



Heriot-Watt University
Research Gateway

Network Connectivity Dynamics, Cognitive Biases, and the Evolution of Cultural Diversity in Round-Robin Interactive Micro-Societies

Citation for published version:

Segovia-Martín, J, Walker, B, Fay, N & Tamariz, M 2020, 'Network Connectivity Dynamics, Cognitive Biases, and the Evolution of Cultural Diversity in Round-Robin Interactive Micro-Societies', *Cognitive Science*, vol. 44, no. 7, e12852. <https://doi.org/10.1111/cogs.12852>

Digital Object Identifier (DOI):

[10.1111/cogs.12852](https://doi.org/10.1111/cogs.12852)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Cognitive Science

Publisher Rights Statement:

© 2020 The Authors.

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Cognitive Science 44 (2020) e12852

© 2020 The Authors. Cognitive Science published by Wiley Periodicals LLC on behalf of Cognitive Science Society (CSS). All rights reserved.

ISSN: 1551-6709 online

DOI: 10.1111/cogs.12852

Network Connectivity Dynamics, Cognitive Biases, and the Evolution of Cultural Diversity in Round-Robin Interactive Micro-Societies

José Segovia-Martín,^a  Bradley Walker,^b  Nicolas Fay,^b 
Monica Tamariz^c 

^a*Faculty of Arts and Humanities, Universitat Autònoma de Barcelona*

^b*School of Psychological Sciences, University of Western Australia*

^c*Psychology, School of Social Sciences, Heriot-Watt University*

Received 8 November 2019; received in revised form 18 March 2020; accepted 24 April 2020

Abstract

The distribution of cultural variants in a population is shaped by both neutral evolutionary dynamics and by selection pressures. The temporal dynamics of social network connectivity, that is, the order in which individuals in a population interact with each other, has been largely unexplored. In this paper, we investigate how, in a fully connected social network, connectivity dynamics, alone and in interaction with different cognitive biases, affect the evolution of cultural variants. Using agent-based computer simulations, we manipulate population connectivity dynamics (early, mid, and late full-population connectivity); content bias, or a preference for high-quality variants; coordination bias, or whether agents tend to use self-produced variants (egocentric bias), or to switch to variants observed in others (allocentric bias); and memory size, or the number of items that agents can store in their memory. We show that connectivity dynamics affect the time-course of variant spread, with lower connectivity slowing down convergence of the population onto a single cultural variant. We also show that, compared to a neutral evolutionary model, content bias accelerates convergence and amplifies the effects of connectivity dynamics, while larger memory size and coordination bias, especially egocentric bias, slow down convergence. Furthermore, connectivity dynamics affect the frequency of high-quality variants (adaptiveness), with late connectivity populations showing bursts of rapid change in adaptiveness followed by periods of relatively slower change, and early connectivity populations following a single-peak evolutionary

Correspondence should be sent to Monica Tamariz, Psychology, School of Social Sciences, Heriot-Watt University, Mary Burton Building, Room MB G1, Edinburgh, UK. E-mail: m.tamariz@hw.ac.uk

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

dynamic. We evaluate our simulations against existing data collected from previous experiments and show how our model reproduces the empirical patterns of convergence.

Keywords: Cultural evolution; Adaptiveness; Connectivity; Network topology; Content bias; Coordination bias; Memory; Punctuational evolution

1. Introduction

Human life is shaped by our culture, that is, by socially transmitted information that determines our behavior, beliefs, attitudes, and values (Richerson & Boyd, 2008). Cultural variants such as technology, language, and beliefs propagate in populations following evolutionary dynamics (Boyd & Richerson, 1988; Cavalli-Sforza & Feldman, 1981; Neiman, 1995)—individuals inherit cultural traits from ancestors or peers and occasionally generate new trait variants. Over generations, cultures evolve: Some variants are lost, while others spread in a population, sometimes to the point of fixation, when we can say the population has converged on a variant.

Drift models have been used to explain cultural evolution (Bentley, Hahn, & Shennan, 2004) and the evolution of human communication systems (e.g., Blythe, 2012; Reali & Griffiths, 2009). When evolution is neutral (drift dynamics), all variants are equally likely to be adopted by individuals, and therefore to propagate in the population. In the absence of innovation, drift may cause cultural variants to disappear and, in turn, reduce cultural variation. Drift models can thus lead, over generations, to convergence on a single variant. This is particularly true for smaller populations because random sampling can remove variants faster, resulting in a change of variant distributions over generations. A variant has been fixed when its frequency is 1, and it has been extinguished when its frequency is 0. Therefore, since the influence of stochastic variation is higher for small populations, the probability of fixation is also higher in smaller populations (Frankham, Briscoe, & Ballou, 2002). Interestingly, it is well known that drift may explain the propagation of cultural variants, including baby names, pottery decorations, and patents (Bentley et al., 2004); dog breeds (Herzog, Bentley, & Hahn, 2004); and some diachronic changes in language (DeGraff, 2001; Komarova & Nowak, 2003; Kroch, 1989). These applications and properties make drift models usable as null models against which other models can be tested (Hahn & Bentley, 2003; Lipo, Madsen, Dunnell, & Hunt, 1997; Neiman, 1995; Reali & Griffiths, 2009; Shennan & Wilkinson, 2001).

1.1. Cognitive biases and memory

Some variants spread more rapidly than others. In these cases, evolution is not neutral, but subject to biases, or selection pressures. *Content bias*, also termed *direct bias* by Boyd and Richerson (1985), relates to individuals' sensitivity to intrinsic properties of traits, and results in the more learnable, efficient, or effective variants having a higher probability of being adopted by others (Hagen & Hammerstein, 2006; Henrich & McElreath, 2007; Vale et al., 2017), and therefore spreading faster through a population than a

neutral, drift model would predict (Gong, Shuai, Tamariz, & Jäger, 2012; Tamariz, Ellison, Barr, & Fay, 2014). *Coordination biases* may involve a preference to use variants we have used before (egocentric bias) or variants produced by others (allocentric bias). In communicative tasks, for instance, an allocentric bias can be observed under certain mechanistic accounts of dialogue, since speakers tend to adopt labels used by their interlocutors (Garrod & Pickering, 2007; Pickering & Garrod, 2004), which, in turn, favors cooperation and coordination (Fusaroli et al., 2012; Fusaroli & Tylén, 2016). Content and coordination biases also interact with each other: Egocentric bias maintains variation, which improves the chances that content bias will select for the most adaptive variant in a population (Segovia-Martín, Walker, Fay, & Tamariz, 2019).

Cultural transmission is also affected by the memory record of cultural variants. Some authors claim that the type of variation that learners produce can be explained by memory limitations: For example, memory can affect language regularization (Hudson Kam & Chang, 2009; Hudson Kam & Newport, 2005), compressibility (Chater & Vitányi, 2003), or conventionalization (Tamariz & Kirby, 2015). In general, memory limitations reduce variation (Tamariz & Kirby, 2015). Because frequency learning is a prominent aspect of social learning in linguistic and nonlinguistic tasks, regularization behavior is consistent with a domain-general account of the observed regularization bias and it might be attributable to limited working memory (Ferdinand, Thompson, Kirby, & Smith, 2013). Ferdinand et al. (2013) also suggest that a tendency to produce representative variants can lead to high-fidelity reproduction of the training set of variants under low memory constraints. This idea is very relevant to the scenarios that we aim to model, in which there is production and reproduction of a discrete number of variants with high fidelity. In addition, it allows us to model memory as a cognitive feature that constraints regularization on the basis of frequency learning.

The effects of content and coordination biases on variant propagation were studied by Tamariz et al. (2014), who constructed a parameterized model of cultural variant transmission to analyze the patterns of variant spread obtained in an experimental study by Fay, Garrod, Roberts, and Swoboda (2010). Fay et al. (2010) had groups of eight individuals playing a Pictionary-like communication game in pairs. During each game, the *director* produced a drawing to represent each of 16 meanings, one at a time. For each of these, the *matcher* tried to guess which of 20 possible meanings the director was trying to communicate. Participants played this game six times with each partner, with roles reversing for each game, so each participant drew and matched each meaning three times with each partner. For a given meaning, directors could invent their own ways to depict the meaning (i.e., produce a novel variant) or produce a variant that was produced by a partner or by themselves in a previous game. After six games, participants swapped partners within their group and played another six games with their new partner. This partner-swapping was repeated until the populations were fully connected, with every participant having played with every other participant. This meant that, by the end of the experimental simulation, for each meaning, a particular variant could spread to the whole population. In other words, the population could converge on a single variant for each meaning. Tamariz et al. (2014) found that this convergence was best explained by a

combination of egocentric bias and content bias, where participants would stick with variants they had produced previously, unless they encountered a better variant, in which case they would switch to that.

1.2. Population structure, network dynamics, and specific social learning scenarios

Demographic factors also add selection pressures (e.g., Henrich, 2004; Mesoudi, 2011; Mesoudi, Whiten, & Laland, 2006; Richerson & Boyd, 2008; Shennan & Wilkinson, 2001; Vaesen, 2012): The degree of adaptiveness, complexity, and cumulative cultural evolution of cultural and communicative variants can be affected by population size (Cuskley, Loreto, & Kirby, 2018; Derex, Beugin, Godelle, & Raymond, 2013; Henrich, 2004; Kempe & Mesoudi, 2014; Kline & Boyd, 2010; Kobayashi & Aoki, 2012; Shennan & Wilkinson, 2001), by the degree of contact and migration between populations (Creanza, Kolodny, & Feldman, 2017; Muthukrishna, Shulman, Vasilescu, & Henrich, 2014; Powell, Shennan, & Thomas, 2009), and by the structure of the social network (Gong, Minett, & Wang, 2008; Lee, Stabler, & Taylor, 2005; Lupyán & Dale, 2010; Mueller-Frank, 2013; Olfati-Saber & Murray, 2004). An additional demographic variable, namely the connectivity between individuals within or across populations, also enhances adaptiveness and complexity because it affects the degree of diversity each individual has access to (Henrich, 2004; Kobayashi & Aoki, 2012; Powell et al., 2009; Shennan & Wilkinson, 2001). However, a recent study suggests that access to diversity is not the only variable at play, and that if we take into consideration the potential for an innovation to be adopted and spread, then an intermediate degree of connectivity may be optimal for cumulative culture, as too much connectivity stifles innovation, while too little cannot maintain complex traits (Derex & Boyd, 2016; Derex, Perreault, & Boyd, 2018).

In the evolution and history of human populations, the structure of the population might have played an important role in cultural change in ancestral and historical periods (Derex & Boyd, 2016). Critically, it should be noted that population fragmentation and cultural isolation have been identified as crucial factors to explain the spread of cultural variants such as high-quality ideas (Björk & Magnusson, 2009), technology (Hovers & Belfer-Cohen, 2006), or research (March, 2005). Furthermore, it has been suggested that inter-population connectivity may be more than just a simple reflection of cultural accumulation, and that it may be a critical driver of cultural change (Creanza et al., 2017).

Important efforts to model the effects of social structures and network dynamics on the spread of cultural variants have been made in recent years. Particularly interesting for the present study is the work that follows Axelrod's model of dissemination of culture (Axelrod, 1997). This model was based on the assumption that people are more likely to interact with others who share the same cultural variants, and this, in turn, tends to increase the number of variants they share. These mechanisms, named homophily and influence, have been shown to be prominent explanations for the persistence of cultural diversity. Expanding on Axelrod's (1997) model, researchers have explored social learning and network connectivity effects by implementing new manipulations, which in general involve complex interactions of agents' cognition and social structures. Particularly relevant for

us are agent-based models that contemplate an interaction between imitation choice rules and dynamic networks, such as the modeling of the co-evolution of networks and agent's preferences (Centola, Gonzalez-Avella, Eguiluz, & San Miguel, 2007), network efficiency and conformity biases (Barkoczi & Galesic, 2016), and network connectivity and group performance (Lazer & Friedman, 2007). All these models manipulate specific features of the network topology and agent's cognition to predict the spread of cultural variants and the convergence on shared cultural variants. It is important to recognize that each of these agent-based models is based on a number of tractability assumptions that, depending on the focus of the study, lead to paying more attention to either macro-processes or micro-processes. The ability of many of these models to make global predictions takes focus away from micro-processes such as the composition of the pairings at the most elementary level, limiting the tractability of each of the agents that take part in the cultural process.

In the real world, the order in which individuals pair over time is sometimes determined by stochastic events that govern path dependencies, that is to say, a limited record of experiences that one has experienced in the past. When individuals in a community are carriers and transmitters of information, different orderings can yield different levels of sub-population isolation at different times, which, in turn, affects the probability of dissemination of information in social groups. For example, the specific pair composition within a micro-society determines the number of agents that can potentially share the same variant at different times. To be more specific, the present study can be applied to a number of sociocultural scenarios that are governed by turn-based interventions (Sacks, Schegloff, & Jefferson, 1978), in which the organization of turn taking might play an important role in, for example, the formation of sides during jury deliberations (Manzo, 1996; Stone, 1969), computer-mediated communication (Garcia & Baker Jacobs, 1999), the evolution of communicative conventions in the lab (Bloom, Russell, & Wassenberg, 1987; Fay, Garrod, & Roberts, 2008; Fay et al., 2010), or communication in group decision-making (Bormann, 1996; Stasser & Taylor, 1991). This might include certain group dynamics in the workplace, job selection processes, or strategies in team games. What these scenarios have in common is that they usually

1. occur in a relatively short period of time;
2. take a relatively small population;
3. consist of a relatively small number of interactions between actors;
4. are systematically structured by a prefixed scheduled algorithm; and
5. have a tendency to require convergence as an outcome (e.g., jury verdict or final outcome of a decision-making process).

These properties make these sociocultural scenarios sensitive to subtle changes in the network structure, which can potentially enhance the impact of the formation of short-term clusters (i.e., groups of agents with strong ties or who interact more frequently with one another) and individual biases on the outcome of the process. The formation of short-term clusters can lead to differential information concentration that can dramatically

affect the spread of highly valued variants within the group (Axelrod, 1997; Centola, 2018; González-Avella, Cosenza, Klemm, Eguíluz, & San Miguel, 2007).

It is important to stress that given similar outcomes, differential temporal pathways toward convergence might have dramatic effects on specific social learning scenarios. When there is an increasing pressure on actors to accelerate decision-making, the importance of partial states increases. For example, time-dependent decision-making models have shown partial outcomes to be crucial to determine the quality of research when scholars face the exploration–exploitation dilemma (Chavalarias, 2017). These time-dependent learning metrics can be amplified in multiple-choice tasks in which opinion is based on one-shot experience per option, such as wine tasting (Mantonakis, Rodero, Lesschaeve, & Hastie, 2009), or even affect juror’s decision-making after a sequence of systematically organized sequence of speeches (Shteingart, Neiman, & Loewenstein, 2013; Stone, 1969). Although these scenarios can vary widely, in this study we explore how the time to convergence can be affected by connectivity dynamics.

Using the computer model described below, we attempt to model social network structures such that they are characterized by a different schedule, which leads to the formation of differential short-term clusters that might serve as a reasonable proxy for scenarios as those mentioned above. While clustering and social influence in social networks have been studied empirically (Becker, Brackbill, & Centola, 2017; Centola & Baronchelli, 2015), there are relatively few formal models that examine the combined effect of cluster formation and individual influence (Centola et al., 2007; Muthukrishna & Schaller, 2019). These models have successfully shown how the co-evolutionary dynamics of social influence and network structures can affect cultural change. In our model, we explore specific evolutionary dynamics of network structures and individual biases. We implement content biases, coordination biases, and memory, three parameters that have been shown as crucial in the acquisition of variants under controlled conditions when using small-scale societies (Tamariz et al., 2014).

1.3. Connectivity dynamic and path dependencies

A key innovation of the present study is that we aim to address the effect of different orderings by focusing on a particular microscale manipulation of the social network that remains largely unexplored: the order in which connections between individuals unfold over time. We call this the network connectivity dynamic. The most basic network topology is a fully connected network, in which all nodes are interconnected. In a population of individuals, this means that, over time, each individual interacts with every other individual. However, the same fully connected network may follow different temporal patterns of connectivity, and this may have consequences for variant spread. For instance, a pattern may lead to temporary isolation of one or more sub-populations, generating path dependencies that affect the dissemination of information in the social network, in a similar way to how long-range ties might affect the population dynamics of beliefs, opinions, and polarization (Turner & Smaldino, 2018) or stochastic events in the evolution of cooperation (Smaldino & Schank, 2012). Importantly for the present study, Turner and Smaldino (2018)

show that the order in which agents interact, all else being equal, has a dramatic influence on the long-term polarization of the population. This finding motivates us to propose a re-evaluation of a specific small-scale cultural evolutionary model of dissemination of variants by systematically manipulating the effect of the temporal network connectivity dynamics on a range of various outcomes and how this interacts with other previously established factors, such as content-based bias, coordination bias, and memory size. In such small sub-populations, drift can reduce diversity, disproportionately favoring variants that happen to be present in the population, and which are not necessarily adaptive (Henrich, 2004). A different connectivity pattern may never yield pockets of isolation, which would lead to different evolutionary dynamics. In this paper, we address how the connectivity dynamic, alone and in combination with content and coordination biases, affects the spread of cultural variants in a population. To manipulate the connectivity dynamic, we draw on a well-known pairing methodology called round-robin tournament, which allows us to optimally control that the probability of contribution of each agent to the evolutionary process is the same. This strict control of the connectivity dynamic condition is important to then examine the robustness of our findings by experimenting with the manipulation of agents' cognition. Three different connectivity dynamics are used in this study, each of which yields a different level of temporal sub-population isolation at a given time (see Section 2).

To our knowledge, the specific manipulation that we propose here has not been taken into account in experimental work or models on cultural and language evolution, and it would be especially relevant to those researchers that use dynamic interactive micro-societies of agents switching partners over time (e.g., communicative games, cooperative games, or tournaments; Baum, Richerson, Efferson, & Paciotti, 2004; Byun, De Vos, Roberts, & Levinson, 2018; Caldwell & Smith, 2012; Fay et al., 2008, 2010; Mesoudi & Whiten, 2008; Raviv, Meyer, & Lev-Ari, 2019; Tamariz et al., 2014). In most cases, experimental designs of micro-societies of interacting actors only include one pair composition out of all the possible combinations of pair shuffling, and therefore, outcomes are related with only one specific population connectivity dynamic, potentially affecting the accuracy of the generalizations made by these studies.

We use the experimental design of Fay et al. (2008, 2010) and Tamariz et al. (2014) as a starting point for constructing our model. Tamariz et al.'s (2014) study was designed to test whether the observed variant distributions obtained by Fay et al. (2010) were best explained by neutral drift or showed evidence of selection and adaptation mediated by content and/or coordination biases. Their results indicated an interaction of both biases: Participants displayed egocentric bias and content bias; they tended to produce the variants they had previously produced themselves except when they encountered a better (simpler, cleverer, etc.) variant (through mutation/innovation or via a partner), in which case they tended to adopt it. Additionally, participants seemed to have *short memory* and tended to produce mostly variants that they had seen or produced in the preceding couple of games. Although the model in Tamariz et al. (2014) was useful to show an interaction between content and coordination biases and depict the number of data structures explained by the biases in the experiment, predictions on the specific weight of each bias on the spread of variants were not determined due to the explanatory nature of the model

implemented, which took as input the history of the representational variants that the participants had used or seen during the experiment. The interactions between memory and the cognitive biases examined were not explored either.

For the present study, we extend Tamariz et al.'s (2014) study in two important ways. First, we add a new manipulation: population connectivity dynamic, in addition to content, coordination biases, and memory size. We will evaluate the importance of individual parameters and their interactions on the evolution of convergence. Our model does not take experimental data as input, but will be initialized with a random computationally generated distribution of cultural variants. Importantly, we use our model to produce data that can be later used to make predictions or compared to the experimental datasets that have been mirrored. Second, our model aims to be a valid tool to find the best model fit of connectivity conditions associated with existing experimental datasets that use interactive pairwise micro-societies based on round-robin tournaments applied to the dissemination of cultural variants. This is done by exactly imitating the schedule algorithm of existing experimental designs with human participants. These two innovations allow us to establish causal links between properties at the level of individual agents (content and coordination biases, memory size) and of the population (connectivity dynamics and population size), on the one hand, and properties of the culture (evolution and adaptiveness of variants), on the other.

In light of previous research, two important questions are addressed in this paper. On the one hand, to what extent do the interactions between content bias, coordination bias, and memory affect the evolution of cultural diversity in round-robin interactive micro-societies? These analyses replicate and extend previous work. On the other hand, do connectivity dynamics affect cultural diversity in a population, and do they interact with the cognitive biases and memory size, modifying the rate of convergence and the adaptiveness of cultural variants during cultural evolution?

2. Methods

We constructed an agent-based model that simulates micro-societies of agents who interact in pairs for a number of rounds. We look at the spread of n competing variants, each of which is originally unique and produced by each agent in the simulation. We systematically manipulated the values of several global and individual parameters, including connectivity dynamics and cognitive biases (see Section 2.4), and quantified the resulting changes in the evolution of the convergence and diversity of cultural variants (see Section 2.5). We also provide comparison of our model by testing our simulated data against the real experimental data collected by Fay et al. (2010) and coded by Tamariz et al. (2014) (see Section 4).

2.1. Purpose

The purpose of the model is to understand how the interaction between cognitive biases, memory, and the order in which agents pair with each other over time affects convergence. It also aims at evaluating the relative importance of each parameter

combination and making predictions on the evolution of cultural diversity. Finally, it is tested against experimental data to find the best model fit associated with different connectivity conditions.

2.2. The model

We consider a simplified micro-society of agents, each of whom is characterized by a number of state variables as described in Table 1. The micro-society initially contains N agents, who pairup and interact in pairs for a number of rounds (R). Each interaction consists of a cultural variant exchange (Fig. 1). We simulated pairwise interactive micro-societies of $N = 8, 16,$ and 32 agents, allowing us to track all agent pairings during $N - 1$ rounds according to the pair composition illustrated in Fig. 2. This type of scheduling algorithm is called round-robin tournament, and it allows every agent to be paired with every other agent for exactly one round. In our micro-societies, N is even, which means that in each $N - 1$ rounds, $N/2$ games (or pairwise interactions) can be run concurrently.

Different algorithms can be used to create a schedule for a round-robin tournament and all of them contain the same number of games $N(N - 1)/2$. In our model, agents interact in each game by exchanging information; therefore, not all the schedules allow the same spread of adaptive information because different schedules can pairup agents forming different levels of sub-population isolation at different times. We implement three different connectivity dynamics that differ with regards to how many agents could potentially share the same variant at a given time. In other words, each connectivity dynamic examined (early, mid, and late) differs from the others in how fast a cultural variant could potentially spread out among the agents of the micro-society (Fig. 2).

In our model, each agent i has its individual cultural attributes defined as a matrix of cultural variants. The state of an agent i is a $2 \times r$ matrix of cultural variants, which we call A ,

$$A = \begin{bmatrix} \sigma_{ie1} & \sigma_{ie2} & \dots & \sigma_{ier} \\ \sigma_{ia1} & \sigma_{ia2} & \dots & \sigma_{iar} \end{bmatrix}, \tag{1}$$

where σ_{ier} corresponds to a cultural variant produced by agent i in round r , and σ_{iar} corresponds to a cultural variant observed by agent i in round r . Thus, the state of agent i is its record of produced and observed cultural variants, which are stored in a two-row and r -column matrix. Produced variants are stored in the memory as egocentric, designated as e . Observed variants are stored in the memory as allocentric, designated as a .

2.3. Model dynamics

The dynamics of the model are defined by the following steps. In the initial state, each agent i is randomly assigned a cultural variant σ_i selected from a pool of n distinct cultural variants without repetition. Each variant originally assigned to each agent is unique, and it is the first variant they produce in their first interaction at round $0, r_0$. The initial pool of variants in the micro-society can be defined as a V vector of length N that

Table 1
Parameters, state variables, and scales

Model parameters				
Entity	Parameter	Symbol	Number of levels	Value(s)
Agent	Content bias	β	11	0.0 to 1.0 in steps of 0.1
	Coordination bias	c	11	0.0 to 1.0 in steps of 0.1
	Memory	m	4	1, 3, 5, 7
	Innovation	μ	1	0.02
	Variant quality	s	2	[0,1]
	Agents' sensitivity to variant quality	b	11	0.0 to 1.0 in steps of 0.1
	Variant in agent's memory record	d	2	[0,1]
	Flat distribution of variants	$\varphi(x)$		
	Egocentric entries for a given history	$h E,m$		
	Allocentric entries for a given history	$h A,m$		
	Distribution of egocentric variants for a given history	$f(h E,m)$		
	Distribution of allocentric variants for a given history	$f(h A,m)$		
	Probability distribution of variants for a given history	$P(x h)$		
	Agent state (agent cultural attributes)	A		
	Cultural variant	σ		
	Agent ID	i, j, \dots, n		
	Produced variant by agent i	σ_i		
	Stored variant in egocentric memory by agent i	σ_{ie}		
	Stored variant in allocentric memory by agent i	σ_{ie}		
	Global	Number of agents per group	N	3
Initial pool of variants		V	3	8,16,32
Number of competing variants		n	3	8,16,32
Number or rounds		$R = N - 1$	3	7,15,31
Round		R		
Initial round		r_0		
Number of games per round		$N/2$	3	4,8,16
Number of games per run		N	3	28,120,496
		$(N - 1)/2$		
Connectivity dynamic		$Early (E)$ $Mid (M)$ $Late (L)$	3	E, M, L

contains one cultural variant for each agent ($\sigma_{i0}, \sigma_{j0}, \dots, \sigma_{n0}$), where σ_{i0} is the variant initially assigned to agent i , σ_{j0} the variant initially assigned to agent j , and so on.

At the beginning of each round r , agents are paired using one of the three connectivity dynamics (*early*, *mid*, and *late*) represented below (see *population connectivity dynamics*),

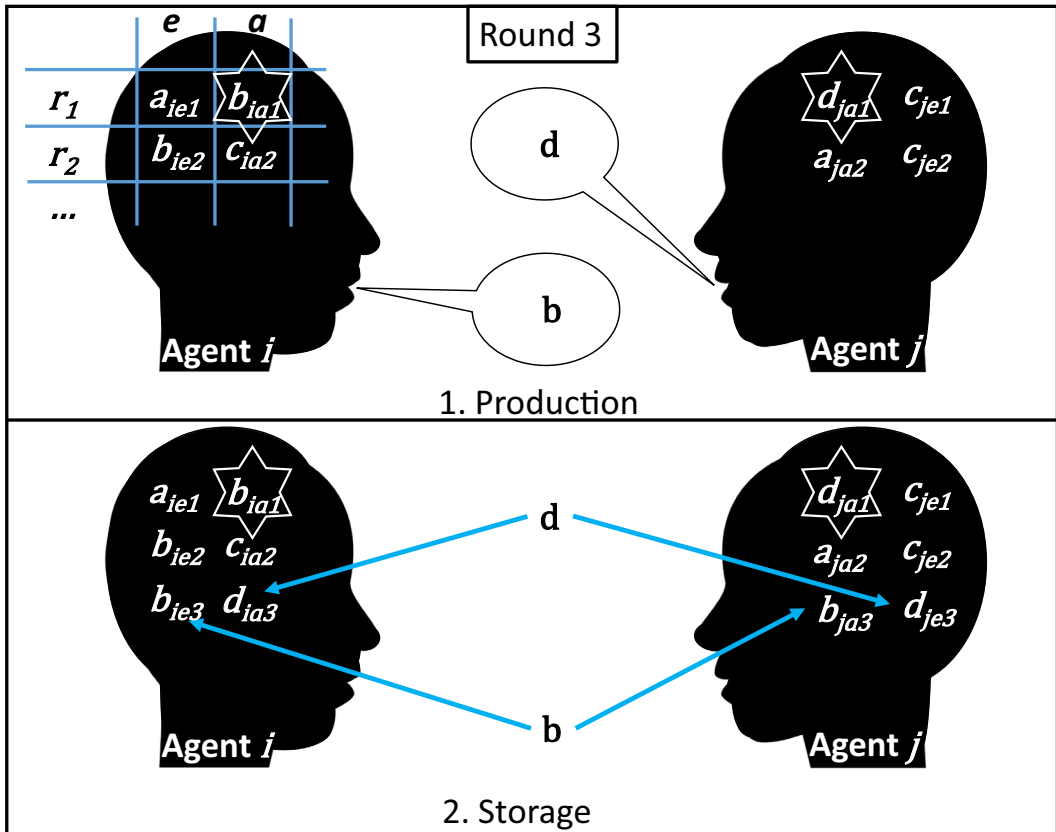


Fig. 1. Illustration of the interaction between a pair of agents at one round of the simulation. At round 3, agents have already some variant tokens in memory, which they have stored in previous rounds. For each variant in memory, the letter represents the type (e.g., a, b). The first index represents the agent’s identity; the second index represents whether the variant token is stored as allocentric (variant produced by other agent) or egocentric (produced by self). (Note that if an agent produces a token of a variant type that in the past was stored as allocentric, the token is stored as egocentric in the present round.) The third index indicates the round at which the variant was stored. During the round depicted (round 3), first, one variant token in each agent’s memory is selected for production (in the figure, the tokens surrounded by a star). Once both agents have produced a new token of the selected variant, they proceed to store them in their memory. Their own variant is marked as egocentric and as produced at round 3; the other agent’s variant is marked as allocentric and also as produced at round 3.

each of which describes a different schedule to pairup agents. Each connectivity dynamic examined differs from the others with respect to how many agents can potentially share the same cultural variant at one given round.

Once agents are paired, at each round r , they interact by presenting and observing one cultural variant. Within each pair (Fig. 1), each agent, in turn, samples its memory to produce a variant (Fig. 1, top) according to the probabilistic function defined in Equation 2. At this point, there is a small probability of innovation, in which case the variant

(A)

Round	Early				Round	Mid				Round	Late			
1	1&2	3&4	5&6	7&8	1	1&2	3&4	5&6	7&8	1	1&2	3&4	5&6	7&8
2	1&4	3&2	5&8	7&6	2	1&4	2&7	3&6	5&8	2	1&4	3&2	5&8	6&7
3	1&6	3&8	5&2	7&4	3	1&6	4&7	2&5	3&8	3	1&3	2&4	5&7	6&8
4	1&8	3&6	5&4	7&2	4	1&5	3&7	2&6	4&8	4	1&5	2&6	3&7	4&8
5	1&3	2&4	5&7	6&8	5	1&7	5&3	2&8	6&4	5	1&6	3&8	5&2	7&4
6	1&5	2&6	3&7	4&8	6	1&8	3&2	7&6	5&4	6	1&7	2&8	3&5	4&6
7	1&7	2&8	3&5	4&6	7	1&3	5&7	2&4	6&8	7	1&8	3&6	5&4	7&2

(B)




Fig. 2. (A) Examples of agent pairings for each type of connectivity dynamic, in a micro-society of 8 agents. In red, agents who may, in the current round, potentially have in their memory the variant produced by agent 1 in round 1. (B) Three different network connectivity dynamics can be described within a pairwise interaction account for a dynamic fully connected network of eight agents. By round 3, potentially eight (in the early connectivity dynamic), six (in the mid connectivity dynamic), or four (in the late connectivity dynamic) agents share the red variant.

produced will be randomly sampled from one of the n variants the population was initialized with. Then, both agents add both variants to their memories (Fig. 1, bottom). That is to say, at round r , when agent i and agent j interact, i produces variant σ_i and j variant σ_j . Both agents store the two variants in memory. Agent i stores variant σ_i in its egocentric memory as σ_{ier} and σ_j in its allocentric memory as σ_{iar} . Agent j stores variant σ_j in its egocentric memory as σ_{jer} and σ_i in its allocentric memory as σ_{jar} .

2.4. State variables and probability distribution of variants

The model takes several parameters as described below:

1. Number of agents (N): we simulate micro-populations of $N = 8$, $N = 16$, and $N = 32$ agents.

2. Number of rounds (R): This depends on the number of agents. A run of the simulation includes $N - 1$ rounds, which allows every agent to be paired with every other agent for exactly one round. A single round is designated as r .
3. Population connectivity dynamics. This reflects the order in which the agents are paired with each other. Different orderings yield different levels of sub-population isolation at different times (Fig. 2). For example, in our eight-agent micro-societies, three different connectivity dynamics can be described with reference to how many agents could potentially share the same variant by round 3: In the early connectivity condition, all eight agents could share the same variant by round 3; in the mid connectivity condition, six agents could share same variant by round 3; in the late connectivity condition, four agents could share the same variant by round 3. We examine three connectivity dynamics: early (E), mid (M), late (L), in 8-, 16-, and 32-agent micro-societies. Due to the wide range of possible permutations for mid connectivity in the larger micro-societies, in this case, for each run, we use a random selection among possible permutations.
4. Coordination bias (c) captures the extent to which an agent has a preference for self-produced variants or partner-produced variants. It fixes the likelihood of a variant being produced depending on whether it originates in others or the agent itself. It takes values from 0 (fully egocentric: preferring self-produced variants over other-produced variants) to 1 (fully allocentric: preferring other-produced variants over self-produced variants). The strength of the coordination model is determined by the combined mathematical complement of the content bias and the innovation rate. That is, for each agent and round, variants are selected for production using the probability of the complement of each of the other cognitive features. When coordination bias is 0.5, we have a neutral coordination model where variants in the egocentric and the allocentric memory are equally likely to be produced. Coordination bias values from 0 to 1 in steps of 0.1 are examined.
5. Content bias (β) and variant quality (s) are two closely related parameters that determine variants selection due to their intrinsic value. Parameter s corresponds to the intrinsic value of each communicative variant, and it indicates to what extent the variant is preferred over the other variants (s is 1 if the signal is preferred over the others, 0 otherwise). Variant quality is a binary trait and the population is initialized with one preferred variant, which is determined at random at the initial state of each simulation. This is motivated by the types of scenarios we aim to simulate, in which, in general, there is a tendency to converge on one or few shared variants. We assume all the agents have the same initial preference. However, this assumption does not mean that agents always choose high-quality variants, since variant choice is defined by a probability distribution function which, in addition to content biases, is determined by other cognitive parameters and the frequency of variants in agents' history (see Equation 2). In the present study, we assume that the preferred variant is adaptive; that is, it has (by definition) higher probability of being produced. In real life, adaptiveness may be due to a variant being easy to produce, to memorize, or to process, attractive, effective, or efficient for a desirable

function, etc. For adaptiveness measures, the frequency of high-quality variants ($s = 1$) is considered (see Section 2.5). Parameter β captures agents' sensitivity to variant quality (s) and ranges from 0—not sensitive at all—to 1—fully sensitive—in steps of 0.1. Parameter β is operationalized only if the target variant σ_i has been produced or seen at least once—in other words, one cannot prefer to re-produce a variant due to its quality until one has been exposed to that particular representation. Thus, given a variant σ_i , the product $\beta \times s_i$ determines the probability that variant σ_i is selected due to its intrinsic value, and this product ranges from 0 to 1, in steps of 0.1. Note that, at each round, as explained in the model dynamics, each agent stores two variants in its memory (the one that was produced by the agent itself and the one that was produced by its partner). Thus, content bias does not affect the storage process in our model (agents store in memory all the variants they encounter); it only affects agents' production. In simple words, the higher the value of β , the higher the probability of producing a cultural variant due to its quality. When content bias is 0, we have a content drift model. Content bias values from 0 to 1 in steps of 0.1 are examined.

6. Memory size. We manipulate memory size (m) by limiting agents' access to their memory store. It corresponds to the maximum amount of history that can affect the variant choice. At each round, when an agent has to produce a variant, it can only choose between the variants that were stored in the preceding rounds. A memory size of m means that the model remembers at most the last m egocentric entries $h|E,m$ and the last m allocentric entries $h|A,m$ from any history h . The relative frequencies of variants in $h|E,m$ define the egocentric initial distribution $f(h|E,m)$ and in $h|A,m$ the allocentric distribution $f(h|A,m)$.
7. Innovation rate (μ). Agents can generate novel variants. We fix the innovation probability at 0.02 using a flat distribution $\phi(x)$ weighted by innovation rate, which means that 98% of variant choices reflect the probability distribution yielded by all the parameters described above, while 2% is a random choice among all eight initial variants. We have chosen this value for two main reasons. First, for consistency with the innovation rate found in Fay et al. (2010) and in Tamariz et al. (2014). Second, because in the present study we attempt to find the best model fit associated with that same dataset.

For each round in the simulation, for each agent, the model yielded a probability distribution of variants (x) for a given history (h) of previous rounds, according to the following equation. We use the prime ($'$) to denote the probabilistic complement: $a' = 1 - a$.

$$\Pr(x|h) = \mu\beta c' f(x|h|_{E,3}) + \mu\beta t f(x|h|_{A,3}) + \mu\beta s + \mu f(x). \quad (2)$$

We run the simulation with 1,452 different parameter value combinations. For each parameter combination, we ran the simulation 1,000 times. For each model run, we assume that all agents have the same connectivity dynamic, biases, and memory sizes. The results below show the average and standard deviations of the number of runs of

each parameter combination examined. All parameters and state variables can be found in Table 1.

2.5. Quantifying convergence and adaptiveness

At its most general level, evolution is defined as a change in the frequencies of different variants in a population over time. We are therefore interested in the composition of the pool of 8, 16, or 32 variant tokens produced by the agents at each round, and how it changes over rounds. We examine the level of convergence in this pool, or the extent to which agents used the same variant. Following others (e.g., Fehér, Wonnacott, & Smith, 2016; Smith & Wonnacott, 2010), we quantify convergence using the information-theoretic notion of entropy (H) (Shannon, 1948):

$$H(V) = - \sum_{v_i \in V} p(v_i) \log_2 p(v_i), \quad (3)$$

where V corresponds to the set of variants, and $p(v_i)$ is the probability of i th variant in that set. Entropy is a well-established alpha diversity index that has been used to measure cultural diversity. High entropy corresponds to low convergence.

Evolution, even by drift, may increase convergence (and decrease entropy), as random sampling at each round gradually eliminates variants from the pool (and our low level of innovation is not enough to compensate for that). For example, at round 0, where each agent produced its own unique variant, the probability distribution over the eight variants was flat (each variant had a probability of 1/8) and the entropy was maximal ($H = 3$ bits). Over time, as agents converged, entropy would decrease; if, say, by round 7, the probability of 1 variant was 0.75, the probability of 2 variant was 0.25 and the probability of the remaining variants was 0, the entropy would be 0.811 bits.

To better understand the mechanics of evolutionary algorithms, a number of studies have investigated time to convergence (TC). The number of rounds until convergence has been used to analyze convergence properties of genetic algorithms in studies about population sizing, network structures, and theory of convergence (e.g., Mueller-Frank, 2013; Olfati-Saber & Murray, 2004; Pelikan, Goldberg, & Cantú-Paz, 2000). This additional measure is important because it allows us to predict more accurately the moment at which one population will reach convergence under different conditions. Additionally, it gives us more information about how relevant agents' choices were in the first rounds. Therefore, in some analyses, we will also present time to convergence (TC) or the number of rounds it takes for the population to reach full convergence (defined as $H = 0$ bit) for the first time.

Researchers on cultural evolution have developed models that link demography and cultural adaptiveness, using a variety of mathematical approaches (Henrich, 2004; Mesoudi, 2011; Shennan & Wilkinson, 2001). These models describe how a trait changes in frequency over time. We calculated the adaptiveness (A) of the cultural system at each round (t) as the frequency of high-quality variants (quality is measured by s , see above)

at that round,

$$A(t) = \frac{n_h(t)}{N(t)}, \quad (4)$$

where $n_h(t)$ is the number of high-quality variants at round t and $N(t)$ is the total set of variants produced in round t .

Two considerations apply to our study when we calculate average change in high-quality variants. First, we use relative fitness equations to account for the adaptiveness of high-quality variants; that is to say, we only consider the frequency of cultural variants having $s = 1$. Second, we assume that cultural variants are distributed, and therefore transmitted at each round, according to our parametrized model, as defined in Equation 2. Thus, in our case, the change in the adaptiveness (ΔA) of high-quality variants due to selection follows immediately from our definition of adaptiveness. This tractability assumption simplifies our equation considerably, because now the change in adaptiveness equation reduces to:

$$\Delta A = A(t + 1) - A(t), \quad (5)$$

where change in adaptiveness (ΔA) takes the difference between the adaptiveness in the subsequent round $A(t + 1)$ and the adaptiveness in the earlier round $A(t)$. Therefore, a change in adaptiveness above 0 ($\Delta A > 0$) indicates that the fitness of high-quality variants produced by agents increased from one round to the next. When $\Delta A = 0$, variant frequency was stable from round to round.

3. Results

In this section, we offer a summary of the results of two selection models (content bias and coordination bias) against a drift model, and how they interact with each other (Section 3.1). Next, we show the effects of memory limitations (Section 3.2). Figures for these analyses can be found at the end of each section. Additional analyses on conditional entropy distributions can be found in the Supplementary Material. In Section 3.3, we focus on the effects of population connectivity dynamics on entropy, time to convergence, and change in adaptiveness of the cultural system from round to round. We also pay special attention to the interplay between connectivity and two strong drivers of convergence: content bias and population size. We use statistical inference to fit models to our data (see Supplementary Material). However, following White, Rassweiler, Samhoury, Stier, and White (2014), we use frequentist statistical models only to calculate effect sizes in our multifactorial simulations, but we do not report p -values, which can be meaningless when applied to simulation model output. In the following analyses, we show mean values and standard deviations (mean \pm SD).

3.1. Cognitive biases

We ran simulations manipulating the level of content bias (β). When compared with a neutral content bias model ($\beta = 0$), content bias increased convergence (decreasing entropy). Mean entropy was greatest when $\beta = 0$ (2.451 ± 0.448 bits) and lowest when $\beta = 1$ (1.020 ± 1.136 bits). Similarly, when keeping coordination bias at a neutral level (coordination bias = 0.5), mean entropy was greatest when $\beta = 0$ (2.277 ± 0.478 bits) and lowest when $\beta = 1$ (1.015 ± 1.137 bits).

When we considered coordination bias (c) alone ($\beta = 0$), egocentric bias reduced convergence. Mean entropy was greatest when $c = 0$ (strongest egocentric bias) (2.013 ± 1.060 bits) and lowest when $c = 0.5$ (neutral coordination) (1.768 ± 0.956 bits). Similarly, when keeping a neutral content bias ($\beta = 0$), mean entropy was greatest when coordination bias was fully egocentric, 0 (2.906 ± 0.171 bits) and lowest when coordination bias was neutral, 0.5 (2.277 ± 0.478 bits).

Both β and c had effects on convergence that differ from a drift model; see Fig. 3. The effect of c on entropy was different for each level of β , revealing an interaction. Average entropy was highest when $\beta = 0$ and c was also 0 (weakest content bias and strongest egocentric bias) (2.906 ± 0.171 bits) and lowest when $\beta = 1$ and $c = 0.5$ (1.015 ± 1.136 bits). When agents' behavior was strongly content-biased, the rate of convergence increased, masking the effect of c . Conversely, weaker content biases allowed c to show its effect on convergence, which can be characterized by a distinctive asymmetric distribution. The slowing effect of c on the rate of convergence becomes hidden as β rises.

Our main result here, which establishes a baseline for posterior analyses, is that the population reaches more convergence when agents are more sensitive to the intrinsic value of variants. It is obvious that a content bias, as implemented in our model, dramatically affects the spread of variants in a population and that this result is consistent with previous studies showing the crucial effects of the individuals' direct assessment of the value of traits on adoption and transmission (Fay et al., 2010; Henrich & McElreath, 2007; Stubbersfield, Tehrani, & Flynn, 2015; Vale et al., 2017); for a review, see Kendal et al. (2018). It is also interesting to note how the effect of β is modulated by c , which plays an important role in maintaining variation. There are two ways in which c acts against convergence:

1. A strong egocentric bias preserves sign variation at the level of the population by inhibiting the adoption of others-produced variants (i.e., fully egocentric agents keep producing their original egocentric variant indefinitely, unless innovation occurs). This increases diversity globally because each agent tends to overproduce its own individual variant, reducing the probability of local convergence.
2. A strong allocentric bias preserves sign variation at the level of the population by inhibiting the production of variants stored as egocentric (i.e., fully allocentric agents keep producing other-produced variants). This increases diversity globally because each agent tends to overproduce variants that are stored in its allocentric memories.

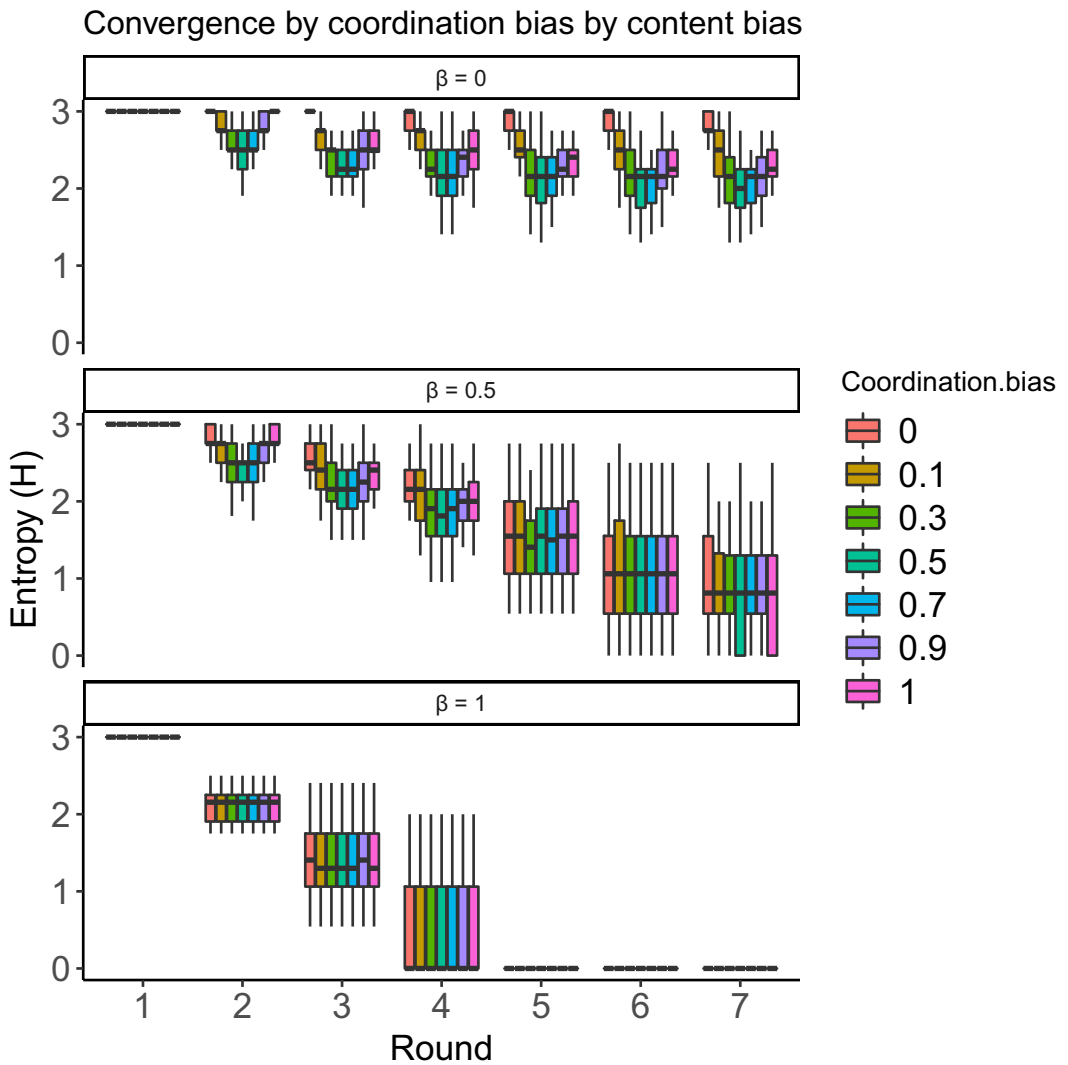


Fig. 3. Convergence (measured as Shannon entropy H) by round by each combination of biases. Examples for content bias 0, 0.5, and 1. A drift model has a content bias of 0 and a coordination bias of 0.5. X-axis represents rounds from 1 to 7, and Y-axis represents entropy in bits. In this and subsequent boxplots: middle line is median, 50% quantile; lower hinge, 25% quantile; upper hinge, 75% quantile; lower whisker is smallest observation greater than or equal to lower hinge $- 1 * IQR$; upper whisker is largest observation less than or equal to upper hinge $+ 1 * IQR$.

Here, it is important to note that an intermediate level of coordination bias facilitates convergence by allowing agents to align more easily on common shared representations. Too allocentric or egocentric scenarios make these conventions unlikely. However, fully egocentric agents are more efficient than allocentric agents at preserving variation. This is because the allocentric memory of allocentric agents contains more shared variants with

other agents (there are seven other agents) than the egocentric memory of egocentric agents (there is only one ego). Consequently, the probability that allocentric agents converge is higher than the probability that egocentric agents converge, which explains the asymmetry of the coordination model (Fig. 3).

However, as β increases, it overrides the capacity of egocentric and allocentric biases to preserve diversity. This effect is important as it highlights both the importance of coordination when the adoption of variants is not based on competition for variant quality, and its irrelevance in scenarios that encourage variant adoption based on the intrinsic properties of the variants. Our model, therefore, assigns different weights and roles to each selection pressure, with β as the main driver of convergence—encouraging selection owing to intrinsic variant quality, and c as a modulating pressure—favoring variation when β is weaker. Given these baseline dynamics, in the next sections, we will analyze how model outcomes are modulated by agents' memory size and network connectivity dynamics.

3.2. Memory

Memory size increased entropy and therefore decreased convergence. Average entropy was greatest in the absence of memory limitations, when agents kept in memory all the variants they had been exposed to (1.935 ± 0.920 bits). In contrast, when we limited the agents' memory to the most recent five rounds (1.920 ± 0.930 bits), three rounds (1.834 ± 0.990 bits), or one round (1.675 ± 1.081 bits), entropy decreased noticeably.

This effect of memory on entropy was more noticeable for intermediate values of content bias ($\beta = 0.5$). When content bias was strongest, memory effects tended to be masked by a floor effect at the lower end of the entropy distribution in the last rounds. On the other hand, when content bias was neutral ($\beta = 0$), memory effects were greater for neutral values of coordination bias ($c = 0.5$), when compared with strongly egocentric or allocentric bias (Fig. 4).

In our analysis of how memory interacted with cognitive biases and influenced convergence, our initial expectation was that, in general, memory limitations would reduce variation. These predictions were based on studies showing that memory limitations can lead to regularization in concurrent frequency learning tasks, which is possibly attributable to limited working memory (Ferdinand et al., 2013; Hudson Kam & Chang, 2009; Hudson Kam & Newport, 2005; Kareev, Lieberman, & Lev, 1997). Our model is consistent with those findings, suggesting that a larger memory size decreases convergence by delaying social learning and, in turn, the spread of variants with high intrinsic value. Interestingly, however, our results show that a reduction in variation seems to be true in all scenarios except when the population is fully allocentric in a null content-biased model. This pattern of inversion suggests that allocentric agents with larger memory sizes might increase the probability of convergence on shared variants in scenarios where variant adoption is not encouraged on the basis of intrinsic variant quality.

As with coordination bias, the effect of memory fades when β rises to levels close to 1. This is because when the population is fully biased toward the current value of a given

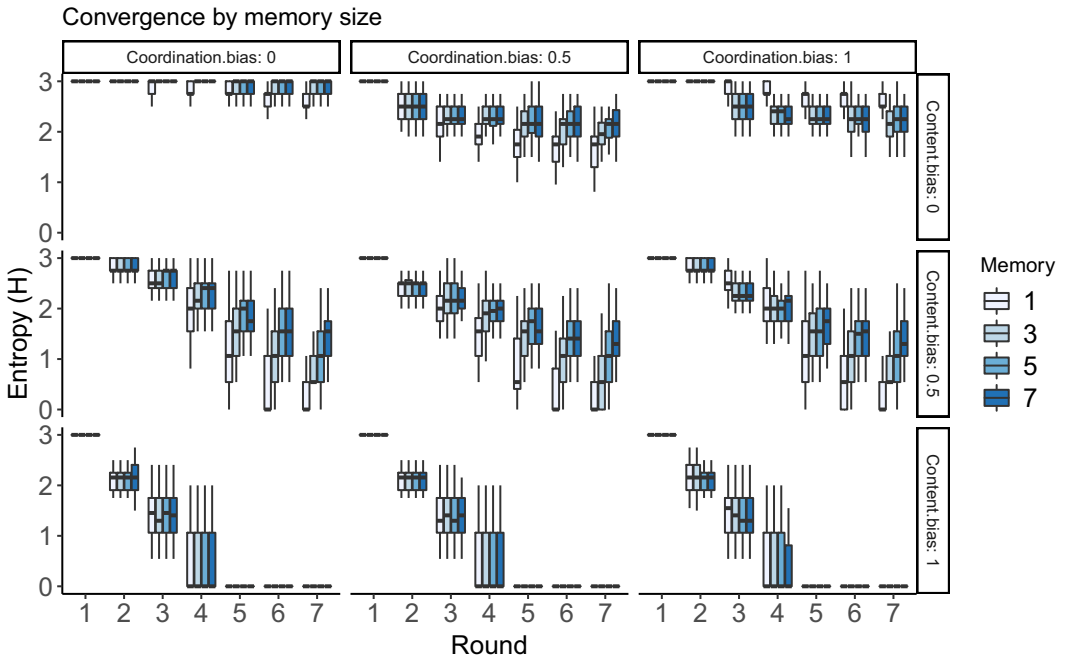


Fig. 4. Convergence (measured as Shannon entropy H) by round by memory. Examples for content bias 0, 0.5, and 1 and for coordination bias 0, 0.5, and 1. The x -axis represents rounds from 1 to 7, and the y -axis represents entropy in bits. In this and subsequent boxplots: middle line is median, 50% quantile; lower hinge, 25% quantile; upper hinge, 75% quantile; lower whisker is smallest observation greater than or equal to lower hinge $- 1*IQR$; upper whisker is largest observation less than or equal to upper hinge $- 1*IQR$.

variant, memory size became less important as a selection mechanism for variants adoption. In contrast, as β decreases, agents begin to increasingly activate coordination as a selection mechanism, by paying more attention to whether variants are stored in egocentric and allocentric memories. The effect of memory size, thus, becomes crucial to determine which variant is chosen by each agent at each time step since it determines the size of the pool of variants that each agent can track back in time in its memory from a given moment. Since variation in convergent processes is in general greater in larger pools, a reduction in memory size yields higher levels of convergence in our model, suggesting an increased reliance in social learning, which helps variants with high intrinsic value to propagate.

As we have shown in the previous section, extreme coordination biases preserve variation by inhibiting the adoption of variants produced by others or by the agent itself. This explains why convergence is lower when $\beta = 0$, regardless of the level of memory implemented. However, the effect of memory is more noticeable for intermediate values of β . This is simply because in this intermediate scenario in which both β and c coexist, β is high enough to substantially increase convergence (yielding a large range of possible mid-convergence outcomes) but not strong enough to eliminate the effect of memory.

3.3. Connectivity dynamic of the population

A key finding in this study is that population connectivity dynamic affected the spread of variants. Convergence was delayed in populations that took longer to reach full connectivity. But this sociostructural effect was only manifested under certain conditions related to the cognition of individual agents. When running simulations using a null content bias model, entropy remained similar for all levels of connectivity (Fig. 5). Interestingly, however, increasing content bias in the agents revealed a substantial effect of the connectivity dynamic on convergence. Mean entropy differences between conditions were greatest at round 3: 0.805 ± 0.730 bits under late connectivity, 0.464 ± 0.498 bits under mid connectivity, and 0.133 ± 0.286 bits under early connectivity. Contrary to what happens with memory and coordination bias, the modulating effect of the network connectivity dynamic was stronger for high–moderate content bias. This is because the spread of variants with high intrinsic value is facilitated by the assessment of variants value based on the nature of the information itself but restricted by the conditions of accessibility to that same information imposed by the network connectivity dynamic. Thus, connectivity dynamics are almost irrelevant in scenarios with null or very low content bias, while they yield substantial effects on the pathways toward convergence for moderate and high content bias, with late connectivity delaying convergence and maintaining variation in the population for a longer time period. The effect of connectivity dynamic was more pronounced in larger populations, where we can observe that the delay due to late connectivity dynamic lasts for a larger number of rounds when we increase population size (Fig. 5). In content-biased populations, as population size increases, mean entropy differences between conditions remained significantly high for longer time periods (e.g., when $\beta = 0.8$, the relative difference between conditions remained above 10% for more than six rounds in 16-agent micro-societies, and for more than 11 rounds in 32-agent micro-societies).

Summarizing, both content bias (by strengthening the selection of high-quality variants) and population size (by lengthening the time to convergence) amplified the effect of the connectivity dynamic, and this, in turn, resulted in a deep alteration of the evolutionary trajectory of convergence. In these scenarios, late connectivity populations clearly show periods of rapid convergence followed by periods of relatively slower convergence, resembling punctuational evolutionary dynamics. In contrast, convergence in early connectivity populations was not affected by these evolutionary bursts and tended to be shaped by a monotonic sigmoid curve.

Although the effects of network connectivity dynamics examined here tend to vanish in the long run, it is important to characterize how fast a group of agents reach the state of convergence. Time to convergence (TC) is a widely used performance indicator in processes that require convergence as an outcome. Differential temporal pathways toward convergence between conditions might have dramatic effects on social learning scenarios in which outcomes are

1. either time-dependent or based on intermediate states (e.g., strategic or argumentative scenarios in which the first consensus is the one that is considered) or

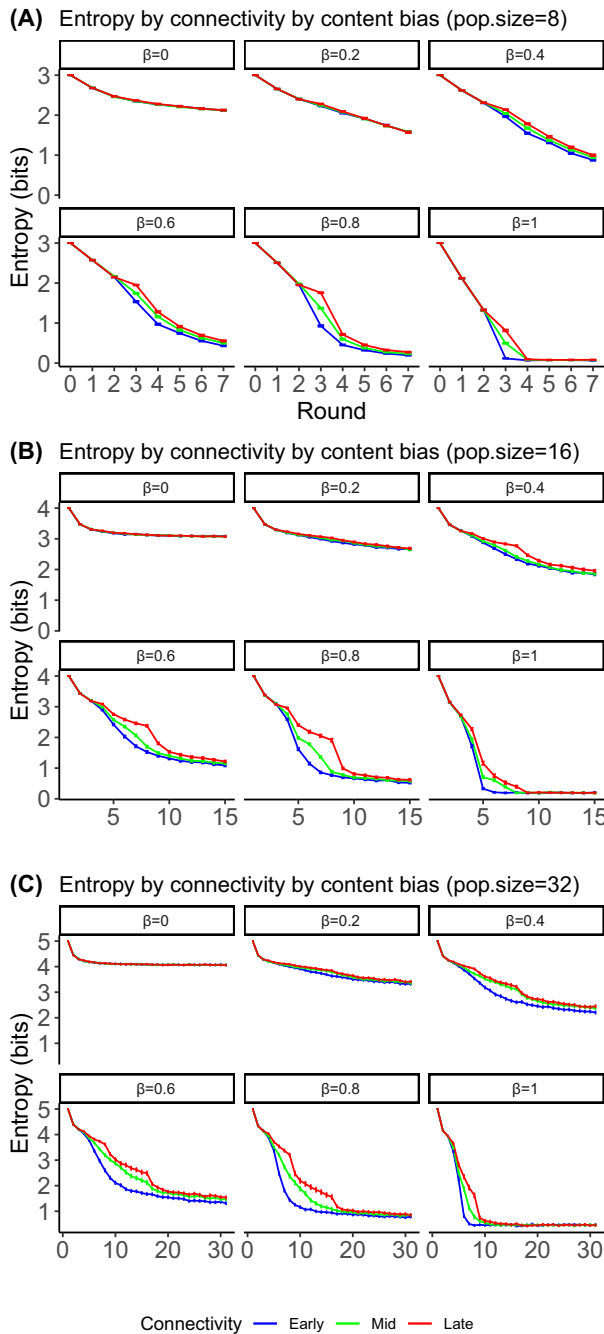


Fig. 5. Convergence (measured as Shannon entropy H) averaged over each level of connectivity and content bias, for population size = 8 (A), population size = 16 (B), and population size = 32 (C). The x -axis represents rounds, and the y -axis represents entropy in bits. Drift models are shown in the top-left ($\beta = 0$). We ran 1,000 simulations for each parameter combination. Error bars indicate 95% CIs.

2. affected by outcome primacy (e.g., the jurors' decision after a sequence of argumentative speeches where earlier data have more impact on behavior than later data).


Although these scenarios can vary widely, we hypothesized that *TC* can inform how connectivity dynamics may alter the outcome of such social processes. We examined the effect of connectivity on the *TC* in those simulations where full convergence ($H = 0$) was reached. Fig. 6 depicts how fast micro-societies converged on a shared cultural convention in those simulations. The time required for a population to reach full convergence was longer in the late connectivity condition when compared to the mid and early connectivity conditions. This shows that the effects of the network connectivity dynamics examined here are strong enough to alter convergence performance in scenarios in which a common consensus is reached in the population. In other words, the outcome of social learning processes that are strongly restricted by time is substantially affected by these changes in the network connectivity dynamic. This is particularly true for moderate–high levels of β . However, coordination bias and memory do not seem to interact with the connectivity dynamic when it comes to explaining time to convergence (Supplementary Material, Figs. S.12 and S.13). This is because the effect of these cognitive parameters on convergence is very limited compared to that of content bias, which is the main driver of convergence.

Similarly, the adaptiveness (*A*) of the cultural system increased more rapidly in populations with early connectivity. The change in adaptiveness of high-quality variants remained above 0 across rounds, indicating that the proportion of high-quality variants always increased from round to round. However, changes in adaptiveness followed different patterns in populations with early, mid, and late connectivity (Fig. 7). Populations with late connectivity evolved in punctuated bursts of change followed by periods of slower change. For instance, in eight-agent micro-societies, at least two rapid bursts of change in the proportion of high-quality variants can be observed before the population became a fully connected network in seven rounds. As above, these patterns can be better observed when we increase population size. Bursts of rapid change are related to the evolutionary moments in which the pockets of isolated agents created by the late connectivity become connected. On the other hand, populations with early connectivity dynamics followed a single-peak evolutionary dynamic. This is due to the fact that high-quality variants could spread in the system continuously (and without any additional restriction imposed by the connectivity dynamic) until the population reached its equilibrium.


4. Model comparison against experimental data

We collect simulated data that include the parameter combination of the best fit models for the 64 data structures coded by Tamariz et al. (2014). The coded data structures correspond to the experimental data collected by Fay et al. (2010) from four distinct eight-person communities and 16 concepts used in a Pictionary-like game, yielding a total of


(A) Time to convergence by connectivity by content (pop size=8)



(B) Time to convergence by connectivity by content (pop size=16)



(C) Time to convergence by connectivity by content (pop size=32)



Connectivity ■ Late ■ Mid ■ Early


Fig. 6. Time to convergence—measured as the natural logarithm of the number of rounds it takes for the population to reach full convergence (defined as $H = 0$ bit) for the first time—averaged over each level of connectivity and content bias, for population size = 8 (A), population size = 16 (B), and population size = 32 (C). We only considered micro-societies that reached full convergence. Scenarios like this did not occur during the round-robin simulation when content bias was relatively low and population size was 16 or 32 agents. Late and mid connectivity delay time to convergence when compared to early connectivity. This is particularly true for high–moderate levels of content bias. We ran 1,000 simulations for each parameter combination.

64 data structures. A verification of the model and its explanatory power is provided in Tamariz et al. (2014). The data coded in Tamariz et al. (2014) for each generation are equivalent to the data produced by the simulator just after the interaction of the agents in each round, that is to say, once the agents have in memory the variants produced by their partners. This can be easily verified by observing the relatively high levels of convergence ($2 < \text{entropy} < 3$) in the first generation of the graphical examples of data structures provided in Tamariz et al. (2014). For this reason, the first generation of Tamariz et al.'s data is equivalent to the first round of ours, and so on for subsequent generations. The initial state of our data ($r = 0$) (where all variants are unique) is therefore excluded from the model comparison. Absolute entropy was calculated for both datasets (simulated and experimental) using the information-theoretic notion of entropy (H) (Shannon, 1948). For a detailed explanation of the coding process, see Tamariz et al. (2014).¹


We run 1,000 simulations under three different conditions: early (E), mid (M), and late (L) connectivity, and test them against the experimental data coded (T). Our model with early connectivity exactly mirrors the pair composition and connectivity dynamic used in Fay et al. (2010). Thus, we predict to find the best fit model for the 64 data structures in the early connectivity. According to Tamariz et al. (2014), most data structures (95%) were best accounted for without $\beta = 0$, and using a wide combination of content-biased models. Thus, using this model fit as a reference for our analysis, we also predict that a distribution of entropy yielded by models including a range of content bias parameters will fit the experimental dataset better than models using extreme configurations of biases. Our simulated data structure includes 1,452 different parameter value combinations in each round. In this analysis, we choose to evaluate the predictive power of our model applied to Fay et al. (2010) by testing, against the experimental data, simulations using (A) the best fit parameter combinations in Tamariz et al. (2014) (without $\beta = 0$), (B) all parameter combinations, (C) a content-biased model of $\beta = 1$, and (D) drift model. As a benchmark for illustration, a graphical representation of these model tests against the experimental data (T) can be found in Fig. 7. Table 2 summarizes the model abbreviations.

EA model was associated with better model fit, in particular from round 2 onward, when the three different connectivity dynamics begin to diverge. Mean distance between EA simulated data and experimental data was $M = 0.056$, $SD = 0.14$. Figs. 8 and 9 show that (a) our model reproduces the experimental patterns of convergence and (b) the version of the model with parameter combination EA that most closely matches the assumptions of the experiment is the best fit model of the versions examined.

(A) Change in adaptiveness by connectivity (pop.size=8)



(B) Change in adaptiveness by connectivity (pop.size=16)



(C) Change in adaptiveness by connectivity (pop.size=32)




Fig. 7. Change in adaptiveness ΔA of high-quality variants by round, averaged over each level of connectivity and content bias, for population size = 8 (A), population size = 16 (B), and population size = 32 (C). Results above 0 indicate that the proportion of high-quality variants increased relative to the previous round. When $\Delta A = 0$, variant frequency was stable from round to round.


These results suggest that *EA* model qualitatively reproduces the empirical data pattern for the acceptable range of accuracy under the set of experimental observations. As predicted, the model with early connectivity, which mirrors the connectivity dynamic and pair composition used in Fay et al. (2010), outperforms its alternatives. Additionally, our results are consistent with Tamariz et al. (2014), who found that 95% of data structures were best accounted for within a range of content biases. We show that as long as we add $\beta = 0$ to the models (see model B in Fig. 8), they lose predictive power, because they underestimate the decrease of entropy in the experiment. Similarly, but due to an overestimation of entropy decline, models with only $\beta = 1$ (model C in Fig. 8) are far from the reality shown by Fay et al. (2010). Thus, we think that there is substantial evidence to conclude that content bias (in a wide range of levels) was an important driver of the spread of the culturally transmitted variants in the experiment. Our results are also consistent with a variety of studies about the role of content biases. In particular, two studies (Fay & Ellison, 2013; Fay et al., 2008) have used the same corpus of data collected by Fay et al. (2010) to examine the transmission of the communication systems that we model in this paper. These studies found functional adaptations of the selected variants that are exactly the same that would be predicted if a critical level of content bias were operating on the communication system (Tamariz et al., 2014). We know that content biases have a strong impact on the adoption of variants. Crucially, for the purpose of the present study, connectivity dynamics, amplified by content biases, may have important consequences for the evolution of cultural variants in populations. Failing to take into account the role of connectivity dynamics in experiments and in real life may preclude a full understanding of the data observed.

5. Discussion

Using a computational approach, the present study extends formal and experimental findings about learning in social networks by simulating pairwise interactive micro-

Table 2
Model abbreviations

	Simulated best fit (without $\beta = 0$)	Simulated (all parameter combinations)	Simulated ($\beta = 1$)	Simulated (Drift model $\beta = 1$)	Experimental data
Early connectivity	EA	EB	EC	ED	T
Mid connectivity	MA	MB	MC	MD	—
Late connectivity	LA	LB	LC	LD	—



*Exp. Data source:
Tamariz et al. 2014*

Fig. 8. Comparison between experimental and simulated data. Simulated data from 16 models (four parameter combinations with three connectivity dynamics). Entropy (H) (y-axis) is plotted against time in rounds (x-axis). Simulated data using early connectivity (E) (blue), mid connectivity (M) (green), and late connectivity (L) (red) is tested against the experimental data collected by Fay et al. (2010) and coded by Tamariz et al. (2014) (pink). The line plots here show the evolution of entropy, with increasing convergence, over rounds. The simulated early connectivity (EA) model was associated with better model fit, in particular from round 2 onward, when the three different connectivity dynamics begin to diverge. Models with $\beta = 0$ (B) lose predictive power when compared to A because they underestimate the decrease in entropy in the experimental data. Models with only $\beta = 1$ (C) overestimate the decrease in entropy, while drift-only models (D) underestimate it. Lines indicate mean entropy, and ribbons indicate 95% CIs.

societies based on round-robin tournaments, where individual cognitive biases, memory constraints, and population connectivity dynamics are systematically manipulated. Our results extend previous studies on social learning showing that content biases are




Fig. 9. (A), Mean distance between simulated data and experimental data. Mean distance between simulated early connectivity (EA, in blue) and experimental data (T) was $M = 0.056$, $SD = 0.14$ and it was associated with better model fit. (B) Entropy by model averaged over each connectivity condition. (C, D, E, and F), Mean entropy by model by round. Error bars indicate 95% CIs. EA model, which mirrors the connectivity dynamic and pair composition used in Fay et al. (2010), outperforms its alternatives.

important drivers of convergence. We also show that connectivity dynamics affect the time-course of the spread of variants in moderate–high content-biased populations: When populations take longer to reach full connectivity, convergence onto a single cultural variant is slowed. Content bias accelerates convergence and amplifies the effects of connectivity dynamics. Larger memory size and coordination bias, especially egocentric bias, are also shown to slow down convergence, especially in moderate–low content-biased populations. Finally, connectivity dynamics are shown to affect the frequency of high-quality variants (adaptiveness), with late connectivity populations showing bursts of rapid change in adaptiveness followed by periods of relatively slower change, and early connectivity populations following a single-peak evolutionary dynamic.

While content bias is the main predictor of convergence, in some circumstances its effect can be modulated by the population connectivity dynamic (Figs. 5 and 6). This means that adding connectivity dynamics may improve the predictive power of models

based on cognitive biases and social networks, especially in cases of strong content biases. Population convergence on shared cultural conventions is driven by the agents' content bias, and the time required to reach a certain degree of convergence (or time to convergence) can be deeply affected by the specific order of interactions between agents, that is, by the population connectivity dynamic: In general, the less connectivity the more time is needed to converge. Furthermore, the effects of these different dynamics in the order of interactions of the agents can be observed even if we maintain the same network topology, in our case a fully connected network.

It is important to note that even though previous work in the field has frequently used fully connected networks (Fay et al., 2010; Komarova, Niyogi, & Nowak, 2001; Tamariz et al., 2014), this type of network topology is unrealistic because it restricts the interaction between agents to a particular pattern of interconnectedness, reducing the complexity of the system. Therefore, population connectivity dynamics might play a different role in other networks, such as scale-free networks (Barabási, Albert, & Jeong, 1999; Barabási et al., 2002). Furthermore, it is well established that people do not contribute equally to group discussions, leading to different degrees of network connectivity (Fay, Garrod, & Carletta, 2000; Stasser & Taylor, 1991). Both factors might potentially motivate future extensions of our current investigation.

Our results also identify a general tendency for adaptiveness to change over time and for cultural variants to converge on high-quality variants, in such a way that it is possible to identify causal links between connectivity dynamics and evolutionary trajectories. In this way, in late connectivity dynamics, several punctuational bursts occur in the course of a complete cycle of interactions between agents. In contrast, early connectivity dynamics follow a single-peak evolutionary trajectory. These computational results extend a number of studies that, under a variety of assumptions, have proposed punctuational or rapid bursts of change as a feature of cultural and language evolution (Atkinson, Meade, Venditti, Greenhill, & Pagel, 2008; Dixon & Ward, 1997; Fitch, 2008; Janda & Joseph, 2003; Sabherwal, Hirschheim, & Goles, 2001). Punctuational changes in our model may provide insight into processes underlying the human ability to adapt quickly to cultural variants introduced by new agents (e.g., due to migration), showing that these changes can be induced merely by manipulating the order of interactions in a population.

In high content-biased populations, the effect of the connectivity dynamic is amplified (Figs. 5 and 6), while coordination bias and memory size effects are masked (Figs. 3 and 4). Interestingly, in low content-biased populations, the effect of these parameters became visible: When compared with a drift model, egocentric and allocentric biases both reduce convergence (Fig. 3). This is because cultural diversity can more easily be maintained over time in the presence of behaviors that maximize the occurrence of either self-produced signals—in the case of egocentric agents—or partner-produced signals—in the case of allocentric agents. The effect of egocentric bias is stronger than that of allocentric bias. This is due to the fact that fully egocentric agents stick to their own variant, which is always the same in the egocentric memory, unless there is an innovation. At the population level, this means that each agent has a different variant, returning maximum entropy—which can only decrease through mutation. Fully allocentric agents, in contrast, always

adopt variants produced by others. High variation is thus maintained, but to a lesser extent than in the egocentric case because allocentric agents choose variants from among all the variants stored in their allocentric memories—variants produced by current or previous partners. This sometimes leads to more than two agents converging on the adoption of a variant, and therefore reducing entropy.

Memory also shows its effect more markedly when content bias is low or intermediate (Fig. 4). The longer the memory span, the more variation is maintained, as variants from earlier rounds, that might not appear at one round, are kept in memory and may reappear. Our study is consistent with previous literature showing that memory limitations lead to a reduction in variation (Ferdinand et al., 2013; Tamariz & Kirby, 2015). Nevertheless, we show that this reduction could be masked in high content-biased populations, when agents have a strong preference for signals with high intrinsic value.

Our results also agree with recent studies showing that population structure and population interaction can be strong predictors of cultural evolution (Creanza et al., 2017; Derex & Boyd, 2016; Derex et al., 2018). In addition, our model shows that cognitive biases and population connectivity dynamics may interact in important ways. When content-biased populations evolve in high isolation (late connectivity dynamic), convergence is slower than when they evolve in moderate (mid connectivity dynamic) or low isolation (early connectivity dynamic). This is because more isolated subpopulations cannot benefit from wider cultural exchange during the first rounds, those in which agents are acquiring the basis of their culture and storing it into their respective memories. This suggests that population structure and, in particular, the connectivity dynamics of the population can have important effects on cultural convergence and should be taken into account when it comes to research on the interactions between cognitive biases, network structures, and cultural evolution.

6. Implications

Our results are relevant to social learning scenarios governed by turn-based interventions in which convergence on shared conventions is crucial (e.g., Bloom et al., 1987; Fay et al., 2008, 2010; Manzo, 1996; Shteingart et al., 2013; Stone, 1969). Our computational model can be used to fit real data obtained from turn-based cultural processes and might be helpful to improve the organization of the turn taking by mitigating undesirable effects linked with one particular connectivity.

In light of evolutionary theory applied to cultural evolution (Atkinson et al., 2008; Fitch, 2008), our results also suggest that, in some scenarios, bursts of change in the cultural system may be partly explained by the order in which individuals interact over time. This demonstrates, for the first time, a direct connection between convergence, adaptiveness, and population connectivity dynamics for a fixed range of combinations of individual cognitive biases.

Our findings are consistent with scientific models and with theoretical and experimental studies of human communication showing that convergence is driven by content biases

(Gong et al., 2008; Tamariz et al., 2014), and also agree with studies on rational learning in social networks showing that the level of convergence is partially determined by the degree of connectivity in the social network (Barkoczi & Galesic, 2016; Centola & Baronchelli, 2015; Centola et al., 2007; Lazer & Friedman, 2007; Mueller-Frank, 2013; Olfati-Saber & Murray, 2004). The specific manipulation of the network connectivity dynamic that we have studied here has not been taken into account in previous experimental work and it would be especially relevant to those researchers that use dynamic interactive micro-societies of agents switching partners over time (e.g., communicative games, cooperative games, or tournaments; Baum et al., 2004; Byun et al., 2018; Caldwell & Smith, 2012; Fay et al., 2008, 2010; Mesoudi & Whiten, 2008; Raviv et al., 2019; Tamariz et al., 2014). In most cases, experimental designs of micro-societies of interacting actors only include one pair composition out of all the possible combinations of pair shuffling, and therefore, outcomes are related with only one specific population connectivity dynamic, potentially affecting the accuracy of the generalizations made by these studies. Our results suggest that this type of research would benefit from experimental designs that control the probability of occurrence of each possible connectivity dynamic.

Our agent-based model is a simplification of a specific problem. Each agent is characterized by a combination of biases toward the quality and origin of a set of variants. The network topology is complete and organized in dyads. Thus, following Rubinstein (2006), in the dilemma of responding to reality, we regard our model as a very limited set of assumptions which is inevitably inapplicable to many contexts.

In this study, we have presented an innovative small-scale simulation project and a comparison of the simulations against experimental data collected from existing experiments, which allowed us to run a high-resolution test of the specific parameters examined within a small but highly controlled dataset. Future research on the topic of this paper should go through experimentation with human participants, with the aim of verifying assumptions and conclusions. For example, further experiments could be oriented toward data collection from connectivity dynamics modeled in this paper, in particular those that have not yet been used in experimental work. Our first effort in this experimental line of research using four-agent micro-societies can be found in Segovia-Martín and Tamariz (2020).

Properties of populations can be important predictors of cultural evolution, and our model has shown that convergence can be altered by the connectivity dynamic. This may help improve the experimental design of ongoing research in the field of cultural evolution and better explain the interactions between network topologies, cognitive biases, and cultural transmission.

Acknowledgments

We thank our colleagues from the Centre for Language Evolution (University of Edinburgh) and the Centre de Lingüística Teórica (Universitat Autònoma de Barcelona). We are grateful to Paul Smaldino, Asha Sato, and Sergio Balari for their comments during

the course of this research. This paper has benefitted from a grant of the Spanish Government (FEDER/ Ministerio de Ciencia, Innovación y Universidades – Agencia Estatal de Investigación) (Ref. FFI2017-87699-P).

Conflict of interests

The authors declare no competing interests.

Open Research badges



This article has earned an Open Data badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. The data is available at https://github.com/jsegoviamartin/network_connectivity_dynamics_model.

Data Availability Statement

Electronic supplementary material and simulation code are available at: https://github.com/jsegoviamartin/network_connectivity_dynamics_model.

Note

1. For reproducibility reasons, the experimental entropy dataset used for the present analysis can be found at https://github.com/jsegoviamartin/network_connectivity_dynamics_model.

References

- Atkinson, Q. D., Meade, A., Venditti, C., Greenhill, S. J., & Pagel, M. (2008). Languages evolve in punctuational bursts. *Science*, *319*(5863), 588.
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, *41*(2), 203–226.
- Barabási, A.-L., Albert, R., & Jeong, H. (1999). Mean-field theory for scale-free random networks. *Physica A: Statistical Mechanics and Its Applications*, *272*(1–2), 173–187.
- Barabási, A.-L., Jeong, H., Neda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and Its Applications*, *311*(3–4), 590–614.
- Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature Communications*, *7*, 13109.
- Baum, W. M., Richerson, P. J., Efferson, C. M., & Paciotti, B. M. (2004). Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evolution and Human Behavior*, *25*(5), 305–326.

- Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, 114(26), E5070–E5076.
- Bentley, R. A., Hahn, M. W., & Shennan, S. J. (2004). Random drift and culture change. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(1547), 1443–1450.
- Björk, J., & Magnusson, M. (2009). Where do good innovation ideas come from? Exploring the influence of network connectivity on innovation idea quality. *Journal of Product Innovation Management*, 26(6), 662–670.
- Bloom, K., Russell, A., & Wassenberg, K. (1987). Turn taking affects the quality of infant vocalizations. *Journal of Child Language*, 14(2), 211–227.
- Blythe, R. A. (2012). Neutral evolution: A null model for language dynamics. *Advances in Complex Systems*, 15(03n04), 1150015.
- Bormann, E. G. (1996). Symbolic convergence theory and communication in group decision making. *Communication and Group Decision Making*, 2, 81–113.
- Boyd, R., & Richerson, P. (1985). *Culture and the evolutionary process*. Chicago, IL: University Press, Chicago.
- Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
- Byun, K.-S., De Vos, C., Roberts, S. G., & Levinson, S. C. (2018). Interactive sequences modulate the selection of expressive forms in cross-signing. In C. Cuskley, M. Flaherty, H. Little, L. McCrohon, A. Ravignani & T. Verhoef (Eds.), *The 12th international conference on the evolution of language: (evolang xii)* (pp. 67–69). EvoLang Conference Proceedings.
- Caldwell, C. A., & Smith, K. (2012). Cultural evolution and perpetuation of arbitrary communicative conventions in experimental microsocieties. *PLoS One*, 7(8), e43807.
- Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach (No. 16)*. Princeton, NJ: Princeton University Press.
- Centola, D. (2018). *How behavior spreads: The science of complex contagions* (Vol. 3). Princeton, NJ: Princeton University Press.
- Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, 112(7), 1989–1994.
- Centola, D., Gonzalez-Avella, J. C., Eguiluz, V. M., & San Miguel, M. (2007). Homophily, cultural drift, and the co-evolution of cultural groups. *Journal of Conflict Resolution*, 51(6), 905–929.
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7(1), 19–22.
- Chavalarias, D. (2017). What's wrong with science? *Scientometrics*, 110(1), 481–503.
- Creanza, N., Kolodny, O., & Feldman, M. W. (2017). Greater than the sum of its parts? Modelling population contact and interaction of cultural repertoires. *Journal of the Royal Society Interface*, 14(130), 20170171.
- Cuskley, C., Loreto, V., & Kirby, S. (2018). A social approach to rule dynamics using an agent-based model. *Topics in Cognitive Science*, 10(4), 745–758.
- DeGraff, M. (2001). *Language creation and language change: Creolization, diachrony, and development*. Cambridge, MA: MIT Press.
- Dere, M., Beugin, M.-P., Godelle, B., & Raymond, M. (2013). Experimental evidence for the influence of group size on cultural complexity. *Nature*, 503(7476), 389.
- Dere, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, 113(11), 2982–2987.
- Dere, M., Perreault, C., & Boyd, R. (2018). Divide and conquer: Intermediate levels of population fragmentation maximize cultural accumulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1743), 20170062.
- Dixon, R. & Robert Malcolm Ward, D. (1997). *The rise and fall of languages*. Cambridge, UK: Cambridge University Press.

- Fay, N., & Ellison, T. M. (2013). The cultural evolution of human communication systems in different sized populations: Usability trumps learnability. *PLoS One*, 8(8), e71781.
- Fay, N., Garrod, S., & Carletta, J. (2000). Group discussion as interactive dialogue or as serial monologue: The influence of group size. *Psychological science*, 11(6), 481–486.
- Fay, N., Garrod, S., & Roberts, L. (2008). The fitness and functionality of culturally evolved communication systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509), 3553–3561.
- Fay, N., Garrod, S., Roberts, L., & Swoboda, N. (2010). The interactive evolution of human communication systems. *Cognitive science*, 34(3), 351–386.
- Fehér, O., Wonnacott, E., & Smith, K. (2016). Structural priming in artificial languages and the regularisation of unpredictable variation. *Journal of Memory and Language*, 91, 158–180.
- Ferdinand, V., Thompson, B., Kirby, S., Smith, K. (2013). Regularization behavior in a nonlinguistic domain. *Proceedings of the annual meeting of the cognitive science society*, 35.
- Fitch, W. T. (2008). Glossogeny and phylogeny: Cultural evolution meets genetic evolution. *Trends in Genetics*, 24(8), 373–374.
- Frankham, R., Briscoe, D. A., & Ballou, J. D. (2002). *Introduction to conservation genetics*. Cambridge, UK: Cambridge University Press.
- Fusaroli, R., Bahrami, B., Olsen, K., Roepstorff, A., Rees, G., Frith, C., & Tylén, K. (2012). Coming to terms: Quantifying the benefits of linguistic coordination. *Psychological Science*, 23(8), 931–939.
- Fusaroli, R., & Tylén, K. (2016). Investigating conversational dynamics: Interactive alignment, interpersonal synergy, and collective task performance. *Cognitive Science*, 40(1), 145–171.
- Garcia, A. C., & Baker Jacobs, J. (1999). The eyes of the beholder: Understanding the turn-taking system in quasi-synchronous computer-mediated communication. *Research on Language and Social Interaction*, 32(4), 337–367.
- Garrod, S., & Pickering, M. J. (2007). Alignment in dialogue. In G. Gaskell (Eds.), *The Oxford handbook of psycholinguistics* (pp. 443–451). Oxford: Oxford University Press.
- Gong, T., Minett, J. W., & Wang, W. S.-Y. (2008). Exploring social structure effect on language evolution based on a computational model. *Connection Science*, 20(2–3), 135–153.
- Gong, T., Shuai, L., Tamariz, M., & Jäger, G. (2012). Studying language change using price equation and pólya-urn dynamics. *PLoS One*, 7(3), e33171.
- González-Avella, J. C., Cosenza, M. G., Klemm, K., Eguíluz, V. M., & San Miguel, M. (2007). Information feedback and mass media effects in cultural dynamics. *Journal of Artificial Societies and Social Simulation*, 10(3), 9. Available at: <http://jasss.soc.surrey.ac.uk/10/3/9.html>.
- Hagen, E. H., & Hammerstein, P. (2006). Game theory and human evolution: A critique of some recent interpretations of experimental games. *Theoretical Population Biology*, 69(3), 339–348.
- Hahn, M. W., & Bentley, R. A. (2003). Drift as a mechanism for cultural change: an example from baby names. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(suppl 1), S120–S123.
- Henrich, J. (2004). Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses—The Tasmanian case. *American Antiquity*, 69(2), 197–214.
- Henrich, J., & McElreath, R. (2007). Dual-inheritance theory: The evolution of human cultural capacities and cultural evolution. In L. Barrett & R. Dunbar (Eds.), *Oxford handbook of evolutionary psychology*, (555–570). Oxford: Oxford University Press.
- Herzog, H. A., Bentley, R. A., & Hahn, M. W. (2004). Random drift and large shifts in popularity of dog breeds. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(suppl 5), S353–S356.
- Hovers, E., & Belfer-Cohen, A. (2006). “Now you see it, now you don’t”—Modern human behaviour in the middle paleolithic. In E. Hovers & S. L. Kuhn (Eds.), *Transitions before the transition* (pp. 295–304). New York: Springer.
- Hudson Kam, C. L., & Chang, A. (2009). Investigating the cause of language regularization in adults: Memory constraints or learning effects? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3), 815.

- Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language Learning and Development*, 1(2), 151–195.
- Janda, R. D., & Joseph, B. D. (2003). On language, change, and language change; or, of history, linguistics, and historical linguistics. In B. D. Janda & R. D. Joseph (Eds.), *The handbook of historical linguistics*, 3–180. Malden, MA: Blackwell.
- Kareev, Y., Lieberman, I., & Lev, M. (1997). Through a narrow window: Sample size and the perception of correlation. *Journal of Experimental Psychology: General*, 126(3), 278.
- Kempe, M., & Mesoudi, A. (2014). An experimental demonstration of the effect of group size on cultural accumulation. *Evolution and Human Behavior*, 35(4), 285–290.
- Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social learning strategies: Bridge-building between fields. *Trends in Cognitive Sciences*, 22(7), 651–665.
- Kline, M. A., & Boyd, R. (2010). Population size predicts technological complexity in Oceania. *Proceedings of the Royal Society B: Biological Sciences*, 277(1693), 2559–2564.
- Kobayashi, Y., & Aoki, K. (2012). Innovativeness, population size and cumulative cultural evolution. *Theoretical Population Biology*, 82(1), 38–47.
- Komarova, N. L., Niyogi, P., & Nowak, M. A. (2001). The evolutionary dynamics of grammar acquisition. *Journal of Theoretical Biology*, 209(1), 43–59.
- Komarova, N. L., & Nowak, M. A. (2003). Language dynamics in finite populations. *Journal of Theoretical Biology*, 221(3), 445–457.
- Kroch, A. S. (1989). Reflexes of grammar in patterns of language change. *Language Variation and Change*, 1(3), 199–244.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667–694.
- Lee, Y., Stabler, T. C. C. E. P., & Taylor, C. E. (2005). The role of population structure in language evolution. *Language*, 22(23), 24–25.
- Lipo, C. P., Madsen, M. E., Dunnell, R. C., & Hunt, T. (1997). Population structure, cultural transmission, and frequency seriation. *Journal of Anthropological Archaeology*, 16(4), 301–333.
- Lupyan, G., & Dale, R. (2010). Language structure is partly determined by social structure. *PLoS One*, 5(1), e8559.
- Mantonakis, A., Rodero, P., Lesschaeve, I., & Hastie, R. (2009). Order in choice: Effects of serial position on preferences. *Psychological Science*, 20(11), 1309–1312.
- Manzo, J. F. (1996). Taking turns and taking sides: Opening scenes from two jury deliberations. *Social Psychology Quarterly*, 59(2), 107.
- March, J. G. (2005). Parochialism in the evolution of a research community: The case of organization studies. *Management and Organization Review*, 1(1), 5–22.
- Mesoudi, A. (2011). Variable cultural acquisition costs constrain cumulative cultural evolution. *PLoS One*, 6(3), e18239.
- Mesoudi, A., & Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509), 3489–3501.
- Mesoudi, A., Whiten, A., & Laland, K. N. (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, 29(4), 329–347.
- Mueller-Frank, M. (2013). A general framework for rational learning in social networks. *Theoretical Economics*, 8(1), 1–40.
- Muthukrishna, M., & Schaller, M. (2019). Are collectivistic cultures more prone to rapid transformation? Computational models of cross-cultural differences, social network structure, dynamic social influence, and cultural change. *Personality and Social Psychology Review*, 24(2), 103–120.
- Muthukrishna, M., Shulman, B. W., Vasilescu, V., & Henrich, J. (2014). Sociality influences cultural complexity. *Proceedings of the Royal Society B: Biological Sciences*, 281(1774), 20132511.

- Neiman, F. D. (1995). Stylistic variation in evolutionary perspective: inferences from decorative diversity and interassemblage distance in Illinois woodland ceramic assemblages. *American Antiquity*, 60(1), 7–36.
- Olfati-Saber, R., & Murray, R. M. (2004). Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9), 1520–1533.
- Pelikan, M., Goldberg, D. E., & Cantú-Paz, E. (2000). Bayesian optimization algorithm, population sizing, and time to convergence (Tech. Rep.). CA, USA: Lawrence Livermore National Lab.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(2), 169–190.
- Powell, A., Shennan, S., & Thomas, M. G. (2009). Late Pleistocene demography and the appearance of modern human behavior. *Science*, 324(5932), 1298–1301.
- Raviv, L., Meyer, A., & Lev-Ari, S. (2019). Larger communities create more systematic languages. *Proceedings of the Royal Society B: Biological Sciences*, 286(1907), 20191262.
- Reali, F., & Griffiths, T. L. (2009). Words as alleles: Connecting language evolution with Bayesian learners to models of genetic drift. *Proceedings of the Royal Society B: Biological Sciences*, 277(1680), 429–436.
- Richerson, P. J., & Boyd, R. (2008). *Not by genes alone: How culture transformed human evolution*. Chicago, IL: University of Chicago Press.
- Rubinstein, A. (2006). Dilemmas of an economic theorist. *Revista de Economía Institucional*, 8(14), 191–213.
- Sabherwal, R., Hirschheim, R., & Goles, T. (2001). The dynamics of alignment: Insights from a punctuated equilibrium model. *Organization Science*, 12(2), 179–197.
- Sacks, H., Schegloff, E. A., & Jefferson, G. (1978). A simplest systematics for the organization of turn taking for conversation. In J. Schenkein (Eds.), *Studies in the organization of conversational interaction* (pp. 7–55). New York: Academic Press.
- Segovia-Martín, J., Walker, B., Fay, N., & Tamariz, M. (2019). Network connectivity dynamics affect the evolution of culturally transmitted variants. arXiv e-prints, arXiv:1902.06598.
- Segovia-Martín, J., & Tamariz, M. (2020). Testing early and late connectivity dynamics in the lab: an experiment using 4-agent micro-societies. PsyArXiv. doi: <https://doi.org/10.31234/osf.io/nuf78>
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.
- Shennan, S. J., & Wilkinson, J. R. (2001). Ceramic style change and neutral evolution: A case study from Neolithic Europe. *American Antiquity*, 66(4), 577–593.
- Shteingart, H., Neiman, T., & Loewenstein, Y. (2013). The role of first impression in operant learning. *Journal of Experimental Psychology: General*, 142(2), 476.
- Smaldino, P. E., & Schank, J. C. (2012). Movement patterns, social dynamics, and the evolution of cooperation. *Theoretical Population Biology*, 82(1), 48–58.
- Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition*, 116(3), 444–449. <https://doi.org/10.1016/j.cognition.2010.06.004>.
- Stasser, G., & Taylor, L. A. (1991). Speaking turns in face-to-face discussions. *Journal of Personality and Social Psychology*, 60(5), 675.
- Stone, V. A. (1969). A primacy effect in decision-making by jurors. *Journal of Communication*, 19(3), 239–247.
- Stubbersfield, J. M., Tehrani, J. J., & Flynn, E. G. (2015). Serial killers, spiders and cybersex: Social and survival information bias in the transmission of urban legends. *British Journal of Psychology*, 106(2), 288–307.
- Tamariz, M., Ellison, T. M., Barr, D. J., & Fay, N. (2014). Cultural selection drives the evolution of human communication systems. *Proceedings of the Royal Society B: Biological Sciences*, 281(1788), 20140488.
- Tamariz, M., & Kirby, S. (2015). Culture: Copying, compression, and conventionality. *Cognitive Science*, 39(1), 171–183.
- Turner, M. A., & Smaldino, P. E. (2018). Paths to polarization: How extreme views, miscommunication, and random chance drive opinion dynamics. *Complexity*, 2018, 1–17.

- Vaesen, K. (2012). Cumulative cultural evolution and demography. *PLoS One*, 7(7), e40989.
- Vale, G. L., Flynn, E. G., Kendal, J., Rawlings, B., Hopper, L. M., Schapiro, S. J., Schapiro, S. J., Lambeth, S. P., & Kendal, R. L. (2017). Testing differential use of payoff-biased social learning strategies in children and chimpanzees. *Proceedings of the Royal Society B: Biological Sciences*, 284(1868), 20171751.
- White, J. W., Rassweiler, A., Samhouri, J. F., Stier, A. C., & White, C. (2014). Ecologists should not use statistical significance tests to interpret simulation model results. *Oikos*, 123(4), 385–388.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article:

Data S1. Content bias, coordination bias, and memory.