Do Sigmoidal Acquisition Curves Indicate Conformity?

Paul E. Smaldino[1,+,*], **Lucy M. Aplin**[2,+], **and Damien R. Farine**[2,3,4,+]

[1]Cognitive and Information Sciences, University of California, Merced, CA 95340, U.S.A.
[2]Edward Grey Institute of Field Ornithology, University of Oxford, Oxford OX1 3PS, U.K.
[3]Department of Collective Behaviour, Max Planck Institute for Ornithology, Konstanz 78457, Germany
[4]Chair of Biodiversity and Collective Behaviour, Department of Biology, University of Konstanz, Konstanz 78457, Germany
[*]corresponding author: paul.smaldino@gmail.com
[+]all authors contributed equally to this work

ABSTRACT

The potential for behaviours to spread via cultural transmission has profound implications for our understanding of social dynamics and evolution. Several studies have provided empirical evidence that local traditions can be maintained in animal populations via conformist learning (i.e. copying the majority). A conformist bias can be characterized by a sigmoidal relationship between a behavior's prevalence in the population and an individual's propensity to adopt that behavior. For this reason, the presence of conformist learning in a population is often inferred from a sigmoidal acquisition curve in which the overall rate of adoption for the behavior is taken as the dependent variable. However, the validity of sigmoidal acquisition curves as evidence for conformist learning has recently been challenged by models suggesting that such curves can arise via alternative learning rules that do not involve conformity. We review these models, and find that, in every case, the proposed alternative learning mechanisms relied on faulty or unrealistic assumptions. We therefore recommend that sigmoidal acquisition curves continue to be taken as evidence for conformist learning. Our paper also highlights the importance of understanding the generative processes of a model, rather than only focusing solely on the patterns produced. By studying these processes, our analysis suggests that current practices by empiricists have provided robust evidence for conformist transmission in both humans and non-human animals.

Introduction

There is increasing evidence that many animals, including humans, use social learning strategically; preferentially copying certain individuals, behaviours, or in particular contexts, in order to optimise the quality of information gained[1,2]. These learning rules will not only determine what behaviours individuals acquire, but will also shape transmission dynamics as information moves through populations, with implications for the formation and patterning of traditions[1,3–5]. One such social learning strategy is conformist learning, where individuals acquire behaviour with a positive frequency-dependent learning bias, preferentially copying the most common behaviour or the majority of individuals[6,7]. Conformity has been of particular interest in study of cultural evolution, as it is considered one of the central pillars in the emergence of culture[8]. That is, if individuals copy the majority, groups should fix on a single socially learnt tradition that will remain stable over time and be resistant to erosion or invasion by alternative variants, allowing for group specific cultures to establish and persist. Once thought of as a potentially 'human only' trait[6], numerous models have now argued that conformity could potentially evolve in various species under a wide range of environmental conditions[7,9]. For example, a conformist learning strategy can be adaptive in spatially varying environments where local knowledge is costly to acquire[3,8,10], as copying the behaviour of the majority of residents will thus be a short-cut to ascertaining best locally adapted behaviour. However, experimentally demonstrating the use of conformist learning in non-human animals has been challenging, particularly in natural populations. A key question has thus become whether conformist learning strategies can be inferred from behavioural data on the acquisition and expression of learnt behaviour.

Recently, two studies provided evidence for conformity in very different systems. van de Waal et al. (2013)[11] trained troops of vervet monkeys (*Chlorocebus aethiops*) to prefer one of two colours of otherwise identical corn. They found that naïve monkeys adopted the majority behaviour of their group, and that dispersing individuals would switch their choice to match the majority behaviour in the group that they joined. Aplin et al. (2015)[12] trained great tits (*Parus major*) to access a food reward by solving a puzzle box by sliding a door either left or right, where puzzle boxes could be opened in either direction. They then observed the spread of this new behaviour from these trained demonstrators across five replicated sub-populations and

over two generations. Naïve birds disproportionately copied the most common solving technique (opening direction) in their local foraging group, with traditions for left or right persisting across generations. Similarly to the vervet monkeys[11], birds that moved between sub-populations with different traditions tended to switch their preference to match the locally common behaviour of the area they immigrated into.

These studies reflect the major lines of evidence used to suggest conformist biases in non-human animals. Firstly, that individuals will change their existing behaviour when faced with a conflict between their preferences and those of the rest of their group, and secondly, that the probability of a naïve individual adopting the most common behaviour will be disproportionately greater than that behaviour's frequency in the group it observes. This second line of evidence can be characterized by a sigmoidal relationship between an individual's probability of acquiring a behaviour and the frequency of its demonstration. When individuals can choose between two behaviours (sometimes including non-adoption as a behaviour), such a sigmoidal relationship indicates how individuals become disproportionately biased towards the common behaviour[3]. Scaling up, the presence of such a sigmoidal acquisition curve at the *population* level—in which the relationship between the frequency of a behavior and the overall rate at which it is adopted is sigmoidal—is generally thought to provide evidence of conformist learning in the population in question[⋆].

Yet is this approach robust? In other words, does a sigmoidal acquisition curve necessarily indicate conformist learning? Several recent papers have argued that the answer is a firm 'no'[14–16]. These critiques center on the claim that sigmoidal curves can be generated without conformist learning, and therefore that the presence of such curves is not evidence for conformist learning. In the remainder of this paper, we evaluate these claims.

The evidence for conformist learning in wild populations of non-human primates and birds was first challenged by van Leeuwen et al. (2015)[14]. In this paper, they disputed the methodological validity of using the frequency at which a behaviour is being demonstrated when testing for a conformist bias (for example, as in Aplin et al. 2015a[12]). Instead they argued that a valid test for conformity must compare the probability of adopting behaviours against the overall proportion of individuals performing each behaviour; for example, if a demonstrator performs a behaviour twice, its contribution to the overall frequency of behaviour in a group is only counted once. In response to this critique, Aplin et al. (2015b)[17] reanalysed their previous data to show that their results were robust to either definition of conformist transmission. Aplin et al.[17] further argued that, in empirical practice, there were no cases in which the two definitions conflicted, suggesting that they are functionally equivalent in all but the most exceptional of cases[17]. This argument was further supported by re-analyses of data from van de Waal et al. (2013) by Whiten and colleagues[18].

In two more recent papers, van Leeuwen et al. (2016)[15] and Acerbi et al. (2016)[16] have issued a more direct challenge to the validity of identifying conformist transmission from population-level sigmoidal acquisition curves. In these papers, they developed a simple computational model in which populations of 100 individuals were randomly initialized with one of two behaviours. Individuals were then allowed to either copy a random demonstrator[15] or choose to copy a demonstrator as determined by a range of different learning rules[16]. In several cases, their modelling results suggested that plotting the frequency of the behaviour against the probability of adoption could lead to sigmoidal acquisition curves even in the absence of conformist learning. In particular, Acerbi et al. (2016)[16] purported to demonstrate two non-conformist learning rules that nevertheless generate conformist curves: one in which individuals have an intrinsic bias for one behaviour ('variant preference'), another in which individuals copy a behaviour from a randomly chosen member of a small subset of the population ('demonstrators subgroup').

Here, we revisit the modelling approaches used by van Leeuwen et al (2015, 2016)[14,15] and Acerbi et al (2016)[16]. Specifically, we reanalyse the models in two regards. First, we examine the evidence that proposed non-conformist learning mechanisms can yield sigmoidal curves. We then consider the argument that the majority should be represented as the frequency of individuals exhibiting a behaviour, rather than by the frequency of the behaviour itself. This reanalysis highlights that in both cases, a number of problematic assumptions and implementation errors distort or exaggerate the results produced. Our aim in this paper is therefore to correct the overly strong claims made in those papers. We conclude that, with proper methodological implementation, a sigmoidal relationship between the adoption of a behaviour and its frequency in the group remains a valid and robust way to identify conformist social learning in wild populations.

Results

Do non-conformist social learning strategies generate sigmoidal curves?

Acerbi et al. (2016)[16] used an individual-based model to investigate when a sigmoidal relationship arises between the proportional frequency of a behavioural variant A and the probability of that a random individual will adopt A over a competing

⋆It is important to differentiate *acquisition curves* from *diffusion curves*, as both are common topics in the study of social learning. A diffusion curve is the instantaneous or cumulative rate of adopting a behavior as a function of time[13]. Our focus is on the acquisition curve, which is the instantaneous rate of adopting a behavior as a function of that behavior's prevalence in the population.

variant B. The authors examined 10 possible learning biases, and concluded that two non-conformist learning biases could also produce sigmoidal curves: 'variant preferences', and 'demonstrator subgroups'. We address each of these in turn (see also Table 1 for a summary). Before doing so, we first address an underlying assumption of their model that has important implications for any interpretation of its results: the 50/50 variant starting condition.

Starting conditions

In the model of Acerbi et al.[16], a population of 100 individuals is initialized in which approximately 50% adopt behavioural variant A, and the rest adopt variant B. That is, the probability that an individual performs behaviour A at time = 0 is, on average, 0.5. Data is then aggregated across all model runs for the entire length of each simulation in order to describe the full range of variant frequencies, as either A or B goes to fixation. Two potential problems arise from this methodology. First, in the absence of group structure or anti-conformist drives, any random copying rule will, in expectation, monotonically converge to a single dominant variant. Thus, data for frequencies of variant A between 0.5 and 1 will come exclusively from runs in which uniform adoption of variant A was the absorbing state, while data for frequencies of variant A between 0 and 0.5 will come from runs in which the absorbing state was variant B (with some exceptions for stochastic noise around frequencies of 0.5). A corollary is that the dynamic in which a rare variant arises and becomes dominant is never explored. Second, because data is aggregated from across all time steps, more common variant frequencies will be overrepresented in the data. This should not matter if all aggregated populations are drawn from the same distributions. However, as we show below, errors can arise when populations using disparate rules are averaged together.

Variant preference

Acerbi et al. (2016)[16] claim that a non-conformist copying rule, which they call variant preference, can generate sigmoidal curves. The rule is instantiated in the model as follows:

> Individuals had a preference for one of the two variants, which was operationalized by endowing individuals with different copying probabilities for the two variants. When presented with a demonstrator showing it, one of the variants was always copied, while the other variant (the less preferred) was copied according to the parameter *pLess* (*pLess* was varied between 0 and 1). All individuals preferred the same variant, which was randomly selected at the beginning of each repetition. (Acerbi et al. 2016, p. 7)[16].

The sigmoidal curves generated stem from the two faulty assumptions described above: averaging data across all runs, and always starting with a 50-50 mix of behavioural variants[16]. We first replicate their result (Fig. 1, top row middle) of a sigmoidal curve. We then separate runs in which variant A is preferred (Fig. 1, top row right) from runs in which variant B is preferred (Fig. 1, top row left). This indicates that the sigmoidal curve is actually an artefact formed by joining together two lines with different slopes. What explains this?

When each individual prefers variant A, the probability of adopting A is simply equal to the probability of encountering an individual exhibiting variant A, which is just the frequency of A. For each individual preferring variant B, the probability of adopting A is the frequency of A times *pLess*: the probability of encountering an A times the probability that A is adopted in those cases. The sigmoidal curve results from these two lines being stuck together. Of course, two lines being averaged together should not result in a sigmoidal curve. This feature stems from the fact that the frequency of each variant is always initialized at 50%. Because adoption dynamics are always monotonic in expectation (i.e., the frequency of one variant always steadily increases, stochastic noise notwithstanding), the left and right sides of the graphs are generally from completely different batches of simulations, as illustrated by the top row of Fig. 1. We ran simulations with a uniform distribution of initial adoption frequencies for variant A ranging from 0% to 100%. The results (Fig. 1, bottom row) illustrate how preferences for either variant are associated with a straight line with slope depending on preference strength. Combining data from both preferences reveals a concave-up curve rather than sigmoidal curve. Such a curve, rather than a straight line, is a result of more data for the dominant variant preference at either end of the frequency spectrum. A simple test supports this interpretation (Fig. 2).

An analytical model of variant preference

The variant preference rule is equivalent to what Boyd and Richerson (1985)[3] called a direct bias. It is easy to derive a simple mathematical model of such a bias, and doing so will help to illustrate the dynamics of the individual-based model discussed above. Each individual observes another individual at random and decides whether to adopt their behaviour. Let p be the frequency of variant A. Upon observation, the probability of adopting one's preferred variant is 1, and the probability of adopting the less-preferred variant is q (equivalent to *pLess* in the Acerbi et al. model[16]). Finally, let r be the frequency of individuals who prefer variant A, in order to accommodate a mixed population of preference types. It now follows that the probability that a randomly chosen individual adopts variant A is given by:

$$\mathrm{P}rob(A) = p[p + r(1-p)(1-q)] + (1-p)[rp + (1-r)pq] \tag{1}$$

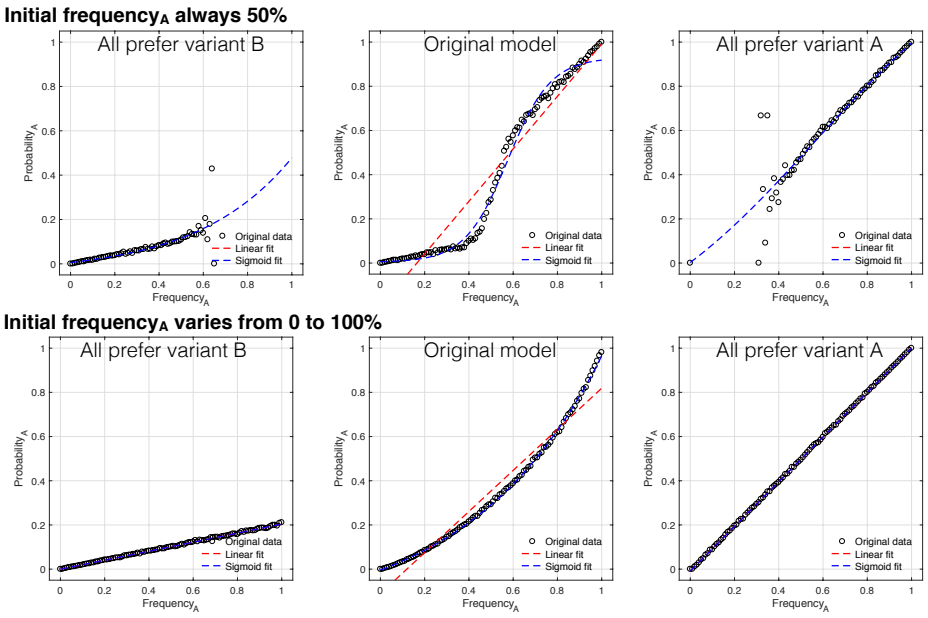

Figure 1. Variant preferences. A replication of the model of Acerbi et al.[16] in which individuals prefer one of the two variants. Includes best fit linear and sigmoid curves to data. Top row: Initial frequency of variant A always initialized to 50%, as in Acerbi et al. Bottom row: Initial frequency varied across runs from 0% to 100% in increments of 0.1%. Centre column: Runs in which variant A is preferred and runs in which variant B is preferred are averaged together. Right and left columns: variant A or B is always preferred, respectively.

Focusing on variant A, we see that if everyone in the population favours the same variant, we should observe strict concave up or down curves, as in Fig. 3A. If there is an even mix of both types, or if we average over an even number of runs of both types, the curves will average out, and we end up with a straight line. However, a sigmoidal curve can once again be generated from the assumption that the population is always initialized with 50% of each variant. Whenever one variant is favoured, that variant will always increase, and so no data will ever be recorded for frequencies less than 0.5. Rather, frequencies below 0.5 of variant A will only be observed when variant B is preferred ($r < 0.5$). Combining these two situations leads to the union of two distinct curves, and can generate an artefactual sigmoid curve, as in Fig. 3B. However, this bears little resemblance to any realistic data that would be collected in an empirical study.

Why does this analytical model give an r-shaped curve for the case where all individuals prefer variant A, while the model of Acerbi et al.[16] gives a straight line? The difference allows us to notice yet another important modelling artefact. Acerbi et al. took as their dependent measure the probability that any randomly selected copying event in which the focal individual was exposed to another individual with variant A resulted in that focal individual adopting variant A. This counts events in which the focal individual was already exhibiting variant A and then encountered an A, and so "copied" the variant she already had. However, it does *not* count events in which the focal individual started with variant A and encountered a B, but refrained from copying the newly encountered variant.

In our analytical model (Fig. 3), the dependent variable is, instead, the probability that a random individual will end up demonstrating variant A, given the current frequency of individuals exhibiting variant A in the population. In other words, the analytical model looks at the behaviours expressed in the population regardless of who encountered whom, and so counts the case where a focal individual started with variant A, encountered a B, and refrained from copying the newly encountered variant. Thus, the probability of demonstrating variant A when it is preferred is higher in our model than in the model of Acerbi et al.[16]. We view this as a more reasonable assumption, since the probability of adopting a behavioural variant should be enhanced by exhibiting an internal bias for that particular variant. We note also that our analytical result mirrors a very similar result obtained previously by Boyd and Richerson (1985, p. 207)[3].

We also believe that the dependent measure of adoption used in our analytical model is more in line with empirical methodology than the measure of adoption used by Acerbi et al.[16]. Firstly, the sample sizes that empirical studies deal with will often be substantially smaller than the size of the populations being modelled. A consequence of this is that researchers will use all of the data available to them (i.e. not ignore some cases—as when individuals starting with variant A encounter B and refrains from copying it). Second, in the vast majority of studies, individuals in the population will not be initiated with a

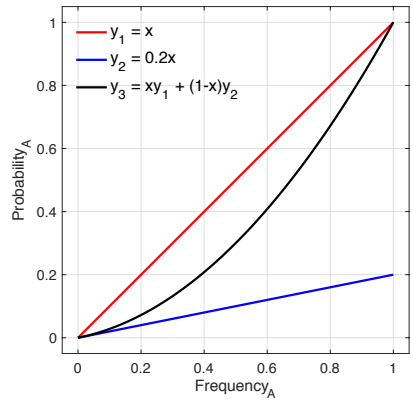

Figure 2. Explanation of concave-up curve. Curves y_1 and y_2 are simple lines with slopes of 1 and 0.2, reflecting the probability of adopting variant A or B, respectively. Curve y_3 is the weighted average of the two lines, in which y_2 for lower values of x and y_1 is dominant for higher values. Note that y_3 replicates the curve seen in Fig. 1, bottom centre.

particular behaviour at random, but rather will start out naïve. When only one individual acquires the task, then the probability of encountering a behavior will start close to 100%, rather than 50%.

In summary, the variant preference copying rule does not generate sigmoidal curves when unrealistic modelling artefacts are removed. That is, the variant preference rule is not a valid competing hypothesis for conformist learning when considering the causes of sigmoidal acquisition curves in results from empirical studies.

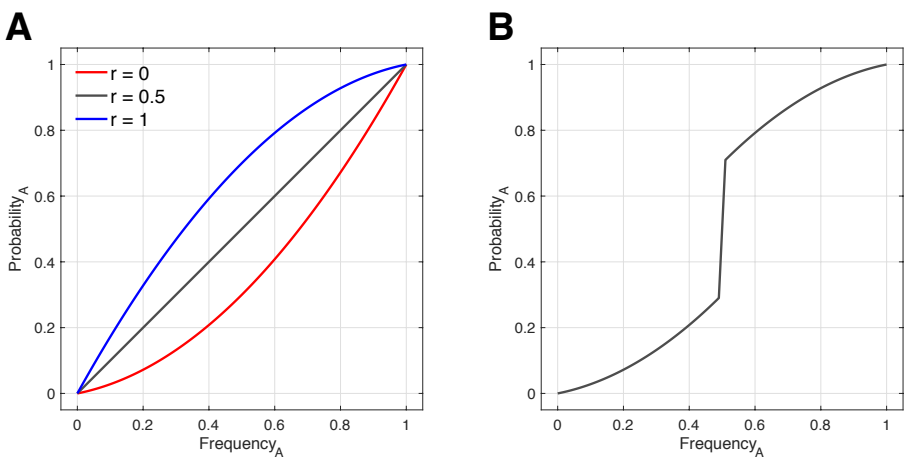

Figure 3. Analytical model for variant preferences. (A) When all individuals prefer variant A ($r = 1$), we replicate the r-shaped curve of Boyd and Richerson (1985)[3]. (B) If the environment is always initialized with approximately 50% of each variant and the preferred variant always increases in frequency, an artefactual sigmoidal curve can emerge if all data for both variant preferences are combined.

Demonstrator subgroups

The second copying rule purported by Acerbi et al. (2016)[16] to yield sigmoidal acquisition curves was called demonstrators subgroup. This rule corresponds to a situation in which individuals copy dominant or prestigious individuals, and works as follows:

> At the beginning of each repetition, a subset S of individuals in the population, of size Dm ($Dm = 5, 10, 20, 50$), was randomly chosen. At each time step, each individual was paired with a demonstrator chosen from the subset S and always copied it. Demonstrators themselves changed their state according to the same rule. (Acerbi et al., p. 7)[16].

We first replicate the results of Acerbi et al.[16] (Fig. 4A). Unlike variant preference, the sigmoidal curve generated by this rule is robust to variation in the initial frequency of each variant when averaged across simulations using a uniform distribution

of initial frequencies (Fig. 4B). If, however, the behavioural variant is initially rare, the sigmoidal curve becomes increasingly skewed, approaching an r-shaped curve (Fig.4C). The reason demonstrators subgroup generates a sigmoidal curve is as follows. Due to the well-known phenomenon of drift in small groups[19,20], random copying among the demonstrators will soon create a majority and then a uniform bloc. When one variant has a majority in the demonstrators subgroup, that variant will be copied at a rate higher than chance, because its frequency will increase more rapidly in the small subgroup than in the population at large. Once the variant reaches fixation in the subgroup, it will be copied with certainty by members of the population, evidenced by the horizontal lines for small and large frequencies of variant A in Fig.4A (the severity of these lines is reduced in Fig.4B because the initial frequencies of the variant in the subgroup and in the larger population are less strongly coupled). By restricting analysis to conditions in which variant A is initially rare, we lose the strong horizontal line for low frequencies of A (Fig.4C), in part because adopting variant A will in those cases be a rare event.

The acquisition curve generated by the 'demonstrators subgroup' rule is more severe (characterized by larger jumps in the second derivative) than the curve generated by conformist transmission; however, it can be classified as a sigmoidal curve. Should researchers therefore be concerned that their sigmoidal data may reflect a strategy of copying dominant individuals rather than the majority? Probably not. Acerbi et al. (2016)[16] showed that the demonstrators subgroup rule reliably generated a sigmoidal curve only when both the total group size and the relative number of demonstrators were very small. Specifically, a sigmoidal curve provided a better fit than a line only when $N \leq 100$ and $Dm <= 5$). In such cases, it should be relatively easy for researchers to identify this sort of dynamic and to test the demonstrators subgroup hypothesis against more direct conformist biases absent dominant leaders.

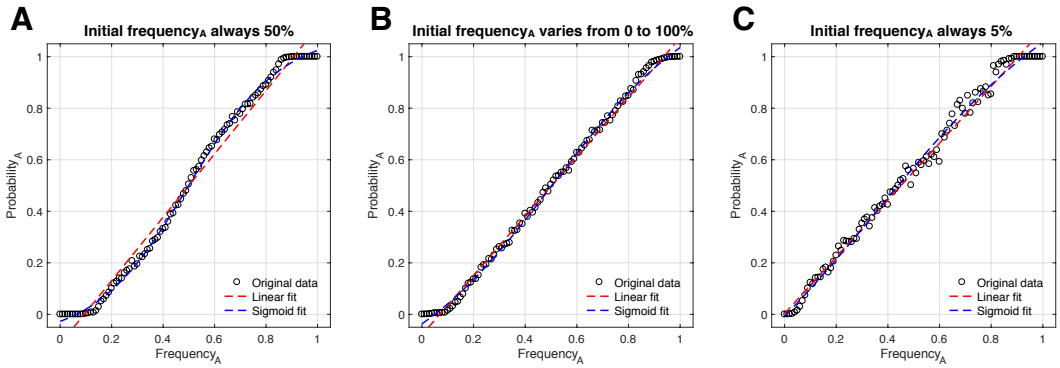

Figure 4. Demonstrators subgroup. Individuals choose a random member of a fixed group of size *Dm* and copy that individual's behavioural variant. For all simulations, *Dm* = 5. (A) Initial frequency of variant A is fixed at 50% for all runs, replicating the results of Acerbi et al. (2016)[16]. (B) The sigmoidal curve remains apparent when initial frequency of variant A is varied across runs from 0% to 100% in increments of 0.1%. (C) When variant A is always rare initially, the sigmoidal curve largely disappears, replaced with an r-shaped curve.

How to measure the majority

Conformist learning occurs when the more common behavioural variant is preferentially adopted. Empirically, there are at least two ways to measure just *how* common a particular variant is. You can count the number of individuals exhibiting each variant, or you can count the total instances of each variant. In the latter case, some individuals may be counted more than once. Two papers by van Leeuwen et al.[14,15] purported to show how using the frequency of observed behaviours as a measure of the majority, rather than the frequency of individuals exhibiting those behaviours, can lead to incorrect conclusions about the presence of conformist learning. To illustrate this, van Leeuwen et al.[15] constructed a model similar to the above, and compared acquisition curves produced using the frequency of behaviours and the frequency of individuals. They found that measuring the frequency of behaviours could produce artefactual sigmoidal curves even when individuals copied the behaviors of randomly chosen individuals, and so claimed that a reliable indicator of conformity could be derived only by measuring the frequency of individuals.

Firstly, it is worth noting that when implementing their model, the authors recorded the complete history of behaviours expressed by the population up to the point where the observer made their decision. This implementation suggests that observers use the complete history of all behaviours exhibited by all individuals when choosing between variants, or at least that empirical researchers analyse their data under this assumption. We would argue that this is approach is highly unlikely to be used by empiricists. Instead, most researchers will record the frequency of behaviours (or of individuals undertaking behaviours) in some constrained time prior to the focal individual making its decision (e.g., estimating what the individual might have plausibly

been able to observe in its current visit to the patch, rather than the what it might have observed across all previous history). For example, Aplin et al. (2015)[12] used a 245 second period prior to each bird's first solve (a mean of 6 previous solves) as their 'history of behaviour' time-window, and based this value on behavioural patterns observed from extensive studies in their population. Here we re-ran the baseline model from van Leeuwen et al (2016)[15] and Acerbi et al (2016)[16], in which individuals randomly select one individual to copy (Fig. 5A). Taking the entire history of behaviours produced by individuals in the population replicated the same sigmoidal response that was produced in the original studies (Fig. 5B). However, if the history was limited to the previous 10 behaviours, we found that the relationship was perfectly linear, as expected under a random copying strategy (Fig. 5C).

However if one is not going to use the complete history of behaviours, then how long should the time window for data collection be? A limitation of using a time window approach for considering the behavioural history of demonstrations is that the size of the time window is not always obvious, as highlighted in the critique by van Leeuwen et al (2016)[15]. To address this concern, we tested sensitivity of the relationship between frequency of behaviours observed and the probability of performing the behaviour to the size of the time window (Fig. 6). We found that a linear model provided a much better fit to the data than a sigmoidal model up to window sizes that were still quite large (more than 2000 observations). From this, we can conclude that an incorrect inference of conformist transmission when copying is in fact random will be very unlikely if any time-windowing is applied when calculating frequency of behaviours.

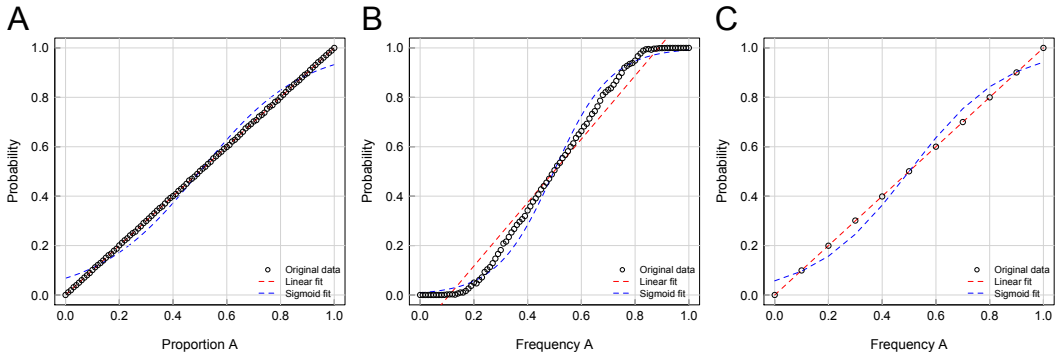

Figure 5. Implementations of the model described in van Leeuwen et al., counting the behavioural variants from the current state of the population (A), and from the complete history of actions performed (B). If the history of actions performed is limited to a recent window (here the last 10 solves, (C)), then there is no evidence that a sigmoidal curve could result.

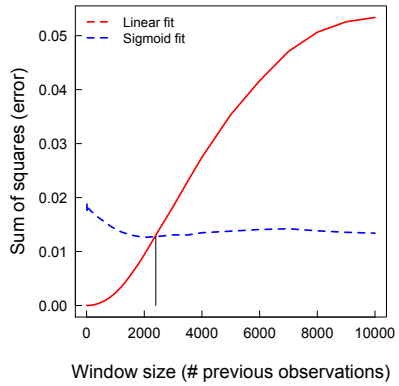

Figure 6. Relative fit of a linear versus sigmoidal model based on the r^2 value of the model fit. Here a window of size x is used to calculate the variant frequencies, where x is related to the number of previous observations from the population. As x increased, the fit of a linear model (solid line) decreases. However, a linear model remains a better fit than a sigmoid model until a very large window size.

Table 1. Summary of problematic issues found in the models by Acerbi et al (2016) and Van Leeuwin et al (2016). Severe issues result in the incorrect conclusions drawn from the simulations that we highlight in this paper. Issues with medium severity do not significantly affect the shapes of the curve in the results of the papers, but they incorrectly implement the biological process that the papers seek to investigate (for example, animals do not copy themselves).

Issue	**Severity**	**Explanation**
Starting conditions	High	The population was always initiated with a probability of 0.5 for each behaviour. Unrealistic from empirical perspective. Exacerbates the potential for an artefactual sigmoidal response. A range of initial conditions should be used to assess model robustness.
Pooling data	High	All of the frequency data across each of the 1000 replicates were pooled, without noting variation across replicates. This aggregates multiple categories of response into a single category. This practice can eliminate discovery of multimodal data, and in this case grafts two linear responses together in artefactual sigmoid. Data with categorically different starting conditions should not be pooled. Ideally, variation within each condition should also be reported.
Full behaviour history	High	The full history of all behaviours was used to calculate frequency of behaviours used by individuals to make behavioural choices. Does not represent what has been done in existing empirical studies, and generates errors by assuming individuals respond to a distant past they have not discounted. Using recent history alleviates any potential mistaken results.
What counts as 'copying' events	Medium	Every event was counted in which an individual encountered behavioural variant A, including when an individual did not change what action it performs, but events in which an individual *retained* variant A after encountering variant B were not counted. This increases the appearance of a sigmoidal acquisition curve (as opposed to either a linear or r-shaped curve) as individuals that continuously perform variant A when the population behaviour becomes entrenched are all counted as copying action A. Defining social learning events is often difficult, so striving for consistency and robustness is important.
Individuals are their own targets for social learning	Medium	Individuals could select themselves as demonstrators, which confounds data on social learning. Similarly, a focal individual's behaviours were counted when it assessed the historical frequency of observed events. Whilst in a population of 100 individuals these issues had relatively little impact on the results, the effects may be much greater in a smaller population.

Discussion

Three recent studies have questioned the empirical evidence for conformist copying in non-human animals, and re-opened the debate as to whether sigmoidal acquisition curves can be used as evidence for conformity. In this paper, we examined the different processes that generated sigmoidal curves in these published models, and find that non-conformist learning processes for sigmoidal outcomes were almost exclusively underpinned by flawed or unrealistic modelling assumptions. Our work further highlights how these assumptions misrepresent the types of analyses conducted by empiricists. More specifically, we first examined the non-conformist learning mechanisms purported to yield sigmoidal curves. We showed that one of them, variant preference, failed to do so when more realistic behavioural and analytical assumptions were used. The second mechanism, demonstrator subgroup, can generate weakly sigmoidal curves under some conditions, but is unlikely to be mistaken for conformist transmission and is easy to control for empirically. Second, we considered the argument that the majority should be represented as the frequency of individuals exhibiting a behaviour, rather than by the frequency of the behaviour itself, as the latter can also generate sigmoidal patterns under random copying. We showed that this view result stems from a modelling artefact and bears little relevance to the overwhelming majority of empirical research.

Our motivation for this paper is neither to provide a general critique of an individual-based modelling approach for understanding social learning and animal behaviour, nor to critique the general caution that alternative generative mechanisms at the individual level may be responsible for population-level patterns. Formal models, including individual-based models, have been critically important in the development of rich theories for understanding social behaviour and evolution. We are also sympathetic to the view that there can exist multiple causal pathways for population-level phenomenal; see, for example, Frank (2009)[21] on why power laws appear ubiquitous in complex systems. Rather we support publishing guidelines highlighting common errors, in line with the current movement towards a more robust scientific literature[22–25]. Indeed, our study highlights the importance of carefully exploring the factors that drive patterns in a model. A mathematical or computational model acts as a sort of logical engine for turning assumptions into conclusions[26,27]. Those assumptions concern the delineation of the parts of a system (often individual organisms), the properties and initial distributions of those parts, and the dynamic interactions by which the population changes[28]. The validity of the conclusions drawn from a model therefore depends fundamentally on the extent to which its assumptions adequately represent reality. All models must make unrealistic assumptions if they are to act as useful abstractions. However, the robustness of any model's results to idiosyncratic assumptions must be carefully assessed if we wish to avoid drawing erroneous conclusions when constructing analogies between models and reality.

Individuals can exhibit various attentional biases during the process of social learning that will have consequences both for acquisition patterns and for what traditions become established at the group level[1,2,29]. While a conformist learning bias is one the clearest examples of such an effect, multiple different social learning strategies have been identified across a wide variety of species, many of which could potentially also influence transmission dynamics[1,4,13]. It is therefore a valuable exercise to explore whether different social learning strategies could produce similar acquisition patterns, and we applaud the effort to do so. In our reanalysis of Acerbi et al.[16], we conclude that, indeed, one such social learning rule *can* generate sigmoidal acquisition curves superficially similar to conformist learning: the exclusive copying of a subgroup of potential demonstrators (the 'demonstrator subgroup' rule). However, it is important to note that this rule only generates a sigmoidal curve when both total group size and the number of demonstrators are small—that is, for populations of 100 individuals or fewer, with five or fewer demonstrators—and even then the sigmoidal nature of the curve is apparent only for very low and very high variant frequencies. For example, a social learning biases such as 'copy adult males' would be very unlikely to generate such a acquisition pattern, while 'copy the single most dominant male' might. This rule is therefore important to consider as an alternative explanation. However, we believe that empirical researchers have shown themselves to be capable of identifying candidate social learning strategies that might be operating in their system, and devising experiments to test hypotheses against one another. Two recent examples include movement direction in baboons[30], for which conformist learning was supported, and nest site navigation in honeybees (reviewed in Seeley 2010[31]), in which the preferentially copied individuals were readily identifiable. Finally, 'copy the dominant' biases have been explored in vervet monkeys[32] and thoroughly discussed in the context of additional evidence for conformity in this species[11,18].

Does a sigmoidal acquisition curve at the population-level uniquely indicate conformist transmission? We cannot say for certain. However, we have shown that there is, to our knowledge, no credible alternative mechanism for the generation of such curves. This does not mean that no such mechanism exists, nor does it eliminate the possibility that environmental or network constraints could contribute to the emergence of sigmoidal acquisition curves in some circumstances. Indeed, it remains a fascinating question as to whether alternative mechanisms can produce those stable and resilient group-level cultures that are so indicative of conformity[3]. However, if such conditions exist, they are not any that have been previously proposed for the social acquisition of information. In the end, a population-level acquisition curve is a somewhat crude tool with which to investigate the cognitive and behavioural strategies used by individuals. More complex methods and models that can incorporate things like individual heterogeneity, population structure, and the possibility that individuals use multiple learning strategies simultaneously are surely needed[33–35]. Yet for the time being, in the presence of a sigmoidal acquisition curve, conformist

transmission should be viewed as a highly probable antecedent.

Methods

In this study, we re-evaluated the modelling results described in Acerbi et al.[16] and Van Leeuwen et al.[15]. Our implementations were based on the source code provided at[36]. We modified this code in only two ways. The first modification was to relax the starting conditions in which the probability for each variant was 0.5 (Figures 1, 4). The second modification is to record, after each copying event, the frequency of behaviours over the previous T solves, where $T = 10$ (Figure 5) or ranged from 0 to 10000 (Figure 6). In Table 1, we list some additional implementation issues that we identified while working these models. However, because we were interested in isolating and demonstrating the primary mechanisms underpinning the apparent emergence of sigmoidal curves under random copying, and we did not otherwise alter the original code provided by the authors. Following Acerbi et al., all acquisition data were plotted along with their best-fit linear and sigmoid curves.

References

1. Laland, K. N. Social learning strategies. *Learn. & behavior* **32**, 4–14 (2004).
2. McElreath, R. *et al.* Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Philos. Transactions Royal Soc. Lond. B: Biol. Sci.* **363**, 3515–3528 (2008).
3. Boyd, R. & Richerson, P. *Culture and the evolutionary process* (University of Chicago Press, Chicago, IL, 1985).
4. Rendell, L. *et al.* Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends cognitive sciences* **15**, 68–76 (2011).
5. Muthukrishna, M., Morgan, T. J. & Henrich, J. The when and who of social learning and conformist transmission. *Evol. Hum. Behav.* **37**, 10–20 (2016).
6. Claidière, N. & Whiten, A. Integrating the study of conformity and culture in humans and nonhuman animals. *Psychol. bulletin* **138**, 126 (2012).
7. Morgan, T. J. H. & Laland, K. N. The biological bases of conformity. *Front. neuroscience* **6**, 87 (2012).
8. Henrich, J. & Boyd, R. The evolution of conformist transmission and the emergence of between-group differences. *Evol. human behavior* **19**, 215–241 (1998).
9. Nakahashi, W. The evolution of conformist transmission in social learning when the environment changes periodically. *Theor. population biology* **72**, 52–66 (2007).
10. Nakahashi, W., Wakano, J. Y. & Henrich, J. Adaptive social learning strategies in temporally and spatially varying environments. *Hum. Nat.* **23**, 386–418 (2012).
11. van de Waal, E., Borgeaud, C. & Whiten, A. Potent social learning and conformity shape a wild primate's foraging decisions. *Sci.* **340**, 483–485 (2013).
12. Aplin, L. M. *et al.* Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nat.* **518**, 538–541 (2015).
13. Hoppitt, W., Kandler, A., Kendal, J. R. & Laland, K. N. The effect of task structure on diffusion dynamics: Implications for diffusion curve and network-based analyses. *Learn. & Behav.* **38**, 243–251 (2010).
14. Van Leeuwen, E. J., Kendal, R. L., Tennie, C. & Haun, D. B. Conformity and its look-a-likes. *Animal Behav.* **110**, e1–e4 (2015).
15. Van Leeuwen, E. J., Acerbi, A., Kendal, R. L., Tennie, C. & Haun, D. B. A reappreciation of 'conformity'. *Animal Behav.* **122**, e5–e10 (2016).
16. Acerbi, A., Van Leeuwen, E. J., Haun, D. B. & Tennie, C. Conformity cannot be identified based on population-level signatures. *Sci. Reports* **6** (2016).
17. Aplin, L. M. *et al.* Counting conformity: evaluating the units of information in frequency-dependent social learning. *Animal Behav.* **110**, e5–e8 (2015).
18. Whiten, A. & van de Waal, E. Identifying and dissecting conformity in animals in the wild: Further analysis of primate data. *Animal Behav.* **122**, e1–e4 (2016).
19. Wright, S. Size of population and breeding structure in relation to evolution. *Sci.* **87**, 430–431 (1938).

20. Kimura, M. & Ohta, T. The average number of generations until fixation of a mutant gene in a finite population. *Genet.* **61**, 763 (1969).

21. Frank, S. A. The common patterns of nature. *J. evolutionary biology* **22**, 1563–1585 (2009).

22. Ioannidis, J. P. How to make more published research true. *PLoS Med* **11**, e1001747 (2014).

23. McElreath, R. & Smaldino, P. E. Replication, communication, and the population dynamics of scientific discovery. *PLoS One* **10**, e0136088 (2015).

24. Smaldino, P. E. & McElreath, R. The natural selection of bad science. *Royal Soc. Open Sci.* **3**, 160384 (2016).

25. Munafò, M. R. *et al.* A manifesto for reproducible science. *Nat. Hum. Behav.* **1**, 0021 (2017).

26. Gunawardena, J. Models in biology: 'accurate descriptions of our pathetic thinking'. *BMC Biol.* **12**, 29 (2014).

27. Smaldino, P. E. Models are stupid, and we need more of them. In Vallacher, R. R., Read, S. J. & Nowak, A. (eds.) *Computational Social Psychology*, 311–331 (Routledge, New York, 2017).

28. Kauffman, S. A. Articulation of parts explanation in biology and the rational search for them. In Buck, R. C. & Cohen, R. S. (eds.) *PSA 1970*, 257–272 (Philosophy of Science Association, 1971).

29. Rendell, L. *et al.* Why copy others? insights from the social learning strategies tournament. *Sci.* **328**, 208–213 (2010).

30. Strandburg-Peshkin, A., Farine, D. R., Couzin, I. D. & Crofoot, M. C. Shared decision-making drives collective movement in wild baboons. *Sci.* **348**, 1358–1361 (2015).

31. Seeley, T. D. *Honeybee democracy* (Princeton University Press, 2010).

32. van de Waal, E., Renevey, N., Favre, C. M. & Bshary, R. Selective attention to philopatric models causes directed social learning in wild vervet monkeys. *Proc. Royal Soc. Lond. B: Biol. Sci.* **277**, 2105–2111 (2010).

33. Aplin, L. M., Sheldon, B. C. & McElreath, R. Conformity does not perpetuate suboptimal traditions in a wild population of songbirds. *Proc. Natl. Acad. Sci.* **00** (2017).

34. Barrett, B. J., McElreath, R. & Perry, S. E. Pay-off-biased social learning underlies the diffusion of novel extractive foraging traditions in a wild primate. *Proc. Royal Soc. B* **284**, 20170358 (2017).

35. Kandler, A., Wilder, B. & Fortunato, L. Inferring individual-level processes from population-level patterns in cultural evolution. *bioRxiv* (2017). URL http://biorxiv.org/content/early/2017/02/25/111575.

36. Acerbi, A. Codes for simulations described in 'conformity cannot be identified based on population level signatures'. https://osf.io/jy3xk/ (2016).

Acknowledgements

We are grateful to Alberto Acerbi and Edwin van Leeuwen for their gracious correspondence, and also for sharing their code, without which our analysis would have been considerably more difficult. We thank Brendan Barrett, William Baum, Bret Beheim, Karl Frost, Richard McElreath, Matt Miller, Tom Morgan, and reading groups at UC Davis and MPI-EVA for helpful feedback on an earlier version of this manuscript. L. M. A. was supported by a Junior Research Fellowship at St John's College, University of Oxford, and both L.M.A. and D.R.F. were supported by a grant from the Biotechnology and Biosciences Research Council (BB/L006081/1).

Author contributions statement

All authors conceived of the project and helped write the manuscript. P.E.S. and D.R.F. performed the analyses. All authors reviewed the manuscript.

Additional information

None of the authors have any competing financial interests.