

Social Coordination with Locally Observable Types

Ennio Bilancini*

Leonardo Boncinelli[†]

March 1, 2017

Abstract

In this paper we study the typical dilemma of social coordination between a risk-dominant convention and a payoff-dominant convention. In particular, we consider a model where a population of agents play a coordination game over time, choosing both the action and the network of agents with whom to interact. The main modelling novelties with respect to the existing literature are: (i) agents come in two distinct types, (ii) the interaction with a different type is costly, and (iii) an agent's type is unobservable prior to interaction. We show that when the cost of interacting with a different type is small with respect to the payoff of coordination, the payoff-dominant convention is the only stochastically stable convention; instead, when the cost of interacting with a different type is large, the only stochastically stable conventions are those where all agents of one type play the payoff-dominant action and all agents of the other type play the risk-dominant action.

JEL classification code: C73, D83, D85.

Keywords: coordination; conventions; evolution; stochastic stability; learning; network formation.

*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. ORCID: 0000-0002-2027-3992

[†]Dipartimento di Scienze per l'Economia e l'Impresa, Università degli Studi di Firenze, Via delle Pandette 9, 50127 Firenze, Italia. Tel.: +39-055-275-9578, fax: +39-055-275-9910, email: leonardo.boncinelli@unifi.it. ORCID: 0000-0003-0626-5133.

1 Introduction

Motivation. Social and economic interactions often require participants to coordinate their actions. Conventions specifying on what side to drive or how to share the product of joint work, and standards such as software or hardware platforms are examples, among many, of successful coordination. In this kind of situations coordination is an inherently strategic issue, giving rise to coordination games with a multiplicity of Nash equilibria, which can be seen as distinct conventions. A social dilemma that often arises in this setting is between a payoff-dominant action, which pays a higher payoff if the associated equilibrium is actually played, and a risk-dominant action, which performs better if out-of-equilibrium play happens (Harsanyi and Selten, 1988). Which of these equilibria is more likely to emerge in the long run has been a matter of study in evolutionary game theory.

Most of the related research so far has focused on identifying conditions that lead to the emergence of either the payoff-dominant or the risk-dominant convention. However, in many real situations both conventions persist for a long time (as discussed, for instance, in Sugden, 1995 and Goyal and Janssen, 1997). Also, agents are rarely fully homogeneous, being often characterized by a variety of traits that can affect the result of the social interaction. Sometimes such traits are not easily observable before interaction takes place, as it may be the case for preferences, religion or cultural heritage. This feature can help to explain the coexistence of conventions.

Indeed, what we typically observe in real life interactions is the formation of clusters of agents that are homogeneous with respect to these traits, despite their unobservability prior to interaction. Agents often adopt distinctive actions that end up to be conventions in their cluster, e.g., dress, appearance, jargon, diet, rituals, meeting places, etc. Some of these conventions can appear to be inferior choices but, nevertheless, they are adopted. The fact that these distinctive actions are endogenously chosen by the agents (in contrast with types that are instead exogenously given) and, in many cases, are also easily observable prior to interaction, suggests a novel explanation of why different conventions can coexist: they may act as signals allowing coordination across types. In other words, if it is known that a specific way of speaking/dressing is adopted mainly by agents with a certain trait, an individual might prefer to interact with the people showing that way of speaking/dressing – even if this forces such individual to coordinate on an action that yields low payoffs – because this allows him to interact with people that possess his own trait. Importantly, this kind of intuition points to a joint explanation of the co-existence of different conventions and the emergence of homogeneous clusters, at least for those situations where agents can

choose with whom to interact.

Our contribution. We explore social coordination in a population made of two different types that have a preference for own-type but can observe others' type only after first interaction, and that interact on a network which they can shape but with a constraint on the maximum number of neighbors.

More precisely, we consider a variant of the model in [Staudigl and Weidenholzer \(2014\)](#). With respect to the latter, we introduce the novelty that agents are of two different types, x and y , and that types are payoff relevant, in the sense that a penalty $d > 0$ is suffered when interacting with a different type. Importantly, actions taken by agents are globally observable, while types are not observable prior to interaction: each agent knows the type of agents with whom he is connected, and ignores the type of other agents, being only able to form expectations on the basis of the distribution of choices at the current population state.

Our main results concern the long-run prediction obtained by applying stochastic stability (see [Young, 1993](#), and [Kandori et al., 1993](#)), and are twofold. When d is low the payoff-dominant convention is the only stochastically stable outcome: in the long run all agents of both types end up choosing the payoff-dominant action. This result can be interpreted as a robustness check of [Staudigl and Weidenholzer \(2014\)](#). When instead d is sufficiently high, then the stochastically stable states are those where all agents of one type choose an action, and all agents of the other type choose the other action. To our knowledge, we are the first in the stream of literature on social coordination and stochastic stability to introduce agents' heterogeneity with local observability of types. As a result, we obtain that the risk-dominant action and the payoff-dominant action coexist in the long run for a reason that is substantially different from restrictions to agents' mobility (see the paragraph on related literature for a more articulated discussion). The intuition is the following. One single mutation *can* be enough to escape from a state where an agent, once he has deleted all of his links by mistake, would be confronted with the risk of interacting with a different type if he chooses to reconnect to agents choosing his same action. Moreover, one single mutation *is* enough provided that the penalty for interacting with a different type is sufficiently large. In our model the former condition holds for all states where actions and types are not perfectly correlated; instead, it never holds for the states where all agents of one type choose one action and all agents of the other type choose the other action. Therefore, when d is large the states where actions and types are perfectly correlated are harder to leave in terms of mutations, which makes them stochastically stable.

We observe that a single mutation would not be sufficient to leave an absorbing state

if the mutant is able to remember the agents with whom he was previously connected and to re-establish a link with them. Indeed, in our model agents are prevented from tracing back their previous mates. This can be interpreted as a lack of memory regarding past interactions. However, what is actually needed for our results is that agents have decaying memory, so that they will *eventually* forget the identity of previous mates. Coupled with positive inertia, this would imply that, with positive probability, the mutant has forgotten the identities of his previous interaction partners when he gets the opportunity to update his strategy.

The rest of the paper is organized as follows. The next paragraph surveys the relevant literature, contrasting our contribution with the existing ones. Section 2 introduces the basic elements of the model. Section 3 discusses the induced Markov chain, and provides some simple results for the unperturbed dynamics. Section 4 considers the perturbed dynamics and gives the main results concerning stochastic stability. Section 5 concludes by discussing the assumptions, and providing directions for future research. An appendix collects the proofs of the propositions on the unperturbed dynamics and of all lemmas, while the proofs of the propositions on the perturbed dynamics are given in the main text.

Related literature. Most papers on social coordination in the long run consider agents who follow myopic best reply rules and occasionally make mistakes.¹ The main message in this literature is that, when the interaction structure is exogenous, inefficient risk-dominant conventions emerge in the long run.^{2,3} However, when the interaction structure is endogenous this result does not necessarily hold and the payoff-dominant action can be selected in the long-run. The endogeneity of the interaction has been modeled mainly in two ways: (i) the agents can choose with whom to form an interaction network, and (ii) the agents can select a location among a number of locations available and then interact with agents in the same location.

In approach (i), to which our model belongs, network formation is typically associated to

¹See, e.g., Eshel et al. (1998), Alós-Ferrer and Weidenholzer (2008), Alós-Ferrer and Shi (2012) and Cui (2014), and references therein, for models of local interaction where agents follow imitative behavior.

²For global interaction models see, e.g., Kandori et al. (1993), Kandori and Rob (1995), Young (1993). For local interaction models see Blume (1993, 1995), Ellison (1993, 2000), Alós-Ferrer and Weidenholzer (2007) and Jiang and Weidenholzer (2016); for a general framework for local interaction models with an exogenous interaction structure see Peski (2010); finally, see Weidenholzer (2010) for a recent survey on local interaction models focusing on social coordination.

³Neary (2012) studies a model of social coordination where the interaction structure is exogenous and global, but agents are heterogeneous in their preferences about the action upon which to coordinate. In this setup only payoff-efficient conventions are selected.

a cost to maintain the existing links. In a non-cooperative setup, [Goyal and Vega-Redondo \(2005\)](#) show that when interaction is unconstrained (i.e., there is no bound to the number of agents one can form a link with) and costs to maintain a link are low, then the risk-dominant convention still emerges in the long run, while for relatively high costs to maintain a link the payoff-dominant convention does emerge.^{4,5} Our model follows the version of [Goyal and Vega-Redondo \(2005\)](#) where both the cost to maintain a link and the payoff flow is asymmetric (i.e., the agent who pays the cost is the only one to receive the payoff of interaction).⁶ With respect to it, our model differs for three distinctive features: (1) agents have a maximum number of interactions that can maintain at the same time, as in [Staudigl and Weidenholzer \(2014\)](#), (2) agents come in two distinct types with a preference for interacting with agents of their own type, and (3) an agent's type can be observed only if already connected to him. Indeed, the paper that is most related to ours is [Staudigl and Weidenholzer \(2014\)](#). They show that, when interaction is constrained in the sense that agents can only support a small number of links with respect to the population size, then the payoff-dominant convention emerges in the long run. Our paper shows that this result is robust to the introduction of features (2) and (3), provided that the preference for interacting with one's own type is not too strong; however, if the cost of interacting with agents of a different type is large enough, then efficiency is lost and both conventions co-exist in the long run.

In approach (ii), the interaction structure is constrained by the fact that an agent can interact only with those agents choosing the same location. [Oechssler \(1997\)](#), [Ely \(2002\)](#), and [Bhaskar and Vega-Redondo \(2004\)](#) are all nice examples of models in which agents play a coordination game and have to choose one location among the locations where the coordination game is played. In these models, the payoff-dominant convention typically emerges in the long run. In the light of our results, a case of particular interest is when locations are subject to a capacity constraint, thus impeding or forcing the movements of agents across them. The important fact here is that under this constraint the payoff-dominant convention is no longer the only one selected in the long run and, in addition,

⁴[Hojman and Szeidl \(2006\)](#) develop a related model with uni-directional payoff flows that accrue from all path-connected agents.

⁵[Jackson and Watts \(2002\)](#) study a cooperative (pairwise) network formation model and show that for low costs to maintain a link the risk-dominant convention is selected, while for high costs both the payoff-dominant and the risk-dominant conventions can be selected. The difference with [Goyal and Vega-Redondo \(2005\)](#) is mainly due to the fact that the transition from one convention to the other is stepwise, while in the non-cooperative setup it is all at once when a sufficient number of agents become mutants.

⁶In the main model of [Goyal and Vega-Redondo \(2005\)](#), instead, the payoff is earned by players on both sides of the link, independently of who is paying to maintain the link.

non-monomorphic states can be stochastically stable. More precisely, the co-existence of both payoff-dominant and risk-dominant conventions can be obtained in the long run. In this regard, [Anwar \(2002\)](#) studies a model where there are constraints on locations and each location has a certain number of patriots (i.e., agents who never want to leave their current location). The author shows that when the constraints on locations are tight (i.e., a small number of agents can/want to move) then the risk-dominant convention emerges in the long run, while if sufficient movement is possible across locations (capacity is large and/or patriots are few) then different conventions emerge at different locations. Further, in the case where the size of location is asymmetric, the location with the smaller size will have agents coordinating on the payoff-dominant convention.⁷

A similar model to that in [Anwar \(2002\)](#) is studied by [Blume and Temzelides \(2003\)](#), who also draw the conclusion that restricted mobility may lead to the coexistence of different conventions. In addition, the analysis of [Blume and Temzelides \(2003\)](#) allows to understand more clearly which are the long-run effects of increased mobility on the adopted convention: both a higher number of locations, and a higher fraction of agents that are mobile, have a positive impact on the likelihood to observe the emergence of the payoff-dominant convention at one location (with all mobile agents being at that location), as opposed to the risk-dominant convention at every location. Furthermore, [Blume and Temzelides \(2003\)](#) study the long-run payoffs of mobile and immobile agents as a function of the exogenous restrictions on mobility: they find that mobile agents enjoy a higher payoff and always benefit from increased mobility, while immobile agents benefit from increased mobility at low levels of mobility only.

We stress that we obtain the co-existence of conventions without such constraints on mobility. Instead, we rely on the cost of type mismatch (that requires some degree of agents heterogeneity) and on the risk of mismatch (that requires some degree of imperfect observability of types). We also stress that in our model, differently from [Anwar \(2002\)](#), when the total population size is large, the relative size of the two populations of types does not affect which population plays which convention, provided that the cost of mismatch is large enough.⁸

⁷See [Pin et al. \(2016\)](#) for a generalization to the case of more than two locations.

⁸Another relevant contribution is [Dieckmann \(1999\)](#), although it is in a sense less related because agents are supposed to follow imitation rules instead of myopic best reply rules. [Dieckmann \(1999\)](#) presents a location model where, besides capacity constraints, the movement across locations is subject to frictions (in the form of the possibility that only the action or only the location is revised as desired) and the play outside the current location is imperfectly observable. The main finding – substantially in line with [Anwar \(2002\)](#) and [Blume and Temzelides \(2003\)](#) – is that imperfect observability and frictions alone cannot block

Finally, [Carvalho \(2016\)](#) shows that the co-existence of conventions can emerge also under global random interaction, provided that the population is divided into cultural groups that have distinct taboos (some actions are never played voluntarily) and different cultural preferences (for the non-taboo actions). In our model too preferences are type-dependent, although in a substantially different way: heterogeneity is not with regard to actions but with regard to the cultural identity of the interacting partner.

2 The model

Network structure. We consider a set $N = \{1, 2, \dots, n\}$ of agents. Each $i \in N$ can choose the subset of other agents with whom to play a fixed bilateral social game. Formally, let $g_i = (g_{i1}, \dots, g_{in})$ be the n -dimensional vector collecting i 's connections; in particular, $g_{ij} \in \{0, 1\}$, and we say that agent i maintains a link with agent j if $g_{ij} = 1$. We assume that $g_{ii} = 0$ for every $i \in N$. Connections are directed, so that $g_{ij} = 1$ does not necessarily imply $g_{ji} = 1$. The maximum number of links that an agent can maintain at any given time is $k \geq 1$, and the cost to maintain any single link is c . An agent i is said to be *isolated* if $g_{ij} = 0$ for every $j \in N$. A profile of link formation choices, one for each agent in N , is denoted by $g = (g_1, g_2, \dots, g_n)$. We will refer to g as the network of interactions.

Social game. Agents play a 2×2 symmetric game in strategic form with common action set. Each agent plays only with the agents with whom he is directly connected.

The table below describes the payoffs associated to the bilateral social game. Payoffs are given only for the row player, since the game is symmetric.

	A	B
A	$\pi(A, A)$	$\pi(A, B)$
B	$\pi(B, A)$	$\pi(B, B)$

where the following inequalities hold:

- i. $\pi(A, A) > \pi(B, A)$ and $\pi(B, B) > \pi(A, B)$;

the emergence of the payoff-dominant convention, while restricted mobility does. In our model the imperfect observability (of types) can prevent the emergence of the payoff-dominant convention.

- ii. $\pi(A, A) < \pi(B, B)$;
- iii. $\pi(B, B) + \pi(B, A) < \pi(A, A) + \pi(A, B)$;

which imply that two coordination equilibria exist (by i.), B is the payoff-dominant action (by ii.) and A is the risk-dominant action (by iii.). We note that $\pi(B, A) < \pi(A, B)$. We further assume that all payoffs are positive, i.e., $\pi(B, A) > 0$.

Each agent $i \in N$ has to choose an action $a_i \in \{A, B\}$ which is played against the choice of each j for which $g_{ij} = 1$. A profile of action choices, one for each agent in N , is denoted by $a = (a_1, a_2, \dots, a_n)$.

Strategies. A strategy for agent i is $s_i = (a_i, g_i)$, where $a_i \in \{A, B\}$ denotes the action chosen by i , and $g_i \in \{0, 1\}^n$, with $g_{ii} = 0$ and $\sum_j g_{ij} \leq k$, denotes the agents with whom agent i is connected (see network structure and social game). A profile of strategies for the whole population – to which we also refer as *state* – is $s = (s_1, s_2, \dots, s_n)$. We will write $s = (a, g)$, where a is the action profile of the entire population and g is the interaction network.⁹

Agents' types. There are two types of agents in the population, x types and y types. The type of agent $i \in N$ is denoted by $w_i \in \{x, y\}$. With some abuse of notation, we use $\neg w_i$ to denote the type other than w_i . We also define the indicator function of type dissimilarity $\delta : N^2 \rightarrow \{0, 1\}$ such that $\delta(i, j) = 1$ if and only $w_i \neq w_j$, and $\delta(i, j) = 0$ otherwise. The number of agents of type x is denoted by n_x , while the number of agents of type y is denoted by $n_y = n - n_x$. Without loss of generality, we assume that $n_x \geq n_y$. We also assume that the number of agents of each type is large relative to the number of maximum connections for each agent; in particular, $n_y \geq 2k + 1$.

If agent i is of type w_i and interacts with agent j whose type is $w_j \neq w_i$, then i incurs a cost of $d > 0$; if instead $w_j = w_i$, then no cost is incurred.

Time. Agents repeatedly interact over time, which is assumed to be discrete, and indexed by the natural numbers, i.e., $t = 0, 1, 2, \dots$. The state of the system at time t is denoted with $s^t = (s_1^t, s_2^t, \dots, s_n^t)$, where s_i^t is the strategy adopted by agent i at time t .

⁹We note that in this setting the number of strategies available to an agent easily gets extremely large when n and k are not so small. For more details, we refer the interested reader to [Staudigl and Weidenholzer \(2014\)](#), where the same model of constrained network formation is presented and analyzed.

Revision protocol. In each round, every agent has the opportunity to revise his strategy with probability $\gamma \in (0, 1)$. When an agent receives a revision opportunity, he selects with positive probability any strategy that maximizes the interim utility, as formally provided in a subsequent paragraph.¹⁰

Information. If agent i receives a revision opportunity at time $t + 1$ then, for any agent j such that $g_{ij}^t = 1$, agent i knows both a_j^t and w_j , while for any agent h such that $g_{ih}^t = 0$, agent i knows only a_h^t . Also, agent i is informed of the summary statistics of the population state at time t ; in particular, he knows the number of x types and y types that are playing action A , as well as those who are playing action B , at state s^t . We stress, hence, that information about others' type accrues to an agent at two levels: locally, the agent knows the type of each of his neighbors, while, globally, he only knows the averages of x -type and y -type agents who are choosing A and B .

On the whole, a revising agent i at time $t + 1$ knows the pieces of information listed in the following table, where we also introduce specific notation to make the following exposition easier.

Notation	Number of agents...
$n(w, a s^t)$	of type w who are choosing action a
$n_{i1}(s^t)$	with whom agent i maintains a link
$n_{i0}(s^t)$	with whom agent i does not maintain a link
$n_{i1}(a s^t)$	playing action a with whom agent i maintains a link
$n_{i0}(a s^t)$	playing action a with whom agent i does not maintain a link
$n_{i1}(w s^t)$	of type w with whom agent i maintains a link
$n_{i0}(w s^t)$	of type w with whom agent i does not maintain a link
$n_{i1}(w, a s^t)$	of type w playing action a with whom agent i maintains a link
$n_{i0}(w, a s^t)$	of type w playing action a with whom agent i does not maintain a link

¹⁰We are aware that the details of the revision protocol may have important consequences, especially for the selection results obtained by stochastic stability. [Alós-Ferrer and Netzer \(2010\)](#), for instance, obtain that 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 \(2015\)](#).

Utilities. The interim utility of agent i who chooses strategy s_i at time $t + 1$, when the previous state was s^t , can be formally written as follows:

$$u_i(s_i^{t+1}, s^t) = \sum_{j \in N} \left[\underbrace{(\pi(a_i^{t+1}, a_j^t) - c) g_{ij}^{t+1}}_{\text{payoff of social game}} - \underbrace{d g_{ij}^t g_{ij}^{t+1} \delta(i, j)}_{\text{cost of type mismatch if old}} - \underbrace{d \frac{n_{i0}(\neg w_i, a_j^t | s^t)}{n_{i0}(a_j^t | s^t)} g_{ij}^{t+1} (1 - g_{ij}^t)}_{\text{expected cost of type mismatch if new}} \right].$$

The first term of $u_i(s_i^{t+1}, s^t)$ is the payoff from the social game net of the link maintenance cost, and it is made by the sum of payoffs accruing to i from the links he chooses to maintain at time $t + 1$. The second term is the cost of maintaining links at time $t + 1$ with agents of a type different from w_i and with whom i was already maintaining a link at time t ; we note that such a cost of type mismatch is known to i , since those agents are already his neighbors and hence he knows their type. The third term is the expected cost due to the formation of new links; differently from the previous term, these agents are new neighbors, and as such their type is unknown to i , who can only compute the expected cost of type mismatch on the basis of the fraction of agents of type different from w_i (i.e., $\neg w_i$) among the agents with whom i is not connected and who choose the observed action.¹¹

Once the decision is taken and the types of agents with whom new links are formed become known, agent i obtains the following ex-post utility at time $t + 1$:

$$v_i(s^{t+1}) = \underbrace{\sum_{j \in N} [\pi(a_i^{t+1}, a_j^{t+1}) - c] g_{ij}^{t+1}}_{\text{total payoff of social game}} - \underbrace{d n_{i1}(\neg w_i | s^{t+1})}_{\text{total cost of type mismatch}}.$$

The first term of $v_i(s^{t+1})$ is, again, the payoff from the social game net of the link maintenance cost, while the second term is the cost of maintaining links with agents of a type different from w_i , independently of whether i already had links towards them or not.

3 Unperturbed dynamics

Markov chain. The process described in Section 2 formally defines a Markov chain (S, P) , where S is the state space (i.e., the set containing all possible states) and P is the transition matrix, with $P_{ss'}$ denoting the probability to move directly from state s to state s' , with $s, s' \in S$.

¹¹We note that the number of agents with whom i is not connected is computed at time t , so that i 's choice to sever links at time $t + 1$ does not affect these numbers.

Neighborhood heterogeneity. We start by providing two results on the short run behavior of the system.

PROPOSITION 1. *If agent i is given a revision opportunity at time $t + 1$, and there exists $j \in N$ such that $g_{ij}^t = 1$, $\delta(i, j) = 1$, and $n(w_i, a_j^t | s^t) \geq k + 1$, then $g_{ij}^{t+1} = 0$ with probability 1.*

Proposition 1 states that a revising agent, say i , maintaining a link towards an agent of a different type, say j , will choose for sure to replace such a link with a new link towards an agent h choosing the same action as j , provided that at least some agent exists who chooses the same action as j and has the same type of i : indeed, h cannot grant a lower payoff than j , and grants a higher payoff with positive probability.

PROPOSITION 2. *If at time t there exist $i, j \in N$ such that $g_{ij}^t = 1$, $\delta(i, j) = 0$ and $a_i^t \neq a_j^t$, and if $c < \pi(B, A)$, $n_{i0}(\neg w_i, a_j^t | s^t) \geq 1$, and either $n_{i0}(a_i^t | s^t) = 0$ or $d \frac{n_{i0}(\neg w_i, a_i^t | s^t)}{n_{i0}(a_i^t | s^t)} > \pi(a_i^t, a_i^t) - \pi(a_i^t, a_j^t)$, then $g_{ij}^{t+1} = 1$ with probability 1.*

Proposition 2 states that a link between agents of the same type who play different actions can be stable, at least in the short run. Indeed, an agent, say i , will keep maintaining a link with another agent, say j , that plays a different action if the act of replacing j with some other agent h who plays i 's action implies a severe risk of type mismatch, i.e., if the penalty d is large enough and if the pool of potential new neighbors playing i 's action has a high enough fraction of agents who are not of i 's type.

The above two results can be understood intuitively as follows. A revising agent can condition his decision to form a link on the action of the new neighbor, but not on his type. As a consequence of this, a link with an agent choosing the same action has no value: such link can always be replaced by another link with an agent choosing the same action. Instead, a link with an agent of the same type has a value: if such link is replaced by a new link then the risk of a mismatch of types is sustained.

Absorbing sets. An absorbing set (otherwise called a recurrent class) is a minimal set of states with respect to the property that the system has probability zero of moving from any state in the set to any state out of the set.

We introduce a number of definitions that will be useful in the subsequent analysis, where we discuss the variety of possible absorbing sets. Consider a state $s = (a, g)$. State s is k -regular if $n_i = k$ for every $i \in N$. State s is *type-segregated* if $g_{ij} = 1$ implies that $\delta(i, j) = 0$. We say that state s is *monomorphic* if $a_i = a_j$ for all $i, j \in N$. In particular, we distinguish

between states that are *A-monomorphic* and states that are *B-monomorphic*, i.e., states where all agents choose *A* and states where all agents choose *B*, respectively. States that are not monomorphic are called *polymorphic*. Among polymorphic states, an important role in our analysis is played by type-monomorphic states; a state is called *type-monomorphic* if $a_i = a_j$ for all i, j such that $w_i = w_j = w \in \{x, y\}$, and it is not monomorphic. We sometimes refer to states that are *type-monomorphic with x on A and y on B* (which means that all agents of type x play A and all agents of type y play B), and states that are *type-monomorphic with x on B and y on A* (with an analogous definition). All remaining polymorphic states are called *type-polymorphic*. Finally, we denote with S_A^A the union of all absorbing sets that contain only *A-monomorphic* states. We define S_B^B analogously. We use S_B^A to denote the union of absorbing sets that contain only type-monomorphic states with x on A and y on B ; the set S_A^B is defined analogously, with the general rule that the apex refers to the choice of agents of type x and the pedix refers to the choice of agents of type y .

PROPOSITION 3.

- (a) *If $c < \pi(A, A)$, then (i) there exist absorbing sets containing *A-monomorphic* states, and (ii) there exist absorbing sets containing *B-monomorphic* states.*
- (b) *If $c < \pi(A, A)$ and $d > \pi(B, B) - \pi(A, A)$, then (i) there exist absorbing sets containing type-monomorphic states with x on A and y on B , and (ii) there exist absorbing sets containing type-monomorphic states with x on B and y on A .*
- (c) *If $c < \pi(A, A)$ and $d > 2[\pi(B, B) - \pi(A, A)]$, then there exist absorbing sets containing type-polymorphic states.*
- (d) *If $c < \pi(A, A)$ and $d > \pi(B, B)(n - 1)/n_x$, then there exist absorbing sets containing states where some agent is isolated.*

Proposition 3 allows us to make a remark about the effect that the introduction of a payoff relevant heterogeneity in types has on the possibility to observe the coexistence of both actions in the long run.¹² If all agents had the same type, or if in our model we set $d = 0$ so that types are payoff irrelevant, essentially we would obtain the model of [Staudigl and Weidenholzer \(2014\)](#), where the only absorbing sets (provided that c is not too large) are those in which all agents choose the same action. This can be easily understood by





¹²We stress that the conditions provided are only sufficient. Since this result does not play an important role for the following investigation of stochastic stability, we have chosen not to give tighter conditions, which would have required a more cumbersome analysis.

noting that, in such a case, if it is optimal for an agent to maintain his current action, then it is optimal for agents using a different action to switch to this one. By introducing agents of different types, and unobservability of types prior to linking, we obtain that agents at different locations in the network are different because they have different local information about the types of their neighbors. Such information is valuable if types are payoff relevant, i.e., when $d > 0$. This leads to a substantially richer variety of states belonging to absorbing sets. Monomorphic states still belong to absorbing sets (point (a)). In addition to them, when d is sufficiently large we find absorbing sets containing polymorphic states. In particular we have that both type-monomorphic states and type-polymorphic states can belong to absorbing sets (points (b) and (c), respectively). This leads to the following observation: even if interactions between agents of different types are unlikely to last for a long time (see Proposition 1), so that the two populations of agents may apparently live totally independent lives, they in fact influence each other in an indirect way. Finally, we have that, when d is large enough, there are absorbing sets containing states where some agent is isolated, even if c is not very large (point (d)). More importantly, as it will become clearer in the following examples, there is the possibility that some agents are isolated and others are not. We stress that this heterogeneity in the network structure can not be observed if isolation is obtained by raising c .

Graphical illustrations. In Fig. 1–4 we provide some graphical representations with the aim of illustrating the variety of possible absorbing sets. In all the examples, we have $n_x = 10$, $n_y = 9$, $k = 3$. Circles identify x types, squares identify y types. Arrows represent interactions. Agents choosing action A are colored in light green, agents choosing B are colored in dark blue.

Fig. 1 depicts monomorphic states; more precisely, in Subfig. 1a we have an A -monomorphic state, and in Subfig. 1b we have a B -monomorphic state. Also, the subfigures represent states that are k -regular, since every agent has exactly 3 connections (in particular, the same connections are in place in the two states). Consistent with point (a) of Proposition 3, such states belong to absorbing sets for any value of d . In particular, each state belongs to a singleton absorbing set: as long as d is positive, every agent prefers not to reshuffle his links, in order to avoid the risk of a type mismatch.





Fig. 2 depicts type-monomorphic states; more precisely, in Subfig. 2a we have a type-monomorphic state with x on A and y on B , and in Subfig. 2b we have a type-monomorphic state with x on B and y on A . Both states are k -regular; this must necessarily be the case, since there is no risk of mismatch due to the perfect correlation between actions and types.



(a) An A -monomorphic state.

(b) A B -monomorphic state.

Figure 1: Examples of monomorphic states.



(a) A type-monomorphic state with x on A and y on B .

(b) A type-monomorphic state with x on B and y on A .

Figure 2: Examples of type-monomorphic states.

We observe that each of these states belongs to an absorbing set comprising many states; this is so because agents are indifferent between keeping current mates and substituting them with other agents choosing the same action. In order for type-monomorphic states to belong to an absorbing set, we must have that agents choosing A do not find it profitable to switch to B and cast links to agents choosing B ; this happens when $d > \pi(B, B) - \pi(A, A)$, which is the same inequality given in point (b) of Proposition 3.




Figure 3: A type-polymorphic state.

Fig. 3 represents a state that is type-polymorphic. In particular, there are agents of each type choosing A , and agents of the same type choosing B . It is easy to understand that, given the current network of interactions, no agent wants to change action. Also, agents who choose B will never reshuffle links, because there is a risk of type mismatch if doing so, while there is no benefit. Agents who choose A face an expected penalty due to type mismatch that is equal to $d/2$; if such a cost is larger than $\pi(B, B) - \pi(A, A)$, which is the same inequality given in point (c) of Proposition 3, then these agents will never change strategy as well, so that the state in Fig. 3 actually belongs to a singleton absorbing set.

We stress that, even if the state under consideration is k -regular, there are type-polymorphic states belonging to absorbing sets where this is not the case (the same can occur for monomorphic states, while type-monomorphic states are necessarily k -regular). Imagine that, starting from the state represented in Fig. 3, a link is removed. Intuitively, if d is large enough, then the expected penalty of a type mismatch is sufficiently large to discourage any attempt to form a new link. Finally, a link between two agents of the same type who choose different actions might also be in place, in case such an interaction brings a positive payoff (which happens if c is not that large), and d is again sufficiently large (see Proposition 2).

Fig. 4 depicts a state where one agent is isolated. In particular, agent i has no link outgoing from him. Furthermore, agent i plays action A , while all other agents (including




Figure 4: A state where an agent is isolated.

those having the same type as i) play B , so the state represented is type-polymorphic. If the expected penalty for i of a new link towards an agent playing B (which is equal to $10d/18$) is larger than the largest benefit coming from the new interaction (which is equal to $\pi(B, B)$) then casting a new link is unprofitable. We observe that the arising inequality is the same as the inequality in point (d) of Proposition 3, once we consider that $n_x = 10$ and $n_y = 9$. We also observe that agent i will keep on switching from action B to A and vice versa, since both actions grant him the same (null) utility. Moreover, all other agents strictly prefer not to change their strategies, since they currently earn the maximum attainable utility and reshuffling links comes with the risk of a type mismatch. Therefore, the state in Fig. 4 actually belongs to an absorbing set. We finally notice that agent i has no incoming links. We remark that this is something specific to this example; indeed, there exist states belonging to absorbing sets where an agent exists who has some incoming links and no outgoing links (so that such an agent is still isolated).

4 Perturbed dynamics

Regular perturbed Markov chain. We are ready to introduce perturbations in the unperturbed dynamics considered in Section 3, and to apply concepts and tools developed by Foster and Young (1990), Young (1993), Kandori et al. (1993) and Ellison (2000).

We adopt the so-called uniform error model for mistakes.¹³ In particular, when an agent

¹³It is well known that the way in which perturbations are modeled has, in general, important consequences on the outcome of the perturbed dynamics (see Bergin and Lipman, 1996). While the uniform error model is a specific way among many possibilities, we stress that it is very common in the literature (starting

is given a revision opportunity, with probability $1 - \epsilon$ he will update his strategy by using the myopic best-reply rule described in the previous section, while with probability ϵ the agent is hit by a perturbation (or mutation, mistake, etc.) and chooses at random one strategy in his strategy set. The arising transition matrix is denoted with P^ϵ , and we refer to (S, P^ϵ) as the perturbed Markov chain resulting from (S, P) . For any positive level of ϵ , the system can move with positive probability from any state to any other state, i.e., it is ergodic. This implies that the perturbed Markov chain is irreducible and aperiodic and hence, by known results, there exists a unique invariant distribution μ^ϵ over states in S that describes the long-run behavior of the system. As ϵ tends to zero, we have that P^ϵ tends to P ; in particular, $P_{ss'}^\epsilon \sim \epsilon^{r(s,s')}$ as $\epsilon \rightarrow 0$, where $r(s, s')$ is the so-called *resistance* of the transition from s to s' , which basically counts how many perturbations (or mutations, mistakes, etc.) are required to complete such a transition in one period of time. A family of perturbed Markov chains for ϵ going to zero which satisfies the above properties is called a regular perturbed Markov chain. For a regular perturbed Markov chain, the limit of the invariant distribution μ^ϵ for ϵ going to zero is known to exist, and the states having positive probability in that limiting distribution are called *stochastically stable*. The following characterization of stochastically stable states will be useful for our subsequent analysis (see [Young, 2001](#), for a more detailed exposition).

The notion of resistance can be extended by relaxing the constraint that the transition must occur in one period, and can be usefully applied to absorbing sets instead of states. Given two absorbing sets S' and S'' , the resistance between S' and S'' is given by the minimum sum of resistances between states over paths that start in a state belonging to S' and end in a state belonging to S'' . Now, for any conceivable tree (i.e., a graph such that any two vertices are connected by exactly one path) having the absorbing set S' as root and all absorbing sets as nodes, consider the sum of resistances assigned to each edge of the tree, and take the minimum over trees of such a sum. This number represents the *stochastic potential* of S' . Intuitively, the stochastic potential tells us how difficult is to reach an absorbing set starting from other absorbing sets. A fundamental result in this literature asserts that a state is stochastically stable if and only if it belongs to an absorbing set with minimum potential: stochastically stable states are those that are relatively easiest to reach in terms of the minimum number of mutations required to reach such states starting from other states.

Two other notions are useful in the following analysis: the radius and coradius ([Ellison](#), from [Young, 1993](#) and [Kandori et al., 1993](#)), despite the fact that recent experimental evidence points to alternative error models ([Mäs and Nax, 2016](#); [Lim and Neary, 2016](#)).

2000). If Q is a union of absorbing sets, consider all possible paths – i.e., sequences of states – starting from a state in Q and ending in a state belonging to an absorbing set that is not part of Q . The radius of Q , denoted with $R(Q)$, is defined as the minimum sum of resistances between states over all such paths. Now consider all possible paths starting from a state belonging to an absorbing set Q' and ending in a state in Q . For each Q' , consider the minimum sum of resistances between states over all such paths. The coradius of Q , denoted with $CR(Q)$, is the maximum among Q' of such minimum numbers. Intuitively, $R(Q)$ and $CR(Q)$ provide measures of how difficult it is, respectively, to leave Q and to reach Q .

Stochastic stability: low cost of mismatch in types. Our first main result on stochastic stability is Proposition 4, and it addresses the case where the cost of interacting with an agent of a different type is low relative to the gain of coordinating on the payoff-dominant action instead of the risk-dominant one. Preliminarily, Lemma 1 provides a characterization of the set S_B^B that is then exploited in the proof of the proposition.

LEMMA 1 (Characterization of S_B^B). *If $c < \pi(A, A)$, $d < \pi(B, B) - \pi(A, A)$, then a state $s \in S_B^B$ if and only if (i) s is B -monomorphic, (ii) s is k -regular, and (iii) s is type-segregated. Moreover, every $s \in S_B^B$ belongs to a singleton absorbing set.*

Lemma 1 states that S_B^B is made of the union of all and only the singleton absorbing sets where each agent plays B and has k links towards agents of his own type.

PROPOSITION 4. *If $c < \pi(A, A)$, $d < \pi(B, B) - \pi(A, A)$, then a state s is stochastically stable if and only if $s \in S_B^B$.*

Proof. We remind that S_B^B is defined as a union of absorbing sets. We then compute its radius and coradius, denoted with $R(S_B^B)$ and $CR(S_B^B)$, respectively. We first observe that in every state where at least $k + 1$ agents choose B , then all agents find it optimal to choose B and have k connections to agents choosing B : indeed, such a strategy grants in the worst case of all mismatches a payoff of $k(\pi(B, B) - c - d)$, while the highest payoff with any other strategy is $k(\pi(A, A) - c)$, which is lower by the assumption $d < \pi(B, B) - \pi(A, A)$. So, starting from a state in S_B^B , at least $n - k$ mutations must occur to switch $n - k$ agents from B to A , hence leaving less than $k + 1$ agents choosing B . Our assumptions guarantee that $n \geq 4k + 2$. This implies that $R(S_B^B) \geq 3k + 2$.

We now suppose to start from a state outside S_B^B . With k mutations, we are sure to reach a state where at least k agents choose B . In such a state, all other agents find it optimal to choose B and connect to agents choosing B (for the same reasons discussed above).

This means that a state in S_B^B can be reached with positive probability in the unperturbed dynamics, thus implying that $CR(S_B^B) \leq k$. Since $R(S_B^B) > CR(S_B^B)$, we can conclude by Theorem 1 of Ellison (2000) that all stochastically stable states belong to S_B^B .

The last step is to show that, for any two states $s = (a, g), s' = (a', g') \in S_B^B$, there exists a sequence of states belonging to absorbing sets $s_1, \dots, s_i, \dots, s_\ell$ such that $s_1 = s, s_\ell = s'$, and a single mutation is sufficient to move from s_i to s_{i+1} for $i = 1, \dots, \ell - 1$. If, for agent i , we have that $s_i \neq s'_i$, then a single mutation can change s_i to s'_i . The state reached in this way forms an absorbing set by Lemma 1, since all agents play B , it is k -regular and type-segregated. With at most n of such steps, we are sure to have reached state s' . Therefore, S_B^B is a mutation-connected component (in the words of Samuelson, 1994), and we apply Theorem 2 in that paper to conclude that all states in S_B^B are stochastically stable. \square

Let us make a remark to better contrast our results with those in Staudigl and Weidenholzer (2014). In our model, the assumption that $n_x \geq n_y \geq 2k+1$ implies that $k < (n-1)/2$. In the model of Staudigl and Weidenholzer (2014) (see Theorem 1) this condition guarantees that the payoff-dominant convention is the stochastically stable outcome. In this respect, our Proposition 4 represents a robustness check of their result: the introduction of types that are not globally observable and that determine a penalty in case of a mismatch does not affect the long-run prediction in favor of the payoff-dominant convention, provided that the penalty for a mismatch in types is sufficiently low. The only difference between our prediction and theirs concerns the shape of the interaction network: we obtain that stochastically stable states are type-segregated, while this feature is clearly absent in Staudigl and Weidenholzer (2014). This is something expected – especially in monomorphic states – as the interaction between agents of different types bears a cost.

Stochastic stability: high cost of mismatch in types. The prediction obtained by stochastic stability drastically changes when d is high. Surprisingly, neither the payoff-dominant convention nor the risk-dominant convention is selected in this case. Rather, we obtain that both actions will co-exist in the long run.

The proof of the result relies on a tree surgery argument, exploiting the techniques by Young (1993). Basically, we will show that, starting from any absorbing set that is not included in $S_B^A \cup S_A^B$, we can build sequences of absorbing sets ending in S_B^A and S_A^B such that each step has resistance 1. Then, we show that the radius of S_B^A and the radius of S_A^B are larger than 1, which implies that all other absorbing sets have a stochastic potential that is higher than S_B^A and S_A^B . The proof of Proposition 5, where our main result is provided, makes use of the preliminary results stated in Lemmas from 2 to 6.

We start by providing a characterization of S_B^A and S_A^B . In particular, Lemma 2 states that S_B^A and S_A^B are two absorbing sets, each made of all and only the states where all agents of one type play an action, all agents of the other type play the other action, and every agent has k links towards agents of his own type.

LEMMA 2 (Characterization of S_B^A and S_A^B). *If $c < \pi(A, A)$, $d > \pi(B, B)\frac{n-1}{n_y}$, then:*

- (a) *a state $s \in S_B^A$ if and only if (i) s is type-monomorphic with x on A and y on B , (ii) s is k -regular, and (iii) s is type-segregated;*
- (b) *a state $s \in S_A^B$ if and only if (i) s is type-monomorphic with x on B and y on A , (ii) s is k -regular, and (iii) s is type-segregated.*

Moreover, S_B^A and S_A^B are two absorbing sets.

With the following lemma we establish that at least 2 mutations are required to leave S_B^A (or S_A^B) reaching, with positive probability, another absorbing set. Moreover, we establish that if the population of agents of each type is sufficiently large, then the number of mutations to leave S_B^A is equal to the number of mutations to leave S_A^B . We draw attention to note that the result holds when $k \geq 2$,¹⁴ and this limits the applicability of Proposition 5 as well.

LEMMA 3 (Away from S_B^A and S_A^B). *If $c < \pi(A, A)$, $d > \pi(B, B)\frac{n-1}{n_y}$, and $k \geq 2$, then $R(S_B^A) \geq 2$ and $R(S_A^B) \geq 2$. If, additionally, $n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}$, then $R(S_B^A) = R(S_A^B)$.*

By Lemma 4 we establish that, starting from any state in an absorbing set, there exists a sequence of absorbing sets that reaches an absorbing set belonging to $S_A^A \cup S_B^B \cup S_B^A \cup S_A^B$, with each step in the sequence requiring a single mutation.

LEMMA 4 (Towards $S_A^A \cup S_B^B \cup S_B^A \cup S_A^B$). *If Q is an absorbing set, then there exists a sequence of absorbing sets Q_1, \dots, Q_ℓ , where $Q_1 = Q$ and $Q_\ell \subseteq S_A^A \cup S_B^B \cup S_B^A \cup S_A^B$, such that a single mutation is sufficient to move from Q_i to Q_{i+1} for $i = 1, \dots, \ell - 1$.*

¹⁴We are indebted to an anonymous referee who has provided a simple example where $k = 1$ and a single mutation is sufficient to leave S_B^A and S_A^B : think of all agents of one type as located around a circle, with each agent who is linked only to the agent on his right; they are all playing A , and a mutation hits an agent so that his action changes from A to B while his neighbor remains the same; now, the agent on his left, if given the possibility, would switch from A to B ; after that, also the agent on the left of the latter agent would like to switch from A to B , and so on; this clockwise path of adjustments, which occurs with positive probability, ends in a state belonging to S_B^B .

Lemma 5 clarifies how the system can pass from S_B^A and S_A^B to S_B^B with a minimum number of mutations; in particular, it shows the existence of a sequence of absorbing sets that, starting from S_B^A (respectively, S_A^B), involves $R(S_B^A)$ (respectively, $R(S_A^B)$) mutations to do the first step, and then proceeds with all steps requiring a single mutation until an absorbing set in S_B^B is reached.

LEMMA 5 (From S_B^A to S_B^B and from S_A^B to S_B^B). *If $Q = S_B^A$ or $Q = S_A^B$ and $Q' \subseteq S_B^B$ are absorbing sets, then:*

- (a) *there exists a sequence of absorbing sets Q_1, \dots, Q_ℓ , where $Q_1 = Q$ and $Q_\ell = Q'$, such that $R(S_B^A)$ mutations are required to move from Q_1 to Q_2 , and a single mutation is sufficient to move from Q_i to Q_{i+1} for $i = 2, \dots, \ell - 1$;*
- (b) *there also exists a sequence of absorbing sets Q'_1, \dots, Q'_ℓ , where $Q'_1 = Q'$ and $Q'_\ell = Q$, such that $R(S_A^B)$ mutations are required to move from Q'_1 to Q'_2 , and a single mutation is sufficient to move from Q'_i to Q'_{i+1} for $i = 2, \dots, \ell - 1$.*

Finally, the following lemma provides the last result to be used in the proof of Proposition 5: for any absorbing set belonging to $S_A^A \cup S_B^B$, we can find a sequence of absorbing sets where each step of the sequence has resistance 1, which originates from the absorbing set under consideration and reaches S_B^A (or S_A^B).

LEMMA 6 (From $S_A^A \cup S_B^B$ to S_B^A and to S_A^B). *If $Q \subseteq S_A^A \cup S_B^B$ and $Q' = S_B^A$ or $Q' = S_A^B$ are absorbing sets, then there exists a sequence of absorbing sets Q_1, \dots, Q_ℓ , where $Q_1 = Q$ and $Q_\ell = Q'$, such that a single mutation is sufficient to move from Q_i to Q_{i+1} for $i = 1, \dots, \ell - 1$.*

We are now ready to state and to prove our main result.

PROPOSITION 5. *If $c < \pi(A, A)$, $d > \pi(B, B) \frac{n-1}{n_y}$, and $k \geq 2$, then every stochastically stable state is contained in $S_A^B \cup S_B^A$. If, additionally, $n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}$, then a state s is stochastically stable if and only if $s \in S_A^B \cup S_B^A$.*

Proof. By Lemma 2 we know that S_B^A and S_A^B are two absorbing sets, and by nesting Lemma 4 with Lemma 5 and with Lemma 6 we can find, starting from any absorbing set, a path between absorbing sets that leads to S_B^A and such that every step in the path involves only 1 mutation, except if the path goes through S_A^B , which instead requires $R(S_A^B)$ mutations to leave. For every absorbing set, take one such path with the absorbing set as the starting point, and consider the arising directed graph; for every absorbing set that has more than

one out-going link, all such links but one are deleted; by doing so, we are able to construct an S_B^A -tree over absorbing sets where every link between any two nodes involves only 1 mutation, except for the link out-going from S_A^B , which involves $R(S_A^B)$ mutations.

We observe that 1 is the minimum number of mutations required to move between absorbing sets, and $R(S_A^B)$ is the minimum number of mutations required to exit S_A^B , by Lemma 3. Therefore, the stochastic potential of S_B^A is equal to $\xi - 2 + R(S_A^B)$, where ξ denotes the total number of absorbing sets. With an analogous reasoning we obtain that the stochastic potential of S_A^B is equal to $\xi - 2 + R(S_A^A)$.

Now we consider an absorbing set $Q \not\subseteq S_B^A \cup S_A^B$. We observe again that 1 is the minimum number of mutations required to move between absorbing sets, $R(S_A^B)$ is the minimum number of mutations required to exit S_A^B , and $R(S_B^A)$ is the minimum number of mutations required to exit S_B^A . Therefore, the stochastic potential of Q cannot be lower than $\xi - 3 + R(S_B^A) + R(S_A^B)$, which is higher than the stochastic potentials of S_B^A and S_A^B , because $R(S_B^A) \geq 2$ and $R(S_A^B) \geq 2$ by Lemma 3. This allows us to conclude, by Theorem 2 in Young (1993), that every stochastically stable state is contained in $S_B^A \cup S_A^B$. In case the additional assumption $n_y > \frac{dk}{\pi(B, B) - \pi(B, A)}$ is satisfied, then the stochastic potentials of S_B^A and S_A^B are equal, and hence we can conclude, again by Theorem 2 in Young (1993), that a state s is stochastically stable state if and only if $s \in S_B^A \cup S_A^B$. \square

The results stated in Proposition 5 can be better understood if contrasted with the results provided in Anwar (2002) and Blume and Temzelides (2003). As mentioned in the Introduction, Anwar (2002) and Blume and Temzelides (2003) study models where agents interact at specific locations in the presence of restrictions to mobility; the main finding is that the risk-dominant and the payoff-dominant convention can co-exist in the long run, provided that there is enough freedom of mobility across locations. However, freedom of mobility must not be too much – i.e., some agents are not allowed to go to the location that they prefer – otherwise only the payoff-dominant convention emerges – as shown by, e.g., Bhaskar and Vega-Redondo (2004). So, co-existence can be understood as the result of imperfections – or frictions – that makes this case lie in between absence of mobility – when the risk-dominant convention emerges in the log-run – and full mobility – when the payoff-dominant convention does emerge.

On the contrary, in the present model the co-existence of distinct conventions is the consequence of a new effect that favors directly type-monomorphic states, rather than mediating between two extreme cases where the two conventions are globally adopted. Type-monomorphic states are indeed the only states where the information on types can be per-

fectly derived from the observation of actions, allowing a mutant to come back with certainty to having interactions with agents of own type; this makes type-monomorphic states relatively more resilient to mutations.

5 Discussion

Network formation model. The non-cooperative model of network formation that we have considered is non-cooperative and asymmetric in both cost bearing and payoff flows (fundamentally, the one-way flow model in [Bala and Goyal, 2000](#)). The assumption that connections are formed unilaterally can be realistic for some cases, while for other cases a cooperative model of network formation might be more appropriate – like that in [Jackson and Watts \(2002\)](#), where the consent of both agents involved is needed to form a link. Further, the assumption that payoffs flow unilaterally to the agent who has established the link – and bears the maintenance cost – is another feature of our model that may fit some cases, but not others. Both assumptions are also found in [Staudigl and Weidenholzer \(2014\)](#), which is the natural benchmark against which to compare our results. In their analysis, as in ours, the two assumptions described above allow a reasonable treatment of the model that, otherwise, would have been by far more complicated – and this, admittedly, is the main reason why we adopt them.

However, we do want to stress that the main intuition underlying our results does not depend crucially on the details of the network formation model. Indeed, it is a general observation that, when types are not observable outside one’s own neighborhood and interactions with different types are costly, there emerges an implicit cost of adding a new neighbor. This is due to the risk of linking to an agent of a type different from one’s own; and such a cost can be large enough to prevent an agent – who has been hit by a mutation – from going back to his status quo. This allows paths made of single mutations to move the system from one absorbing set to another absorbing set, making it rather easy to exit most absorbing sets – including the ones with the payoff-dominant convention and the risk-dominant convention. Similarly, we observe that there is only one situation where such a risk is totally absent: when there is perfect correlation between actions and types, i.e., in type-monomorphic states; indeed, if we start from a type-monomorphic state then going back to the status quo after a single mutation comes at no implicit cost and, hence, it is quite more difficult to leave a type-monomorphic state. This is the core reason why the absorbing sets collecting type-monomorphic states are stochastically stable, and in this argument there is no substantial role played by a specific network formation model (even if

details can well and substantially affect transition periods and patterns).

Type unobservability and homophily. The ingredients which are crucial in the above reasoning are related to two distinct aspects: the unobservability of the types of agents with whom no link is currently maintained, and the magnitude of the cost that has to be incurred to interact with an agent of a different type. The fact that types are unobservable prior to interaction seems a plausible assumption in many situations, at least if we think of the type as private piece of information which is learned only after interaction – possibly inferring the type from the payoff earned (which seems natural if type is a payoff relevant characteristic).¹⁵ The cost to be incurred for type mismatch can be seen as a form of homophily, i.e., as the result of preferences for interacting with one’s own type. In this regard, we observe that our model exhibits homophily in the long run, and this is less obvious than it might appear at first because agents cannot choose directly to interact with agents of similar type – exactly because types are not observable prior to interaction. In this sense, the present paper can be seen as a marginal contribution to the recent literature on homophily in social interactions (see Currarini et al., 2009 and Bramoullé et al., 2012): social coordination plus weak homophily (i.e., d is small enough) leads to the global adoption of the payoff-dominant convention (like in Staudigl and Weidenholzer, 2014) with the additional feature that interactions take place only between agents of the same type (see Proposition 4 and comments thereafter), while if homophily is strong enough (i.e., d is large enough) in the long-run we observe segregation in both types *and* actions (see Proposition 5). As a final remark on the cost of type mismatch, we observe that such cost can well be interpreted as a part of the link maintenance cost. With this interpretation, while a connection with a agent of similar type has a cost of c , a connection with an agent of different type has a cost of $c + d$, where d measures some additional type-related cost of interaction (e.g., extra communication costs due to the use of a different language).

Lack of memory. There is another assumption of our model that is crucial for our results, as anticipated in the Introduction: the agents’ lack of memory. The important consequence of the lack of memory is that agents who are hit by a mutation, and change the network of interactions, cannot choose to go back and connect with the same agents with whom they

¹⁵We observe that our main results could be extended with few adjustments to a model where agents also observe the types of those they are passively linked with; intuitively, our arguments in the proofs remain substantially valid provided that, when constructing outgoing paths from absorbing sets, we make sure that the mutants do not have any incoming links, which can be obtained easily (at most with sequences where each step requires a single mutation) since populations have been assumed large enough (i.e., $k < (n_y - 1)/2$).

were previously connected. In fact, if this were possible, there would be no risk of type mismatch when an agent is hit by a mutation and considers to go back to his previous status quo – i.e., to undo what a mutation has done. So, without this assumption our results would not be warranted. However, we think that assuming agents’ lack of memory does fit some relevant cases. For instance, it fits the case where the effects of a mutation persist over several periods and the mutants can receive the opportunity to readjust their strategy only after some time, but in the meanwhile old neighbors might have become untraceable or unrecognizable. In general, we can think of the lack of memory as not only due to agents’ cognitive bounds, but also to the actual technology of interactions which might make it impossible to retain with certainty the access to old neighbors (for instance, this is something quite widespread in social interactions over the internet). Most importantly, as stressed in the Introduction, what is really required for our main results is a weak version of the lack of memory, i.e., a positive probability that agents are not able to trace back their old neighbors once they disconnect from them.¹⁶

Extensions. We end this discussion with a couple of remarks concerning the coordination game. First, we observe that the mechanism behind our results does not rest crucially on the fact that one action is payoff-dominant and the other action is risk-dominant. What is important in order to establish that absorbing sets made of type-monomorphic states are more resilient to mutations – i.e., that at least two mutations are required to leave these absorbing set (see Lemma 3) – is that, once a mutation has turned a single agent from A to B , the agents who are connected to him must not find it profitable to switch from A to B . In our model this condition is implied by the fact that A is risk-dominant (we remind also that in S_B^A and S_A^B all agents are k -regular with agents of their own type). However, if action B were both payoff-dominant and risk-dominant, and provided that at least 2 out of k neighbors must play B to have that playing B is better than playing A , then stochastic stability would still select, under the assumptions of Proposition 5, states where agents of one type coordinate on B , and agents of the other type coordinate on A .

Moreover, the inspection of the mechanism driving our results makes us confident that similar conclusions hold in more general coordination games, and with more than 2 types of agents. Our belief is that in any coordination game with a number of actions greater than or equal to the number of types, if the cost of interacting with any type different from one’s own is large enough, then different types will coordinate on different actions. Indeed, states

¹⁶This is reminiscent of the results found by Horvath et al. (2012) for the case of prisoner dilemma played on a network, where a too long memory could disrupt superior equilibria.

where at least one action is played by agents of two different types admit paths leading to new absorbing sets which are made of steps involving a single mutation, on the ground of the same intuition used in this paper: going back to the status quo after a mutation involves the risk of a type mismatch for such agents. But this does not hold in states where types and actions are perfectly correlated, and hence such states are more resilient to mutations; of course, the existence of states where types and actions are perfectly correlated requires that the existence of at least as many actions as types.

Another dimension along which our results could be extended concerns the error model, even if a more careful analysis is required because not all mutations receive the same weight once we abandon the uniform error model.¹⁷

We think that these observations reinforce the main message that can be drawn from our contribution: in a setting where a population of agents has to form interactions and coordinate on some action, if agents differ for some unobservable characteristic and interactions between agents with dissimilar characteristics are costly enough, then actions will end up playing the role of signals, allowing the formation of clusters of agents who are type-homogeneous, each cluster coordinating on a different action.

References

- 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 F. Shi (2012). Imitation with asymmetric memory. *Economic Theory* 49(1), 193–215.
- Alós-Ferrer, C. and S. Weidenholzer (2007). Partial bandwagon effects and local interactions. *Games and Economic Behavior* 61(2), 179–197.
- 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.
- Bala, V. and S. Goyal (2000). A noncooperative model of network formation. *Econometrica* 68(5), 1181–1229.
- Bergin, J. and B. L. Lipman (1996). Evolution with State-Dependent Mutations. *Econometrica* 64, 943–956.

¹⁷See [Hwang et al., 2016](#) for a study on the relationships between error models and long-run conventions in a class of coordination games, and [Hwang and Newton, 2016](#) for a result that allows a simplified analysis when mistakes are payoff-dependent.

- Bhaskar, V. and F. Vega-Redondo (2004). Migration and the evolution of conventions. *Journal of Economic Behavior & Organization* 55(3), 397–418.
- 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. (1995). The statistical mechanics of best-response strategy revision. *Games and Economic Behavior* 11(2), 111–145.
- Blume, L. E. (1997). Population games. In D. S. L. D. Arthur, B. (Ed.), *The economy as an evolving complex system II*, pp. 425–460. Reading, MA: Addison-Wesley.
- Bramoullé, Y., S. Currarini, M. O. Jackson, P. Pin, and B. W. Rogers (2012). Homophily and long-run integration in social networks. *Journal of Economic Theory* 147(5), 1754–1786.
- Carvalho, J.-P. (2016). Coordination and culture. *Economic Theory*, 1–27.
- Cui, Z. (2014). More neighbors, more efficiency. *Journal of Economic Dynamics and Control* 40, 103–115.
- Currarini, S., M. O. Jackson, and P. Pin (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica* 77(4), 1003–1045.
- Dieckmann, T. (1999). The evolution of conventions with mobile players. *Journal of Economic Behavior & Organization* 38(1), 93–111.
- 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).
- Eshel, I., L. Samuelson, and A. Shaked (1998). Altruists, egoists, and hooligans in a local interaction model. *American Economic Review*, 157–179.
- Foster, D. and P. Young (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology* 38(2), 219–232.
- Goyal, S. and M. C. Janssen (1997). Non-exclusive conventions and social coordination. *Journal of Economic Theory* 77(1), 34–57.
- 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.

- Hojman, D. A. and A. Szeidl (2006). Endogenous networks, social games, and evolution. *Games and Economic Behavior* 55(1), 112–130.
- Horvath, G., J. Kovářík, and F. Mengel (2012). Limited memory can be beneficial for the evolution of cooperation. *Journal of Theoretical Biology* 300, 193–205.
- Hwang, S.-H., W. Lim, P. Neary, J. Newton, et al. (2016). Conventional contracts, intentional behavior and logit choice: equality without symmetry. Technical report.
- Hwang, S.-H. and J. Newton (2016). Payoff-dependent dynamics and coordination games. *Economic Theory*, 1–16.
- 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.
- Jiang, G. and S. Weidenholzer (2016). Local interactions under switching costs. *Economic Theory*, 1–18.
- 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.
- Lim, W. and P. R. Neary (2016). An experimental investigation of stochastic adjustment dynamics. *Games and Economic Behavior* 100, 208–219.
- Mäs, M. and H. H. Nax (2016). A behavioral study of noise in coordination games. *Journal of Economic Theory* 162, 195–208.
- Neary, P. R. (2012). Competing conventions. *Games and Economic Behavior* 76(1), 301–328.
- 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.
- Pin, P., E. Weidenholzer, and S. Weidenholzer (2016). Constrained mobility and the evolution of efficient outcomes. *mimeo*.
- Samuelson, L. (1994). Stochastic stability in games with alternative best replies. *Journal of Economic Theory* 64(1), 35–65.
- Staudigl, M. and S. Weidenholzer (2014). Constrained interactions and social coordination. *Journal of Economic Theory* 152, 41–63.
- Sugden, R. (1995). The coexistence of conventions. *Journal of Economic Behavior & Organization* 28(2), 241–256.

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.

Appendix - Proofs of Propositions and Lemmas

Proof of Proposition 1

Proof. Suppose that agent i is given a revision opportunity at time $t + 1$, and suppose ad absurdum that there exists a strategy $s_i = (a_i, g_i)$ that maximizes the interim utility of i and tells him to maintain the link with agent j , i.e., $g_{ij} = 1$. We construct another strategy $s'_i = (a'_i, g'_i)$ such that $a'_i = a_i$ and g'_i is equal to g_i with the only difference that in g'_i the link with agent j is removed and a new link is formed with an agent ℓ such that $a_\ell^t = a_j^t$. We observe that the assumption that $n(w_i, a_j^t | s^t) \geq k + 1$ implies that there exists at least an agent having the same type as i and choosing the same action as j . This in turn implies (i) that a new link can actually be formed with an agent ℓ choosing $a_\ell^t = a_j^t$, and (ii) that the overall change in utility for i by playing s'_i instead of s_i is strictly positive: this is so because the change of utility due to the play of the social game with all neighbors is trivially equal to zero (since $a_\ell^t = a_j^t$ and all other neighbors have remained the same, while the change in the expected number of different types is negative (since there is a positive probability that ℓ has the same type as i , and agent j is known for sure to be of different type).

We have proven that $u_i(s'_i, s^t) > u_i(s_i, s^t)$, and this suffices to show that agent i cannot choose a strategy such that $g_{ij}^{t+1} = 1$. \square

Proof of Proposition 2

Proof. Let us suppose, ad absurdum, that strategy $s_i = (a_i, g_i)$ such that $g_{ij} = 0$ maximizes i 's interim utility. We first observe that g_i cannot tell i to have less than k links, because otherwise i might increase his utility by simply adding the link with j , obtaining an additional utility of $\pi(a_i, a_j^t)$, which is surely positive because of the assumption that $c < \pi(B, A)$, which implies that every payoff in the social game is positive even after subtracting the maintenance cost. Therefore, g_i tells i to have k links, which in turn implies that a new link with some agent ℓ has been formed, since the link with j has been removed. We construct another strategy $s'_i = (a'_i, g'_i)$ such that $a'_i = a_i$ and g'_i is equal to g_i with the only

difference that in g'_i the link with agent j is maintained and the link with agent ℓ is not formed. We now argue that $u_i(s'_i, s^t) > u_i(s_i, s^t)$. We first note that the expected number of mismatches in types is lower with s'_i than with s_i , because $w(i) = w(j)$ for sure while ℓ 's type is different from i 's type with positive probability: indeed, if $a_\ell^t = a_j^t$ then we have the assumption that $n_{i0}(\neg w(i), a_j^t | s^t) \geq 1$, while if $a_\ell^t \neq a_j^t$ (and hence $a_\ell^t = a_i^t$) then we have that $n_{i0}(a_i^t | s^t) \geq 1$ and hence the assumption that $d \frac{n_{i0}(\neg w(i), a_i^t | s^t)}{n_{i0}(a_i^t | s^t)} > \pi(a_i^t, a_i^t) - \pi(a_i^t, a_j^t)$ implies $n_{i0}(\neg w(i), a_i^t | s^t) \geq 1$.

Therefore, if i obtains the same or a larger utility in the social game by interacting with j than with ℓ , then we have obtained that $u_i(s'_i, s^t) > u_i(s_i, s^t)$. The only case where i obtains a larger utility interacting with ℓ than with j is if $a_j^t \neq a_i$ and $a_\ell^t = a_i$. Even in such a case, the assumption that $d \frac{n_{i0}(\neg w(i), a_i^t | s^t)}{n_{i0}(a_i^t | s^t)} > \pi(a_i^t, a_i^t) - \pi(a_i^t, a_j^t)$ ensures us $u_i(s'_i, s^t) > u_i(s_i, s^t)$.

We can conclude that no strategy s_i such that $g_{ij} = 0$ can maximize i 's interim utility, and this means that i will surely maintain the link with j if he has a revision opportunity at time $t + 1$ (and no change of course happens if i is not given a revision opportunity). \square

Proof of Proposition 3

Proof. We first prove point (a). Consider a state s that is A -monomorphic, k -regular, and type-segregated. We check that an agent who receives a revision opportunity at s would see his utility reduced if he changes strategy. Indeed, having less than k links is suboptimal, since $\pi(A, A) > c$. Moreover, removing a link and casting a new one brings a neighbor who still plays A (since the state is A -monomorphic) but possibly is of a different type, hence generating an expected loss. Finally, switching from A to B is clearly detrimental, due to $\pi(A, A) > \pi(B, A)$. Therefore we can conclude that no agent will ever change strategy, and hence state s belongs to a singleton absorbing state. An analogous reasoning can be made for a state that is B -monomorphic, k -regular, and type-segregated, where $\pi(B, B) > c$ holds because $\pi(B, B) > \pi(A, A)$.

We now prove point (b). Consider a state s that is type-monomorphic with x on A and y on B , k -regular and type-segregated. Take an agent of type x who receives a revision opportunity. Maintaining less than k links is sub-optimal for him, since $\pi(A, A) > c$ and, in addition, all agents playing A are of type x so that there is no risk of a type mismatch. Replacing an existing link with a new one has no effect on utility if the new link is cast towards an agent who currently plays A , since all agents playing A are of type x and hence there is no risk of a type mismatch. Replacing an existing link with a new one has a negative effect on expected utility if the new link is cast towards an agent playing B , since all agents

currently playing B are of type y and $d > \pi(B, B) - \pi(A, A)$, which means that the penalty for the type mismatch is larger than the maximum attainable gain. Finally, switching from A to B without changing neighbors is detrimental because of $\pi(A, A) > \pi(B, A)$. A similar argument holds a fortiori if we consider an agent y who receives a revision opportunity. We can conclude that any revising agent will at most reshuffle his links among agents playing his same action, who are surely of his same type (due to the perfect correlation between actions and types). Therefore, starting from s we can only reach other states that are type-monomorphic with x on A and y on B , k -regular and type-segregated. This shows that an absorbing set exists, containing type-monomorphic states with x on A and y on B . Clearly, if we invert x with y we obtain that the same reasoning applies to type-monomorphic states with x on B and y on A .

Then, we prove point (c). Consider a state s where $k + 1$ agents of type x play B , $k + 1$ agents of type y play B , and all other agents play A ; also, s is k -regular and type-segregated (which is possible, since $n_x \geq n_y \geq 2k + 2$). Any agent playing B will never change strategy, because he is attaining the maximum possible payoff (i.e., $k(\pi(B, B) - c)$), which is not reachable if he switches to A , and changing neighbors comes with the risk of a type mismatch – since some agents currently playing B are of type x and some are of type y . Any agent playing A will never change his strategy as well. Indeed, the maximum gain which can be obtained by removing an existing link and connecting to someone playing B is $\pi(B, B) - \pi(A, A)$, which is lower than the expected cost of a type mismatch (which is equal to $d/2$ since $n(x, B|s) = k + 1 = n(y, B|s)$) because of the assumption that $d > 2(\pi(B, B) - \pi(A, A))$. Moreover, removing an existing link and connecting to someone playing A brings no benefit and an expected cost due to type mismatch; deleting any of the k links is suboptimal, since $\pi(A, A) > c$; and switching from A to B without changing neighbors is detrimental because of $\pi(A, A) > \pi(B, A)$. Therefore, state s belongs to a singleton absorbing state.

Finally, we prove point (d). Consider a state s where agent i of type y plays A and maintains no link, while all other agents play B and maintain k links towards agents of the same type different from i . Any agent other than i will never change strategy, because he is attaining the maximum possible payoff (i.e., $k(\pi(B, B) - c)$), which is not reachable if he switches to A , and changing neighbors comes with the risk of a type mismatch because some agents currently playing B are of type y . If agent i chooses to connect towards an agent playing B , he will earn at most $\pi(B, B)$, but has to suffer an expected cost of type mismatch equal to $dn_x/(n - 1)$, which is larger than $\pi(B, B)$ due to the assumption that $d > \pi(B, B)(n - 1)/n_x$. If agent i is isolated, then he is indifferent between playing A and

B . Therefore, we have found an absorbing set that is made of states s and s' , where s' is identical to s with the only difference that agent i plays B instead of A . \square

Proof of Lemma 1

Proof. The requirement that s is B -monomorphic is trivially a necessary condition for $s \in S_B^B$. We first show that, given (i), if we are not in a state such that also (ii) and (iii) hold, then it must be the case that with positive probability we reach a state where (i), (ii), and (iii) hold. Suppose that $a_i = B$ for all $i \in N$, but s is not k -regular and/or not type-segregated. We observe that, for a revising agent, choosing action A would clearly be suboptimal. Moreover, the expected payoff of forming a link with a new neighbor who plays B is both higher than not forming that link at all (because $\pi(B, B) - \pi(A, A) > d$ and $\pi(A, A) > c$ imply $\pi(B, B) - c > d$) and higher than maintaining an existing connection with a type different from one's own (because the resulting match cannot be worse, and possibly better). Therefore, with positive probability any agent who has less than k links and/or links with agents of a type different from his own will form new links with agents who play B , and with positive probability these new agents are of his own type.

We now show that a state satisfying (i), (ii) and (iii) forms a singleton absorbing set. To do so, it is enough to observe that any agent who receives a revision opportunity would see his payoff decreased, in expectation, by changing strategy. Indeed, by choosing action A , the agent would obtain a utility that is surely lower than his current utility $k(\pi(B, B) - c)$, and the same is true if he chooses to maintain less than k links; also, substituting an existing link with a new one comes with the risk of linking to an agent of a type different from one's own, which leads to a lower expected payoff. \square

Proof of Lemma 2

Proof. We provide the proof for point (a) only, as the proof for point (b) is almost identical to the proof of point (a).

By definition, if $s \in S_B^A$, then s is type-monomorphic with x on A and y on B , so (i) is trivially necessary. We first show that, given (i), if we are not in a state such that also (ii) and (iii) hold, then it must be the case that with positive probability we reach a state where (i), (ii), and (iii) hold. Suppose that s is type-monomorphic with x on A and y on B , but not k -regular and/or not type-segregated. Consider a revising agent currently playing B . We observe that choosing A is clearly suboptimal since, at most, he can obtain $k(\pi(A, A) - c - d) < 0$ because, by (i), all agents playing A are of a type different from his own and, by assumption, $d > \pi(B, B)(n - 1)/n_y > \pi(A, A)$. Moreover, the expected

payoff of forming a link with a new neighbor who plays B is the same of keeping an existing link with a neighbor who also plays B (because, again by (i), B is played only by agents of similar type), it is strictly greater than the expected payoff of forming a link with an agent who plays A (because, by (i), A is played only by agents of a different type), and it is strictly greater than not forming that link at all (because $\pi(B, B) > \pi(A, A)$ and $\pi(A, A) > c$ imply $\pi(B, B) - c > 0$). Consider a revising agent currently playing A . We observe that choosing B is suboptimal since, at most, he can obtain $k(\pi(B, B) - c - d) < 0$ because, by (i), all agents playing B are of a type different from his own and, by assumption, $d > \pi(B, B)(n - 1)/n_y > \pi(B, B)$. Moreover, the expected payoff of forming a link with a new neighbor who plays A is the same of keeping an existing link with a neighbor who also plays A (because, by (i), A is played only by agents of similar type), it is strictly greater than the expected payoff of forming a link with an agent who plays B (because, by (i), A is played only by agents of a different type and, by assumption, $d > \pi(B, B)(n - 1)/n_y > \pi(A, B)$, so that $\pi(A, B) - d < 0 < \pi(A, A)$), and it is strictly greater than not forming that link at all (because $\pi(A, A) - c > 0$). Therefore, with positive probability any agent who plays B (respectively, A) and that has less than k links and/or links with agents of a type different from his own will form new links with agents who play B (respectively, A) up to k connections in total, and with certainty these new agents are of his own type.

We now show that the set of states satisfying (i), (ii) and (iii) forms an absorbing set. To do so, we first observe that, for the same arguments described above, any agent who receives a revision opportunity would see his payoff certainly decreased by changing action, and/or by choosing to maintain less than k links, and/or by linking to new agents who play a different action (since, by (i), they must be of a different type). Therefore, if we start from a state where conditions (i), (ii) and (iii) are satisfied, we will always remain in states where those conditions are satisfied. We finally show that, taken any two distinct states s and s' satisfying (i), (ii) and (iii), we can move from one to the other with positive probability. Indeed, $s = (a, g)$ and $s' = (a', g')$ can only differ because $g_i \neq g'_i$ for some agent i ; every such agent can receive with positive probability a revision opportunity, and he can choose with positive probability to reshuffle all his links as long as links are cast towards agents choosing his own action, since by (i) there is no risk of forming a link with an agent of a type different from one's own. \square

Proof of Lemma 3

Proof. We first show that 1 mutation is not sufficient to move from S_B^A to another absorbing set. Consider a state $s \in S_B^A$, and suppose that a single mutation hits an agent possibly

changing both his action and his network of interactions. Suppose that an agent different from the mutant is given a revision opportunity. We claim that such an agent will not change action and will not form new links with agents choosing an action different from his own. To see why this is so, we observe five facts. First, forming new links with an agent who is currently playing a different action is suboptimal, as the expected payoff is negative due to the high penalty for a mismatch in type (because the expected payoff from such a link is at most $\pi(B, B) - \frac{n_y}{n_y-1}d$, which is negative due to the assumption that $d > \pi(B, B)^{\frac{n-1}{n_y}}$). Second, if the mutant switched from A to B , then any revising agent who is maintaining a connection with the mutant will not switch to B since he has $k-1$ other neighbors playing A and so by switching he would get $(k-1)\pi(A, A) + \pi(A, B) > \pi(B, B) + (k-1)\pi(B, A)$ (the inequality holding because $k \geq 2$ and A is the risk-dominant action). Third, if instead the mutant switched from B to A , then any neighboring revising agent will not switch to A since he can keep playing B , remove the link with the mutant and form a new link with another agent playing B (who exists because $n_y \geq 2k+1$) which gives him $k\pi(B, B) - kc > \pi(A, A) + (k-1)\pi(A, B) - kc$. Fourth, changing action is clearly suboptimal for an agent who is not maintaining a connection with the mutant. Finally, we observe that when the mutant is given a revising opportunity, he will certainly choose to have k links towards agents playing the same action he was playing before the mutation, because only doing so he can avoid to pay the cost $d > \pi(B, B)$ since all other agents' action is perfectly correlated with their type; given this, it follows that the mutant will also choose to play the action he was playing before the mutation, since this allows him to coordinate. From these five observations it follows that, after that a single mutation has occurred, the system will surely go back to a state where conditions (i), (ii) and (iii) in point (a) of Lemma 2 are satisfied, hence belonging to S_B^A .

A similar reasoning can be applied considering S_A^B in the place of S_B^A , thus obtaining that $R(S_B^A) \geq 2$ and $R(S_A^B) \geq 2$. In what follows we show that, if in addition we have that:

$$n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}, \quad (1)$$

then $R(S_B^A)$ and $R(S_A^B)$ can be computed considering only mutations that hit agents of a single type – either x or y .

We start by providing a sufficient condition to have that m mutations hitting agents of type x are enough for the system to leave S_B^A with positive probability. By Lemma 2, we know that S_B^A is a single absorbing set. So we can choose a specific state in S_B^A to start from, and in particular we can choose a state where there exists a cluster made of $k+1$ agents of type x (i.e., $g_{ij} = 1$ for any i and j in the cluster). If m mutations hit m distinct agents

in the cluster inducing them to switch action from A to B (and with no change to their interaction networks), and if the following inequality is satisfied:

$$m(\pi(B, B) - c) + (k - m) \max\{0, \pi(B, A) - c\} > k(\pi(A, A) - c), \quad (2)$$

then all non-mutants in the cluster who receive a revision opportunity will find it optimal to choose to perform action B while not changing their interaction networks. So, with positive probability the system reaches a new state that belongs to an absorbing set different from S_B^A (indeed, at least the $k + 1$ agents of type x in the cluster will never go back to action A). We also observe that other agents of type x might then find it profitable to switch to action B , and that all agents of type y will keep playing action B . Finally, we note that (2) is surely satisfied when $m = k$.

Suppose now that, starting from a state $s \in S_B^A$, m_x mutations hit agents of type x , m_y mutations hit agents of type y , so that state s' is reached from which another absorbing set can be reached with positive probability. Since we know that (2) is satisfied when $m = k$, then in the following we focus on the case where $m_x + m_y < k$.

We first observe that at least one of the following two inequalities must be satisfied in s' :

$$m_x(\pi(B, B) - c) + (k - m_x) \max\{0, \pi(B, A) - c\} \geq k \left[\pi(A, A) - c - \frac{dm_y}{n_x - 2k + m_y} \right]; \quad (3)$$

$$m_y(\pi(A, A) - c) + (k - m_y) \max\{0, \pi(A, B) - c\} \geq k \left[\pi(B, B) - c - \frac{dm_x}{n_y - 2k + m_x} \right]. \quad (4)$$

To see why, suppose that both (3) and (4) are not satisfied. Then, no agent of type x who is given a revision opportunity finds it profitable to play action B (due to the failure of (3)), and no agent of type y who is given a revision opportunity finds it profitable to play action A (due to the failure of (4)). For the mutant this is true a fortiori, since he can interact with smaller number of mutants – being himself one of the mutants. Hence, sooner or later all agents of type x will go back to play A and all agents of type y will go back play B . When we have perfect correlation between types and actions, it is obvious (because of the assumption that $d > \pi(B, B)(n - 1)/(n_y)$) that revising agents will choose to maintain k links with agents choosing their same action (and hence having their same type). Therefore, if both (3) and (4) are not satisfied, then from s' no other absorbing set can be reached.

We now show that, under the assumption that $n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}$, inequality (4) is false. To understand this, it is enough to re-write (4) with $<$ instead of \geq , and to obtain an explicit bound on n_y (getting rid of m_x by making the new inequality harder to be satisfied). Therefore, inequality (3) must hold.

We then observe that, if inequality (3) holds, then inequality (2) is implied by the assumption that $n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}$. To understand this, we fix $m = m_x + m_y$, we take the difference between the left-hand side of (2) and the left-hand side of (3), and we set it larger than the difference between the right-hand side of (2) and the right-hand side of (3). Working out such inequality we obtain a bound on n_x that is implied by $n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}$ (once it is noted that $n_y < n_x$). This means that, if m_x mutations hitting agents of type x and m_y mutations hitting agents of type y allow the system to leave S_B^A with positive probability, then $m_x + m_y$ mutations hitting agents of type x only are also sufficient for the system to leave S_B^A with positive probability.

We can repeat all the previous arguments with S_A^B in the place of S_B^A . The only difference is that m_x and m_y have inverted roles, and the same occurs for n_x and m_y . Summing up, it is true that, if $n_y > 2k + \frac{kd}{\pi(B, B) - \max\{\pi(A, A), \pi(A, B)\}}$, then we are allowed to focus on mutations hitting only agents of one type, in the attempt to determine $R(S_B^A)$ and $R(S_A^B)$.

Finally, we focus on mutations hitting only one type of agents, and we determine $R(S_B^A)$ and $R(S_A^B)$. We consider a state $s \in S_B^A$, and we observe that we already know that m mutations that hit agents of type x inducing them to switch from A to B are enough to leave S_B^A with positive probability, if inequality (2) is satisfied. We remark that such inequality is also necessary for such an exit from S_B^A . Indeed, it is immediate to observe that, if (2) is not satisfied, then no agent of type x who is given a revision opportunity will find it profitable to play action B , and agents of type y will clearly keep playing action A . Furthermore, once a mutant goes back to action A , the gain (potentially negative) of choosing B over A for agents of type x is further reduced, while agents of type y never find it profitable to play A over B . Sooner or later, perfect correlation between types and actions will be restored, and a k -regular and type-segregated state will be reached, belonging to S_B^A . We denote with \bar{m} the minimum m such that inequality (2) is satisfied. We note that $\bar{m} \geq k$.

We now consider mutations hitting agents of type y who switch from B to A . As long as at least $k + 1$ agents of type B keep choosing B , then mutants will sooner or later go back to B (due to the perfect correlation between B and y , and the fact that B is the payoff-dominant action). Since $n_y \geq 2k + 1$, this means that at least $k + 1$ mutations hitting agents of type y are required to leave S_B^A . Since $\bar{m} \geq k$, we can conclude that $R(S_B^A) = \bar{m}$. We can repeat exactly the same arguments with S_A^B in the place of S_B^A , thus obtaining that $R(S_A^B) = \bar{m}$. Therefore, $R(S_B^A) = R(S_A^B)$. \square

Proof of Lemma 4

Proof. The proof begins by showing that, starting from a generic state $s \in Q$, another state \hat{s} can be reached with positive probability such that it satisfies properties that help in the following construction of a path from \hat{s} to an absorbing set $Q_\ell \subseteq S_A^A \cup S_B^B \cup S_B^A \cup S_A^B$.

Preliminarily, we define $\beta_x(s) \subseteq N_x$ as the set of agents of type x at state s who are playing A and have at least one best reply strategy where either action B is played, or a new link towards an agent currently playing B is cast, or both. Similarly, define $\alpha_y(s)$ as the set of agents of type y such that, at state s , they are playing B and have at least one best reply strategy where either action A is played, or a new link towards an agent currently playing A is cast, or both. We now show how the system can move with positive probability from state s to a state \hat{s} where $\beta_x(\hat{s}) = \emptyset$ and $\alpha_y(\hat{s}) = \emptyset$.

Starting from state s , with positive probability all and only the agents in $\beta_x(s) \cup \alpha_y(s)$ will receive a revision opportunity, and will choose a best reply strategy where, either action B (respectively, action A) is played, or a new link towards an agent currently playing B (respectively, A) is cast, or both. Call s' the state reached after these updates. If $\beta_x(s') = \emptyset$ and $\alpha_y(s') = \emptyset$ then we are done; otherwise we iterate the updating process, giving revision opportunities to all and only the agents in $\beta_x(s') \cup \alpha_y(s')$. We observe that this iteration will yield a state \hat{s} with $\beta_x(\hat{s}) = \emptyset$ and $\alpha_y(\hat{s}) = \emptyset$ in a finite number of repetitions. This is so because agents who switch from A to B (respectively, from B to A) exit definitely set $\beta_x(s)$ (respectively, set $\alpha_y(s)$), and agents who cast a new link towards an agent currently playing B (respectively, A) can be part of set $\beta_x(s)$ (respectively, set $\alpha_y(s)$) for at most k times (since k is the maximum number of links that each agent can maintain).

Let us denote with $N_{xA}(\hat{s})$ the set of agents of type x who are playing action A at state \hat{s} , and with $N_{yB}(\hat{s})$ the set of agents of type y who are playing action B at state \hat{s} . The proof now proceeds by considering 4 possible cases concerning the emptiness/non-emptiness of the sets $N_{xA}(\hat{s})$ and $N_{yB}(\hat{s})$. For each case, we construct the needed sequence of states from $\hat{s} \in Q$ to a state in Q_ℓ that is either monomorphic or type-monomorphic.

Case 1. Suppose first that $N_{xA}(\hat{s}) \neq \emptyset$ and $N_{yB}(\hat{s}) \neq \emptyset$. We apply the following path-building procedure.

Consider a single mutation that hits an agent j of type x who is playing action B at \hat{s} . If no such agent exists, we are done. Otherwise, suppose that after the mutation agent j copies the strategy $\hat{s}_i = (\hat{a}_i, \hat{g}_i)$ of an agent $i \in N_{xA}(\hat{s})$; in particular, j will adopt strategy $s'_j = (a'_j, g'_j)$ such that $a'_j = \hat{a}_i$, $g'_{jh} = \hat{g}_{ih}$ for every $h \neq i, j$, and $g'_{ji} = \hat{g}_{ij}$. We now observe that 3 properties hold for every agent $h \in N_{xA}(s')$: agent h has no best reply where (i) action B is chosen, or (ii) a new link towards an agent playing B is cast, or (iii) an existing link towards an agent of type x playing A is removed unless he is certain to find an agent of type

x when casting a new link towards an agent playing A . (i) and (ii) come from the fact that agents belonging to $N_{xA}(\hat{s})$ have, by construction, no best reply where action B is chosen or a new link towards an agent choosing B is cast; the same holds for the same agents a fortiori at state s' (where agent j has switched from B to A), and it also holds for agent j , who is copying agent i after mutation. (iii) comes from the simple observation that, given the optimality of choosing A , it cannot be optimal to remove a link from an agent of type x playing A unless he is certain to find an agent of type x when casting a new link towards an agent playing A .

For similar reasons, analogous properties hold for the agents belonging to $N_{yB}(s') (= N_{yB}(\hat{s}))$: no agent of type y who is playing B has a best reply where action A is chosen, or a new link towards an agent playing A is cast, or an existing link towards an agent of type y playing B is removed unless he is certain to find an agent of type y when casting a new link towards an agent playing A .

The above three properties imply that any state s'' that is reachable with positive probability in the next period of the unperturbed dynamic is such that, for every agent $h \in N_{xA}(s')$, (1.) the number of h 's neighbors of type x choosing action B has not increased, i.e., $n_{hxB}(s'') \leq n_{hxB}(\hat{s})$, (2.) the number of neighbors of type x choosing action A has not decreased, i.e., $n_{hxA}(s'') \geq n_{hxA}(\hat{s})$, (3.) the probability of mismatch for a new link towards an agent choosing action B has not decreased, i.e., $\frac{n_{yB}(s'') - n_{hyB}(s'')}{n_B(s'') - n_{hB}(s'')} \geq \frac{n_{yB}(\hat{s}) - n_{hyB}(\hat{s})}{n_B(\hat{s}) - n_{hB}(\hat{s})}$, and (4.) the probability of type mismatch for a new link towards an agent choosing action A has not increased, i.e., $\frac{n_{yA}(s'') - n_{hyA}(s'')}{n_A(s'') - n_{hA}(s'')} \leq \frac{n_{yA}(\hat{s}) - n_{hyA}(\hat{s})}{n_A(\hat{s}) - n_{hA}(\hat{s})}$.

Analogous inequalities, of course appropriately adjusted, hold for agents belonging to $N_{yB}(s')$. Altogether these inequalities imply that, for the agents in $N_{xA}(s') \cup N_{yB}(s')$, the 3 properties holding at state s' also hold at state s'' . By induction, we can conclude that the same properties will hold forever, and hence an absorbing set must be reached where the number of agents of type x playing A never falls below $n_{xA}(s') = n_{xA}(\hat{s}) + 1$, and the number of agents of type y playing B never falls below $n_{yB}(s') = n_{yB}(\hat{s})$.

Starting from any state s' in this absorbing set, and following the reasoning done at the beginning of the proof for state s , a state \hat{s}' where $\beta_x(\hat{s}') = \emptyset$ and $\alpha_y(\hat{s}') = \emptyset$ can be reached with positive probability. At state \hat{s}' , there exist at least $n_{xA}(\hat{s}) + 1$ agents of type x playing A , and $n_{yB}(\hat{s})$ agents of type y playing B . Then, following the above argument, a single mutation allows to reach another absorbing set where the number of agents of type x playing A is at least $n_{xA}(\hat{s}) + 2$, and the number of agents of type y playing B is at least $n_{yB}(\hat{s})$. Given the finiteness of the set N_x , an absorbing set where all agents of type x play A must

eventually be reached and the number of agents of type y playing B is at least $n_{yB}(\hat{s})$. This completes the procedure.

The same path-building procedure can now be repeated, constructing a sequence of absorbing sets, with each step requiring a single mutation, and where the minimum number of agents of type y playing B increases by at least 1 at each step, while the minimum number of agents of type x playing A always remains n_x . Given the finiteness of the set N_y , at the end of this procedure a state is reached where all agents of type x play A , and all agents of type y play B ; at such a state, all agents find it optimal to have exactly k links with agents choosing the same action, and hence will end up having k connections with agents of the same type; by Lemma 2, we know that the absorbing set S_B^A has been reached.

Case 2. We now suppose that $N_{xA}(\hat{s}) = \emptyset$ and $N_{yB}(\hat{s}) \neq \emptyset$. We apply the procedure described above to agents in $N_{yB}(\hat{s})$, thus obtaining that a single mutation per step is sufficient to move along a sequence of absorbing sets, where the number of agents of type y playing B increases by at least 1 at each step, until all agents of type y play B . Starting from any state \hat{s}'' in the absorbing set that has been reached, and following the reasoning done at the beginning of the proof for state \hat{s} , a state \hat{s}'' where $\beta_x(\hat{s}'') = \emptyset$ and $\alpha_y(\hat{s}'') = \emptyset$ can be reached with positive probability. We know for sure that $N_{yB}(\hat{s}'') = N_y$. If $N_{xA}(\hat{s}'') \neq \emptyset$, then we can apply the path-building procedure to agents in $N_{xA}(\hat{s}'')$, and reason analogously to what done for case 1, so reaching the absorbing set S_B^A .

If instead $N_{xA}(\hat{s}'') = \emptyset$, then all agents of type x are playing B at state \hat{s}'' . The only possibility that some agents of type x are indifferent between playing A and playing B is that they are isolated (an agent i is isolated if $g_{ij} = 0$ for all $j \in N$). If they want to cast new links with agents choosing B , with positive probability they will do so and will be lucky enough to link to agents of type x . These agents now strictly prefer B over A . If more than one agent of type x remains isolated, then all such agents can jointly switch from B to A with positive probability; in the subsequent period, these agents will find it optimal to connect among themselves as playing A now implies to be an x and there is no risk of type mismatch; this leads with positive probability to a state \hat{s}''' where $\beta_x(\hat{s}''') = \emptyset$, $\alpha_y(\hat{s}''') = \emptyset$, and $N_{xA} \neq \emptyset$. Then, the path-building procedure described in case 1 can be applied starting from \hat{s}''' , and the absorbing set S_B^A is eventually reached. Finally, if at most one agent of type x is isolated and indifferent between A and B , then a single mutation can hit such an agent and let him connect with other agents of type x playing B , so that an absorbing set belonging to S_B^B is reached.

Case 3. Suppose that $N_{xA}(\hat{s}) \neq \emptyset$ and $N_{yB}(\hat{s}) = \emptyset$. This case runs as in case 2, with reversed roles between x and y and, when only one agent of type y is isolated, leading to an

absorbing set belonging to S_A^A .

Case 4. Finally, we consider the case in which $N_{xA}(\hat{s}) = \emptyset$ and $N_{yB}(\hat{s}) = \emptyset$. All agents of type x find it optimal to choose B and to have k links towards agents playing B (who are surely of type x), while all agents y find it optimal to choose A and to have k links towards agents playing A (who are surely of type y). The absorbing set S_A^B is so necessarily reached. \square

Proof of Lemma 5

Proof. Suppose to start from S_B^A . As shown in the proof of Lemma 3, after the formation of a cluster of $k + 1$ agents of type x (which happens with positive probability starting from any $s \in S_B^A$) it is enough to have $R(S_B^A)$ mutations hitting the agents in such a cluster (in particular, making them switch from action A to action B while keeping their interaction network fixed) to move the system to a state from which, with positive probability, a new absorbing set \tilde{Q} is reached where at least those $k + 1$ agents of type x choose B , and all agents of type y keep choosing B .

If $\tilde{Q} \subseteq S_B^B$, we are done. Otherwise, consider a single mutation hitting an agent of type x who currently plays A , making him choose action B and cast all his connections towards agents choosing action B . With positive probability, the mutant casts all his links towards agents of type x (which, by construction, are at least $k + 1$). This leads the system to either S_B^B or to another absorbing set where the number of agents of type x playing B has increased by at least 1, while all agents of type y keep playing B . By repeating this argument, S_B^B is surely reached within a finite number of steps each of which requires 1 mutation only.

The same reasoning can be applied to S_A^B in the place of S_B^A , completing the proof. \square

Proof of Lemma 6

Proof. We show in the following that, starting from state $s \in S_A^A \cup S_B^B$, we can reach S_B^A following a path of absorbing sets such that a single mutation is sufficient to move from one absorbing set to its successor in the path. The same arguments can be repeated for S_A^B instead of S_B^A , completing the proof.

Suppose that to be in state $s \in S_B^B$. Suppose also that a single mutation hits an agent, say i , of type x making him switch to action A and no links. Call this new state s' . Since $d > \pi(B, B) \frac{n-1}{n_y}$, at s' agent i does not want to cast new links towards agents playing B ; so all states which are reachable with positive probability from s' with one round of revision opportunities are such that i maintains no links. Moreover, if i is the only isolated agent of type x and no other agent wants to switch to action A , then s' must belong to an absorbing

set; otherwise, we reach a new state s'' which belongs to a new absorbing set where either i forms links with other isolated agents of type x who (with positive probability) switch to play A or some agents currently maintaining a link towards i switch to action A . With a further single mutation, another agent of type x who is currently playing B can be made switch to A , severe all his current links, and connect to and only to agents of type x who are playing A ; this leads to a new state s'' belonging to a new absorbing set where the number of agents of type x playing A has increased. We can iterate the last passage until we get to some state in S_B^A .

Suppose now that $s' \in S_A^A$. We can apply an argument similar to the one just described (with the only difference that mutations affect agents of type y) and draw an analogous conclusion. \square