Journal of Economic Dynamics & Control 124 (2021) 104074

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

When are efficient conventions selected in networks?☆

Carlos Alós-Ferrer[a,*], Johannes Buckenmaier[a], Federica Farolfi[b]

[a] *Zurich Center for Neuroeconomics (ZNE), Department of Economics, University of Zurich, Blümlisalpstrasse 10, Zurich CH-8006, Switzerland*

[b] *Center for Research in Experimental Economics (CREED), Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, Amsterdam 1018 WB, the Netherlands*

A R T I C L E   I N F O

Article history:
Received 17 August 2020
Revised 17 November 2020
Accepted 12 January 2021
Available online 23 January 2021

JEL classification:
C63
C72
C73
D83

Keywords:
Agent-based models
Coordination games
Local interactions
Learning
Networks

A B S T R A C T

We study the determinants of convergence to efficient conventions in coordination games played on networks, when agents focus on past performance (imitative play). Previous theoretical results provide an incomplete picture and suggest potentially-complex interactions between the features of dynamics and behavior. We conducted an extensive simulation study (with approximately 1.12 million simulations) varying network size, interaction and information radius, the probability of actual interaction, the probability of mistakes, tie-breaking rules, and the process governing revision opportunities. Our main result is that "more interactions," be it in the form of larger interaction neighborhoods or of a higher interaction probability, lead to less coordination on efficient conventions. A second observation, confirming previous but partial theoretical results, is that a large network size *relative* to the size of neighborhoods (a "large world") facilitates convergence to efficient conventions. Third, a larger information neighborhood helps efficiency because it increases visibility of efficient payoffs across the network. Last, technical details of the dynamic specification as tie-breaking or inertia, while often relevant for specific theoretical results, appear to be of little empirical relevance in the larger space of dynamics.

1. Introduction

Coordination games epitomize a host of economic problems, from the adoption of technological standards or social conventions to trade policies and international migration flows. The selection of specific equilibria (or "conventions," in the terms of Young, 1993) for certain dynamics or certain kinds of individual-level behavior can be viewed as a stylized summary of the prediction that a society will eventually achieve efficiency or not. It is well-known, however, that even when Pareto-efficient equilibria are available, in many cases, alternative ones are selected. This is especially true when the alternative equilibria are risk-dominant (Harsanyi et al., 1988), which for 2×2 coordination games simply means that they maximize payoffs under the assumption that opponents randomize uniformly.

☆ Financial support from the German Research Foundation (DFG) through project AL-1169/5-1 is gratefully acknowledged. We also thank the Regional Computing Center of the University of Cologne (RRZK) for providing computing time on the High Performance Computing (HPC) system CHEOPS, funded by the German Research Foundation (DFG).

* Corresponding author.
E-mail address: carlos.alos-ferrer@econ.uzh.ch (C. Alós-Ferrer).

https://doi.org/10.1016/j.jedc.2021.104074

Fixing a particular behavioral rule is not enough to derive a prediction. In a seminal contribution, Kandori et al. (1993) showed that imitate-the-best rules lead to risk-dominant conventions under global interactions (interacting with all other agents), but Robson and Vega-Redondo (1996) found that the very same rules lead to Pareto-efficient conventions if interactions are governed by random matching instead. That is, with frequent interactions (play against all other agents each period), risk dominance obtains, while with sporadic interactions (play against a randomly-sampled single agent each period), Pareto-efficiency is selected. Also, natural extensions of a given rule can dramatically change selection results. For example, even under global interactions, the addition of long-enough memory leads to the selection of Pareto efficient equilibria for imitative rules (Alós-Ferrer, 2008).

The nature of interactions is particularly consequential as well. A large literature has studied selection results when coordination games are played on networks, reflecting the sensible assumption that economic interactions are typically local in nature (see Weidenholzer, 2010, for a review of earlier results in this literature). It is well-known that myopic best reply favors risk-dominant conventions or generalizations thereof in networks, starting with the circular city or checkerboard models (Ellison, 1993; 2000).[1] Going back to imitation, in the case of games on networks, Alós-Ferrer and Weidenholzer (2006) showed that whether risk-dominant or Pareto-efficient conventions are selected depends on the interaction radius, with a larger radius favoring risk-dominance. Alós-Ferrer and Weidenholzer (2008) introduced the distinction between interaction and information and provided a selection result for general networks. That work considered the case of informational spillovers, where interaction is local, and information extends at least slightly beyond the interaction neighborhood. Specifically, for a large class of networks, Pareto efficiency wins the day provided information flows beyond the confines of interaction neighborhoods. To date, this remains the only selection result for general networks in this area. A few papers (Cui and Wang, 2016; Khan, 2014) have identified conditions for the selection of efficient conventions in a diametrically opposed situation, where information is limited to the closest neighbors in a network, but interactions are global and follow random matching in the full population.

Putting together all these partial results paints an incomplete picture of the determinants of the selection of efficient or risk-dominant equilibria. While we understand the role of certain dimensions in isolation and for particular cases (imitation vs. best reply, local vs. global interactions, informational spillovers or lack thereof, etc.), those dimensions interact in complex ways with each other and with other features of the dynamics, the game, and the network. In this work, we seek to shed light on the conditions leading to long-run efficiency in coordination games played on networks, when agents follow imitation rules. We focus on the *imitate-the-best rule* where players imitate the best-performing strategy that they observe.[2] We consider two main dimensions: First, the information-interaction structure agents operate in. Second, the frequency of interactions within one's network. For the latter, one can view the contributions of Kandori et al. (1993) and Robson and Vega-Redondo (1996) as the extreme points of the spectrum, which yield vastly different selection results (risk dominance and Pareto efficiency, respectively). One of our goals is to more gradually explore the space between those end points.

We conducted extensive (over one million) agent-based simulations using a supercomputer, covering a wide range of information-interaction structures and interaction frequencies. Compared to the classic stochastic stability analysis that relies on a double-limit approach, a computational approach has the advantage that one can study and compare results for a much larger space within a single study and, hence, it is useful when the goal is to identify general drivers of selection instead of considering specific examples. A drawback is, however, that computational "selection" results are always probabilistic in nature, and hence, we can only speak of the likelihood of survival and not of the clear-cut selection that is typical of theoretical results.

Our main result is that, generally, "more interactions" decrease the likelihood of full coordination on the efficient equilibrium. That is, a larger interaction neighborhood or more frequent interactions within a given neighborhood shift selection away from efficiency and toward risk-dominant conventions. The results also extend and confirm previous theoretical evidence suggesting that a large network relative to the size of the interaction neighborhoods supports efficiency.

The analysis also allows us to examine the robustness of the results with respect to a number of additional dimensions. One of the weaknesses of the literature is that some of the punctual results obtained over the years depend on specific details of the dynamics. Two such details are particularly worrying, namely tie-breaking assumptions and the specification of revision opportunities. The former refers to the behavior of agents when the behavioural rule identifies more than one option, e.g. several different strategies yield the highest observed payoff (in the context of imitation). In this case, most models specify that any acceptable option be tried with positive probability ("no reason not to deviate"), while other contributions (e.g. Oechssler, 1997) argue that, on the basis of e.g. unmodeled switching costs, agents who are currently playing an "optimal" action should not switch to other such actions. Revision opportunities refer essentially to the possibility of inertia and the speed of the learning dynamics in terms of updating. Models as Kandori et al. (1993) specify that every agent has an

[1] Allowing for more than two strategies, and for particular versions of the best reply dynamics (asynchronous updating), Peski (2010) has shown that strengthening risk-dominance to strict 1/2-dominance guarantees unique selection for a large class of networks (including all where each agent has an even number of neighbors). However, strict 1/2-dominance is a very demanding concept, and Alós-Ferrer and Weidenholzer (2007) showed that weaker generalizations of risk dominance fail to guarantee unique selection even for circular cities for games with four or more strategies.

[2] This rule has several interesting features. First, it can be seen as naive in the sense that for large neighborhoods it poses a minimal computational burden on agents. For example, alternative rules such as imitating the strategy which yields the highest average payoffs in the neighborhood typically are computationally more demanding. Second, it nicely captures the well-documented human tendency to focus on salient outcomes such as those leading to high payoffs (Barron and Erev, 2003; Erev and Barron, 2005).

independent probability of being able to revise in any given period ("independent inertia"). With the remaining probability, the agent upholds the current strategy. Other models use simultaneous learning (Alós-Ferrer, 2008), where there is no inertia, or asynchronous learning (e.g. Blume, 1993; 2003; Peski, 2010), where each period one and only one agent has the opportunity to revise. Different specifications of revision opportunities and tie-breaking rules can give rise to crucial differences in the dynamics (see Alós-Ferrer, 2003; Alós-Ferrer and Netzer, 2015). In our simulations, we use different tie-breaking rules and different revision processes and find no effect on selection. We conclude that although revision opportunities and tie-breaking assumptions are sometimes relevant for theoretical knife-edge results, their empirical relevance seems limited.

The remainder of the paper is structured as follows. Section 2 presents the coordination game and the learning environment. Section 3 describes the agent-based simulations. Section 4 presents the main results. Section 5 discusses the robustness analysis. Section 6 concludes.

2. The model

We study behavior in a general 2×2 coordination game. Following the literature (e.g., Alós-Ferrer and Weidenholzer, 2008), we normalize the payoffs as follows,

	P	R
P	1,1	$0, \alpha$
R	$\alpha, 0$	β, β

where $0 < \alpha < \beta < 1$ and $\alpha + \beta > 1$. This game has two strict equilibria given by (P, P) and (R, R). The former outcome is Pareto-efficient, whereas the latter one is *risk-dominant* in the sense of Harsanyi et al. (1988). We also write $\pi(s, s')$ for the payoff when a player plays s against an opponent adopting s', with $s, s' \in \{P, R\}$.

There is a population I of N players arranged on a circle so that the immediate neighbors of a player i are $i-1$ and $i+1$ (modulo N). Players play the game recurrently in discrete time $t = 1, 2, \ldots$ against a subset of their $2k$ closest neighbors with $k \in \{1, \ldots, \lfloor \frac{N}{2} \rfloor\}$. For a given k the potential interaction neighborhood of player i is $K(i) = \{i-k, \ldots, i-1, i+1, \ldots, i+k\}$ for $k \neq \lfloor \frac{N}{2} \rfloor$ and $K(i) = I \setminus \{i\}$ for $k = \lfloor \frac{N}{2} \rfloor$. That is, the set $K(i)$ describes the set of feasible interactions of player i, and the collection of potential interaction neighborhoods $(K(1), \ldots, K(N))$ defines a graph with undirected links. However, agent i does not necessarily interact with all neighbors in $K(i)$, but only with a random subset $R^t(i) \subseteq K(i)$ of realized interaction neighbors and plays the above coordination game against each of them. The size of the subset $R^t(i)$ essentially captures the frequency of interactions. Interactions are symmetric in the sense that any realization $R^t = \left(R^t(i)\right)_{i \in I}$ satisfies $j \in R^t(i)$ if and only if $i \in R^t(j)$ for all i, j.

The model described so far corresponds to the $2k$-model within a linear city as in Ellison (1993), Alós-Ferrer and Weidenholzer (2006), or Alós-Ferrer and Weidenholzer (2008, Section 4), with the addition of random sampling within the interaction neighborhoods. We now further specify the latter as follows. Each potential link between two agents i, j (that is, such that $j \in K(i)$ and hence $i \in K(j)$) is actually realized with (independent) probability $p \in (0, 1]$. In particular, a given agent i actually interacts with any given potential neighbor in $K(i)$ with probability p. For $p \in (0, 1)$, every given agent i has a strictly positive probability of interacting with any subset of $K(i)$. For $p = 1$, there is no random sampling and player i always interacts with all his potential interaction neighbors in every period, that is, $R^t(i) = K(i)$ for all t.

Given a strategy profile $\omega = (s_j)_{j \in I}$ the payoff of player i is

$$\Pi_i(\omega) = \frac{1}{|R^t(i)|} \sum_{j \in R^t(i)} \pi(s_i, s_j).$$

That is, players are concerned about relative average payoffs or in other words about the payoff per interaction. Agent i observes only the strategies chosen and the payoffs obtained by all players in her information neighborhood $M(i)$. For $m \in \{1, \ldots, \lfloor \frac{N}{2} \rfloor\}$ we define the information neighborhood of player i as $M(i) = \{i-m, \ldots, i-1, i, i+1, i+m\}$ for $m \neq \lfloor \frac{N}{2} \rfloor$ and $M(i) = I$ for $m = \lfloor \frac{N}{2} \rfloor$. Players follow an *imitate-the-best rule* as in, e.g., Robson and Vega-Redondo (1996) or Alós-Ferrer and Weidenholzer (2008), with player i choosing a strategy in the set

$$B_i(\omega) = \{s'_i \mid s'_i = s_j \text{ with } j \in M(i), \Pi_j(\omega) \geq \Pi_l(\omega)\ \forall l \in M(i)\}.$$

Potential ties are broken according to a tie-breaking rule $\mathcal{T}(\omega)$ specifying the probability that P is chosen in case $B_i(\omega) = \{P, R\}$. We consider two distinguished rules: The *cautious tie-breaking* rule $\mathcal{T}^C$ specifies $\mathcal{T}^C(\omega) = 1$ if $s_i = P$ and $\mathcal{T}^C(\omega) = 0$ otherwise. The *random tie-breaking* rule $\mathcal{T}^R$ specifies that all ties are broken randomly, that is, $\mathcal{T}^R(\omega) = \frac{1}{2}$. Both rules apply only if there is actually a tie.

Revision opportunities determine which players receive the chance to update their strategies in a given period. Specifically, a *revision process* q is a probability measure on $\mathcal{P}(I)$ that specifies for each $J \subseteq I$ the probability $q(J)$ that exactly the players in J receive an opportunity to revise (Alós-Ferrer and Netzer, 2010; 2015). We consider two types of revision processes: the *simultaneous learning* process specifies $q(I) = 1$ and $q(J) = 0$ for $J \neq I$, that is, in each period all players receive a revision opportunity. The *independent inertia* process specifies $q(J) = (1-\rho)^{|J|} \rho^{N-|J|}$ where the inertia parameter $0 < \rho < 1$ gives the probability that any given player does not receive a revision opportunity.

Set parameters: N, ε, $\mathcal{T}$, ρ, k, p, m.

Period 0: Select an initial population profile $\omega_0 \in \{P, R\}^N$ where each $s_i(0)$, $i = 1, \dots, N$, is randomly drawn from the uniform distribution over $\{P, R\}$.

Period t:

1. *Random Sampling:* Determine set of realized interactions $R^t(i) \subseteq K(i)$.
2. *Payoffs:* Determine $\Pi_i(\omega_t) = \frac{1}{|R^t(i)|} \sum_{j \in R^t(i)} \pi(s_i(t), s_j(t))$.
3. *Imitate-the-best*: Determine set of best-performing strategies $B_i(\omega_t)$.
4. *Revision opportunities:* Determine set of revising players J_t according to ρ.
5. *Mistakes:* Draw $\varepsilon_i \sim B(1, \varepsilon)$ (Bernoulli distribution). If $\varepsilon_i = 0$, then $i \in J_t$ randomly imitates one of the strategies in $B_i(\omega_t)$ (with ties broken according to $\mathcal{T}$), and $i \in N \setminus J_t$ chooses $s_i(t+1) = s_i(t)$. If $\varepsilon_i = 1$, agent i experiments and chooses one of the strategies $\{P, R\}$ at random, with equal probabilities.
6. *Population profile for $t+1$:* Set $\omega_{t+1} = (s_i(t+1))_{i=1}^N$.
7. *Termination:* If $t \geq T$, then stop. Otherwise, increase the period counter t by one and proceed to the next period.

Fig. 1. Pseudo-code of an agent-based simulation of length T.

Table 1
Example of a full factorial design.

Description	Parameter	Values
Network size	N	$\{20, 21, \dots, 59, 60\}$
Interaction radius	k	$\{1, \dots, \lfloor \frac{N-1}{2} \rfloor\}$
Information radius	m	$\{1, \dots, \lfloor \frac{N-1}{2} \rfloor\}$
Random Matching	p	$\{0.25, 0.5, 0.75, 1\}$
Tie-breaking	$\mathcal{T}$	$\{\mathcal{T}^R, \mathcal{T}^C\}$
Revision opportunities	ρ	$\{0, 0.05, 0.1\}$
Mutation probability	ε	$\{0.01, 0.05, 0.1\}$

The behavioral rule together with the tie-breaking specification and the revision process defines a dynamic adjustment process described by a Markov chain in discrete time, $t = 1, 2, \dots$. This process is then perturbed by allowing players to make mistakes with a given probability $\varepsilon > 0$. Specifically, if this event occurs ("a mutation happens") the agent selects one of the available strategies randomly, with uniform probabilities. The stochastically stable states are the long-run equilibria of the process as noise vanishes ($\varepsilon \longrightarrow 0$; Kandori et al., 1993; Young, 1993).

3. Agent-based simulations

We translate the dynamics described above into an agent-based simulation protocol described in pseudo-code in Fig. 1. Even for a given game with fixed payoff parameters α and β, the relevant space of the model described in the previous section is seven-dimensional. The information-interaction network is characterized by the network size N, the (potential) interaction radius k, the information radius m, and the random sampling probability p. Fixing the behavioral rule of imitate-the best, the dynamics is characterized by the tie-breaking rule $\mathcal{T}$, the revision process q, and, avoiding the double-limit approach, the mutation probability ε.

To explore this parameter space there are two fundamentally different approaches. One approach would entail a full factorial design using a small number of different values for each parameter, and a fixed number of simulations, say 150, for each parameter combination. Even if such a design is highly discretized, e.g. like the example in Table 1, this approach would require simulations for $3 \times 2 \times 3 \times \frac{N}{2} \times 4 \times \frac{N}{2} = 18 \times N^2$ different parameter combinations for any given network size

Table 2
Range and distribution for the simulation parameters.

Description	Parameter	Values
Network size	N	$N \sim U(\{20, \ldots, 60\})$
Interaction radius	k	$k \sim U(\{1, \ldots, \lfloor \frac{N-1}{2} \rfloor\})$
Information radius	m	$m \sim U(\{1, \ldots, \lfloor \frac{N}{2} \rfloor\})$
Random Matching	p	$p(s,k) = \frac{s}{2k},\ s \sim U(\{\lceil \frac{k}{2} \rceil, \lceil \frac{k}{2} \rceil + 1, \ldots, 2k-1\})$.
Tie-breaking	$\mathcal{T}$	$\mathcal{T} \sim U(\{\mathcal{T}^{R}, \mathcal{T}^{C}\})$
Revision opportunities	ρ	$\rho \sim U(\{0, \delta\}),\ \delta \sim U((0, 0.1])$
Mutation probability	ε	$\varepsilon \sim U([0.01, 0.1])$

Notes: $U(S)$ denotes the uniform distribution over $S \subseteq \mathbb{R}$. $\mathcal{T}^{R}$ denotes Random-Tie-Breaking. $\mathcal{T}^{C}$ denotes Cautious-Tie-Breaking. 1,118,149 simulations were conducted, with randomly-sampled parameter values.

N. Thus this example would require a total of 192,618,000 (about 193 million) simulations. The drawback of this approach is that it requires a high degree of discretization, and even then an exhaustive computational analysis using a full factorial design might remain computationally unpractical.

We rely on an alternative *parameter-sampling approach* that economizes on computation time and requires a much smaller degree of discretization. That is, instead of committing to an ex-ante fixed set of parameters, each simulation uses a set of randomly drawn parameters from the space of interest.

We conducted about 1.12 million simulations following the parameter-sampling approach. Table 2 describes the parameter space and the exact distribution for each parameter. Payoff parameters of the coordination game were fixed at $\alpha = 0.5$ and $\beta = 0.7$.[3]

The simulation length, T, was chosen large enough such that for all conceivable states the empirical distribution $\hat{\mu}_t : \{P, R\}^N \longrightarrow [0, 1]$ does not change by more than $\delta = 10^{-5}$. To ensure this level of precision a length of $T(\delta) \geq \frac{2}{\delta} = 200,000$ rounds suffices. To see this, let $\hat{\mu}_{t-1} : \{P, R\}^N \longrightarrow [0, 1]$ and $\omega_t \in \{P, R\}^N$, and note that $\hat{\mu}_t = \frac{t-1}{t}\hat{\mu}_{t-1} + \frac{1}{t}\mathbb{1}_{\omega_t}$. Then

$$
\begin{aligned}
||\hat{\mu}_t - \hat{\mu}_{t-1}||_1 &= \left|\left| \tfrac{t-1}{t}\hat{\mu}_{t-1} + \tfrac{1}{t}\mathbb{1}_{\omega_t} - \hat{\mu}_{t-1} \right|\right|_1 \\
&= \tfrac{1}{t}||\hat{\mu}_{t-1} - \mathbb{1}_{\omega_t}||_1 = \tfrac{1}{t}\left(\textstyle\sum_{\omega \neq \omega_t} \hat{\mu}_{t-1}(\omega) + |1 - \hat{\mu}_{t-1}(\omega_t)|\right) < \tfrac{2}{t}
\end{aligned}
$$

Hence $||\hat{\mu}_t - \hat{\mu}_{t-1}||_1 < \delta$ holds if $\frac{2}{t} \leq \delta$, which yields the sufficient condition $t \geq \frac{2}{\delta}$. For a precision of $\delta = 10^{-5}$, a length of $T = 200,000$ periods suffices.

4. Results

In this section we present the results of approximately 1.12 million agent-based simulations. The objective of the simulations was to obtain estimates of the (limit) invariant distribution μ^*, which determines the long-run equilibria of the dynamics (see, e.g. Ellison, 2000). By the Ergodic Theorem (Karlin and Taylor, 1975) μ^* can be approximated by the average time spent by the dynamics in each state for (long enough) simulations (as in Alós-Ferrer and Buckenmaier, 2017; Alós-Ferrer et al., 2021). Formally, μ^* is a distribution over the set of absorbing states $\Omega \subseteq \{P, R\}^N$. The monomorphic states $\bar{P} = (P, \ldots, P)$ and $\bar{R} = (R, \ldots, R)$ are always absorbing for imitation-based dynamics. However, there might be other, non-monomorphic, absorbing states. Since we are interested in whether coordination on P or R occurs and to keep the analysis tractable, we focus on the two monomorphic states and aggregate all non-monomorphic ones (absorbing or not) in a residual, denoted *res* (Section 5.2 below briefly discusses non-monomorphc absorbing states and shows that they are empirically inconsequential). Specifically, for each simulation we record the fraction of time spent in each of the monomorphic states as well as the fraction of time spent in non-monomorphic ones. Formally, our main output is the relative probability distribution $f : \{\bar{P}, \bar{R}\} \cup \{\text{res}\} \longrightarrow [0, 1]$.

Note that the simulations are of finite length and rely on a positive (although small) mutation probability. Hence, in contrast to theoretical results that rely on a double-limit approach, mechanistic differences (e.g. due to differences in the mutation probability ε or the overall size of the network N) in the amount of time spent in non-monomorphic states are to be expected. To illustrate this relation, we consider the probability that at least one mutation occurs, $\eta(\varepsilon, N) = 1 - (1 - \varepsilon)^N$, as an indicator for the overall level of noise in the system. Indeed, this measure is highly correlated with the time spent in non-monomorphic states $f(res)$ (Pearson's $\rho = 0.9704,\ N = 1,118,149,\ p < 10^{-7}$).

To control for such mechanistic differences, we focus on the restriction of the relative frequency distribution to $\{\bar{P}, \bar{R}\}$, which we denote by $F : \{\bar{P}, \bar{R}\} \longrightarrow [0, 1]$. That is, we consider the relative amount of time spent at a given monomorphic state conditioning on the total amount of time spent in both monomorphic states. We interpret $F(\bar{P})$ as a measure of convergence toward $\bar{P}$ relative to $\bar{R}$, and take it as a probabilistic version of "selection."

[3] We chose to fix the payoff matrix instead of also varying the payoff parameters for two reasons. First, none of the theoretical results, whose limits we seek to explore, depend on the specifics of the payoff matrix beyond P being Pareto-efficient and R being risk-dominant. Second, adding two further dimensions to the already seven-dimensional parameter space is computationally expensive.

4.1. Interaction vs. information

We consider the distinction between interaction and information introduced in Alós-Ferrer and Weidenholzer (2008) (see also Alós-Ferrer and Weidenholzer, 2014; Cui, 2014).[4] That work considered coordination games with an efficient strategy and showed that, under information spillovers, that is, when the information neighborhood contains and exceeds the interaction neighborhood ($m > k$), a simple condition requiring the network to be large relative to the interaction neighborhood ($N > (2k+1)^2$) guarantees that agents will be able to coordinate on the Pareto-Efficient equilibrium. In contrast, in the absence of information spillovers risk-dominant equilibria will often prevail. Khan (2014) finds that under global random matching, where in each period every agent interacts only with a single randomly selected partner, risk-dominant equilibria may prevail if information is sufficiently restricted. However, the efficient outcome is the unique long-run equilibrium under full observability or if each individual observes at least four other individuals.

An examination of the models described above and others (as e.g. Robson and Vega-Redondo, 1996) suggests that efficiency is related to a particular property. In all models selecting efficient conventions, the combination of behavior, interaction, and information is such that the payoffs of the latter are "sticky" in an informational sense. Informational spillovers help efficiency because the payoffs of efficiency are observed across the network. In contrast, limited observability prevents the flow of information and, hence, favours risk-dominant outcomes. Imitate the best max takes advantage of this fact because it translates high payoffs into high probability of adoption more than linearly, as opposed to, say, imitate the best average.

In the context of our agent-based setting, we thus expect that the key determinants for the selection of payoff-efficient equilibria will be the network size N, the interaction radius k, the information radius m, and the sampling probability p. In light of the theoretical results outlined above, apart from investigating the overall effects of N, k, m, and p, we also ask whether those effects differ within and outside specific regions in that four-dimensional space. Specifically, we consider the regions defined by the following conditions.

Definition 1. For a set of parameters (N, k, m, p) we say that

(1) there are *informational spillovers* if $m > k$,
(2) the network is *relatively large* if $N > (2k+1)^2$,
(3) there is *random sampling* if $p < 1$, and
(4) interactions are *global* if $p = 1$.

The interpretation of conditions (1), (3), and (4) is straightforward. In particular, we remark that our set of simulations contains cases with $m > k$ (as in Alós-Ferrer and Weidenholzer, 2008), with $m = k$ (as in any work not distinguishing information and interaction), and with $m < k$ (analogously to Cui and Wang, 2016; Khan, 2014). Concerning (2), a network is relatively large if it is large relative to the size of its neighborhoods. The idea is that a network with many agents, all of whom are however directly linked or can be linked in a few steps, is a small world, while a network with a limited number of agents who interact with relatively few other agents is a large world (Alós-Ferrer and Weidenholzer, 2008; 2014).

Our objective is twofold. On the one hand, we seek to explore long-run behavior in the vast space between the results of Alós-Ferrer and Weidenholzer (2008) and Khan (2014), which are in some sense opposites. On the other hand, we want to study the tightness of the theoretical conditions (1)–(4) outlined above and, in particular, whether and how network size, information structure, interaction structure, and sampling probability systematically affect selection of the efficient outcome.

We find that on average coordination on P occurs 74.7% of the time when $m > k$, whereas $F(\bar{P})$ is only 39.0% when there are no informational spillovers. This difference is larger when interactions are global (63.9% vs. 26.4%) than when there is random sampling (85.5% vs. 51.7%). When the network is relatively large the efficient outcome is selected 93.6% of the time, whereas this percentage drops to 51.3% when this condition is violated. Again, this difference is larger when interactions are global (global, 93.0% vs. 38.1%; random sampling, 94.4% vs 64.6%). Unsurprisingly when there are informational spillovers and the network is relatively large, the efficient outcome is selected almost universally (99.2%) in line with the theoretical prediction (Alós-Ferrer and Weidenholzer, 2008). In contrast, when at least one of those conditions is violated $F(\bar{P})$ drops drastically to 51.1%. This difference is larger with global interactions (98.3% vs 38.1%) compared to random sampling (100% vs 64.3%).

To confirm these observations and to uncover possible monotonicities, we turn to a regression analysis. Since $F(\bar{P})$ is a frequency, we use fractional logit regressions (Papke and Wooldridge, 2008). Table 3 shows the results of these regressions where we include N, k, m, and p as independent variables.[5] For the full sample of 1.12 million simulations, we find that coordination on P increases in the size of the network (in terms of number of agents) and decreases with the interaction radius. Both effects are in agreement with the results of Alós-Ferrer and Weidenholzer (2008, 2014), which suggest that a large network (relative to the size of the interaction neighborhoods) facilitates coordination on efficient conventions. We also find that coordination on P increases with the size of the information neighborhood, in alignment with the idea sketched

[4] Alós-Ferrer and Weidenholzer (2006) consider the circular city model under local information, that is, interaction and information neighborhood coincide, characterizing whether a risk dominant convention or an efficient convention will be established in the long run. The answer to this question depends on the interaction radius of the individual agents, i.e. on "how local" interactions are.

[5] The results are robust when we additionally control for i) inertia and tie-breaking, ii) $\eta(\varepsilon, N)$, or iii) all of those. See also Section 4.2.

Table 3
Fractional logit regression on $F(\bar{P})$.

$F(\bar{P})$	Full Sample	Spillovers Yes	Spillovers No	Large Network Yes	Large Network No	Random Sampling Yes	Random Sampling No
N	0.022***	0.044***	0.018***	0.003	0.019***	0.019***	0.024***
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
k	−0.188***	−0.236***	−0.148***	−0.362***	−0.167***	−0.178***	−0.198***
	(0.000)	(0.001)	(0.001)	(0.021)	(0.001)	(0.001)	(0.001)
m	0.078***	−0.023***	0.120***	0.637***	0.072***	0.105***	0.057***
	(0.000)	(0.001)	(0.001)	(0.020)	(0.000)	(0.001)	(0.001)
p	−3.430***	−4.737***	−3.073***	−0.608***	−3.502***	−3.752***	
	(0.010)	(0.020)	(0.012)	(0.041)	(0.010)	(0.017)	
Obs.	1,118,149	527,487	590,662	121,919	996,230	558,977	559,172

Notes: Regressions on columns 2–7 are restricted to subsamples as given. Spillovers refer to condition (1); Large Network refers to condition (2); Random Sampling refers t conditions (3) (Yes) and (4) (No). Robust standard errors in parentheses. * $p < 10^{-3}$, ** $p < 10^{-5}$, *** $p < 10^{-7}$.

above that informational spillovers help efficiency because it is easier to observe the efficient payoffs across the network. Last, a larger sampling probability reduces coordination on the efficient outcome, suggesting that more frequent interactions (as in the comparison between Kandori et al., 1993 and Robson and Vega-Redondo, 1996) hinder efficiency. The result that larger neighborhoods lead to a decrease in the probability of $\bar{P}$ can also be cast in the latter terms.

Next, we partition the parameter space into the two regions with and without informational spillovers.[6] In the absence of informational spillovers the results mirror those on the full sample. However, when there are informational spillovers, surprisingly, a further increase in the information radius leads to *less* coordination on P. This negative effect, however, should not be overinterpreted. Since the subsample is restricted to simulations with $m > k$, there is a mechanical effect where small values of m are only possible in this subsample for small values of k, creating a correlation between m and k. The next two models consider the parts of the parameter space where the network is relatively large and where it is not, respectively. For the area where $N < (2k+1)^2$ the results again closely resemble the ones obtained on the full sample, whereas when the network is relatively large a further increase in the network size does not lead to an additional increase in coordination on P. Also, in this case, the effect of the sampling probability becomes much smaller than in other models. Finally, when restricting only to situations with random sampling or global interactions, we again see a positive effect of the network size and the information radius on $F(\bar{P})$, whereas the effect of both the interaction radius and the sampling probability is negative.

Summarizing, we find that "more interactions," be it in the form of a larger interaction neighborhood or of a higher sampling probability, universally lead to less coordination on the payoff-efficient equilibrium. The evidence is also in line with previous results suggesting that a large network relative to the size of the interaction neighborhoods supports efficiency. The effects of the information radius appear to be more nuanced, suggesting that the key network feature is whether information spills over the interaction radius or not, but once this is guaranteed, a larger information radius brings no further positive effect.

4.2. Revision processes and tie-breaking

Different specifications of revision opportunities and tie-breaking rules can give rise to crucial differences in the dynamics (see, e.g. Alós-Ferrer, 2003; Alós-Ferrer and Netzer, 2015). For instance, Alós-Ferrer and Netzer (2010) showed that the selection of potential maximizers for the logit dynamics is actually knife-edge and can vanish if revision opportunities do not follow an asynchronous process, as originally postulated by Blume (1993) and Blume et al. (1997). If a given prediction depends on such details, its strength is greatly diminished.

Revision opportunities are one such detail. Another is given by tie-breaking rules. Typically, behavioral rules determine a set of strategies a player may choose, but additional assumptions are required to specify which strategy in that set is chosen. For example, one might assume that a player randomizes uniformly among all strategies in the set (random tie-breaking) or sticks with her current strategy if it is in the set (cautious tie-breaking). Crucially, how exactly ties are broken may affect the selection of the long-run equilibrium (Alós-Ferrer, 2003; Oechssler, 1997; Sandholm, 1998).

In what follows we study whether the revision process, specifically inertia, and tie-breaking are empirically relevant for the selection between Pareto-efficient and risk-dominant equilibra in finite time. Under simultaneous learning coordination on P occurs 55.8% of the time, which is essentially identical to the time spent at P when there is independent inertia (55.9%). In the latter case, we find also no correlation between the level of inertia and convergence to the efficient outcome (Pearson's $\rho = 0.0301,\ N = 1,118,149,\ p = 0.0301$). Turning to the tie-breaking assumptions, we find that the degree of coordination on P is very similar under random and cautious tie-breaking ($\mathcal{T}^R$, 55.9%; $\mathcal{T}^C$, 55.8%).

[6] An alternative approach would be to consider larger regressions on the full sample including interaction terms. However, the latter essentially deliver the same messages while considerably complicating the discussion.

Table 4
Fractional logit regression on $F(\bar{P})$.

	Full	Spillovers		Large Network		Random Sampling	
$F(\bar{P})$	Sample	Yes	No	Yes	No	Yes	No
SimLearn	0.001	0.001	−0.001	−0.108	0.002	−0.002	0.007
	(0.006)	(0.010)	(0.008)	(0.037)	(0.006)	(0.009)	(0.009)
ρ	0.142	0.058	0.134	−0.743	0.145	0.187	0.134
	(0.093)	(0.154)	(0.130)	(0.572)	(0.098)	(0.139)	(0.132)
RTB	0.002	0.011	−0.005	0.011	0.002	0.004	0.003
	(0.004)	(0.006)	(0.005)	(0.023)	(0.004)	(0.006)	(0.005)
Obs.	1,118,149	527,487	590,662	121,919	996,230	558,977	559,172

Notes: Regressions on columns 2–7 are restricted to subsamples as given. Spillovers refer to condition (1); Large Network refers to condition (2); Random Sampling refers t conditions (3) (Yes) and (4) (No). Robust standard errors in parentheses. * $p < 10^{-3}$, ** $p < 10^{-5}$, *** $p < 10^{-7}$.

Table 5
Fractional logit regression on $F(\bar{P})$.

$F(\bar{P})$	1	2	3	4
N	0.025***	0.022***	0.022***	0.022***
	(0.000)	(0.000)	(0.000)	(0.000)
k	−0.188***	−0.181***	−0.188***	−0.188***
	(0.000)	(0.001)	(0.000)	(0.000)
m	0.078***	0.078***	0.076***	0.078***
	(0.000)	(0.000)	(0.001)	(0.000)
p	−3.431***	−3.431***	−3.431***	−3.136***
	(0.010)	(0.010)	(0.010)	(0.023)
ε	3.524***	2.554***	0.829**	5.706***
	(0.297)	(0.168)	(0.166)	(0.323)
$N \times \varepsilon$	−0.060***			
	(0.007)			
$k \times \varepsilon$		−0.130***		
		(0.015)		
$m \times \varepsilon$			0.035	
			(0.014)	
$p \times \varepsilon$				−5.419***
				(0.375)
Obs.	1,118,149	1,118,149	1,118,149	1,118,149

Notes: Robust standard errors in parentheses. * $p < 10^{-3}$, ** $p < 10^{-5}$, *** $p < 10^{-7}$.

Table 4 shows the results of a series of fractional logit regressions on $F(\bar{P})$. As independent variables we include the probability of inertia (ρ) and a dummies indicating simultaneous learning ($\rho = 0$) or random tie-breaking. On the full sample we find that none of these variables shows a significant effect. It is conceivable, however, that revision opportunities and/or inertia are relevant in specific areas of the parameter space. Hence, we again consider the regions specified by conditions (1)–(4). We find no significant effect of simultaneous learning, inertia, or random-tie breaking in any region. We conclude that although revision opportunities and tie-breaking assumptions are sometimes relevant in theory, empirically their relevance seems limited when time is finite.

4.3. Non-vanishing noise

Theoretical results in the literature typically are concerned with the limit case where the probability of mistakes (ε) vanishes. In contrast, our computational approach is based on a setting with small but positive and non-vanishing noise. In that sense, our work also allows to explore the robustness of previous results (that only hold in the limit as ε goes to zero) for non-vanishing, small levels of noise. This is important because in real settings the environment is constantly evolving and, hence, positive mutation rates are often optimal as they facilitate adaptation to a new optimum following a change in the environment (see Ben-Porath et al., 1993, for a theoretical argument).

In this section, we further explore how the mutation probability affects our results, and in particular coordination on the efficient outcome. To that end, we reconsider the regression on the full sample reported in Table 3 (first column), adding the mutation probability and its interactions with our main variables of interest as regressors. Table 5 reports the results of those regressions. Overall, a higher mutation probability increases coordination on the efficient outcome. Importantly, the main effects of an increase in N, k, m, and p remain unchanged when we additionally control for the mutation probability. The effect of network size, however, becomes smaller as the likelihood of mutation increases (model 1). In contrast, larger mutation rates strengthen the negative effect of neighborhood size on coordination on the efficient convention (model 2).

Table 6
Fractional logit regression on $F(\bar{P})$ controlling for the initial share of P-players.

	Full	Spillovers		Large Network		Random Sampling	
$F(\bar{P})$	Sample	Yes	No	Yes	No	Yes	No
N	0.030***	0.057***	0.035***	0.002	0.030***	0.020***	0.041***
	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
k	−0.281***	−0.387***	−0.233***	−0.363***	−0.269***	−0.245***	−0.324***
	(0.001)	(0.002)	(0.001)	(0.021)	(0.001)	(0.001)	(0.001)
m	0.117***	−0.035***	0.188***	0.641***	0.116***	0.144***	0.094***
	(0.001)	(0.001)	(0.001)	(0.020)	(0.001)	(0.001)	(0.001)
p	−5.291***	−7.740***	−5.123***	−0.614***	−5.887***	−5.361***	
	(0.016)	(0.031)	(0.020)	(0.042)	(0.017)	(0.023)	
InitShareP	23.168***	28.229***	25.625***	2.612***	26.567***	20.105***	26.915***
	(0.060)	(0.100)	(0.090)	(0.136)	(0.069)	(0.077)	(0.094)
Obs.	1,118,149	527,487	590,662	121,919	996,230	558,977	559,172

Notes: Regressions on columns 2–7 are restricted to subsamples as given. Spillovers refer to condition (1); Large Network refers to condition (2); Random Sampling refers t conditions (3) (Yes) and (4) (No). Robust standard errors in parentheses. * $p < 10^{-3}$, ** $p < 10^{-5}$, *** $p < 10^{-7}$.

Also the effect of the sampling probability becomes larger as ε increases (model 3), whereas we observe no significant interaction effect for the size of the information neighborhood.

Summarizing, we find that larger mutation rates increase coordination on the efficient outcome. At the same time, the effects of the size of the interaction neighborhood and the sampling probability increase as the mutation probability increases. In contrast, the effect of network size dampens as ε increases, whereas we find no moderating effect of the mutation probability on the effect of the size of the information neighborhood.

5. Robustness analysis

In the theoretical double limit as $t \to \infty$ and $\varepsilon \to 0$, selection does not depend on the initial conditions. However, this is not necessarily true for finite-horizon agent-based simulations, since the mutation probability ε does not vanish and the time horizon T, although long, is still finite. In the simulations, the initial conditions were random, that is, every position on the circle had an independent probability of 50% of being initially occupied by a P- or an R-player. In this section we report two robustness checks where we take into account the initial distribution of play (Section 5.1), as well as the prediction of the noise-free, unperturbed dynamics (Section 5.2).

5.1. Initial distribution of strategies

A first way to account for path dependence is to consider the initial share of P-players at $t = 0$. Since initial conditions were determined randomly, the average share of P-players across all initial states is exactly 50.0%. We find that the convergence on P obtains 74.0% of the time when the share of P-players in the initial distribution is above one half, whereas it is only 37.1% when there are more R-players than P-players initially. Indeed, we find a clear correlation between the initial share of P-players and convergence towards the payoff-dominant equilibrium (Pearson's $\rho = 0.442,\ p < 10^{-7}$).

Table 6 shows the analogous regressions to Table 3 but additionally controlling for the initial share of P-players. Although the initial distribution is highly predictive of the average time spent at full coordination on P, the results obtained previously are all qualitatively unchanged. Hence, we conclude that our results are robust to the specifics of the initial condition.

5.2. Deterministic prediction

It is important to establish that our results capture essential features of the long-run behavior of the dynamics, and not just whether the initial state lies in the basin of attraction of $\bar{P}$ or $\bar{R}$ for the unperturbed, noise-free dynamics. For, in the latter case, we would merely be capturing the short-run of the dynamics, which has limited explanatory power for selection in the long-run. To that end, we first determined whether an initial state lies in the basin of attraction of $\bar{P}$ or $\bar{R}$ by running a deterministic version of each of our 1.12 million simulations without mutations until either full coordination on P or R was achieved, or 200,000 periods had passed (the exercise is analogous to Lee and Valentinyi, 2000). We then repeated the regressions reported in Table 3 while controlling for the deterministic prediction of the unperturbed dynamics, and hence for potential path-dependence in our finite-horizon simulations. That is, we re-ran all regressions including, as an additional control, the state where the dynamics converges to given the initial condition when $\varepsilon = 0$.

In the deterministic simulations, we find that, overall, 55.1% of the initial states lie in the basin of attraction of $\bar{P}$, whereas the unperturbed dynamics converges to $\bar{R}$ in 44.8% of the cases. In only 881 out of the 1,118,149 initial states ($< 0.1\%$)

Table 7
Fractional logit regression on $F(\bar{P})$ controlling for deterministic prediction.

	Full	Spillovers		Large Network		Random Sampling	
$F(\bar{P})$	Sample	Yes	No	Yes	No	Yes	No
N	-0.004^{***}	0.023^{***}	-0.005^{***}	-0.027^{***}	-0.002^{*}	-0.001	-0.008^{***}
	(0.000)	(0.001)	(0.001)	(0.002)	(0.000)	(0.000)	(0.001)
k	-0.111^{***}	-0.213^{***}	-0.063^{***}	-0.182^{***}	-0.108^{***}	-0.106^{***}	-0.167^{***}
	(0.001)	(0.002)	(0.001)	(0.027)	(0.001)	(0.001)	(0.002)
m	0.105^{***}	-0.014^{***}	0.134^{***}	0.462^{***}	0.102^{***}	0.114^{***}	0.068^{***}
	(0.001)	(0.002)	(0.001)	(0.015)	(0.001)	(0.001)	(0.002)
p	-6.571^{***}	-9.049^{***}	-5.986^{***}	0.207^{*}	-7.080^{***}	-5.135^{***}	
	(0.017)	(0.033)	(0.021)	(0.054)	(0.018)	(0.022)	
DetState	6.115^{***}	6.858^{***}	5.855^{***}	3.328^{***}	6.449^{***}	4.259^{***}	7.701^{***}
	(0.013)	(0.027)	(0.016)	(0.035)	(0.015)	(0.016)	(0.023)
Obs.	1,117,268	527,487	589,781	121,402	995,866	558,977	55,8291

Notes: Regressions on columns 2–7 are restricted to subsamples as given. Spillovers refer to condition (1); Large Network refers to condition (2); Random Sampling refers t conditions (3) (Yes) and (4) (No). Robust standard errors in parentheses. * $p < 10^{-3}$, ** $p < 10^{-5}$, *** $p < 10^{-7}$.

convergence was not achieved within the time limit.[7] Conditional on convergence to a given state, the average time required to reach full coordination was 261 periods for P and 474 periods for R (these relatively short times are unsurprising for the deterministic dynamics).

The results of the regressions are reported in Table 7. After controlling for the deterministic prediction, we see that an increase in the network size tends to have a negative effect on the time spent in the payoff-dominant equilibrium except in the region with informational spillovers. This is in stark contrast to the results obtained previously in Table 3 where the effect of N was positive in all but one model. On the other hand, for the interaction radius k, the information radius m, and the sampling probability p, we find that all results are robust. In particular, the previous conclusion that “more interactions” (larger interaction neighborhood or higher sampling probability) lead to less coordination on the payoff-efficient equilibrium is unaffected. However, the regressions suggest that the effect of absolute network size (that is, in terms of number of agents) is channeled through the unperturbed part of the dynamics. This might be less surprising than it seems. Theoretical results (e.g. Alós-Ferrer and Weidenholzer, 2008) show that, at least in certain subclasses of networks, the size of the basin of attraction of $\bar{R}$ for the unperturbed dynamics is monotonically decreasing in N, while population size does not have a direct effect on the basin of attraction of $\bar{P}$. Thus, the effect of N observed in previous regressions might primarily reflect the differences in the relative sizes of the basins of attraction for the unperturbed dynamics.

6. Conclusion

We ran 1.12 million agent-based simulations to study the determinants of convergence to efficient conventions in coordination games played on networks, when agents focus on past performance (imitative play). Our motivation was the fact that previous theoretical results indicate potentially-complex interactions between the features of dynamics and behavior, and an extensive simulation analysis allows us to study and compare results for a large part of the space of dynamics instead of considering restricted sets thereof for analytical tractability.

Our main result is that “more interactions” lead to less coordination on the payoff-efficient equilibrium. This can be in the form of larger interaction neighborhoods or of a higher sampling probability for a given interaction neighborhood. Interestingly, this is in agreement with the comparison between two classical studies, Kandori et al. (1993) and Robson and Vega-Redondo (1996), although none of those considered networks. In the first, interactions are frequent in the sense that, each period, each agent plays against every other agent, and the dynamics converges to the risk-dominant convention. In the second, interactions are sporadic in the sense that, each period, each agent plays against a single, randomly-sampled other agent, and the dynamics converges to the Pareto-efficient convention.

A second observation is that, as suggested by previous theoretical results (Alós-Ferrer and Weidenholzer, 2008; 2014), a large network size *relative* to the size of neighborhoods (a “large world”) facilitates convergence to efficient conventions. However, controlling for the prediction of the deterministic, noise-free dynamics suggests that the effect of the *absolute* size of the network might be channeled by the deterministic dynamics, possibly due to the relative size of the basins of attraction in the latter.

The third result is that the specification of tie-breaking rules and revision opportunities (inertia), which are often consequential for particular theoretical results, are empirically of little relevance in our set of simulations. This is reassuring, as predictions should be robust to technical details of the dynamic specification.

[7] In particular, this suggests that absorbing states or sets other than $\bar{P}$ and $\bar{R}$ are not empirically relevant in our simulations. The 881 cases where convergence to $\bar{P}$ or $\bar{R}$ did not obtain for the unperturbed dynamics were excluded from the regression analysis.

Our work is also related to a strand of literature that studies evolutionary game dynamics on general networks using simulations and approximation techniques (e.g. pair approximation, see Matsuda et al., 1992; Van Baalen, 2000). That literature has typically concentrated on the role of spatial structure for the emergence and maintenance of cooperation in social dilemmas (Allen et al., 2017; Hauert and Doebeli, 2004; Nowak and May, 1992; Ohtsuki et al., 2006; Santos et al., 2006; Zhang et al., 2016). However, a few contributions have also examined coordination games (Ohtsuki and Nowak, 2006b). For instance, Ohtsuki and Nowak (2006a) study coordination games on a circle and find that imitation of randomly-sampled agents can lead to efficient outcomes. Our work can also be interpreted as providing a link between the "approximate" results in that literature and the "exact" results obtained in the stochastic stability literature.

Of course, in spite of the large number of simulations, our study is still limited, since, for feasibility and concreteness, we have made a number of specific decisions on the set of simulations. For instance, our study has considered only networks created out of the basic $2k$-model on the circular city, following Ellison (1993) and others. It would be desirable to conduct further studies (possibly with a smaller number of simulations) with different sets of networks. Also, we have concentrated on the traditional, highly-stylized binary case. The theoretical results in, e.g., Alós-Ferrer and Weidenholzer (2008) extend to any $n \times n$ game with an efficient strategy, and it would be interesting to consider simulations for a large set of networks and games with more than two strategies. These tasks are left for future research.

References

Allen, B., Lippner, G., Chen, Y.-T., Fotouhi, B., Momeni, N., Yau, S.-T., Nowak, M.A., 2017. Evolutionary dynamics on any population structure. Nature 544 (7649), 227–230.
Alós-Ferrer, C., 2003. Finite population dynamics and mixed equilibria. Int. Game Theory Rev. 5 (3), 263–290.
Alós-Ferrer, C., 2008. Learning, bounded memory, and inertia. Econ. Lett. 101, 134–136.
Alós-Ferrer, C., Buckenmaier, J., 2017. Cournot vs. Walras: a reappraisal through simulations. J. Econ. Dyn. Control 82, 257–272.
Alós-Ferrer, C., Buckenmaier, J., Farolfi, F., 2021. Imitation, network size, and efficiency. Netw. Sci. forthcoming.
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.
Alós-Ferrer, C., Weidenholzer, S., 2006. Imitation, local interactions, and efficiency. Econ. Lett. 93, 163–168.
Alós-Ferrer, C., Weidenholzer, S., 2007. Partial bandwagon effects and local interactions. Games Econ. Behav. 61, 1–19.
Alós-Ferrer, C., Weidenholzer, S., 2008. Contagion and efficiency. J. Econ. Theory 143, 251–274.
Alós-Ferrer, C., Weidenholzer, S., 2014. Imitation and the role of information in overcoming coordination failures. Games Econ. Behav. 87, 397–411.
Barron, G., Erev, I., 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. J. Behav. Decis. Mak. 16, 215–233.
Ben-Porath, E., Dekel, E., Rustichini, A., 1993. On the relationship between mutation rates and growth rates in a changing environment. Games Econ. Behav. 5 (4), 576–603.
Blume, L., 1993. The statistical mechanics of strategic interaction. Games Econ. Behav. 5, 387–424.
Blume, L., 2003. How noise matters. Games Econ. Behav. 44, 251–271.
Blume, L., Arthur, B., Durlauf, S., Lane, D., 1997. Population games. In: The Economy as an Evolving Complex System II. Addison-Wesley, Reading, MA, pp. 425–460.
Cui, Z., 2014. More neighbors, more efficiency. J. Econ. Dyn. Control 40, 103–115.
Cui, Z., Wang, R., 2016. Collaboration in networks with randomly chosen agents. J. Econ. Behav. Organ. 129, 129–139.
Ellison, G., 1993. Learning, local interaction, and coordination. Econometrica 61, 1047–1071.
Ellison, G., 2000. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. Rev. Econ. Stud. 67, 17–45.
Erev, I., Barron, G., 2005. On adaptation, maximization, and reinforcement learning among cognitive strategies. Psychol. Rev. 112, 912–931.
Harsanyi, J., Selten, R., Cambridge, M., 1988. A General Theory of Equilibrium Selection in Games. MIT Press.
Hauert, C., Doebeli, M., 2004. Spatial structure often inhibits the evolution of cooperation in the snowdrift game. Nature 428 (6983), 643–646.
Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. Econometrica 61, 29–56.
Karlin, S., Taylor, H.M., 1975. A First Course in Stochastic Processes, 2nd ed Academic Press, San Diego.
Khan, A., 2014. Coordination under global random interaction and local imitation. Int. J. Game Theory 43 (4), 721–745.
Lee, I.H., Valentinyi, A., 2000. Noisy contagion without mutation. Rev. Econ. Stud. 67 (1), 47–56.
Matsuda, H., Ogita, N., Sasaki, A., Satō, K., 1992. Statistical mechanics of population: the lattice Lotka-Volterra model. Progr. Theor. Phys. 88 (6), 1035–1049.
Nowak, M.A., May, R.M., 1992. Evolutionary games and spatial chaos. Nature 359 (6398), 826–829.
Oechssler, J., 1997. An evolutionary interpretation of mixed-strategy equilibria. Games Econ. Behav. 21, 203–237.
Ohtsuki, H., Hauert, C., Lieberman, E., Nowak, M.A., 2006. A simple rule for the evolution of cooperation on graphs and social networks. Nature 441 (7092), 502–505.
Ohtsuki, H., Nowak, M.A., 2006a. Evolutionary games on cycles. Proc. R. Soc. B: Biol. Sci. 273 (1598), 2249–2256.
Ohtsuki, H., Nowak, M.A., 2006b. The replicator equation on graphs. J. Theor. Biol. 243 (1), 86–97.
Papke, L.E., Wooldridge, J.M., 2008. Panel data methods for fractional response variables with an application to test pass rates. J. Econom. 145 (1), 121–133.
Peski, M., 2010. Generalized risk-dominance and asymmetric dynamics. J. Econ. Theory 145 (1), 216–248.
Robson, A.J., Vega-Redondo, F., 1996. Efficient equilibrium selection in evolutionary games with random matching. J. Econ. Theory 70, 65–92.
Sandholm, W.H., 1998. Simple and clever decision rules for a model of evolution. Econ. Lett. 61, 165–170.
Santos, F.C., Pacheco, J.M., Lenaerts, T., 2006. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. Proc. Natl. Acad. Sci. 103 (9), 3490–3494.
Van Baalen, M., 2000. Pair approximations for different spatial geometries. In: Dieckmann, U., Law, R., Metz, J.A.J. (Eds.), The Geometry of Ecological Interactions: Simplifying Spatial Complexity. Cambridge University Press, Cambridge, UK, pp. 359–387.
Weidenholzer, S., 2010. Coordination games and local interactions: a survey of the game theoretic literature. Games 1 (4), 551–585.
Young, H.P., 1993. The evolution of conventions. Econometrica 61, 57–84.
Zhang, B., Li, C., Tao, Y., 2016. Evolutionary stability and the evolution of cooperation on heterogeneous graphs. Dyn. Games Appl. 6 (4), 567–579.