Emergence of Collective Cooperation and Networks from Selfish-Trust and Selfish-Connections

Korosh Mahmoodi (koroshm@andrew.cmu.edu)
Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University
5000 Forbes Ave,
Pittsburgh PA 15213 USA

Cleotilde Gonzalez (coty@cmu.edu)
Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University
5000 Forbes Ave,
Pittsburgh PA 15213 USA

Abstract

Emergence of collective cooperation in an inherently selfish society is a paradox that has preoccupied biologists, sociologists, and cognitive scientists alike for centuries. We propose a computational model and demonstrate through simulations how collective cooperation can emerge from selfish interests: the goal of improving each individual's own rewards. We also demonstrate how the same selfish interests lead to the dynamic emergence of a network of interconnected agents. Our model includes two simple mechanisms: Selfish-Trust (ST) and Selfish-Connection (SC). ST involves the possibility of relying on others in a society of agents when it is beneficial to the individual, and SC involves the possibility of connecting to other agents when those agents help improve the individual's own benefit. Our simulation results suggest that collective cooperation can emerge from ST and a complex dynamic network can emerge from ST and SC. The simulated data demonstrate an important property of many living organisms: patterns of temporal complexity, which are essential to transfer information among agents of any society of living beings.

Keywords: Altruism Paradox, Emergence of Cooperation, Selfishness, Trust, Networks, Artificial Intelligence

Introduction

For Charles Darwin (Darwin, 1871) altruism remained a paradox: the act of sacrificing an individual's own benefit for the benefit of the collective community of living organisms was regarded as a contradiction to evolutionary theories. The dilemma of emergence of cooperative behavior in situations in which there is a large incentive to defect for the individual benefit has been widely studied in sociology and cognitive sciences. The Prisoner's dilemma (PD) has been a leading metaphor for the study of the evolution of cooperative behavior in populations of selfish in which selfishness is more rewarded in the short-term (M. Nowak & Sigmund, 1993; Gonzalez, Ben-Asher, Martin, & Dutt, 2015).

The PD, dates back to the early development of Game Theory (Rapoport & Chammah, 1965), and it is a common abstraction of the essential elements of many naturalistic situations involving cooperative behavior. It is generally represented with a payoff matrix that provides payoffs according to the actions of two players (see Table 1). When both players cooperate, each of them gains the payoff $\prod(t) = R$, and when both players defect, each of them gains $\prod(t) = P$. If the player i defects and player j cooperates, player i gains the payoff $\prod(t) = T$ and player j gains the payoff $\prod(t) = S$ and

		Player j C	Player j D
Player i	C	(R,R)	(S,T)
	D	(T,S)	(P,P)

Table 1: The general payoffs of PD game. The first value of each pair is the payoff of agent i and the second value is the payoff of the agent j.

vice versa. The constraints on the values of the payoffs in the PD are $T > R > P > S$ and $S + T < 2R$. The temptation to defect is established by setting the condition $T > R$.

The dilemma is that, while the longer-term best mutual action is to cooperate, in the short-term each individual would prefer to defect because it indicates a higher reward to the individual. Assuming that the other player also searches for its own individual maximum reward, the pair will end up in a $D-D$ situation with the minimum payoff for the two players $2P$.

How do individuals realize that cooperation is mutually beneficial in the long-term? this question has been addressed by many researchers, at various levels of inquiry, involving pairs of agents (Gonzalez et al., 2015; Moisan, ten Brincke, Murphy, & Gonzalez, 2018) as well as larger social networks (M. Nowak & Sigmund, 1993). Research suggests that, at the pair level, people dynamically adjust their actions according to their observations of others' actions and outcomes; at the network level, research suggests that the emergence of cooperation may be explained from *network reciprocity*, where individuals play with those agents with whom they are already connected in a network structure. The demonstration of how social networks and structured populations with explicit connections foster cooperation was introduced by Nowak and May (1992). Alternative models based on Network reciprocity assume agents in a network play the PD with the agents with whom they have specific interconnections. Agents act by copying the strategy of the richest neighbor, basing their decisions on the observation of the others' payoffs. Thus, network reciprocity depends on the existence of a network structure (an already predefined set of connections among agents) and on the awareness of the behavior and payoffs of interconnected agents. Network reciprocity assumes that the evolution of cooperation is a function of the difference between the payoffs of the interacting agents.

Thus, past research assumes that the emergence of collective cooperation requires the observation of others' actions and/or outcomes and the existence of predefined connections among agents. Indeed, empirical work suggests that the emergence of cooperation depends on the level of information available to each agent (Martin, Gonzalez, Juvina, & Lebiere, 2014); and the less information about other agents exist, the more difficult, and perhaps the longer it takes, for cooperation to emerge (Martin et al., 2014; Rapoport & Chammah, 1965). However, other experiments suggest that humans do not consider others' payoffs when making their decisions, and that a network structure does not influence the final cooperative outcome (Fischbacher, Gächter, & Fehr, 2001). Indeed, in many aspects of life, we influence others through our choices and others' choices affect us, but we are not necessarily aware of the exact actions and rewards received by others affecting us. For example, when a member of society avoids air travel in order to reduce the individual's carbon footprint, he or she might not be able to observe whether others are reducing their air travel too, yet rely on decisions others make, influencing the community as a whole. It is thus, difficult to explain how behaviors can be self-perpetuating even when the source of influence is unknown (Martin et al., 2014).

In this research, we aim at advancing our understanding of the emergence of collective cooperation in the absence of explicit knowledge of others' actions and outcomes, and in the absence of an explicit predefined network structure that connects agents in a society. We introduce an algorithm (*Living Thing*, LT) to demonstrate that collective cooperation can emerge and survive between agents, out of selfishness (i.e., the individual's need to act on their own personal benefit), and in the absence of others' information (i.e., without a need to any predefined network). We aim at developing hypotheses that can help resolve social dilemmas that exist in the real world. For example, if we understand how collective cooperation emerges only from the decisions of each individual, we could propose solutions that reduce the dilemmas in social problems such as littering in public places or the lack of contributions to a reduction of CO_2 in the atmosphere (Martin et al., 2014).

A LT agent will act according to the reinforcement of its own past actions (Reinforcement Learning, RL), but it will rely on two mechanisms that may overwrite the agent's RL actions: Selfish-Trust (ST) and Selfish-Connection (SC). ST is a decision to follow or rely on other agent's decision expecting that it will improve the own agent's reward with respect to the agent's own previous payoff. ST is expected to turn the initially defector agents to agents that cooperate most of the time. SC is a mechanism that helps agents learn who to play with: agents increase the propensity of playing with the same other agent if the payoff received after playing with that other agent is higher than the agent's own previous payoff.

Past models of network formation rely on a concept of preferential attachment (PA) (Barabási & Albert, 1999), which uses rules according to which an agent would have a higher chance of linking with other agents that already have many links (i.e., high reputation nodes). In contrast, LT demonstrates that such propensities to connect to other agents emerge dynamically, according to the experienced benefits that the other agent brings to the individual's own benefit (SC).

We carry an analysis of the emergence of cooperation from these mechanisms. The simulation results hint at how to explain emergent collective cooperation from individual selfish interests. An important hypothesis emerging from this work is that cooperation can emerge and survive out of the selfishness of agents even when there are no specific awareness of outcomes of other agents, and that a network structure can emerge dynamically from the connections guided by self individual interests.

Living Thing (LT) Algorithm

Figure 1 shows one-time cycle of the LT algorithm from the perspective of one of the agents, agent i, but every step is executed for agent j simultaneously. In Step 1, a pair of randomly selected agents i and j "agree" to play. Only one pair of agents is selected at each time cycle. The following are general notations in the algorithm: V_i is the decision of the agent to Cooperate (C) or Defect (D); r represents a random number in the interval [0 1] which should be generated whenever it is called in the algorithm; Δ is a positive number that represents an increase in three possible cumulative tendencies: to play C or D, to trust the paired agent or not, or to play again with a previous agent. These cumulative tendencies increase by Δ when the benefit of the agent i changes with respect to its previous benefit and if there is no change then we set $\Delta = 0$ which means no change happens in the system. Δ, in general, can be a function of the difference between two past payoffs of the agent and it can be different for different cumulative tendencies, but it does not change the general results presented later (the form of sensitivity to payoffs is important when two systems interact with each other which is out of scope of this paper).

The following steps are executed in each time t of the algorithm:

Pairing Agents (1)

Agent i and agent j get picked randomly. Agent i at time t has the propensity $P_{ij}(t) = M_{ij}(t)/\sum_k M_{ik}(t)$ to play with agent j. $0 < P_{ij}(t) < 1$. $M_{ij}(t)$ is cumulative tendency for agent i to pick agent j to play at time t. This cumulative tendency changes at step 7 according to the last two payoffs received by agent i.

At the same time, agent j has a propensity $P_{ji}(t) = M_{ji}(t)/\sum_k M_{jk}(t)$ to play with agent i. $0 < P_{ji}(t) < 1$. Two agents i and j pair-up if two random numbers, $0 < r_1 < 1$ and $0 < r_2 < 1$, satisfy inequalities $r_1 < P_{ij}(t)$ and $r_2 < P_{ji}(t)$.

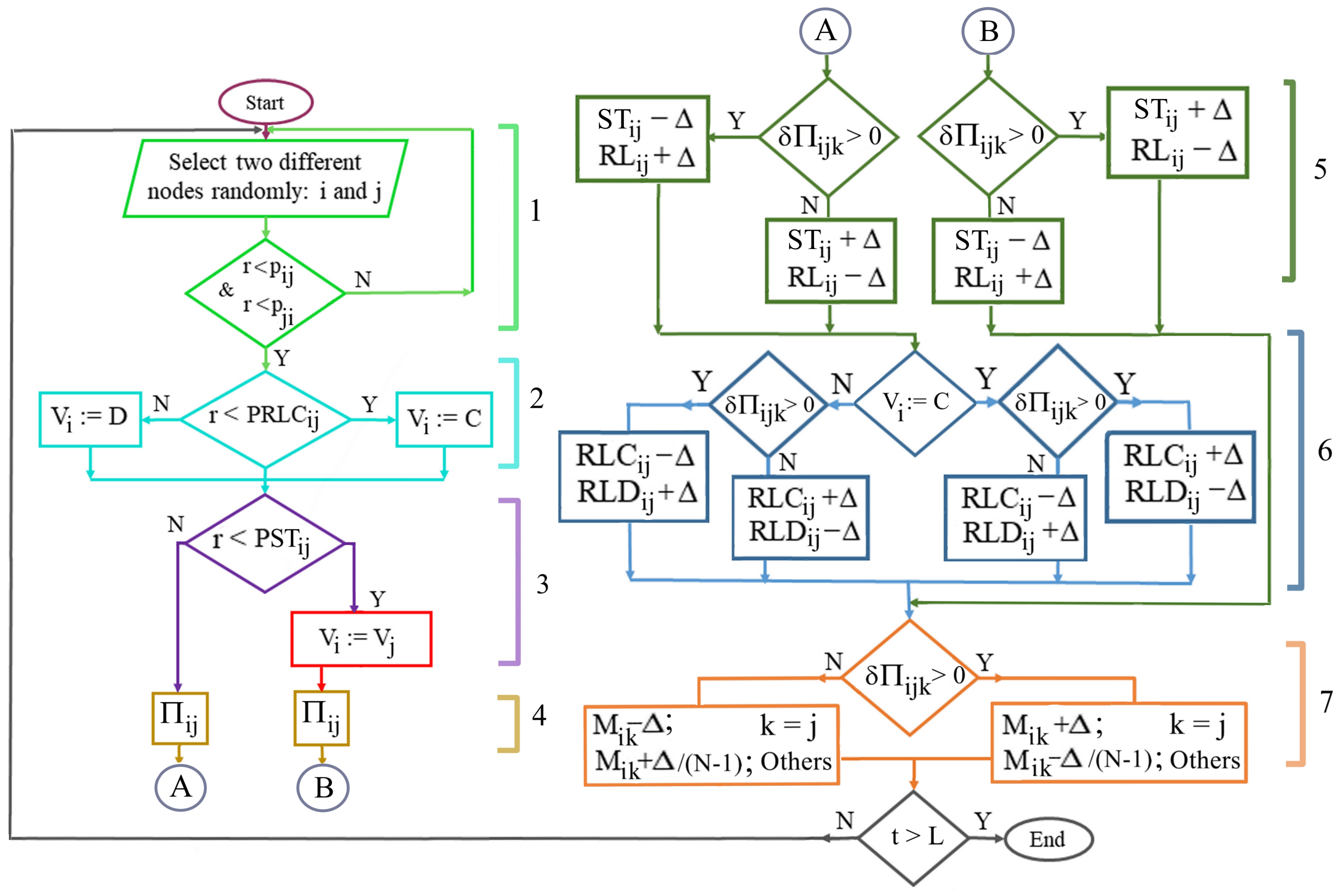

Figure 1: Flowchart of the LT algorithm. ”Y” and ”N” letters represent ”Yes” and ”No”, respectively.

Otherwise another two agents are randomly selected in each time t.

Reinforcement Learning (RL) (2)

Agent i initially selects an action by reinforcement learning (RL): This agent has the propensity $PRLC_{ij}(t) = RLC_{ij}(t)/(RLC_{ij}(t) + RLD_{ij}(t))$ to pick C and has the propensity $PRLD_{ij}(t) = RLD_{ij}(t)/(RLC_{ij}(t) + RLD_{ij}(t))$ to pick D as it's next potential decision.

$RLC_{ij}(t)$ and $RLD_{ij}(t)$ are cumulative tendencies for agent i playing with agent j, at time t, for choice C or D, respectively. These cumulative tendencies change at step 6 based on the last two payoffs of agent i.

To select the action, a random number r gets picked and if $r < PRLC_{ij}(t)$ then it's next decision will be C otherwise will be D. The same process applies for agent j.

Selfish-Trust (ST) (3)

Instead of executing the decision determined by RL in Step 2, agent i has a chance to trust the decision made by agent j made using RL in Step 2, and with whom agent i is paired with. The propensity that agent i relies on the decision of agent j is: $PST_{ij}(t) = ST_{ij}(t)/(ST_{ij}(t) + RL_{ij}(t))$. $ST_{ij}(t)$ and $RL_{ij}(t)$ are cumulative tendencies for agent i to execute its choice based on ST from agent j, at time t, or to execute its choice based on RL respectively. These cumulative tendencies update in step 5 based on the last two payoffs of agent i. Again, if a random number r is less than $PST_{ij}(t)$ then ST happens.

Evaluating Own Payoffs (4)

At time t agent i after executing its C or D action while playing the PD game with agent j, receives the payoff $\prod_{ij}(t)$. The last two payoffs of the agent i are used to determine the changes in its accumulative tendencies: $\delta\prod_{ijk}(t) = \prod_{ij}(t) - \prod_{ik}(t-1)$, where agent (k) is the agent that played with agent i in trial $t-1$. In the flowchart we showed this quantity as $\delta\prod$.

Update of cumulative tendency of ST or RL (5)

If agent i used ST and after playing with agent j its payoff is higher than its previous payoff, $\delta\prod_{ijk}(t) > 0$, then the accumulative tendencies ST_{ij} and RL_{ij} , for the next time agent i and j, change to $ST_{ij} + \Delta$ and $RL_{ij} - \Delta$. The same happens for agent j.

Similarly, if agent i used RL and after playing with agent j its payoff is higher than its previous payoff, $\delta\prod_{ijk}(t) > 0$, then the accumulative tendencies RL_{ij} and ST_{ij}, for the next

		Player j	
		C	D
Player i	C	$(1,1)$	$(0,1.9)$
	D	$(1.9,0)$	$(0,0)$

Table 2: The payoffs of PD game used in the simulations. The first value of each pair is the payoff of agent i and the second value is the payoff of its pair, agent j.

time agent i and j pair up, change to $RL_{ij}+\Delta$ and $ST_{ij}-\Delta$. The same happens for agent j.

Update of cumulative tendency to choose C or D (6)

Step 6 is only active if the agent decided to use RL is step 3. If agent i played with agent j and received a payoff higher than its previous payoff, $\delta\prod_{ijk}(t)>0$, and this happened because agent i played C, then the accumulative tendencies RLC_{ij} and RLD_{ij}, for the next time agent i and j pair up, change to $RLC_{ij}+\Delta$ and $RLD_{ij}+\Delta$. If the increase happened because agent i played D, then the accumulative tendencies RLD_{ij} and RLC_{ij} , for the next time agent i and j pair up, change to $RLD_{ij}+\Delta$ and $RLC_{ij}-\Delta$. The same happens for agent j.

Selfish-Connection (SC) (7)

In this step the cumulative tendency to play with a specific agent changes. If agent i, after playing with agent j, receives higher benefit with respect to its previous payoff, $\delta\prod_{ijk}(t)>0$, then the cumulative tendency of pairing with agent j, M_{ij}, increases to $M_{ij}+\Delta$ and for the rest of the cumulative tendencies decreases to $M_{il}-\Delta/(N-1), l\neq j$. If $\delta\prod_{ijk}(t)<0$, then the cumulative tendency of pairing with agent j, M_{ij}, decreases to $M_{ij}-\Delta$ and for the rest of the cumulative tendencies increases to $M_{il}+\Delta/(N-1), l\neq j$. The same happens for agent j.

Simulation Methods

We studied a system with $N=100$ agents. Initially all the agents are defectors, have payoff of zero, have more chance to stay as defector; $RLC_{ij}(0)=1$, $RLD_{ij}(0)=99$, have more chance to use RL over ST; $ST_{ij}(0)=1$, $RL_{ij}(0)=99$, and have equal chance to pair up with other agents; $M_{ij}(0)=100$. We set $\Delta=10$. Δ is the property of the system and shows the sensitivity of the agent to the feedback from its two last payoffs. Smaller Δ decreases the rate of reaching to cooperation but doesn't change the dynamical properties of the system. The payoffs matrix used has the values shown in Table 2 as suggested by Gintis (2009): $R=1$, $P=0$ and $S=0$. So, the maximal possible value of T is 2. We selected the value $T=1.9$, which gives a very strong incentive to defect.

Results

Emergence of Cooperation from ST

Here we show that simple mechanism of ST can lead agents who play PD game (which has a high tendency to defect) toward cooperation. Figure 2 shows the proportion of cooperation in simulations that rely only on the RL mechanism (blue curve) compared to the emergence of cooperation when the simulations rely on the additional ST mechanism.

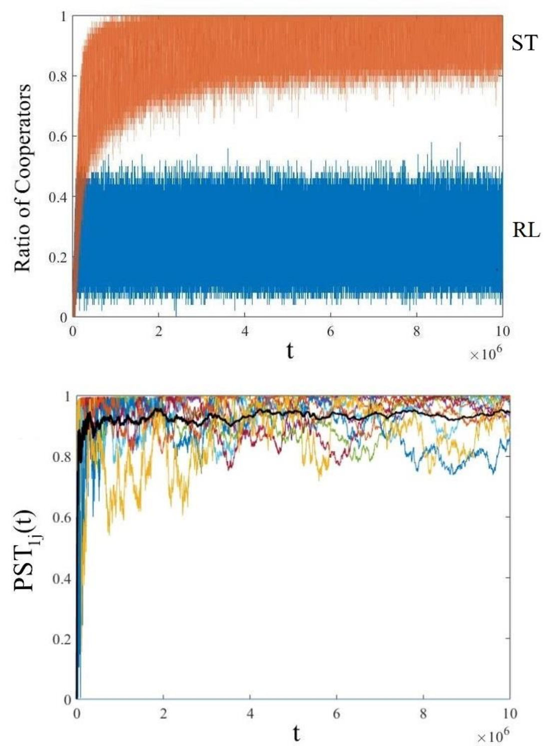

Figure 2: Top panel: the orange curve is the ratio of Cooperators vs. time for $N=100$ agents, randomly paired up to play PD game and used ST (steps 7 in LT algorithm inactive) for updates of the strategy and cumulative tendencies. The blue curve is the ratio of cooperators for $N=100$ agents, randomly paired up to play PD game and just used RL for the decision making process (steps 3 and 7 in LT algorithm inactive). Bottom panel: emergence and evolution of the probability of trust of unit 1 the other 9 agents in a system with 10 agents, used ST to update their strategies and cumulative tendencies. The thicker, black curve in this figure is the average of all the nine probabilities of trust

The blue curve in the top panel of Figure 2 shows the time evolution of the ratio of cooperators when RL is the only mechanism that agents use to update their strategy (steps 3 and 7 are inactive).

The orange curve in the top panel of Figure 2 shows the emergence of cooperation between agents when at any trial two of them (out of 100) paired up randomly and used ST (steps 7 in LT algorithm inactive) to update their strategy (C or D) and their cumulative tendencies. The system reached its dynamic equilibrium after about 2×10^6 and sustained around the average ratio of cooperation of 0.9.

This shows the effect of ST on improving the level of cooperation compared to only RL. In the absence of ST the ratio of cooperators fluctuates around 0.3 which means the majority of agents are defectors.

The nine curves in the bottom panel of Figure 2 are the chances that one of the agents might trust others in a system with 10 agents. The thicker, black curve is the average of the STs between agent 1 and the other nine agents. The chances of ST increased and sustained to about 0.9. This means that the agent learns that ST has benefit for it (and for the whole society). The payoff of the individual is not shown here because it is proportional to the level of cooperation: more cooperation results in more payoff for individual agent and for the emerged group. In conclusion, cooperation emerges and survives because ST lets the strategies to spread between the agents, if it benefits them individually.

Emergence of Complexity over time

The analysis of the fluctuations of a time series gives us a measure for the complexity of the system. We define the events in the time series as the times the time series crosses its mean value. The distribution of the time intervals between the two consecutive events is of interest (Figure 3).

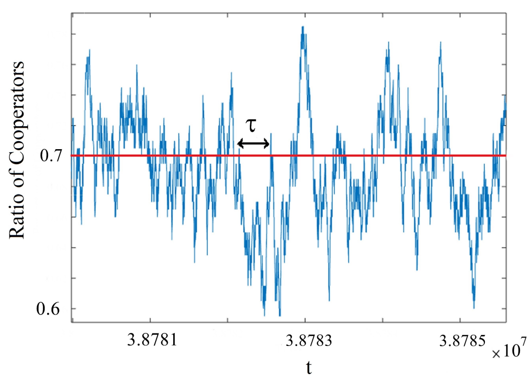

Figure 3: Demonstration of defining events in a time series. The blue curve is a zoomed in part of the Ratio of Cooperators' time series (as an example) and the horizontal red line is its mean value. Whenever the time series crosses the mean value is defined as an event. The distribution of the time intervals between two consecutive events gives a measure for the complexity of the system (complexity index μ).

We collect all the time intervals between two consecutive events (τ's) and evaluate the probability density function (PDF) $\psi(\tau)$. If the resulting distribution is Poisson then the dynamic of the system is random and obeys ordinary statistics. But, if the PDF is a power law; $\psi(\tau) \propto 1/\tau^{\mu}$, then the dynamics falls in the category of the complex systems, for example, the dynamics of the brain. The parameter μ, the slope of the Inverse power law in a log-log plot, is a measure for the complexity of the time series: when $\mu > 3$ then the system is ordinary while for $1 < \mu < 3$ there is Ergodicity Breaking and the system does not obey ordinary statistics.

Temporal criticality is crucial for transfer of information between two intelligence systems (Aquino, Bologna, West, & Grigolini, 2011). To measure the complexity index μ of the time series of the ratio of cooperators on the four cases of emergence of cooperation (top panels in Figure 2 and Figure 6), we studied their fluctuations in the asymptotic regime ($t > 5 \times 10^6$) around their mean value.

Figure 4 shows that the time series of the ratio of cooperators for both cases where agents used ST or just used RL, time series of top panel in Figure 2, have inverse power law PDF with the same complexity index of $\mu = 1.3$ (which falls in the interval $1 < \mu < 3$). The difference is that the linear part of the distribution for the first case, where ST exists, is extended toward larger τ's which means the system is more complex respect to the latter case.

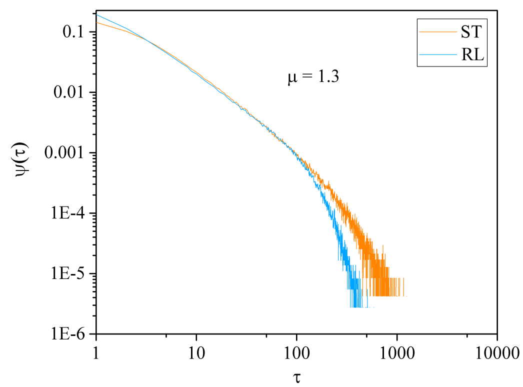

Figure 4: The PDF of the time intervals between the two consecutive events of the time series of Figure 2 (in log-log scale).

Figure 5 shows the PDFs of the time series of the ratio of cooperators for the cases where agents are using ST SC (LT) or RL SC for their evolution, top panel in Figure 6. Both cases have an inverse power law PDF. The PDF of the first case where ST and SC are both active shows very extended linear part with complexity index of $\mu = 1.73$. This complexity is similar to that of the time series of the living things (Allegrini, Paradisi, Menicucci, & Gemignani, 2010).

On the other hand, the complexity of the second case is similar to the system where agents were using just RL (blue curve of Figure 5). This means that SC could increase the complexity of the system where ST was already in action.

Emergence of connections with other agents

In this section we add another level of learning to ST by letting the agent to find the agents which playing with them increases its payoff respect to its previous payoff (all sections of the LT algorithm active). The aim is to show that a dynamic complex network emerges naturally and from the ST and SC mechanisms of the LT model.

The blue curve in the top panel of Figure 6 is the ratio of cooperators when SC is added to RL (section 3 of the LT algorithm inactive). The blue curves in the top panels of Figure 2 and Figure 6 are very similar, which means that adding

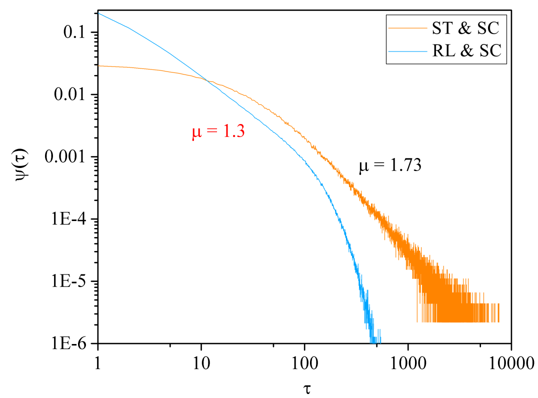

Figure 5: The PDF of the time intervals between the two consecutive events of the time series of Figure 6 (in log-log scale).

the ability to the agents to select their partner is not favoring cooperation when ST is inactive. The reason is that only when ST is active the level of cooperation between the agents increases and because of mutual benefit, which is stable, an agent can rank the links by changing the chance of playing with other agents based on the increase on its last two payoffs.

To illustrate this, we plot the chances of an agent (called agent 1) to trust the other 9 agents using LT algorithm in the bottom panel of Figure 6. This figure shows that the agent trusts some of the agents more than others, most of the time, and only from time to time it trusts other agents. But later on, the agent starts to trust some agents most of the time. The thicker, black curve in this figure is the chance that the agent (1) connects to the agent corresponding to the purple curve. The similarities between the red and purple curves show that the agent 1 learns to connect with the agent which it is most trusting, most of the time. The preferential connections here are dynamic and are based on the perception of the benefit that an agent receives from the other agents. This process creates connections among agents that are dynamic. Some connections become stronger and others become weaker according to the SC mechanism.

Figure 7 demonstrates the chance of a random agent (represented by a dot in the center) pairing with the other 99 agents at two different times: $t = 10^2$ (top panel) and 10^6 (bottom panel). The thickness of the lines represents the chance of the pairings. At $t = 10^2$, we observe an almost uniform distribution of the probability of the connections of an agent to the others (top panel); but later, the agent learned to prefer to connect to some partners more than to others (bottom panel).

Figure 8 shows the probability density function of the pairings for all the agents in Figure 7 (bottom panel). This distribution, plotted in a log-log scale, shows an inverse power law $\propto 1/P^{\beta}$ with complexity index of $\beta = 1.3$, rather than having a Poisson distribution, showing that the emerged network is

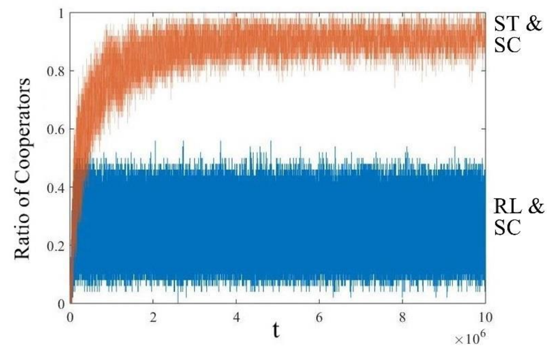

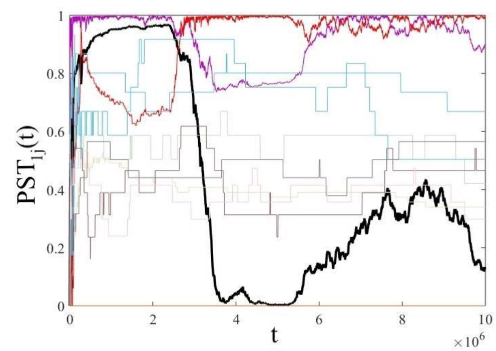

Figure 6: Top panel: orange curve: the ratio of Cooperators vs. time for $N = 100$ agents which in addition to ST they use SC to pick their partner to play PD game (all sections of the LT algorithm active). The Blue curve is the ratio of cooperators, paired using SC, but updated their strategies only with RL (section 3 of the LT algorithm inactive). Bottom panel: emergence and evolution of trust of unit 1 the other 9 agents in a system with 10 agents using the LT algorithm (ST SC) to update their strategies and cumulative tendencies. The thicker, black curve in this figure is the chance of agent 1 to play with the agent which is most trusting (the purple curve close to 1).

complex.

Discussion and Implications of Results

The novelty of the LT algorithm is the demonstration of how collective behavior can emerge from Selfish Trust and how a network can emerge from Selfish connections; in the absence of an explicit a-priory network structure, and in the absence of explicit awareness of others' outcomes. LT uses ST that adapts to increase or decrease the chance that agent i will trust the strategy of agent j, if that strategy is beneficial or detrimental for agent i itself. LT also uses SC that adapts to increase or decrease the chance of agent i to connect to agent j, if agent j has contributed or not to the own benefit of agent i. This means that selfishness of agent i is used as the main learning incentive: If the payoff of the agent i increased with respect to its previous payoff then it will increase the likelihood of repeating its last action. This control of the dynamics is internal and emergent according to the self-interest of the

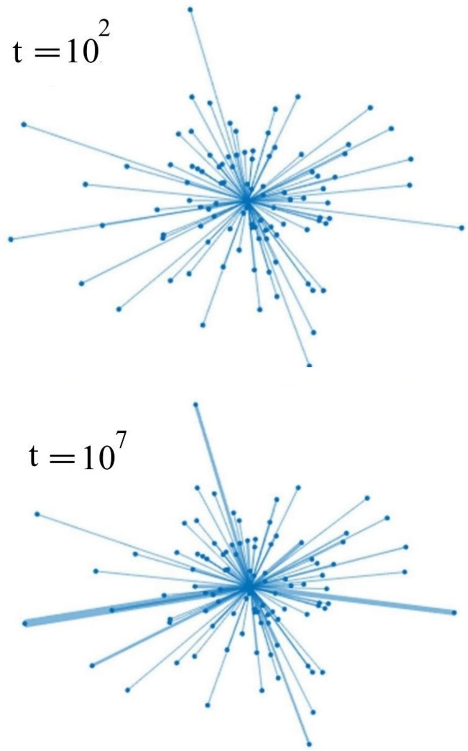

Figure 7: Each dot represents an agent. The dot in the center shows the agent of interest and the other 99 dots are connected to it with lines. The thickness of each line is proportional to the chance of the corresponding agents to pair up at time $t = 10^2$ (top figure) and $t = 10^7$ (bottom figure).

agents, leading the system to self-organization. The role of ST is to spread the strategies between the agents, if it is increasing the payoff of individuals with respect to their previous one. Adding SC to ST lets each agent learn which agents to connect to, in order to increase its own payoff with respect to its last one. SC can improve the complexity of the system by forming a dynamic network of chances of pairings, which results in an inverse power law PDF with complexity index of $\beta = 1.3$. The self-organized system evolved by LT host events with inverse power law PDF of the interval between the consecutive events with complexity index $\mu = 1.73 < 3$. This is the main property of dynamic complex systems which makes them able to transfer information and match to another.

Acknowledgments

This research was supported by the Army Research Office, Network Science Program, Award Number: W911NF1710431.

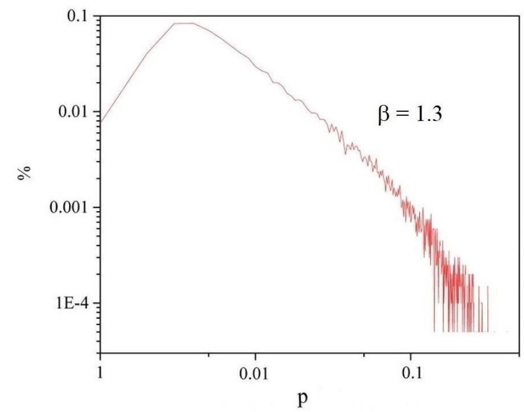

Figure 8: The probability density function of the pairings between all possible pairs at $t = 10^7$, in a log log scale.

References

Allegrini, P., Paradisi, P., Menicucci, D., & Gemignani, A. (2010). Fractal complexity in spontaneous eeg metastable-state transitions: new vistas on integrated neural dynamics. *Frontiers in physiology*, *1*, 128.

Aquino, G., Bologna, M., West, B. J., & Grigolini, P. (2011). Transmission of information between complex systems: 1/f resonance. *Physical Review E*, *83*(5), 051130.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *science*, *286*(5439), 509–512.

Darwin, C. (1871). *The descent of man. the modern library.* Random House Inc., NY.

Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics letters*, *71*(3), 397–404.

Gintis, H. (2009). The bounds of reason game theory and the unification of the behavioral sciences.

Gonzalez, C., Ben-Asher, N., Martin, J. M., & Dutt, V. (2015). A cognitive model of dynamic cooperation with varied interdependency information. *Cognitive Science*, *39*(3), 457-495. doi: 10.1111/cogs.12170

Martin, J. M., Gonzalez, C., Juvina, I., & Lebiere, C. (2014). A description-experience gap in social interactions: Information about interdependence and its effects on cooperation. *Journal of Behavioral Decision Making*, *27*(4), 349-362. doi: 10.1002/bdm.1810

Moisan, F., ten Brincke, R., Murphy, R., & Gonzalez, C. (2018). Not all prisoner's dilemma games are equal: Incentives, social preferences, and cooperation. *Decision*, *5*(4), 306,322.

Nowak, M., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, *364*(6432), 56–58.

Nowak, M. A., & May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, *359*(6398), 826.

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's Dilemma: A Study in Conflict and Cooperation*. Ann Arbor: University of Michigan Press.