

# Globalization and the rise and fall of cognitive control

Mohsen Mosleh <sup>1✉</sup>, Katelynn Kyker <sup>2</sup>, Jonathan D. Cohen<sup>3,4</sup> & David G. Rand <sup>1,5,6</sup>

The scale of human interaction is larger than ever before—people regularly interact with and learn from others around the world, and everyone impacts the global environment. We develop an evolutionary game theory model to ask how the scale of interaction affects the evolution of cognition. Our agents make decisions using automatic (e.g., reflexive) versus controlled (e.g., deliberative) cognition, interact with each other, and influence the environment (i.e., game payoffs). We find that globalized direct contact between agents can either favor or disfavor control, depending on whether controlled agents are harmed or helped by contact with automatic agents; globalized environment disfavors cognitive control, while also promoting strategic diversity and fostering mesoscale communities of more versus less controlled agents; and globalized learning destroys mesoscale communities and homogenizes the population. These results emphasize the importance of the scale of interaction for the evolution of cognition, and help shed light on modern challenges.

<sup>1</sup>Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02138, USA. <sup>2</sup>Department of Psychology, Yale University, New Haven, CT 06520, USA. <sup>3</sup>Princeton Neuroscience Institute, Princeton, NJ 08540, USA. <sup>4</sup>Department of Psychology, Princeton, NJ 08540, USA. <sup>5</sup>Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02138, USA. <sup>6</sup>Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02138, USA. ✉email: [mmosleh@mit.edu](mailto:mmosleh@mit.edu)

Humanity is experiencing a period of unprecedented innovation and technological sophistication, as well as unprecedented interconnectedness and globalization. However, many of these innovations and technological advancements are often misused in a short-sighted fashion (e.g., the over-prescription of antibiotics, abuse of opiates, and over-reliance on fossil fuels<sup>1</sup>). In parallel, a substantial backlash is occurring against the kind of expertise and logical, deliberative thinking that produced such innovations (e.g., an increasing rejection of scientific expertise<sup>2</sup> and disregard for the factualness of news<sup>3</sup>). In this paper, we examine how the uniquely interconnected state of the modern world might influence these societal dynamics.

We do so by combining two traditionally distinct theoretical perspectives<sup>4–6</sup>: cognitive psychology and evolutionary dynamics. With respect to cognition, we focus on a critical dimension of cognitive function that spans from automatic to controlled processing<sup>7–10</sup>. Along this dimension, automatic cognitive processes are more hardwired or over-learned habits, which leads to greater efficiency (e.g., greater speed and less effort) when they are appropriate, but at the cost of reduced flexibility (e.g., the ability to quickly adjust to the details of the current situation). More controlled processes, conversely, involve more deliberation and thought—requiring greater investment of time and effort, but more quickly responding to changes of circumstance and/or the specifics of a particular decision. Despite the immense body of work on automatic versus controlled processing, this distinction has been almost entirely absent from evolutionary game theory until recently<sup>11</sup>. Because humans do not exist in isolation, it is essential to consider population-level dynamics (e.g., the interaction between different agents) and the impact that cognition has on the environment. Evolutionary game theory provides the ideal tool for formally modeling such effects. For example, in the context of cooperation<sup>12–15</sup>, models have demonstrated the evolution of intuitive cooperation and deliberative self-interest. And in the context of cognition-environment feedback effects<sup>4–6</sup>, models have demonstrated the emergence of persistent cycles of automatic versus controlled processing.

Here, we integrate this combination of cognitive psychological<sup>16–19</sup> and evolutionary game theoretic approaches with social network models<sup>20–32</sup> to explore how the evolution of cognitive control is influenced by local versus global interaction structures. We simulate the evolutionary dynamics of a population of agents embedded in a network who vary along the dimension of automatic versus controlled processing (as defined below). Specifically, we add population structure to a multi-process agent model<sup>4</sup> in which agent  $i$  engages in automatic processing with probability  $1-x_i$  and controlled processing with probability  $x_i$ . As described in detail below, agents earn fitnesses based on their strategy, the strategy of other agents in the population, and the state of the environment; and then an evolutionary process occurs in which agents with lower fitness tend to adopt the strategies of agents with higher fitness (or die and be replaced with offspring of higher fitness agents).

Specifically, the fitness of agent  $i$  in timestep  $t$  is given by:

$$f_{i,t} = x_i \pi_t^C + (1 - x_i) \pi_t^A \quad (1)$$

where  $\pi_t^C$  is the payoff of engaging in controlled processing and  $\pi_t^A$  is the payoff of engaging in automatic processing. These terms are then defined as:

$$\pi_t^C = 1 - c + w(1 - \langle x \rangle_t) \quad (2)$$

$$\pi_t^A = 1 - p_t \quad (3)$$

$$p_t = 1 - \langle x \rangle_{t-1} \quad (4)$$

where:

$c$  is the fixed cost of controlled processing.

$w$  is the impact of automatic processing on the cost of controlled processing.

$\langle x \rangle_t$  is the average probability of controlled processing in the population at time  $t$ .

$p_t$  is the state of the environment at time  $t$ .

We now explain the logic behind these definitions. Our modeling of the payoff consequences of automatic versus controlled processing ( $\pi_t^C$  and  $\pi_t^A$ ) are based on the trade-off between efficiency versus flexibility of cognitive processing. One of the central tenets of cognitive psychology and neuroscience is that the defining feature of control-dependent processing (relative to automatic) is slower but more flexible responding; that is, the ability to respond to a stimulus more slowly but in a non-habitual, contextually-relevant way. Automatic processing is therefore conceptualized as supporting efficient and typically effective behavior, achieved by encoding pre-compiled responses that are optimally adapted to a particular set of circumstances, but are slow to develop or adapt in response to environmental change or experience. In contrast, controlled processing is conceptualized as supporting a more flexible form of processing that can adjust more quickly to changes in contingencies and thereby generate advantageous responses under a wider range of conditions. Critically, however, this flexibility comes at a cost<sup>33–38</sup>.

The distinction between automatic and controlled processing that we are capturing in our model can be seen in a wide range of everyday behaviors. While automatic processing is critical to the efficient and often skillful functioning of an agent (e.g., most of our physiological regulatory processes are automatic, and the development of automaticity is fundamental to the acquisition of complex skills, such as typing, playing the piano, driving a car, etc.), nevertheless, automatic behavior can be disadvantageous in many circumstances, especially when its immediate benefits are outweighed by longer term negative consequences. For example, research shows that automatic processing favors immediate gratification whereas the exercise of control promotes delayed gratification; and that controlled processing facilitates planning for the future. Accordingly, reliance on automatic processing can be associated with making poor investment decisions<sup>39</sup>, being overweight<sup>40</sup>, using addictive substances like alcohol, tobacco, and illegal drugs<sup>41</sup>, and preferring simple over nuanced information sources<sup>42</sup>, among many other things. Controlled processing is also central to the reasoning and analytic thinking that enable humans to solve complex problems, to be creative, and to achieve scientific and technological advances<sup>43</sup>.

Across all of these examples of day-to-day decision-making, the common thread is that controlled processing can often lead to better decision-making outcomes, but is more costly (e.g., requires more time and mental effort). To capture the effect of this relationship on average in our model, the basic decision-making payoff of automatic processing ( $1-p_t$  with  $0 < p_t < 1$ ) is less than the basic decision-making payoff of control (1). Thus  $p_t$  describes the relative advantage of controlled decision-making in time period  $t$ . In addition to these payoffs from decision-making,  $\pi_t^C$  expression also includes the fixed cost  $c$  paid to engage in controlling processing.

Our basic payoff specification also includes a social element, captured by the  $w(1 - \langle x \rangle_t)$  term of  $\pi_t^C$ . This describes the payoff that controlled agents receive from engaging with automatic agents. The parameter  $w$  can either be negative (such that interacting with agents acting automatically reduces the payoff of controlled processing) or positive (such that interacting with agents acting automatically increases the payoff of controlled processing). Negative  $w$  reflects situations in which interacting with automatic agents reduces the payoff of controlled agents.

For example, because automatic processing is faster than controlled processing, in circumstances for which timing is a factor, an agent relying on automaticity may outcompete one relying on controlled-processing to acquire resources by acting immediately to acquire the resource while the controlled agent is more carefully considering its best course of action<sup>5</sup>. Alternatively, agents acting automatically can derail the plans of controlled decision-makers due to automaticity's own lack of planning (e.g., irresponsible use of the externalities generated by controlled processing, such as abuse of antibiotics or reckless use of weaponry). Conversely, positive  $w$  reflects situations in which controlled agents benefit from interaction with automatic agents. For example, controlled decision-makers may profit by producing products that agents acting automatically preferentially desire (e.g., alcohol, drugs, and low quality news), or by exploiting the short-sightedness of agents acting automatically (e.g., manipulating automatic agents into false beliefs that benefit the controlled agent).

Finally, our modeling of the dynamics of  $p$  are designed to capture how the relative proportion of the two cognitive processing employed by agents in the population interacts—typically at much longer timescales than the day-to-day decision-making described above—with the environment. Specifically, we assume that the reasoning and analytic thinking enabled by controlled processing produce innovations and corresponding externalities that lead to a more hospitable environment, such as technologies that make the environment less variable and/or higher yield. This in turn makes the future more stable, reduces the importance of careful planning, and thus reduces the payoff advantage of controlled processing  $p$ . That is, the externalities generated by controlled processing help automatics as much as controls (without the automatics having to be the cost of control) because the externalities make it less necessary to exercise control in order to achieve good outcomes (i.e., controls and automatics both arrive at the same, better, outcome, but automatics do it without having to pay cost). For example, the invention of sophisticated irrigation and farming techniques makes food production much less susceptible to environmental variation, and in turn reduces the need to plan for the future by stockpiling against periods of drought. Or, in terms of the social environment, social science discoveries regarding the power of defaults have enabled policy makers to design institutions in which the optimal choice is the default (e.g., making saving for retirement opt-out rather than opt-in)—such that automatic agents are more likely to make optimal choices without having to exert control. Importantly, controlled processing is often necessary for maintaining these advances as well as developing them—without sufficient foresight and planning, and the willingness to invest in upkeep today in order to reap the benefit of functional systems tomorrow, the benefits of the technologies developed using controlled processing can be undermined.

We implement this cognition-environment feedback by having the relative advantage of controlled decision-making  $p$  decrease as the average level of controlled processing in the population  $\langle x \rangle$  increases. Specifically,  $p$  in time period  $t$  is set equal to  $1 - \langle x \rangle$  in period  $t-1$ , yielding the  $p_t$  expression given above. In a more complex version of the model, we also consider the impact of time lags on the interaction between the population and the environment—see “Methods” section for details.

We then study the population dynamics of automatic versus controlled processing using evolutionary game theory. Specifically, agents update their strategies using the death-birth Moran process with exponential fitness. In each generation, an agent (the learner) is randomly selected to update, and her probability of controlled processing is replaced with another agent (the teacher) selected proportional to an exponential function of the neighbors'

fitness; or, with probability  $u$ , a mutation occurs and the learner is assigned a new random strategy, sampled from a uniform distribution over the interval  $[0, 1]$ .

Mathematically, this population dynamic could represent either genetic evolution or cultural evolution (i.e., social learning). In our interpretation of the results, we mostly focus on cultural evolution, which can occur over fairly rapid timescales. Individuals regularly update their strategies via social learning—that is, they compare their payoff to the payoffs of others, and (probabilistically) adopt the strategies of higher payoff others. To be consistent with standard convention in evolutionary game theory, we will refer to each strategy updating round as a “generation”, but remind the reader to not take this term too literally. That is, depending upon the processes and phenomena of interest, and how the model is applied, a “generation” could refer to updating at the cultural level as a consequence of social learning, or at the evolutionary time scale of reproduction.

Importantly, there are three different timescales within the model. The fastest timescale is the timescale of decision-making, which occurs on the timescale of minutes to hours. Each agent  $i$  makes many decisions every day—for example, decisions about whether to consume a resource or save it for the future, to eat an unhealthy meal or resist the temptation, to take a piece of information at face value or more carefully consider its credibility. The fraction of automatic versus controlled decisions is determined by the strategy parameter  $x_i$ . Agents also interact with each other every day, with controlled agents incurring costs or gaining benefits from interacting with automatic agents, depending on the sign of  $w$ . Together, the payoffs from decisions and interactions on this decision making timescale determine the agent  $i$ 's fitness  $f_{i,t}$ . Next is the timescale of strategy updating (i.e., learning), which is somewhat slower—e.g., on the order of months or years. Occasionally, agents compare their own success (i.e., payoff) with the success of other agents. If they find another agent that is regularly doing better than them, they (probabilistically) adopt that agent's strategy. Finally, there is the timescale of cognition-environment feedback, whereby the level of control versus automaticity in the population affects the stability of the environment. With a sufficient investment of control-based analytic thinking, agents can invent and deploy new technologies that make life easier; and with sufficient lack of investment, these technologies can break down. These types of events tend to happen on a timescale that is longer than that of strategy updating. (Although this is less our focus, the model could also describe genetic evolution, in which the timescale of strategy updating—i.e., reproduction—would be longer, and more similar to the timescale of cognition-environment feedback.)

The main theoretical innovation of the present work is to add population structure to this model, in order to examine the impact of local versus global interactions on the evolutionary outcomes. For simplicity, our main analyses capture population structure by embedding agents in a ring-structured population in which they are connected to one neighbor on each side (results are not unique to the ring structure—for analysis of Small-world networks<sup>44</sup> with varying strength of community structure, see Supplementary Information). We then examine the impact of local connectedness (restricted to neighbors) versus global connectedness (with all agents) on three different dimensions in our model: learning, contact between agents, and the environment.

With respect to learning (or reproduction), we vary how teachers are selected. In local learning the teacher is randomly selected from the immediate neighbors of the learner in the network. In global learning, the teacher is randomly selected from the whole population.

With respect to contact between agents, we vary how the fraction of controlled agents is calculated when determining the

payoff consequences of contact between of agents with varying levels of automatic versus controlled processing (i.e.,  $\langle x \rangle_t$  in the  $w(1 - \langle x \rangle_t)$  term in the payoff of controlled processing). In global contact, the fraction of controlled agents in the whole population is used to determine  $\langle x \rangle_t$ , as in the basic model formulation presented above. In local contact, conversely, the fraction of controlled neighbors is used—such that the fitness of controlled processing is impacted only by neighboring agents in the network. As a result, each agent  $i$  has its own fraction of controlled neighbors, such that  $\langle x \rangle_{i,t}$  replaces  $\langle x \rangle_t$  in the fitness function.

With respect to environment, we vary how the fraction of controlled agents is calculated when updating the environmental variable  $p_t$  that stipulates the decision-making advantage of controlled processing. In global environment, all agents experience the same value of  $p_t$ , as in the main model formulation presented above. The value of  $p_t$  is then updated based on the average level of controlled processing across the entire population. In local environment, conversely, each agent  $i$  experiences its own local environment  $p_{i,t}$  which replaces  $p_t$  in the fitness function; and  $p_{i,t}$  is updated based on the average level of controlled processing of agent  $i$  and agent  $i$ 's immediate neighbors.

Together, these three different implementations of local versus global lead to eight ( $2^3$ ) possible combinations of local versus global connectedness, seven of which involve some element of globalization and one of which is entirely local. By examining model outcomes across these eight scenarios, we can explore the impact of differing forms of globalization on the temporal dynamics of how cognitive control evolves in the population.

## Results

**Globalization and average levels of control.** We begin by examining the impact of different forms of globalization (i.e., of learning, contact, and/or environment) on the equilibrium average level of cognitive control within the population by comparing each of the seven globalization scenarios to the entirely local scenario.

Before directly examining the results of agent-based simulations, it is helpful to gain an intuition for the kinds of effects globalization has in the model, and why. To do so, we conduct an analysis using approximations of the fitness of automatic and controlled processing for a simplified case in which an agent is either fully controlled or fully automatic (see Supplementary Note 1). The results provide intuitions regarding the impact of each form of globalization. Globalization of the environment causes controlled agents to improve the environment experienced (and the payoff earned) by automatic agents throughout the population, and thus always results in lower levels of cognitive control in equilibrium. Globalization of contact increases the exposure of controlled agents to automatic agents, the consequences of which depend dramatically on  $w$ : this exposure decreases the payoff of control when automatic agents hurt controlled agents,  $w < 0$ , but increases the payoff of control when automatic agents benefit controlled agents,  $w > 0$ . Globalization of learning homogenizes the population and undercuts the clustering of strategies within the network, thereby weakening the effects of local contact and/or environment. Finally, all three forms of globalization interact, such that control is substantially higher in the fully local scenario compared to any of the globalization scenarios (provided that  $w < 0$ ).

We now present simulations of the full agent-based model. Agents were placed on a ring structure, and  $c$  and  $w$  were systematically varied across simulations. For each set of parameter values, we then simulated the model for  $8 \times 10^4$  generations and plotted the time-averaged frequency of control  $\langle x \rangle$  averaged over 10 replicates for each set of simulation

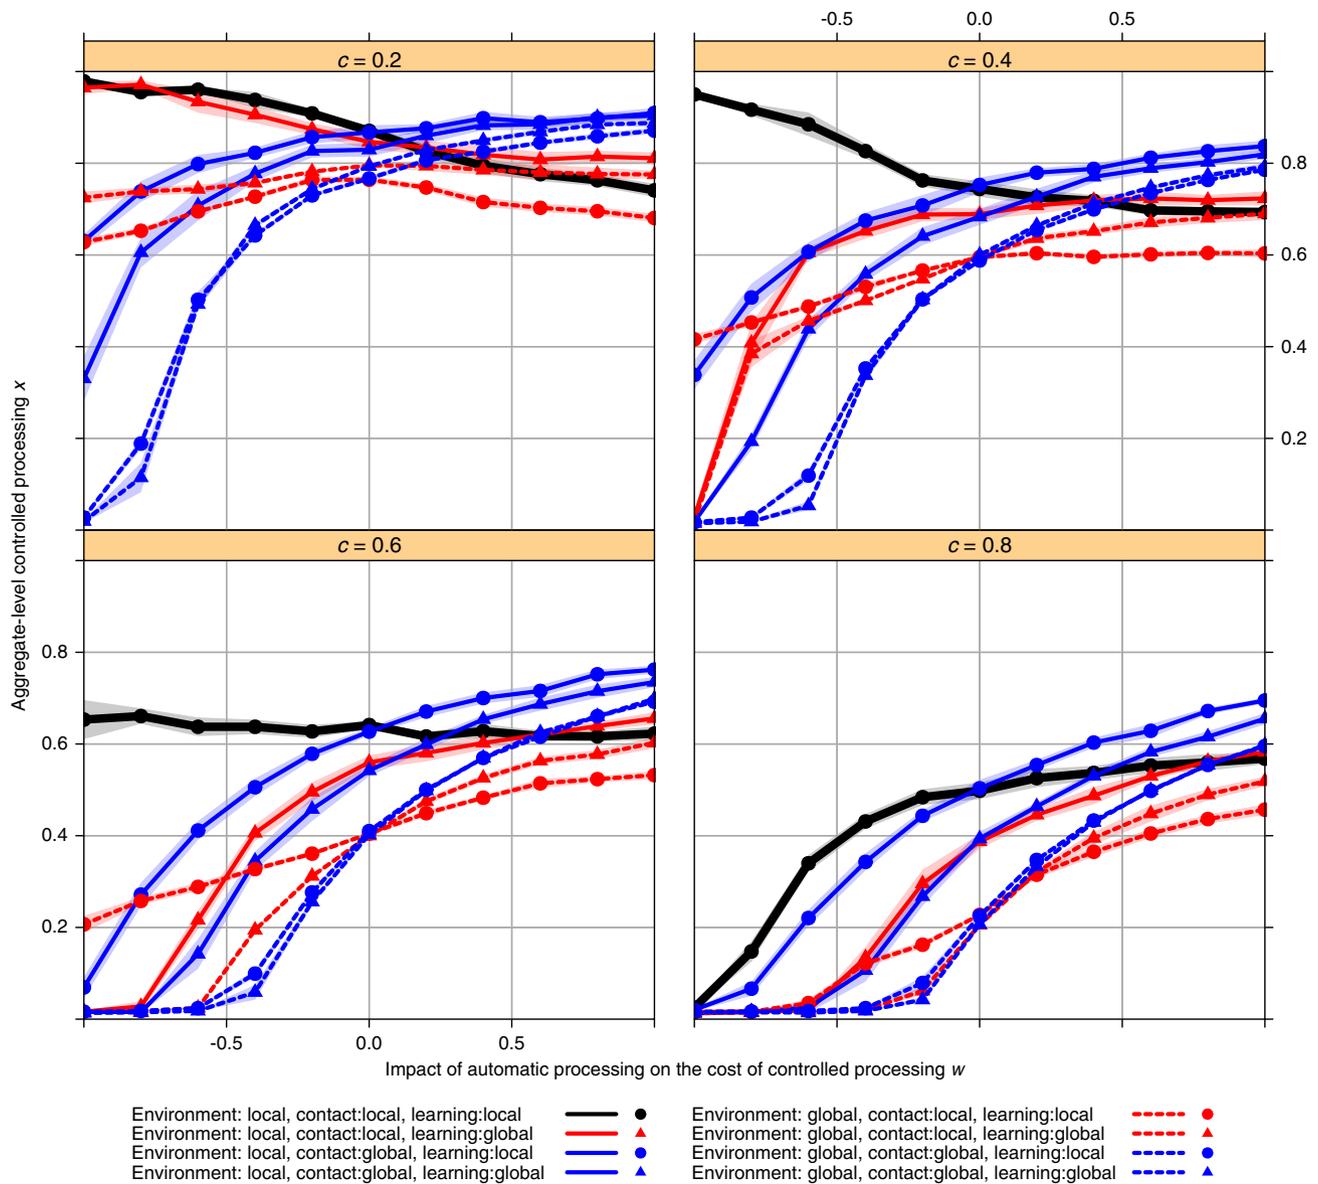
parameters (as shown in Supplementary Figs. 1 and 2, there is no significant variation across simulation replicates in the steady state value of  $\langle x \rangle$ ). In each simulation, we initialized the population from  $x_i = 0.01$  (almost entirely automatic processing). Figure 1 shows that the simulation results generally accord with the predictions arising from the analytical approximation of the simplified system which are described above (and shown in Supplementary Table 1). When  $w$  is negative—such that the presence of automatic processing reduces the payoff of engaging in controlled processing—globalization of virtually any kind substantially reduces the average level of controlled processing in the population.

Conversely, when  $w$  is positive—such that the presence of automatic processing benefits those engaging in controlled processing—certain forms of globalization (specifically, global contact) can somewhat increase the average level of controlled processing in the population (although still maintaining a substantial level of automatic processing). Our results are robust to other network structures (see Supplementary Fig. 3) and larger networks (see Supplementary Fig. 4).

**Globalization and the dynamics of control.** The foregoing section addressed long-run average levels of control in the population. Here, we examine the impact of the scale of interaction on the dynamics of automaticity versus control, both across space and time (Fig. 2). We observe two striking effects.

The first concerns the emergence of mesoscale communities. Qualitatively, when learning is global (second row of Fig. 2), the entire population is well-mixed, and neither spatial nor temporal patterns are apparent. Conversely, as expected based on prior work on population structure<sup>22</sup>, when learning is local (top row of Fig. 2), communities of low versus high cognitive control emerge within the network (and the boundaries of these communities migrate over time). Interestingly, although global interactions tend to homogenize, we find a different pattern when it comes to the scale of environment: the emergence of mesoscale communities is markedly stronger when the environment is global rather than local. To quantify these differences in the extent of mesoscale communities, we calculate the average level of assortment  $a$  within the population (defined as the difference in (i) the fraction of an agent's neighbors that have the same strategy as the agent versus (ii) the fraction of agents in the whole population—excluding the agent itself—that have the same strategy as the agent). Doing so shows that there is no assortment,  $a = 0$ , in any of the global learning scenarios—no mesoscale communities emerge. Conversely, there is substantial assortment when learning is local and environment is local ( $a = 0.65$  for local contact,  $a = 0.65$  for global contact); but even higher levels of assortment when learning is local but environment is global ( $a = 0.70$  for local contact,  $a = 0.74$  for global contact).

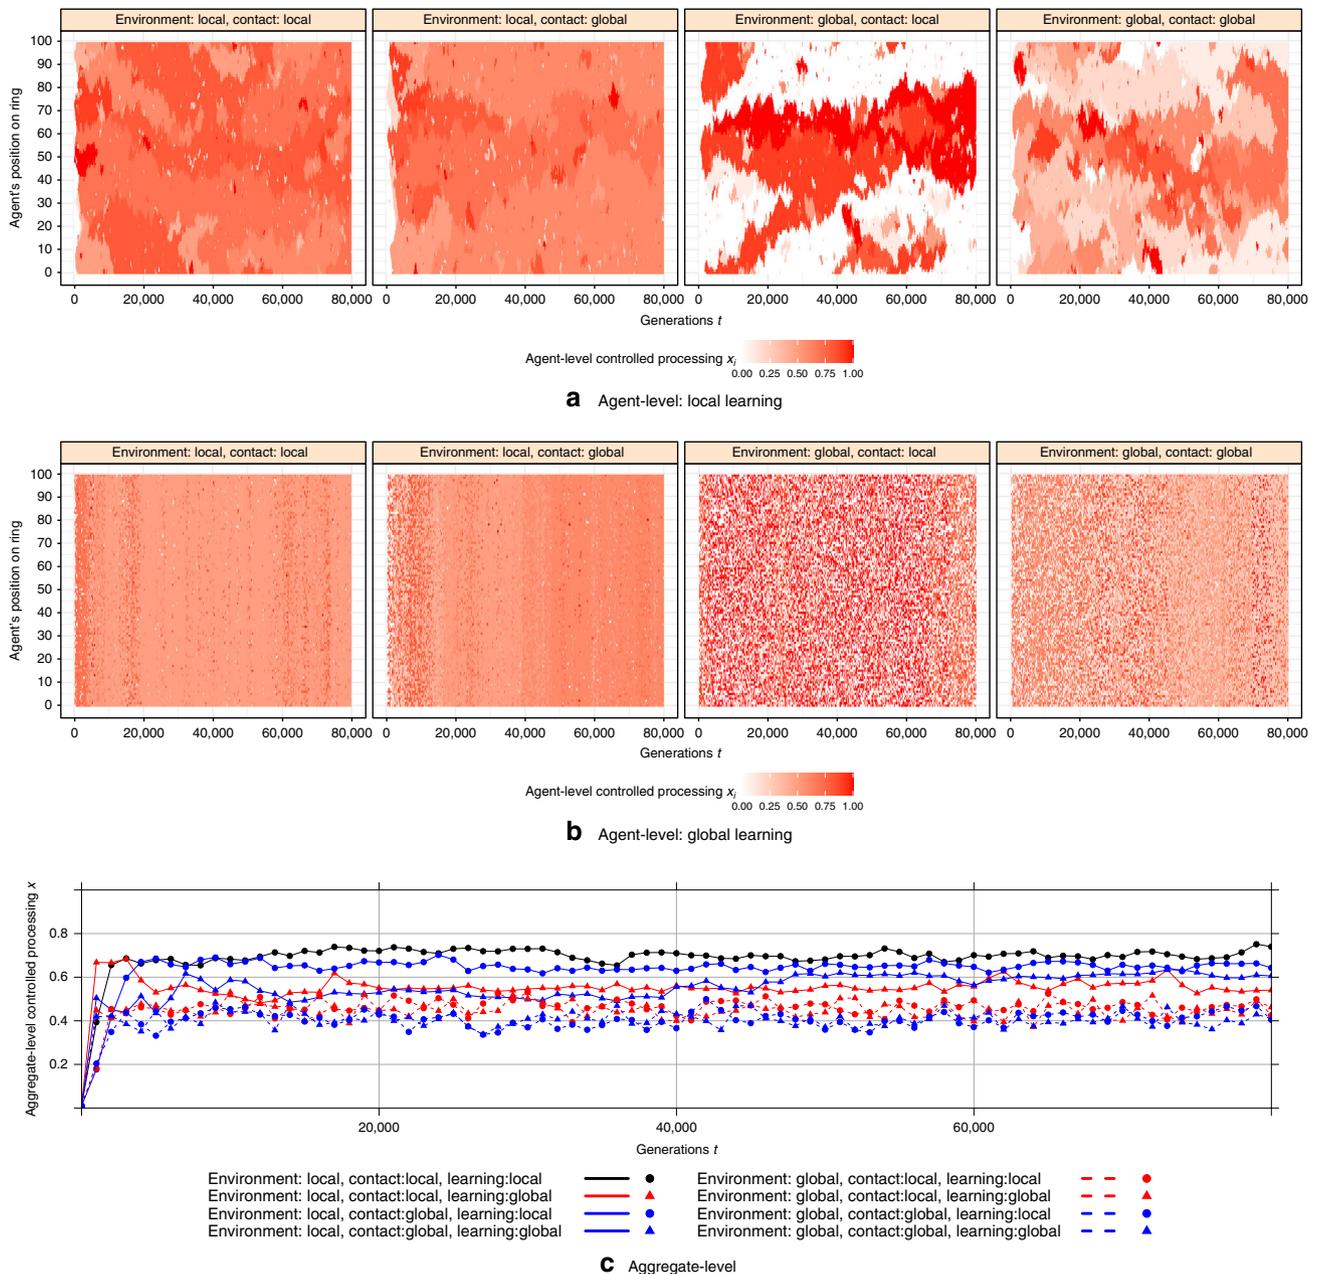
Why does local environment homogenize whereas global environment facilitates clusters? The answer arises from the particular anti-coordination payoff structure of automatic versus controlled agents in our model. When environment is local, the advantage of controlled processing  $p_{i,t}$  that each agent experiences is heavily influenced by that agent's neighbors. If agent  $i$ 's neighbors are controlled, then  $p_{i,t}$  will be low and it will be advantageous for agent  $i$  to use automatic processing; if agent  $i$ 's neighbors are automatic, then the opposite will be true. Thus, it is more difficult to maintain clusters of agents with the same strategy—the local environment creates incentives to diverge from the strategies of one's neighbors. Conversely, when the environment is global, all agents face the same  $p_t$ , and thus the environment does not create an incentive to diverge from one's neighbors.



**Fig. 1** Aggregate-level controlled processing  $x$ . Each line represents a combination of local versus global contact, environment, and learning, across different levels of the impact of automatic processing on the cost of controlled processing ( $w$ ) and fixed cost of controlled ( $c$ ). Contact is indicated by color (red = local, blue = global). Environment is indicated by line type (solid = local, dashed = global). Learning is indicated by symbol (circle = local, triangle = global). All values are averaged over 10 simulation replicates. The shaded area represents the 95% confidence interval of  $\langle x \rangle$  across simulation replicates for each set of parameters and combination of local versus global interaction.

The second striking feature in Fig. 2 involves the variance in  $x$  values across the population at any given time. Qualitatively, when the environment is local, there is relatively little variation in the level of control within the population at any given time—most agents show an intermediate level of control (pink coloration in Fig. 2). Conversely, when the environment is global, there is wide variation in the level of control within the population at any given time – some agents are almost entirely automatic (white) whereas others are almost entirely controlled (red). This is particularly true when contact is local. Quantitatively, the average value across time of the variance in  $x$  values in the population is very low when environment is local (0.000–0.008 depending on scale of contact and learning), higher when environment is global and contact is global (0.052 for local learning, 0.080 for global learning), and highest when environment is global and contact is local (0.171 for local learning, 0.181 for global learning).

The intuition for this variance-increasing effect of global environment is as follows. The basic structure of our model is such that the equilibrium involves either (i) a homogenous population in which all agents play a mixed strategy that randomizes over automatic and controlled processing,  $x = x^*$  (leading to low variance across strategies) or (ii) a heterogeneous population in which fraction  $(1 - x^*)$  of agents are totally automatic,  $x = 0$ , and fraction  $x^*$  are totally controlled,  $x = 1$  (leading to high variance across strategies). When environment is local, agents in different parts of the network face different values of  $p$  (and thus different selective forces on  $x$ ) making it difficult to achieve the population-level coordination necessary for the heterogeneous solution. Conversely, when environment is global, all agents face the same value of  $p$ , and thus are better able to arrive at the heterogeneous solution. (For mesoscale measures on other network structures see Supplementary Fig. 5).

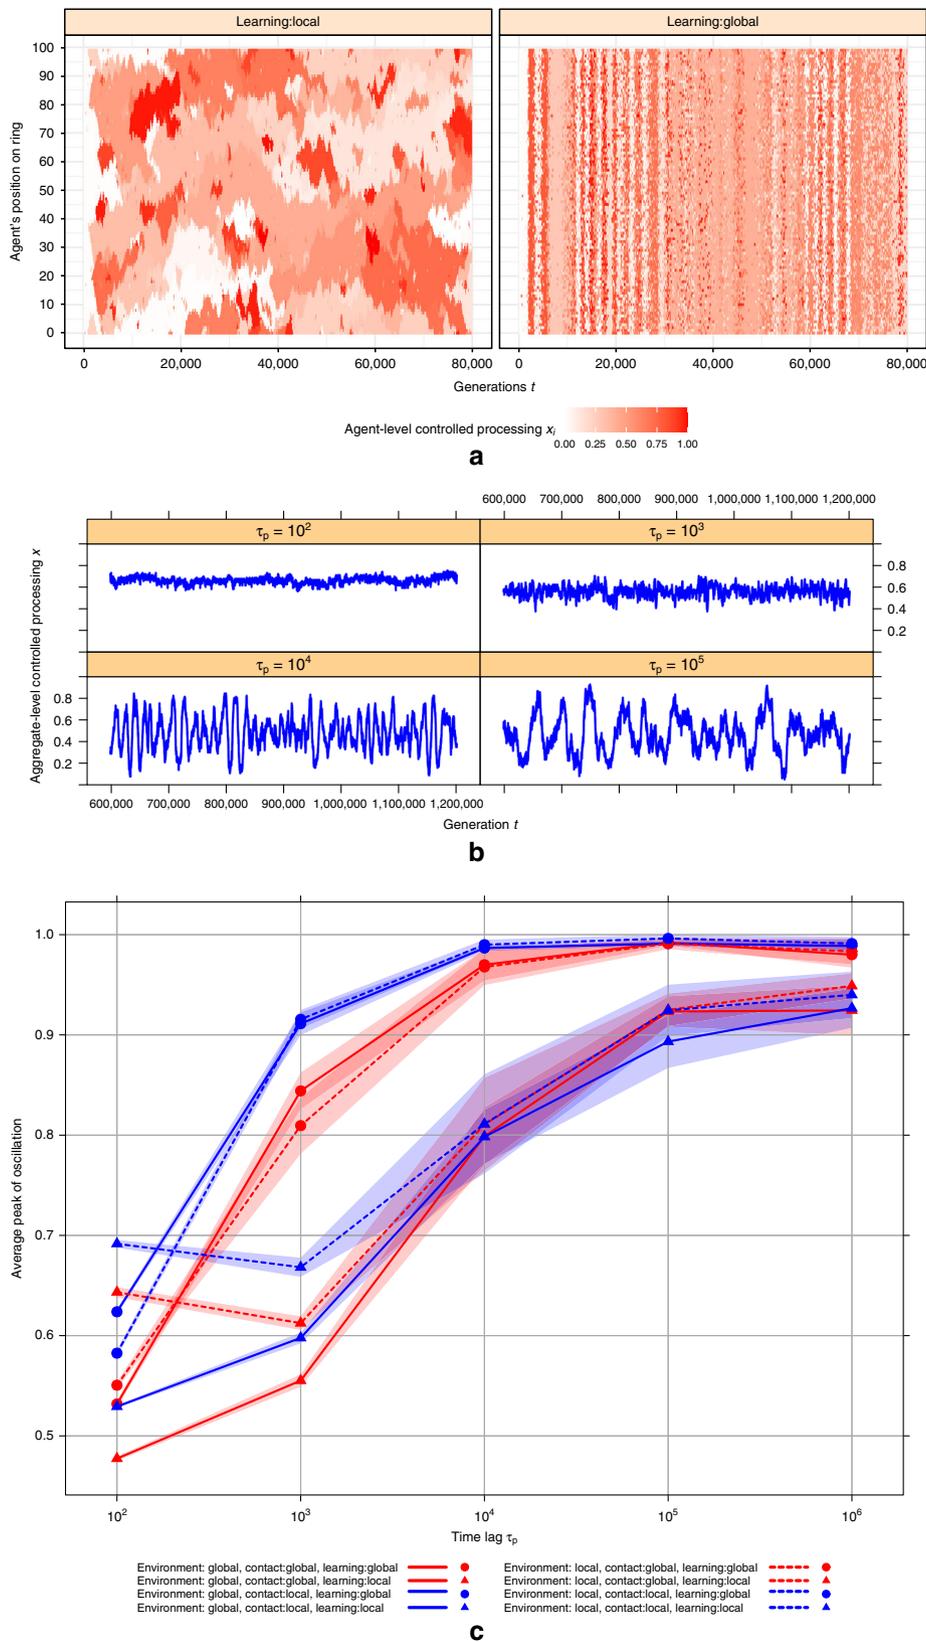


**Fig. 2** Spatiotemporal dynamics of automatic versus controlled processing. **a, b** Represent agent-level controlled processing  $x$  for each condition of localness over  $8 \times 10^4$  generations. Shown is the strategy of the agent in each location on the ring structure (on the y axis, such that pixels above and below each other are neighbors) at each generation (on the x axis), for each globalization scenario. **a** Shows results for global learning and panel (b) shows results for local learning. **c** Shows the aggregate-level controlled processing  $x$  over time for each condition of localness. Results are created for impact of automatic processing on the cost of controlled processing  $w = -0.15$  and fixed cost of controlled  $c = 0.5$ .

Finally, we used the agent-based model to examine the more complex situation in which there is lag in the impact of cognition on the environment, parameterized by  $\tau_p$  (see Methods for details of implementation). First, we examined the spatiotemporal distribution of strategies in the presence of lag for the situation in which all contact, environment, and learning are local versus global (Fig. 3a). Unlike the no-lag case considered in Fig. 2, strong temporal fluctuations are observed even in the fully globalized scenario. These fluctuations, although present, are less stark in the local scenario. When global, the population fluctuates between almost entirely automatic and entirely controlled (as has been shown previously in well-mixed populations<sup>4</sup>). When local, there is more coexistence of low and high control agents (due to the existence of local communities which are, to some extent,

fluctuating out of phase with each other). Nonetheless, dramatic oscillations in which the population is almost entirely controlled or automatic, can still occur in the local scenario, as shown by Fig. 3b which shows the average level of control over time in the entire local scenario. When the lag is sufficiently large, the population fluctuates between almost entirely controlled ( $\langle x \rangle$  near 1) and almost entirely automatic ( $\langle x \rangle$  near 0), as in the global case considered previously.

To quantify how globalization effects the relationship between lag and fluctuations, we examined the magnitude of the oscillations in controlled processing that emerged for local vs. global interactions at different levels of the lag parameter  $\tau_p$ . Specifically, we used the average value of the peak of  $x$  to quantify the size of the fluctuations. Figure 3c (see Supplementary Fig. 6



for results on other network structures) shows that, across different globalization conditions, increasing the lag  $\tau_p$  leads to higher amplitude cycles (as in the well-mixed population<sup>4</sup>). However, when learning is local, a larger lag is required to produce cycles of equivalent magnitude. Thus, although local

learning promotes spatial oscillations (as shown in Fig. 2), it suppresses aggregate-level temporal oscillations.

Why is this so? There are two contributing factors. First, the mesoscale communities created by local learning lead to decoupling, such that even if each section of the network is

**Fig. 3 Effects of time lag  $\tau_p$  on magnitude of oscillations.** **a** Represents agent-level controlled processing  $x$  for global vs local learning (here time lag  $\tau_p = 10^3$  and contact and environment are global, however, the spatial pattern does not qualitatively change for local contact and/or local environment). **b** Shows time series of oscillations of aggregate-level controlled processing  $x$  when contact, environment, and learning are local across different values of time lag. **c** Shows the average peak of oscillations for different condition of localness across different time lag. To cover all values of time delay, we run simulations over  $1.2 \times 10^6$  generations. Results are averaged over 10 simulation replicates. The shaded area represents the 95% confidence interval across simulation replicates for each set of parameter values and combination of local versus global interaction. Across all panels, the impact of automatic processing on the cost of controlled processing  $w = -0.15$  and fixed cost of controlled  $c = 0.5$ .

cycling with the same magnitude of the cycles observed under global learning, they are out of phase with each other. Thus, in aggregate, the cycles in the average level of control across the whole population are of smaller amplitude (see Supplementary Fig. 7). Second, when learning is local, each agent can only imitate a small fraction of the population; thus, it takes longer for the whole population to shift to the currently-optimal strategy. Hence, in order for a full population swing to occur (which causes the environment to move back in the other direction), the environment needs to stay put longer; that is, with local learning, it takes longer for strategic innovation to spread through the population, which mitigates the amplitude of the oscillations.

## Discussion

The simulation results reported above suggest that globalization can have a dramatic—and complex—impact on the evolutionary dynamics of cognitive control. They also reveal important distinctions between different forms of globalization. Although each form involves greater interaction with dissimilar others, the details of that interaction matter.

Consistent with most prior work on networks, we find that globalization of learning (i.e., being able to imitate successful others outside your local neighborhood) fundamentally alters the spatiotemporal dynamics of control, breaking up the mesoscale neighborhoods of individuals with similar cognitive strategies that exist when learning is local. Interestingly, a consequence of this homogenization is that global learning induces population-wide fluctuations in the level of control (because all agents are phase-locked).

The impact of the other forms of globalization, however, are specific to the particular system of automatic versus controlled processing we study here, and not obvious based on prior work. For example, the consequences of globalization of contact between agents depend on the nature of the contact. If automatic processing imposes costs on controlled processing (e.g., faster automatic agents outcompeting more thoughtful controlled agents when directly competing over resources, or automatic agents' lack of planning undermining the plans of controlled agents), then globalization of contact harms control due to increased exposure to agents engaged in automatic processing. But if control benefits from interaction with automaticity (e.g., controlled agents selling products that are particularly helpful—and thus attractive—to automatic agents, or controlled agents exploiting the short-sightedness of automatic agents), then the opposite is true and globalization of contact promotes the evolution of control.

Finally, globalization of environment—that is, an increase in the extent to which one's actions influence the environment experienced by everyone (such as the ozone layer)—uniformly reduces the level of control, by allowing automatic processing to enjoy the environmental benefits created by controlled processing. And surprisingly, it is global environment that facilitates the emergence of mesoscale communities and maintenance of variance in strategies within the population, whereas local environment leads to homogenization. This is the opposite of what one might expect from prior work on, for example, the evolution of

cooperation<sup>28</sup>, in which (local) spatial structure maintains heterogeneity whereas (global) well-mixed populations are homogeneous. Together, these observations highlight the importance of the details when considering the impact of globalization—which form(s) of interaction are affected by globalization (learning, contact, and/or environment), and what underlying game structure is driving the dynamics (automatic/controlled processing, versus, for example, cooperation<sup>45</sup>).

In addition to these insights into the impact of globalization, our results also have implications for previous work on cognition-environment feedback<sup>4–6</sup>. Specifically, they lend support to the generality of prior conclusions about the emergence of cycles in automaticity versus control, which had previously been shown to occur when cognition-environment feedback is sufficiently lagged. First, they provide additional evidence for the robustness of such cycles, extending this observation to a broader and more complex set of conditions. More specifically, they show that under conditions of lagged feedback, cycles can occur in both local and global settings. Moreover, we find that when learning is local, spatio-temporal fluctuations can occur even in the absence of such lags. Thus, the current work strengthens formal theoretical support for the idea that societies characterized by high levels of cognitive control (and associated social and technological complexity) may be at risk of collapse due to cognition-environment feedback.

It is of course important to recognize that our model is, in the game theoretic tradition, an abstract and highly simplified conceptual model. Rather than quantitative predictions, our model makes high-level predictions about the kinds of outcomes that may arise, and the circumstances under which different outcomes might be expected to occur. These conceptual predictions can be explored in future work, using both laboratory experimental systems and examination of the observational data.

Finally, the insights generated by our model have the potential to help inform policy. Our results make several relevant points. First, they suggest that it may be of particular importance to promote interactions whereby control benefits from contact with automaticity, rather than suffers; such beneficial interactions make globalization of contact work for cognitive control instead of against it. For example, in a market context, facilitating the purchase by automatic agents of innovations created by controlled processing, the sale of which benefits the controlled agents (e.g., through marketing that appeals to automatic processing). Second, our results reaffirm the challenges posed by global-level environmental issues such as climate change, which create a free-rider problem that undermines cognitive control. And third, our results demonstrate the potential for a truly global collapse (and risk of extinction) of control when learning is global: as the whole world becomes bound together, there are fewer reservoirs of control to buffer against local fluctuations, and more cataclysmically, to repopulate after a large crash.

Globalization is an extremely complex issue, and it is essential to explore its causes and consequences using a wide range of perspectives and disciplines. We hope that the work presented here will add a new lens with which to consider the issue, provide new insights, and highlight important new directions for future work.

## Methods

**Agent-based model.** Our simulation results are produced using agent-based simulations on a network with population size  $N = 100$  (results are robust to larger population size; see Supplementary Fig. 4). Each agent  $i$  is defined by a single parameter  $x_i$  that represents its probability of controlled processing. The fitness of each agent is calculated according to the following function:

$$f_{i,t} = x_i(1 - c + w(1 - \langle x \rangle_{i,t})) + (1 - x_i)(1 - p_{i,t}) \quad (5)$$

where  $c$  is the fixed cost of control,  $w$  is the impact of automatic processing on the cost of controlled processing,  $\langle x \rangle_{i,t}$  is the average probability of control among the agent's direct neighbors (including the agent itself) in generation  $t$ , and  $p_{i,t}$  is the relative advantage of controlled processing in generation  $t$ .

Agents update their strategies using the Moran death-birth process (our model is robust to the update rule, as additional simulations show that the pairwise comparison process<sup>46,47</sup> leads to qualitatively similar outcomes). In each generation, an agent  $L$  (the learner) is randomly selected to potentially update its strategy. Then the learner is replaced by a copy of agent  $T$  (the teacher) who is selected with probability  $W(T \rightarrow L) = \exp(s f_T)$  where  $s$  is the intensity of selection and  $f_T$  is the fitness of the teacher (reproducer). Alternatively, with probability  $u$  a mutation occurs; in that case, instead of adopting the other agent's strategy, the learner adopts a new strategy sampled from a uniform distribution on the interval  $[0, 1]$ .

Following ref. 4, we also implement cognition-environment feedback by updating the advantage of controlled processing in each generation with a time lag. When environment is local, we specify  $N$  agent-specific values  $p_{i,t}$  and stipulate that only the strategy of immediate neighbors of agent  $i$  (including agent  $i$  itself) influences the change in  $p_{i,t}$  over time:

$$p_{i,t} = p_{i,t-1} + \frac{(1 - \langle x \rangle_{i,t-1}) - p_{i,t-1}}{\tau_p} \quad (6)$$

Conversely, when environment is global, there is a single value of  $p$  for the whole population, which updates according to:

$$p_t = p_{t-1} + \frac{(1 - \langle x \rangle_{t-1}) - p_{t-1}}{\tau_p} \quad (7)$$

$\tau_p$  represents the time lag of cognition-environment feedback. When  $\tau_p = 1$ , the basic model described in the main text with no time lag recovers.

**Numerical simulations.** Our simulations use intensity of selection  $s = 10$  (strong selection) and mutation rate  $u = 0.01$  (as expected, increasing the mutation rate pushes average values of  $x$  towards 0.5; see Supplementary Figs. 8 and 9). We structure the population within a ring in which each agent is connected to  $k/2$  neighbors on each side (for a total of  $k$  neighbors). Simulations were initialized with  $x = 0.01$  for all agents. We investigate the effect of globalization across 8 different scenarios described by the 2 by 3 combination of [local, global]  $\times$  [learning, environment, and contact]. That is, for each dimension of globalization we either consider only neighboring agents (ring structure with  $k = 2$ ; local) or the whole population (full graph; global). In a complimentary analysis (Supplementary Figs. 3, 5, and 6) we explored Small-world networks varying strength of community structure (rewiring rate).

We ran the simulations for  $8 \times 10^4$  generations except for the effect of  $\tau_p$  on average peak of oscillations (Fig. 3) where we ran the simulation for  $1.2 \times 10^6$  generations. For  $\tau_p = 1$  where population reaches an equilibrium and stays there, we averaged over last  $4 \times 10^4$  generations for the stable strategy. Results reported in Fig. 1 and Fig. 3c are averaged over 10 simulation replicates. All simulations were done in parallel on MIT Sloan Engaging cluster.

**Reporting summary.** Further information on research design is available in the Nature Research Reporting Summary linked to this article.

## Data availability

All data required to run the simulations are available at: <https://osf.io/fy94w/> or can be requested from the authors. A reporting summary for this Article is available as a Supplementary Information file.

## Code availability

All scripts necessary to reproduce the results are available at: <https://osf.io/fy94w/> or can be requested from the authors.

Received: 23 April 2019; Accepted: 26 May 2020;

Published online: 18 June 2020

## References

- Hardin, R. Group provision of step goods. *Behav. Sci.* **21**, 101–106 (1976).
- Lewandowsky, S. & Oberauer, K. Motivated rejection of science. *Curr. Directions Psychological Sci.* **25**, 217–222 (2016).
- Pennycook, G. & Rand, D. G. Lazy, not biased: susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition* **188**, 39–50 (2019).
- Rand, D. G., Tomlin, D., Bear, A., Ludwig, E. A. & Cohen, J. D. Cyclical population dynamics of automatic versus controlled processing: an evolutionary pendulum. *Psychological Rev.* **124**, 626 (2017).
- Toupo, D. F. P., Strogatz, S. H., Cohen, J. D. & Rand, D. G. Evolutionary game dynamics of controlled and automatic decision-making. *Chaos* **25**, 073120 (2015).
- Tomlin, D. A., Rand, D. G., Ludwig, E. & Cohen, J. D. The evolution and devolution of cognitive control: The costs of deliberation in a competitive world. *Sci. Rep.* **5**, 11002 (2015).
- Allport, G. W. *The Nature of Prejudice*. (Addison Wesley, 1954).
- Cohen, J. D., Dunbar, K. & McClelland, J. L. On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychol. Rev.* **97**, 332 (1990).
- Shiffrin, R. M. & Schneider, W. Controlled and automatic information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychol. Rev.* **84**, 127–190 (1977).
- Treisman, A. & Kahneman, D. *Changing Views of Attention and Automaticity*. (Academic Press, Orlando, FL, 1984).
- Bear, A. & Rand, D. G. Evolution: the value of information. *Nat. Hum. Behav.* **1**, 0156 (2017).
- Bear, A. & Rand, D. G. Intuition, deliberation, and the evolution of cooperation. *Proc. Natl Acad. Sci. USA* **113**, 936–941 (2016).
- Bear, A., Kagan, A. & Rand, D. G. Co-evolution of cooperation and cognition: the impact of imperfect deliberation and context-sensitive intuition. *Proc. Biol. Sci.* **284**, 20162326 (2017).
- Jagau, S. & van Veelen, M. A general evolutionary framework for the role of intuition and deliberation in cooperation. *Nat. Hum. Behav.* **1**, 0152 (2017).
- Mosleh, M. & Rand, D. G. Population structure promotes the evolution of intuitive cooperation and inhibits deliberation. *Sci. Rep.* **8**, 6293 (2018).
- Slooman, S. A. The empirical case for two systems of reasoning. *Psychological Bull.* **119**, 3 (1996).
- Kahneman, D. A perspective on judgment and choice: mapping bounded rationality. *Am. Psychol.* **58**, 697 (2003).
- Evans, J. S. B. & Stanovich, K. E. Dual-process theories of higher cognition advancing the debate. *Perspect. Psychol. Sci.* **8**, 223–241 (2013).
- Stanovich, K. E. & West, R. F. Individual differences in rational thought. *J. Exp. Psychol.: Gen.* **127**, 161–188 (1998).
- Szolnoki, A. & Perc, M. Evolutionary dynamics of cooperation in neutral populations. *New J. Phys.* **20**, 013031 (2018).
- Szolnoki, A. & Perc, M. Resolving social dilemmas on evolving random networks. *EPL (Europhys. Lett.)* **86**, 30007 (2009).
- Nowak, M. A. & May, R. M. Evolutionary games and spatial chaos. *Nature* **359**, 826–829 (1992).
- Ellison, G. Learning Local interaction, and coordination. *Econometrica* **61**, 1047–1071. <https://doi.org/10.2307/2951493> (1993).
- Lindgren, K. & Nordahl, M. G. Evolutionary dynamics of spatial games. *Phys. D* **75**, 292–309 (1994).
- Killingback, T. & Doebeli, M. Spatial Evolutionary Game Theory: Hawks and Doves Revisited. *Proc. Biol. Sci.* **263**, 1135–1144 (1996).
- Nakamaru, M., Matsuda, H. & Iwasa, Y. The Evolution of Cooperation in a Lattice-Structured Population. *J. Theor. Biol.* **184**, 65–81 (1997).
- Szabo, G. & Fath, G. Evolutionary games on graphs. *Phys. Rep.* **446**, 97–216 (2007).
- Nowak, M. A., Tarnita, C. E. & Antal, T. Evolutionary dynamics in structured populations. *Philos. Trans. R. Soc. B: Biol. Sci.* **365**, 19–30 (2010).
- Hauert, C. & Imhof, L. A. Evolutionary games in deme structured, finite populations. *J. Theor. Biol.* **299**, 106–112 (2012).
- Perc, M. Phase transitions in models of human cooperation. *Phys. Lett. A* **380**, 2803–2808 (2016).
- Mosleh, M. & Heydari, B. Fair topologies: community structures and network hubs drive emergence of fairness norms. *Sci. Rep.* **7**, 2686 (2017).
- Perc, M. & Szolnoki, A. Coevolutionary games—A mini review. *Biosystems* **99**, 109–125 (2010).
- Cohen, J. D. The vulcanization of the human brain: a neural perspective on interactions between cognition and emotion. *J. Economic Perspect.* **19**, 3–24 (2005).
- Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704 (2005).
- Gershman, S. J., Horvitz, E. J. & Tenenbaum, J. B. Computational rationality: a converging paradigm for intelligence in brains, minds, and machines. *Science* **349**, 273–278 (2015).
- Keramati, M., Dezfouli, A. & Piray, P. Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Comput. Biol.* **7**, e1002055 (2011).

37. Shenhav, A., Botvinick, M. M. & Cohen, J. D. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* **79**, 217–240 (2013).
38. Griffiths, T. L., Lieder, F. & Goodman, N. D. Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* **7**, 217–229 (2015).
39. Frederick, S. Cognitive reflection and decision making. *J. Economic Perspect.* **19**, 25–42 (2005).
40. Jarmolowicz, D. P. et al. Robust relation between temporal discounting rates and body mass. *Appetite* **78**, 63–67 (2014).
41. Amlung, M. & MacKillop, J. Further evidence of close correspondence for alcohol demand decision making for hypothetical and incentivized rewards. *Behav. Processes.* **113**, 187–191 (2015).
42. Shenhav, A., Rand, D. G. & Greene, J. D. The relationship between intertemporal choice and following the path of least resistance across choices, preferences, and beliefs. *Judgm. Decis. Mak.* **12**, 1–18 (2017).
43. Pennycook, G., Fugelsang, J. A. & Koehler, D. J. Everyday consequences of analytic thinking. *Curr. Directions Psychological Sci.* **24**, 425–432 (2015).
44. Watts, D. J. & Strogatz, S. H. Collective dynamics of ‘small-world’ networks. *Nature* **393**, 440 (1998).
45. Tarnita, C. E., Antal, T., Ohtsuki, H. & Nowak, M. A. Evolutionary dynamics in set structured populations. *Proc. Natl Acad. Sci. USA* **106**, 8601–8604 (2009).
46. Traulsen, A., Pacheco, J. M. & Nowak, M. A. Pairwise comparison and selection temperature in evolutionary game dynamics. *J. Theor. Biol.* **246**, 522–529 (2007).
47. Szabó, G. & Tóke, C. Evolutionary prisoner’s dilemma game on a square lattice. *Phys. Rev. E* **58**, 69 (1998).

### Acknowledgements

The authors gratefully acknowledge funding from the Ethics and Governance of Artificial Intelligence Initiative of the Miami Foundation, the William and Flora Hewlett Foundation, the John Templeton Foundation, and the Social Sciences and Humanities Research Council of Canada.

### Author contributions

M.M., J.D.C., and D.G.R. designed the research, M.M. and K.K. ran the simulations, M.M. analyzed the results, M.M., J.D.C., and D.G.R. wrote the paper.

### Competing interests

The authors declare no competing interests.

### Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41467-020-16850-0>.

Correspondence and requests for materials should be addressed to M.M.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020

*Supplementary Information*  
**“Globalization and the Rise and Fall of Cognitive Control”**  
Mosleh et al.

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02138  
E-mail: mmosleh@mit.edu

## Contents

<b>1. Supplementary Note 1: Analytical approximation for <math>\tau_p = 1</math></b>	<b>2</b>
<b>2. Supplementary Figures</b>	<b>4</b>
2.1. Supplementary Figure 1. Variations of simulation outcomes for different replicates (local learning).....	4
2.2. Supplementary Figure 2. Variations of simulation outcomes for different replicates (global learning).....	5
2.3. Supplementary Figure 3. Aggregate-level controlled processing $x$ for no lag $\tau_p = 1$ on Small-world network.....	6
2.4. Supplementary Figure 4. Aggregate-level controlled processing $x$ for no lag $\tau_p = 1$ on a ring structure with $N=10^5$ nodes .....	7
2.5. Supplementary Figure 5. Meso-scale communities measures as a function of probability of link rewiring in Small-world network.....	8
2.6. Supplementary Figure 6. Effect of time-lag on amplitude of oscillation for Small-world networks varying probability of rewiring .....	9
2.7. Supplementary Figure 7. Local oscillation of neighboring agents with local learning ...	10
2.8. Supplementary Figure 8. Aggregate-level controlled processing $x$ for no lag $\tau_p = 1$ and mutation $u=0.2$ .....	11
2.9. Supplementary Figure 9. Aggregate-level controlled processing $x$ for no lag $\tau_p = 1$ and mutation $u=0$ .....	12

### Supplementary Note 1: Analytical approximation for $\tau_p = 1$

To give an intuition regarding the impact of globalization on the evolution of cognitive control in the absence of lag, we approximate the fitness of automatic and controlled processing for the simplified case in which an agent is either fully controlled or fully automatic (see Supplementary Table 1).

**Supplementary Table 1. Approximation of fitness of automatic and controlled processing under different conditions of local versus global interaction and learning.**

Environment	Contact	Learning	Fitness of controlled processing $f_c$	Fitness of automatic processing $f_a$
Local	Local	Local	$1 - c + \frac{2}{3}\beta w(1 - \langle x \rangle)(1 - a)$	$\frac{2}{3}\beta \langle x \rangle (1 - a)$
Local	Local	Global	$1 - c + \frac{2}{3}w\beta(1 - \langle x \rangle)$	$\frac{2}{3}\beta \langle x \rangle$
Local	Global	Local	$1 - c + w(1 - \langle x \rangle)$	$\frac{2}{3}\beta \langle x \rangle (1 - a)$
Local	Global	Global	$1 - c + w(1 - \langle x \rangle)$	$\frac{2}{3}\beta \langle x \rangle$
Global	Local	Local	$1 - c + \frac{2}{3}w(1 - \langle x \rangle)(1 - a)$	$\langle x \rangle$
Global	Local	Global	$1 - c + \frac{2}{3}w(1 - \langle x \rangle)$	$\langle x \rangle$
Global	Global	Local	$1 - c + w(1 - \langle x \rangle)$	$\langle x \rangle$
Global	Global	Global	$1 - c + w(1 - \langle x \rangle)$	$\langle x \rangle$

First, we consider the impact of contact among agents, wherein the frequency of automatic processing impacts the payoff of controlled processing (via the  $w(1 - \langle x \rangle)$  term in  $f_c$ ). When contact is global, the impact of contact is determined by the average level of automaticity in the population as a whole. When contact is local, the impact of contact is instead determined by the average level of control across the focal agent and her neighbors (of which there are two in the ring-structured population we study). Thus, the strategy of the focal agent has much greater impact on her contact payoff in the local case. In particular, if the focal agent is entirely controlled, then the average level of automaticity in her local neighborhood can be at most  $2/3$  (if both other agents were maximally automatic). As a result, the contact term in the payoff of controlled processing becomes  $\frac{2}{3}w(1 - \langle x \rangle)$  (as seen in rows 1, 2, 5 and 6 of Supplementary Table 1). This means that when automatic agents impose costs on controlled agents,  $w < 0$ , global contact will reduce the payoff of control relative to local contact (and thus reduce the equilibrium level of control). Conversely, when controlled agents benefit from contact with automatic agents,  $w > 0$ , globalizing contact will increase the payoff (and thus equilibrium level) of control.

Next, we consider environment, wherein the frequency of controlled processing determines the state of the environment and thereby the payoff of automatic processing ( $f_a = \langle x \rangle$ ). When

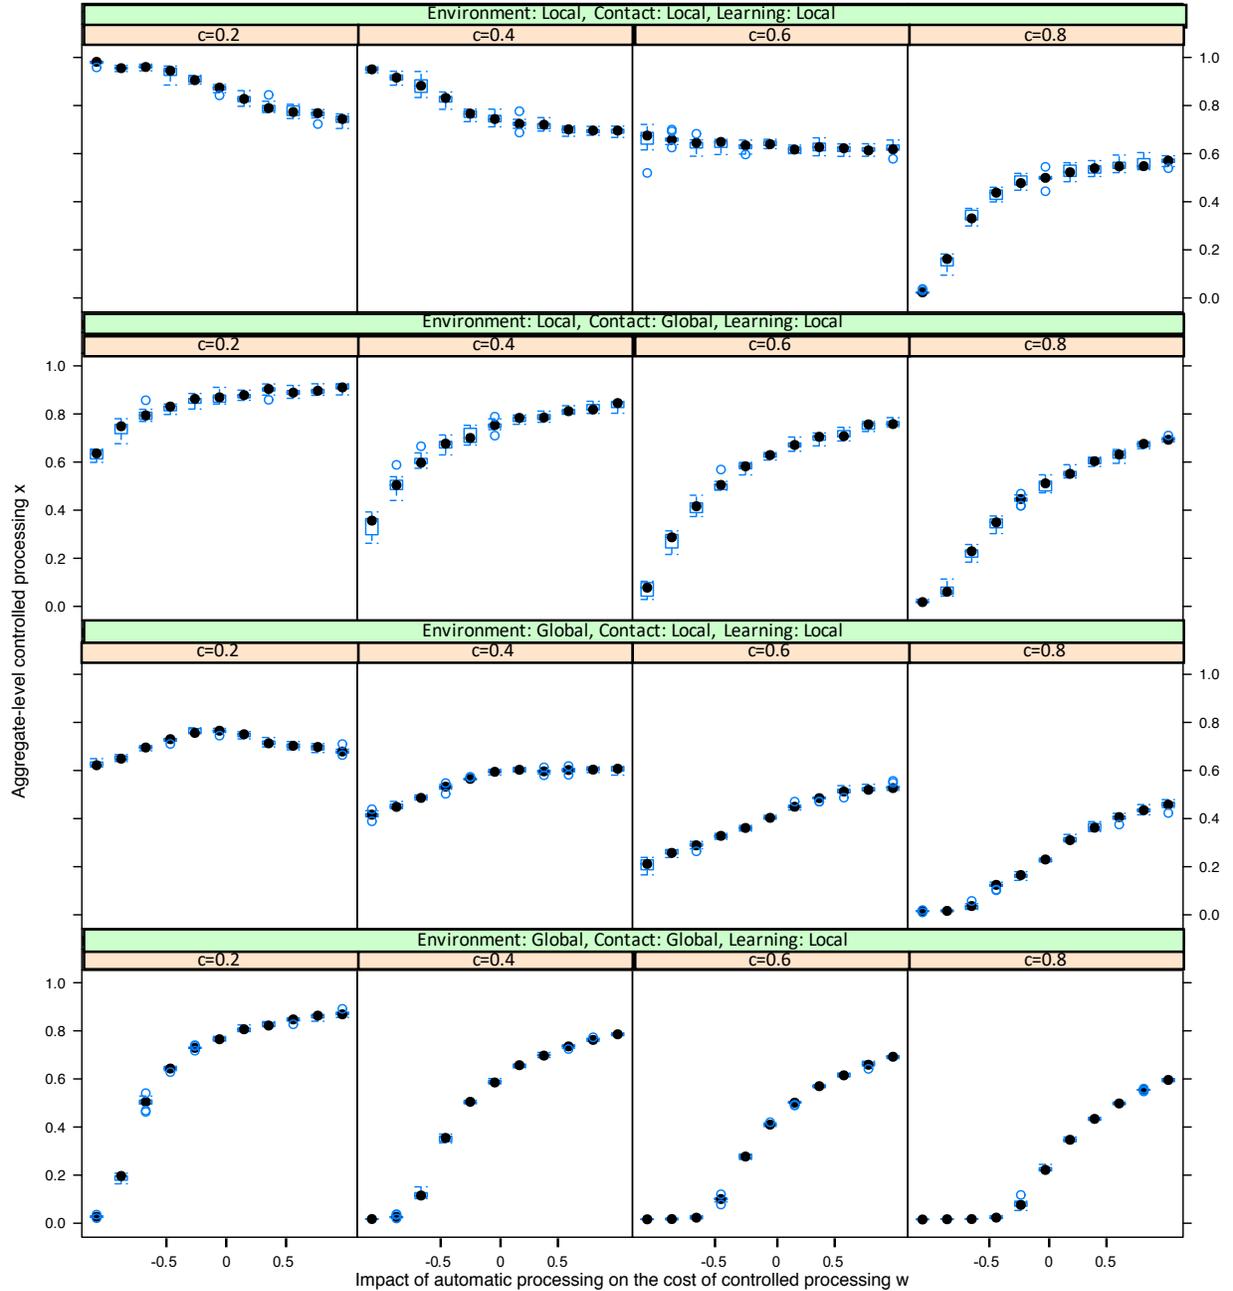
environment is global, all agents share the same environment described by the variable  $p$ , the value of which is determined by the average level of automaticity across the whole population. When environment is local, each agent experiences her own local environment that is determined by the average level of automaticity over her two neighbors and herself. Thus, as with contact, the focal agent's strategy has a much bigger impact on her local environment when environment is local compared to global. In particular, if the focal agent is fully automatic, then the average level of control in her local neighborhood can be at most  $2/3$  (if both other agents were maximally controlled). Furthermore, when environment is local, the focal agent also has a much larger impact on the environments experienced by – and therefore the optimal strategies of – her neighbors. If the focal agent is fully automatic, she decreases the  $p$  experienced by her neighbors, thereby creating selection pressure for them to become more automatic. Although the exact functional form of this impact is complex, we capture it conceptually by introducing a factor  $\beta \leq 1$  that diminishes the level of control of an automatic agent's neighbors. Combining these two consequences of local interaction, the payoff of automatic processing becomes  $\frac{2}{3}\beta < x >$  when contact is local (as seen in rows 1-4 of Supplementary Table 1). This means that regardless of model parameters, the payoff of automatic processing is lower (and therefore the frequency of control is higher) when contact is local relative to global.

These observations also imply a *synergy* between the locality of contact and environment, driven by local environment's impact on the optimal strategy of the focal agent's neighbors. As described in the preceding paragraph, local environment causes clustering of strategies (i.e., more automatic agents' neighbors to be more automatic and, symmetrically, more controlled agents' neighbors to become more controlled). When contact is also local, the neighbors which a controlled agent interacts with directly are more likely to be controlled by a factor  $\beta$ , leading to a further scaling of the  $w$  term in  $f_c$  (as seen in rows 1 and 2 of Supplementary Table 1). Thus, local environment amplifies the impact of contact on the average frequency of control (which, as described above, can be positive or negative depending on the sign of  $w$ ).

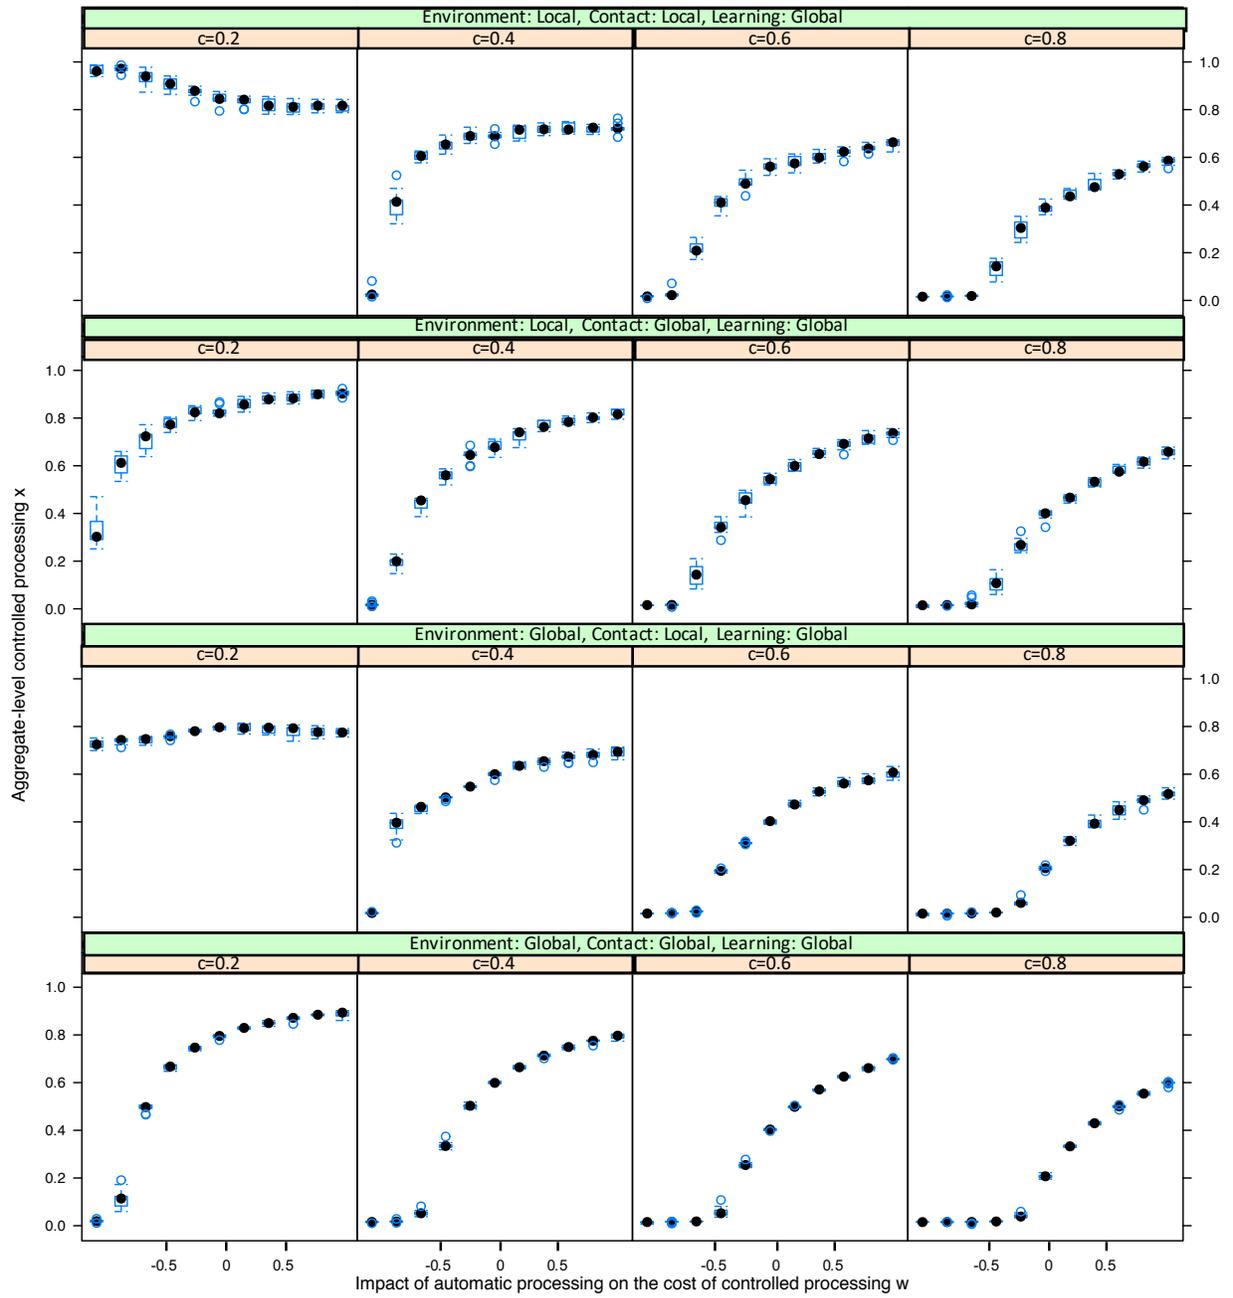
Finally, we consider learning. When learning is global, no clustering or assortment in strategies arises due to the learning process. Thus (except for the assortment that arises from local environment captured by  $\beta$ ), the average level of control among a focal agent's neighbors is equal to the population level of control  $< x >$ . When learning is local, however, this leads to assortment, such that a focal agent's neighbors are more likely than chance to have the same strategy as the focal agent. We follow common convention and define the level of assortment  $\alpha \leq 1$ , such that a controlled agent's neighbors will be controlled with probability  $(1 - \alpha) < x > + \alpha$ , whereas an automatic agent's neighbors will be automatic with probability  $(1 - \alpha) < x >$ . (As above, we forgo an exact calculation of the level of  $\alpha$ , as the specific value of  $\alpha$  does not qualitatively change its impact on the evolutionary outcomes). When both contact and environment are global, assortment (and thus local versus global learning) has no effect on the evolutionary outcomes: agents are no more likely to interact with their neighbors than any other agent, and thus the strategies of neighbors versus the population as a whole is irrelevant. When either or both contact/environment are local, however, local learning amplifies the other local effects: controlled agents' neighbors are more likely to be controlled by a factor  $\alpha$  (appearing in the  $w$  term of  $f_c$ , lines 1 and 5 of Supplementary Table 1), and automatic agents' neighbors are more likely to be automatic by a factor  $\alpha$  (appearing in  $f_a$ , lines 1 and 3).

## Supplementary Figures

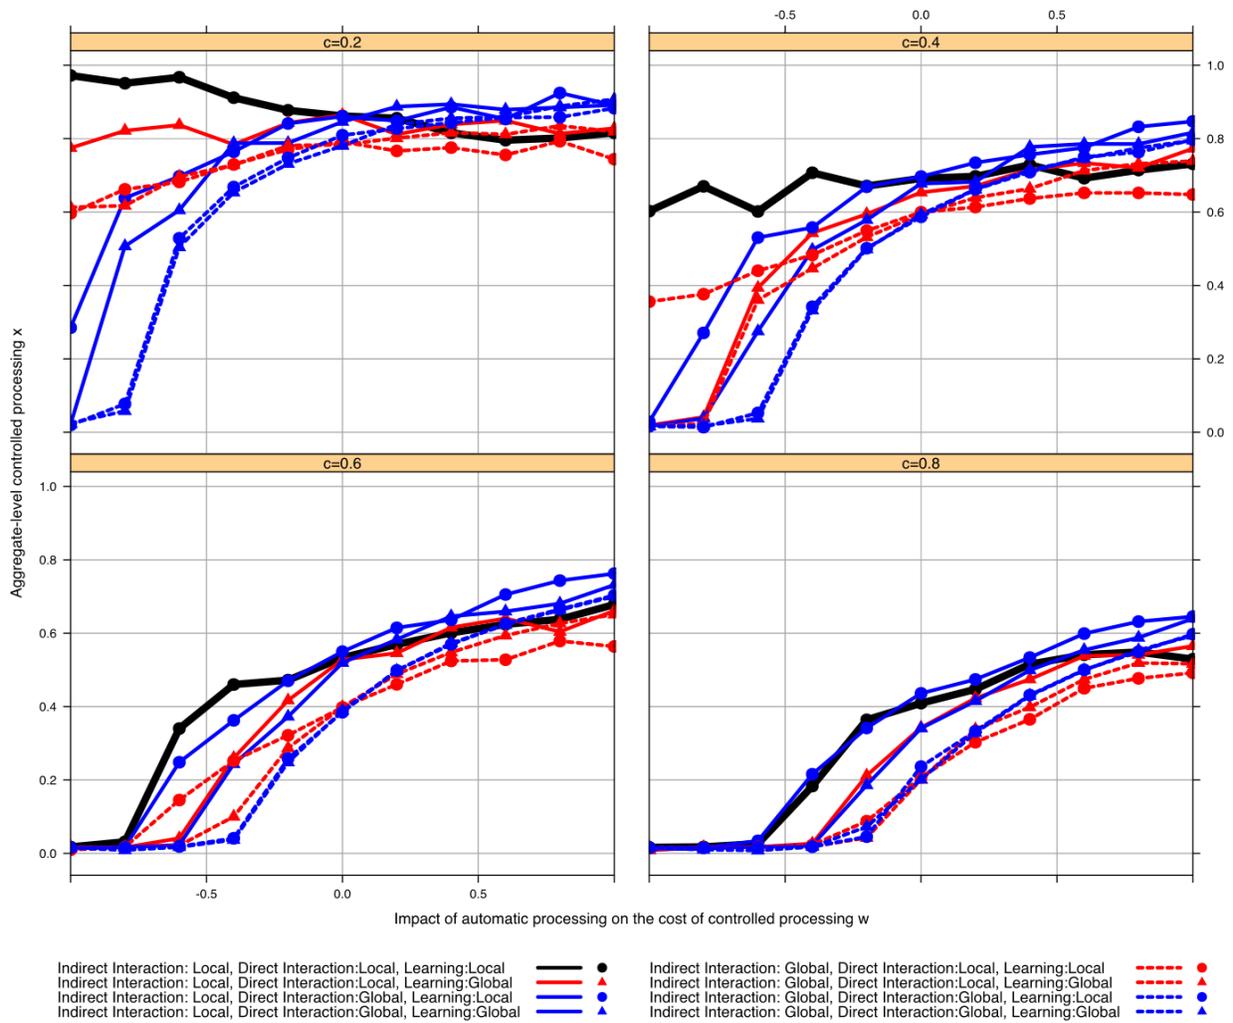
In this section, we provide numerical details of our simulation model and also show additional results for robustness of our result to Small-world networks varying probability of link rewiring.



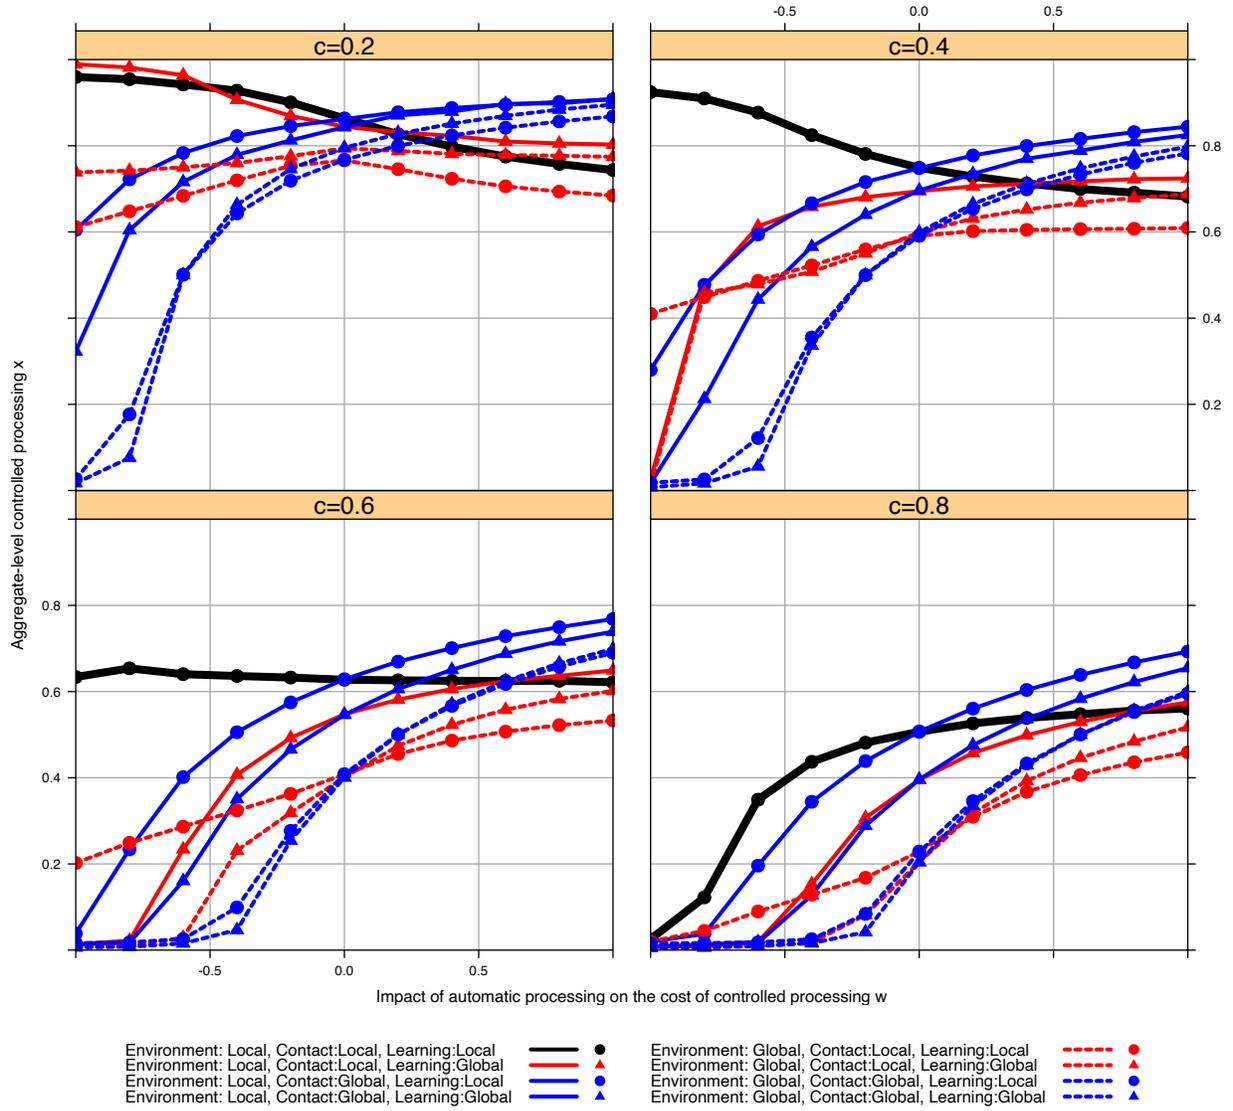
**Supplementary Figure 1. Variations of simulation outcomes for different replicates (local learning).** The boxplots in each panel shows variations of simulation outcome of aggregate-level controlled processing  $x$  for 10 replicates and for each given set of  $w$ ,  $c$ , and localness.



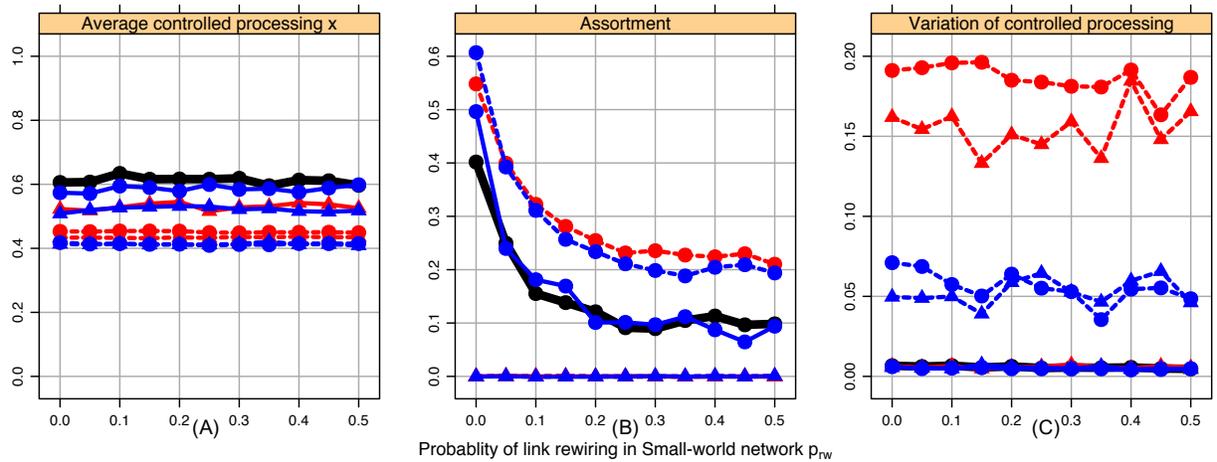
**Supplementary Figure 2. Variations of simulation outcomes for different replicates (global learning).** The boxplots in each panel shows variations of simulation outcome of aggregate-level controlled processing  $x$  for 10 replicates and for each given set of  $w$ ,  $c$ , and localness.



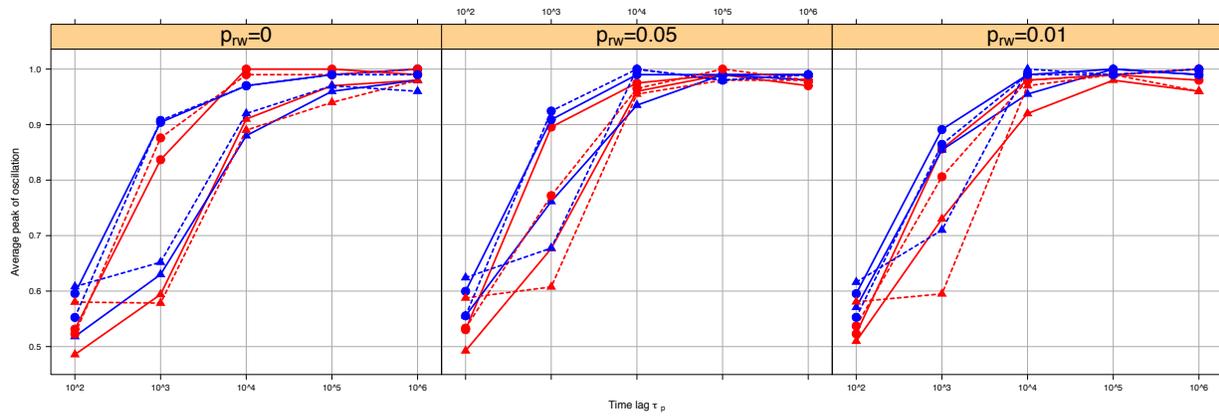
**Supplementary Figure 3.** Aggregate-level controlled processing  $x$  for no lag  $\tau_p = 1$  on Small-world network where the average number of neighbors is 4 and probability of rewiring is 0.2. Each line represents a combination of local versus global contact, environment, and learning, across different levels of the impact of automatic processing on the cost of controlled processing ( $w$ ) and fixed cost of controlled ( $c$ ). Contact is indicated by color (red = local, blue = global). Environment is indicated by line type (solid = local, dashed = global). Learning is indicated by symbol (circle = local, triangle = global).



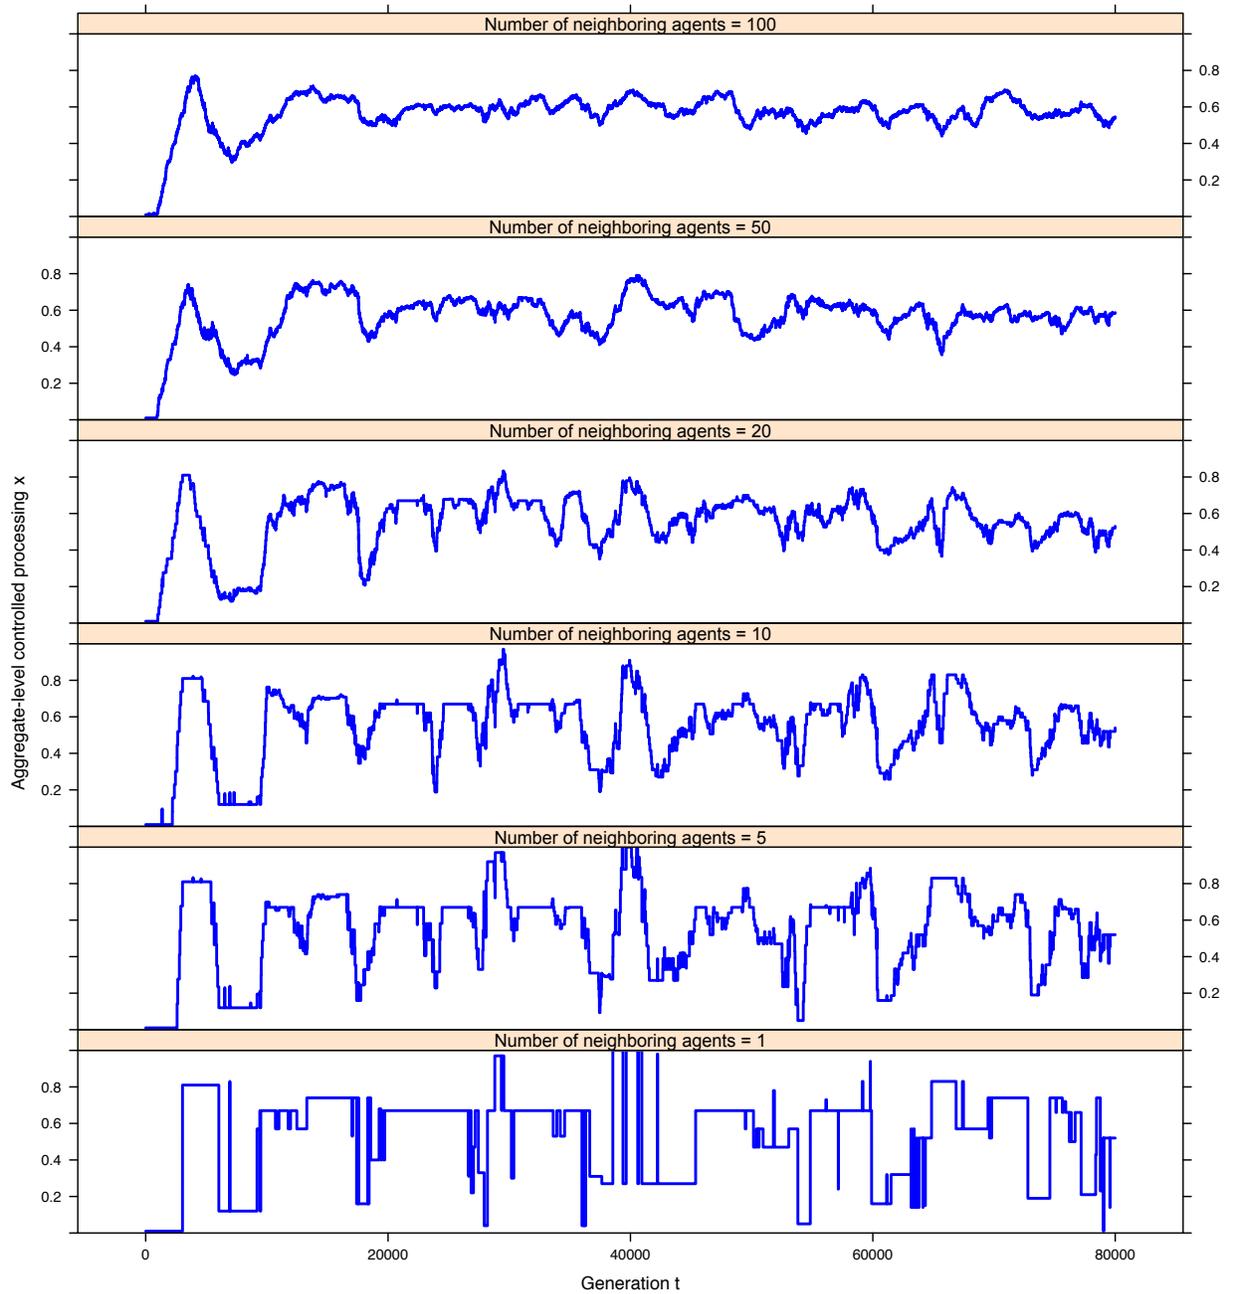
**Supplementary Figure 4. Aggregate-level controlled processing  $x$  for no lag  $\tau_p = 1$  on a ring structure with  $N=10^5$  nodes where the average number of neighbors is 2.** Each line represents a combination of local versus global contact, environment, and learning, across different levels of the impact of automatic processing on the cost of controlled processing ( $w$ ) and fixed cost of controlled ( $c$ ). Contact is indicated by color (red = local, blue = global). Environment is indicated by line type (solid = local, dashed = global). Learning is indicated by symbol (circle = local, triangle = global). Given the size of the network, the simulation was run for  $3.6 \times 10^7$  generations to reach stable outcome for  $x$ .



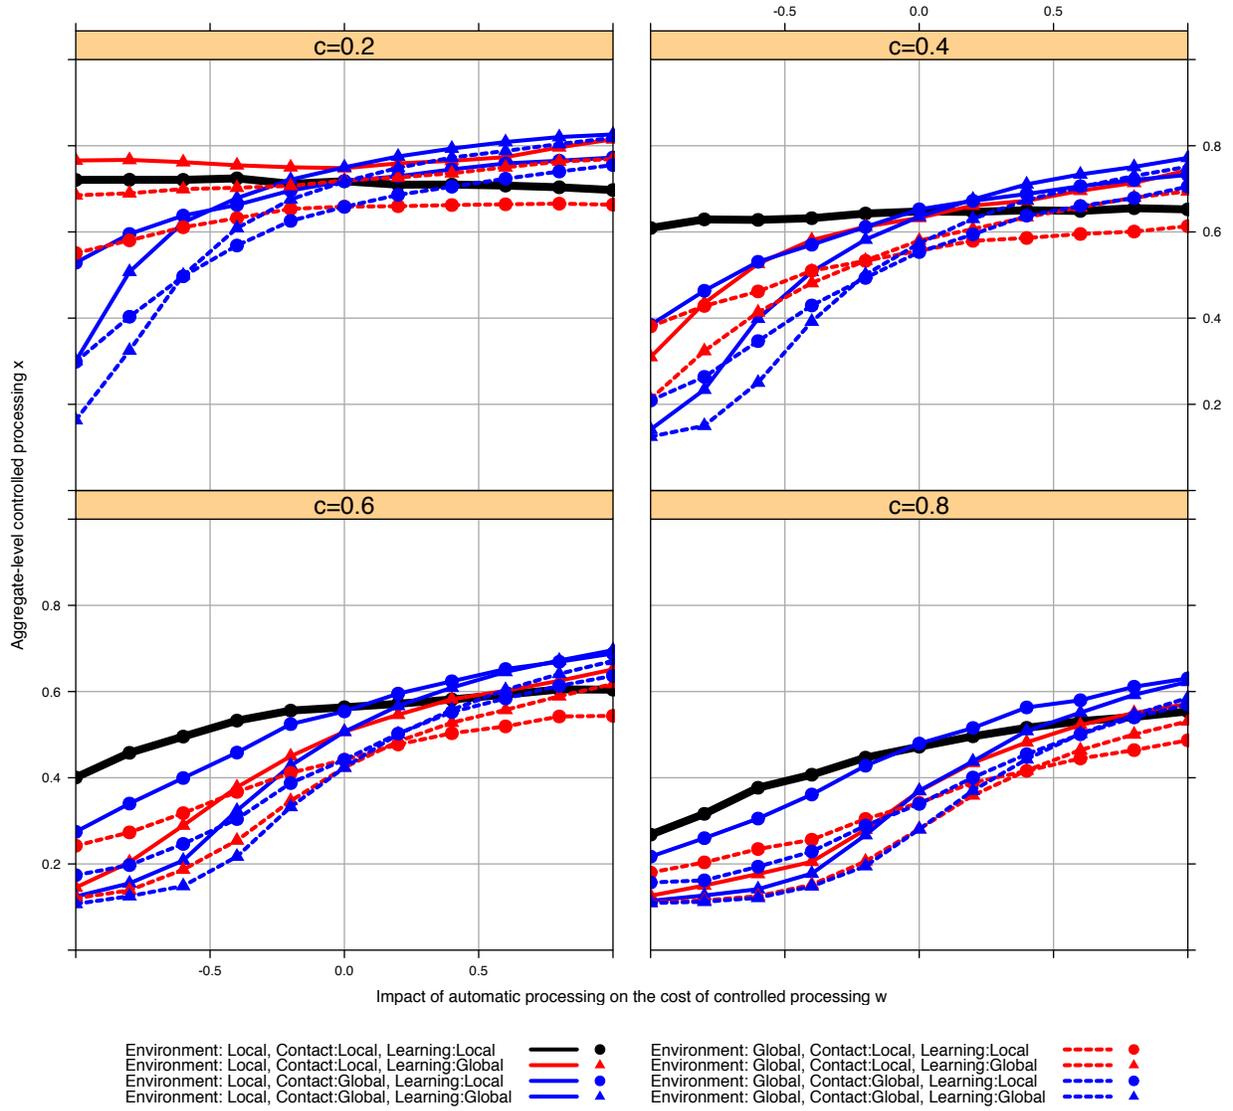
**Supplementary Figure 5. Meso-scale communities measures as a function of probability of link rewiring in Small-world network.** Each panel shows one measure averaged over last  $4 \times 10^4$  generations for 10 simulation replicates. Results are generated on Watts-Strogatz Small-world networks, where each agent was connected to two other neighbors on each side ( $k=4$ ), impact of automatic processing on the cost of controlled processing  $w=-0.15$  and fixed cost of controlled  $c=0.5$ . All results of the ring structure qualitatively generalize to the Small-world networks. Aggregated controlled processing  $x$  and variation of strategies do not change with probability rewiring, yet assortment decreases as the probability of rewiring increases due to weaker community structure of the network.



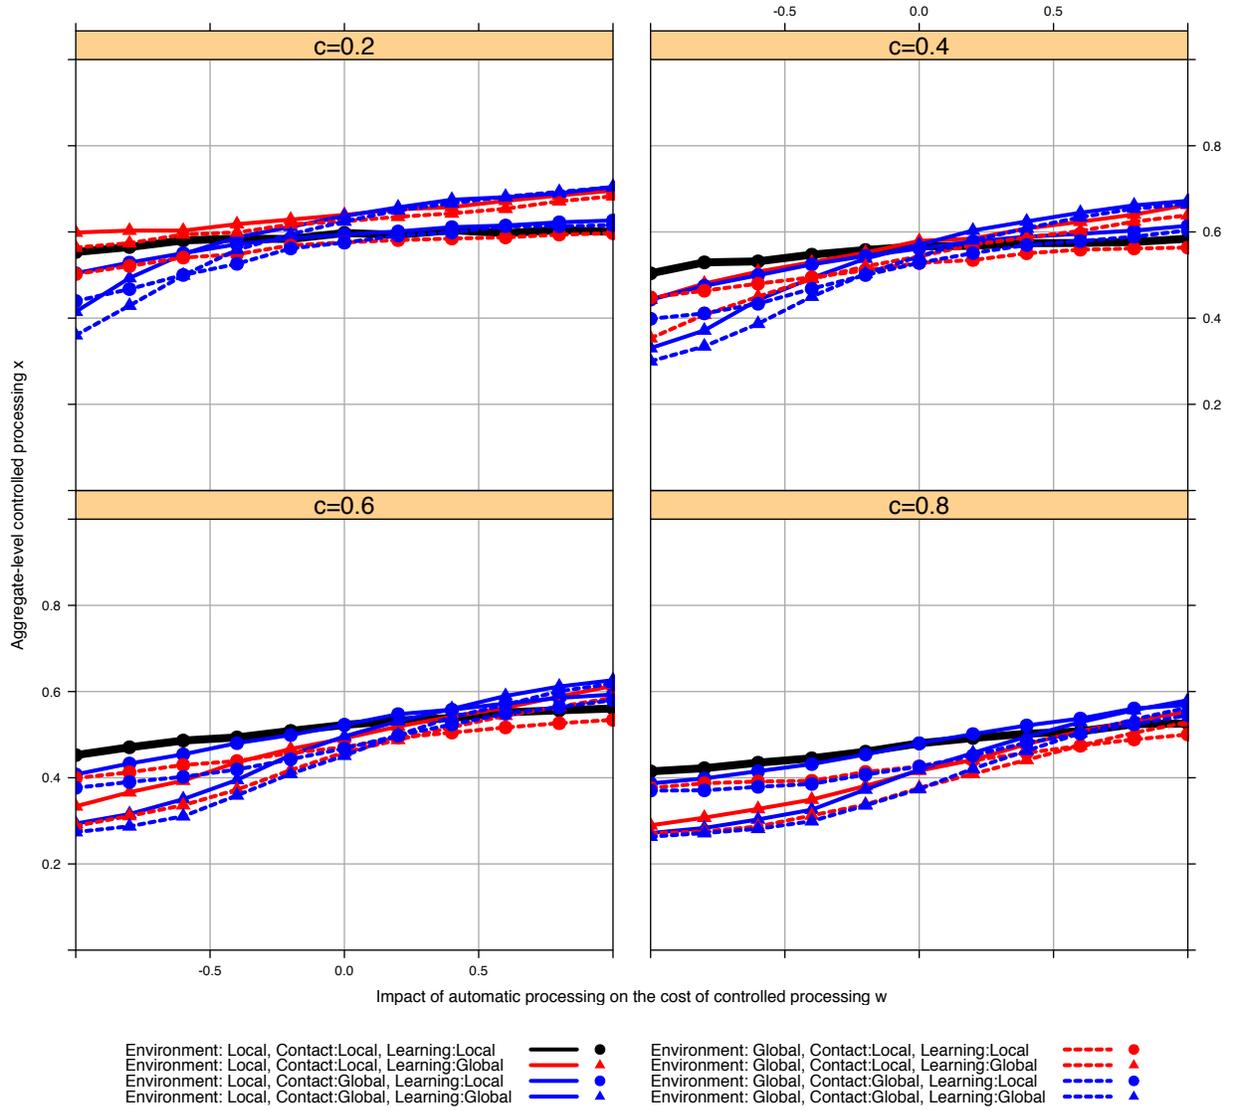
**Supplementary Figure 6. Effect of time-lag on amplitude of oscillation for Small-world networks varying probability of rewiring  $p_{rw}$ .** Increasing probability of rewiring weakens the effect of local learning on amplitude of oscillation as it decreases the strength of community structure in the network.



**Supplementary Figure 7. Local oscillation of neighboring agents with local learning.** Each panel shows oscillation of controlling processing  $x$  for a different number of agents located next to each other on a ring structure. Results are created for time lag  $\tau_p = 10^3$ , the cost of controlled processing  $w = -0.15$ , fixed cost of controlled  $c = 0.5$ , and local learning (the figure is generated with local environment and local contact; however, global environment and/or global contact do not qualitatively change these results).



**Supplementary Figure 8. Aggregate-level controlled processing  $x$  for no lag  $\tau_p = 1$  and mutation  $u=0.2$ .** Each line represents a combination of local versus global contact, environment, and learning, across different levels of the impact of automatic processing on the cost of controlled processing ( $w$ ) and fixed cost of controlled ( $c$ ). Contact is indicated by color (red = local, blue = global). Environment is indicated by line type (solid = local, dashed = global). Learning is indicated by symbol (circle = local, triangle = global).



**Supplementary Figure 9. Aggregate-level controlled processing  $x$  for no lag  $\tau_p = 1$  and mutation  $u=0.5$ .** Each line represents a combination of local versus global contact, environment, and learning, across different levels of the impact of automatic processing on the cost of controlled processing ( $w$ ) and fixed cost of controlled ( $c$ ). Contact is indicated by color (red = local, blue = global). Environment is indicated by line type (solid = local, dashed = global). Learning is indicated by symbol (circle = local, triangle = global).