DISEI - Università degli Studi di Firenze

Working Papers - Economics

The Evolution of Conventions under Condition-Dependent Mistakes

E. Bilancini and L. Boncinelli

Working Paper N. 11/2016

DISEI, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze, Italia
www.disei.unifi.it

The Evolution of Conventions under Condition-Dependent Mistakes

Ennio Bilancini[*] Leonardo Boncinelli[†]

July 4, 2016

Abstract

In this paper we study the long run convention emerging from stag-hunt interactions when errors converge to zero at a rate that is positively related to the payoff earned in the previous period. We refer to such errors as *condition-dependent mistakes*. We find that, if interactions are sufficiently stable over time, then the payoff-dominant convention emerges in the long run. Moreover, if interactions are neither too stable nor too volatile, then the risk-dominant convention is selected in the long run. Finally, if interactions are quite volatile, then the maximin convention emerges even if it is not risk-dominant. We introduce the notion of *condition-adjusted-risk-dominance* to characterize the convention emerging in the long run under condition-dependent mistakes. We contrast these results with the results obtained under alternative error models: *uniform mistakes*, i.e., errors converge to zero at a rate that is constant over states, and *payoff-dependent mistakes*, i.e., errors converge to zero at a rate that depends on expected losses.

JEL classification code: C72; C73.

Keywords: risk-dominant; payoff-dominant; maximin; mistakes; stag hunt; stochastic stability.

[*]Dipartimento di Economia "Marco Biagi", Università degli Studi di Modena e Reggio Emilia, Viale Berengario 51, 43 ovest, 41121 Modena, Italia. Tel.: +39 059 205 6843, fax: +39 059 205 6947, email: ennio.bilancini@unimore.it.

[†]Dipartimento di Scienze per l'Economia e l'Impresa, Università degli Studi di Firenze, Via delle Pandette 9, 50127 Firenze, Italia. Tel.: +39 055 2759578, fax: +39 055 2759910, email: leonardo.boncinelli@unifi.it.

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 (Lewis, 2008).[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 mis-coordination 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 stag-hunt game played by a population of agents. We model agents' interactions as a pairwise match of randomly drawn individuals, but in addition we allow for some degree of stability of the interaction. We derive the emerging long-run convention applying the notion of stochastic stability (Young, 1993; Kandori et al., 1993). Basically, stochastic stability selects those equilibria that are relatively easiest to reach in terms of mutations – also referred to as errors or mistakes – 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 give reliable predictions in terms of equilibrium selection.

In this regard our novelty is the investigation of what convention is stochastically stable when errors converge to zero at a rate that is positively related to the payoff earned in the past. We call such an error model *condition-dependent mistakes*, drawing from the literature on genetics and biology (Agrawal, 2002). We contrast the results under condition-dependent mistakes with the results obtained under alternative error models: 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*.

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 stable over time. Instead, if interactions are mildly volatile, then the selection of the risk dominant convention is obtained, while if interactions are very volatile, then the

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

maximin convention is obtained. 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.

As a consequence of this, 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 last finding 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 is related to the literature on contagion (see, e.g., Morris, 2000 and Alós-Ferrer and Weidenholzer, 2008).

The rest of the paper is organized as follows. Section 2 contrasts this paper with the 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 robustness of our findings, and Section 6 suggests an interpretation of the contribution and draws speculative conclusions.

2 Related Literature

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) and, more generally, is reminiscent of the notion of proper equilibrium applied in game theory (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).[2] 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. This justification can rely on 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. Moreover, 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), 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.

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.,

[2] 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.

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.[3] 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. Notably, if restricted mobility is imposed in these models with locations, the coexistence of conventions can be obtained, with one convention being established in one place and the other convention in another place (Anwar, 2002; Blume and Temzelides, 2003). Bilancini and Boncinelli (2015) 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.

In the present paper we adopt a different model of interaction, where agents are randomly matched to play a stag-hunt, but their interaction does not necessarily terminate after the first round of play. In particular, we introduce a parameter that measures the stability of the interaction and that determines 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. Our model of interaction is very similar to that in Robson and Vega-Redondo (1996), who also stress the sensitivity of stochastically stable conventions to the specification of the matching mechanism. While in our paper agents are myopic best responders, in their model they tend to adopt the strategy that led to the highest average payoff. 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.

[3]For a general framework for local interaction models with an exogenous interaction structure see Peski (2010). See also Weidenholzer (2010) for a recent survey on local interaction models focusing on social coordination.

	A	B
A	a	c
B	d	b

Figure 1: A stag-hunt game.

3 Model

We consider a population of n agents, indexed according to the set of integers $N = \{1, 2, ..., n\}$, 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 Figure 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 $b + d > a + c$.[4] 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.

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 $\mu^t(i) \in N$ 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 $i, j \in N$.

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

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

who is given a revision opportunity at time t takes a choice following a myopic best-reply rule against s^{t-1}, randominzing incase of indifference. This means that actions are taken before an agent knows whether his current match terminates (see Section 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:

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

By using δ, we can now write the payoff that agent i expects to earn at time t if he chooses action $x \in \{A, B\}$ as follows:

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

We say that action $x \in \{A, B\}$ is best reply 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.

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

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 $\prod_{k=0}^{n-1} \frac{1}{n-1-2i}$). 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 reply 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 reply, 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. □

4 Stochastically Stable Conventions

We follow Young (1993) and Kandori et al. (1993) and we add agents' mistakes to the model of 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 the action that is not best reply, 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, where s is the state at time $t - 1$. 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. As a consequence 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.[5] 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 larger 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 $0 < \tilde{\tau}_m < \tilde{\tau}_M < 1$.

The first error model that we consider is one where all mistakes are equally 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.

[5] 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.

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.*

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 linked 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

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

which is the expected utility of choosing A, and

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

which is the expected utility of choosing B.

If (3) is larger than or equal to (4), then A can be chosen by best reply, 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 $k+1$.

By means of simple algebra we hence 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:

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

Since $\delta < \tau$, if $\tau \leq \frac{a-d}{a-d+b-c}$ then $\delta < \frac{a-d}{a-d+b-c}$. Moreover, since $\lim_{n \to +\infty} \delta = \tau$, if $\tau > \frac{a-d}{a-d+b-c}$ then $\delta > \frac{a-d}{a-d+b-c}$ if n is large enough. Hence, for n large enough, the conditions distinguishing the two cases in (5) can be rewritten by substituting τ with δ. By the same token, an analogous substitution can be done for (6).

We now compare ρ^u_A with ρ^u_B to determine which conventions are stochastically stable. If $\tau \leq \tilde{\tau}_m$, then $\rho^u_A = \rho^u_B = 1$, and so both convention A and convention B are stochastically stable. Otherwise, let $x \in \{A, B\}$ be the risk-dominant action and y the other action. If $\tilde{\tau}_m < \tau \leq \tilde{\tau}_M$ then $\rho^u_x = 1 < \rho^u_y$, and so only convention x is stochastically stable. Finally, consider $\tau > \tilde{\tau}_M$. We observe that, for n sufficiently large, the following difference

$$\left(1 - \frac{a-d}{\delta(a-d+b-c)}\right)(n-2) - \left(1 - \frac{b-c}{\delta(a-d+b-c)}\right)(n-2) \tag{7}$$

is strictly greater than 1 if $x = B$ and strictly smaller than -1 if $x = A$. It follows that $1 < \rho^u_x < \rho^u_y$, and so convention x is the unique stochastically stable convention. □

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 reply 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 reply action and the expected payoff of making a mistake.

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

$$r(i, s) = f(\Pi^s_i(x) - \Pi^s_i(y)),$$

where f is a strictly increasing function with positive values, x is best reply 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* $\tau \in (0, 1]$.

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. Indeed, 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 $f\left(\frac{\delta(n-k-1)}{n-k}(b-c) - \left[1 - \frac{\delta(n-k-1)}{n-k}\right](a-d)\right)$, 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. Hence:

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

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:

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

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

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 and 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 very low.

Finally, we turn to consider an error model where the resistance to mistakes is positively related to the payoff earned in the past 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. The following definition aims at capturing the property which is crucial for the long-run evolution in such cases.

Definition **4** (Condition-adjusted-risk-dominance)**.** *Given a stag-hunt game as in Figure 1 and a condition-dependet error model with resistance function g, we say that action $x \in \{A, B\}$ is condition-adjusted-risk-dominant, or c-risk-dominant, if:*

$$x = \arg\max_{\{A,B\}} \;\; \pi(x, A)\frac{\tau g(c) + (1-\tau)g(d)}{g(c) + g(d)} + \pi(x, B)\frac{\tau g(d) + (1-\tau)g(c)}{g(c) + g(d)}.$$

It can be useful to compare c-risk-dominance with risk-dominance. While an action is risk-dominant if it is best reply against a belief that the opponent randomizes between A and B with probability $1/2$, an action is c-risk-dominant if it is best reply against a belief that the opponent chooses A with probability $[\tau g(c) + (1-\tau)g(d)]/[g(c) + g(d)]$ and B with the remaining probability. This can be interpreted as if a player does not have a pure Laplacian belief about the opponent's behavior, but rather he expects that the maximin action is more (or less) likely to be played in environments where the termination probability is larger (or lower) than $1/2$ (to an extent that depends on miscoordination payoffs, c and d, on the resistance to mistakes function g, and on the termination probability τ). We use the term c-risk-dominant convention to refer to either convention A or convention B, depending on which convention is actually c-risk-dominant action.

Figure 2 allows to visualize how the notion of c-risk-dominance relates with the standard notions of payoff-dominance, risk-dominance, and maximin. The results can be derived in a straightforward way from Definition 4. When $\tau < \tau_m$, the c-risk-dominant action coincides with the payoff-dominant action, which is B. When $\tau = \tau_m$, both the payoff-dominant action and the risk-dominant action are c-risk dominant. When $\tau_m < \tau \leq \tau_M$, the c-risk-dominant action coincides with the risk-dominant action. We remind that A is risk-dominant if $a-d > b-c$, while B is risk-dominant if $a-d < b-c$. When $\tau > \tau_M$, the c-risk-dominant action is the maximin action if

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

or the risk-dominant action if the reverse inequality holds. If $g(c)/g(d) = \ell(\tau)$, both the risk-dominant action and the maximin action are c-risk-dominant. We observe that, when $\tau = 1$, the inequality $g(c)/g(d) > (b-c)(a-d)$ ensures that the c-risk-dominant action corresponds to the maximin action.

Proposition 3 provides a characterization of stochastic stability under condition-dependent mistakes in terms of c-risk-dominance.

Proposition 3. *Under condition-dependent mistakes, if n is large enough then the c-risk-dominant convention is the unique stochastically stable convention for all $\tau \in (0,1]$ such that $\tau \neq \tilde{\tau}_m$ and $g(c)/g(d) \neq \ell(\tau)$.*[6]

Proof. The determination of the stochastic potential of convention A under condition-dependent mistakes, 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. Indeed, 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

[6]We opted not to deal with the cases $\tau = \tilde{\tau}_m$ and $g(c)/g(d) = \ell(\tau)$, because they have measure zero in the parameter space, and they would bring in the analysis annoying complications. Indeed, when $\tau = \tilde{\tau}_m$ one mistake is enough to move from one convention to the other and vice versa, so the payoff-dominant convention is the unique stochastically stable convention because mistakes are less likely when starting from it. However, we note that such a result depends on the details of the updating rule for actions: if a revising agent chooses the current action in case both actions are best reply (as if there were negligible costs to change) then it is the risk-dominant convention to be the unique stochastically stable convention. When instead $g(c)/g(d) = \ell(\tau)$, determining which convention is stochastically stable is complicated by the working of the ceiling function, so that either the risk-dominant convention, or the maximin convention, or both can be stochastically stable, depending on specific values of n.

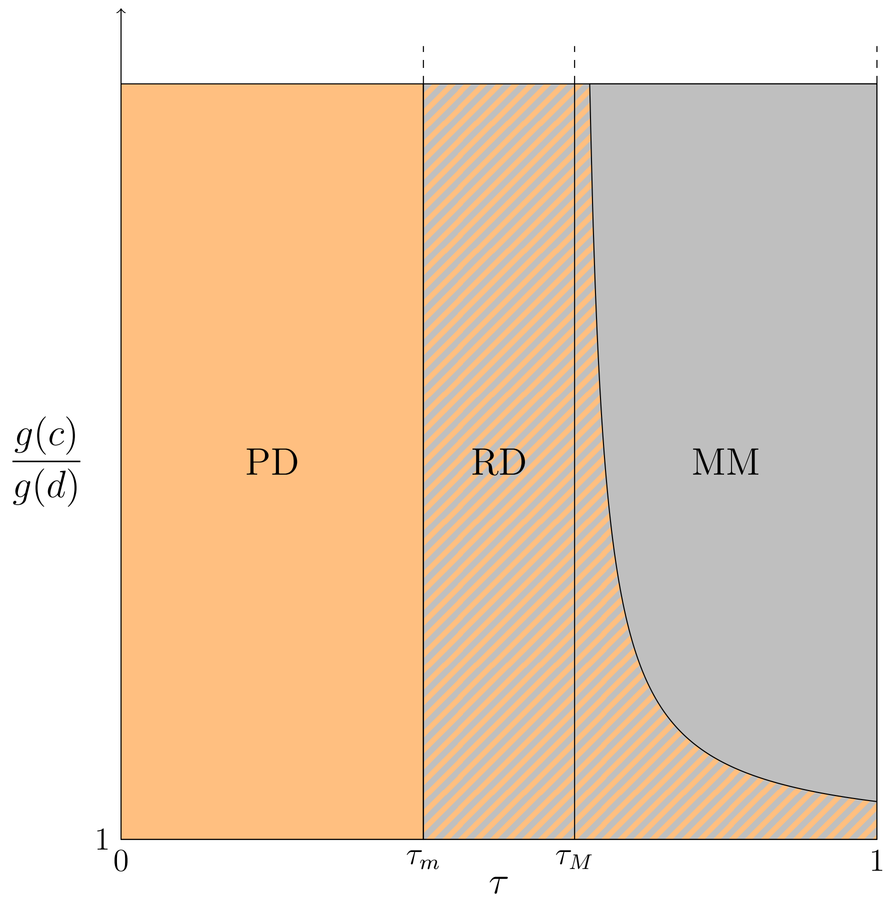

Figure 2: The orange color refers to action B, which is payoff-dominant, and the gray color to action A, which is maximin. The orange-grey stripes pattern referes to either action B or action A, depending on which action is risk-dominant.

by $g(d)$, since d is the payoff earned by an agent playing B who is matched with an agent playing A. Hence:

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

The stochastic potential of convention B under payoff-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:

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

By means of the same argument used in the proof of Proposition 1, we can rewrite the conditions distinguishing the two cases in (8) and in (9) by substituting τ with δ.

We now compare ρ^c_A with ρ^c_B to determine which conventions are stochastically stable. If $0 < \tau < \tilde{\tau}_m$, then $\rho^c_A = g(b) > g(a) = \rho^c_B$, since g is a strictly increasing function and $b > a$, and so convention B is the unique stochastically stable convention. We already know that B is also c-risk-dominant for $0 < \tau < \tilde{\tau}_m$.

We now consider the case where $\tilde{\tau}_m < \tau \leq \tilde{\tau}_M$. Suppose that A is the risk-dominant action, so that $a - d > b - c$. Hence, we have that $\rho^c_A = 1 < 1 + g(a) = \rho^u_y$, and so convention A is the unique stochastically stable convention. 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. In any case, we already know that the risk-dominant action is also c-risk-dominant for $\tilde{\tau}_m < \tau \leq \tilde{\tau}_M$.

Finally, we consider the case where $\tilde{\tau}_M < \tau \leq 1$. For n sufficiently large, it is a matter of computation to see that

$$\frac{g(c)}{g(d)} > (<)\ \ell(\tau) \tag{10}$$

implies that $\rho^c_A < (>) \rho^c_B$, and hence convention A (B) is the unique stochastically stable convention. We note that the inequality in (10) also determines which action is c-risk-dominant, so that in this case as well the stochastically stable convention coincides with the c-risk-dominant convention. □

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. However, a mistake in convention A is weighted by $g(a)$, while a mistake in convention B is weighted by $g(b)$, since a and b are the individual payoffs earned in the two conventions, respectively. Therefore, the minimum stochastic potential is that of convention B, which is hence the stochastically stable convention.

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 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, 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 $g(c)/g(d)$, together with the relative size of the two basins of attractions, which depends on

τ and 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.

Table 1 summarizes all the results of this section, simplifying 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 stable 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).

	termination probability		
mistakes	$0 < \tau < \tilde{\tau}_m$	$\tilde{\tau}_m < \tau < \tilde{\tau}_M$	$\tilde{\tau}_M < \tau \leq 1$
uniform	MM, PD	RD	RD
payoff-dependent	RD	RD	RD
condition-dependent	PD	RD	RD, if $g(c)/g(d) < \ell(\tau)$ MM, if $g(c)/g(d) > \ell(\tau)$

Table 1: Stochastically stable conventions.

5 Discussion

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

A first dimension is the revision protocol. We are aware that the details of the revision protocol may have important consequences for the selection results obtained by stochastic stability.[7] 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 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.

A second dimension is the number of conventions. The substance of our results remain 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

[7] 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).

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.[8]

A third dimension concerns the timing between the choice of actions and the formation of matches. 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 remain stable until benefits are paid. Consistently 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 termination probability is lower than 1. 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 reply. 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 $\tau < \tilde{\tau}_m$ for the baseline model. Only when we have complete random matching (i.e., $\tau = 1$) the above argument does not apply. Results are summarized in Table 2.

	termination probability	
mistakes	$0 < \tau < 1$	$\tau = 1$
uniform	MM, PD	RD
payoff-dependent	RD	RD
condition-dependent	PD	RD, if $g(c)/g(d) < (b-c)/(a-d)$ MM, if $g(c)/g(d) > (b-c)/(a-d)$

Table 2: Stochastically stable conventions when actions are taken after the formation of matches.

[8]The same definition of risk-dominance needs attention when more than two conventions exist. Morris et al. (1995) define an action to be p-dominant if it is best reply against any mixed strategy that assigns at least probability p to that action. Ellison (2000) extends the notion of risk-dominance to games with an arbitrary number of conventions, showing that a $1/2$-dominant convention is the unique long run equilibrium in a noisy best-reply process, both under global and under local interactions where agents are arranged on a circle. See Peski (2010) for a extension of this result to general networks of interaction.

A fourth dimension to consider is about 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 if 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 it is likely to be stochastically stable.

A fifth dimension concerns which of past payoffs determine the condition of an agent, and hence the likelihood of his mistakes. 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.[9]

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,

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

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 stable 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 stable, 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 among human beings have been for most of the time extremely stable – at least back in the period when human societies were typically made of small numbers of hunterers/gatherers (roughly about 10 thousands years ago). So, condition-dependent mistakes might have spread over other mistakes 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 less and less stable, 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.

References

Agrawal, A. (2002). Genetic loads under fitness-dependent mutation rates. *Journal of Evolutionary Biology 15*(6), 1004–1010.

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

Alós-Ferrer, C. and N. Netzer (2010). The logit-response dynamics. *Games and Economic Behavior 68*(2), 413–427.

Alós-Ferrer, C. and N. Netzer (2015). Robust stochastic stability. *Economic Theory 58*(1), 31–57.

Alós-Ferrer, C. and S. Weidenholzer (2008). Contagion and efficiency. *Journal of Economic Theory 143*(1), 251–274.

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

Bergin, J. and B. L. Lipman (1996). Evolution with State-Dependent Mutations. *Econometrica 64*, 943–956.

Bhaskar, V. and F. Vega-Redondo (2004). Migration and the evolution of conventions. *Journal of Economic Behavior & Organization 55*(3), 397–418.

Bilancini, E. and L. Boncinelli (2015). Social coordination with locally observable types. mimeo.

Binmore, K. and L. Samuelson (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory 74*(2), 235–265.

Blume, A. and T. Temzelides (2003). On the geography of conventions. *Economic Theory 22*(4), 863–873.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior 5*(3), 387–424.

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

Blume, L. E. (2003). How noise matters. *Games and Economic Behavior 44*(2), 251–271.

Cotton, S. (2009). Condition-dependent mutation rates and sexual selection. *Journal of Evolutionary Biology 22*(4), 899–906.

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

Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies 67*(1), 17–45.

Ely, J. C. (2002). Local conventions. *Advances in Theoretical Economics 2*(1).

Goyal, S. and F. Vega-Redondo (2005). Network formation and social coordination. *Games and Economic Behavior 50*(2), 178–207.

Harsanyi, J. C. and R. Selten (1988). *A general theory of equilibrium selection in games*, Volume 1. The MIT Press.

Jackson, M. O. and A. Watts (2002). On the formation of interaction networks in social coordination games. *Games and Economic Behavior 41*(2), 265–291.

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

Kandori, M. and R. Rob (1995). Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory 65*(2), 383–414.

Lewis, D. (2008). *Convention: A philosophical study.* John Wiley & Sons.

Maruta, T. (2002). Binary games with state dependent stochastic choice. *Journal of Economic Theory 103*(2), 351–376.

Morris, S. (2000). Contagion. *The Review of Economic Studies 67*(1), 57–78.

Morris, S., R. Rob, and H. S. Shin (1995). p-dominance and belief potential. *Econometrica*, 145–157.

Myerson, R. B. (1978). Refinements of the nash equilibrium concept. *International Journal of Game Theory 7*(2), 73–80.

Naidu, S., S.-H. Hwang, and S. Bowles (2010). Evolutionary bargaining with intentional idiosyncratic play. *Economics Letters 109*(1), 31–33.

Newton, J. (2012). Coalitional stochastic stability. *Games and Economic Behavior 75*(2), 842–854.

Oechssler, J. (1997). Decentralization and the coordination problem. *Journal of Economic Behavior & Organization 32*(1), 119–135.

Peski, M. (2010). Generalized risk-dominance and asymmetric dynamics. *Journal of Economic Theory 145*(1), 216–248.

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

Sharp, N. P. and A. F. Agrawal (2012). Evidence for elevated mutation rates in low-quality genotypes. *Proceedings of the National Academy of Sciences 109*(16), 6142–6146.

Shaw, F. and C. Baer (2011). Fitness-dependent mutation rates in finite populations. *Journal of Evolutionary Biology 24*(8), 1677–1684.

Skyrms, B. (2004). *The stag hunt and the evolution of social structure.* Cambridge University Press.

Staudigl, M. and S. Weidenholzer (2014). Constrained interactions and social coordination. *Journal of Economic Theory 152*, 41–63.

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

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.

Young, H. P. (2001). *Individual strategy and social structure: An evolutionary theory of institutions*. Princeton University Press.

Young, H. P. (2015). The evolution of social norms. *Annual Review of Economics 7*(1), 359–387.