Cultural diffusion dynamics depend on behavioural production rules

Michael Chimento*[1,2,3], Brendan J. Barrett[2,3,4,5], Anne Kandler[5], and Lucy M. Aplin[1,2]

[1] *Cognitive and Cultural Ecology Research Group, Max Planck Institute of Animal Behavior; Am Obstberg 1, 78315 Radolfzell, Germany*

[2] *Centre for the Advanced Study of Collective Behaviour, University of Konstanz; Universitätsstraße 10, 78464 Konstanz, Germany*

[3] *Department of Biology, University of Konstanz; Universitätsstraße 10, 78464 Konstanz, Germany*

[4] *Department for the Ecology of Animal Societies, Max Planck Institute of Animal Behavior; Am Obstberg 1, 78315 Radolfzell, Germany*

[5] *Department of Human Behavior, Ecology and Culture, Max Planck Institute for Evolutionary Anthropology; Deutscher Platz 6 04103 Leipzig, Germany*

December 22, 2021

Abstract

Culture is an outcome of the acquisition of knowledge about behaviour through social transmission, and its subsequent production. Transmission and production are often discussed interchangeably or modeled separately, yet to date, no study has accounted for both processes and explored their interaction. We present a generative model that integrates the two in order to explore how variation in either might shape cultural diffusion dynamics. Agents make behavioural choices that change as they learn from their behavioural productions. Their repertoires also change over time, and the social transmission of behaviours depends on their frequency. We diffuse a novel behaviour through social networks across a large parameter space to demonstrate how accounting for both transmission and production reveals dependencies between individual-level behavioural production rules and population-level diffusion dynamics. We then investigate how such dependencies might affect the performance of two commonly used inferential models for social learning; Network-based Diffusion Analysis (NBDA), and Experienced Weighted Attraction models (EWA). By clarifying the distinction between acquisition and usage, we illuminate often-overlooked theoretical differences between social learning and social influence. These distinctions yield consequences and new considerations for how inferential methods are applied to empirical studies of culture.

Keywords

cultural evolution, social learning, agent based model, network-based diffusion analysis, experience weighted attraction

1 Background

Cultural evolution, or the changes in distributions of socially-learned traits over time, provides a secondary system of inheritance that potentially enables rapid adaptive plasticity across species [1, 2, 3, 4]. Understanding how distributions of cultural traits change requires an understanding of two central processes of iterated learning [5]: the learning of novel behaviour via social transmission, and the subsequent production of these socially learned behaviours. The linkage between transmission and production is encapsulated in a well-accepted definition of social transmission: when a behaviour is produced by an individual (demonstrator), it exerts a positive causal influence on the rate at which that individual's associates (observers) acquire the same behaviour [6]. Precisely which and when cultural behaviours are socially acquired from conspecifics depends on cognitive [7, 8, 9], social [6, 10, 11], and environmental factors [12, 13]. Once knowledge of a behaviour is acquired, the realized production, or usage of that behaviour is also determined by these factors [13, 14], but not necessarily in the same way. Confusingly, the literature has given the label of enquotesocial learning to both the acquisition of novel behaviours via social transmission and the influence of social information on production decisions [4, 15, 16, 17]. We argue in this study that transmission and production must be treated as conceptually distinct, and then carefully disentangled in order to understand how constraints pertaining to either might affect emergent cultural outcomes.

One major goal of empirical studies on the study of animal culture has been to identify when social learning is responsible for the cumulative uptake of a behaviours in a population, as opposed to asocial learning (i.e. individual innovation). A model of social transmission is therefore required to explain whether individuals transition from naïve to knowledgeable because of social or asocial learning. One primary determinant of social transmission is the opportunity for naïve individuals to learn from knowledgeable individuals, generally represented by their social association. This link between social association and the likelihood of behavioural acquisition is at the core of network-based diffusion analysis (NBDA) [6, 11, 18], which aims to statistically test for social transmission against a null hypothesis of asocial learning. For any given behavioural diffusion, NBDA accounts for social network connections to knowledgeable individuals, as well as orders or times of acquisition by naïve individuals to estimate the strength of social learning relative to asocial learning. Since it's introduction over a decade ago, it has proven to be an invaluable tool to identify social learning [18, 19], applied throughout the animal kingdom to study social learning in birds [20, 21, 22, 23], fish [24, 25], cetaceans [26, 27, 28], and primates [29, 30, 31, 32]. It has also been extended several times to include multiple alternative networks [10], dynamic networks in which relationships change over time [30], and complex contagion [33].

As with any model-based inferential framework, NBDA has several short-comings. The event of acquisition is not observable, so in the above studies, the time of first observed production of a behaviour by an individual is assumed to be equivalent to their time of acquisition of knowledge of that behaviour. This is usually a necessary assumption, as inferring knowledge state without the behavioural output is nearly impossible, especially in non-human animals. However this leads to a conflation of transmission and production. Secondly, while a model of social transmission is highly effective at modeling the initial diffusion of a behaviour through a population, it is not suited to address cultural change, where novel cultural variants compete with established variants. Finally, NBDA analyses do not usually account for behavioural frequency information, a determinant of transmission

beyond association—although this is recommended if possible [11, 34]). In empirical studies, behavioural frequency information appears to have only been included in the models of two previously published NBDA analyses, on ravens and house sparrows [35, 36]. In theory, the inclusion of a model of production would complement these shortcomings: it would help differentiate between acquisition of knowledge and production, and define a distribution of behavioural frequencies which would thus influence transmission probabilities.

Another goal of empirical studies is understanding the psychological rules that govern choices between between behaviours, particularly whether and how individuals integrate individual and social information during decision-making. This has implications for broader evolutionary theory, as it could determine if, how and when adaptive behaviours spread through populations, ultimately yielding fitness consequences. These psychological rules have been most commonly studied using experience weighted attraction (EWA) models [7, 37]. EWA is a dynamic learning model originally developed to analyse how different types of information influence behavioural choice in economic games [37, 38]. It is an extended model of reinforcement learning, which is an important cognitive factor influencing patterns of production, determining which behaviours are retained and which are abandoned [39]. In particular, this model of production integrates both individual and social information in decision-making. It is an excellent tool for understanding how choices might change with experience, and can be re-parameterized to estimate the relative contributions of different biases acting on socially observed information [7]. Due to this flexibility, EWA models have become a popular method to identify social learning (*sensu* social influence on production) in human and animal populations, and to analyse social learning biases and strategies.

Variations of this model of production have been used generatively to study social learning strategies, both theoretically [40, 41] and inferentially across many species , including chimpanzees (*Pan troglodytes*) [16], capuchin monkeys (*Cebus capucinus*) [15], vervet monkeys (*Chlorocebus pygerythrus*) [42], great tits (*Parus major*) [4], and humans (*Homo sapiens*) [7, 17, 43, 44]. However, the formalization of this model requires the definition of a repertoire of behaviours which an individual may produce. To date, implementations of the model have not accounted for differences in repertoire size over time or between individuals. Social learning suggests the acquisition of new, previously unknown behaviours by naïve individuals [6], leading to lasting changes in behaviour. Thus, the tacit assumption of homogeneous knowledge of the EWA model could be problematic when studying social learning in the sense of social transmission, rather than social learning in the sense of social influence on production. This encapsulates a second conflation of production and transmission—social information use during production decisions is not equivalent to its use in the acquisition of new behaviours.

Most social learning studies to date have therefore studied either 1) acquisition by assuming that the time of acquisition approximates the time of production, or 2) production by assuming that all individuals have already acquired knowledge of the same number of behaviours. To our knowledge, no model of cultural evolution exists in which both transmission and production interact and inform each other—an illustration of the ongoing disconnect between the consideration of psychological and social mechanisms [45]. To accomplish this, we construct a generative agent-based model of cultural diffusion that synthesizes both processes in networked populations. A sub-model of transmission informs the repertoires of agents, while a sub-model of production determines the frequencies of behavioural usage based on personal experience and social in-

formation. Behavioural frequency information is then fed back into the sub-model of transmission, where it influences the probability of transmission. First, we explore theoretically how changes to production rules might affect the diffusion of a novel behaviour. Second, we apply commonly used inferential models (NBDA and EWA) on the generated data to highlight potential consequences of variation between the order of acquisition and order of first production for inference in empirical studies. Our study emphasizes the importance of differentiating between social learning and social information use during production, two processes that have been often treated as equivalent in wider discussions of cultural evolution.

2 Methods

We assumed a population of constant size N where agents are situated within a social association network that constrains observation and acquisition opportunities to connected nodes (**Figure 1**). All agents initially knew how to produce one behaviour, a, which can be interpreted as an established tradition that has diffused through the population prior to the beginning of the simulation. To remove stochasticity resulting from differences in innovation timing, one randomly selected agent also had knowledge of a novel behaviour b. Both a and b generated an equivalent payoff, and they were assumed to be mutually exclusive (e.g. an agent can use only one behaviour at a time). In each time-step, agents decided which behaviour in their repertoire to produce. This decision was influenced by past personal experience of payoffs, as well as social information provided by an agent's neighbours in the association network (see **Section 2.2** for details). Despite both behaviours yielding the same payoff, reinforcement learning played an important role in the simulations, as agents still had to learn the expected payoffs of both behaviours with experience. Additionally, agents had the opportunity to acquire behaviour b if it had been produced by at least one of their neighbours (see **Section 2.1** for details). Repeating this cycle of production and (potential) acquisition, we recorded the times where each agent has acquired behaviour b and produced it for the first time. The simulation ended when all agents have produced behaviour b.

In brief, our simulation simultaneously modeled the spread of a novel behaviour through a population and the subsequent agents' decision of which behaviour to produce. Below, we describe the assumed dynamics of social transmission and production, highlight their interplay, and detail the simulation set-up as well as the inferential methods used. A list of parameters and their interpretations can be found in **Table 1**.

2.1 Sub-model of behavioural acquisition

The novel behaviour was acquired through social transmission using a model consistent with the dynamic underlying the NBDA framework. The probability of transmission took into consideration an individual's associates and their behavioural productions, a rate of social transmission, and a rate of asocial learning. We estimated a probability of transmission of the novel behaviour b for all naïve individuals i, i.e. individuals that only possess behaviour a, at each time-step t by calculating:

$$P(\text{individual } i \text{ acquires } b|t) = 1 - exp(\lambda_b(T(a_i, z_b(t), w_b(t)) + A)). \quad (1)$$

The parameter λ_b represented the baseline rate of transmission, and quantified how easily learned the novel behaviour is either socially or asocially. A represented the presence of asocial learning, and could take a value of 0 or 1, where 0 is if no asocial learning is assumed to take place. If the agent acquired the behaviour, it was then able to be used by the agent. NB: the agent only acquired knowledge of that behaviour, and did not produce the behaviour upon acquisition.

The transmission function $T(a_i, z_b(t), w_b(t))$ considered agent i's association, $a_i = [a_{ij}]_{j=1,\dots,N}$, their knowledge states, $z_b(t) = [z_{b,j}(t)]_{j=1,\dots,N}$, i.e. whether they have acquired behaviour b at time t or not, and their transmission weights (behavioural frequency), $w_b(t) = [w_{b,j}(t)]_{j=1,\dots,N}$. In more detail, $T(a_i, z_b(t), w_b(t))$ was defined by

$$T(a_i, z_b(t), w_b(t)) = s \sum_{j=1}^{N} a_{i,j} z_{b,j}(t) w_{b,j}(t) \tag{2}$$

where $a_{i,j} = 1$ if i and j were neighbours in the association network and $a_{i,j} = 0$ otherwise. Equally, $z_{b,j}(t) = 1$ if agent j had behaviour b in its repertoire at time t and $z_{b,j}(t) = 0$ if not. Lastly, the transmission weight $w_{b,j}(t)$ represented the proportion of time individual j has produced behaviour b within the memory window, i.e. the last m time-steps

$$w_{b,j}(t) = \frac{\sum_{r=t-m}^{t} n_{b,j}(r)}{m} \tag{3}$$

where $n_{b,j}(r) = 1$, agent j produced b r times steps ago; otherwise $n_{b,j}(r) = 0$. In summary, the transmission function recorded the proportion of the relative usage of behaviour b of all knowledgeable neighbours of agent i. This was weighted by the rate of social transmission per unit connection, relative to the asocial

Table 1: **Summary of notation**. Mathematical symbol, name, range of possible values, and short description of relevant items in our model.

	Symbol	Name	Range	Description
Transmission sub-model	λ_b	baseline rate	$[0,\infty]$	the baseline rate of transmission, or learnability independent of social or asocial mechanisms, for behaviour b.
	A	asocial learning parameter	$\{0,1\}$	binary variable indicating the presence (or absence) of asocial learning
	s	social transmission rate	$[0,\infty)$	the rate of social acquisition per unit connection, relative to A.
Production sub-model	ϕ	recent experience bias	$[0,1]$	Weight given to the most recently received payoff vs. previously received payoffs $(1-\phi)$.
	σ	social info. bias	$[0,1]$	Weight given to socially observed information vs. individually experienced information $(1-\sigma)$.
	τ	conservatism	$[0,\infty)$	Sensitivity to differences in expected payoffs.
	f_P	production conformity	$[0,\infty)$	Sensitivity to the relative proportions of socially observed behaviours during production; anti-conformist (<1), linearly conformist (1), or conformist (>1)
Misc.	m	memory window	$[1,\infty)$	The number of time-steps back for which an agent remembers their socially observed information.

learning rate A. Thus, $s = 10$ can be interpreted as social learning increasing the likelihood of acquisition by a factor of 10 for each knowledgeable associate.

The transmission weight function, $w_{b,i}(t)$, sat at the nexus of our sub-models of transmission and production, and was where we implemented their interaction. The inclusion of production information adjusted the influence that a knowledgeable neighbour might have, reducing their influence in the case in which that neighbour knew b, yet rarely produced it. It allowed the behavioural frequencies, derived from the production sub-model, to influence the diffusion probabilities calculated from the transmission sub-model, linking the two sub-models together.

2.2 Sub-model of behavioural production

Each time-step, agents chose to produce one behaviour from it's repertoire, for which it received a payoff, using the sub-model of production to make a decision. While a variety of parameterizations were possible, the model presented in this paper is consistent with prior applications of EWA inferential models. Our model was characterized by a set of 4 parameters that governed the updating of behavioural choices over time: recent experience bias (ϕ), social information bias (σ), an conservatism bias (τ) and frequency-dependent conformity bias (f_P). Agents updated their probabilities of producing one behaviour over another by accounting for personal experience, as well as socially observed information within memory.

First, each agent i first calculated the attraction score for any behaviour k in its repertoire Z_i using the expression

$$AT_{i,k}(t) = \phi\pi_k(t-1) + (1-\phi)AT_{i,k}(t-1). \tag{4}$$

The variable $\pi_k(t-1)$ described the payoff obtained from behaviour k in the last time-step. If k was not chosen in $t-1$, then $\pi_k(t-1) = 0$, allowing for attraction loss over time. Parameter ϕ determined the importance of the most recently received payoff *vis-a-vis* previously experienced payoffs.

In the second step, each agent i converted its attraction towards any of the behaviours k in its repertoire Z_i into a probability by using a softmax rule

$$I_{i,k}(t) = \frac{exp(\tau AT_{i,k}(t))}{\sum_{k \in Z_i} exp(\tau AT_{i,k}(t))} \tag{5}$$

Parameter τ controlled the sensitivity towards differences in the attraction score: high values of τ result in an almost deterministic choice of the behaviour with the highest payoff (behavioural conservatism) whereas low values of τ lead to choices almost independent from attraction scores. Values of τ between these extremes therefore modelled different degrees of conservatism (or exploration behaviour). **Equation** (5) assigned probabilities to behaviours even when their attraction score is 0, allowing for the possibility of production without the agent having prior experience producing that behaviour.

In the third step, each agent evaluated its social information by counting how many times its neighbours have produced any of the behaviours in its repertoire within the memory window, modulated by the production conformity parameter

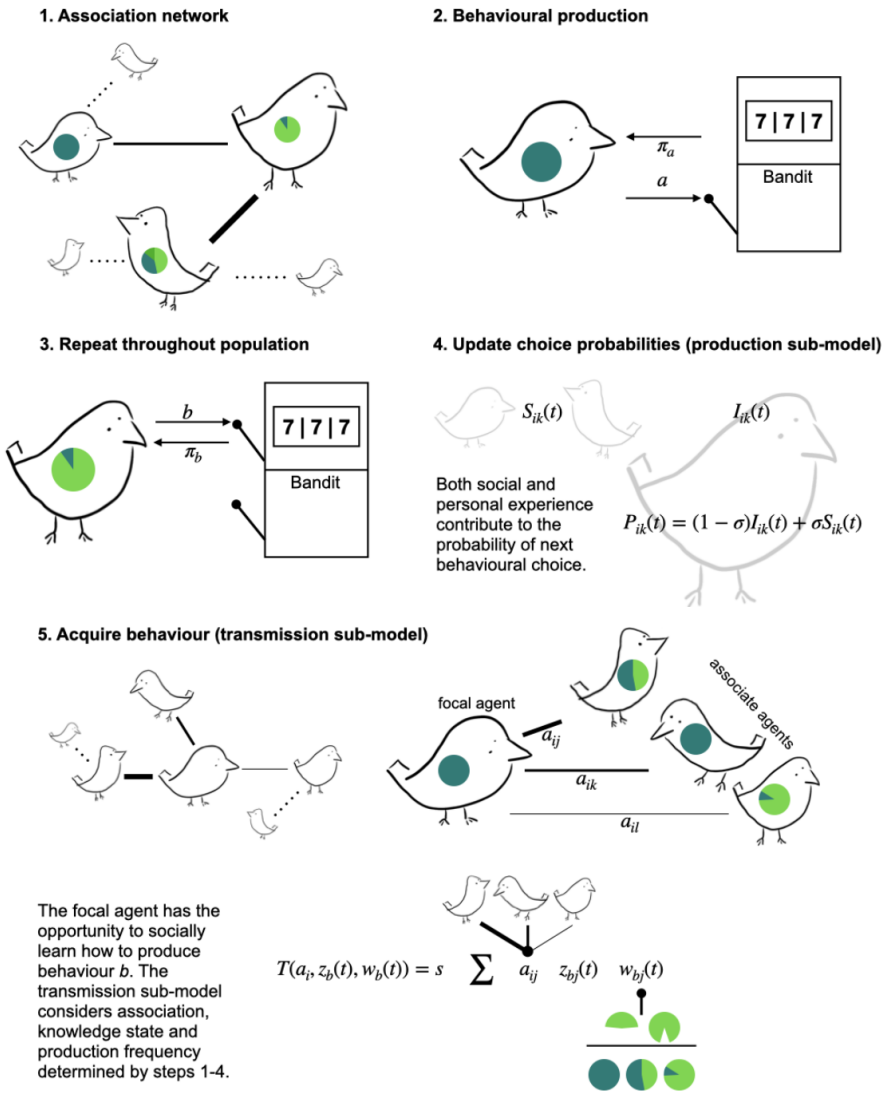

Figure 1: **Model schematic. 1)** Agents, arbitrarily represented by birds, are arranged in an association network, with varying edge-weights. Pie charts represent an agent's repertoire, and the frequencies of behavioural production within memory. **2)** Agents interact with a two-armed bandit, although the number of arms known to the agent depends on the state of their repertoire. This agent knows how to produce behaviour a and b, chooses a, and receives payoff π_a. **3)** Each knowledgeable agent produces one behaviour per time-step. **4)** After all agents produce a behaviour, personal information about the payoff received and observed social information is taken into account to calculate a probability for producing a behaviour in the next time-step. **5)** Before the time-step ends, a diffusion phase occurs where agents can acquire behaviour b and add it to their repertoire. This depends on the social transmission rate, association with neighbours, and their behavioural frequencies within memory. Here, the focal agent already knows a, and could socially learn b from associate j or l to be used in the next time-step.

$$S_{i,k}(t) = \frac{\left(\sum_{j=1}^{N} a_{i,j} \sum_{r=t-m}^{t} n_{k,j}(r)\right)^{f_P}}{\sum_{k \in Z_i} \left(\sum_{j=1}^{N} a_{i,j} \sum_{r=t-m}^{t} n_{k,j}(r)\right)^{f_P}}. \tag{6}$$

The parameter $f_P > 1$ determined the strength of frequency dependent influence on the agent. In order to handle heterogeneous knowledge states, agents only observed behaviours which were in their own repertoire assuming that if an agent did not know how to produce a behaviour, it's observation could not influence what an agent would produce in that time-step. However, this information did influence the potential acquisition of this behaviour as described in section 2.1.

Lastly, each agent i combined their different sources of information, personally experienced payoffs and socially observed behaviours, to generate a probability of producing any of the behaviours k in its repertoire Z_i

$$P_{i,k}(t) = (1-\sigma) I_{i,k}(t) + \sigma S_{i,k}(t) \tag{7}$$

where σ described the preference for individual and social information: values of $sigma$ close to 0 would almost neglect social information while values close to 1 would almost neglect individual information.

In summary, in each time-step each agent i produced a behaviour k from its repertoire with probability $P_{i,k}(t)$. This production of agent i, as well as the productions of all agents in its neighbourhood, influenced the values of $I_{i,k}(t+1)$ (see **Equations** (4) and (5)) and $S_{i,k}(t+1)$ (see **Equation** (6)), potentially leading to a changed production probability $P_{i,k}(t+1)$ at time $t+1$.

2.3 Initialization, measurements & conditions of simulations

We ran simulations with populations of size $N = 24$ agents. Agents were arranged in 4 random network architectures (**Figure S1** for exemplar networks), all with edge weights set to 1 to eliminate stochastic variation arising from differences in association. Random regular networks and small world networks (using Watts-Strogatz algorithm [46]) were fixed degree networks with $k = 6$ neighbours for each agent. Erdős–Rényi networks and Barabási–Albert scale free networks were parameterized such that $\bar{k} = 6$. Degree was important to standardize between networks for any comparison, as the probabilities of behaviour acquisition from the NBDA equations are degree dependent. We did not include complete networks, as complete networks are an unrealistic baseline for cultural models [47].

All agents were initialized as already knowing behaviour a ($Z_a = 1$, $Z_b = 0$). One randomly selected agent had knowledge of b from the first time-step to initiate diffusion. Both behaviours were of equal payoff ($\pi_a = \pi_b = 1$). Despite obtaining the same payoff, reinforcement learning still played an important role in the simulations, as agents had to learn the expected payoffs of both behaviours with experience.

Within the sub-model of behavioural acquisition, the baseline learning rate was set to $\lambda_b = 0.05$. The social transmission rate parameter was set to $s = \{5, 10, 15\}$. The asocial learning parameter was set to $A = 0$, such that the novel behaviour could not be innovated independently (see **Equation** (1)). Within the sub-model of behavioural production, we explored populations in which agents were either strongly biased towards

past payoffs, equally biased towards past and present payoffs, or strongly biased towards present payoffs ($\phi = \{0.01, 0.5, 0.99\}$; **Equation** (4)). Agents could be linearly sensitive (non-conservative) or hyper-sensitive (conservative) to differences in attraction scores ($\tau = 1, 2$; **Equation** (5)). When choosing which behaviour to produce, agents were either unbiased (linearly sensitive) or biased towards high frequency behaviours ($f_P = \{1, 3\}$; **Equation** (6)). Agents were either biased towards personal information, unbiased between social and personal information, or biased towards social information ($\sigma = \{0.25, 0.5, 0.75\}$; **Equation** (7)). Lastly, we tested varying memory window values of $m = \{10, 20, 30\}$, which would directly impact both the transmission sub-model (**Equation** (2)), as well as the production sub-model (**Equation** (6)). Intermediary and extended values of these parameters were also evaluated, although they showed no unexpected or reversed patterns and were thus excluded from the paper for concision.

For each of the 1296 parameter constellations described above, we ran 100 simulations. A simulation ended when the novel behaviour had been acquired and produced by each agent at least once. In each simulation we recorded the time-steps at which each agent acquired the novel behaviour and first produced the novel behaviour. We also recorded frequencies of both behaviours, i.e. the proportion of the population producing a (or b), every 5 time-steps. From these data we derived time-to-diffusion (TTD), time-to-first-production (TTFP), and vectors of order of acquisition (o_a) and order of first production (o_p), and vectors of time of acquisition (t_a) and time of first production (t_p). Vector o_a contains the position of each agent in the acquisition sequence. In a simplified example where $N = 3$ and $o_a = [3, 1, 2]$, this means that agent 1 acquired the novel behaviour last, agent 2 first and agent 3 second. Similarly, vector o_p contains the position of each agent in the first production sequence and the time of acquisition. Vector t_a contains the acquisition time of each agent. For example, again assuming $N = 3$ and $t_a = [20, 5, 15]$, this means that agent 1 acquired the novel behaviour at time-step 20, agent 2 at 5 and agent 3 at 15. Vector t_p contains the times of first production of each agent. To quantify the divergence in the orders of acquisition and first production, we calculated a divergence score as the mean Manhattan distance

$$d_{\text{order}} = \frac{1}{N-1} \sum_{i=1}^{N} |o_a(i) - o_p(i)| \tag{8}$$

representing the mean difference between acquisition and first production position in the population (excluding the first production of the seed agent, thus $N-1$). Further, we calculated the proportion of the population who obtained production positions that differed from their acquisition positions. To quantify the delay between the acquisition and first production timing, we calculated a delay score as the mean Manhattan distance

$$d_{\text{time}} = \frac{1}{N} \sum_{i=1}^{N} |t_a(i) - t_p(i)| \tag{9}$$

representing the mean time delay between acquisition and first production time in the population. In sections 3.1 and 3.2 we use a random regular network and the following parameter constellation as reference setting

$$s = 5, \phi = 0.5, \tau = 1, f_P = 1, \sigma = 0.5, m = 10. \tag{10}$$

We then demonstrate how changing one parameter at a time affects transmission and production dynamics, in terms of these aforementioned metrics.

2.4 Testing performance of NBDA and EWA on generated data

Finally, we explored how NBDA and EWA might perform on data generated by our model. As a proof of concept we first tested a TADAc NBDA model [11] on time-of-acquisition data, generated either by social diffusion (i.e. $A = 1$ and $s = 5$) or asocial learning (i.e. $A = 1$ and $s = 0$) without the influence of production frequencies, to demonstrate that dynamic modelled in the transmission sub-model coincides with the dynamic assumed in the NBDA framework. In this scenario, NBDA should recover the model parameters s and λ_b with high certainty. We then tested the same inferential model on time-of-first-production data, generated by either social diffusion or asocial learning. In other words, we tested how much a linkage of social transmission and production could influence inference results when the inferential method was blind to such a distinction. This blindness to production information is a common assumption in published NBDA analyses, as they have not included behavioural frequency information as a transmission weight.

In detail, we simulated one diffusion at each point in the parameter space described (excluding varying s, set to our reference value of 5) in **Section 2.3** for both scenarios described above, resulting in 432 different simulations. s was not varied, as we are focused on the effect of production parameters on inference, rather than transmission parameters. We increased population size to 100 to increase power during analysis, with mean degree set to 10 to create reasonably sparse networks. Using the recommended inferential steps, we ran both a social and asocial TADAc models to determine support for social learning [6]. For simulations of social diffusions, the seed agent was included as a demonstrator in the model. We compared AICc scores between the two models, along with parameter estimates, to determine relative support for social or asocial transmission [6]. Support for social learning was defined as $\Delta AICc > 0$, where $\Delta AICc = AICc_{asocial} - AICc_{social}$. We then compared how divergence and delays between acquisition and production affected the level of true or false support for the underlying generative process.

We use the same strategy to evaluate the performance of EWA analyses. We estimate only ϕ and σ parameters under three scenarios: 1) homogeneous repertoires and no diffusion, 2) heterogeneous repertoires with social diffusion using the acquisition sub-model, and 3) heterogeneous repertoires with only asocial diffusion using the acquisition sub-model. Scenario 1 is again a proof of concept, and the EWA model is expected to recover accurate estimates of ϕ and σ. Scenarios 2 and 3 will shed light on the question of how accurately a standard EWA model performs if behaviours are allowed to diffuse between individuals. For each scenario, we performed 9 simulations, one for each combination of ϕ and σ values ($\{0.25, 0.5, 0.75\}$ for both parameters), while setting τ and f_P to 1. A small amount of random noise (drawn from $N(0, 0.01)$) was added to these values in each agent to provide variation for the inferential models to work with. All simulations used random regular networks. In scenario 2, s was set to the reference value of 5.

For scenario 1, we recorded behavioural productions for 300 time-steps to provide generous power for the analysis. In scenarios 2 and 3, each simulation ran for twice as long as it's TTD (social diffusion: range $[65 - 199]$, asocial learning: $[89, 235]$ time-steps). This meant that for each simulation, we recorded equal numbers of observations both during diffusion, and after all agents acquired knowledge of the novel behaviour.

We then fit inferential models to the production data from each simulation from these three scenarios using a Hamiltonian Markov Chain Monte Carlo (MCMC). Models were run using 3 chains, 2500 iterations, with 1000 warm-up iterations. All models were fit using R v. 4.0.2 [48] with Stan v. 2.27 [49] via Rstan v. 2.21.2 [50]. Good model convergence was visually confirmed from traceplots and rank histograms.

3 Results

Our model explicitly accounted for the acquisition of knowledge, as well its production, allowing us to compare these two processes. We first explored the dynamics of either sub-model in isolation using the reference parameter constellation. Simulating data with the production sub-model alone required initialising all agents with knowledge of both behaviours, resulting in a flat acquisition curve constant at 100% of the population, and generating an r-shaped production curve that ended quickly (mean TTFP: 6.34 time-steps; **Figure S2A**). We then simulated data with transmission disconnected from production, which led to an s-shaped acquisition curve that completed much faster than the first production curve (mean TTD 6.41 time-steps; **Figure S2B**). Further, the shape of the production curve was not connected to the acquisition curve, leading to a very large divergence between the two in later stages of the simulations (mean TTFP: 18.7 time-steps).

Linking the transmission and production sub-models led to the dynamic depicted in **Figure 2A** (see also Table 2 for the corresponding values of the considered statistics). We observed that the acquisition curve was delayed (mean TTD: 48.90 time-steps) compared to the "transmission only" scenario above, due to the fact that an agent could only become a demonstrator once it had produced the novel behaviour. We also observed that the shape of the production curve was now s-shaped, as production was conditioned on acquisition. The mean time it took for all agents to produce the novel behaviour at least once was $55.60\,[41,79]$ (89% HPDI) time-steps. Thus, the production curve lagged behind the acquisition curve, as agents retained a higher expected payoff for the established behaviour and preferred it's usage by the time full diffusion is reached (**Figure 2A**). We note that if the simulations were allowed to continue running, the behaviours trended towards an equilibrium production ratio of $1:1$ at the population level, as both behaviours yielded equivalent payoffs. This was the case even when social information bias σ was high, which only resulted in larger short term deviations caused by stochastic drift, and did not permanently bias populations towards the established behaviour.

Importantly, production did not simply lag behind knowledge acquisition, as we observed divergences in the ordering of these events throughout the simulations. To show this, we compared the order of acquisition (o_a) to the order of first production of the novel behaviour (o_p) (**Figure 2B**). Beyond the first enquoteseed agent, we found large variation in these orders, with only 25% of agents producing the behaviour in the same order as acquiring knowledge of it (i.e. dark diagonal pattern in **Figure 2B**), and a divergence score $d_{\text{order}} = 2.16$. This can be interpreted as after agent i acquired knowledge of the novel behaviour, approximately 2 more agents acquired knowledge of it before agent i first produced it. Simulations using the reference setting obtained a mean delay score of $d_{\text{time}} = 5.24$, which represents the mean amount of time that passed between an agent acquiring a behaviour and producing it.

Table 2: **Summary of results**. Mean and HPDI (89% highest posterior density interval) values for time-to-diffusion (TTD), time-to-first-production (TTFP), divergence of orders of acquisition and first production (d_{order}), and time delay between acquisition and first production (d_{time}). The reference constellation of parameter values is given first, and rows are arranged by model parameters.

		TTD		TTFP		d_{order} (divergence)			d_{time} (delay)	
	value	mean	HPDI	mean	HPDI	% divergent	mean	HPDI	mean	HPDI
	reference constellation: random regular, $s = 5$, $\phi = 0.5$, $\tau = 1$, $f_P = 1$, $\sigma = 0.5$, $m = 10$									
		48.90	[31,66]	55.60	[41,79]	0.75	2.16	[1.13,3.13]	5.24	[4.04,7.46]
network										
random regular		*reference*								
small world		65.19	[45,81]	71.83	[50,87]	0.64	1.26	[0.78,1.91]	5.01	[3.08,6.54]
Erdős–Rényi		50.43	[31,73]	56.85	[35,77]	0.74	1.96	[1.22,2.87]	5.04	[3.08,6.25]
Barabási–Albert		56.92	[29,81]	62.53	[35,86]	0.72	1.88	[1.04,2.43]	4.82	[3.38,6.21]
s										
	5	*reference*								
	10	33.46	[21,45]	41.78	[29,54]	0.81	2.70	[1.65,3.48]	5.13	[3.46,6.71]
	15	28.97	[17,42]	37.72	[26,50]	0.82	3.03	[2.09,4.09]	5.08	[3.33,6.62]
ϕ										
	0.01	39.48	[27,52]	43.53	[29,55]	0.69	1.66	[0.7,2.17]	3.12	[2.29,4.21]
	0.5	*reference*								
	0.99	49.25	[30,61]	55.97	[42,74]	0.75	2.11	[1.22,2.96]	5.22	[3.88,7]
τ										
	1	*reference*								
	2	75.24	[46,98]	85.92	[56,111]	0.80	2.50	[1.57,3.3]	10.39	[6.62,13.33]
f_P										
	1	*reference*								
	3	54.30	[31,73]	65.85	[42,84]	0.77	2.35	[1.39,3.39]	6.40	[3.62,8.58]
σ										
	0.25	35.82	[23,45]	41.29	[31,53]	0.74	2.00	[1.22,2.78]	3.49	[2.25,4.83]
	0.5	*reference*								
	0.75	76.75	[43,112]	86.43	[55,122]	0.76	2.26	[1.22,3.13]	8.99	[6.12,12.17]
m										
	10	*reference*								
	20	64.77	[41,82]	72.07	[48,89]	0.72	1.74	[1.04,2.35]	5.64	[4.12,7.83]
	30	74.57	[56,95]	82.07	[60,99]	0.69	1.60	[0.7,2.17]	5.75	[4.17,7.58]

3.1 Aspects of transmission influence diffusion and production tempo, divergence and delay

Two factors related to transmission could influence variation: network architecture and the social transmission rate. In particular, node degree set the upper limit to the number of agents the behaviour could spread to from one individual. The four network types we tested were parameterized such that the mean degree for each agent was 6, however we still found slight variation in diffusion tempo and divergence between network types. While holding all other parameters at the reference level (see **Equation** (10)) random regular networks obtained the fastest tempo, followed by Erdős–Rényi networks, Barabási–Albert networks and small world networks (see **Table 2**, network architecture). We found that small-world networks obtained the least divergence between orders of acquisition and production (41.5% lower mean divergence score compared to reference), followed by Barabási–Albert, Erdős–Rényi, and random regular networks. We observed also observed minor differences in delays between acquisition and production: the shortest mean delay score was found in Barabási–Albert networks, which was 8% shorter than the longest, which was random regular.

The social transmission rate s influenced the probability of acquisition in a given time-step, thus having

consequences for diffusion tempo, as well as divergence in orders of production and transmission. Increasing the value of s increased the tempo of diffusion as behaviours took fewer time-steps to spread from knowledgeable to naïve agents (see **Table 2**, s). More interestingly, it also increased amount of variation between order of acquisition and production, with 7% more agents obtaining divergent orders under the strongest social transmission rate, when compared with the reference setting. This was because the higher s was, the more neighbours might have acquired knowledge of the novel behaviour before the next first-production event. Generally, longer periods between acquisition events lead to less divergence between orders of acquisition and production.

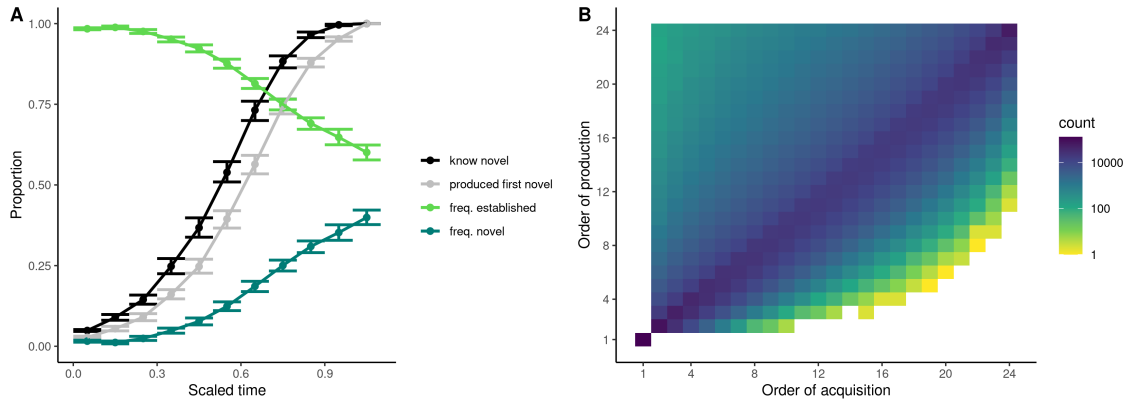

Figure 2: **Diffusion and production are separate processes. A)** Average acquisition curve of novel behaviour (black), first production curve of the novel behaviour (grey), and behavioural frequency data of the established behaviour (light green) and the novel behaviour (dark green). Data presented from the reference level (N=100 simulations, see Methods). Knowledge of the novel behaviour diffuses faster than it's relative popularity. On average, by the time the behaviour has diffused to all individuals, it is used only a quarter of the time. **B)** Density plot comparing the order in which the novel behaviour enters an agent's repertoire (x-axis) and the order in which that agent first produces the novel behaviour. Data is presented for all simulations (N=129,600 simulations), and each count is one agent from one simulation. Measuring the order of first production usually does not reflect the order of acquisition, and across all simulations, only 22% of all agents produced the novel behaviour in the same order that they acquired it.

3.2 Aspects of production influence diffusion and production tempo, divergence and delay

We turned to specifics of the production process, and analysed how they impacted acquisition and production dynamics. We first looked at parameters that influenced how social information was used in production decisions. The novel behaviour diffused the fastest and with least variation when the bias towards social information over personal information σ was weak, obtaining a TTD that was 26.7% faster than the reference setting. This result could not have been predicted without linking production and transmission, as simulating data under the same conditions using only the sub-model of production (i.e. as in **Figure S2A**) showed that increasing the social information bias hastened the time-to-first production (mean TTFP: 5.66 time-steps versus reference level of 6.34 time-steps). Conversely, a strong social information bias, i.e. $\sigma = 0.75$, slowed diffusion by 57% (**Figure 3A**; **Table 2**, σ). Increasing social information bias only slightly increased divergence between orders of acquisition and production, measured by d_{order}, by about 4% compared to the

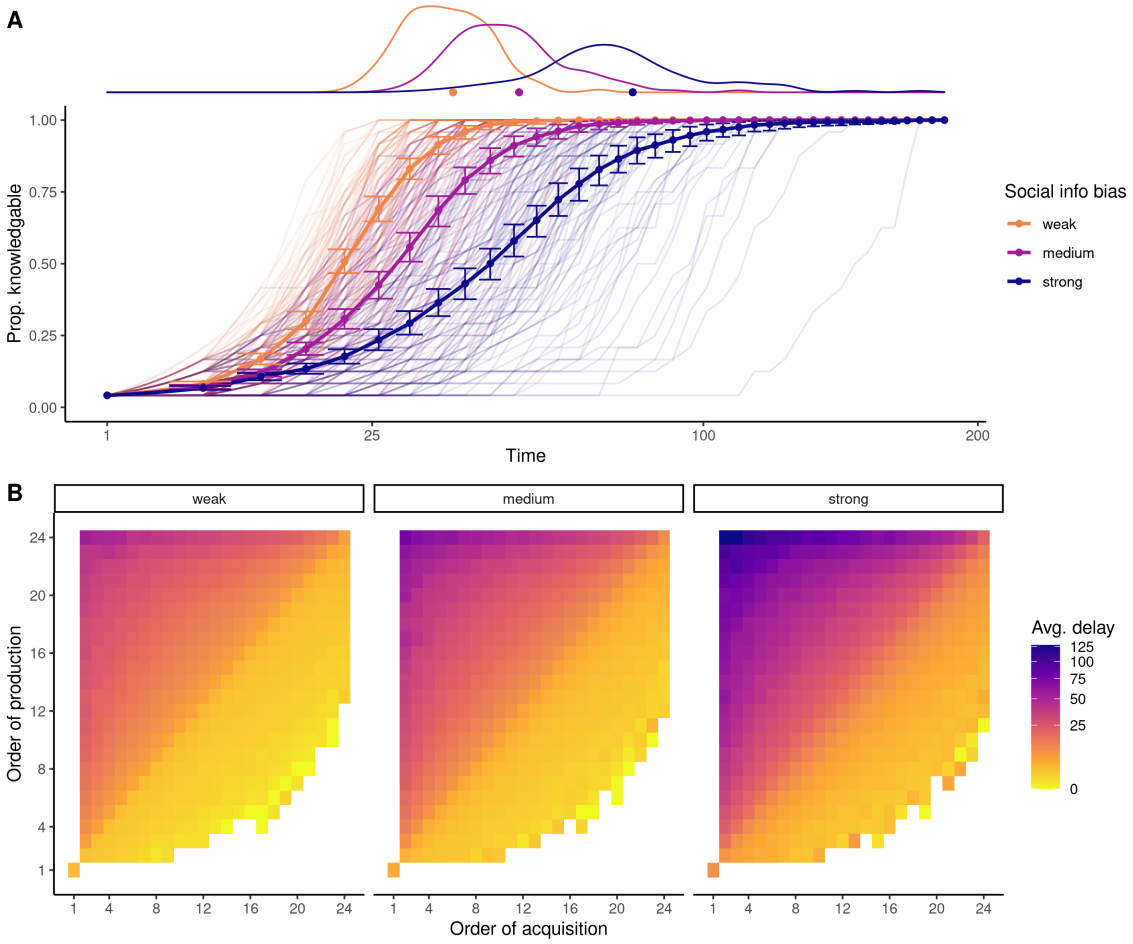

Figure 3: **Strong social information bias slows diffusion A)** Diffusion curves for the novel behaviour with distributions of time-to-diffusion in top margin, pooled within social information bias (colour, mean with 95% bootstrapped CI and traces from individual simulations). Data shown for reference level (N=300 sims) and x-axis is square-root transformed for ease of reading. Strong social information bias within the sub-model of production greatly slows the speed of diffusion. **B)** Mean delay score for all simulations (fill, square-root transformed), with each panel subset to one value of σ (N=43,200 sims each). Stronger social information bias obtains longer delays in the beginning stages of diffusion. Early learners, including the seed agent, have stronger preference for the established behaviour, but this effect is lessened for later learners, as the novel behaviour is better represented in productions by the later stage of the diffusion.

reference setting, but had less of an effect on the percentage of individuals obtaining divergent orders (1% difference from reference). However, increasing social information bias greatly increased the delay between acquisition and production, measured by d_{time}, by 72% compared to the reference setting (**Figure 3B**, **Figure S3**). When social information bias was strong, longer delays were obtained earlier in the diffusion compared to later. However, when social information bias was weak, the delay remained consistent throughout the diffusion (**Figure S3**). These results demonstrate that the slowing effect of social information bias on diffusion was caused by a feedback loop between production of the established behaviour and its social observation. A weaker bias towards social information meant that agents gave less importance to others' behaviours, increasing the probability of the production of the novel behaviour throughout the simulation, and thus hastening its diffusion.

We then explored the effects of production conformity, comparing populations in which all agents were

either linearly sensitive, or hyper sensitive to the relative frequencies of the novel behaviour. Compared to reference, production conformity increased mean TTD by 11% (**Table 2**, f_P). Production conformity increased divergence score by 8.8% compared to no production conformity, and also increased the delay score by 22%. The effects of production conformity were magnified at high values of social information bias ($\sigma = 0.75$), slowing tempo by 29% and increasing delay by 50% relative to no production conformity. Finally, relevant for both social information bias and conformity was the memory window size. We found that doubling the value of memory slowed diffusion by 32%, and tripling it slowed diffusion by 52% compared to the reference setting. Increasing memory decreased the amount of divergence in orders of acquisition and production, yet increased delays between acquisition and production (see **Table 2**, m).

Next, we considered two parameters that are relevant for personal experience. A weak bias towards recent experience resulted in diffusions that were 19.3% faster than reference (**Figure S4**; **Table 2**, ϕ). The advantage conferred by a weak bias was due to initial conditions of our simulations. Agents began with knowledge of how to produce behaviour the establish behaviour a, but had no information about it's expected payoff in their attraction scores ($AT_{ia}(0) = 0$). A weak recent experience bias slowed their learning about the payoffs of both the established behaviour *and* the novel behaviour, once acquired. We re-ran simulations with populations that initially had complete information about the established tradition's payoff ($AT_{ia}(0) = 1$). In these simulations, weak recent experience bias slowed diffusion of the novel tradition by 27% relative to reference (**Figure S4B**; mean TTD [HPDI] for $\phi_{.01}$: 205 [127, 277], $\phi_{.5}$: 161 [111, 226], $\phi_{.99}$: 160 [102, 217]). NB: this condition did not reverse the direction of the effects any other parameters. Also related to personal experience was the effect of conservatism, or hyper sensitivity to differences in expected payoff. This parameter, along with σ, created the largest effects on diffusion dynamics and divergence. Conservatism greatly decreased the tempo of diffusion by 54%, increased divergence by 16% with 80% agents obtaining divergent orders, and nearly doubled the delay between acquisition and first production (**Table 2**, τ).

Summarizing, using the modelling framework described in section 2 we have demonstrated that the processes of social transmission and production are tightly connected, and ignoring either may have consequences for the speed and order of knowledge acquisition and (first) production. Further, we showed that changes in almost any model parameter, regardless whether it affects the social transmission or production process, can influence the cultural dynamics, potentially greatly. We note that this analysis is not an exhaustive exploration of the complex interplay between social transmission and production. This would be the subject of a subsequent study.

3.3 The performance of inferential models on data generated with both transmission and production

Inferential NBDA and EWA models have been applied to data collected from wild populations many times, although always independently. However, if we assume that collected data reflects distinct and interacting processes of transmission and production, not accounting for one or the other might distort inferred parameter estimates in either model used alone. Using simulated data in which we knew the precise parameter values used to generate the data, we test the potential limitations of tempting, "vanilla" versions of both models.

We first performed a sanity check using continuous time-of-acquisition diffusion analysis (TADAc [11]), feeding it "ideal" recorded time-of-acquisition and network data from our generative model without accounting for the potential influence of production. As expected, the AICc score for a social transmission model was preferred against an asocial model for all social diffusion simulations (mean $\Delta AICc$ [HPDI]: $152[109, 197]$), and the inferential model estimated the rate of social transmission close to that which had generated the data (**Figure 4A**). The model did have trouble distinguishing between s and λ_b, and we observed a trade-off between estimates for social transmission rate ($\hat{s}$) and base transmission rate ($\hat{\lambda}_0$), as simulation variance led to differences in how strictly the behaviour strictly followed the network **Figure S5**. This issue of inexact estimation recovery is known [6, 51], and has previously led to the advice of not interpreting the exact social transmission rate, and rather rely on AICc scores and model comparison for interpretation. Under asocial learning (s constrained to 0), we found that when using ideal acquisition data, support was largely ambivalent between asocial models and social models (mean $\Delta AICc$ [HPDI]: $-1.23[-2.08, 0.39]$). In summary, our results mirrored previous results regarding the inferential properties of NBDA applied to acquisition data.

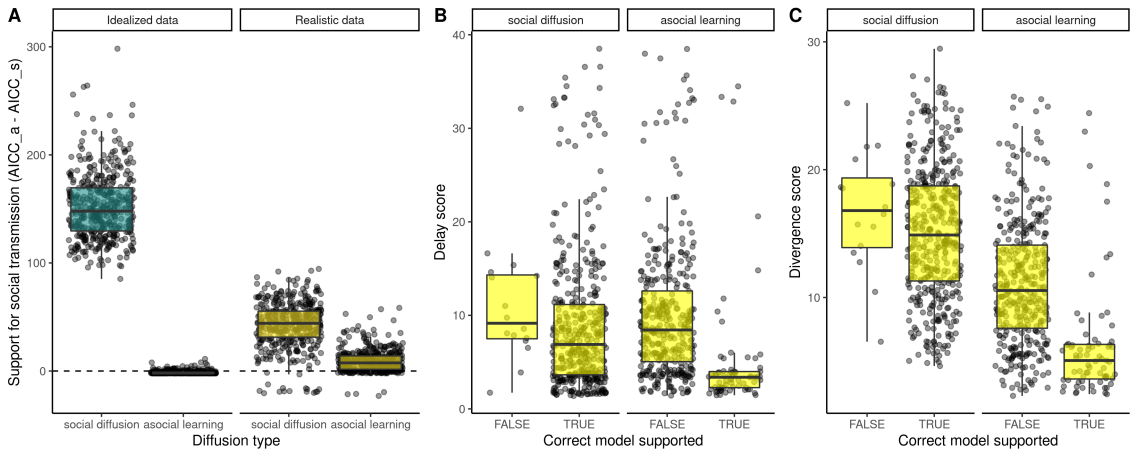

Figure 4: **Delay and divergence can produce false support in NBDA analyses.** $N = 432$ simulations per condition. **A)** Inference on idealized data (teal fill) resulted in expected support levels for both social and asocial generative processes in TADAc models. However, using time-of-first production data where production influences transmission (yellow fill) results in false support for the incorrect generative process, much more-so under conditions of asocial diffusion. Simulations that received support for the correct underlying generative process had shorter delays between acquisition and production (**B**) and less divergence between orders of acquisition and production (**C**).

We then tested TADAc on more realistic, first-production data, generated through the model developed in section 2 by allowing behavioural frequencies to inform transmission probabilities. For social transmission, the distribution of support shifted towards asocial support (mean $\Delta AICc$ [HPDI]: $42.7[10.9, 74.4]$), with 16 false-negatives supporting asocial transmission (**Figure 4A**, 2nd panel). The strongest driver of false negatives appeared to be production conformity, as it was present in 14 of 16 false negatives (**Table S1**). Likewise, the models estimated false-positive support for social transmission under the asocial learning only condition in 361 simulations (mean $\Delta AICc$ [HPDI]: $9.22[-2.08, 22.4]$). Increasing social information bias, recent experience bias and conservatism all increased the prevalence of false positives (**Table S1**). False support, positive or negative was correlated with increased mean delays between acquisition and production (**Figure 4B**), as well as more divergence between the order of acquisition and the order of first production (**Figure 4C**).

With respect to EWA inference, we first generated data from populations that began the simulation with homogeneous repertoires to make sure that reasonable parameter estimates could be recovered. The inferential model recovered point estimates of all parameters that were close to the true values used to generate the data. We then inferred parameter values from simulations that did not begin with homogeneous repertoires, where the novel behaviour either diffused socially from a single knowledgeable individual, and again where the novel behaviour was asocially learned. Estimated posterior density distributions from all three conditions are summarized in **Table S2**, and visualized per simulation in **Figure S6**.

In simulations of heterogeneous repertoires with social learning or asocial learning, the point estimates for recent experience bias ($\hat{\phi}$) were overestimated, with 89% highest posterior density intervals not crossing the true value. In both conditions, social learning bias ($\hat{\sigma}$) was overestimated, and 89% HPDIs did not include the true values. This overestimation was an artifact of agents producing what their neighbours produced while their knowledge state was constrained to the single established behaviour. These results highlight that EWA models are intended to model production of known behaviours, or choice, rather than social transmission, and high $\hat{\sigma}$ values should not be interpreted as evidence of individuals acquiring behaviours through social transmission. σ should instead be interpreted as the influence of social information, relative to personal information, on behavioural choice.

4 Discussion

In this study, we aimed to disentangle processes of transmission and production to understand their respective contribution to cultural diffusion dynamics. To do so, we developed a theoretical simulation framework that accounted for the diffusion of knowledge of a behaviour as well as its production. We illuminated dependencies between these processes, revealing how aspects of either can influence the tempo of diffusion, as well as variation between the population-level order of acquisition of knowledge and usage of behaviour. Understanding the relationship between psychological learning rules and population-level patterns is a long-standing topic of research in cultural evolution [9, 41, 52, 53]. We contribute to this understanding by highlighting the subtle difference between psychological rules that influence production *contra* transmission, thus drawing a deeper line between the acquisition of novel behaviours via social information, and the influence of social information on the production of behaviour.

Previously, social transmission and production biases have been treated as interchangeable in many discussions of social learning strategies [7, 8, 9], although their difference has been briefly mentioned in others [54]. We emphasize that biases that act on production are distinct from those that operate on transmission.The importance of this distinction can be illustrated by using the real-world example of selective attrition in passerine birdsong [55]. Juvenile birds acquire song components from adult conspecifics, meaning that there is an age-dependent bias operating on acquisition. However, as an adult, when producing it's own song, a bird might preferentially use the most frequently heard components, akin to the production conformity bias we explored. Selective attrition has been hypothesized to underlie observed changes in repertoire sizes and compositions in species with a fixed acquisition period [55, 56, 57], and cannot be ruled out in open-ended learners [58, 59]. Transmission and production biases may yield similar effects on diffusion. For example, a previous theoretical

study found that a transmission conformity bias slowed diffusion [60], echoing our finding that production conformity bias slowed diffusion. However, the effects of production biases on transmission are not necessarily predictable from analyzing production alone, as we demonstrated with the slowing effect of social information bias (σ) on diffusion only after production and transmission were linked. This discrepancy suggests that more studies are needed to further explore the consequences of making such distinctions, especially to determine whether the presence of any particular production bias might appear to be a different transmission bias.

Our findings may have consequences for the way we analyse real-world diffusion data. Currently, studies tend to either assume that 1) all individuals have uniform knowledge about a certain set of behaviours, i.e. only the production process is studied (e.g. [4, 15, 16, 42]), or that 2) an individual's acquisition of knowledge of the behaviour coincides with the behaviour's first production, i.e. the process of social transmission is studied with data on observed first production (e.g. [23, 25, 27, 28, 32]). Our results suggests that by ignoring any of these two processes we may over- or underestimate the role of transmission and production processes on observed diffusion dynamics, and we validate this by performing NBDA and EWA analyses on first production data generated through our simulation framework.

We found that NBDA could still correctly identify true positives, although with more ambivalent support. Additionally, it had much more difficulty identifying true negatives, providing model support for social transmission when asocial learning was the generative process. In general, NBDA analyses yielded the best results when there was minimal divergence between orders of acquisition and production, and minimal delay between acquisition and production. However, we demonstrated that in populations that are behaviourally conservative, or highly sensitive to social information, divergence and delay is likely high, resulting in less accurate inference results. Therefore, NBDA is perhaps best suited to to situations where the association networks are large and sparse, because this would lead to longer periods between acquisition events across the network, either between nodes or clusters of nodes, reducing overall potential divergence. Our results also strongly support the inclusion of behavioural frequency data as a transmission weight, in addition to association. Even if NBDA is blind to the processes that generate production distributions, the inclusion of this information would help mitigate their effects on inference.

Inferential EWA models have been previously used to identify social learning strategies, with the social information bias parameter (σ) described as a "social learning" bias [4, 15, 16, 17, 42]. However, these previously studied social learning biases would more accurately be referred to as production biases, as the effect of social information on production is transient, limited by the memory window of the model, and does not deal with the addition of novel behaviours to repertoire. Indeed, the EWA model was originally developed to be applied to experimental designs in which participants had homogeneous repertoires, and made choices between already known behaviours where this assumption was satisfied (i.e. all choices are explicitly presented in a decision environment [7, 17, 37, 43]). Our results show that time spent in heterogeneous knowledge states can cause mis-estimation of parameters, and that EWA alone cannot distinguish between social diffusion and asocial learning of a novel behaviour. We highlight this because the label of "social learning" has been given both social transmission and social influence, which could lead researchers who are less familiar with the EWA framework to draw inaccurate conclusions based on the interpretation of the σ parameter. In light of our results, we believe that EWA might be of greater use in situations where one can be relatively certain of

homogeneous knowledge states. This includes smaller social groups, where the association network is more homogeneous and diffusion might occur rapidly relative to behavioural production rate. If possible, data recorded during the active diffusion of behaviours could be excluded. Task structure should also be taken into account, as open ended tasks with a large range of possible solutions of varying difficulty might result in more heterogeneity than closed tasks, and some tasks might be more susceptible to morphological constraints on production, such as the case where juveniles' jaws were simply too small or weak to perform some variants in a foraging task [15].

In summary, we render a clear distinction between acquisition and expression of behaviour, and thus social learning and social influence. These distinctions are critical for explaining patterns of cultural change and cultural polymorphisms [61], since individuals might know multiple variants and yet only produce a subset as a function of their own experiences and biases. The acquisition and performance of novel behaviours exact opportunity costs, and the decision to express a particular behaviour at a particular time competes alongside every other decision an individual has to make, with consequences for what behaviours are maintained in the repertoire and which are discarded [2, 39]. These individual-level decisions aggregate and percolate, influencing cultural diffusion dynamics at the population level. Finally, our study emphasizes the importance of collecting detailed production data beyond the first observed production of an individual, as well as recording social networks. We also suggest caution when applying inferential methods blind to either the processes of social transmission or production to observational data. We hope this work will encourage further development of new statistical methods to understand data generated through interacting processes of transmission and production.

5 Data Availability

Code and data for statistical analyses and main text figures, as well as code to replicate the agent based models are available at https://github.com/michaelchimento/transmission_production_abm. A DOI will be assigned upon acceptance.

6 Acknowledgements

M.C. thanks Sonja Wild for discussions about NBDA models, and Richard McElreath for discussions about EWA models. M.C. received funding from the International Max Planck Research School for Organismal Biology. B.J.B. was supported by the Max Planck Society and the Alexander von Humboldt Foundation in the framework of the Alexander von Humboldt Professorship endowed by the Federal Ministry of Education and Research awarded to Margaret C. Crofoot. L.M.A was funded by a Max Planck Research Group Leader Fellowship. L.M.A. and M.C. were supported by the Max Planck Society and the Centre for the Advanced Study of Collective Behaviour, funded by the Deutsche Forschungsgemeinschaft (DFG) under Germany's Excellence Strategy (EXC 2117-422037984).

7 Author Contributions

Conceptualisation, M.C., B.B., A.K., and L.M.A.; Methodology, M.C., B.B., A.K.; Software, M.C. (agent based model, analysis), B.B. (Stan); Investigation, M.C.; Resources, M.C., B.B.; Writing - Original Draft, M.C.; Writing - Review & Editing, M.C., B.B., A.K., L.M.A.; Visualization, M.C.; Supervision, L.M.A.; Funding Acquisition, L.M.A.

8 Competing Interests

The authors declare no competing interests, financial or otherwise.

9 Materials & Correspondence

Correspondence should be addressed to Michael Chimento, e-mail: mchimento@ab.mpg.de

References

[1] Andrew Whiten et al. "The extension of biology through culture". In: *Proceedings of the National Academy of Sciences* 114.30 (2017), pp. 7775–7781.

[2] Kevin N Laland, Wataru Toyokawa, and Thomas Oudman. "Animal learning as a source of developmental bias". In: *Evolution & development* 22.1-2 (2020), pp. 126–142.

[3] Eva Jablonka and Marion J Lamb. *Evolution in four dimensions, revised edition: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. MIT press, 2014.

[4] Lucy M Aplin, Ben C Sheldon, and Richard McElreath. "Conformity does not perpetuate suboptimal traditions in a wild population of songbirds". In: *Proceedings of the National Academy of Sciences* 114.30 (2017), pp. 7830–7837.

[5] Simon Kirby, Tom Griffiths, and Kenny Smith. "Iterated learning and the evolution of language". In: *Current opinion in neurobiology* 28 (2014), pp. 108–114.

[6] William Hoppitt and Kevin N Laland. *Social learning: an introduction to mechanisms, methods, and models*. Princeton University Press, 2013.

[7] Richard McElreath et al. "Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 363.1509 (2008), pp. 3515–3528.

[8] Luke Rendell et al. "Cognitive culture: theoretical and empirical insights into social learning strategies". In: *Trends in cognitive sciences* 15.2 (2011), pp. 68–76.

[9] Rachel L Kendal et al. "Social learning strategies: Bridge-building between fields". In: *Trends in cognitive sciences* 22.7 (2018), pp. 651–665.

[10] Damien R Farine et al. "The role of social and ecological processes in structuring animal populations: a case study from automated tracking of wild birds". In: *Royal Society Open Science* 2.4 (2015), p. 150057.

[11] Matthew J Hasenjager, Ellouise Leadbeater, and William Hoppitt. "Detecting and quantifying social transmission using network-based diffusion analysis". In: *Journal of Animal Ecology* 90.1 (2021), pp. 8–26.

[12] Esther Sebastián-González and Patrick J Hart. "Birdsong meme diversity in a habitat landscape depends on landscape and species characteristics". In: *Oikos* 126.10 (2017), pp. 1511–1521.

[13] T Gruber et al. "Efficiency fosters cumulative culture across species". In: *Philosophical Transactions of the Royal Society B* 377.1843 (), p. 20200308.

[14] Alecia J Carter, Miquel Torrents Tico, and Guy Cowlishaw. "Sequential phenotypic constraints on social information use in wild baboons". In: *Elife* 5 (2016), e13125.

[15] Brendan J Barrett, Richard L McElreath, and Susan E Perry. "Pay-off-biased social learning underlies the diffusion of novel extractive foraging traditions in a wild primate". In: *Proceedings of the Royal Society B: Biological Sciences* 284.1856 (2017), p. 20170358.

[16] Rachel Kendal et al. "Chimpanzees copy dominant and knowledgeable individuals: implications for cultural diversity". In: *Evolution and Human Behavior* 36.1 (2015), pp. 65–72.

[17] Dominik Deffner, Vivien Kleinow, and Richard McElreath. "Dynamic social learning in temporally and spatially variable environments". In: *Royal Society Open Science* 7.12 (2020), p. 200734.

[18] Mathias Franz and Charles L Nunn. "Network-based diffusion analysis: a new method for detecting social learning". In: *Proceedings of the Royal Society B: Biological Sciences* 276.1663 (2009), pp. 1829–1836.

[19] William Hoppitt, Neeltje J Boogert, and Kevin N Laland. "Detecting social transmission in networks". In: *Journal of Theoretical Biology* 263.4 (2010), pp. 544–555.

[20] Lucy M Aplin et al. "Experimentally induced innovations lead to persistent culture via conformity in wild birds". In: *Nature* 518.7540 (2015), p. 538.

[21] Damien R Farine, Karen A Spencer, and Neeltje J Boogert. "Early-life stress triggers juvenile zebra finches to switch social learning strategies". In: *Current Biology* 25.16 (2015), pp. 2184–2188.

[22] Virginia K Heinen et al. "Food discovery is associated with different reliance on social learning and lower cognitive flexibility across environments in a food-caching bird". In: *Proceedings of the Royal Society B* 288.1951 (2021), p. 20202843.

[23] Barbara C Klump et al. "Innovation and geographic spread of a complex foraging culture in an urban parrot". In: *Science* 373.6553 (2021), pp. 456–460.

[24] Mike M Webster et al. "Environmental complexity influences association network structure and network-based diffusion of foraging information in fish shoals". In: *The American Naturalist* 181.2 (2013), pp. 235–244.

[25] Matthew J Hasenjager, William Hoppitt, and Lee A Dugatkin. "Personality composition determines social learning pathways within shoaling fish". In: *Proceedings of the Royal Society B* 287.1936 (2020), p. 20201871.

[26] Jenny Allen et al. "Network-based diffusion analysis reveals cultural transmission of lobtail feeding in humpback whales". In: *Science* 340.6131 (2013), pp. 485–488.

[27] Sonja Wild et al. "Multi-network-based diffusion analysis reveals vertical cultural transmission of sponge tool use within dolphin matrilines". In: *Biology letters* 15.7 (2019), p. 20190227.

[28] Sonja Wild et al. "Integrating genetic, environmental, and social networks to reveal transmission pathways of a dolphin foraging innovation". In: *Current Biology* 30.15 (2020), pp. 3024–3030.

[29] Rachel L Kendal et al. "Evidence for social learning in wild lemurs (Lemur catta)". In: *Learning & Behavior* 38.3 (2010), pp. 220–234.

[30] Catherine Hobaiter et al. "Social network analysis shows direct evidence for social transmission of tool use in wild chimpanzees". In: *PLoS Biol* 12.9 (2014), e1001960.

[31] Stuart K Watson et al. "Socially transmitted diffusion of a novel behavior from subordinate chimpanzees". In: *American journal of primatology* 79.6 (2017), e22642.

[32] Charlotte Canteloup, William Hoppitt, and Erica van de Waal. "Wild primates copy higher-ranked individuals in a social transmission experiment". In: *Nature communications* 11.1 (2020), pp. 1–10.

[33] Josh A Firth et al. "Analysing the Social Spread of Behaviour: Integrating Complex Contagions into Network Based Diffusions". In: *arXiv preprint arXiv:2012.08925* (2020).

[34] Will Hoppitt. "The conceptual foundations of network-based diffusion analysis: choosing networks and interpreting results". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 372.1735 (2017), p. 20160418.

[35] Ipek G Kulahci et al. "Social networks predict selective observation and information spread in ravens". In: *Royal Society open science* 3.7 (2016), p. 160256.

[36] Zoltán Tóth et al. "The effect of social connections on the discovery of multiple hidden food patches in a bird species". In: *Scientific reports* 7.1 (2017), pp. 1–9.

[37] Colin Camerer and Teck Hua Ho. "Experience-weighted attraction learning in normal form games". In: *Econometrica* 67.4 (1999), pp. 827–874.

[38] Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games". In: *Journal of Economic theory* 104.1 (2002), pp. 137–188.

[39] Bennett G Galef Jr. "Why behaviour patterns that animals learn socially are locally adaptive". In: *Animal behaviour* 49.5 (1995), pp. 1325–1334.

[40] Michal Arbilly et al. "Evolution of social learning when high expected payoffs are associated with high risk of failure". In: *Journal of the Royal Society Interface* 8.64 (2011), pp. 1604–1615.

[41] Brendan J Barrett. "Equifinality in empirical studies of cultural transmission". In: *Behavioural processes* 161 (2019), pp. 129–138.

[42] Charlotte Canteloup et al. "Processing of novel food reveals payoff and rank-biased social learning in a wild primate". In: *Scientific Reports* 11.1 (2021), pp. 1–13.

[43] Richard McElreath et al. "Applying evolutionary models to the laboratory study of social learning". In: *Evolution and Human Behavior* 26.6 (2005), pp. 483–508.

[44] Wataru Toyokawa, Andrew Whalen, and Kevin N Laland. "Social learning strategies regulate the wisdom and madness of interactive crowds". In: *Nature Human Behaviour* 3.2 (2019), pp. 183–193.

[45] Sabine Coussi-Korbel and Dorothy M Fragaszy. "On the relation between social dynamics and social learning". In: *Animal behaviour* 50.6 (1995), pp. 1441–1453.

[46] Duncan J Watts and Steven H Strogatz. "Collective dynamics of 'small-world'networks". In: *nature* 393.6684 (1998), pp. 440–442.

[47] Mauricio Cantor et al. "Social network architecture and the tempo of cumulative cultural evolution". In: *Proceedings of the Royal Society B* 288.1946 (2021), p. 20203107.

[48] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria, 2018. URL: https://www.R-project.org/.

[49] Stan Development Team. *Stan Modeling Language Users Guide and Reference Manual*. Version 2.27. 2021. URL: http://mc-stan.org/.

[50] Stan Development Team. *RStan: the R interface to Stan*. R package version 2.21.2. 2020. URL: http://mc-stan.org/.

[51] Andrew Whalen and William JE Hoppitt. "Bayesian model selection with network based diffusion analysis". In: *Frontiers in psychology* 7 (2016), p. 409.

[52] Robert Boyd and Peter J Richerson. *Culture and the evolutionary process*. University of Chicago press, 1988.

[53] Thomas JH Morgan et al. "The evolutionary basis of human social learning". In: *Proceedings of the Royal Society B: Biological Sciences* 279.1729 (2012), pp. 653–662.

[54] Kevin N Laland. "Social learning strategies". In: *Animal Learning & Behavior* 32.1 (2004), pp. 4–14.

[55] Peter Marler and Susan Peters. "Developmental overproduction and selective attrition: new processes in the epigenesis of birdsong". In: *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology* 15.4 (1982), pp. 369–378.

[56] Douglas A Nelson. "Song overproduction, selective attrition and song dialects in the white-crowned sparrow". In: *Animal Behaviour* 60.6 (2000), pp. 887–898.

[57] Robert F Lachlan, Oliver Ratmann, and Stephen Nowicki. "Cultural conformity generates extremely stable traditions in bird song". In: *Nature communications* 9.1 (2018), pp. 1–9.

[58] Joanne S Nicholson et al. "Song sharing and repertoire size in the sedge warbler, Acrocephalus schoenobaenus: changes within and between years". In: *Animal Behaviour* 74.5 (2007), pp. 1585–1592.

[59] Louis Ranjard et al. "Geographic patterns of song variation reveal timing of song acquisition in a wild avian population". In: *Behavioral Ecology* 28.4 (2017), pp. 1085–1092.

[60] Joseph Henrich. "Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change". In: *American Anthropologist* 103.4 (2001), pp. 992–1013.

[61] Ofer Tchernichovski et al. "How social learning adds up to a culture: from birdsong to human public opinion". In: *Journal of experimental biology* 220.1 (2017), pp. 124–132.

Supplementary Materials

Tables

- **Table S1**: Contributions of parameters to false negatives and positives in NBDA inference
- **Table S2**: Posterior distributions from EWA inference on simulated data

Figures

- **Figure S1**: Exemplar networks.
- **Figure S2**: Simulation dynamics when production and transmission are disconnected.
- **Figure S3**: Exploration of delay between acquisition and production across social information biases.
- **Figure S4**: The effect of recent experience bias depends on initial conditions.
- **Figure S5**: Parameter estimates from NBDA analyses of simulated data
- **Figure S6**: Posterior density estimates by simulation from EWA inference.

Table S1: **Contributions of parameters to false negatives and positives in NBDA inference**. Summary of the number of simulations (out of 432) resulting in either false negative ($AICc_{asocial} - AICc_{social} \leq 0$ when data was generated with social transmission) or false positive support ($AICc_{asocial} - AICc_{social} > 0$ when data was generated with asocial learning), according to relative AICc scores.

value	false negatives	false positives
graph type		
random regular	4	93
small world	2	93
Erdős–Rényi	3	92
Barabási–Albert	7	83
Social info. bias		
weak	4	104
medium	5	120
strong	7	137
Recent exp. bias		
weak	6	98
medium	7	129
strong	3	134
Conservatism		
non-conservative	7	172
conservative	9	189
Production conformity		
1	2	181
3	14	180
Memory window		
10	3	123
20	3	116
30	10	122

Table S2: **Posterior distributions from EWA inference on simulated data**. Posterior estimates for EWA parameters from three types of simulations (column headers), with each row containing the true parameter value used to generate data, and the estimated mean and 89% highest probability density intervals. In "homogeneous" simulations, all agents began with knowledge of both behaviours, without active diffusion of the novel behaviour. Here, parameter estimates were reasonably matched the true values used to simulate the data. In "heterogeneous social learning" simulations, agents began with knowledge of the establish behaviour, while a novel behaviour diffused via social learning. Recent experience bias (ϕ) and social information bias (σ) were over-estimated. In "Heterogeneous asocial learning" simulations, the behaviour diffused only through asocial learning. Again, here ϕ and σ are over-estimated.

	Homogenous		Heterogenous: social learning		Heterogenous: asocial learning	
true value	est.	HPDI	est.	HPDI	est.	HPDI
phi						
0.25	0.32	[0.2,0.45]	0.64	[0.42,0.87]	0.55	[0.3,0.8]
0.50	0.50	[0.38,0.62]	0.73	[0.55,0.91]	0.71	[0.55,0.88]
0.75	0.74	[0.63,0.86]	0.80	[0.66,0.95]	0.85	[0.75,0.95]
sigma						
0.25	0.27	[0.2,0.35]	0.64	[0.57,0.71]	0.55	[0.42,0.66]
0.50	0.51	[0.43,0.6]	0.80	[0.75,0.86]	0.74	[0.68,0.81]
0.75	0.76	[0.7,0.83]	0.92	[0.88,0.96]	0.88	[0.83,0.93]

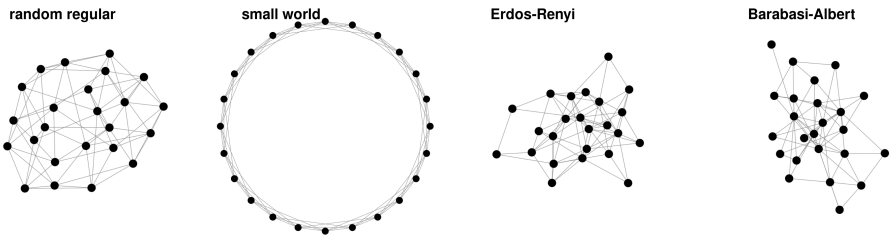

Figure S1: **Exemplar networks.** Four types of network architectures used for simulations. Random regular and small world networks contain nodes which each have 6 connections to neighbours. Erdős–Rényi and Barabási–Albert networks were parameterized such that the mean degree was 6 to make results comparable.

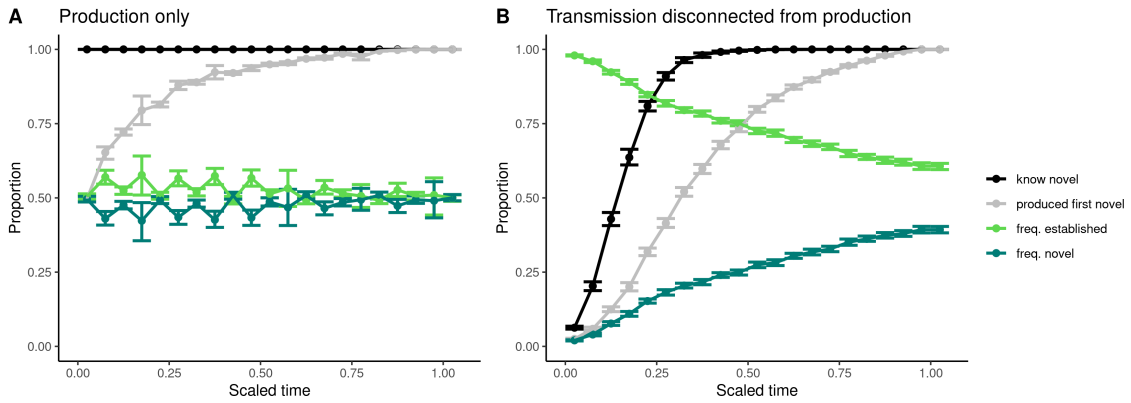

Figure S2: **Simulation dynamics when production and transmission are disconnected.** Mean values (with 95% bootstrapped CI) over time (scaled between 0 and 1) of the proportion of agents which know the novel behaviour (black), the proportion of agents which have produced the novel behaviour at least once (grey), the proportion of behaviours that are the established behaviour (light green), and the novel behaviour (dark green). Parameters set at reference level (N=100 sims for each panel). **A)** Generating data with the sub-model of production only results in an r-shaped curve of first production, with behavioural frequencies oscillating around 50%. The production sub-model is not appropriate for approximating diffusion. **B)** Generating data with the sub-model of transmission disconnected from production results in a diffusion curve that is completely disconnected from the first production curve, as agents can learn the novel behaviour without ever having seen it produced. The first production curve is only constrained by the transmission curve very early in the simulation, and more resembles the r-shaped curve in panel A.

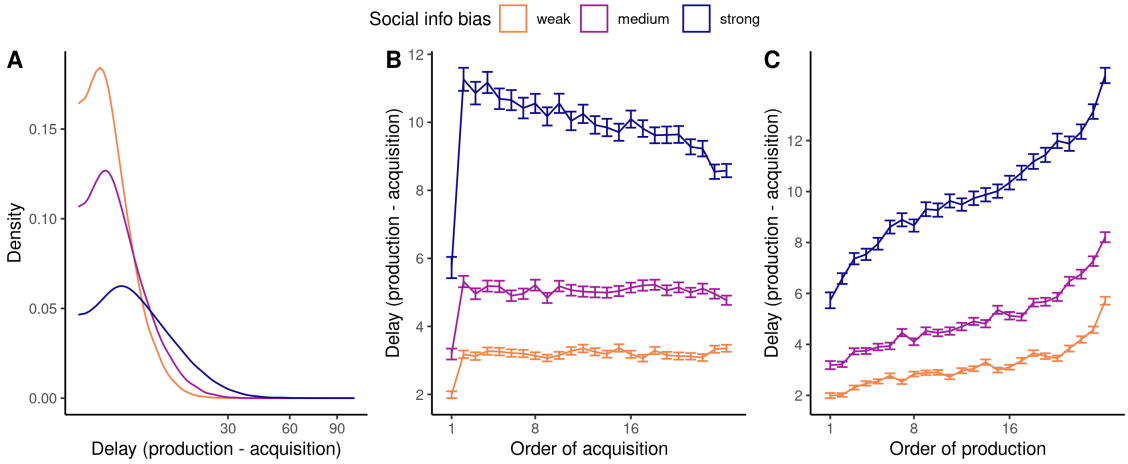

Figure S3: **Exploration of delay between acquisition and production across social information biases. A)** Comparison between social information bias values (colour) of densities of delay values found in agents. X-axis is sqrt transformed for ease of reading. A weak social information bias obtained the highest density of short delays, followed by medium and strong biases. **B)** Mean time difference between acquisition and production across σ values. When social information bias was strong, the seed agent (order of acquisition=1) took the longest to first produce the novel behaviour, and there was also a decreasing relationship between order of acquisition and delays. Delays remained relatively constant over the order of acquisition under weak and medium biased learners. This was because the relative proportion of the novel behaviour would increase over simulation time, and mattered the most for agents which were strongly sensitive to socially observed behavioural frequencies. **C)** Delay values showed similar increasing relationship with order of first production across all 3 levels of bias towards social information. The seed agent was always the fastest to first produce the novel behaviour, and the final agent to produce the behaviour always obtained the longest time difference between acquisition and production.

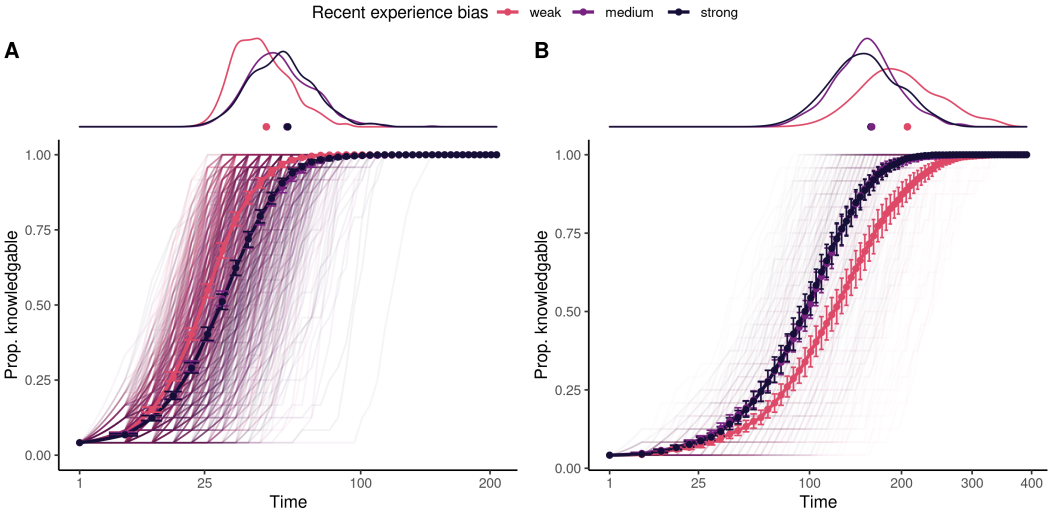

Figure S4: **The effect of recent experience bias depends on initial conditions. A)** Acquisition curves pooled within recent experience bias (colour, mean with 95% bootstrapped CI and traces from individual simulations). When agents have minimal knowledge of the payoffs of the established behaviour, a weak recent experience bias hastens diffusion of novel behaviours. This is because the weak bias prevents agents from learning too much about the established tradition before the novel tradition begins to diffuse. **B)** When agents have full knowledge about the payoffs of the established behaviour, a weak recent experience bias slows diffusion of novel behaviours.

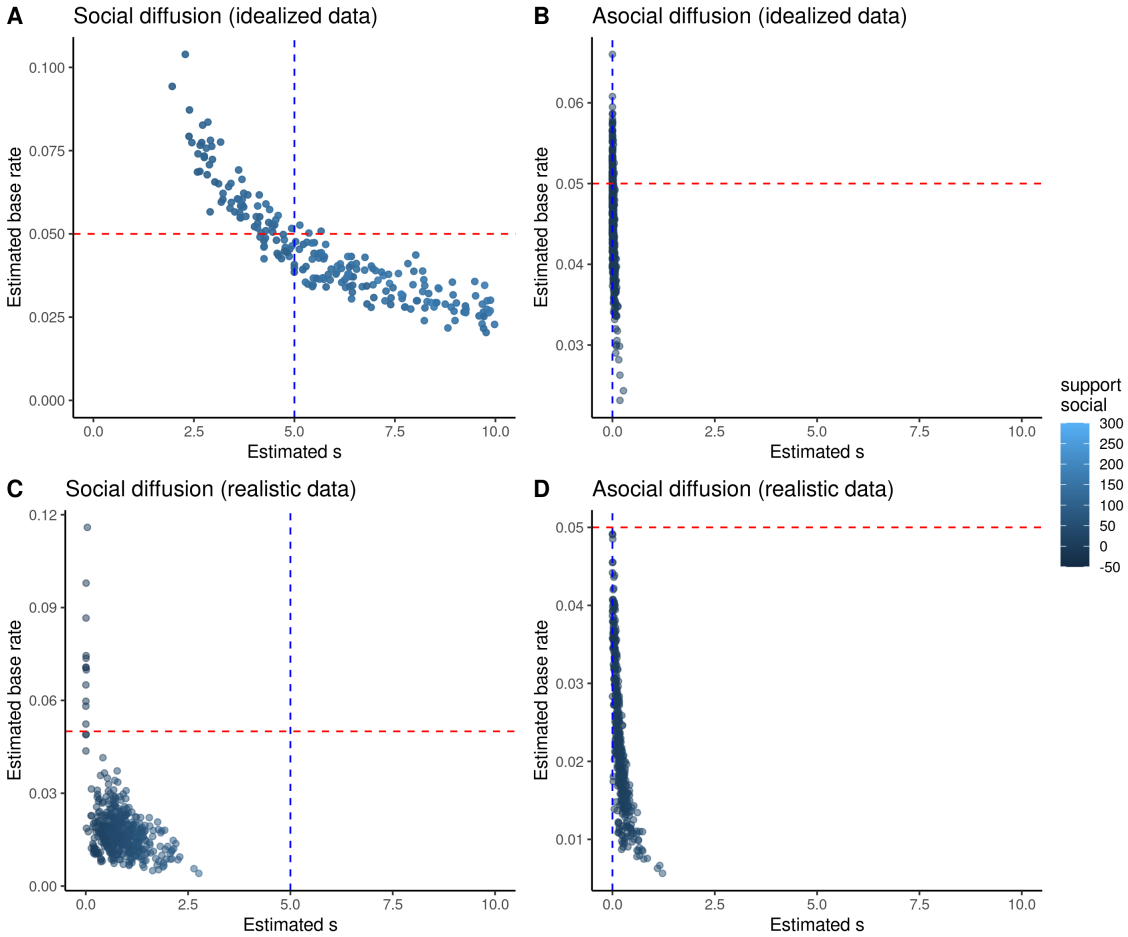

Figure S5: **Parameter estimates from NBDA analyses of simulated data** Each point represents parameter estimates for one simulation: rate of social transmission (x-axis) and baseline transmission rate values (y-axis). colour represents the support given to social transmission (AICc social model - AICc asocial only model). Dashed lines represent true values of parameters used to generate data. **A)** Ideal data from a social diffusion results in correct inference, although simulation variance leads to the curved shape of parameter estimates. If the diffusion follows the network closely, estimated strenghth of social transmission is inflated. **B)** Idealized asocial diffusion also results in correct inference, with very low s estimates and little support for social transmission. **C & D)** Realistic data is generated with transmission and production sub-models interacting. Parameters are mis-estimated in both cases, with many more false negative and positives than the idealized cases.

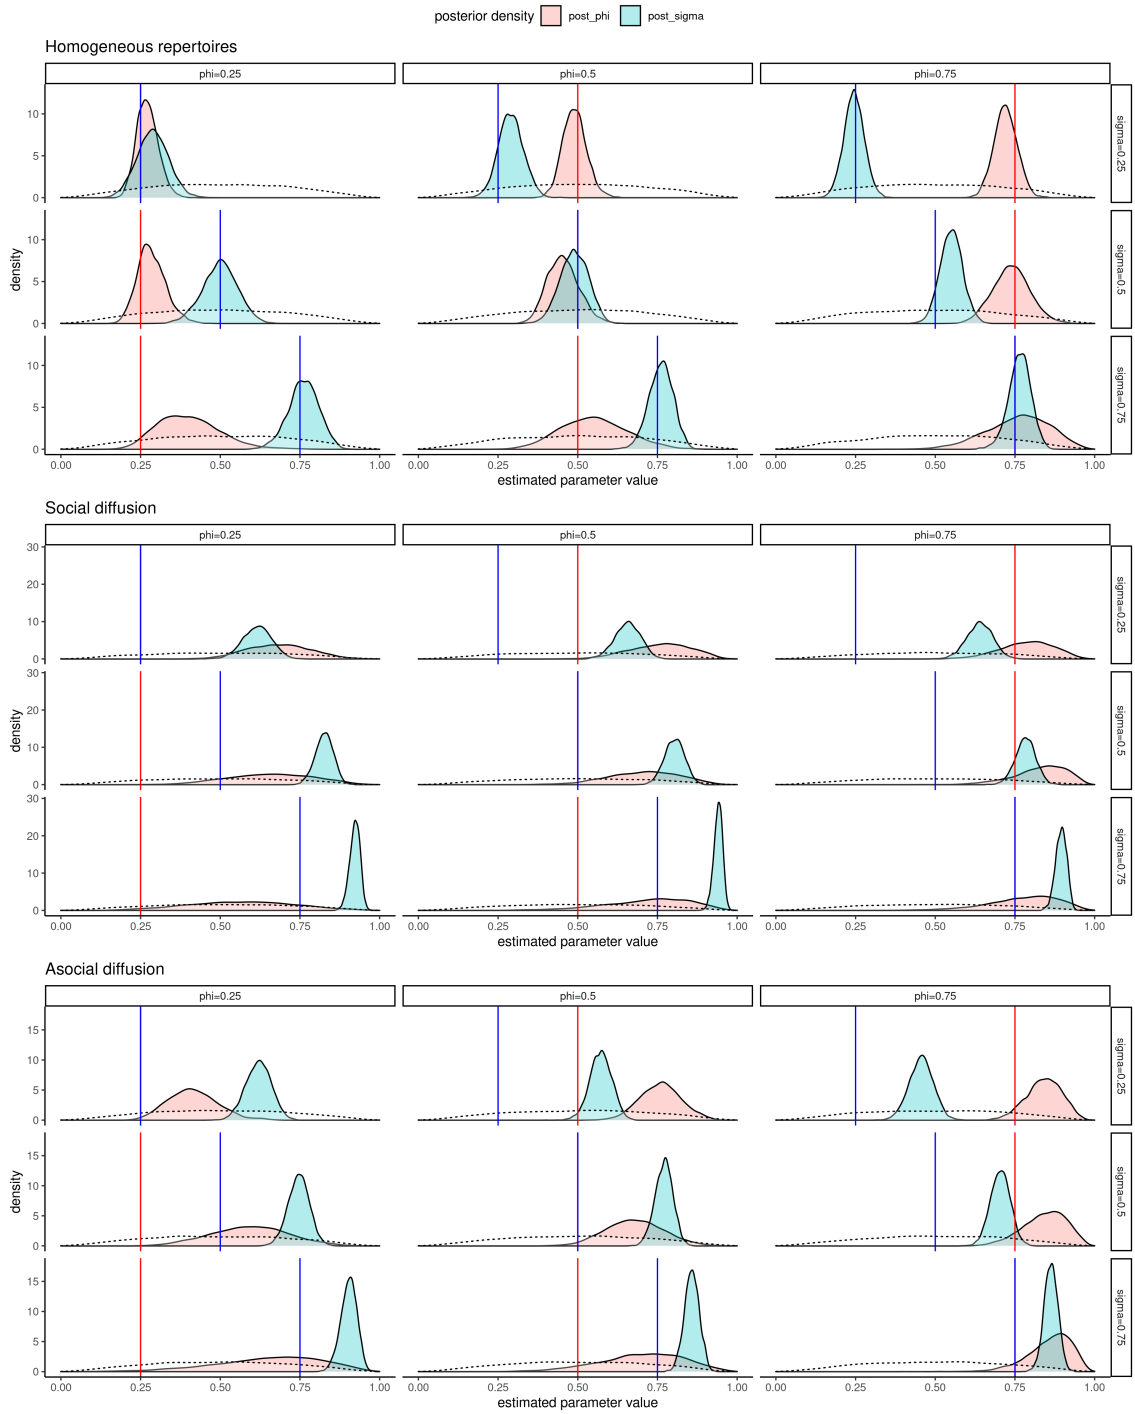

Figure S6: **Posterior density estimates for parameters from EWA inference.** Each panel shows the inferred posterior density estimates for recent experience bias (ϕ) and social information bias (σ) for one simulation (N=1500 draws from posterior per parameter). The distribution with the dashed line represents the prior for both variables. True values used to simulate the data are shown with vertical lines, coloured by parameter. Inference on data simulated with homogeneous repertoires results in parameter estimates consistent with true values. Inference on data simulated with social or asocial diffusion results in general over estimation of both parameters.