Economic Theory (2020) 69:497–521
https://doi.org/10.1007/s00199-019-01174-y

RESEARCH ARTICLE

The evolution of conventions under condition-dependent mistakes

Ennio Bilancini[1] · Leonardo Boncinelli[2]

Received: 4 December 2017 / Accepted: 10 January 2019 / Published online: 24 January 2019

Abstract
We study the long-run conventions emerging in a stag-hunt game when agents are myopic best responders. Our main novel assumption is that errors converge to zero at a rate that is positively related to the payoff earned in the past. To fully explore the implications of this error model, we introduce a further novelty in the way we model the interaction structure, assuming that with positive probability agents remain matched together in the next period. We find that, if interactions are sufficiently persistent over time, then the payoff-dominant convention emerges in the long run, while if interactions are quite volatile, then the maximin convention can emerge even if it is not risk-dominant. We contrast these results with those obtained under two alternative error models: uniform mistakes and payoff-dependent mistakes.

Keywords Risk-dominant · Payoff-dominant · Maximin · Stag hunt · Stochastic stability

JEL Classification C72 · C73

1 Introduction

A convention can be understood as an equilibrium of a coordination game played by the individuals belonging to a group or a society: once a particular way of doing things becomes established as a rule, it continues in force because we prefer to conform

✉ Leonardo Boncinelli
leonardo.boncinelli@unifi.it

Ennio Bilancini
ennio.bilancini@imtlucca.it

1 IMT School of Advanced Studies, Piazza S. Francesco 19, 55100 Lucca, Italy

2 Dipartimento di Scienze per l'Economia e l'Impresa, Università degli Studi di Firenze, Via delle Pandette 9, 50127 Florence, Italy

to the rule given the expectation that others are going to conform (Lewis 1969).[1] An important issue that can be studied in this setting is the tension between Pareto dominance and safety: while a convention pays a higher payoff to every individual if coordination is actually achieved, another convention is less risky since it performs better if miscoordination occurs. The stylized game that is used to capture this conflict is the *stag-hunt* game, which is often seen as a paradigmatic representation of the obstacles to achieve social cooperation (Skyrms 2004).

In this paper, we study the evolution of conventions in a population of agents who are matched in pairs to play a stag-hunt game. Each agent only occasionally revises his action, and when he does so, he chooses a myopic best reply. We derive the emerging long-run convention applying the notion of stochastic stability, which was introduced by Foster and Young (1990) and then developed by Young (1993) and Kandori et al. (1993) building on technical results in Freidlin and Wentzell (1998). Basically, when agents are allowed to make mistakes with a small probability, stochastic stability selects those equilibria that are relatively easiest to reach in terms of mistakes—also referred to as errors or mutations—starting from other equilibria. Importantly, Bergin and Lipman (1996) have highlighted the strong dependence of predictions on the error structure: if mutation rates are state-dependent, which equilibrium is selected by stochastic stability depends on exactly how these rates vary from state to state. Therefore, having an appropriate error model turns out to be crucial for stochastic stability to give reliable predictions in terms of equilibrium selection. Moreover, which convention is stochastically stable is also sensitive to the specification of the matching mechanism, as shown by, e.g., Robson and Vega-Redondo (1996).[2] Importantly, the implications of an error model may depend on the interaction structure.

The novelty of our contribution stems from considering two innovative features: (1) errors converge to zero at a rate that is positively related to the payoff earned in the past and (2) agents are randomly matched to play a stag-hunt, but their interaction does not necessarily terminate after the first round of play.

We refer to the error model indicated in (1) as *condition-dependent mistakes*, drawing from the literature on genetics and biology (Agrawal 2002). We contrast the results obtained under condition-dependent mistakes with those obtained under alternative error models, where the risk-dominant convention typically emerges: errors converging to zero at a rate that is constant over states, referred to as *uniform mistakes*, and errors converging to zero at a rate that depends on expected losses, referred to as *payoff-dependent mistakes*.

In regard to the interaction structure referred to in (2), we introduce a parameter that measures the probability that a formed pair is broken after each play of the stag-hunt. This model of interaction encompasses both the case of random pairwise encounters—when the probability that a formed pair breaks is 1—as well as the case of fixed exogenous neighborhood—when the probability that a formed pair breaks is 0. We find that under condition-dependent mistakes such parameter crucially determines which convention is selected in the long run. While our model of interaction is similar in spirit

[1] See Young (2015) for an overview of other mechanisms that can sustain conventions.

[2] Their main conclusion is that, differently from Kandori et al. (1993), the payoff-dominant convention is selected in the long run when the population is large enough.

to that in Robson and Vega-Redondo (1996), there are also substantial differences. In particular, in their model there is no role for termination probability and agents follow an imitative behavioral rule, in that they tend to adopt the strategy that led to the highest average payoff.

At a first glance, one might expect that condition-dependent mistakes always favor the selection of the payoff-dominant convention. Intuitively, one can think of a self-reinforcing mechanism at work here: a larger equilibrium payoff induces less mistakes that in turn sustain the current equilibrium. In fact, we find that this is so only if interactions are sufficiently persistent over time, i.e., long-lasting. Instead, if interactions are characterized by intermediate levels of volatility, i.e., they are somewhat liable to being terminated from period to period, then the selection of the risk-dominant convention is obtained; moreover, if interactions feature high levels of volatility, i.e., they are quite unlikely to last until next period, then the maximin convention is obtained when mistake rates are sufficiently sensitive to payoffs (and the risk-dominant convention is obtained otherwise). Notably, in the latter case a maximin convention can be stochastically stable even if it is not risk-dominant (Harsanyi and Selten 1988), that is when its basin of attraction is smaller than that of the other convention.[3] We find it also remarkable that the long-run convention emerging under condition-dependent mistakes crucially depends on the interaction structure, in particular on the extent to which agents remain matched together over time as opposed to be frequently reshuffled.

As a consequence, it turns out that when interactions are rather volatile, the inefficiency of the stochastically stable convention can be particularly severe under condition-dependent mistakes. The reason behind this is related to a form of *contagion of mistakes*. If relationships are volatile and the population size is large, many mistakes are required to move from one convention to the other. When we start from a convention and we try to reach the other convention, an agent who has changed action by mistake can be matched with someone who still takes the conventional action and hence obtains a low payoff due to miscoordination. Under condition-dependent mistakes, such an agent is in turn more likely to make a mistake, thus fostering the propagation of mistakes. Since the maximin action pays a larger payoff in case of miscoordination, under condition-dependent mistakes we have that each mistake (other than the first mistake, which cannot be the result of contagion) is less likely when we start from the maximin convention. This makes the maximin convention more resilient, allowing it to be stochastically stable even if its basin of attraction is the smallest. In this perspective, our contribution may be seen as related to the literature on contagion (see, e.g., Ellison 1997; Morris 2000; Alós-Ferrer and Weidenholzer 2008), with the difference that contagion does not take place via best responses but via mistakes.

The rest of the paper is organized as follows. Section 2 contrasts this paper with the relevant literature. Section 3 presents the social game and defines the unperturbed dynamics of behavior in the population. Section 4 introduces errors with the aim of selecting between conventions, studying how this selection depends on the error model used and the termination probability of a relationship. Section 5 discusses the

[3] Maximin plays a role also in the recent work by Sawa and Wu (2018), where a risk-dominant strategy is no longer guaranteed to be stochastically stable if it is not maximin too.

robustness of our findings, and Sect. 6 suggests an interpretation of the contribution and draws suggestive conclusions.

2 Related literature

Myopic best response is a typical choice rule in the literature on stochastic stability,[4] and independent inertia—i.e., at every time each agent has an independent, strictly positive probability not to receive the opportunity to revise his action—is a widely applied revision protocol as well.[5] Our model is more innovative for what concerns the interaction structure and, in particular, the error model. Hence we focus and organize our review of the literature on such issues, and we then conclude with a few recent contributions that have started investigating human error-making from an experimental perspective.[6]

2.1 Interactions

The stability and shape of the structure of interactions has been shown to have important consequences on group behavior (Macy and Willer 2002). The literature on the evolution of conventions in the stag-hunt game has explored a variety of interaction structures. Some models consider random pairwise encounters (Kandori et al. 1993; Kandori and Rob 1995; Young 1993). In other models, the interactions occur with exogenously given neighbors—as in Ellison (1993), where players are arranged on a circle and interact with the two immediate neighbors.[7] Typically, in these models the risk-dominant convention is selected in the long run. Another way to model agents' interactions is to consider endogenous network formation, where agents choose with whom to interact (Goyal and Vega-Redondo 2005; Jackson and Watts 2002; Staudigl and Weidenholzer 2014). In these models, the payoff-dominant convention is shown to emerge in the long run if the single interaction is sufficiently costly or the total number of interactions per agent is sufficiently constrained. An alternative to network formation is that agents can choose where to interact—not with whom—selecting a location among a number of locations available and then interacting randomly with agents choosing the same location (Oechssler 1997; Ely 2002; Bhaskar and Vega-Redondo 2004). The possibility to "vote by their feet" helps agents to coordinate on the payoff-dominant action (Efferson et al. 2016, provide experimental evidence of this). Notably, if restricted mobility is imposed in these models with locations, the

[4] Other possible choice rules encompass—among others—imitation, reinforcement learning and better response (see Young 2001, Section 3.1, for a discussion of different varieties of learning behavior).

[5] The role of the revision protocol on the speed of convergence in stochastic evolutionary models has been investigated by Norman (2009). We discuss in Sect. 5 the robustness of our results to different revision protocols.

[6] For a thorough review of evolutionary models and coordination games, the interested reader can refer to the recent survey in Newton (2018b).

[7] See Peski (2010) for a general framework to study local interaction with an exogenous interaction structure, and Newton (2018a) for a recent extension. See also Weidenholzer (2010) for a recent survey on local interaction models focusing on social coordination.

coexistence of conventions can be obtained, with one convention being established in one place and the other convention in another place (Anwar 2002; Shi 2013a; Blume and Temzelides 2003). Bilancini and Boncinelli (2018) obtain coexistence of conventions also with network formation and constrained interaction, in a model where agents have different types, types are only locally observable, and type mismatches are costly.[8]

2.2 Error model

Some decision errors, or mistakes, have been shown to be inevitable in any decision-making process (Haselton and Buss 2000; Nesse 2005). Different error models have been explored by the literature on stochastic stability. Most of the research in economics has focused on mistake rates that are state independent, i.e., uniform across states. The plausibility of such an error model can be argued on the basis of the result in Van Damme and Weibull (2002), who consider a model where agents with some effort can control the probability to make a mistake, and obtain that mistake rates approach zero at the same rate across states when control costs become vanishingly small. In some circumstances, however, it may be reasonable to expect that mistakes generating smaller payoff losses occur with a higher rate. This is what happens, for instance, in the logit response dynamics introduced by Blume (1993), which is related to the notion of quantal response equilibrium (McKelvey and Palfrey 1995) and, to some extent, to the notion of proper equilibrium (Myerson 1978).

In biology, a different form of state-dependent mutations has been considered, as a function of the condition of the organism: lower condition organisms have a higher mutation rate (Agrawal 2002).[9] We believe that such a class of error models can be relevant for economics as well, at least for a couple of reasons. The first one applies whenever higher payoffs can be related to better physiological conditions: Individuals that are better fed and healthier are less likely to make mistakes. This justification seems particularly relevant for the evolutionary analysis of the stag-hunt game, which aims at representing typical interactions in societies over a very long time span and hence may also consider the remote past. The second reason hinges on the interpretation of mistakes as experimentation: Agents that earn higher payoffs are less inclined to change action and see if, by chance, coordination can be obtained elsewhere.[10] On a more general ground, the importance of realized payoffs in contrast with expected payoffs as a driver of behavior relates our contribution also to the literature on learning focusing on reinforcement of successful behavior (see Young 2009 and Pradelski and Young 2012 on trial and error, Huttegger and Skyrms 2012 on probe and just—both based on a similar idea by Thuijsman et al. 1995 on the behavior of foraging bees—and Nax et al. 2016 for recent experimental evidence). What distinguishes our contribution

[8] Coexistence of conventions can emerge also under global random interaction when agents are heterogeneous, as in the case they belong to distinct cultural groups (see, e.g., Carvalho 2017).

[9] See also the extensions in Shaw and Baer (2011) and Cotton (2009). See Sharp and Agrawal (2012) for evidence on bacteria and Agrawal and Wang (2008) for evidence on insects.

[10] This justification is possibly related to the existence of an aspiration level, like in Binmore and Samuelson (1997), where the realized payoff is given by the expected payoff plus a random shock, and the individual changes action only if the realized payoff falls below the aspiration level.

from this literature is mainly that experienced payoffs only affect the likelihood of mistakes, and in particular the level of experienced payoff, not its change (in other words, the benchmark in our model is non-adaptive, while it is adaptive in the models belonging to the literature under consideration).

The idea that a lower past payoff makes mistakes more likely is also partly related with the notion of intentional mistakes introduced by Naidu et al. (2010),[11] who study contract games with the assumption that only actions that can lead to higher payoffs can be chosen by mistake, and more generally with coalitional stochastic stability (Newton 2012), where mistakes that can be rationalized as profitable deviations for a coalition of individuals are assumed to be more likely than other mistakes. We stress that the latter kind of experimentation involves a certain degree of group-level rationality, while condition-dependent mistakes are purely based on individual past satisfaction.

The general formulation of errors in Blume (2003) and Maruta (2002) encompasses uniform mistakes and payoff-dependent mistakes, but not condition-dependent mistakes. Those contributions investigate under which conditions on the error model the risk-dominant convention is selected in the stag-hunt game. In doing this, error probabilities are assumed to be general functions of current expected payoffs, but they are not allowed to depend on past payoffs, as is the case in condition-dependent mistakes. Moreover, while Blume (2003) and Maruta (2002) conduct their analysis under the assumption of random matching only, we also explore the role of the termination probability. In the context of matching problems, Newton and Sawa (2015) consider an extremely large class of error models (encompassing all ours) and identify a necessary condition for a matching to be stochastically stable, namely being most robust to one-shot deviation.[12]

2.3 Experimental evidence

Experiments have recently been conducted in order to establish which error model better fits actual individual behavior. Mäs and Nax (2016) analyze data from a laboratory experiment on coordination in networks and find that the vast majority of decisions constitute myopic best responses, and deviations are less frequent when implying larger payoff losses. In addition, deviation rates vary with patterns of realized payoff, with error rates going up when payoff goes down—indicating the potential presence of condition-dependent mistakes. Lim and Neary (2016) have conducted a laboratory experiment where subjects play a language game (see Neary 2012), and they also find that behavior is highly consistent with the myopic best response rule. In particular, deviations from the learning rule occur with a likelihood that depends negatively on the payoff of the myopic best response, but not on the deviation payoff, which, again, suggests that condition-dependent mistakes may be present. They also find autocorre-

11 Recently, Hwang et al. (2016) make use of intentional mistakes to show that unequal conventions may persist over long periods of time despite being inefficient and not supported by formal institutions.

12 There is a growing literature that investigates the impact of error models on stochastic stability results applied in matching and assignment problems, showing that different error models lead to different predictions (see, e.g., Klaus and Newton 2016; Nax and Pradelski 2015; Boncinelli and Pin 2018).

Fig. 1 A stag-hunt game

	A	B
A	a	c
B	d	b

lation in deviations from the best response action. They conclude that experimental data are more in line with intentional and payoff-dependent mistakes, rather than uniform mistakes.[13] Hwang et al. (2018) study two populations of players who are matched to play a two-player coordination game with zero payoffs off-diagonal[14]—which makes it hard to detect condition-dependent mistakes, since an equilibrium where an agent obtains a higher payoff is associated with a more severe expected loss for such an agent in case of deviation. Their experimental data are consistent with high levels of best response play, as well as payoff-dependent and intentional mistakes. No direct test of condition-dependent mistakes has been conducted so far.

3 Model

We consider a population of n agents, indexed according to the set of integers $N = \{1, 2, \ldots, n\}$, with n even. Time is discrete and denoted with $t = 0, 1, 2, \ldots$. Agents are matched in pairs to interact repeatedly. Once a match between two agents is formed, there is a probability $\tau > 0$ that it gets broken after each interaction. We refer to τ as the *termination probability*. Agents without a partner are randomly matched among themselves.

The one-shot interaction between two matched agents takes the form of a 2-player stag-hunt game. Each agent plays only with the agent with whom he is currently matched. The table in Fig. 1 describes the payoffs associated with the game. (Payoffs are given only for the row player, since the game is symmetric.) We use $\pi(x, x')$ to indicate the payoff earned by an agent who plays x against an agent who plays x'.

We assume that $b > a$, that is, B is the *payoff-dominant* action. We also assume that $c > d$, so that A is the *maximin* action. We do not make an assumption about which action is *risk-dominant*: If $a + c > b + d$, then A is the risk-dominant action, while B is the risk-dominant action if $b + d > a + c$.[15] We assume that either A or B are risk-dominant, i.e., $a + c \neq b + d$. We further assume that $a \geq c$, and we note that $b > d$ is implied by previous assumptions.

[13] Frey et al. (2012) conduct an experiment to study the emergence of conventions in a networked coordination game and find evidence that the behavior dynamics following an initial deviation leads agents' choices away from the risk-dominant equilibrium more often than it does from the payoff-dominant equilibrium.

[14] In a similar setting, see also the analysis and approximation results by Hwang and Newton (2017).

[15] The game, in case action A is not risk-dominant, is sometimes called an assurance game.

The population state at time t is described by $s^t = (\alpha^t, \mu^t)$, with $\alpha^t = (\alpha^t(1), \ldots, \alpha^t(n))$ and $\mu^t = (\mu^t(1), \ldots, \mu^t(n))$, where $\alpha^t(i) \in \{A, B\}$ is the action of agent i at time t, and $\mu^t(i) \in N \backslash \{i\}$ is the partner interacting with i at time t, with $\mu^t(i) = j$ implying $\mu^t(j) = i$ and $\mu^t(i) \neq \mu^t(j)$ for all $i, j \in N, i \neq j$.

The action revision protocol is as follows. In each time period, every agent has an independent probability $\gamma \in (0, 1)$ to be given the opportunity to revise his action. An agent who is given a revision opportunity at time t takes a choice following a myopic best response rule against s^{t-1}, randomizing in case of indifference. This means that actions are taken before an agent knows whether his current match terminates (see Sect. 5 for a discussion of this point).

Preliminarily to give an expression for the expected payoff, it is useful to observe that, given the termination probability τ, every agent is matched in the next period with a partner that is different from the current one with a probability that depends on n and that is equal to[16]

$$\delta = \tau - \tau \sum_{k=0}^{\frac{n}{2}-1} \frac{(\frac{n}{2}-1)!}{k!(\frac{n}{2}-1-k)!} (1-\tau)^{\frac{n}{2}-1-k} \tau^k \frac{1}{2k+1}, \tag{1}$$

where δ is expressed as the probability that the current match terminates, i.e., τ, minus the probability to get rematched with the same partner. To better understand the expression used for this probability, we note that the binomial coefficient measures the number of combinations of k other matches terminating out of $n-1$ overall matches, while $(1-\tau)^{\frac{n}{2}-1-k}\tau^k$ is the probability of any such combination to occur, and finally $\frac{1}{2k+1}$ is the probability to be rematched exactly with the same partner when $2k+1$ other agents are in the pool of unmatched agents.

We observe that in the event that agent i is matched with a new partner, which happens with probability δ, any other agent j (that is neither i nor his previous partner) has the same probability to become the next partner. This allows us to write the payoff that agent i expects (with adaptive expectations implied by myopic best response) to earn at time t if he chooses action $x \in \{A, B\}$ as follows:

$$\Pi_i^{s^t}(x) = \delta \left(\sum_{\substack{j=1,\ j\neq i, \\ j\neq\mu^{t-1}(i)}}^{n} \frac{\pi(x, \alpha^{t-1}(j))}{n-2} \right) + (1-\delta)\pi\left(x, \alpha^{t-1}\left(\mu^{t-1}(i)\right)\right). \tag{2}$$

The first term in (2) is the expected payoff of interacting with a new mate, weighted by the probability δ that this happens. We note that the $n-2$ agents different from both $\mu^{t-1}(i)$ and i have all the same probability to terminate the match with their partner at $t-1$ and then be rematched with i. The second term in (2) is the expected payoff

16 We note that $0! = 1$ and $0^0 = 1$. So, when $\tau = 1$ we have that $\delta = 1 - 1/(n-1)$ which means that, even if partnerships terminate with certainty, there is always a positive probability to interact with the previous partner through re-matching.

of i interacting with $\mu^{t-1}(i)$, weighted by the probability $(1-\delta)$ that i and $\mu^{t-1}(i)$ keeps on interacting at next time.

We say that action $x \in \{A, B\}$ is best response for agent i at state s^t if $\Pi_i^{s^t}(x) \geq \Pi_i^{s^t}(y)$, with $y \in \{A, B\}$, $y \neq x$.

We call *convention A* the set of states $s = (\alpha, \mu)$ such that $\alpha(i) = A$ for every $i \in N$. Analogously, we call *convention B* the set of states $s = (\alpha, \mu)$ such that $\alpha(i) = B$ for every $i \in N$. We refer to convention A as the maximin convention, since A is the maximin action. Analogously, we refer to convention B as the payoff-dominant convention, since B is the payoff-dominant action. Finally, by risk-dominant convention we refer to either convention A or convention B, depending on whether A or B is the risk-dominant action.

The dynamic system under consideration is a Markov chain (see Young 2001, for an overview of Markov chain theory). The definition of recurrent class is worth remembering. A set of states $\mathcal{C}$ is a recurrent class if: (i) every pair of states in $\mathcal{C}$ communicate with each other—meaning that there is a positive probability to move from one state to the other in a finite number of steps—and (ii) no state in $\mathcal{C}$ communicates with a state not in $\mathcal{C}$—so that the probability of leaving $\mathcal{C}$ is zero.

The following lemma establishes that the recurrent classes of this Markov chain coincide with the two conventions:

Lemma 1 *Convention A and convention B are the only two recurrent classes for all* $\tau \in (0, 1]$.

It may be worth to note that from any state s in a convention, there is always a positive probability to get to any other state s' in the same convention: since $\tau > 0$, there is a positive probability that all matches in s terminate and that agents are rematched according to s'.

4 Stochastically stable conventions

In the spirit of Young (1993) and Kandori et al. (1993), we add agents' mistakes to the model described in the previous section. In particular, the probability that an agent who has received a revision opportunity in state $s = (\alpha, \mu)$ makes a mistake, i.e., selects an action that is not a best response, approaches zero at the same rate as $\epsilon^{r(i,s)}$ when ϵ goes to zero. We refer to $r(i, s) > 0$ as the *resistance to mistake* of agent i at s. We remark that this formulation allows for mistake rates that depend on the agent who takes the decision, and the state at which the decision is taken. Due to the introduction of mistakes, we obtain a Markov chain that is irreducible, i.e., has a unique recurrent class. This in turn implies that there exists a unique invariant distribution that describes the fraction of time spent on each state in the long run, irrespectively of the initial state.

As ϵ goes to zero, mistakes become rarer and rarer, and the invariant distribution converges to the so-called stochastically stable distribution. We say that a convention is *stochastically stable* if its states have positive probability in the stochastically stable distribution.[17] We rely on the techniques developed by Young (1993, 2001), which allow to characterize stochastically stable conventions in terms of minimum stochastic potential. In our model, the *stochastic potential* of convention A is the minimum total number of mistakes over paths of states starting from convention B and reaching convention A, with each mistake weighted by $r(i, s)$ if made by agent i at state s. The stochastic potential of convention B is analogously defined. A convention is stochastically stable if its stochastic potential is smaller than or equal to the stochastic potential of the other convention.

We proceed to determine which conventions are stochastically stable by considering three different error models in turn. In doing so, the following thresholds on the termination probability turn out to be important for our results:

$$\tilde{\tau}_m := \min\left\{\frac{a-d}{a-d+b-c}, \frac{b-c}{a-d+b-c}\right\};$$
$$\tilde{\tau}_M := \max\left\{\frac{a-d}{a-d+b-c}, \frac{b-c}{a-d+b-c}\right\}.$$

Since $a + c \neq b + d$, it follows that $0 < \tilde{\tau}_m < \tilde{\tau}_M < 1$. We note that these thresholds measure the sizes of the two basins of attractions in the standard case of random matching (corresponding to $\tau = 1$). They turn out to be relevant in our analysis because they tell us when a single mistake is enough to trigger the transition from one convention to the other. The following argument aims at illustrating the claim. Consider a state in which all agents but one are playing B. The partner of the agent choosing A has the highest incentive to choose A. For him, A is a best response if $(1-\delta)a + \delta c \geq (1-\delta)d + \delta b$. This rearranges to $\delta \leq \frac{a-d}{a-d+b-c}$, which is implied by $\tau < \frac{a-d}{a-d+b-c}$ if n is large enough. Any following switch from B to A by another agent turns out to be easier. Finally, we note that proceeding in the same way we can conclude that a single mistake is enough to trigger the passage from convention A to convention B when $\tau < \frac{b-c}{a-d+b-c}$ (and the population is large).

The first error model that we consider is one where all mistakes are equal likely; in the sense that all mistakes have the same resistance.

Definition 1 (*Uniform mistakes*) In the uniform mistakes model,

$$r(i, s) = h,$$

where h is a positive number.

The following proposition states which conventions are stochastically stable depending on the termination probability.[18]

[17] It is known that only states belonging to recurrent classes of the model without mistakes can receive positive probability in the stochastically stable distribution. Moreover, if one state receives positive probability then all states in the same recurrent class receive positive probability as well.

[18] Since the precise expression for the threshold value of n is cumbersome, we opted to omit it from the proposition, aiming at an easier statement which points to the main result.

Proposition 1 *Under uniform mistakes, (a) if* $\tau \in (0, \tilde{\tau}_m]$*, then both the maximin convention and the payoff-dominant convention are stochastically stable, (b) if* $\tau \in (\tilde{\tau}_m, 1]$*, then the risk-dominant convention is the unique stochastically stable convention provided that n is large enough.*

Proposition 1 provides a generalization of a standard result (see, e.g., Kandori et al. 1993) by taking into account the termination probability. More precisely, if the termination probability is sufficiently large, then the unique stochastically stable convention is the risk-dominant convention, as in the case of pure random matching analyzed in the literature (which corresponds to $\tau = 1$). However, and quite intuitively, if the termination probability is low enough, then both conventions turn out to be stochastically stable. Indeed, by best response an agent would conform his action to the current partner's action even if no one else in the population takes that action, and this is so because such an interaction is likely to last long over time. Hence, by means of random matching of agents who remain unpaired, a single mistake can be enough for an action to spread in the population, so that a convention is replaced by the other convention, irrespectively of the initial convention.

We now turn to analyze what happens if mistakes are payoff-dependent, i.e., $r(i, s)$ is positively related to the difference between the expected payoff of playing the best response action and the expected payoff of making a mistake.

Definition 2 (*Payoff-dependent mistakes*) In the payoff-dependent mistakes model,

$$r(i, s) = f\left(\Pi_i^s(x) - \Pi_i^s(y)\right),$$

where f is a strictly increasing function with positive values, x is a best response for agent i at state s, and y is the other action.

The following sharp result on stochastically stable conventions is obtained if mistakes are payoff-dependent.

Proposition 2 *Under payoff-dependent mistakes, the risk-dominant convention is the unique stochastically stable convention for all* $\tau \in (0, 1]$.

When the resistance to mistakes is increasing in the expected payoff loss, then stochastic stability selects the convention that is risk-dominant. Proposition 2 states that this result, which is well known in the literature for the case of uniform matching (see Van Damme and Weibull 2002; Blume 2003), holds even if the termination probability is lower than 1. With respect to the error model where the resistances to mistakes are uniform, here there is an additional effect to consider in order to assess which convention is stochastically stable: Mistakes are particularly unlikely when starting from the risk-dominant convention because they lead to larger expected losses. This further effect allows to conclude that the risk-dominant convention is the unique stochastically stable convention even when a single mistake is sufficient to move from one convention to the other, that is even when the termination probability is low.

Finally, we turn to consider an error model where the resistance to mistakes is positively related to the payoff earned in the immediately preceding period.

Definition 3 (*Condition-dependent mistakes*) In the condition-dependent mistakes model,

$$r(i,s) = g(\pi(\alpha(i), \alpha(\mu(i)))),$$

where g is a strictly increasing function with positive values.

It turns out that, under condition-dependent mistakes, it is possible that neither the payoff-dominant nor the risk-dominant convention are selected in the long run. In particular, it may happen that when B is not only payoff-dominant but also risk-dominant, the maximin convention A is selected instead.

Preliminarily, it is useful to define the following:

$$\beta(\tau) := \frac{\tau(a-d+b-c)-(a-d)}{\tau(a-d+b-c)-(b-c)},$$

which turns out to be equal to the ratio between the stochastic potential of convention A and the stochastic potential of convention B under uniform mistakes (see the proofs of Propositions 1 and 3).

Proposition 3 provides a characterization of stochastic stability under condition-dependent mistakes.[19]

Proposition 3 *Under condition-dependent mistakes, (a) if* $\tau \in (0, \tilde{\tau}_m]$*, then the payoff-dominant convention is the unique stochastically stable convention, (b) if* $\tau \in (\tilde{\tau}_m, \tilde{\tau}_M]$ *then the risk-dominant convention is the unique stochastically stable convention provided that* n *is large enough, (c) if* $\tau \in (\tilde{\tau}_M, 1]$ *then the risk-dominant convention is the unique stochastically stable convention for* $g(c)/g(d) < \beta(\tau)$*, while the maximin convention is the unique stochastically stable convention for* $g(c)/g(d) > \beta(\tau)$*, in both cases provided that* n *is large enough.*

Figure 2 allows to visualize the results of Proposition 3 for n sufficiently large.

The intuition behind Proposition 3 is as follows. If the termination probability τ is sufficiently small, then a single mistake is sufficient to move both from convention A to convention B, and from convention B to convention A. A mistake in convention A is weighted by $g(a)$, while a mistake in convention B is weighted by $g(b)$, as a and b are the individual payoffs earned in the two conventions, respectively. Since $b > a$ the minimum stochastic potential is that of convention B, which is hence the stochastically stable convention. Here we stress that c and d play no role at all.

If the termination probability τ is sufficiently large, then more mistakes are required to move from one convention to the other. In particular, the number of mistakes grows

[19] We opted not to deal with the case $g(c)/g(d) = \beta(\tau)$, which has measure zero in the parameter space and would bring in the analysis annoying complications. Indeed, when $g(c)/g(d) = \beta(\tau)$, determining which convention is stochastically stable is complicated by the working of the ceiling function, so that either the risk-dominant convention, the maximin convention or both can be stochastically stable, depending on specific values of n. Also, we opted not to provide a precise expression for the threshold value of n, for the same reason as in Footnote 18. We note that, in general, such threshold here is different than the analogous threshold in Proposition 1.

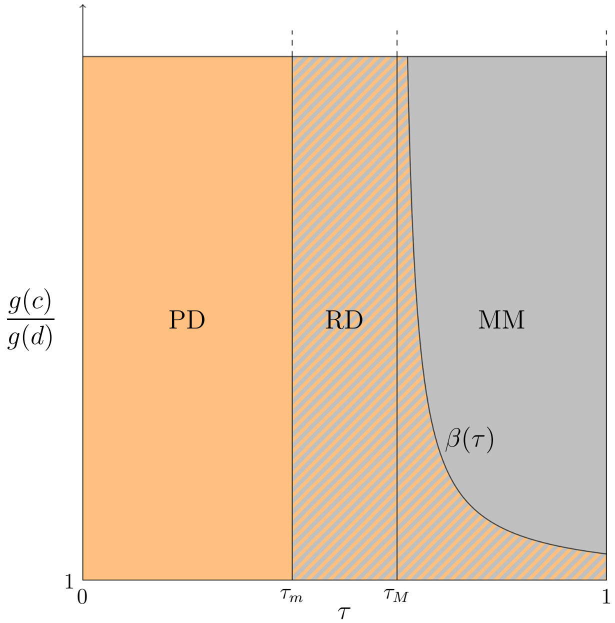

Fig. 2 The orange color refers to action B, which is payoff-dominant, and the gray color to action A, which is maximin. The orange-gray stripes pattern refers to either action B or action A, depending on which action is risk-dominant (color figure online)

unboundedly in the population size n. Moreover, any mistake other than the first is weighted by $g(d)$ when moving from convention B to convention A, and by $g(c)$ when moving from convention A to convention B,[20] since the agents who are most likely to make a mistake are those who end up matched with someone who has already changed action. When n is sufficiently large, what matters to establish which convention is stochastically stable is $g(c)/g(d)$ compared with $\beta(\tau)$, which gives the relative number of mistakes required to move from one convention to the other and which tends to $(b-c)/(a-d)$ when τ approaches 1. We stress that, if $g(c)/g(d)$ is large enough, then the stochastically stable convention corresponds to the maximin convention.

We observe that, under condition-dependent mistakes, all mistakes other than the first one depend on the payoff of a miscoordinating agent. Indeed, when moving from an established convention to a novel convention, any agent who is still choosing the established action can be matched with an agent choosing the novel action in the next period, obtaining a payoff that is quite low due to miscoordination and, hence, being more likely to mistakenly choose the novel action in the following period. This form of contagion of mistakes is allowed by the specific way in which the interaction structure evolves over time in our model, which is the result of combining a positive termination probability and the random matching protocol for agents who have remained unpaired.

[20] We remind that we have assumed $c \geq a$. In Sect. 5, we comment on the case $c < a$ when discussing the payoff structure.

Table 1 Stochastically stable conventions

Mistakes	Termination probability		
	$0 < \tau \leq \tilde{\tau}_m$	$\tilde{\tau}_m < \tau \leq \tilde{\tau}_M$	$\tilde{\tau}_M < \tau \leq 1$
Uniform	MM, PD	RD	RD
Payoff-dependent	RD	RD	RD
Condition-dependent	PD	RD	RD, if $g(c)/g(d) < \beta(\tau)$ MM, if $g(c)/g(d) > \beta(\tau)$

Table 1 summarizes the results of this section, facilitating the comparison of the stochastically stable conventions in the three different error models that we have considered. We stress how the termination probability affects this comparison: relatively to the cases of uniform mistakes and payoff-dependent mistakes, condition-dependent mistakes favor the emergence of the efficient (i.e., payoff-dominant) convention when relationships are quite persistent over time, while they favor the emergence of the inefficient (i.e., maximin) convention when relationships are quite volatile, even if such convention is not risk-dominant (and, hence, the inefficiency can be particularly severe).

5 Discussion

In this section, we discuss the robustness of our findings along a few dimensions.

5.1 Revision protocol

We are aware that the details of the revision protocol may have important consequences for the selection results obtained by stochastic stability.[21] In this paper, we have considered the case of independent inertia, where each agent has an independent, strictly positive probability of not being able to switch action at each period. Some procedures in the proofs rely on the presence of inertia, and we remark that the stochastic potentials under payoff-dependent mistakes and condition-dependent mistakes would change if an agent who switched action by mistake is not allowed to maintain the same action by inertia in the following period. The importance of inertia is particularly evident under condition-dependent mistakes, because the contagion of mistakes triggered by an initial mistake requires that miscoordinated interactions occur over time, so that in turn mistakes must occur over time, with agents who already changed action maintaining by inertia the same action until the threshold fraction of the population is reached. We stress, however, that our results remain true with another prominent revision protocol

[21] Alós-Ferrer and Netzer (2010) show how the result that stochastic stability under the logit dynamics of Blume (1993, 1997) selects potential maximizers in exact potential games depends crucially on the assumption of asynchronous learning, i.e., one and only one agent can revise strategy at each time. A recent strengthening of the notion of stochastic stability which is robust to the specification of revision opportunities and tie-breaking assumptions can be found in Alós-Ferrer and Netzer (2015).

considered in the literature, i.e., asynchronous learning, where each period one and only one agent is randomly drawn to revise his action. Basically, our proofs can be understood as the construction of sequences of agents who revise their choice, one per period of time, thanks to independent inertia. It is immediate to observe that the same procedure can be followed with asynchronous learning, where by assumption only one agent per period of time is allowed to change action.

5.2 Matching protocol

The interaction structure assumed in our model includes random matching as a special case when $\tau = 1$, and considers relationships which become more persistent over time as τ decreases. While random matching has a prominent role in the literature, many other interaction structures have been considered as well, which are not included as special cases of the interaction structure of this paper. We here provide a brief discussion on the role which condition-dependent mistakes could play under two alternative matching protocols which are common in the literature.

When the interaction structure is a round robin matching (i.e., each agent meets every other agent in turn), realized payoffs coincide with expected payoffs for every agent in each period. This prevents the occurrence of series of unlucky events which yield very low payoffs to some agents, who would be the ones making mistakes in a minimum resistance path exiting from a convention under condition-dependent mistakes. Since unlucky events yield a relatively higher payoff for the maximin action (which would make the maximin convention stabler relatively to the other convention under condition-dependent mistakes), the elimination of uncertainty due to a round robin matching sterilizes a possible reason why the maximin convention can be selected as stochastically stable. We may be led to think that risk dominance, which captures the advantage in terms of the relative number of mistakes to move from one convention to the other, remains the only determinant of stochastic stability, as it happens in the case of uniform mistakes when the risk-dominant convention is always stochastically stable. Under condition-dependent mistakes, however, there is an additional determinant to be considered: Since error rates negatively depend on absolute payoffs, mistakes are less likely when starting from the payoff-dominant convention than they are when starting from the other convention. The result is that the payoff-dominant convention may be stochastically stable even if it is not risk-dominant, provided that its basin of attraction is not too small. This is in sharp contrast with what we obtain with condition-dependent mistakes and random matching ($\tau = 1$), while it resembles to some extent what happens when τ is low.

When the interaction structure is as in Ellison (1993) with two neighbors (i.e., agents are arranged in a circle, and each agent interacts with his two closest neighbors), a single mistake is enough to reach the risk-dominant convention starting from the other one, while all agents but one have to make a mistake starting from the risk-dominant convention in order to reach the other convention. Despite this imbalance in the number of mistakes, stochastic stability may select the convention that is not risk-dominant, if it is payoff-dominant, because the single mistake that is required to leave it is the most unlikely under condition-dependent mistakes. However, fixed everything else, if the population size is above some threshold level then the risk-dominant convention becomes stochastically stable. The contagion that risk-dominance triggers in a local interaction model is the driver of this result, showing that the selection of the risk-dominant convention, which is known to be quite robust in many respects (Lee et al. 2003; Weidenholzer 2012), is also essentially robust to mistakes that are condition-dependent.

5.3 Myopia

In our model, agents are myopic in the sense that they best reply to the current state, applying simple adaptive expectations. However, the same agents are assumed to be sophisticated enough to take into account the termination probability τ when forming beliefs on the opponent's play. One may wonder what would happen if agents were more myopic, failing to take τ into consideration.

A natural approach here is to assume that agents misperceive the termination probability, forming beliefs by considering that the termination probability is equal to 1 when τ is close to 1, and equal to 0 when τ is close to 0. In other words, agents myopically believe that either interactions are totally random or that they last forever. We stress that our analysis applies quite naturally to these cases as well. In particular, when agents behave as if $\tau = 0$ the same results hold as for the case $0 < \tau \leq \tilde{\tau}_m$ in Table 1; if instead agents behave as if $\tau = 1$, the same results hold as for the case $\tilde{\tau}_m < \tau \leq 1$ in Table 1 with $\beta(1)$. The reason is that any sequence of matchings has still positive probability to occur, since the actual τ is strictly between 0 and 1, while the individual behavior is driven by the misperceived τ, so that agents behave conforming to the partner when they think $\tau = 0$ (allowing a single mistake to move the system from any convention to the other one), and best-replying to the frequency of actions in the population when they think $\tau = 1$ (so that many mistakes have to accumulate to allow a transition between conventions, the same as in the case of random matching).

5.4 Number of conventions

The substance of our results remains unaffected when there are more than two equilibria in the underlying social game. For the sake of the argument, consider the case of symmetric games where there is one payoff-dominant convention. On one extreme, when the termination probability is sufficiently low, the payoff-dominant convention emerges in the long-run under condition-dependent mistakes, for essentially the same reason as in the model of this paper: A single mutation is sufficient to move from one convention to any another, and since the resistance to mistakes is larger the higher the payoff earned, it turns out that the payoff-dominant convention is the most resilient in terms of mistakes, and hence, it is stochastically stable. On the other extreme, when the termination probability is sufficiently high, many mistakes are required to move from one convention to another, and each mistake but the first one must be weighted by the resistance to mistakes of miscoordinating agents—since they interact with someone who already changed action by mistake. Therefore, and analogously to what happens with two actions only, the safest convention tends to be more resilient to mistakes, since it yields larger payoffs in case of miscoordination, even if the precise condition to identify which convention is stochastically stable would require some dedicated effort.

5.5 Long-term versus short-term yields of actions

We believe that most interactions with an underlying stag-hunt structure that are economically relevant have, at least to some extent, the nature of investment. This means that, when an action is taken, there is not certainty that the relationship will be maintained in the future until benefits are paid. Consistent with this opinion, we have considered actions that cannot be conditioned to the event that the current match survives or terminates. Admittedly, there are social interactions where actions yield short-term consequences. In such situations, it is natural to assume that actions are taken after the formation of matches. With this assumption, a single mistake is always sufficient to move from one convention to the other, as long as the probability that a pair remains in place in the next period is positive, which is always the case (even when $\tau = 1$) because of the probability of being rematched with the same partner after the match has terminated. Indeed, with positive probability a match survives after an agent has made a mistake, and this leads the other agent of the pair, who knows that the match has not terminated, to change action by best response. By means of subsequent reshuffles of pairs, the action originally chosen by mistake can spread to the whole population. Hence the same results hold as in the case where $\tau \leq \tilde{\tau}_m$ for the baseline model.

5.6 Group interactions

In line with most of the literature, we have focused on pairwise interactions. Reasonably, interactions of the stag-hunt type may involve more than two agents. A relevant question to ask is then to what extent our results carry on to multiple-agent interactions. We argue that the quality of our results is easily preserved as long as not all interactions take place among $k > 2$ individuals, but there is a positive probability of interactions taking place among just two, just three, etc. On one extreme, consider the case where the termination probability is sufficiently low. A group of agents, once formed, is quite likely to last for a long time and hence what matters for stochastic stability is the relative within-group stability to mistakes of the two conventions. We observe that, in a coordination game, the pair is the group formation that is more likely to switch from one convention to the other. Hence, analogously to what happens in the model of this paper, a single mistake is enough to overturn convention in a group, and then in the whole population by means of random matching, implying that the payoff-dominant convention is stochastically stable under condition-dependent mistakes. On the other extreme, consider the case where the termination probability is sufficiently high. If the payoff loss when one group member deviates from a convention is largest in case of a pair, then the pair is again the group formation that is relevant for stochastic stability, when we construct paths from one convention to the other. As for the model

in this paper, all mistakes other than the first one receive a larger weight when we start from the maximin convention, which hence is likely to be stochastically stable.
Condition dependence on past payoffs. We have assumed that the resistance to mistakes of an agent at time t depends only on the payoff he earned at time $t-1$. Another reasonable possibility is to consider a weighted sum of a sequence of past payoffs. It is easy to see that our results are maintained in such a case. Indeed, if the length of the sequence determining the condition is equal to k, it is enough to exploit inertia, considering that with positive probability no agent receives a revision opportunity for k periods.[22]

5.7 Payoff structure

Lastly, we remark that the assumption $a \geq c$ has been made for simplicity only. In case $a < c$, the only difference that arises concerns the analysis under condition-dependent mistakes: When we move from the maximin convention to the payoff-dominant convention, the path of minimum overall resistance is now made of mistakes by agents who still coordinate on the risk-dominant action, since coordination pays less than miscoordination. Therefore, all mistakes in the path are weighted by $g(a)$. Since $d > a$ and hence $g(d) > g(a)$, the maximin convention maintains an advantage in terms of robustness to mistakes over the payoff-dominant (irrespectively of which convention is risk-dominant), and this allows us to conclude that our results remain unchanged, with the only adjustment that $g(c)/g(d)$ is replaced by $g(a)/g(d)$ in the analysis.

6 Conclusions

In this paper, we have carried out an extensive analysis on the role of condition-dependent mistakes in the evolution of conventions in stag-hunt interactions. We have found that when interactions are quite persistent over time, i.e., the same two individuals are likely to interact also in the future, then condition-dependent mistakes select the payoff-dominant convention in the long run. However, as interactions become less persistent, safety motives kick in and condition-dependent mistakes tend to select the risk-dominant convention, as done by uniform mistakes and payoff-dependent mistakes. Furthermore, and notably, when interactions are extremely volatile the safety motives become even stronger under condition-dependent mistakes, leading to the selection of the maximin convention even if it is not risk-dominant, provided that it is sufficiently safer than the payoff-dominant convention, or that mistake probabilities are sufficiently sensitive to payoffs earned.

These results suggest a potential reason why a trait embedding a model of condition-dependent mistakes might itself have experienced evolutionary success. Considering human history in a very long-run perspective, it seems fair to say that interactions

[22] If asynchronous learning is assumed as revision protocol, we might simulate inertia by giving revision opportunities to agents who are already best responding and hence would keep choosing the same action by best response.

among human beings have been for most of the time extremely persistent—at least back in the period when human societies were typically made of small numbers of hunters/gatherers (roughly before the advent of agriculture about ten thousand years ago). So, condition-dependent mistakes might have spread over other mistake behaviors exactly because those groups following condition-dependent mistakes were experiencing average higher condition.

The speculation above, whose soundness should be explored by specific research, also suggests—by the same token—that in modern societies condition-dependent mistakes might be quite detrimental. To the extent that interactions have become more and more volatile, if individuals still behave according to condition-dependent mistakes, then individuals might have been coordinating more and more on the safe convention, even when the payoff-dominant convention is also risk-dominant.

Appendix

Proof of Lemma 1

Proof We first show that conventions A and B are recurrent classes. Consider two states, $s = (\alpha, \mu)$ and $s' = (\alpha', \mu')$, that belong to convention A. Starting from s, we note that with positive probability all matches terminate (probability $\tau^{\frac{n}{2}}$) and new pairs are formed exactly as described by μ' (probability $\prod_{k=0}^{n-1} \frac{1}{n-1-2k}$). Since $\alpha = \alpha'$, and any revising agent will maintain action A (so that it is irrelevant which agents actually receive a revision opportunity), we can conclude that s' can be reached from s with positive probability. If s belongs to convention A and s' does not, then it means that $\alpha'(i) = B$ for some agent i. Starting from s, we simply observe that an agent who receives a revision opportunity will never change from A to B, and hence, state s' cannot be reached with positive probability from s. Altogether, convention A is shown to be a recurrent class. An analogous argument shows that convention B is also a recurrent class.

We now show that no other recurrent class exists. Consider a state s that belongs neither to convention A nor to convention B. We proceed to construct a path of states starting from s and ending in either convention A or convention B, with each step in the path having positive probability to occur. At time t we are in state s, and we consider agent i, with $x \in \{A, B\}$ being a best response action for i. We denote with K the set of agents who are different from i and choose an action different from x in state s. With probability $\gamma(1-\gamma)^{n-1}$, the only agent who receives a revision opportunity at time t is agent i, who will choose x with positive probability (either 1 if x is the unique best response or $1/2$ otherwise). If set K is empty, then we have reached a state in convention x. If set K is non-empty, then consider agent $j \in K$: With probability at least equal to $\tau^2 \frac{1}{n-1}$ the matches involving i and j terminate, and i and j are then matched together. At time $t+1$, with probability $\gamma(1-\gamma)^{n-1}$ the only agent who receives a revision opportunity is agent j, who will choose x with positive probability, since j is matched with an agent choosing x, and the fraction of other agents choosing x has not decreased with respect to i's decision at previous time. As a result, the cardinality of set K gets reduced by 1. If set K is now empty, then we have reached a state in convention x; otherwise, we repeat the above procedure. In a finite number of iterations, equal to the cardinality of K, set K is cleared, and convention x is reached. □

Proof of Proposition 1

Proof Without loss of generality, we set $h = 1$. We start by deriving the stochastic potential ρ_A^u of convention A under uniform mistakes. Starting from a state in convention B, one mistake is required to turn one agent's action from B to A. Suppose that k agents choose A, with $k \geq 1$. The agents who have the highest incentive to switch from B to A are those who are matched to an agent choosing A. We note that we always have a positive probability that an agent playing B is matched with someone who already changed to A. The decision of one such agent involves the comparison between

$$(1-\delta)a + \delta\frac{n-k-1}{n-2}c + \delta\frac{k-1}{n-2}a, \tag{3}$$

which is the expected utility of choosing A, and

$$(1-\delta)d + \delta\frac{n-k-1}{n-2}b + \delta\frac{k-1}{n-2}d, \tag{4}$$

which is the expected utility of choosing B.

If (3) is larger than or equal to (4), then A can be chosen by best response, increasing the number of agents choosing A from k to $k+1$. We observe that, if (3) is larger than or equal to (4) when k agents choose A, then this holds a fortiori when the number of such agents is $k+1$. Hence, the stochastic potential ρ_A^u is given by the minimum k such that (3) is larger than or equal to (4).

By means of simple algebra we obtain:

$$\rho_A^u = \begin{cases} 1 & \text{if } \delta \leq \frac{a-d}{a-d+b-c}, \\ 1 + \left\lceil \left(1 - \frac{a-d}{\delta(a-d+b-c)}\right)(n-2) \right\rceil & \text{if } \delta > \frac{a-d}{a-d+b-c}, \end{cases} \tag{5}$$

where $\lceil \cdot \rceil$ denotes the ceiling function.

With an analogous reasoning, we can derive an expression for the stochastic potential of convention B under uniform mistakes, which we denote by ρ_B^u:

$$\rho_B^u = \begin{cases} 1 & \text{if } \delta \leq \frac{b-c}{a-d+b-c}, \\ 1 + \left\lceil \left(1 - \frac{b-c}{\delta(a-d+b-c)}\right)(n-2) \right\rceil & \text{if } \delta > \frac{b-c}{a-d+b-c}. \end{cases} \tag{6}$$

We now compare ρ_A^u with ρ_B^u to determine which conventions are stochastically stable. If $0 < \tau \leq \tilde{\tau}_m$, then $0 < \delta < \tilde{\tau}_m$ since $\delta < \tau$. Hence $\rho_A^u = \rho_B^u = 1$, which means that convention A and convention B are both stochastically stable. If $\tilde{\tau}_m < \tau \leq \tilde{\tau}_M$, then $\tilde{\tau}_m < \delta < \tilde{\tau}_M$ for n sufficiently large, since $\delta < \tau$ and $\lim_{n \to +\infty} \delta = \tau$. Suppose that A is the risk-dominant action and, hence, $(b-c)/(a-d+b-c) < \delta < (a-d)/(a-d+b-c)$. Therefore $\rho_A^u = 1 < \rho_B^u$, and so only convention A is stochastically stable. By the same token, if B is the risk-dominant action then it is the only stochastically stable convention. Finally, consider $\tilde{\tau}_M < \tau \leq 1$, which implies $\tilde{\tau}_M < \delta < 1$ if n is large enough (for reasons analogous to those given above). We observe that, for n sufficiently large, the following difference,

$$\rho_B^u - \rho_A^u = \left(\frac{(a-d) - (b-c)}{\delta(a-d+b-c)} \right)(n-2), \tag{7}$$

is strictly greater than 1 if A is the risk-dominant action, implying that convention A is the only stochastically stable convention. By the same token, if B is the risk-dominant action then it is the only stochastically stable convention. □

Proof of Proposition 2

Proof The determination of the stochastic potential of convention A under payoff-dependent mistakes, which we denote by ρ_A^p, proceeds as for the case with uniform mistakes considered in the proof of Proposition 1, with a few adjustments. We note that the path with minimum total resistance to move from convention B to convention A requires that mistakes are made, if possible, by agents who are matched to someone already choosing A (which is always possible but for the first mistake) and also that mistakes are made sequentially; the reason is that the expected loss of choosing A is lower for an agent linked to someone choosing A, and it decreases in the number of agents choosing A. Hence, the first mistake is weighted by $f(b-c)$, since $b-c$ is the payoff loss when A is chosen instead of B. Moreover, each of the following mistakes, until the threshold of ρ_A^u agents choosing A is reached, is weighted by $f\left(\frac{\delta(n-k-1)}{n-2}(b-c) - \left[1 - \frac{\delta(n-k-1)}{n-2}\right](a-d)\right)$, where the argument of function f is the expected loss of choosing A instead of B when k agents are choosing A, and the current partner is choosing A. Therefore:

$$\rho_A^p = \begin{cases} f(b-c) & \text{if } \delta \leq \frac{a-d}{a-d+b-c}, \\ f(b-c) + \sum_{k=1}^{\rho_A^u - 1} f\left(\frac{\delta(n-k-1)}{n-2}(b-c) - \left[1 - \frac{\delta(n-k-1)}{n-2}\right](a-d)\right) & \text{if } \delta > \frac{a-d}{a-d+b-c}. \end{cases}$$

The stochastic potential of convention B under payoff-dependent mistakes, which we denote by ρ_B^p, is obtained similarly, considering what done for ρ_B^u and adjusting the weights of mistakes analogously:

$$\rho_B^p = \begin{cases} f(a-d) & \text{if } \delta \leq \frac{b-c}{a-d+b-c}, \\ f(a-d) + \sum_{k=1}^{\rho_B^u - 1} f\left(\frac{\delta(n-k-1)}{n-2}(a-d) - \left[1 - \frac{\delta(n-k-1)}{n-2}\right](b-c)\right) & \text{if } \delta > \frac{b-c}{a-d+b-c}. \end{cases}$$

We compare ρ_A^p with ρ_B^p to determine which conventions are stochastically stable. Suppose that A is the risk-dominant action, $a-d > b-c$. Since f is a strictly increasing function, and noting that $\rho_A^u \le \rho_B^u$ always, we can conclude that $\rho_A^p < \rho_B^p$, which means that convention A is the unique stochastically stable convention irrespectively of δ, and hence, for every $\tau \in (0, 1]$. If, instead, B is assumed to be the risk-dominant action, then $a-d < b-c$, and we obtain that convention B is the unique stochastically stable convention for every $\tau \in (0, 1]$. □

Proof of Proposition 3

Proof The determination of the stochastic potential of convention A when mistakes are condition-dependent, which we denote by ρ_A^c, proceeds as for the case with uniform mistakes considered in the proof of Proposition 1, with a few adjustments. We note that the path with minimum total resistance to move from convention B to convention A requires that mistakes are made, if possible, by agents who are matched to someone already choosing A (which is always possible but for the first mistake); the reason is that the resistance to mistake of an agent decreases with decreasing payoff currently earned by the agent, and a miscoordinating agent earns the lowest payoff. Hence, the first mistake is weighted by $g(b)$, since b is the payoff earned by every agent when all agents choose B. Moreover, each of the following mistakes, until the threshold of ρ_A^u agents choosing A is reached, is weighted by $g(d)$, since d is the payoff earned by an agent playing B who is matched with an agent playing A. Therefore:

$$\rho_A^c = \begin{cases} g(b) & \text{if } \delta \le \frac{a-d}{a-d+b-c}, \\ g(b) + g(d)(\rho_A^u - 1) & \text{if } \delta > \frac{a-d}{a-d+b-c}. \end{cases} \tag{8}$$

The stochastic potential of convention B under condition-dependent mistakes, which we denote by ρ_B^c, is obtained similarly, considering what done for ρ_B^u and adjusting the weights of mistakes analogously:

$$\rho_B^c = \begin{cases} g(a) & \text{if } \delta \le \frac{b-c}{a-d+b-c}, \\ g(a) + g(c)(\rho_B^u - 1) & \text{if } \delta > \frac{b-c}{a-d+b-c}. \end{cases} \tag{9}$$

We now compare ρ_A^c with ρ_B^c to determine which conventions are stochastically stable. If $0 < \tau \le \tilde{\tau}_m$, then $0 < \delta < \tilde{\tau}_m$ since $\delta < \tau$; hence, $\rho_A^c = g(b) > g(a) = \rho_B^c$, since g is a strictly increasing function and $b > a$, and so convention B is the unique stochastically stable convention.

We now consider the case where $\tilde{\tau}_m < \tau \le \tilde{\tau}_M$, which implies $\tilde{\tau}_m < \delta < \tilde{\tau}_M$ for n sufficiently large, since $\delta < \tau$ and $\lim_{n\to+\infty} \delta = \tau$. Suppose that A is the risk-dominant action and, hence, $(b-c)/(a-d+b-c) < \delta < (a-d)/(a-d+b-c)$. Therefore, we have that $\rho_A^c = g(b)$ and $\rho_B^c = g(a) + g(c)(\rho_B^u - 1)$. Since ρ_B^u grows unboundedly in n, we can conclude that convention A is the unique stochastically stable convention when n is large enough. If, instead, B is assumed to be the risk-dominant action, then by the same token $\rho_B^c = g(a)$ and $\rho_A^c = g(b) + g(d)(\rho_A^u - 1)$, and we obtain that convention B is the unique stochastically stable convention.

Finally, we consider the case where $\tilde{\tau}_M < \tau \leq 1$, which implies $\tilde{\tau}_M < \delta < 1$ if n is large enough (for reasons analogous to those given above). Taking the difference between (8) and (9), and considering n sufficiently large, we can write

$$\rho_A^c - \rho_B^c < (>)\ 0 \Longleftrightarrow \frac{g(c)}{g(d)} > (<)\ \frac{\rho_A^u}{\rho_B^u}, \tag{10}$$

since both ρ_A^u and ρ_B^u grow unboundedly when $n \to \infty$. By using (5) and (6) to substitute for ρ_A^u and ρ_B^u, and noting that $\lim_{n\to+\infty} \delta = \tau$, we obtain

$$\rho_A^c - \rho_B^c < (>)\ 0 \Longleftrightarrow \frac{g(c)}{g(d)} > (<)\ \frac{\tau(a-d+b-c)-(a-d)}{\tau(a-d+b-c)-(b-c)} = \beta(\tau). \tag{11}$$

Therefore, $g(c)/g(d) > (<)\beta(\tau)$ implies that $\rho_A^c < (>)\ \rho_B^c$, and hence, convention A (B) is the unique stochastically stable convention. □

References

Agrawal, A.: Genetic loads under fitness-dependent mutation rates. J. Evol. Biol. **15**(6), 1004–1010 (2002)

Agrawal, A.F., Wang, A.D.: Increased transmission of mutations by low-condition females: evidence for condition-dependent dna repair. PLoS Biol. **6**(2), e30 (2008)

Alós-Ferrer, C., Netzer, N.: The logit-response dynamics. Games Econ. Behav. **68**(2), 413–427 (2010)

Alós-Ferrer, C., Netzer, N.: Robust stochastic stability. Econ. Theory **58**(1), 31–57 (2015). https://doi.org/10.1007/s00199-014-0809-z

Alós-Ferrer, C., Weidenholzer, S.: Contagion and efficiency. J. Econ. Theory **143**(1), 251–274 (2008)

Anwar, A.W.: On the co-existence of conventions. J. Econ. Theory **107**(1), 145–155 (2002)

Bergin, J., Lipman, B.L.: Evolution with state-dependent mutations. Econometrica **64**, 943–956 (1996)

Bhaskar, V., Vega-Redondo, F.: Migration and the evolution of conventions. J. Econ. Behav. Organ. **55**(3), 397–418 (2004)

Bilancini, E., Boncinelli, L.: Social coordination with locally observable types. Econ. Theory **65**(4), 975–1009 (2018). https://doi.org/10.1007/s00199-017-1047-y

Binmore, K., Samuelson, L.: Muddling through: noisy equilibrium selection. J. Econ. Theory **74**(2), 235–265 (1997)

Blume, L.E.: The statistical mechanics of strategic interaction. Games Econ. Behav. **5**(3), 387–424 (1993)

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

Blume, L.E.: How noise matters. Games Econ. Behav. **44**(2), 251–271 (2003)

Blume, A., Temzelides, T.: On the geography of conventions. Econ. Theory **22**(4), 863–873 (2003). https://doi.org/10.1007/s00199-002-0350-3

Boncinelli, L., Pin, P.: The stochastic stability of decentralized matching on a graph. Games Econ. Behav. **108**, 239–244 (2018)

Carvalho, J.-P.: Coordination and culture. Econ. Theory **64**(3), 449–475 (2017). https://doi.org/10.1007/s00199-016-0990-3

Cotton, S.: Condition-dependent mutation rates and sexual selection. J. Evol. Biol. **22**(4), 899–906 (2009)

Efferson, C., Roca, C.P., Vogt, S., Helbing, D.: Sustained cooperation by running away from bad behavior. Evol. Hum. Behav. **37**(1), 1–9 (2016)

Ellison, G.: Learning, local interaction, and coordination. Econometrica **61**, 1047–1071 (1993)

Ellison, G.: Learning from personal experience: one rational guy and the justification of myopia. Games Econ. Behav. **19**(2), 180–210 (1997)

Ely, J.C.: Local conventions. Adv. Theor. Econ. **2**(1) (2002)

Foster, D., Young, P.: Stochastic evolutionary game dynamics. Theor. Popul. Biol. **38**(2), 219–232 (1990)

Freidlin, M.I., Wentzell, A.D.: Random perturbations. In: Random Perturbations of Dynamical Systems, pp. 15–43. Springer, Berlin (1998)
Frey, V., Corten, R., Buskens, V.: Equilibrium selection in network coordination games: an experimental study. Rev. Netw. Econ. **11**(3) (2012)
Goyal, S., Vega-Redondo, F.: Network formation and social coordination. Games Econ. Behav. **50**(2), 178–207 (2005)
Harsanyi, J.C., Selten, R.: A General Theory of Equilibrium Selection in Games, vol. 1. The MIT Press, Cambridge (1988)
Haselton, M.G., Buss, D.M.: Error management theory: a new perspective on biases in cross-sex mind reading. J. Pers. Soc. Psychol. **78**(1), 81 (2000)
Huttegger, S., Skyrms, B.: Emergence of a signaling network with "probe and adjust". In: Sterelny, K., Calcott, B., Joyce, R. (eds.) Signaling, Commitment, and Emotion. MIT Press, Cambridge (2012)
Hwang, S.-H., Newton, J.: Payoff-dependent dynamics and coordination games. Econ. Theory **64**(3), 589–604 (2017)
Hwang, S.-H., Naidu, S., Bowles, S.: Social conflict and the evolution of unequal conventions. Technical report, Santa Fe Institute Working Paper (2016)
Hwang, S.-H., Lim, W., Neary, P., Newton, J.: Conventional contracts, intentional behavior and logit choice: equality without symmetry. Games Econ. Behav. **110**, 273–294 (2018)
Jackson, M.O., Watts, A.: On the formation of interaction networks in social coordination games. Games Econ. Behav. **41**(2), 265–291 (2002)
Kandori, M., Rob, R.: Evolution of equilibria in the long run: a general theory and applications. J. Econ. Theory **65**(2), 383–414 (1995)
Kandori, M., Mailath, G.J., Rob, R.: Learning, mutation, and long run equilibria in games. Econometrica **61**, 29–56 (1993)
Klaus, B., Newton, J.: Stochastic stability in assignment problems. J. Math. Econ. **62**, 62–74 (2016)
Lee, I.H., Szeidl, A., Valentinyi, A.: Contagion and state dependent mutations. Adv. Theor. Econ. **3**(1) (2003)
Lewis, D.: Convention: A Philosophical Study, p. 2008. Wiley, Hoboken (1969)
Lim, W., Neary, P.R.: An experimental investigation of stochastic adjustment dynamics. Games Econ. Behav. **100**, 208–219 (2016)
Macy, M.W., Willer, R.: From factors to actors: computational sociology and agent-based modeling. Annu. Rev. Sociol. **28**, 143–166 (2002)
Maruta, T.: Binary games with state dependent stochastic choice. J. Econ. Theory **103**(2), 351–376 (2002)
Mäs, M., Nax, H.H.: A behavioral study of noise in coordination games. J. Econ. Theory **162**, 195–208 (2016)
McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. Games Econ. Behav. **10**(1), 6–38 (1995)
Morris, S.: Contagion. Rev. Econ. Stud. **67**(1), 57–78 (2000)
Myerson, R.B.: Refinements of the Nash equilibrium concept. Int. J. Game Theory **7**(2), 73–80 (1978)
Naidu, S., Hwang, S.-H., Bowles, S.: Evolutionary bargaining with intentional idiosyncratic play. Econ. Lett. **109**(1), 31–33 (2010)
Nax, H.H., Pradelski, B.S.: Evolutionary dynamics and equitable core selection in assignment games. Int. J. Game Theory **44**(4), 903–932 (2015)
Nax, H.H., Burton-Chellew, M.N., West, S.A., Young, H.P.: Learning in a black box. J. Econ. Behav. Organ. **127**, 1–15 (2016)
Neary, P.R.: Competing conventions. Games Econ. Behav. **76**(1), 301–328 (2012)
Nesse, R.M.: Natural selection and the regulation of defenses: a signal detection analysis of the smoke detector principle. Evol. Hum. Behav. **26**(1), 88–105 (2005)
Newton, J.: Coalitional stochastic stability. Games Econ. Behav. **75**(2), 842–854 (2012)
Newton, J.: The Deconstruction of Conventions, 19 July 2018. Available at SSRN https://ssrn.com/abstract=3216269 (2018a)
Newton, J.: Evolutionary game theory: a renaissance. Games **9**(2), 31 (2018b)
Newton, J., Sawa, R.: A one-shot deviation principle for stability in matching problems. J. Econ. Theory **157**, 1–27 (2015)
Norman, T.W.: Rapid evolution under inertia. Games Econ. Behav. **66**(2), 865–879 (2009)
Oechssler, J.: Decentralization and the coordination problem. J. Econ. Behav. Organ. **32**(1), 119–135 (1997)
Peski, M.: Generalized risk-dominance and asymmetric dynamics. J. Econ. Theory **145**(1), 216–248 (2010)

Pradelski, B.S., Young, H.P.: Learning efficient Nash equilibria in distributed systems. Games Econ. Behav. **75**(2), 882–897 (2012)

Robson, A.J., Vega-Redondo, F.: Efficient equilibrium selection in evolutionary games with random matching. J. Econ. Theory **70**(1), 65–92 (1996)

Sawa, R., Wu, J.: Prospect dynamic and loss dominance. Games Econ. Behav. **112**, 98–124 (2018)

Sharp, N.P., Agrawal, A.F.: Evidence for elevated mutation rates in low-quality genotypes. Proc. Natl. Acad. Sci. **109**(16), 6142–6146 (2012)

Shaw, F., Baer, C.: Fitness-dependent mutation rates in finite populations. J. Evol. Biol. **24**(8), 1677–1684 (2011)

Shi, F.: Comment on on the co-existence of conventions [J. Econ. Theory 107 (2002)]. J. Econ. Theory **1**(148), 418–421 (2013)

Skyrms, B.: The Stag Hunt and the Evolution of Social Structure. Cambridge University Press, Cambridge (2004)

Staudigl, M., Weidenholzer, S.: Constrained interactions and social coordination. J. Econ. Theory **152**, 41–63 (2014)

Thuijsman, F., Peleg, B., Amitai, M., Shmida, A.: Automata, matching and foraging behavior of bees. J. Theor. Biol. **175**(3), 305–316 (1995)

Van Damme, E., Weibull, J.W.: Evolution in games with endogenous mistake probabilities. J. Econ. Theory **106**(2), 296–315 (2002)

Weidenholzer, S.: Coordination games and local interactions: a survey of the game theoretic literature. Games **1**(4), 551–585 (2010)

Weidenholzer, S.: Long-run equilibria, dominated strategies, and local interactions. Games Econ. Behav. **75**(2), 1014–1024 (2012)

Young, H.P.: The evolution of conventions. Econometrica **61**, 57–84 (1993)

Young, H.P.: Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press, Princeton (2001)

Young, H.P.: Learning by trial and error. Games Econ. Behav. **65**(2), 626–643 (2009)

Young, H.P.: The evolution of social norms. Annu. Rev. Econ. **7**(1), 359–387 (2015)