Trusting the experts: the domain-specificity of prestige-biased social learning

C O Brand^{1*}

A Mesoudi¹

TJH Morgan^{2,3}

¹ Human Behaviour and Cultural Evolution Group, Department of Biosciences, College of Life & Environmental Sciences, University of Exeter. * Corresponding author: c.brand@exeter.ac.uk

² School of Human Evolution and Social Change, Arizona State University,

³ Institute of Human Origins, Arizona State University

Abstract

Prestige-biased social learning (henceforth “prestige-bias”) occurs when individuals predominantly choose to learn from a prestigious member of their group, i.e. someone who has gained attention, respect and admiration for their success in some domain. Prestige-bias is proposed as an adaptive social-learning strategy as it provides a short-cut to identifying successful group members, without having to assess each person’s success individually. Previous work has documented prestige-bias and verified that it is used adaptively. However, the domain-specificity and generality of prestige-bias has not yet been explicitly addressed experimentally. By domain-specific prestige-bias we mean that individuals choose to learn from a prestigious model only within the domain of expertise in which the model acquired their prestige. By domain-general prestige-bias we mean that individuals choose to learn from prestigious models in general, regardless of the domain in which their prestige was earned. To distinguish between domain specific and domain general prestige we ran an online experiment (n=397) in which participants could copy each other to score points on a general-knowledge quiz with varying topics (domains). Prestige in our task was an emergent property of participants’ copying behaviour. We found participants overwhelmingly preferred domain-specific (same topic) prestige cues to domain-general (across topic) prestige cues. However, when only domain-general or cross-domain (different topic) cues were available, participants overwhelmingly favoured domain-general cues. Finally, when given the choice between cross-domain prestige cues and randomly generated Player IDs, participants favoured cross-domain prestige cues. These results suggest participants were sensitive to the source of prestige, and that they preferred domain-specific cues even though these cues were based on fewer samples (being calculated from one topic) than the domain-general cues (being calculated from all topics). We suggest that the extent to which people employ a domain-specific or domain-general prestige-bias may depend on their experience and understanding of the relationships between domains.

Keywords: social learning, prestige-bias, prestige-cues, hierarchy, expertise, domain-specific, domain-general

Data and analyses available: https://github.com/lottybrand/Prestige_2_Analysis

Preregistration available: <https://osf.io/93az8>

INTRODUCTION

Prestige-biased social learning (henceforth prestige-bias) occurs when learners preferentially learn from individuals who are copied by others, attended to by others, or who generally receive freely conferred deference from others [1–7]. These individuals are said to have ‘prestige’. Prestige-bias has been proposed as an adaptive social learning strategy as it provides an efficient short-cut to acquiring adaptive social information when more direct means (e.g. identifying and copying successful skills/knowledge, or highly skilled/knowledgeable individuals) are unavailable or costly [5]. Consistent with this claim, previous experimental work shows that participants will use prestige information when selecting a model from whom to socially learn, but only when a) the prestige information is related to success and b) direct success information is unavailable [1].

However, the domain-specificity or -generality of prestige-bias remains little understood. By these terms we mean whether a prestigious individual’s influence is limited to the specific domain in which they are successful (domain-specific prestige) or whether prestigious individuals who are successful in one domain become influential even outside their domain of expertise (domain-general prestige). To this end, we can ask whether social learners specifically copy individuals who are prestigious in the domain of interest (domain-specific prestige-bias), or if they copy individuals who are prestigious regardless of the source of this prestige (domain-general prestige-bias). For instance, is a successful and prestigious boat-builder sought out for advice on all kinds of general matters, or is their influence limited to their specific craft? Henrich & Gil-White [5] suggest that “...prestige hierarchies can be domain-specific. For example, if I defer to you because of your superior computer skills and you defer to Bob because he is an excellent grass hockey player, I may not give Bob any special deference if grass hockey is not my thing,” (Henrich & Gil-White, pp.170). However, they go on to suggest that “prestigious individuals are influential even beyond their domain of expertise,” (Henrich & Gil-White pp.184), and it is often claimed in the literature that prestigious celebrities who have acquired fame via a specific domain, such as sport or acting, have influence beyond their domain of expertise (see [8] for examples and critical discussion). As such, the domain-specificity of prestige-bias remains unclear.

We suggest that domain-general prestige-bias will be adaptive when success across different domains is correlated. As noted by Henrich & Gil-White, “much of the information that leads to success in one domain will often be transferable to others... acquiring skill in one domain (e.g. a martial art) is often touted as promoting success in many other areas. For example, problem-solving methods, goal-achieving strategies, eye-hand coordination, control over one’s emotions, etc. are useful across several domains,” (Henrich & Gil-White, pp.184). Henrich & Gil-White go further to suggest that “because figuring out which combinations of ideas, beliefs, and behaviours make someone successful is costly and difficult, selection favoured a general copying bias, which also tends to make prestigious individuals generally influential (as people copy and internalise their opinions),” (p.184). We suggest that the domain-generality of prestige-bias is an open empirical question, and is dependent on not only how correlated the domains in question are, but also whether observers have an accurate perception of these correlations.

Previous research shows some evidence of the domain-specificity of prestige-bias. An experiment found that 3-4 year old children preferentially copied the artefact choice of a prestigious demonstrator, defined as the demonstrator to whom bystanders had attended rather than ignored, when their prestige was acquired during an artefact task (i.e. in the same domain), but not when it was acquired on a food-preference task (i.e. in a different domain) [4]. The reverse preference was seen when the demonstrator acquired their prestige on the food-preference task rather than the artefact choice task. Other studies from naturally-occurring groups have provided indirect evidence that prestige-bias may not be as domain-general as suggested by Henrich & Gil-White. For example, Reyes-Garcia et al. [9] found that prestige was unrelated to ethnobotanical medicinal knowledge in an indigenous Tsimane community, but instead was related to having a formal position in that society. Their measure of prestige was the number of nominations received when others were asked to list all the important men in the village, suggesting that being highly skilled and knowledgeable in the domain of medicine does not translate to being an important man in the village generally, or gaining a formal position. Similarly, a study of naturally-occurring volunteer groups in Cornwall, such as chess or kayaking clubs, found that prestige ratings were not related to success in a quiz that the group completed together, but were related to formal positions that the individuals already had in their groups, such as teacher, team captain, or group organiser/secretary [10].

Here we present an experiment to directly assess the degree to which participants use domain-specific or domain-general prestige cues when choosing from whom to learn. We use participant generated prestige cues and manipulate the choices participants have to directly compare the different types of cues. Participants answered quiz questions from four explicitly different topics: weight estimation, world geography, art history, and language identification. We treat each of these topics as a 'domain', and ask whether people use success-derived prestige in one domain/topic as a guide for who to copy in a different domain/topic.

On each of 100 questions, participants could answer for themselves or copy another individual from their group (henceforth 'demonstrator'). Depending on the experimental condition (see Table 1), participants could choose who to copy based on (1) the demonstrators' scores on questions from the current topic (domain-specific success) or (2) the number of times the demonstrators were previously copied (our experimental proxy for prestige) on (a) questions from the current topic (domain-specific prestige), (b) all topics (domain-general prestige) or (c) a different topic (cross-domain prestige).

A previous experiment using the same quiz and topics [1] found that participants' scores across both rounds of questions from the same topic were more strongly associated (R^2 between 0.3 and 0.6), than scores across both rounds of different topics (R^2 between 0.06 and 0.36). The associations between scores on a particular topic and overall score across rounds were intermediate between these two ranges. These associations therefore indicate that domain-specific scores (scores on questions within the same topic) are the most reliable cue to choose potential demonstrators from whom to copy when answering questions from the same topic/domain, compared to cross-domain or domain-general cues (see Table 2). Participants,

however, were not informed of these correlations, nor had any way to calculate them during the experiment.

Given that domain-specific prestige should be most highly correlated with score on the topic being answered, and that cross-domain prestige will be least correlated with score on the topic being answered, we formulated the following hypotheses, preregistered at <https://osf.io/93az8>

1. Domain-specific prestige-bias will be employed more often than domain-general prestige-bias and cross-domain prestige
2. Domain-general prestige-bias will be employed more often than cross-domain prestige-bias
3. Cross-domain prestige-bias is only employed when access to domain-general or domain-specific prestige information is not available.

METHODS

Participants:

Consistent with our previous study [1], we aimed to recruit ten groups of ten participants via Mechanical Turk for each of the four conditions, totalling 400 participants. Due to difficulties in coordinating multiple groups of participants simultaneously, some participants dropped out at various stages, giving a starting sample size of 397, with 332 participants reaching the final demographics page. We ensured a minimum of 5 participants in each networked group throughout. All participants were above the age of 18 (age range 20 – 69, mean age = 39.1), with 113 men and 78 women (not every participant specified their gender). All were given a monetary reward for their time of USD\$10, and had the opportunity of winning a bonus payment of \$20 if they scored over 85% in the quiz. 33 participants received the \$20 bonus payment. All participants provided informed consent before being able to proceed to the task. The consent form informed them that their participation was entirely voluntary, their data were entirely anonymous, and they could withdraw their involvement at any time by closing their browser.

Materials:

The experimental automation platform Dallinger (dallinger.readthedocs.io) was used to create an online game in which groups of players can play and interact simultaneously. Participants answered 100 questions with two alternative answers each, one correct and one incorrect. The 100 questions were split into four topics of 25 questions each: geography, weight estimation, language, and art (see Fig. 1 for an example).

Figure 1: Four example screenshots representing questions from the Geography, Art, Language and Weight topics as seen by participants. Half the questions from each topic included images and half did not. Participants could either click one of the two blue buttons showing two possible answers (one correct, one incorrect), or click the red button labelled “Ask Someone Else” which allows participants to copy someone else within their group. The number at the bottom is a countdown timer that forces participants to answer within 15 seconds. Player ID numbers are generated randomly for each participant when they begin the quiz. Here, three example screenshots come from one player (ID 538) and one from another player (ID 753). In reality, all participants view each question simultaneously. Note: all map images used in the quiz were sourced from <https://www.eduplace.com/ss/maps/> and edited accordingly.

Procedure:

Participants were given 100 binary choice questions based on four different general-knowledge or ‘trivia’ style topics, 25 in each category. Participants had fifteen seconds to answer each question, and they scored one point for every question they answered correctly. On each question, instead of answering themselves, participants could choose to “Ask Someone Else” by clicking the corresponding button. This allowed them to see information about other participants (‘demonstrators’) who had answered that question themselves (if everyone chose to “Ask Someone Else” then participants were shown a message saying “sorry, everyone chose to ‘ask someone else’ so no one can score points for this question”). The information they saw depended on the condition, detailed in Table 1. They then chose a demonstrator whose answer they used for that question. If the chosen demonstrator answered the question correctly, the copying participant also scored a point for that question. If the demonstrator was incorrect, they did not. No one received feedback on whether their answer was right or wrong at any point.

In all conditions, the 100 questions were split into two rounds, 60 questions in Round 1 (15 from each topic), and 40 in Round 2 (10 from each topic). In Round 1, when participants chose to “Ask Someone Else”, they always saw the current domain-specific score of all available

demonstrators. However, in Round 2, participants could choose between two kinds of information, which varied according to the experimental condition (see Table 1). The information available included domain-specific prestige (the number of times each demonstrator was copied on Round 1 questions from this topic), domain-general prestige (the number of times copied on all Round 1 topics), cross-domain prestige (the number of times copied on Round 1 questions from a randomly selected, different topic) and a random cue (each demonstrator's Player ID, which was a randomly chosen 3 digit number). The four conditions pitted different pairs of information against each other, to establish participants' preferences, and across the four conditions we constructed a hierarchy of information, from most to least preferred.

Table 1: Information displayed when choosing to "Ask Someone Else" in Rounds 1 and 2 across conditions, with our predicted choice for Round 2 in bold.

	Condition A	Condition B	Condition C	Condition D
Round 1	Domain-specific score			
Round 2	Cross-domain prestige or Domain-specific prestige	Cross-domain prestige or Domain-general prestige	Domain-general prestige or Domain specific prestige	Random Cue or Cross-domain prestige

Pre-registered Predictions:

The following predictions were preregistered at <https://osf.io/93az8>.

Assumption Checks:

1. Domain-specific prestige (i.e. the number of times a participant was copied in Round 1 of the current topic) is more tightly correlated with domain-specific score than is domain-general prestige or cross-domain prestige. This is an assumption check based on our previous findings and the experimental set-up. If this prediction is not upheld it suggests the topics were more tightly correlated than anticipated, and our following predictions may not hold.
2. Across all conditions, when choosing to "Ask Someone Else" in Round 1, participants preferentially copy the highest-scoring demonstrator. This is an assumption-check, to make sure that subsequent copying frequency cues are genuine signals of performance, and a replication of our previous findings.
3. Across all conditions, when choosing to "Ask Someone Else" in Round 2, participants preferentially copy the most-copied demonstrator. This is an assumption-check to make sure that people are actually employing prestige-bias when it is potentially useful, and a replication of our previous findings.

Main Predictions:

4. In Condition A, participants will predominantly choose to copy based on domain-specific prestige, rather than cross-domain prestige, as domain-specific prestige will be most correlated to score on the relevant domain.
5. In Condition B, participants will predominantly choose to copy based on domain-general prestige, rather than cross-domain prestige, as domain-general prestige will be more correlated to score on the relevant domain than cross-domain prestige (as it contains the topic-specific score within it).
6. In Condition C, participants will predominantly choose to copy based on domain-specific prestige rather than domain-general prestige, as domain-specific prestige will be more correlated to score on the relevant domain.

Follow-up predictions:

7. The overall frequency of copying (i.e. choosing to “Ask Someone Else”) is higher in Conditions A and C than in B. This is because Conditions A and C both provide domain-specific prestige information which is a more direct cue to domain-specific success, whereas Condition B provides only indirect cues to domain-specific score.
8. Participants score higher in Conditions A and C than B. This is because Conditions A and C provide more tightly correlated cues to success with the relevant domain.

Unregistered predictions:

As laid out in the preregistration, after collecting and analysing data for Conditions A-C and confirming that participants showed a clear preference for domain-specific information over cross-domain and domain-general prestige, we then collected additional data for Condition D which compared cross-domain prestige (the least preferred information) with a random cue (Player ID number). Thus for Condition D, we predicted that participants would predominantly choose to copy based on cross-domain prestige rather than the random cue, as cross-domain prestige contains some correlation with score, whereas the random cue does not. Because Condition D contains the least favoured information, we also predicted lower copying rates and lower scores than the other conditions.

Pre-registered Analyses:

We ran a series of Bayesian multi-level mixed models using the Rethinking package version 1.90 [11] in R version 3.6.0 [12]. Our raw data and analysis scripts can be found at https://github.com/lottybrand/Prestige_2_Analysis. Each model corresponding to each prediction is also included in the supplementary material. Model parameters are interpreted as providing evidence of an effect on the outcome if their 89% credible interval (CI) did not cross zero, with 89% CIs the default value in the Rethinking package. Priors were chosen to be weakly regularising, in order to control for both under and overfitting the model to the data. Convergence criteria such as effective sample sizes and Rhat values were used to check for appropriate model convergence throughout, and trace plots were inspected for signs of incomplete mixing when necessary.

Assumption Checks (Predictions 1, 2 and 3):

To test Prediction 1 we used *cor.test()* to generate Pearson's correlation coefficients. We report the R^2 coefficients in Table 2 to give a more intuitive sense of the proportion of variation explained.

To test Predictions 2 and 3 (participants preferentially copy the highest scoring or most copied demonstrator), the demonstrators are coded as "highest scoring/copied" or not, to check if participants were consistently choosing the top scorer/most copied out of those that were available for that copying instance. "Highest scorer" or "Most copied" was the outcome variable. We used a multilevel binomial model with varying intercepts for group and participant. To test Prediction 2, data from all instances of copying in Round 1 was used (Model 1 in supplementary material). To test Prediction 3, data from all instances of copying in Round 2 was used (Model 2 in supplementary material). In these models, the intercept terms tell us whether participants tended to choose the highest scorer / most copied demonstrator.

Main Analysis (Predictions 4, 5 and 6):

To test Predictions 4, 5 and 6 a generalised linear mixed model (Model 3 in supplementary material) of all data from Round 2 was used with "chose predicted" (i.e. chose to view the information that we predict in each condition) as the binomial outcome variable (yes or no), with varying intercepts for participant, group and condition.

Follow-up analysis (Predictions 7 and 8):

To test Prediction 7 a generalised linear mixed model of all data from Round 2 was used with 'copied' as the binomial outcome variable (chose to copy, i.e. "Ask Someone Else", or not). To test Prediction 8, a general linear mixed model of data from each condition was used with final score of each participant as the outcome variable and condition as the predictor variable, with a varying intercept for participant and group.

RESULTS

Prediction 1 (Assumption check): As predicted, scores were more tightly associated between Rounds 1 and 2 of the same topic than they were between different topics, except for the Geography topic in which the within-topic R^2 (.22) was lower than two cross-topic R^2 (Art and Language) (Table 2). Also as predicted, the overall Round 1 score was intermediate between the within-topic associations and the cross-topic associations for Art and Weight, but not for Geography and Language where the overall R^2 was higher than the within-topic R^2 . This latter unexpected finding may reflect the fact that overall score contains four times as much data as does the score for any specific topic.

Table 2: R^2 between Round 1 and Round 2 topic scores representing total variance of Round 2 topic score explained by Round 1 score of the same topic. All scores are participants' *asocial* score, so only includes scores they achieved without copying others. *Bold* R^2 shows highest association for that topic.

Round 1 asocial score	Round 2 asocial score	R^2
Art	Art	0.51
Geography	Art	0.19
Language	Art	0.25
Weight	Art	0.15
Overall Round 1 Score	Art	0.41
Geography	Geography	0.22
Art	Geography	0.31
Language	Geography	0.26
Weight	Geography	0.15
Overall Round 1 Score	Geography	0.35
Language	Language	0.34
Art	Language	0.30
Geography	Language	0.18
Weight	Language	0.14
Overall Round 1 Score	Language	0.37
Weight	Weight	0.32
Art	Weight	0.22
Geography	Weight	0.09
Language	Weight	0.18
Overall Round 1 Score	Weight	0.30

Predictions 2 and 3 (Assumption checks): As predicted, when participants chose to copy others in Round 1 they preferentially copied the most successful (i.e. highest scoring) demonstrator out of all those available (mean coefficient estimate: 3.15, 89% CI: 2.77, 3.55, this corresponds to participants choosing the highest scoring model 95.9% of the time on the probability scale). When participants chose to copy in Round 2, participants preferentially copied the most prestigious (i.e. most copied) demonstrator out of all those available (mean coefficient estimate: 4.29, 89% CI: 3.59, 5.10, this corresponds to participants choosing the highest scoring model 98.6% of the time on the probability scale).

Predictions 4, 5 and 6 (Main hypothesis tests): As shown in the raw data of Figure 2, participants preferred the domain-specific prestige cue when available, followed by domain-general, cross-domain, and finally random cues, as predicted. Our models confirmed that participants preferentially chose the predicted information as opposed to the alternative information in each condition we tested, showing a strong preference for the predicted information (see Figure 3; **Condition A:** mean: 3.05 89% CI: 2.42, 3.73 **Condition B:** 2.51, 89% CI: 1.90, 3.16, **Condition C:** mean: 2.84, 89% CI: 2.24, 3.46 **Condition D:** 2.10, 89% CI: 1.34, 2.87).

353

354
355
356
357
358
359

Figure 2: Raw counts of the information chosen when participants chose to copy someone else's answer in Round 2. Total possible copying instances for each condition in Round 2 were: Condition A = 3240, Condition B = 3437, Condition C = 3377, Condition D = 3416.

360
361
362
363
364
365

Figure 3: Model predictions for participants choosing the predicted information compared to the alternative information in Round 2 of the four Conditions, on the probability scale.

Prediction 7: Our model did not provide strong evidence of a difference in copying frequency between Conditions A (18% copying rate, $n=3240$), B (14% copying rate, $n=3437$) and C (20% copying rate, $n=3377$), however there is evidence of participants copying less in Condition D compared to the other conditions (7% copying rate, $n=3416$). This was confirmed by computing contrasts between the estimates of the four conditions (mean difference between Conditions A and D: 1.77, 89% CI: 0.72, 2.81. Mean difference between Conditions B and D: 1.26, 89% CI: 0.29, 2.24. Mean difference between Conditions C and D: 1.83, 89% CI: 0.76, 2.86).

Prediction 8: Similarly to prediction 7, there was not strong evidence that participants scored differently between Conditions A (mean score 71.1%), B (71.6%) and C (72.5%). However, there was strong evidence that participants scored lower in Condition D (mean score 63.5%) compared to Condition C only. This was confirmed by computing contrasts between the conditions (mean difference between Condition D and C: -1.15, 89% CI: -2.23, 0.04). When looking at just the Round 2 scores (in which there was a difference in the information participants could choose), participants scored lower in Condition D (mean score 24.8) compared to Condition B (mean score 27.4, mean difference -1.18 89% CI: -2.02 -0.33) and Condition C (mean score 27.4, mean difference -1.19 89% CI: -2.03, -0.32), but not Condition A (mean score 26.5).

The supplementary material contains model specifications, a comparison with our previous prestige study that used a similar design [1] and further exploratory analyses, including results showing that copying rate predicts score, and plots showing individual variation in prestige score and copying rate.

DISCUSSION

We experimentally tested whether participants adopted a domain-specific prestige-bias or a domain-general prestige-bias when both options were available to them during an online quiz. This quiz contained different topics that we used to represent different ‘domains’ of knowledge. Participants were able to copy other players’ quiz answers based on their domain/topic-specific scores in Round 1 of the quiz, and were then able to copy other players based on how many times they had previously been copied by others (our measure of ‘prestige’) in Round 2 of the quiz. In a series of pair-wise comparisons between prestige cue types, we found that participants overwhelmingly chose to use domain-specific prestige cues (times copied in the same topic) over both domain-general (time copied overall) and cross-domain (times copied in a different randomly chosen topic) prestige cues, that they preferred domain-general to cross-domain cues, and that they preferred cross-domain to a random cue entirely unconnected to success in the task. We therefore revealed a ‘hierarchy’ of prestige cues in which the most favoured cue, when available, was domain-specific prestige, followed by domain-general, then cross-domain, and lastly random cues.

This study adds to the already extensive body of work which shows that people tend to use social information in an adaptive and flexible way, depending on the information that is available to them [13–15]. We also provide further evidence of success-biased social learning, in that

participants copied the highest scoring player available to them when copying based on score in Round 1 [1,14,16]. Participants also copied the most-copied player available to them when copying based on prestige in Round 2. Our results also support previous evidence suggesting that participants are sensitive to the domain-specificity of prestige when copying or learning from others [4,9,10].

Our study used different topics of a quiz to represent different domains of knowledge, and participants' scores in Round 1 of a given topic were generally better predictors of their score in Round 2 of the same topic than of any other topic. Thus, copying based on domain-specific prestige was most adaptive for increasing a player's chances of selecting the correct answer on any given question, and ultimately in achieving a high score (and bonus payment) in the quiz. As the quiz was played live by groups of participants with no feedback on their scores, participants did not have access to how correlated scores were within or between topics. Participants did have experience of each topic in Round 1, presumably aiding their assessment of how related the topics were. Similarly, in real life people do not have direct access to data on how correlated different domains of skill or knowledge are, but presumably gain an intuitive understanding based on their experience or exposure to different domains during their lifetimes. Our predictions assume that people have an intuitive sense of which domains should be correlated, based on their experience. Thus, an individual's tendency to prefer domain-specificity over generality relies on their own understanding or assessment of how correlated the domains might be. How accurate their perceptions of the correlations are is likely to depend on their own experience and expertise in those domains, and emerge either through unconscious associative learning, conscious deliberation, or a mixture of the two. If direct experience of a domain is lacking, intuitions might be based on stereotypes or folk understanding, leading to inaccurate assessment of relationships between domains, and maladaptive use of prestige-bias.

We found strong evidence that participants prefer domain-specific prestige cues to cross-domain prestige-cues, i.e. they would rather copy the answer of someone who was copied extensively on the current topic, than someone who was copied extensively on a different topic. Our topics were explicitly labelled as different to each other, and participants had experience of each topic in Round 1. This explicit labelling and direct experience may be responsible for participants' strong domain-specific bias. However, if topics were more alike, participants may not show such a strong preference for domain-specific prestige. Importantly, if they had less experience of the topics or if the topics were unfamiliar to them, they may be less able to make a judgement about their correlations. For example, an individual unfamiliar with science may see both a physicist and a biologist as belonging to the same domain of expertise, 'scientist', and perhaps incorrectly choose to learn about physics from a biologist. The biologist may be a better model compared to an historian or an artist, but another scientist might choose a more domain-specific model, i.e. a physicist. This may help to explain why prestigious individuals often influence others on topics outside of their domain of expertise, especially within celebrity culture and social media platforms such as Twitter. It may be more informative for someone to listen to a prestigious children's author on a topic related to biology if they have no other access to biology experts. However another biologist may choose not to listen to a children's author on

this topic as they have access to plenty of other biologist models. Perhaps if people had wider and more direct access to experts on particular topics, they would preferentially weight the opinion of the expert in that same topic over an expert in a different topic.

As noted by Henrich & Gil-White [5], the emergence of domain-general prestige depends on whether success in multiple domains depends on success on an underlying trait. In our quiz, the topics were chosen to be as dissimilar to each other as possible (within the constraints of an online quiz), however we showed that performance across topics was to some extent correlated, either because the domains are drawing on similar cognitive abilities, or because success across domains is correlated with an underlying trait such as education, socioeconomic status, or 'general intelligence' such as 'g' [17–20]. The theoretical justifications and debates surrounding the construct of *g* are beyond the scope of this paper, but if theoretical and statistical models of *g* are reliable and replicable, then this would support the argument that domain-general prestige-bias can often be as adaptive as, if not more adaptive than,, domain-specific prestige. Just as a general intelligence factor may underpin general cognitive abilities, success across many broad areas of societal interest, such as politics, science or ethics, could plausibly be correlated with underlying traits such as reasoning, decision-making, or evidence-based judgement. For this reason it could be adaptive to use domain-general prestige-bias for large-scale decisions. For example, for successful decision-making on climate change policy, an understanding of data, evidence gathering, political landscapes and policy logistics may all be necessary, and may all be reliant on common training, expertise or experience. In this way, a renowned "thinker" or "intellectual" (e.g. a podcast host such as Sean Carrol) may be just as informative on climate policy as a climate scientist; the climate scientist may know the climate change data exceptionally well but may not have any experience of the political landscape or policy logistics.

Interestingly, domain-specific prestige was overwhelmingly preferred to domain-general prestige, even though in our experimental setup the former actually contained less information, being based on 15 topic-specific questions from Round 1, than the latter, which was based on all 60 questions from Round 1 (and indeed overall score in Round 1 was a better predictor of topic specific score in Round 2 in two out of four topics). Thus the domain-general prestige information included copying instances from four times as many questions compared to the domain-specific prestige information. This difference in information was also reflected in the prestige scores visible to participants, in that more copying instances would have occurred across all topics compared to within each topic. This suggests that people might be less sensitive to the amount of information than might be expected. Whether our experimental set-up reflects real life depends on a trade-off between depth and breadth of expertise. For example, someone may have a long history of prestige in biology generally, thus lots of information has been gathered about their success as a biologist. Someone else may have recently gained expertise in a specific area of biology, such as molecular genetics. Our results suggest people would rather learn molecular genetics from the recently trained molecular genetics expert than the long-standing general biology expert. Similarly, if someone is renowned for being generally successful or knowledgeable in a variety of domains, do we have more information about their expertise than about the expertise of someone who is a long-time expert in a narrower field?

This trade-off between depth and breadth of expertise is an open-question and one worth investigating when trying to understand who people choose to learn from and why.

In summary, we find evidence that domain-specific prestige-bias is preferred to both cross-domain and domain-general prestige-bias, at least when the domains of interest are sufficiently dissimilar to each other, and when individuals have had experience of each domain. Participants revealed a hierarchy of prestige in that domain-specific prestige was most preferred, followed by domain-general, and cross-domain prestige was least preferred. This preference was present despite the fact that the domain-general prestige scores contained more information (were generated based on all topics) compared to domain-specific prestige scores (generated from one topic). This apparent trading of depth over breadth of expertise warrants further experimental investigation to understand who people learn from and why.

References

1. Brand CO, Heap S, Morgan TJH, Mesoudi A. The emergence and adaptive use of prestige in an online social learning task. *Sci Rep.* 2020;10: 12095. doi:10.1038/s41598-020-68982-4
2. Cheng JT, Tracy JL, Foulsham T, Kingstone A, Henrich J. Two ways to the top: Evidence that dominance and prestige are distinct yet viable avenues to social rank and influence. *J Pers Soc Psychol.* 2013;104: 103–125. doi:10.1037/a0030398
3. Cheng JT, Tracy JL, Ho S, Henrich J. Listen, follow me: Dynamic vocal signals of dominance predict emergent social rank in humans. *J Exp Psychol Gen.* 2016;145: 536–547. doi:10.1037/xge0000166
4. Chudek M, Heller S, Birch S, Henrich J. Prestige-biased cultural learning: bystander's differential attention to potential models influences children's learning. *Evol Hum Behav.* 2012;33: 46–56. doi:10.1016/j.evolhumbehav.2011.05.005
5. Henrich J, Gil-White FJ. The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evol Hum Behav.* 2001;22: 165–196. doi:10.1016/S1090-5138(00)00071-4
6. Jiménez ÁV, Mesoudi A. Prestige-biased social learning: current evidence and outstanding questions. *Palgrave Commun.* 2019;5: 1–12. doi:10.1057/s41599-019-0228-7
7. Lenfesty HL, Morgan TJH. By Reverence, Not Fear: Prestige, Religion, and Autonomic Regulation in the Evolution of Cooperation. *Front Psychol.* 2019;10. doi:10.3389/fpsyg.2019.02750
8. Acerbi A. *Cultural Evolution in the Digital Age.* Oxford University Press; 2019.
9. Reyes-Garcia V, Molina JL, Broesch J, Calvet L, Huanca T, Saus J, et al. Do the aged and knowledgeable men enjoy more prestige? A test of predictions from the prestige-bias model of cultural transmission. *Evol Hum Behav.* 2008;29: 275–281. doi:10.1016/j.evolhumbehav.2008.02.002

10. Brand CO, Mesoudi A. Prestige and dominance-based hierarchies exist in naturally occurring human groups, but are unrelated to task-specific knowledge. *R Soc Open Sci.* 2019;6: 181621. doi:10.1098/rsos.181621
11. McElreath R. *Statistical Rethinking*. CRC Press; 2016.
12. R: The R Project for Statistical Computing. [cited 15 Dec 2020]. Available: <https://www.r-project.org/>
13. Laland KN. Social learning strategies. *Anim Learn Behav.* 2004;32: 4–14.
14. Morgan TJH, Rendell LE, Ehn M, Hoppitt W, Laland KN. The evolutionary basis of human social learning. *Proc R Soc B.* 2012;279: 653–662. doi:10.1098/rspb.2011.1172
15. Rendell L, Boyd R, Cownden D, Enquist M, Eriksson K, Feldman MW, et al. Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science.* 2010;328: 208–213. doi:10.1126/science.1184719
16. Mesoudi A. An experimental simulation of the “copy-successful-individuals” cultural learning strategy: adaptive landscapes, producer–scrounger dynamics, and informational access costs. *Evol Hum Behav.* 2008;29: 350–363. doi:10.1016/j.evolhumbehav.2008.04.005
17. Conway ARA, Kane MJ, Engle RW. Working memory capacity and its relation to general intelligence. *Trends Cogn Sci.* 2003;7: 547–552. doi:10.1016/j.tics.2003.10.005
18. Kovacs K, Conway ARA. Process Overlap Theory: A Unified Account of the General Factor of Intelligence. *Psychol Inq.* 2016;27: 151–177. doi:10.1080/1047840X.2016.1153946
19. Kovacs K, Conway ARA. What Is IQ? Life Beyond “General Intelligence.” *Curr Dir Psychol Sci.* 2019;28: 189–194. doi:10.1177/0963721419827275
20. Savi AO, Marsman M, van der Maas HLJ, Maris GKJ. The Wiring of Intelligence. *Perspect Psychol Sci.* 2019;14: 1034–1061. doi:10.1177/1745691619866447

Author Contributions: COB, TJHM & AM designed the study and preregistered the predictions together. COB coded the Dallinger experiment and analysed the data. COB, TJHM and AM wrote the manuscript together.

Ethics: All tasks were given ethical permission by the University of Exeter Biosciences ethics committee.

Competing interests: All authors declare no competing interests.

Funding: This research was supported by The Leverhulme Trust (grant no. RPG-2016-122 awarded to AM)

Data and analyses available: https://github.com/lottybrand/Prestige_2_Analysis

Preregistration available: <https://osf.io/93az8>