


# WORKING PAPERS

Carlos Alós-Ferrer  
Simon Weidenholzer

## Imitation and the Role of Information in Overcoming Coordination Failures

September 2010

Working Paper No: 1008



**DEPARTMENT OF ECONOMICS**

**UNIVERSITY OF VIENNA**

All our working papers are available at: <http://mailbox.univie.ac.at/papers.econ>

# Imitation and the Role of Information in Overcoming Coordination Failures\*

Carlos Alós-Ferrer<sup>†</sup>

Simon Weidenholzer<sup>‡</sup>

September 2010

## Abstract

We model the structure of a firm or an organization as a network and consider minimum-effort games played on this network as a metaphor for cooperations failing due to coordination failures. For a family of behavioral rules, including Imitate the Best and the Proportional Imitation Rule, we show that inefficient conventions arise independently of the interaction structure, if information is limited to the interaction neighborhoods. However, in the presence of informational spillovers, a minimal condition on the network guarantees that efficient conventions will eventually dominate. An analogous result is established for average opinion games.

**Keywords:** Minimum Effort Games, Local Interactions, Learning, Imitation.

**JEL Classification Numbers:** C72, D83.

---

\*We thank Larry Blume and Maarten Janssen for helpful comments. Financial support from the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017 is gratefully acknowledged.

<sup>†</sup>University of Konstanz, Department of Economics. Box 150, D-78457 Konstanz (Germany). Phone: (0049) 7531 88 2340. Email. carlos.alos-ferrer@uni-konstanz.de

<sup>‡</sup>University of Vienna, Department of Economics. Hohenstaufengasse 9, A-1010 Vienna (Austria). Phone: (0043) 4277 37424 Email: simon.weidenholzer@univie.ac.at

There are many situations where the performance of a group depends on the effort exerted by its weakest member. For instance, in synchronized swimming or in rowing one poorly performing team member will jeopardize the team's chance of success. Likewise, in an orchestral concert, a single violin out of tune may spoil an entire performance. Whereas in these examples a good coach or a conductor may help the team overcome coordination failures and achieve a high level of group performance, under many circumstances it is impossible to contract the effort levels chosen by the individual team members. Moreover, as Knez and Camerer (1994) argue, the recent trend of flattening hierarchies has led organizations to coordinate different business functions through informal mechanisms rather than formal lines of authorities. For instance, consider a group of computer programmers who jointly write on a computer program, which consists of several subroutines, each written by one programmer. The performance of the subroutine is determined by the effort exerted by the respective programmer and the performance of the joint program is determined by its weakest subroutine. In this sense, a bug in a subroutine causes the entire program to malfunction. This situation gives rise to a minimum effort or weakest link game, as analyzed by e.g. Van Huyck, Battalio and Beil (1990), henceforth VHBB, where the payoff<sup>1</sup> of an agent depends on the minimum of all effort levels chosen and a cost associated to the agent's actual effort level. Note that under these premises, a given agent will optimally choose an effort level equal to the minimum of all effort levels chosen. This implies that i) each profile where each agent inserts the same effort level corresponds to a Nash equilibrium and ii) that those Nash equilibria where teams manage to coordinate on high effort level are very "fragile" in the sense that one single agent deviating to a lower effort level will prompt other agents to follow. Thus, this element of strategic uncertainty inherent in minimum effort games might eventually cause entire work groups or cooperations to fail.

In the present paper we use social networks to model the organizational structure of a firm or industry that is confronted with such a weakest link problem.<sup>2</sup> The social network determines who interacts with whom in the firm. In fact, by modeling social interactions within the firm as a network, we are able to capture *any* form of interaction between agents in the firm. Note that in general these interaction structures will be rather local, meaning that agents will only interact with a small subset of the population, so that not everybody will interact with everybody else. Figure 1 provides an example of a possible interaction structure in an IT-firm, which consists of a marketing group, a group of programmers, agents working in a financial division, and a board. Note that while there will be interaction within, say, the group of programmers, there will be hardly any interaction between the group of programmers and the work group in the marketing department. At the behavioral level, we follow Herbert Simon's (1947) classic view that an organization is composed of boundedly rational agents. That is, when deciding on the effort to be invested in some task, agents do not use highly sophisticated forms of reasoning but instead rely on simple behavioral rules.<sup>3</sup> Our main focus is on behavioral rules based on imitation, where agents essentially mimic the behavior of other agents who are perceived as successful. Imitation seems to be a well justified behavioral rule in e.g. circumstances where the game itself is not properly understood, agents lack computing capacities, or simply want to economize on decision costs (see Alós-


---

<sup>1</sup>Throughout the paper, we will interpret payoffs as observable performance levels rather than, say, privately known wages. Assuming that performance will eventually influence actual monetary payoffs, focusing on the former is a reasonable simplification when studying learning in organizations.

<sup>2</sup>In Section I.A we provide several examples of weakest link games appearing in organizational theory and in social interactions in general. See also Knez and Camerer (1994) for several examples of weakest link games appearing in firms.

<sup>3</sup>See Sobel (2000) for a survey of learning models in economics.

Ferrer and Schlag (2009) for a broader view and a review of the literature on imitation rules and see Apesteguia, Huck and Oechssler (2007) for experimental evidence on imitation learning). Thus, it seems reasonable to expect that imitation plays an important role in a complex environment, such as a large firm or organization. Further, note that several widespread business practices, such as benchmarking, focusing on “best practices”, or employee-of-the-month programs, are essentially of an imitative nature. In particular, we assume that each agent uses some imitation rule belonging to a general class of behavioral rules. This class is characterized by a few sensible properties capturing the impact of high observed payoffs on the behavior of agents. Prominent examples in this class are the imitate the best max rule and the proportional imitation rule. Note that our framework also allows us to accommodate situations where different agents use different imitation rules.



**Figure 1.** An example of an organizational structure. Agents  $\{2, 5, 6\}$  work in marketing, agent 7 manages acquisitions and supplies, agent 8 is the accountant, and agents  $\{4, 9, 10, 11\}$  form a group of programmers. Agent 2 is the head of marketing, agent 3 is the financial officer, and agent 4 leads the group of programmers. All division managers sit in the board of the company, i.e. the board is made up of agents  $\{1, 2, 3, 4\}$ .

Our main research question is, under which conditions on the flow of information among subgroups of the firm’s workers, and under which conditions on the organizational structure, would firms achieve coordination at high effort conventions.

The answers to these questions could provide new insights into the circumstances under which failing companies might be turned around.

In particular, we analyze and compare two scenarios. In a first scenario, we consider the case where agents only receive information from other agents they interact with, i.e. information is a local matter. In this case we find that regardless of the interaction structure the (inefficient) lowest effort convention will be the only long run outcome. The simple reason for this is that if only one agent (by accident) deviates to a lower effort level, he will earn the highest payoff among his interaction partners and hence will never switch back. As other agents might either follow the bad example or make mistakes themselves, we will end up in a convention where everybody exerts the lowest possible effort. Conversely, once in this bad equilibrium, moving to a more efficient equilibrium takes more than one agent deviating to a higher effort level. Consequently, if information is a local matter, it will be very difficult for cooperations to achieve

coordination on high effort levels. We underscore this inefficiency result by showing that if agents use more sophisticated best reply rules instead of imitation we will also expect to observe coordination on the lowest effort level. The reason behind this second inefficiency result is that each agent always has the lowest effort level in her interaction neighborhood as her best response. Thus only one agent deviating to a lower effort level will trigger a chain reaction and we reach the state where everybody chooses the lowest effort level.

We then move on to discuss a scenario where agents who learn by imitation not only know what is going on in their interaction neighborhood but may, from time to time, also receive pieces of information not stemming from their direct interaction partners. The idea behind these “information spillovers” is that there is either some kind of institutionalized exchange of information (e.g. an employee of the month award or some other form of benchmarking or best practice) or there is simply some informal way of exchanging information between different work groups (e.g. chitchat in the company’s coffee room). Under information spillovers we find two opposing effects: First, note that once a cluster of agents chooses the highest effort level, at least one player will earn the highest payoff and may be copied by other agents with whom he is not interacting. In a next step, other agents might receive the high payoff from coordinating at the high effort convention and might be imitated themselves. In this manner, the efficient strategy may spread out contagiously from an initially small group. Secondly, note that if one player in each disjoint neighborhood switches to the lowest effort level, those players will earn the highest payoff and hence will be copied. If the number of disjoint neighborhoods exceeds the size of the smallest interaction neighborhood, the first effect will dominate and we obtain efficient outcomes. Conversely, if the number of disjoint neighborhoods exceeds the size of the smallest neighborhood, the second effect will dominate and the selection of inefficient conventions remains. Consequently, both visibility of success throughout the company and a relatively decentralized organizational structure of a firm will be essential for obtaining efficiency. We remark that the conditions we provide are *exhaustive*, in the sense that we characterize the set of long run outcomes for *any* given organizational structure.

Our work is related to the literature on learning in coordination games (see e.g. Crawford and Haller (1990), Crawford (1991, 1995), Kandori, Mailath and Rob (1993), or Young (1993)) and in particular to the literature of learning on networks (see e.g. Ellison (1993), Blume (1993), Anderlini and Ianni (1996), Bala and Goyal (1998), Eshel, Samuelson and Shaked (1998), Morris (2000), or Alós-Ferrer and Weidenholzer (2006, 2008)). Conceptually, the present paper is most closely related to Eshel et al. (1998) and to Alós-Ferrer and Weidenholzer (2006, 2008). Eshel et al. (1998) study imitation learning in prisoners’ dilemma games played by a population of agents situated around a circle. They find that imitation might help populations to reach cooperative outcomes. Alós-Ferrer and Weidenholzer (2006, 2008) study agents using the imitate the best max rule playing a  $2 \times 2$  coordination games against each other. Alós-Ferrer and Weidenholzer (2006) consider the circular city model, without information spillovers. They examine whether a risk dominant convention or an efficient convention will be established in the long run. They find that this depends on the interaction radius or the individual agents. As shown in the present paper, in minimum effort games, the particular interaction structure will be irrelevant as the inefficient convention is *always* the long run outcome. Alós-Ferrer and Weidenholzer (2008) consider coordination games information spillovers in general networks. Unlike the present paper where we provide an answer to the question of which convention will be adopted in the long run for every network, Alós-Ferrer and Weidenholzer (2008) only provide sufficient conditions for an efficient strategy to be selected.

The issue of coordination in minimum effort games has also received much attention in the experimental literature. In their seminal paper VHBB show that whereas small groups are able to coordinate on high effort outcomes, it becomes virtually impossible for larger groups to coordinate on effort levels higher than the lowest effort level.<sup>4,5</sup> Given these results, the question of how cooperations might eventually reach efficient outcomes has been the focus of recent experimental research. Brandts and Cooper (2006, 2007) study how a bonus rate can help individuals to coordinate on higher effort conventions. In their “corporate turnaround game”, agents play a minimum effort game against each other and initially face relatively low incentives to coordinate at the high effort equilibrium. By exogenously introducing a bonus rate Brandts and Cooper (2006) modify the underlying base game, which still belongs to the class of minimum effort games, to offer larger incentives to coordinate at the high effort equilibrium. By doing so, they observe coordination at higher effort levels. Interestingly, they also find that the magnitude of the bonus does not play a role and that individuals manage to maintain high effort outcomes once the bonus rate is removed. Thus, a temporary “shock therapy” can leave permanent marks. Brandts and Cooper (2007) study the interaction between employees and employers when confronted with a weakest link structure. They find that communication between managers and employees is a more effective tool for overcoming coordination failures than changes in the incentive structure. Blume and Ortmann (2007) demonstrate that costless pre play communication can also drastically increase the effort levels chosen by individuals in the absence of a manager. Weber (2006) uses a different approach to foster coordination at high effort equilibria in minimum effort games. Starting with initially small groups of players who are more likely to overcome the coordination problem, the group size is slowly increased, thereby also reaching efficient outcomes for larger groups. It turns out that for this mechanism to work, it is crucial that the newly entering players are already familiar with the history of play. Feri, Irlenbusch and Sutter (2009) study a situation where small groups instead of single agents play a minimum effort game against each other. Within this context, they find that team decision making significantly increases effort levels compared to the scenario where the effort choices are made by single individuals. Cabrales, Miniaci, Piovesan and Ponti (2009) study the interplay between fairness considerations and strategic uncertainty in the context of optimal contracts between employers and employees when confronted with a weakest link structure.

The remainder of the paper is organized as follows. Section I introduces the elements of the model and the main techniques employed. Those elements are the minimum effort games (Section I.A), the networks (Section I.B), and the (imitative) behavioral rules (Section I.C). In Section A we briefly present and discuss inefficiency results for the case of local information, including implications for best reply learning. Section III contains our main result on learning the efficient convention under information spillovers. Section IV briefly comments on average opinion games, a different class of games where a similar efficiency result can be established. Section V concludes.

---

<sup>4</sup>Crawford (1991) and Robles (1997) offers an explanation of these results rooted in evolutionary game theory. The former model presents an adaptive process, which tracks the empirical data of VHBB quite well and the latter paper presents a model of best response learning in minimum effort games.

<sup>5</sup>See also Cooper, Dejong, Forsythe and Ross (1990) and Battalio, Samuelson and Van Huyck (2001) for experimental results on two player coordination games and Van Huyck, Battalio and Beil (1991) for experimental results on average opinion games. See Devetag and Ortmann (2007) for a detailed literature review of experimental results on coordination games.

## I. The Model

We model the structure of a firm or an organization as a network and consider a model of social learning in discrete time. In each period, each agent plays a minimum effort game against the agents he interacts with. When deciding on how much effort to invest we assume that players follow simple behavioral rules based on imitation. Consequently our model consists of three ingredients: i) the minimum effort games played by agents; ii) the network which specifies the interaction and information structure; and iii) the behavioral rules used by agents. We will now discuss each of these concepts in turn.

### A. *Minimum Effort Network Games*

We model strategic interaction within the organization through *minimum effort* or *weakest link games*. The relevance of these games and the basic strategic problems arising in them have been known for quite some time. David Hume (1739, Bk. III, Pt.II, Sec. VII.), writing in the 18th century, provided the following observation, which captures the dilemma faced by individuals confronted with a minimum effort game.<sup>6</sup>

Two neighbors may agree to drain a meadow, which they possess in common; because 'tis easy for them to know each others mind, and each may perceive that the immediate consequence of failing in his part is the abandoning of the whole project. But 'tis difficult, and indeed impossible, that a thousand persons shou'd agree in any such action.

This example illustrates well the two most important strategic factors present in minimum effort games: First, minimum effort games exhibit strategic complementariness, i.e. the incentives to put in high effort levels are non-decreasing in the effort level provided by the others. Second, uncertainty about the other players' choices makes it very difficult to achieve coordination at high effort levels in large populations.

The strategic structure of a minimum effort game is encountered in a wide range of both social and economic interactions. For instance, in 1803 in his *Speeches in Parliament*, William Windham noted on the subject of defense that “The strength of a chain, according to an old observation, was the strength of the weakest link.” More recently, Hirshleifer (1983) and Cornes (1993) studied the private provision of public goods when agents face a weakest link structure as present e.g. in the problem of building dykes to protect against flooding. Bryant (1983) and Cooper and John (1988) argue that coordination failures in minimum effort production technologies might be the source of underemployment in rational expectations models. Kremer (1993) argues that many production processes consist of several tasks or subcomponents and that one poorly undertaken task or one imperfect subcomponent might lead to a severe reduction of the product's final quality. Among the examples he lists are the Challenger catastrophe, where a single malfunctioning o-ring led to the loss of the space shuttle, or companies failing due to errors in marketing despite producing perfectly good products. In his model, the skill of employees determines the project's chance of success. In equilibrium, firms will seek to either only hire highly skilled workers or only low skilled workers, and both wages and output will rise steeply in skill. Thus, minimum effort technologies might explain the differences in earnings among high and low skilled worker but may also account for the gap between industrialized and developing countries. Further, Knez and Camerer (1994) provide us with two neat examples of minimum effort games in labor economics. In their first example, agents prefer

---

<sup>6</sup>We owe this quote to Skyrms (2006).

to shirk in a group production process, thereby resembling a prisoners' dilemma or public goods game. Following Holmstrom (1982), a principal can overcome this problem by penalizing every individual if the group falls short of a certain target level. Under such a regime, agents prefer to shirk if only one agent in their group shirks and prefer not to shirk if everybody else does not shirk, thus giving rise to a minimum effort game. In their second example, Knez and Camerer (1994) argue that, due to peer effects, it might be the case that in certain situations, agents would like to keep their effort levels in line with the least productive team member, thus again creating a weakest link structure. Within the context of international relations, Barrett (2007) argues that the problem of countries coordinating to eradicate infectious diseases actually constitutes a minimum effort game, as one country not participating in the joint effort might provide a safe haven from which the disease might spread back.<sup>7</sup>

In this paper, we consider minimum effort games as discussed by VHBB with the difference that they are played on a network. In a nutshell, the payoff (interpreted as performance) of each agent  $i$  out of a given population  $I$  depends on the minimum effort among the agent's own effort level and all effort levels chosen in a subset of agents  $K(i) \subseteq I$ , called the *interaction neighborhood* of  $i$ , and a cost associated to the agent's effort level (we will give more details on the network structure below).

Formally, each agent chooses an effort level  $e$  from the set

$$E = \{e_{\min}, e_{\min+1}, \dots, e_{\max}\} \subseteq \mathbb{R},$$

with  $e_{\min} < \dots < e_{\max}$ . We denote by  $\omega = (e_j)_{j \in I}$  a strategy profile in the overall population and we write  $\vec{e}$  for the "monomorphic" state where all players choose the same effort level  $e$ , i.e. where  $e_j = e$  for all  $j \in I$ .

In particular, the payoff of agent  $i$  is given by the minimum effort chosen in  $K(i) \cup \{i\}$  minus a cost  $\delta e_i$  with  $0 < \delta < 1$  associated with choosing effort level  $e_i$ . So, given a strategy profile  $\omega = (e_j)_{j \in I}$  player  $i$  will earn a payoff of

$$u_i(e_i, \omega) = \min_{j \in K(i) \cup \{i\}} e_j - \delta e_i$$

Note that this payoff does not depend on the number of players in player  $i$ 's interaction neighborhood. In addition, note that in minimum effort games, best-response behavior is particularly simple. Every agent will obtain the highest payoff if he simply chooses the lowest effort level chosen in his interaction neighborhood. Hence,  $\vec{e}$  is a Nash equilibrium for every  $e \in E$ .

## B. The Network

We now specify the underlying interaction and information structure. Rather than assuming any specific topology, we will keep our analysis as general as possible and consider arbitrary networks satisfying a minimal number of properties. We use a characterization of a *local interaction-information system* similar

---

<sup>7</sup>Applications of minimum effort games are not confined to interaction between humans but also extend to non-human biology. A particularly nice example is provided by the cooperative hunting technique of orcas and bottlenose dolphins, known as carousel feeding, where a group of orcas encircles their prey and stuns them with their tails (see Steiner, Hain, Winn and Perkins (1979)). Note that this situation gives rise to a minimum effort game as individual whales not participating in this cooperative hunt would allow the prey to escape.



to the one used by Alós-Ferrer and Weidenholzer (2008) which in turn is based on Morris (2000).<sup>8,9</sup> A *local interaction system* consists of a finite population of agents, such that each of them interacts with a subset of the population only. Formally,

**Definition 1.** A local interaction system is a pair  $(I, (K(i))_{i \in I})$  where  $I$  is a finite set of players and  $K(i) \neq \emptyset$  is a subset of  $I$  for each  $i \in I$  such that

(K1) Irreflexivity: for all  $i \in I$ ,  $i \notin K(i)$ .

(K2) Symmetry: for all  $i, j \in I$ ,  $j \in K(i) \Rightarrow i \in K(j)$ .

We refer to  $K(i)$  as the *interaction neighborhood* of  $i$ . If  $j \in K(i)$ , we say that  $j$  is a neighbor of  $i$ . In addition, most local interaction systems of interest will additionally be connected, i.e. for any pair of players, there is some path connecting them. That is, starting from any agent the iteration of  $K(\cdot)$  will eventually cover the whole population. We do not impose a connectedness assumption, allowing us to encompass models where agents can interact at a number of alternative, predetermined locations (e.g. branch offices, subsidiaries, or divisions), as in e.g. Anwar (2002) and Ely (2002).

We are particularly interested in two parameters associated with a given local interaction system: The maximum number of disjoint neighborhoods and the size of its smallest - and its largest - interaction neighborhood. Formally, if we define  $\mathcal{V}$  to be the set of all population subsets whose neighborhoods are pairwise disjoint, i.e.

$$\mathcal{V} \equiv \left\{ V \subseteq I \mid (K(i) \cup \{i\}) \cap (K(j) \cup \{j\}) = \emptyset \forall i, j \in V, i \neq j \right\},$$

we can characterize the maximum number of disjoint neighborhoods of a local interactions as

$$w^* = \max \{|V| \mid V \in \mathcal{V}\}.$$

Further, we let

$$Q_{\min} \equiv \min \{|K(i)| \mid i \in I\}.$$


denote the size of the smallest interaction neighborhoods in the network.

As an example of a local interaction-information system, consider the circular city model of local interaction as discussed in Ellison (1993) or Eshel et al. (1998) where a finite population of players is arranged around a circle and each of the agents only interacts with his  $k$  closest neighbors to the left and to the right. See Figure 2 for an illustration. So the interaction neighborhood of agent  $i$  is given by  $K(i) = \{i - k, \dots, i - 1, i + 1, \dots, i + k\}$  and contains  $2k$  players. In the circular city model, we have  $w^* = \lfloor \frac{|I|}{2k+1} \rfloor$  and  $Q_{\min} = 2k$ .

In addition to whom a given agent interacts with, it is important to specify from whom he receives information from. We assume that each agent  $i$  samples a random subset of the population and may only observe the actions adopted and the payoffs obtained by the agents in his sample (details follow in Section

<sup>8</sup>Morris's (2000) definition applies to infinite populations whereas Alós-Ferrer and Weidenholzer (2008) consider both finite and countably infinite populations.

<sup>9</sup>Alternatively, we could specify the local interaction system as a graph where the edges are the players and the links represent interactions between those players, as e.g. in Jackson and Wolinsky (1996) or Bala and Goyal (1998).



**Figure 2.** The circular city model of local interaction.

IC). In particular, we assume that each agent draws his sample from an *information neighborhood*  $M(i)$ . We adopt the following definition

**Definition 2.** An information system for a local interaction system  $(I, (K(i))_{i \in I})$  is a collection  $(M(i))_{i \in I}$  such that, for all  $i, j \in I$ ,

(M1) Observed Play:  $K(i) \cup \{i\} \subseteq M(i)$ .

(M2) Symmetry:  $j \in M(i) \Rightarrow i \in M(j)$ .

This is similar to the definition given in Alós-Ferrer and Weidenholzer (2008) with the main difference being that the information neighborhood here describes *potential sources* of information and is interpreted in a probabilistic sense. The case where each agent always observes his complete information neighborhood will be encompassed as a particular case.

Within our context, the potential source of the sample will turn out to play an important role. We will consider two different natural scenarios. In the first scenario, agents may only receive information from their own interaction neighborhood, i.e.  $M(i) = K(i) \cup \{i\}$  for all  $i \in I$ . In this scenario, information is a strictly local matter. In the second scenario, agents may also receive information from agents outside the narrow bounds of their own interaction neighborhoods. So, we are interested in cases where the information neighborhoods extend at least “a bit” beyond the interaction neighborhoods. In the circular city model, this is naturally the case if the information neighborhood of agent  $i$  is given by  $K(i) = \{i - m, \dots, i, \dots, i + m\}$  with  $m > k$ . In order to translate this condition to the case of irregular networks Alós-Ferrer and Weidenholzer (2008) introduce the notion of *contacts*, who correspond to “closest acquaintances”. More specifically, the contacts  $K^*(i)$  of agent  $i$  are agents  $j$  such that agent  $i$  may potentially observe all their interactions, i.e.

$$K^*(i) = \{j \in I \mid K(j) \cup \{j\} \subseteq M(i)\}.$$

With this definition, each agent may potentially observe what is happening in the interaction neighborhoods of all of his contacts, i.e.

$$M(i) \supseteq \bigcup_{j \in K^*(i)} K(j).$$

For example, in the circular city model with  $m = k + 1$ , the contacts of agent  $i$  are the agents  $i - 1$  and  $i + 1$ . In order to model the idea of the information neighborhoods extending at least slightly beyond the interaction neighborhoods we assume that the relation “to be a contact of” is connected, so that iteration of  $K^*(\cdot)$  eventually covers the whole population, i.e.

**Assumption 1.** For each  $i, j \in I$ , there exists  $\{i_1, i_2, \dots, i_L\} \subseteq I$  such that  $L \geq 1$ ,  $i_1 = i, i_L = j$  and  $i_{l+1} \in K^*(i_l)$  for each  $l = 1, \dots, L - 1$ .

In the circular city model this assumption *exactly* translates into  $m > k$ , and in this sense can be considered as the “minimal condition” such that information may smoothly flow through the network. One can easily think of very plausible information systems such that assumption 1 holds. For instance, think of situations where agents exchange information with their interaction partners on what is going on in their interaction neighborhoods. Thus, in such a scenario, the contacts of agent  $i$  include all his neighbors, i.e.  $K^*(i) \supseteq K(i)$  for all  $i \in I$ . If the underlying interaction system itself is connected, it follows that the contact relationship is also connected, meaning that assumption 1 holds in this scenario. Likewise, if all agents always observe all other agents in the population, i.e.  $M(i) = I$  for all  $i \in I$  the set of contacts of a given agent  $i$  is given by  $K^*(i) = I$  and assumption 1 trivially holds.

### C. Behavioral Rules

The last element of our model concerns the behavioral rules used by agents. Again we will adopt a general approach and merely require those rules to satisfy a number of general properties. In particular, we allow for rule heterogeneity, i.e. different agents might be endowed with different behavioral rules. These rules in turn will give rise to a dynamics in discrete time which we will describe now.

**Inertia.** At each period in time  $t = 0, 1, 2 \dots$  with strictly positive probability  $0 < \rho_i < 1$ , each agent  $i$  receives the opportunity to revise his strategy.<sup>10</sup> That is, with probability  $1 - \rho_i$ , an agent is not able to revise his strategy at a given period. Revision opportunities are assumed to be independent across agents and time.

**Information Sampling.** If information is either costly to obtain or costly to process (or both), it is very likely that agents will not necessarily gather or evaluate all information available. Hence we assume that when an agent receives the opportunity to update his strategy at time  $t$  he draws a random sample  $\mathcal{M}(i, t - 1) \subseteq M(i)$  of the information available in his information neighborhood on which his decision will be based.<sup>11</sup> Samples are independent across agents and time. Further, we assume the following properties.

(S1) The ex-ante probability of drawing any agent in the information neighborhood is positive, i.e.

$$\Pr(j \in \mathcal{M}(i, t)) > 0 \text{ for all } j \in M(i) \text{ and } t.$$

(S2) In each period, each agent  $i$  observes at least himself and the pattern of play in his interaction neighborhood, i.e.

$$\mathcal{M}(i, t) \supseteq K(i) \cup \{i\} \text{ for all } t.$$

<sup>10</sup>I.e. we are considering a model of positive inertia.

<sup>11</sup>Durieu and Solal (2003) analyze random sampling in Ellison’s (1993) circular city model of best reply learning. They assume that agents only receive a random sample of the information available in their interaction neighborhood, though.

Given an information (random) sampling model as just described, denote by  $\mathcal{J}_i \subseteq I$  the set of all possible samples which agent  $i$  can draw. Note that random sampling also encompasses the scenario where agents always observe all agents in their information neighborhoods, i.e. when  $\mathcal{M}(i, t) = M(i)$  for all  $t$ .

Let us clarify what we mean by “sampled information”. Given a state  $\omega = (e_i)_{i \in I}$  and a sample  $J \subseteq I$ , let  $\omega_J = (e_j)_{j \in J}$ . Further, let  $u(\omega) = (u_i(e_i, \omega))_{i \in I}$  and  $u_J(\omega) = (u_j(e_j, \omega))_{j \in J}$ . Sampling the set  $J$  means that the agent is aware of  $\omega_J$  and  $u_J(\omega)$ , but not of the parts of  $\omega$  and  $u(\omega)$  not contained in the former vectors. Of course, depending on the shape of the network, it might be possible for a sophisticated agent to infer information on actions outside  $\omega_J$  from the payoff information in  $u_J(\omega)$ . The behavioral rules we will consider below, however, only make use of information in  $\omega_J$  and  $u_J(\omega)$ .

**Behavioral rules.** When a revision opportunity arises, an agent takes a decision on the basis of the sampled information. Let  $\Omega = E^I$  be the set of all possible population profiles, and let  $\Delta(E)$  denote the probability distributions on the set of actions (effort levels)  $E$ . In the most general formulation, a behavioral rule with informational sampling for agent  $i$  is any mapping

$$R_i : \Omega \times \mathcal{J}_i \mapsto \Delta(E)$$

such that  $R_i(\omega, J) = R_i(\omega', J)$  for all  $J \in \mathcal{J}_i$  and  $\omega, \omega' \in \Omega$  such that  $\omega_J = \omega'_J$  and  $u_J(\omega) = u_J(\omega')$ ; that is, decisions only depend on the information that an agent actually has.

For our purposes, it will be enough to characterize a behavioral rule by the sets of actions  $S_i(\omega, J) \subseteq E$ , which are chosen with strictly positive probability by agent  $i$  after observing the sample of agents  $J \subseteq M(i)$  given that the current state is  $\omega$ . Then, if  $\omega^{t-1}$  is the profile of effort levels chosen by agents at time  $t - 1$ , the behavioral rule prescribes that agent  $i$  choose an action  $s_i^t$  from the set of actions  $S_i^t = S_i(\omega^{t-1}, \mathcal{M}(i, t - 1))$  at period  $t$  with positive probability.

**Mutation.** Following a standard approach to determine the stability of outcomes, we assume that with probability  $\epsilon > 0$  (independent across agents and time), an agent ignores the prescription of the behavioral rule and chooses an action at random, i.e. he makes a mistake or he “mutates”.

We now turn to the behavioral rules in more detail. We will focus on behavioral rules based on *imitation*, i.e. on rules where only actions observed in the previous period are adopted.<sup>12</sup> Let  $\omega = (e_i)_{i \in I} \in \Omega$  and denote the *carrier* of  $\omega$  in  $J$  by  $C(\omega, J) = \{e \in E \mid e_j = e \text{ for some } j \in J\}$ . A behavioral rule for agent  $i$  is imitative if  $S_i(\omega, J) \subseteq C(\omega, J)$  for all  $\omega \in \Omega$  and  $J \in \mathcal{J}_i$ . Hence, the strategy adopted by agent  $i$  in period  $t$  is always one of the strategies adopted by agents in  $i$ ’s sample in period  $t - 1$ , i.e.  $S_i^t \subseteq C(\omega^{t-1}, \mathcal{M}(i, t - 1))$ .

We will now discuss some possible properties of imitation rules, based on simple behavioral principles, which will play a major role in our analysis.

**Definition 3.** For every  $\omega = (e_i)_{i \in I}$ , let  $B(\omega, J) = \{e_j \mid j \in \arg \max_{k \in J} u_k(e_k, \omega)\}$  denote the set of actions which have given the largest payoffs in the sample  $J$ . We say that an imitation rule for agent  $i$  is

(B1) *salience-based* if  $S_i(\omega, J) \cap B(\omega, J) \neq \emptyset$ ;

(B2) *optimistic at the top* if  $S_i(\omega, J) \subseteq B(\omega, J)$  whenever  $e_i \in B(\omega, J)$ ; and

(B3) *cautious at the top* if  $S_i(\omega, J) \subseteq B(\omega, J)$  whenever  $i \in \arg \max_{j \in J} u_j(e_j, \omega)$ ,

<sup>12</sup>That is, we are considering imitation rules with memory of length one. One could of course extend the focus by considering rules based on longer periods of memory as done by e.g. Alós-Ferrer (2008).

where all conditions are understood for all  $\omega \in \Omega$  and  $J \in \mathcal{J}_i$ .

Let us briefly discuss these conditions. For the interpretations, it is convenient to keep in mind that an agent might observe a single action yielding different payoffs in his observed sample.

Condition B1 (salience-based) states that some action which has always been observed to yield maximum payoffs is imitated with some positive probability. However, it does not require the agent to imitate all actions that have yielded the highest payoffs with positive probabilities, or to imitate only such actions. This amounts to a weak focus on the salience of the highest observed payoffs while potentially encompassing many different behavioral phenomena.

Condition B2 (optimistic at the top) requires an agent to imitate only among actions which have been observed to yield maximum payoffs, if his own action is among them. Note that this is different from the “conservativeness” condition to stay with the current action if it yields maximum payoffs. First, the agent is only required to focus on some subset of actions that have yielded maximum payoffs. This need not include the agent’s current action. That is, the agent is required not to consider actions that did not yield maximum payoffs anywhere in his sample. Second, the agent is required to focus on highest-payoff actions if he observes that his strategy yielded maximum payoffs from some other agent in his sample, even if his own payoff from this strategy is lower. That is, if the agent observes that he is using one of the “best ideas around”, he will focus on some such best ideas even if his current choice did not quite work for him. The condition is optimistic in the sense that the agent implicitly focuses on the best payoffs of actions, even if an action is associated with several different payoffs in the sample.

Last, B3 (cautious at the top) is a weak version of B2, stating that if an agent earns the maximum payoffs in his sample, he will only consider switching to actions that are also associated with maximum payoffs somewhere in the sample. Note that the agent compares the payoffs of other actions with the payoff that he actually has obtained, and not with the best payoff he has observed associated with his own action (which might have been attained by a different agent).

To see that B2 implies B3, let  $\omega \in \Omega$ ,  $J \in \mathcal{J}_i$ , and  $i \in \arg \max_{j \in J} u_j(e_j, \omega)$ . By definition,  $e_i \in B(\omega, J)$  and hence, if B2 holds, we obtain  $S_i(\omega, J) \subseteq B(\omega, J)$ .

Properties B1 to B3 are stylized conditions capturing the impact of high payoffs on the behavior of agents with a particular focus on highest payoffs. The relevance of high observed or experience payoffs (rather than, say, average payoffs) on human decisions is well established in psychology (e.g. Erev and Barron (2005)). While condition B1 focuses on the relevance of the highest observed payoffs, conditions B2 and B3 pin down the importance of having the agent’s own action attain the highest payoffs. Although it is easy to give stronger, data-driven arguments pinning down what “reasonable” properties of imitation rules should be, we choose to keep the behavioral assumptions at the necessary minimum for our results. In particular, some of our results will require assuming B3, while other results only require the weaker B2.

Apart from the obvious benefit of added generality, there is a further reason for relying on general properties of the imitation rules and not adopting a particular one. We specifically want to allow for agent heterogeneity, that is, agents might be endowed with different imitation rules, as long as all rules respect a few basic principles.

In order to better understand these properties, we now enumerate a few examples of imitation rules proposed in the literature.

*Example 1.* As a first example, consider the “imitate the best” or **Imitate the Best Max Rule (IBM)**, as e.g. used by Robson and Vega-Redondo (1996), Vega-Redondo (1997), or Alós-Ferrer and Weidenholzer (2006, 2008), which prescribes that an agent choose the action that has yielded the highest payoff in his information neighborhood. Formally, agent  $i$  chooses

$$e_i^t = e_j^{t-1} \quad \text{with } j \in \arg \max_{j' \in \mathcal{M}(i, t-1)} u_{j'}(e_j^{t-1}, \omega_{t-1})$$

randomizing in case there are several strategies yielding the maximum payoff in the observed sample; in other words,  $S_i^t = B(i, t-1)$ . Clearly, the Imitate the Best Max Rule is salience-based and optimistic at the top.

*Example 2.* A further example of a salience-based imitation rule is the **Proportional Imitation Rule (PIR)** proposed by Schlag (1998). Under PIR, agents choose actions with probability proportional to the positive part of the payoff difference between this action’s payoff and the agent’s own payoff in the previous period. Within our local interaction context, however, it might be the case that the same action earns different payoffs for different agents. A direct translation of PIR could be to imitate a sampled *agent* with a probability proportional to (the positive part of) the payoff difference between the sample agent’s and the imitator’s. This rule is salience-based and cautious at the top. A different alternative is to assume that agents evaluate each action according to the maximum payoff it earned in the sample, giving rise to a “max-PIR”, which turns out to be salience based and optimistic at the top. This is because some action that earns the highest payoff will be imitated with positive probability and if an agent’s action earns the highest payoff somewhere he will not switch.

*Example 3.* As an alternative behavioral rule we will also consider myopic best response learning (as e.g. used by Ellison (1993, 2000) or Morris (2000)) where agents play a best response to the distribution of play in their interaction neighborhood in the previous period.

$$e_i^t \in \arg \max u_i(e_i, \omega_{t-1})$$

Note that under best reply learning, the distinction between information and interaction is hard to defend, as best response learning implicitly postulates that agents are aware of the strategic situation they are confronted with.

Interestingly, in our framework, myopic best reply is actually an imitative rule. To see this, fix an agent  $i \in I$  and let  $e$  be the minimum effort in  $K(i) \cup \{i\}$ . There are three cases. If  $e \neq e_i$ , the best response of player  $i$  is to adopt  $e$ , which is an element of  $C(\omega, J)$  by (S2). If  $e = e_i$  and some other agent in  $K(i)$  also plays  $e$ , the best response is to keep  $e_i = e$ . Finally, if  $e = e_i$  and  $e_j > e$  for all  $j \in K(i)$ , the best response is to adopt  $\min \{e_j \mid j \in K(i)\}$ , which again is an imitative choice.

Note, however, that the third case just discussed shows that best reply is neither salience-based nor cautious at the top. For, if e.g.  $e_i = e_{\min}$ ,  $e_j = e' > e_{\min}$  for all  $j \in K(i)$ , and  $e_j = e_{\min}$  for all  $j \in J \setminus K(i)$ , agent  $i$ ’s best reply is imitating  $e'$  even though  $B(\omega, J) = \{e_{\min}\}$ .

An imitation rule closely related to best-reply in this framework, and specifically tailored to minimum-effort games, could be termed *lazy imitation*: Imitate the lowest observed effort. This rule is cautious at the top.

*Example 4.* An additional imitational rule that deserves mentioning is the *Imitate the Best Average, IBA*

rule as used e.g. by Eshel et al. (1998). IBA prescribes that agents choose the action that has earned the on average highest payoff in their information neighborhood in the previous period. In our opinion, this rule should be thought of more as a normative rule than as one with good empirical, behavioral foundations. Indeed, it is easy to see that IBA does not fulfill any of the properties B1 to B3. Suppose that there are three players in player  $i$ 's information sample,  $i, j$ , and  $k$ , and  $e_i = e_k = e$ ,  $e_j = e' \neq e$ . Suppose  $u_i(e, \omega) = 1$ ,  $u_k(e, \omega) = \alpha$  and,  $u_j(e', \omega) = \beta$  with  $\alpha < \beta < 1$ , i.e.  $e$  earns a high payoff for player  $i$  and a low payoff for player  $k$  whereas  $e'$  earns an intermediate payoff. For  $2\beta > 1 + \alpha$  action  $e'$  on average earns a higher payoff than  $e$  and is prescribed to be chosen by player  $i$  under IBA even though his original action earned him the highest payoff. Hence, under IBA it might happen that all observed actions that yield maximal payoffs receive probability zero, meaning that IBA is *not* salience-based. The same example shows that IBA is not cautious at the top, and hence also not optimistic at the top.

There are of course many other examples of imitation rules capturing particular behavioral phenomena. For instance consider a *Satisficing Imitation Rule*, *SIR* where agents stick to their own action if and only if it yields themselves the highest payoff but randomize on the full set of observed actions otherwise. *SIR* is salience-based and cautious at the top by construction, but is not optimistic at the top. Note that if agents were also to stick to their action if it yields the highest payoff to somebody, the resulting version of *SIR* would be in addition optimistic at the top.

#### D. The Learning Process

The dynamics without mistakes give rise to a Markov process (the *unperturbed process*) for which the standard tools apply (see e.g. Karlin and Taylor (1975)). Given two states  $\omega, \omega'$  denote by  $\text{Prob}(\omega, \omega')$  the probability of transition from  $\omega$  to  $\omega'$  in one period.

An *absorbing set* (or recurrent communication class) of the unperturbed process is a minimal subset of states which, once entered, is never abandoned. An *absorbing state* is an element which forms a singleton absorbing set, i.e.  $\omega$  is absorbing if and only if  $P(\omega, \omega) = 1$ . States that are not in any absorbing set are called *transient*.

Every absorbing set of a Markov chain induces an *invariant distribution*, i.e. a distribution over states  $\mu \in \Delta(\Omega)$  which, if taken as initial condition, would be reproduced in probabilistic terms after updating (more precisely,  $\mu \cdot P = \mu$ ). The invariant distribution induced by an absorbing set  $W$  has support  $W$ . By the Ergodic Theorem, this distribution describes the time-average behavior of the system once (and if) it enters  $W$ . That is,  $\mu(\omega)$  is the limit of the average time that the system spends in state  $\omega$ , along any sample path that eventually gets into the corresponding recurrent class.

The process with experimentation is called *perturbed process*. Since experiments make transitions between any two states possible, the perturbed process has a single absorbing set formed by the whole state space (such processes are called *irreducible*). Hence, the perturbed process is ergodic. The corresponding (unique) invariant distribution is denoted  $\mu(\varepsilon)$ .

The *limit invariant distribution* (as the rate of experimentation tends to zero)  $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu(\varepsilon)$  exists and is an invariant distribution of the unperturbed process  $P$  (see e.g. Freidlin and Wentzell (1988), Young (1993), or Ellison (2000)). That is, it singles out a stable prediction of the original process, in the sense that, for any  $\varepsilon$  small enough, the play approximates that described by  $\mu^*$  in the long run.

The states in the support of  $\mu^*$ ,  $\{\omega \in \Omega \mid \mu^*(\omega) > 0\}$  are called *Long Run Equilibria (LRE)* or

stochastically stable states. The set of stochastically stable states is a union of absorbing sets of the unperturbed process  $P$ . LRE have to be absorbing sets of the unperturbed dynamics, but many of the latter are not LRE; we can consider them “medium-run-stable” states, as opposed to LRE.

Ellison (2000) presents a powerful method to determine the stochastic stability of long run outcomes. In a nutshell, a state is a long run equilibrium if it is more robust to mistakes, compared to others. The particular result we will rely on states that if the *radius* of a union of absorbing sets exceeds its (*modified*) *coradius* then the long run equilibrium is contained in this set.

In this context, let  $\tilde{\Omega}$  be a union of absorbing sets of the unperturbed model. The radius of  $\tilde{\Omega}$  is defined as the minimum number of mutations needed to leave the basin of attraction of  $\tilde{\Omega}$ . The coradius of  $\tilde{\Omega}$  is defined as the maximum over all other states of the minimum number of mutations needed to reach  $\tilde{\Omega}$ . The modified coradius is obtained by subtracting a correction term from the coradius that accounts for the fact that large evolutionary changes will occur more rapidly if the change takes the form of a gradual step-by-step evolution, rather than the form of a single evolutionary event (which would require more simultaneous mutations).

More formally, the basin of attraction of  $\tilde{\Omega}$  is given by

$$D(\tilde{\Omega}) = \{\omega \in \Omega \mid \text{Prob}(\exists \tau \text{ such that } \omega^\tau \in \tilde{\Omega} \mid \omega^0 = \omega) = 1\}.$$

where probability refers to the unperturbed dynamics. Let  $c(\omega, \omega')$  denote the minimum number of simultaneous mutations required to move from state  $\omega$  to  $\omega'$ . Now, a path is defined as a finite sequence of distinct states  $(\omega^1, \omega^2, \dots, \omega^k)$  with associated cost

$$c(\omega^1, \omega^2, \dots, \omega^k) = \sum_{\tau=1}^{k-1} c(\omega^\tau, \omega^{\tau+1}).$$

The radius of a union of absorbing sets  $\tilde{\Omega}$  is defined by

$$R(\tilde{\Omega}) = \min \left\{ c(\omega^1, \dots, \omega^k) \mid (\omega^1, \dots, \omega^k) \text{ such that } \omega^1 \in \tilde{\Omega}, \omega^k \notin \tilde{\Omega} \right\}.$$

The coradius of a union of absorbing sets  $\tilde{\Omega}$  is defined by

$$CR(\tilde{\Omega}) = \max_{\omega^1 \notin \tilde{\Omega}} \min \left\{ c(\omega^1, \dots, \omega^k) \mid (\omega^1, \dots, \omega^k) \text{ such that } \omega^k \in \tilde{\Omega} \right\}.$$

If the path passes through a sequence of absorbing sets  $L_1, L_2, \dots, L_r$ , where no absorbing set succeeds itself, we can define the modified cost of the path as

$$c^*(\omega^1, \omega^2, \dots, \omega^k) = c(\omega^1, \omega^2, \dots, \omega^k) - \sum_{i=2}^{r-1} R(L_i).$$

Let  $c^*(\omega^1, \tilde{\Omega})$  denote the minimum (over all paths) modified cost of reaching the set  $\tilde{\Omega}$  from  $\omega_1$ . The modified coradius of a collection  $\tilde{\Omega}$  of absorbing sets is defined as

$$CR^*(\tilde{\Omega}) = \max_{\omega \notin \tilde{\Omega}} c^*(\omega, \tilde{\Omega}).$$



Ellison (2000) shows that

**Lemma 1.** (Ellison 2000). *If  $R(\tilde{\Omega}) > CR^*(\tilde{\Omega})$  the long run equilibrium (LRE) is contained in  $\tilde{\Omega}$ .*

Note that since  $CR^*(\tilde{\Omega}) \leq CR(\tilde{\Omega})$  also  $R(\tilde{\Omega}) > CR(\tilde{\Omega})$  is sufficient. Furthermore, Ellison (2000) provides us with an elegant bound on the expected waiting time until we first reach the LRE. In particular, one can show that the expected waiting time until  $\tilde{\Omega}$  is first reached is of order  $O\left(\varepsilon^{-CR^*(\tilde{\Omega})}\right)$  as  $\varepsilon \rightarrow 0$ .

Following Ellison (2000), it is easy to give slightly sharper results for the case that  $\tilde{\Omega}$  is a single absorbing set rather than a union of sets. We will make use of the following particular version for convenience.

**Lemma 2.** (Alós-Ferrer and Kirchsteiger 2010). *Let  $A$  be an absorbing set. Then:*

- (i) *If  $R(A) = CR^*(A)$ , the states in  $A$  are LRE.*
- (ii) *If  $R(A) > CR^*(A)$ , the only LRE are those in  $A$ .*

## II. Inefficiency

We start our analysis by considering circumstances under which agents are unable to overcome the coordination problem and will end up in situations where everybody chooses the lowest effort level. We first show that when agents use behavioral rules based on imitation and if information is a local matter we will always expect to observe the lowest effort convention. We highlight this result by showing that even in the case where agents use more sophisticated best reply rules we will expect companies to be caught in this “coordination trap”.

### A. Local Imitation and Inefficiency

We will first analyze the case where the source of potential information is restricted to the bounds of the interaction neighborhood. So, each agent may only observe strategies and the payoffs associated with these strategies of players in his own interaction neighborhood. Thus, the information neighborhood of player  $i$  is given by  $M(i) = K(i) \cup \{i\}$  for all  $i \in I$ .

**Theorem 1.** *Assume local information. If all agents use imitation rules fulfilling B3 (cautious at the top), the inefficient convention  $\vec{e}_{\min}$  is the unique LRE in any local interaction system.*

*Proof:* Consider some player  $i$  who is choosing the lowest effort level in his interaction neighborhood,  $e_i = \min_{j \in K(i) \cup \{i\}} e_j$ . This player will receive the highest payoff in his interaction neighborhood. To see this note that  $\min_{k \in K(K(i))} e_k \leq e_i$ . This implies that

$$u_j \leq e_i - \delta e_j \leq e_i - \delta e_i = u_i$$

for all  $j \in K(i)$ . Note that the second inequality is a strict inequality if  $e_j > e_i$ . This implies that a player with the lowest effort level in his interaction neighborhood will never switch to a higher effort level under any imitation rule that is cautious at the top.

Consider now any state  $\omega$  in an absorbing set, such that  $\omega \neq \vec{e}_{\min}$ . By the previous argument, we can always leave those states if one player mutates to  $e_{\min}$ . The dynamics will then lead to some absorbing

set such that, in all states in the set, at least one additional player (in comparison to the original state) chooses  $e_{\min}$ . Thus we can reach the inefficient convention  $\vec{e}_{\min}$  by a series of single mutations, implying that  $CR^*(\vec{e}_{\min}) = 1$ . In addition, we cannot leave the basin of attraction of  $\vec{e}_{\min}$  with one mutation. This implies that  $R(\vec{e}_{\min}) > 1$ .

Since B2 implies B3, this result automatically holds for rules which are optimistic at the top, such as imitate the best max or the proportional imitation rule. The argument, however, applies for any imitation rule where actions that only earn payoffs lower than the own action are never imitated. Further, note that the scope of the theorem could also be broadened by considering situation where agents do not observe all of their interaction partners, i.e. when  $M(i) \subseteq K(i) \cup \{i\}$ .

The main reason behind this result is that one player deviating to a lower effort level is always enough to leave the basin of attraction of any absorbing state other than the inefficient convention. This implies that we can move into the basin of attraction of the inefficient convention by means of a single mutation chain. Note that the selection of the inefficient convention is not (necessarily) due to it spreading out contagiously (as in Ellison (1993), Morris (2000), or Alós-Ferrer and Weidenholzer (2008)) but rather due to a step-by-step transition (as in Ellison (2000) or Alós-Ferrer and Weidenholzer (2006)).

It is noteworthy that Alós-Ferrer and Weidenholzer (2006) have shown that in  $2 \times 2$  coordination games in the circular city model, whether an inefficient (risk dominant) action, or an efficient action will be selected depends on the relative size of the information neighborhood. If agents only interact with – and receive information from – their two closest neighbors, i.e.  $k = 1$ , under the imitate-the-best rule, the risk-dominant equilibrium is uniquely selected. Similarly to the result presented above, the transition into the basin of attraction of the risk dominant strategy works via a stepwise transition. Interestingly, for larger interaction neighborhoods ( $k > 1$ ), it can be the case that the efficient convention is selected. The main reason for this result is that for larger neighborhoods, agents at the boundary of a cluster of agents playing the efficient strategy observe players near the center of the efficient cluster who earn a relatively high payoff and hence may be imitated. That is, in  $2 \times 2$  coordination games, whether inefficient or efficient conventions will be observed in the long run depends on the interaction structure. In contrast, the prediction that only inefficient outcomes will be observed in the long run is independent of the interaction structure in minimum effort games. This demonstrates that if agents do not look beyond the bounds of their interaction neighborhood there is no hope of achieving efficiency.

### B. Best Reply Learning and Inefficiency

In a next step we will also consider a cognitively more demanding process, namely best reply. As commented above, in our framework, best reply is imitative, although it does not fulfill any of the properties B1 to B3. Still, the same inefficiency result can be proven.

**Proposition 3.** *In any local-interaction system the inefficient convention  $\vec{e}_{\min}$  is the unique LRE under best reply learning.*

*Proof:* If a single agent adopts  $e_{\min}$ , the best response of all his interaction partners is to switch to  $e_{\min}$ . If the underlying network is connected, it follows that the inefficient action  $e_{\min}$  will spread contagiously to the entire population following initial adoption by a single agent. If the network is not connected, we know that one agent switching to the lowest effort level  $e_{\min}$  may prompt all his neighbors to switch.

Thus, we can exhibit a single mutation chain leading to the inefficient convention  $\vec{e}_{\min}$ . Consequently,  $CR^*(\vec{e}_{\min}) = 1$ . Conversely, one mutation is never enough to leave the inefficient convention. Hence,  $R(\vec{e}_{\min}) > 1$  and the claim follows from Lemma 1.

Thus, under best reply learning, there does not seem to be any hope for obtaining efficient outcomes. Further, note from the proof that the inefficiency result is actually stronger than the previous one. In any connected network, the lowest effort level is contagious under best reply learning in the sense of Morris (2000). That is, if just one agent (a “rotten apple”) adopts it, it will spread to the entire population.

### III. Information Spillovers and Efficiency

In the previous section, we considered settings where agents may only observe information stemming from their interaction partners. We have shown that, under local information, agents are not able to coordinate at a high effort convention, regardless of the interaction structure. We now discuss information spillovers, i.e. situations where information may also originate from agents who are not direct interaction partners. In this section, we will show that information spillovers might facilitate outcomes where agents choose high effort levels. Thus, the possibility of observing agents who are not direct interaction partners might help to overcome the coordination problem.

Hence, suppose that agents may, from time to time, observe agents who they are not interacting with, in addition to their interaction partners. As discussed in Section I.B, we model these information spillovers by assuming that the contact relationship is connected, i.e. assumption 1 is satisfied.

We will now introduce a lemma which will be key for our main results, and which relies exclusively on property B1.

**Lemma 4.** *Suppose all agents use salience-based imitation rules. If some agent  $i$  and all of his neighbors choose the highest effort level present in the population then under information spillovers there exists a positive probability path leading to the convention  $\vec{e}_i$ .*

*Proof:* As agent  $i$  and all his neighbors choose the highest effort level present in the population, he will receive the highest payoff in the whole network. Further, no other action can reach this payoff. Hence, with positive probability agent  $i$  will be sampled and his action imitated (by B1) by agents in his information neighborhood. This implies that, with positive probability, the system shifts us to a state where all agents in  $M(i)$  choose  $e_i$ . Now all agents in the set  $K^*(i)$  will earn the highest payoff and hence may be sampled and imitated (again by B1, since no other action can yield this payoff) with positive probability by agents in the set  $M(K^*(i))$ . Now all agents  $K^*(K^*(i))$  earn the highest payoff. Iterating this argument and appealing to connectedness of the contact relationship  $K^*(\cdot)$  we arrive at the convention  $\vec{e}_i$ .

Lemma 4 basically tells us that once there is a group of agents choosing a high effort level then the news that this high effort level is capable of earning high payoffs spreads through the network and we *may* arrive in a state where everybody is choosing a high effort level. Thus, Lemma 4 exhibits a way in which a population might move towards efficient outcomes. In the following lemma, we show that B1 and B2 are sufficient conditions on the imitation rule to guarantee that, starting from a small cluster, the efficient action will spread to the whole population *for sure*.

**Lemma 5.** *Suppose all agents use imitation rules that are salience-based and optimistic at the top. If some agent  $i$  and all of his neighbors choose the highest effort level present in the population, then under information spillovers, the convention  $\vec{e}_i$  will be reached with probability 1 under any salience-based optimistic imitation rule.*

*Proof:* Again, agent  $i$  will earn the highest payoff, and this payoff can only be reached by agent  $i$ 's effort level. By the definition of random sampling, all agents in  $K(i)$  will always sample agent  $i$ . Thus if agents use imitation rules that are optimistic at the top, neither agent  $i$ , nor any agent in  $K(i)$  will ever switch to a different strategy. From Lemma 4, we know that under a salience-based imitation rule there exists a positive probability path leading to the convention  $\vec{e}_i$ . Thus, eventually the system will converge to the convention  $\vec{e}_i$ .

Lemma 5 expresses the idea that once we have a cluster of agents that is such that one agent earns the highest payoff, this cluster of agents cannot be invaded from outside under any optimistic imitation rule. Under a salience-based imitation rule this high effort strategy will have to take over the entire population at some point in time.

Before we proceed to our main result, we need to characterize the absorbing states of the dynamics. The following lemma states that these are actually singletons formed by monomorphic states (conventions).

**Lemma 6.** *Consider the imitation dynamics with positive inertia where all players adopt salience-based imitation rules. Then the only absorbing sets are singletons containing conventions.*

While this result is intuitively plausible, the proof is neither trivial nor intuitive and we relegate it to the appendix. The reason why straightforward, intuitive arguments fail is that in general, several different effort levels might easily lead to the maximal payoff (as an example, see Figure 3) and hence, unless one assumes a particular behavioral rule for all agents, our conditions allow for different reactions to the same environment. This is, of course, unavoidable if one wants to obtain results for general behavioral rules and allow for heterogeneity among agents.

With the help of the previous lemmata we are able to prove the following theorem.


**Theorem 2.** *Consider any local-interaction system with information spillovers. Suppose all agents use imitation rules that are salience-based (B1) and optimistic at the top (B2), and assume positive inertia. Then,*

- i) if  $w^* > Q_{\min} + 1$  the efficient convention  $\vec{e}_{\max}$  is the unique LRE;*
- ii) if  $w^* < Q_{\min} + 1$  the inefficient convention  $\vec{e}_{\min}$  is the unique LRE; and*
- iii) if  $w^* = Q_{\min} + 1$  all conventions  $\vec{e}$  are LRE.*

*Proof:* By Lemma 6, we only need to consider transitions among conventions.

*Step 1.* Consider a convention with effort level  $e$ . Then,  $Q_{\min} + 1$  mutations are necessary and sufficient for a transition to  $\vec{e}^{up}$  with  $e^{up} > e$ .

Let us first show sufficiency. Suppose the process starts at the convention with effort level  $e$  and consider agents mutating to  $e^{up}$ . Once  $e^{up}$  is played by a player  $i$  and all of his neighbors  $j \in K(i)$



**Figure 3.** Non-monomorphic state where multiple effort levels earn the highest payoff on the circular city with  $k = 1$ ,  $E \supseteq \{4, 5, 6, 8\}$ , and  $\delta = \frac{1}{2}$ . Utilities  $u_i$  in parenthesis.

player  $i$  will receive the highest payoff. By Lemma 4 there exists a positive probability path leading to the convention  $\vec{e}^{up}$ . Note that since the number of neighbors may be different across agents in general, we choose the agent with the fewest neighbors, implying that we can move from any convention to a convention with a higher effort conventions at a cost of  $Q_{\min} + 1$ .

On the other hand, no transition to a convention with a higher effort level  $e^{up} > e$  is possible unless some player  $i$  and all his neighbors mutate to  $e^{up}$ , otherwise all mutants will have strictly lower payoffs than the incumbents and hence the incumbents will not switch strategy due to B2. Eventually, the mutants will return to effort level  $e$  by B1.

*Step 2.* Consider a convention with effort level  $e$ . A transition to a convention with effort level  $e^{down} < e$  requires exactly  $w^*$  mutations.

Consider mutations to lower effort levels. By Lemma 5, under any salience-based optimistic imitation rule satisfying B1 and B2, we will move back to the convention  $\vec{e}$  from any state where some player  $i$  and his neighbors play  $e$ . Hence, in order to leave the basin of attraction of  $\vec{e}$  we need to destabilize any cluster of agents that is such that somebody receives the maximum payoff. Since there are  $w^*$  disjoint neighborhoods, we need  $w^*$  players mutating to a lower effort level  $e^{down}$ . Note that the agents who have mutated to lower effort levels will now earn a payoff of  $(1 - \delta)e^{down}$ , which is strictly higher than the payoff of those players who did not switch,  $e^{down} - \delta e'$ . Hence the mutants will be imitated with positive probability by B1. So  $w^*$  mutations are also sufficient for a transition from  $\vec{e}$  to a lower effort convention.

We can now complete the proof. By Step 1, the convention  $\vec{e}^{up}$  can be reached from any other convention with  $Q_{\min} + 1$  mutations, and actually just as many mutations are needed. Hence  $CR(\vec{e}_{\max}) = Q_{\min} + 1$ . By Step 2, the convention  $\vec{e}_{\max}$  cannot be left with less than  $w^*$  mutations and hence  $R(\vec{e}_{\max}) = w^*$ . Part (i) of the statement now follows from Lemma 2(ii).

Consider the convention  $\vec{e}_{\min}$ . By Step 2, it can be reached from any other convention with  $w^*$  mutations, and actually this many mutations are also necessary. Hence  $CR(\vec{e}_{\min}) = w^*$ . By Step 1, the convention  $\vec{e}_{\min}$  cannot be left with less than  $Q_{\min} + 1$  mutations and hence  $R(\vec{e}_{\min}) = Q_{\min} + 1$ . Part

(ii) of the statement now follows from Lemma 2(ii).

Finally, by Steps 1 and 2, for every  $\vec{e} \neq \vec{e}_{\min}, \vec{e}_{\max}$  we have  $CR(\vec{e}) = \max\{Q_{\min} + 1, w^*\}$  and  $R(\vec{e}) = \min\{Q_{\min} + 1, w^*\}$ . Part (iii) now follows from Lemma 2(i).

As Theorem 1, this result holds e.g. for the imitate the best max or the proportional imitation rule.

The intuition behind this result is the following. Essentially we have two opposing effects at work: First, the efficient strategy may spread from the smallest subgroup to the entire population and, second, if one agent in each disjoint neighborhood deviates to a lower effort level, all other agents will follow. If  $Q_{\min} + 1 > w^*$  the first effect will dominate and we will observe the efficient convention. However, if  $Q_{\min} + 1 < w^*$  the second effect will dominate and we will observe the inefficient convention.

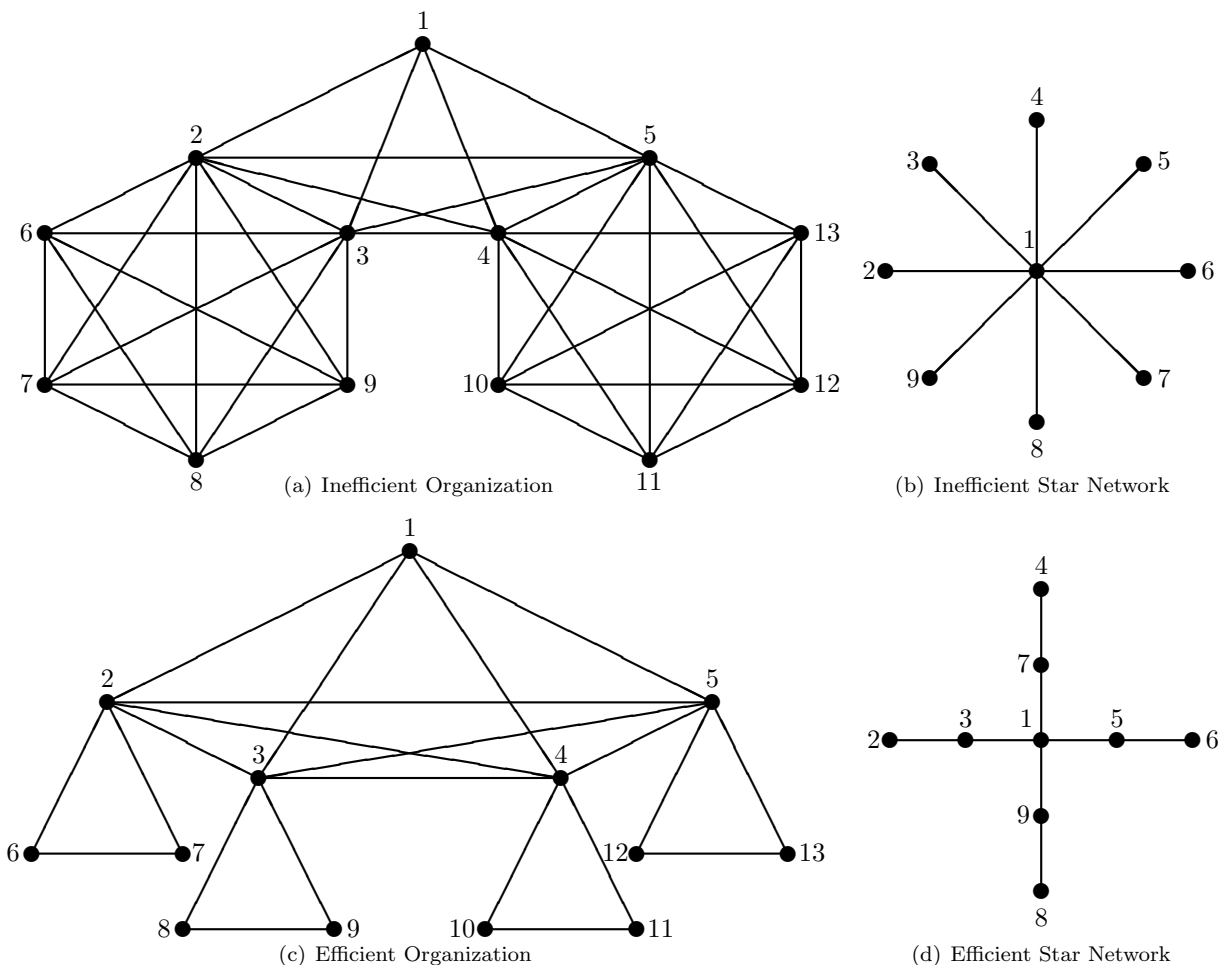
Note that the conditions identified in Theorem 2 provide a complete characterization of LRE in the case of information spillovers, i.e. for every network we are able to provide an answer to the question of which effort level will be observed in the long run, depending on both the smallest size of the interaction neighborhood  $Q_{\min}$  and on the number of disjoint neighborhoods  $w^*$ .

Let us now discuss Theorem 2 by means of a few examples. First, reconsider the introductory example depicted in Figure 1. Here we have  $w^* = 3$  and  $Q_{\min} + 1 = 2$  and, thus, that under information spillovers, the efficient convention will be adopted in the long run. The underlying idea is the following: If agents 3 and 7 (or agents 3 and 8) choose the highest effort level  $e_{\max}$  agent 7 (or agent 8) will earn the highest payoff. Under information spillovers, the efficient effort level will spread through the firm and we reach the efficient convention  $\vec{e}_{\max}$ . Further, note that as we have three disjoint neighborhoods, it takes at least three agents to change their strategy in order to leave the efficient convention  $\vec{e}_{\max}$ . Thus, the combination of “decentrality” of the underlying network and information spillovers guarantee efficient outcomes in the present example.

If, however, the underlying network is instead “central”, the organization will get stuck at the inefficient convention. In Figure 4(a)) we plot an example of a firm where this is the case. There are two work groups in the firm consisting of agents  $\{2, 3, 6, 7, 8, 9\}$  and  $\{4, 5, 10, 11, 12, 13\}$ . In addition, agents  $\{1, 2, 3, 4, 5\}$  sit on the board of the company. Inspecting Figure 4(a)) reveals that we have  $w^* = 2$  and  $Q_{\min} + 1 = 5$  and, thus, that the inefficient convention is selected. One possibility for overcoming this coordination failure would be to split up the existing work groups into smaller groups. For instance, one could split up the two large work groups into four smaller work groups, as done in Figure 4(c)). Now, we have  $w^* = 4$  and  $Q_{\min} + 1 = 3$  indicating that the efficient convention will be adopted in the long run.

Note, however that it might not always be technologically feasible to move to an organizational design with smaller work groups. Nevertheless, in these circumstances it might help to avoid overly centralized network structures. In Figure 4(c)) we exhibit a star network with  $w^* = 1$  and  $Q_{\min} + 1 = 2$ , which leads to the adoption of the inefficient convention. Figure 4(d)) plots another star network, which is obtained by inserting an additional level of hierarchy while leaving the number of agents constant. This ensures that the number of disjoint neighborhoods increases. Indeed, we have  $w^* = 4$  and  $Q_{\min} + 1 = 2$  and, thus, expect to observe the efficient convention in the long run.

We will now consider the speed of convergence of our model, i.e. the speed with which the dynamics approaches its long run outcome. First, consider the case of local information where the inefficient convention is selected. Here we have that  $CR^*(\vec{e}_{\min}) = 1$ . It follows by Ellison’s (2000) Radius-Coradius Theorem that the expected waiting time until the inefficient convention is first observed is of



**Figure 4.** Various organizational structures

order  $O(\varepsilon^{-1})$  as  $\varepsilon \rightarrow 0$ . Second, consider the case of information spillovers. If the inefficient convention is selected, i.e.  $w^* < Q_{min} + 1$  we have  $CR(\vec{e}_{max}) = w^*$ , and, hence, the expected waiting time is of order  $O(\varepsilon^{-w^*})$  as  $\varepsilon \rightarrow 0$ . Now suppose that the organizational structure of the firm is such that  $w^* > Q_{min} + 1$  holds and the efficient convention is selected. As we have  $CR(\vec{e}_{max}) = Q_{min} + 1$  it follows that the expected waiting time is of order  $O(\varepsilon^{-Q_{min}-1})$  as  $\varepsilon \rightarrow 0$ . Hence, in all cases the expected waiting time is independent of the population size, meaning that path dependence plays a minor role and that the long run prediction will be observed at a rather early stage of play. This implies that our model retains its predictive power even in very large populations.

Let us now assume that the interaction structure is such that the efficient convention is selected. Note that in the absence of information spillovers, the speed of convergence towards the inefficient convention is always faster than the speed of convergence towards the efficient convention in the presence of information spillovers. Hence, it will take more time to shift an organization to an efficient convention, by e.g. introducing a system of best practice or benchmarking, than it will take a company to succumb to coordination failure in the absence of such a system.<sup>13</sup>

<sup>13</sup>A similar phenomenon has been observed by Janssen and Mendys-Kamphorst (2004) in the context of - fast crowding

#### IV. Average Opinion Games

We have developed a framework encompassing general network structures and a large class of (possibly heterogeneous) behavioral rules. We have focused on minimum effort games due to their intrinsic interest as a model of economic activity.

Our framework, however, can readily be applied to broader classes of games. One can of course just consider network games where every player plays a qualitatively different game. Naturally, little can be said in such a situation. A more interesting framework is one where an underlying economic activity is captured by essentially the same game being played by every agent in the network. Formally, however, one must operate with a family of games, for, unless the network is very regular, agents will typically have different numbers of neighbors. That is, the capability of accommodating different numbers of players is an essential characteristic of a game to be played within an arbitrary network.

Let the available strategies belong to a given finite, nonempty set  $S$ . Let a strategy profile in the population be given by  $\omega = (s_j)_{j \in I}$ . Retain the notation  $\omega_J = (s_j)_{j \in J}$  for any given  $\emptyset \subsetneq J \subseteq I$ . The payoff of each agent  $i$  out of a given population  $I$  will be given by

$$u_i(s_i, \omega) = u(s_i, \omega_{K(i)})$$

where  $u : S \times \{\omega_J \mid \omega \in S^I, \emptyset \subsetneq J \subseteq I\} \mapsto \mathbb{R}$  is a player-independent function determining the payoff of a player given his own action and the actions of players within a given set. It is reasonable to further assume that  $u$  is symmetric in the sense that the payoff does not change if one permutes the actions in  $\omega_J$ .

Minimum effort network games are a first example of such games. Another example are Cournot oligopolies (and other forms of oligopolistic competition), where each firm might compete against a different number of competitors. More generally, one can consider families of *aggregative games* as outlined in Alós-Ferrer and Ania (2005). These are families of games where the payoff of a player depends on his own strategy and an aggregate of all players' strategies. Hence, it is easy to consider *aggregative network games*, where each player plays against an aggregate of the strategies in his interaction neighborhood.

In this paper, we are interested in an equilibrium selection problem when several Nash equilibria are Pareto-ranked. This does not in general apply to general aggregative games. However, there is an interesting class of aggregative games, disjoint from minimum effort games, where the efficiency part of Theorem 2 can be similarly established. Those are "average opinion games", which generalize a game analyzed by Van Huyck et al. (1991). The defining characteristic of such a game is that payoffs are increasing in an average of all players and decreasing in the distance between the player's choice and that average. Van Huyck et al. (1991) provide an experimental investigation of the salience of payoff dominance in average opinion games.

Average opinion games stylize an important aspect of social interactions, namely the tendency of agents to keep their actions in line with the actions chosen by their peers. This tendency, named variously "peer effects", "conformism bias", or "popularity weighting", plays a role in a variety of economic models, e.g. Ellison and Fudenberg (1993), Brock and Durlauf (2001), and many others. We refer to Van Huyck et al. (1991) for a more detailed discussion of these games.

---

out and slowly crowding in - social norms.



Let the common (finite) strategy space  $S$  be a subset of  $\mathbb{R}$ . Order the strategies in  $S$  so that  $S = \{s_{\min}, \dots, s_{\max}\}$  with  $s_{\min} < \dots < s_{\max}$ . Let  $\text{co}(S) = [s_{\min}, s_{\max}]$ , i.e. the convex hull (interval) of the finite set  $S$ . For each  $n \in \mathbb{N}$ , consider a function  $g^n : \text{co}(S)^n \mapsto \text{co}(S)$ . We will call  $g^n$  a *generalized average* for  $n$  strategies if  $g^n$  is symmetric (i.e. the value does not change if the arguments are permuted) and weakly increasing in all coordinates, and  $g(s, s, \dots, s) = s$  for all  $s \in \text{co}(S)$ .

Natural examples of generalized averages are the power means (also called generalized means) of Hardy, Littlewood and Polya (1952), used e.g. by Cornes (1993) to discuss generalized public good games:

$$g(s_1, \dots, s_n) = \left[ \left( \frac{1}{n} \right) \sum_{i=1}^n s_i^r \right]^{1/r}$$

with  $r \neq 0$  being a parameter. This corresponds to a symmetric CES production function and includes the arithmetic mean ( $r = 1$ ), the geometric mean (or symmetric Cobb Douglas) as a limit when  $r \rightarrow 0$ , and the maximum and minimum functions (as limits when  $r \rightarrow \infty$  and  $r \rightarrow -\infty$ , respectively).<sup>14</sup> However, we will only require the properties stated above and not a specific functional form.

If  $n$  is odd, another well-defined generalized average is given by the median. For arbitrary  $n$ , the largest (or smallest) of the medians also qualifies as a generalized average.

We will consider games where the payoff of a player only depends on the generalized average of all players' actions and the absolute difference between that average and the player's own action. Formally, an *average opinion game* with  $n$  players is given by a common strategy set  $S$  (the opinions), a generalized average  $g^n$  as above, and payoff functions given by

$$u_i(s_i, s_{-i}) = u(g^n(s_i, s_{-i}), |g^n(s_i, s_{-i}) - s_i|)$$

where  $u : \text{co}(S) \times \mathbb{R} \mapsto \mathbb{R}$  is strictly increasing in the average  $g^n$  and strictly decreasing in the difference. In particular, the maximum attainable payoff is given by  $u(s_{\max}, 0)$ .

For example, Van Huyck et al. (1991) consider  $n$  odd and the payoff function

$$u_i(s_i, s_{-i}) = aM(s_i, s_{-i}) - b[M(s_i, s_{-i}) - s_i]^2 + c$$

where  $a, b > 0$  and  $M(\cdot)$  denotes the median.

Average opinion games share with minimum effort games the quality that every symmetric profile is a Nash equilibrium, with the one where all players choose  $s_{\max}$  Pareto-dominating all others. Indeed, if a player  $i$  and all his neighbors adopt  $s_{\max}$ , then  $i$  receives the payoff  $u(g^n(s_{\max}, \dots, s_{\max}), |g^n(s_{\max}, \dots, s_{\max}) - s_{\max}|) = u(s_{\max}, 0)$  which is the maximum attainable payoff.

However, average opinion games describe a qualitatively different situation than minimum effort games. "Actions" (opinions) have no direct cost, but a deviation itself carries a negative payoff.

Consider now average opinion games on networks. For each agent  $i \in I$ , fix a generalized average  $g_i$  for  $n_i = |K(i)| + 1$  strategies (that is, different players might weight opinions following different methods). The payoff to player  $i$  is then given by  $u(g_i(\omega_{K(i)} \cup \{i\}), |g_i(\omega_{K(i)} \cup \{i\}) - s_i|)$  where the function  $u$  is as above; alternatively,  $u$  might vary across agents, as long as the properties given above are fulfilled.

<sup>14</sup>These four particular cases are referred to by Cornes (1993) as average, weaker link, best shot, and weakest link, respectively.

**Theorem 3.** *Consider an average opinion network game played on a local-interaction system with information spillovers. Suppose all agents use imitation rules which are salience-based (B1) and optimistic at the top (B2). Then, if  $w^* > Q_{\min} + 1$  the efficient convention  $\vec{s}_{\max}$  is the unique LRE.*

*Proof:* Consider any element of an absorbing set. Then,  $Q_{\min} + 1$  mutations are sufficient for a transition to  $\vec{s}_{\max}$ . The argument is analogous to Lemma 4: If a player and all his neighbors mutate to  $s_{\max}$ , there exists a positive probability path leading to the efficient convention.

Consider now the efficient convention. Leaving the basin of attraction of this absorbing state requires at least  $w^*$  mutations. The argument is analogous to Lemma 5: Under any salience-based, optimistic at the top imitation rule, we will move back to the convention  $\vec{s}_{\max}$  from any state where some player  $i$  and his neighbors play  $s_{\max}$ .

We conclude that  $CR(\vec{s}_{\max}) \leq Q_{\min} + 1$  and  $R(\vec{e}_{\max}) \geq w^*$ , and hence the statement follows from Lemma 2(ii).

This result hence illustrates a different framework where the same conditions as in Theorem 2 ensure that learning by imitation will be sufficient to guarantee coordination at an efficient level.

## V. Discussion

It is a well-known but seldom formulated stylized fact that, in the right interaction structures, imitation rules tend to lead to efficient outcomes in situations where purely strategic behavior (say, through myopic best-reply rules) leads to inefficient equilibria. This point underlies e.g. the comparison of Kandori et al. (1993) and Robson and Vega-Redondo (1996). In the former, imitation rules lead to risk-dominant equilibria in round-robin tournaments. In the latter, it is shown that if interaction occurs through true random matching, then Pareto efficient equilibria are selected for the same games and behavioral rules. For general network structures and coordination games, the connection between imitation and efficiency is made explicit by Alós-Ferrer and Weidenholzer (2008). There, it is shown that this connection depends on a further element, the presence of information spillovers in the network, and on structural restrictions on the network.

This paper seeks a better understanding of the imitation-leads-to-efficiency stylized fact within a more inherently economic framework, as captured by minimum-effort production games played in networks. We find that, first, in the absence of informational spillovers, efficient outcomes will never be selected by any salience-based imitation rule. Once such spillovers are introduced, we have identified conditions under which agents who learn by imitation might overcome coordination failures.

Our findings could add new insights under which circumstances organizations caught in a “coordination trap” might be turned around. From the viewpoint of organizational design, we have identified two important determinants for achieving coordination at high effort convention: decentrality and information spillovers. While the former can be achieved by splitting up tasks and delegating them to relatively small work teams, the latter can be tackled by introducing a system of benchmarking or best practices or by purposefully designing the organization’s environmental characteristics (from established information channels to the very design of offices) in a way that facilitates the exchange of information.

There are two natural extensions of the research presented here. First, it would be interesting to extend the results to general return functions, not necessarily of the minimum-effort type (or the average

opinion type considered in Section IV). Under reasonable conditions, this will strengthen the selection of efficient outcomes, at the price of losing a general characterization result. Second, it would be desirable to obtain a full characterization of the family of imitation rules which lead to efficiency for appropriate return functions. As observed here, such a family has to contain all salience-based and optimistic imitation rules, but cannot contain all salience-based ones.

## APPENDIX

*Proof of Lemma 6:* For each state  $\omega$ , denote by  $u^*(\omega)$  the maximum payoff realized in  $\omega$ , and by  $E(\omega)$  the set of effort levels employed by players who obtain maximum payoff in  $\omega$ .

We only need to show that, for every absorbing set, we can find a state in this set and a positive probability path from this state to a convention.

*Step 1.* Let  $\hat{\Omega}$  be an absorbing set. Let  $\omega$  be a state in  $\hat{\Omega}$  such that  $u^*(\omega)$  is maximum (among states in the absorbing set). Let  $e$  be the smallest effort level in  $E(\omega)$ , and let  $i$  be a player with  $e_i = e$  and  $u_i(e, \omega) = u^*(\omega)$ . It follows that  $u^*(\omega) \leq (1 - \delta)e$ .

With positive probability, every player in  $K(i)$  imitates some effort level in  $E(\omega)$  by (B1), while (by inertia), player  $i$  stays with  $e$ . We reach a new state  $\omega'$  where all neighbors of  $i$  adopt effort levels larger than or equal to  $e$  (because  $e$  was the minimum of  $E(\omega)$ ). Hence  $u_i(e, \omega') = (1 - \delta)e \geq u^*(\omega)$ . But, since  $\omega' \in \hat{\Omega}$  and  $u^*(\omega)$  was maximum in  $\hat{\Omega}$ , it follows that  $u^*(\omega) = (1 - \delta)e$ .

*Step 2.* Let us return to  $\omega$ . We now know that all neighbors of  $i$  have effort levels larger than or equal to  $e$  and that  $u^*(\omega) = (1 - \delta)e$ . Let  $e'$  be the smallest effort level in  $E(\omega)$  larger than  $e$ . Let  $J$  be the set of players  $j$  with  $e_j = e'$  and  $u^*(e', \omega) = u^*(\omega) = (1 - \delta)e$ . It follows that there exists an effort level  $f' < e'$  such that all  $j \in J$  have at least a neighbor  $k \in K(j)$  with  $e_k = f'$ , which lowers the payoff of  $j$  causing  $u_j(e', \omega) = f' - \delta e' = (1 - \delta)e$ . Thus  $e < f' < e'$  and, by definition of  $e'$ , we have that  $f' \notin E(\omega)$ .

Now consider a state  $\omega''$  reached from  $\omega$  as follows. All players uphold their strategies by inertia except those  $k \in K(j)$  with  $e_k = f'$  for  $j \in J$ . These players will adopt some effort level in  $E(\omega)$  by (B1). Notice that again the payoff of  $i$  remains  $(1 - \delta)e$  (even if some neighbor has changed strategy) and hence  $u^*(\omega'') = (1 - \delta)e$ .

Suppose that for *some* player  $j \in J$ , *all* of their neighbors previously playing  $f'$  have adopted effort levels different from (and hence strictly larger than)  $e$ . This implies that the minimum effort level in  $K(j)$  has strictly increased. Hence in the new state  $\omega'$  we have that  $u_j(e', \omega') > u_j(e', \omega) = u^*(\omega)$ , a contradiction.

Hence, for *all* players  $j \in J$ , at least *one* neighbor previously playing  $f'$  has adopted  $e$ , the minimum effort level in  $E(\omega)$ . Then the minimum effort level in  $K(j)$  has decreased strictly and  $u_j(e', \omega'') < u^*(\omega) = u^*(\omega'')$  for all  $j \in J$ . Hence  $e \in E(\omega'')$  but  $e' \notin E(\omega'')$ .

*Step 3.* Since inertia applies independently to each player, there exists a further state  $\omega'''$ , which can be reached from  $\omega$  as follows. Only the players who changed from  $f'$  in  $\omega$  to  $e$  in  $\omega''$  receive the revision opportunity, and indeed switched to  $e$ . That is, in state  $\omega'''$  the only changes in effort levels have involved effort reductions and hence *no payoff can have increased*. Since also in  $\omega'''$  the maximum payoff is  $(1 - \delta)e$  and is reached by  $i$ , it follows that  $E(\omega''') \subseteq E(\omega)$  but  $e' \notin E(\omega''')$ . That is,  $\omega'''$  has the same maximum payoff but strictly less effort levels leading to it.

Applying the argument to  $\omega'''$ , we can further reduce the set of payoff-maximizing effort levels. Iterating, we find a state  $\omega^* \in \hat{\Omega}$  such that  $u^*(\omega^*) = (1 - \delta)e$  and  $E(\omega^*) = \{e\}$ .

Start from this state  $\omega^*$ . With positive probability, all players in  $M(i) \subseteq K(i)$  receive revision opportunities (and are the only ones to receive them) and imitate  $e$  by (B1). Hence next period  $i$  still earns  $(1 - \delta)e$ , and, by maximality of  $u^*(\omega)$  in  $\hat{\Omega}$ , this must also be the maximum payoff in the new state. Now all agents in  $K^*(i)$  play  $e$  and have only neighbors playing  $e$ , hence earn  $(1 - \delta)e$ , which is the maximum payoff. Since the only change in effort level is to  $e$ , in the new state we also have that the only effort level leading to maximum payoffs is  $e$ . Now all agents in  $M(K^*(i))$  (and only these agents) receive revision opportunity and imitate  $e$ . Thus all agents in  $K^*(K^*(i))$  receive maximum payoff and still the only effort level leading to maximum payoffs is  $e$ . Iterating, by Assumption 1 we eventually reach the convention where every player adopts  $e$ .

## References

- Alós-Ferrer, Carlos**, “Learning, Bounded Memory, and Inertia,” *Economics Letters*, 2008, *101*, 134–136.
- **and Ana B. Ania**, “The Evolutionary Stability of Perfectly Competitive Behavior,” *Economic Theory*, 2005, *26*, 179–197.
- **and G. Kirchsteiger**, “General Equilibrium and the Emergence of (Non) Market Clearing Trading Institutions,” *Economic Theory*, 2010, *44*, 339–360.
- **and Karl Schlag**, “Imitation and Learning,” in P. Anand, P. Pattanaik, and C. Puppe, eds., *The Handbook of Rational and Social Choice*, Oxford University Press, 2009.
- **and Simon Weidenholzer**, “Imitation, Local Interactions, and Efficiency,” *Economics Letters*, 2006, *93*, 163–168.
- **and —**, “Contagion and Efficiency,” *Journal of Economic Theory*, 2008, *143*, 251–274.
- Anderlini, Luca and Antonella Ianni**, “Path Dependence and Learning from Neighbors,” *Games and Economic Behavior*, 1996, *13*, 141–177.
- Anwar, Ahmed**, “On the Co-existence of Conventions,” *Journal of Economic Theory*, 2002, *107*, 145–155.
- Apestequia, Jose, Steffen Huck, and Jörg Oechssler**, “Imitation—theory and experimental evidence,” *Journal of Economic Theory*, September 2007, *136* (1), 217–235.
- Bala, Venkatesh and Sanjeev Goyal**, “Learning from Neighbours,” *Review of Economic Studies*, 1998, *65* (3), 595–621.
- Barrett, Scott**, *Why Cooperate? The Incentive to Supply Global Public Goods.*, Oxford University Press, 2007.
- Battalio, Raymond C., Larry Samuelson, and John Van Huyck**, “Optimization Incentives and Coordination Failure in Laboratory Stag Hunt Games,” *Econometrica*, May 2001, *69* (3), 749–64.
- Blume, Andreas and Andreas Ortmann**, “The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria,” *Journal of Economic Theory*, January 2007, *132* (1), 274–290.
- Blume, Lawrence**, “The Statistical Mechanics of Strategic Interaction,” *Games and Economic Behavior*, 1993, *5*, 387–424.

- Brandts, Jordi and David J. Cooper**, “A Change Would Do You Good . . . An Experimental Study on How to Overcome Coordination Failure in Organizations,” *American Economic Review*, June 2006, *96* (3), 669–693.
- **and** —, “It’s What You Say, Not What You Pay: An Experimental Study of Manager-Employee Relationships in Overcoming Coordination Failure,” *Journal of the European Economic Association*, December 2007, *5* (6), 1223–1268.
- Brock, William A. and Steven N. Durlauf**, “Discrete Choice with Social Interactions,” *Review of Economic Studies*, 2001, *68*, 235–260.
- Bryant, John**, “A Simple Rational Expectations Keynes-Type Model,” *The Quarterly Journal of Economics*, August 1983, *98* (3), 525–28.
- Cabrales, Antonio, Raffaele Miniaci, Marco Piovesan, and Giovanni Ponti**, “Social Preferences and Strategic Uncertainty: An Experiment on Markets and Contracts,” *American Economic Review*, 2009, *forthcoming*.
- Cooper, Russell W. and Andrew John**, “Coordinating Coordination Failures in Keynesian Models,” *The Quarterly Journal of Economics*, August 1988, *103* (3), 441–63.
- , **Douglas V. Dejong, Robert Forsythe, and Thomas W. Ross**, “Selection Criteria in Coordination Games: Some Experimental Results,” *American Economic Review*, March 1990, *80* (1), 218–33.
- Cornes, Richard**, “Dyke Maintenance and Other Stories: Some Neglected Types of Public Goods,” *The Quarterly Journal of Economics*, 1993, *108* (1), 259–271.
- Crawford, Vincent P**, “An ‘Evolutionary’ Interpretation of Van Huyck, Battalio, and Beil’s Experimental Results on Coordination,” *Games and Economic Behavior*, 1991, *3*, 25–59.
- , “Adaptive Dynamics in Coordination Games,” *Econometrica*, January 1995, *63* (1), 103–43.
- **and Hans Haller**, “Learning How to Cooperate: Optimal Play in Repeated Coordination Games,” *Econometrica*, May 1990, *58* (3), 571–95.
- Devetag, Giovanna and Andreas Ortmann**, “When and why? A critical survey on coordination failure in the laboratory,” *Experimental Economics*, September 2007, *10* (3), 331–344.
- Durieu, Jacques and Philippe Solal**, “Adaptive Play with Spatial Sampling,” *Games and Economic Behavior*, 2003, *43*, 189–195.
- Ellison, Glenn**, “Learning, Local Interaction, and Coordination,” *Econometrica*, 1993, *61*, 1047–1071.
- , “Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution,” *Review of Economic Studies*, 2000, *67*, 17–45.
- **and Drew Fudenberg**, “Rules of Thumb for Social Learning,” *Journal of Political Economy*, 1993, *101*, 612–643.
- Ely, Jeffrey C.**, “Local Conventions,” *Advances in Theoretical Economics*, 2002, *2*, 1–30.
- Erev, Ido and Greg Barron**, “On Adaptation, Maximization, and Reinforcement Learning among Cognitive Strategies,” *Psychological Review*, 2005, *112*, 912–931.
- Eshel, Ilan, Larry Samuelson, and Avner Shaked**, “Altruists, Egoists, and Hooligans in a Local Interaction Model,” *American Economic Review*, 1998, *88*, 157–179.
- Feri, Francesco, Bernd Irlenbusch, and Matthias Sutter**, “Efficiency Gains from Team-Based Coordination - Large-Scale Experimental Evidence,” *American Economic Review*, 2009, *forthcoming*.
- Freidlin, Mark I. and Alexander D. Wentzell**, *Random Perturbations of Dynamical Systems*, 2nd

- Ed.*, New York: Springer Verlag, 1988.
- Hardy, Godfrey H., John E. Littlewood, and George Polya**, *Inequalities*, Cambridge, England: Cambridge University Press, 1952.
- Hirshleifer, Jack**, “From weakest-link to best-shot: The voluntary provision of public goods,” *Public Choice*, January 1983, *41* (3), 371–386.
- Holmstrom, Bengt**, “Moral Hazard in Teams,” *Bell Journal of Economics*, Autumn 1982, *13* (2), 324–340.
- Hume, David**, *A Treatise of Human Nature* 1739.
- Jackson, Matthew O. and Asher Wolinsky**, “A Strategic Model of Social and Economic Networks,” *Journal of Economic Theory*, 1996, *71*, 44–74.
- Janssen, Maarten C. W. and Ewa Mendys-Kamphorst**, “The price of a price: on the crowding out an in of social norms,” *Journal of Economic Behavior and Organization*, 2004, *55*, 377–395.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob**, “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 1993, *61*, 29–56.
- Karlin, Samuel and Howard M. Taylor**, *A First Course in Stochastic Processes, 2nd Ed.*, San Diego: Academic Press, 1975.
- Knez, Marc and Colin Camerer**, “Creating Expectational Assets in the Laboratory: Coordination in “Weakest-Link Games,”” *Strategic Management Journal*, 1994, *15*, 101–119.
- Kremer, Michael**, “The O-Ring Theory of Economic Development,” *The Quarterly Journal of Economics*, August 1993, *108* (3), 551–75.
- Morris, Stephen**, “Contagion,” *Review of Economic Studies*, 2000, *67*, 57–78.
- Robles, Jack**, “Evolution and Long Run Equilibria in Coordination Games with Summary Statistic Payoff Technologies,” *Journal of Economic Theory*, 1997, *75*, 180–193.
- Robson, Arthur J. and Fernando Vega-Redondo**, “Efficient Equilibrium Selection in Evolutionary Games with Random Matching,” *Journal of Economic Theory*, 1996, *70*, 65–92.
- Schlag, Karl**, “Why Imitate, and if so, how? A Boundedly Rational Approach to Multi-armed Bandits,” *Journal of Economic Theory*, 1998, *78*, 130–156.
- Simon, Herbert**, *Administrative Behavior: A Study of Decision Making-Processes in Administrative Organizations.*, New York: Macmillan, 1947.
- Skyrms, Brian**, “Trust, risk, and the social contract,” *Synthese*, January 2006, *160* (1), 21–25.
- Sobel, Joel**, “Economists’ Models of Learning,” *Journal of Economic Theory*, 2000, *94* (2), 241–261.
- Steiner, William W., James H. Hain, Howard E. Winn, and Paul J. Perkins**, “Vocalizations and feeding behavior of the killer whale (*Orcinus orca*),” *Journal of Mammalogy*, 1979, *60*, 823–827.
- Van Huyck, John, Raymond C. Battalio, and Richard O. Beil**, “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure,” *American Economic Review*, 1990, *80* (1), 234–48.
- , – , and – , “Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games,” *The Quarterly Journal of Economics*, August 1991, *106* (3), 885–910.
- Vega-Redondo, Fernando**, “The Evolution of Walrasian Behavior,” *Econometrica*, 1997, *65*, 375–384.
- Weber, Roberto A.**, “Managing Growth to Achieve Efficient Coordination in Large Groups,” *American Economic Review*, March 2006, *96* (1), 114–126.
- Young, Peyton**, “The Evolution of Conventions,” *Econometrica*, 1993, *61*, 57–84.