: equality without symmetry. Working Papers. In preparation (2016)

Hwang, S.H., Newton, J.: A classification of bargaining solutions by evolutionary origin. Working Papers 2014-02. University of Sydney, School of Economics (2014)

Josephson, J.: Stochastic better-reply dynamics in finite games. Econ. Theory **35**, 381–389 (2008)

Kalai, E.: Proportional solutions to bargaining situations: interpersonal utility comparisons. Econometrica **45**, 1623–1630 (1977)

Kalai, E., Smorodinsky, M.: Other solutions to nash's bargaining problem. Econometrica **43**, 513–18 (1975)

Kandori, M., Mailath, G.J., Rob, R.: Learning, mutation, and long run equilibria in games. Econometrica **61**, 29–56 (1993)

Lagos, R., Rocheteau, G., Wright, R.: Liquidity: a new monetarist perspective. J. Econ. Lit. Forthcoming (2016) (**Forthcoming**)

Lundberg, S., Pollak, R.A.: Separate spheres bargaining and the marriage market. J. Political Econ. **101**(6), 988–1010 (1993)

Manser, M., Brown, M.: Marriage and household decision-making: a bargaining analysis. Int. Econ. Rev. **21**, 31–44 (1980)

McElroy, M.B., Horney, M.J.: Nash-bargained household decisions: toward a generalization of the theory of demand. Int. Econ. Rev. **22**(2), 333–349 (1981)

Monderer, D., Shapley, L.S.: Potential games. Games Econ. Behav. **14**, 124–143 (1996)

Morris, S., Ui, T.: Generalized potentials and robust sets of equilibria. J. Econ. Theory **124**, 45–78 (2005)

Mortensen, D.T., Pissarides, C.A.: Job creation and job destruction in the theory of unemployment. Rev. Econ. Stud. **61**, 397–415 (1994)

Mortensen, D.T., Pissarides, C.A.: New developments in models of search in the labor market. Handbook Labor Econ. **3**, 2567–2627 (1999)

Naidu, S., Hwang, S.H., Bowles, S.: Evolutionary bargaining with intentional idiosyncratic play. Econ. Lett. **109**, 31–33 (2010)

Nash, John F.J.: The bargaining problem. Econometrica **18**, 155–162 (1950)

Newton, J.: Coalitional stochastic stability. Games Econ. Behav. **75**, 842–854 (2012)

Newton, J., Angus, S.: Coalitions, tipping points and the speed of evolution. University of Sydney Economics Working Paper Series 2013-02 (2013)

Newton, J., Angus, S.D.: Coalitions, tipping points and the speed of evolution. J. Econ. Theory **157**, 172–187 (2015)

Okada, D., Tercieux, O.: Log-linear dynamics and local potential. J. Econ. Theory **147**, 1140–1164 (2012)

Peski, M.: Generalized risk-dominance and asymmetric dynamics. J. Econ. Theory **145**, 216–248 (2010)

Rogerson, R., Shimer, R., Wright, R.: Search-theoretic models of the labor market: a survey. J. Econ. Literature **43**, 959–988 (2005)

Sandholm, W., Staudigl, M.: Large deviations and stochastic stability in the small noise double limit. Theor. Econ. **11**, 279–355 (2016)

Sandholm (2010) should appear before “Sandholm et al. (2010)”. Please arrange correctly
Staudigl, M.: Stochastic stability in asymmetric binary choice coordination games. Games Econ. Behav. **75**, 372–401 (2012)
Young, H.P.: The evolution of conventions. Econometrica **61**, 57–84 (1993)
Young, H.P.: Conventional contracts. Rev. Econ. Stud. **65**, 773–792 (1998a)
Young, H.P.: Individual strategy and social structure. Princeton University Press, Princeton (1998)