PROCEEDINGS B

royalsocietypublishing.org/rspb

Research

Cite this article: Jiménez ÁV, Chang L, Jensen K, Lu HJ, Mesoudi A. 2026 Limited evidence that reputation-based partner choice facilitates information sharing in humans. *Proc. R. Soc. B* **293**: 20260876.
https://doi.org/10.1098/rspb.2026.0876

Received: 5 April 2026
Accepted: 19 May 2026

Subject Category:
Behaviour

Subject Areas:
behaviour, cognition

Keywords:
cooperation, cultural evolution, human behaviour, partner choice, social learning

Author for correspondence:
Alex Mesoudi
e-mail: a.mesoudi@exeter.ac.uk

Electronic supplementary material is available online at https://doi.org/10.6084/m9.figshare.c.8502444.

THE ROYAL SOCIETY PUBLISHING

Limited evidence that reputation-based partner choice facilitates information sharing in humans

Ángel V. Jiménez[1], Lei Chang[2], Keith Jensen[3], Hui Jing Lu[4] and Alex Mesoudi[1]

[1]Centre for Ecology and Conservation, University of Exeter, Penryn, UK
[2]Department of Psychology, University of Macau, Macau
[3]Divison of Psychology, Communication and Human Neuroscience, University of Manchester, Manchester, UK
[4]Department of Applied Social Sciences, Hong Kong Polytechnic University, Hong Kong

LC, 0000-0001-6457-0254; AM, 0000-0002-7740-1625

A necessary prerequisite for the accumulation of beneficial knowledge, or 'cumulative cultural evolution', is the sharing of information via social learning. Yet little work in the field of cultural evolution has examined the mechanisms that support information sharing in the face of exploitative information free-riding and information hoarding. We ran a series of online interactive experiments (N = 716) combined with computational reinforcement and social learning models to test whether the mechanism of reputation-based partner choice can effectively support information sharing. Participants in groups chose whether to (i) engage in costly innovation and (ii) whether to share the resultant knowledge. Sharers received increased reputations for sharing and participants could use reputations to select recipients of knowledge. We found strong priors for information sharing that persisted throughout the experiments in participants from both the UK and China (study 1). However, partner choice was generally too weak to explain the presence of widespread information sharing, which persisted even when we reduced the benefit of innovation (study 2), introduced exploitative bots (study 3) and removed reputations altogether (study 4). Our results suggest that indiscriminate, group-based sharing is more important for facilitating cumulative cultural evolution than discriminate reputation-based sharing.

1. Introduction

Humans are distinctive in our capacity for cumulative cultural evolution (CCE), where populations accumulate ever-improving knowledge, technologies and customs over successive generations [1–3]. A key facilitator of CCE is high-fidelity social learning: learners must be able to acquire accumulated cultural knowledge from demonstrators with sufficient accuracy [4–6]. Often neglected, however, is the demonstrators' role in choosing whether to allow potential learners to access their cultural knowledge in the first place. This relates to an often-overlooked cooperative dilemma at the heart of CCE [7]. CCE requires both innovation (asocial learning) to create new knowledge and copying (social learning) to preserve and accumulate knowledge across generations [3]. The former is assumed to be more costly than the latter: it takes more time and effort to invent something new than copy it from someone else. Consequently, social learners can act as 'information free-riders' by copying beneficial knowledge from innovators without bearing the costs of innovation [7,8]. To avoid exploitation, innovators can refuse to innovate and themselves become free-riding social learners. But if no-one is innovating, then CCE stops. Alternatively, innovators might become 'information hoarders', refusing to let others copy their beneficial knowledge [7]. Again, though, CCE stops, as others cannot build on this beneficial knowledge.

Many examples through history show how information free-riding and hoarding can inhibit CCE. Medieval guilds and apprenticeships, such as those controlling glass-making in Venice [9] or woollen yarn in Germany [10], stifled innovation by hoarding knowledge [11]. Mokyr [12] argues that the Enlightenment occurred when information free-riding and hoarding were replaced with widespread information sharing amongst European networks of scientists and inventors. The recent 'open science' movement similarly represents a shift from scientists hoarding information for their own benefit to the open release of data and methods, facilitating the identification of replicable results that can be built on by others [13].

Our aim here is to examine when and why demonstrators share information with potential learners and the consequences for CCE. We test the mechanism of reputation-based partner choice [14], building on previous models [7] which explored how mechanisms from the evolution of cooperation literature might solve informational cooperative dilemmas. Under partner choice, demonstrators who share information with learners increase their reputation for sharing, and learners with high reputations for sharing are preferentially chosen by demonstrators as recipients of sharing. Consequently, the costs of innovation borne by demonstrators are paid back by preferentially receiving beneficial information from others. Information free-riders and information hoarders are excluded from these benefits as they have nothing to share or refuse to share, respectively, hence have low reputations and are not chosen as recipients of beneficial information.

Reputation-based partner choice has likely been a significant driver of innovation and information sharing in human CCE [7], which typically involves non-kin (precluding kin selection) interacting in wide networks of cooperators (precluding direct reciprocity) who can choose with whom to interact rather than being randomly paired up (precluding indirect reciprocity [15]). Reputation-based partner choice has been identified as crucial for information sharing by both historians [12] and in contemporary settings such as contributions to open source projects like Wikipedia [16].

Most previous CCE models and experiments preclude information free-riding and hoarding by design. Learners can typically automatically access demonstrators' knowledge, and demonstrators have no choice whether or not to share their knowledge. In one exception [17], participants could set 'informational access costs' that others had to pay in order to access their knowledge. There, high-scoring participants set excessively high access costs that no other participant was willing to pay, effectively blocking social learning. While this illustrates a strong tendency for information hoarding, the experiment did not incorporate any mechanism to facilitate information sharing, nor did it feature true open-ended CCE. Elsewhere, a recent study of advice-giving [18] allowed participants to choose whether to share advice with others, and in one condition incorporated reputation. However, advice sharing occurred indiscriminately with all other participants by design and could not be targeted at specific learners (precluding partner choice), while reputations reflected advice accuracy rather than sharing behaviour and had no real consequences for the sharer. Hence neither the fundamental cooperative dilemma of information hoarding, nor solutions to that dilemma, were possible.

Here we ran a series of CCE experiments in which information free-riding and information hoarding were possible, and where reputation-based partner choice could be used to maintain innovation and information sharing and facilitate CCE. While models [7] show that partner choice can in theory maintain open innovation despite the temptation to free-ride and hoard, experimental evidence is crucial for determining whether real people actually do this, and under what conditions. This is particularly important given findings from behavioural economics that people often fail to behave 'rationally' in non-informational cooperative dilemmas [19].

Study 1 provided an initial test of partner choice in participants from the UK and China. Study 2 addressed several potential confounds in study 1. In both studies participants predominantly engaged in costly innovation and information sharing, facilitating CCE. Yet partner choice was too weak to explain this, with most participants sharing indiscriminately rather than based on sharing reputation. Study 3 introduced artificial bots who exclusively engaged in either information hoarding (study 3a) or free-riding (study 3b), allowing participants to see that these strategies do better than indiscriminately sharing. Study 4 removed reputation entirely, preventing partner choice from being used. Yet in both cases information sharing remained the most common strategy amongst our participants. We end by discussing potential reasons for this mismatch between our participants' behaviour and theoretical expectations.

2. Study 1

(a) Methods

Participants were presented with a fictional, unfamiliar scenario to avoid triggering pre-existing biases regarding real-world information sharing. Participants played the role of space explorers who had to identify new elements on an unfamiliar planet (see figure 1 for a schematic of the task, and electronic supplementary material, appendix S1 for screenshots and instructions). Participants were randomly allocated to groups of n = 8 interacting in real time on oTree [20]. On each trial participants made two primary independent and binary decisions. First, whether to innovate or not (figure 1, panel 1). This was framed as using a Rock Analysis Unit to attempt to discover a new element. Innovation had a probability p = 0.75 of success. Participants were informed that innovation was not certain, but not the exact probability. If successful, participants were informed that one randomly selected element from a set of 1000 had been added to their element repository (figure 1, panel 2). Innovation cost c = 10 credits whether successful or unsuccessful. There was no cost to not innovating.

The second decision was whether to share information or not (figure 1, panel 3). Participants choosing to share could select one other group member to receive knowledge of their elements (figure 1, panel 4). Here participants could see every other

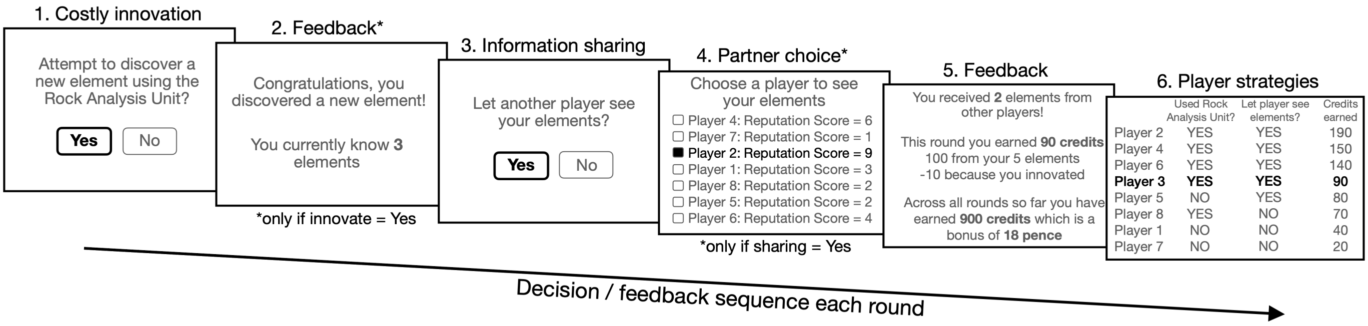

Figure 1. Schematic of the decisions taken and feedback received by participants on each round of the experiments (NB these are simplified versions of what participants saw; for actual screenshots see electronic supplementary material, appendix S1). Participants (1) chose whether to engage in costly innovation or not, and (2) those who did potentially discovered a new element. They then (3) chose whether to let another group member copy their elements, and (4) those who did could choose which group member based on reputations for past sharing. All participants then (5) received feedback regarding elements received and payoff on this and all rounds. Finally, participants could (6) see the decisions and round payoff of other group members, allowing payoff-biased copying of strategies: in studies 2–4 they saw all other group members as per the figure; in study 1, they saw only one randomly selected group member.

group member's randomly-generated ID number and their reputation score, which in study 1 was the total number of elements that the group member had shared with anyone else in all previous trials. The chosen recipient acquired knowledge of all elements that the sharer knew but the recipient did not, and the sharer's reputation increased by $r = N_E$ where N_E is the number of elements successfully copied. In study 1, choosing not to share incurred a hoarding cost of $d = 5$ credits. Note the terms 'share' or 'sharing' were not used in the experiment to avoid priming participants that their decisions were necessarily acts of cooperation, and were instead explained neutrally as 'Do you want to let another participant see your elements?'. Following the sharing choice, participants were informed how much their own reputation increased due to their decision, and how many elements they learned via others sharing information with them.

Following these two decisions, participants received $b = 20$ credits for every known element in their element repository. This benefit was received on the trial the element was discovered or copied and every subsequent trial. Participants were told their payoff on each trial and cumulative payoff across all trials in both credits and real money (figure 1, panel 5). Finally, they were shown the decisions (innovate or not; share or not) of one other randomly chosen group member, and whether that participant received a higher, lower, or equal payoff to them (figure 1, panel 6). This was framed as using a Scanner to eavesdrop on another planetary miner. The Scanner allowed participants to copy other participants' strategies and potentially copy the strategy of participants performing better than them (i.e. payoff-biased strategy copying: [7]). Reputations were not shown by the Scanner. Conversely, payoffs and decisions were not shown during information sharing, so that sharing decisions were based on reputation rather than performance. To keep each group synchronized in real time, participants had 30 s to make each decision or read feedback, otherwise the computer made a default decision for them (not innovate/not share).

The two independent binary decisions (innovate or not; share or not) gave four possible strategies describing a participant's behaviour on each trial: Open Innovators both innovated and shared; Closed Innovators innovated but did not share; Open Non-innovators did not innovate but shared; and Closed Non-innovators neither innovated nor shared. Open Innovators are information sharers who also bear the cost of innovation and thus contribute to CCE; Closed Innovators are information hoarders who benefit from private knowledge but do not share that knowledge with others; and Open and Closed Non-innovators are information free-riders who do not bear the cost of innovation.

Study 1 comprised 16 groups of 8 participants ($n = 128$; gender: 73 male, 52 female, 3 other; age: mean = 38.72 years, s.d. = 11.91, range = 19–78) recruited online in the UK via Prolific and another 16 groups of 8 participants ($n = 128$; gender: 53 male, 72 female, 1 other; age: mean = 21.5, s.d. = 4.98, range = 18–52) in Macau, China who completed the same experiment online in Cantonese. Participants were assigned to groups on a first-come-first-served basis: as each participant joined the session, they were placed into the next available group until the group of eight was complete. Prolific provides higher quality data than other online recruitment services [21]. UK participants were paid £6 per hour for an estimated time of 1 h plus a bonus based on their final payoff. Chinese participants were paid 40 RMB for an estimated time of 1 h plus a performance bonus. Comprehension of task instructions was assessed via five multiple choice questions (see electronic supplementary material, appendix S1), which were answered with high accuracy (UK: mean = 4.30, s.d. = 0.97; China: mean = 4.27, s.d. = 0.73).

(b) Computational model and predictions

Before running study 1, we created computational models in which simulated participants ('agents') played the experimental task (see preregistration at https://osf.io/fk8wy). These models extend previous models of partner choice and payoff-biased strategy copying [7] to incorporate reinforcement learning [22] as per previous experiments [23,24]. The models were used to generate preregistered predictions and ensure that reputation-based partner choice could in theory maintain information sharing. Following data collection, the same models were fit to real participants' data using multilevel Bayesian parameter optimization in stan [25] via cmdstanr [26] in R [27]. This generated posterior distributions for each participant's, and the entire sample's, parameter values, allowing parameters to vary across participants and incorporating non-independence of participants within groups.

The models (see electronic supplementary material, appendix S2) contain nine parameters determining agents' behaviour. α ($0 \le \alpha \le 1$) and β ($\beta \ge 0$) are the updating rate and inverse temperature of the Q-learning reinforcement learning algorithm [22], where weights (Q values) associated with each strategy (Open Innovation, Closed Innovation, Open Non-innovation and Closed Non-innovation) determined the probability of choosing each strategy on each trial and were updated according to payoffs received. Higher α gives faster updating of Q values due to payoffs. Higher β makes agents more likely to choose the strategy with the highest Q value. Lower β makes agents more exploratory with $\beta = 0$ giving random choice. p_{S1} and p_{S0} control payoff-biased strategy copying via the Scanner. p_{S1} is the probability of switching to the demonstrator's strategy if they scored higher than the agent. p_{S0} is the probability of switching to the demonstrator's strategy if they scored lower than or equal to the agent. Payoff-biased strategy copying occurs if $p_{S1} > 0$ and $p_{S1} > p_{S0}$. β_{rep} ($\beta_{\text{rep}} \ge 0$) controls the strength of reputation-based partner choice. Higher β_{rep} makes agents more likely to share with the player with the highest reputation, lower β_{rep} causes agents to also pick lower-reputation players, and $\beta_{\text{rep}} = 0$ makes partner choice random and independent of reputation. Four parameters, $Q1_{OI}$, $Q1_{CI}$, $Q1_{ON}$ and $Q1_{CN}$, give agents' prior probabilities of using each strategy at trial 1. These were all equal in the simulations given a lack of knowledge about real participants' priors. All parameters were also identical across agents given lack of prior information about how parameters vary across individuals.

Model analysis (electronic supplementary material, figure S1) showed that Open Innovation emerges when there is a combination of: (i) reinforcement learning, where α and β are sufficiently high that agents learn which strategy optimizes payoffs, with α greater than zero and β greater than around 0.4; (ii) partner choice, where β_{rep} is greater than around 0.4 such that agents select recipients of sharing based on reputation rather than randomly; and (iii) payoff-biased strategy copying, where the frequency of Open Innovation when favoured by partner choice increases with p_{S1}. Our preregistered predictions were therefore that (i) posterior distributions of α, β, β_{rep} and p_{S1} all have 95% credible intervals (CIs) greater than zero; (ii) the frequency of Open Innovation in the final 50% of trials is higher than any other strategy as indicated by non-overlapping 95% CIs; and (iii) the frequency of Open Innovation across participants should predict the total number of accumulated elements, given that information sharing plus innovation should facilitate CCE. We also predicted that (iv) Open Innovation should initially be more frequent in the Chinese than UK sample, given stronger tendencies for social learning in collectivistic East Asian societies than individualistic Western societies [28,29]. Hence, the 95% CIs for the initial Q value for Open Innovation, $Q1_{OI}$, should be higher in the UK than Chinese sample. Given that partner choice should be a universal mechanism, we predicted no differences in subsequent use of reputation or frequency of Open Innovation (i.e. all other parameters should have overlapping UK and Chinese 95% CIs).

The model highlights a key difference between our information sharing game and the public goods game (PGG) [19]. In a PGG, players divide resources (e.g. money) between a public pot and their private holding, with the pot shared equally across all players. The optimal strategy is to free-ride and contribute nothing to the pot: both free-riders and contributors get the same equal share of the pot, but free-riders always have greater private holding than contributors. Our game is different because of the non-rivalrous nature of information [30]. Participants who share elements still have knowledge of, and benefit from, those elements, unlike resources in a PGG which can either be contributed or kept, not both. Consequently, in the absence of partner choice, non-sharers (e.g. Closed Innovators) do not always have higher absolute payoffs than sharers (e.g. Open Innovators), because the latter do not lose their elements when they share. Assuming no partner choice and indiscriminate sharing, and with no payoff-biased strategy copying, players should be indifferent between sharing and non-sharing (electronic supplementary material, figure S1B, bottom-left grey area). With partner choice, Open Innovation should be favoured as Closed Innovators will be excluded from receiving others' elements (electronic supplementary material, figure S1B, upper orange area). Adding payoff-biased strategy copying favours Closed Innovation when partner choice is weak or absent (electronic supplementary material, figure S1B, bottom-right blue area) because, while Open and Closed Innovation have equal absolute payoffs, Closed Innovators have higher relative payoffs due to their private elements. To maintain a similar incentive structure to a PGG, our participants were paid in proportion to absolute rather than relative payoffs. Nevertheless, we expect participants to be sensitive to the performance of other players and engage in payoff-biased strategy copying, consistent with previous cooperation experiments incorporating social learning [31,32]. To reiterate, even without payoff-biased strategy copying based on relative payoffs, Open Innovation should still not be the dominant strategy unless partner choice is sufficiently strong.

(c) Results

As predicted, Open Innovation was the most frequent strategy in the final 50% of trials in both the UK (figure 2A) and China (figure 2B), as indicated by non-overlapping CIs with other strategies (electronic supplementary material, table S1). The frequency of Open Innovation was equally high at the start of the experiment and throughout all 20 trials, indicating strong priors for Open Innovation that changed little over time.

Partner choice was not random, but it was not strong. Figure 2C,D shows that while the group member with the highest reputation was chosen as the recipient of sharing more than chance, lower ranking players were also often chosen, even players with the lowest reputations. This is reflected in model estimates of β_{rep} greater than zero (UK: $\beta_{\text{rep}} = 0.14[0.06,0.23]$; China: $\beta_{\text{rep}} = 0.29[0.19,0.40]$; see figure 2E,F and electronic supplementary material, table S3) but lower than 0.4 which our simulations indicated were needed to maintain Open Innovation. However, the sample estimates of β_{rep} conceal considerable variation across participants. Electronic supplementary material, figure S3 shows the distribution of β_{rep}, plus example reputation rank plots for participants with high ($\beta_{\text{rep}} \gg 0.4$), threshold ($\beta_{\text{rep}} \approx 0.4$) and low ($\beta_{\text{rep}} \approx 0$) values of β_{rep}. Most participants (UK: 79.7%; China: 68%) had β_{rep} well below the 0.4 threshold, with the remaining minority having large β_{rep} denoting effective partner choice.

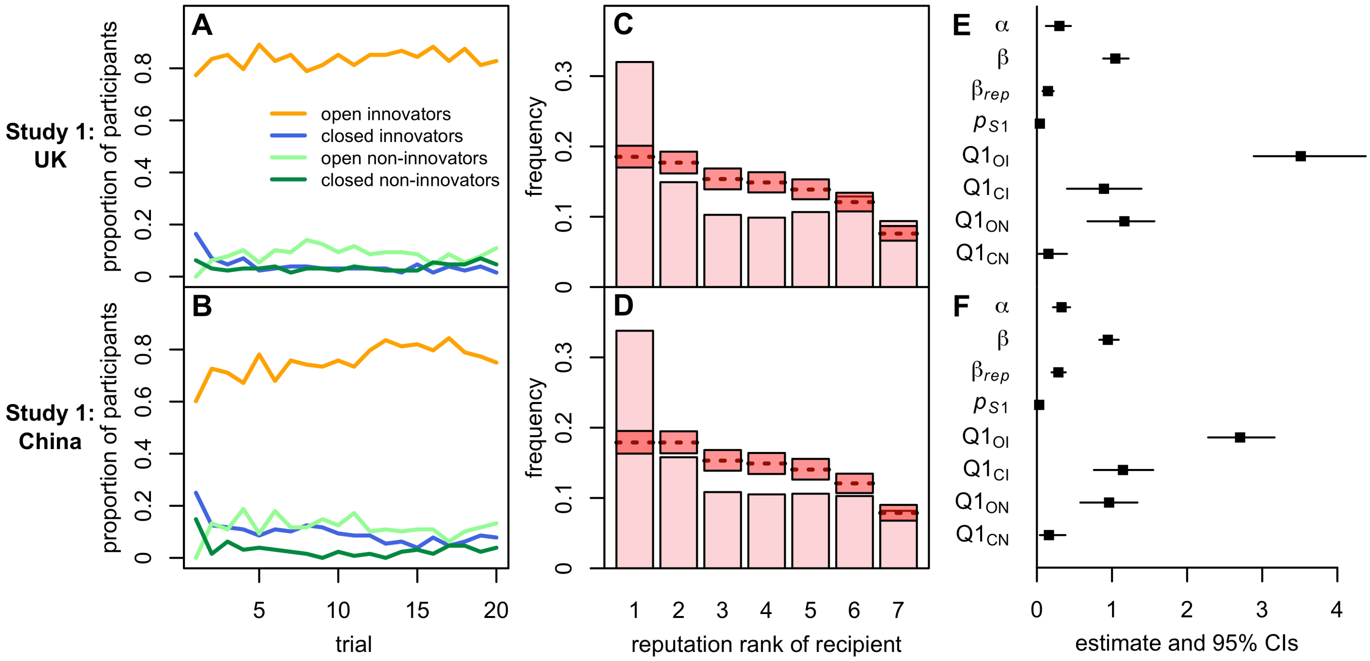

Figure 2. The proportion of participants in study 1 who engaged in each strategy over 20 trials for the (A) UK and (B) China sample. See electronic supplementary material, figure S2 for frequencies in each group. (C) and (D) show, for the UK and China respectively, the frequency with which participants chose recipients of sharing based on the reputational rank of other group members, where 1 is the player with the highest reputation and 7 the lowest reputation on that trial. Light pink bars show the actual frequencies, while dotted lines and dark pink boxes show the mean and 95% CIs, respectively, for the frequencies that would be expected if choice was random. Expected frequencies are not the same for all ranks because of ties; where two or more group members had the same reputation score then there were fewer than seven unique ranks, so lower ranks would be expected less irrespective of reputation score. Expected frequencies were generated via simulation assuming random choice but with the same tie structure as the real data. (E) and (F) show mean estimates and 95% CIs of the eight key parameters describing participant behaviour as derived from model fitting (see text for details and electronic supplementary material, table S2 for parameter estimates).

Figure 2E,F shows evidence for reinforcement learning with updating rate (UK: $\alpha = 0.33[0.17,0.52]$; China: $\alpha = 0.33[0.20,0.48]$) and inverse exploration (UK: $\beta = 1.04[0.86,1.28]$; China: $\beta = 0.94[0.82,1.11]$) parameters greater than zero, although only β unambiguously differed from zero using explicit model comparison (electronic supplementary material, table S3). This suggests little updating due to payoffs, and a preference for the highest-weighted strategy, consistent with the constant selection of Open Innovation. Contrary to predictions, there was virtually no payoff-biased strategy copying, with p_{S1} indistinguishable from zero (UK: $p_{S1} = 0.04[0.01,0.08]$; China: $p_{S1} = 0.03[0.01,0.06]$; electronic supplementary material, table S3) and overlapping almost exactly with p_{S0} (UK: $p_{S0} = 0.03[0.01,0.06]$; China: $p_{S0} = 0.03[0.01,0.06]$). There were higher initial Q values for Open Innovation ($Q1_{OI}$) than for any other strategy. The lack of variation in strategy choice meant that the frequency of Open Innovation did not predict the number of accumulated elements across participants (electronic supplementary material, table S4). Also contrary to predictions, Chinese participants did not have stronger initial preferences for Open Innovation than UK participants given overlapping CIs of $Q1_{OI}$ (UK: $Q1_{OI} = 3.48[2.68,4.44]$; China: $Q1_{OI} = 2.72[2.26,3.27]$). There were no meaningful differences in any behavioural parameters, nor strategy frequencies, between the UK and China samples as indicated by overlapping CIs. This suggests that these findings are not specific to a society with particular norms regarding information sharing or innovation, and are replicable at least across these two countries.

3. Study 2

(a) Methods

Results from study 1 were ambiguous. While Open Innovation was the most frequent strategy, partner choice was not sufficiently strong to explain this behaviour, with most participants sharing indiscriminately rather than based on reputation. Moreover, there was virtually no payoff-biased strategy copying. Under these conditions we would expect participants to be indifferent between Open and Closed Innovation (electronic supplementary material, figure S1B, bottom-left grey area), yet Open Innovation dominated.

Study 2 addressed several methodological issues that potentially explain these results. Due to resource constraints, studies 2–4 were conducted only in the UK. The revised study design and hypotheses were preregistered (https://osf.io/rhefc). Hypotheses were the same as for study 1 minus the cross-cultural comparison. The changes were to (i) remove the cost of hoarding paid when participants chose not to share (i.e. $d = 0$), which while models showed did not substantively affect strategy payoffs, may have been a psychological barrier to hoarding and explain why Closed Innovation was rare; (ii) increase the number of trials from 20 to 30, allowing more time to explore strategy payoffs; (iii) limit sharers' reputation increase per trial to $r = 1$ point when at least one element was learnt, so that participants are not incentivized to share with low-reputation players with few elements in order to increase their reputation by several points in a single trial; (iv) remove feedback regarding how much participants' own reputations increased, because when participants chose to share with recipients who already knew all the sharer's elements then no elements were actually shared and the sharer's reputation did not increase, giving the misleading signal that sharing does not increase reputation; (v) increase the potential for payoff-biased strategy copying by showing on each trial a leaderboard containing the payoffs and choices of all group members, not just one randomly-chosen player. Finally, (vi) we added a condition reducing the benefit of innovation, b. Study 2a had high benefit, $b = 20$, as in study 1. Study 2b had low benefit, $b = 5$. Reducing the benefit of innovation should reduce the payoffs of Open and Closed Innovation, given that innovation yields lower benefits. Hence, we can test whether Open Innovation emerges despite lower benefits of innovation, and where Open or Closed Non-innovation (i.e. information free-riding) represents the alternative exploitative strategy, rather than Closed Innovation (i.e. information hoarding).

Our models were revised with new parameter values and the new payoff-biased strategy copying. The latter now has an overall probability p_s of engaging in strategy copying and a parameter β_{PB} ($\beta_{PB} \geq 0$) controlling the extent to which the strategy of high-scoring players was copied rather than copying randomly. Electronic supplementary material, figure S4 shows revised model results. Due to the smaller magnitude and variation of reputation scores, β_{rep} must now be larger to support Open Innovation, around $\beta_{\mathrm{rep}} > 1.25$ for $b = 20$ but increasing for higher p_s, and $\beta_{\mathrm{rep}} > 2$ for $b = 5$ again increasing for higher p_s. Generally, more-potent payoff-biased strategy copying means that a much smaller region of the parameter space now favours Open Innovation.

Ten groups of 8 participants completed study 2a ($n = 80$; gender: 41 male, 38 female, 1 other; age: mean = 39.36, s.d. = 11.85, range = 20–77), and another 10 groups of 8 participants completed study 2b ($n = 80$, gender: 42 male, 37 female, 1 other; age: mean = 39.62, s.d. = 11.60, range = 20–67), all run online via oTree and recruited via Prolific. No participant took part in more than one session, either within or across studies. Comprehension of task instructions was very good (study 2a mean = 4.50, s.d. = 0.76; study 2b mean = 4.26, s.d. = 0.95).

(b) Results

Despite the methodological changes, results from study 2a with high benefit of innovation ($b = 20$) were virtually identical to study 1. The most common strategy was again Open Innovation (figure 3A) which had higher, non-overlapping 95% CIs in the final 50% of trials (electronic supplementary material, table S1) and higher initial Q values (figure 3C). As before, while partner choice was not random ($\beta_{\mathrm{rep}} = 0.54[0.19,1.07]$; figure 3C; electronic supplementary material, table S3), the sample-wide β_{rep} estimate was still below the predicted threshold of 1.25 from the models. Reputation barplots (figure 3B) show lower absolute frequencies of picking low-reputation group members, although this is partly because with less variance in reputations there were more ties so very low ranks were often not present. Even taking this into account, lower reputation players were still picked at around chance, rather than never as partner choice requires. The distribution of β_{rep} across participants (electronic supplementary material, figure S8) showed a majority (67.5%) with low $\beta_{\mathrm{rep}} < 1.25$ inconsistent with partner choice. There was again no relationship between frequency of Open Innovation and accumulated elements (electronic supplementary material, table S4). Reinforcement learning was evident ($\alpha = 0.29[0.21,0.37]$; $\beta = 0.92[0.79,1.10]$) and payoff-biased strategy copying was higher than in study 1 but still not strong ($p_s = 0.07[0.02,0.12]$) and did not differ from zero according to model comparison (electronic supplementary material, table S3). When it was used, demonstrators were chosen based on payoff, albeit with much uncertainty ($\beta_{PB} = 2.85[0.15,11.03]$).

Reducing the benefit of innovation to $b = 5$ in study 2b reduced the frequency of Open Innovation (figure 3D) compared to study 2a. Nevertheless, the frequency of Open Innovation in the final 50% of trials was still higher than any other strategy as indicated by non-overlapping CIs (electronic supplementary material, table S1). In the final trials Open Non-innovation became more frequent than Open Innovation, likely because participants knew the experiment ended at trial 30 and it took at least three trials for the accumulated benefit of innovation $b = 5$ to exceed the cost of innovation $c = 10$. Model fits otherwise revealed similar parameter values for study 2b as for study 2a, except a lower initial Q value for Open Innovation in study 2b ($Q1_{OI} = 1.52[1.06,2.12]$) than study 2a ($Q1_{OI} = 3.39[2.65,4.35]$). Thus, participants did not just learn that Open Innovation was less effective in study 2b, they anticipated it from the start. The strength of partner choice ($\beta_{\mathrm{rep}} = 0.65[0.33,1.07]$) was slightly higher in study 2b than study 2a but with an identical upper CI, suggesting that partner choice was still not sufficiently strong to explain the initial emergence of Open Innovation. The distribution of β_{rep} across participants in study 2b (electronic supplementary material, figure S8) showed a majority (65%) with low $\beta_{\mathrm{rep}} < 1.25$ inconsistent with partner choice. The frequency of Open Innovation did not predict the number of accumulated elements in study 2b (electronic supplementary material, table S4).

Overall, study 2 replicated the findings of study 1 but with several methodological improvements ruling out artefacts and alternative explanations. Open Innovation was again the most frequent strategy despite weak partner choice. Reducing the benefit of innovation ($b = 5$) reduced the frequency of Open Innovation, especially in the final trials when the cost of innovation exceeded its expected future benefits. However, Open Innovation remained the most common strategy, and partner choice remained weak, even when innovation was less beneficial making Open Innovation in theory less likely to emerge. Payoff-biased strategy copying was slightly more likely, but still very weak, and so did not facilitate the spread of Closed Innovation or Closed Non-innovation as predicted.

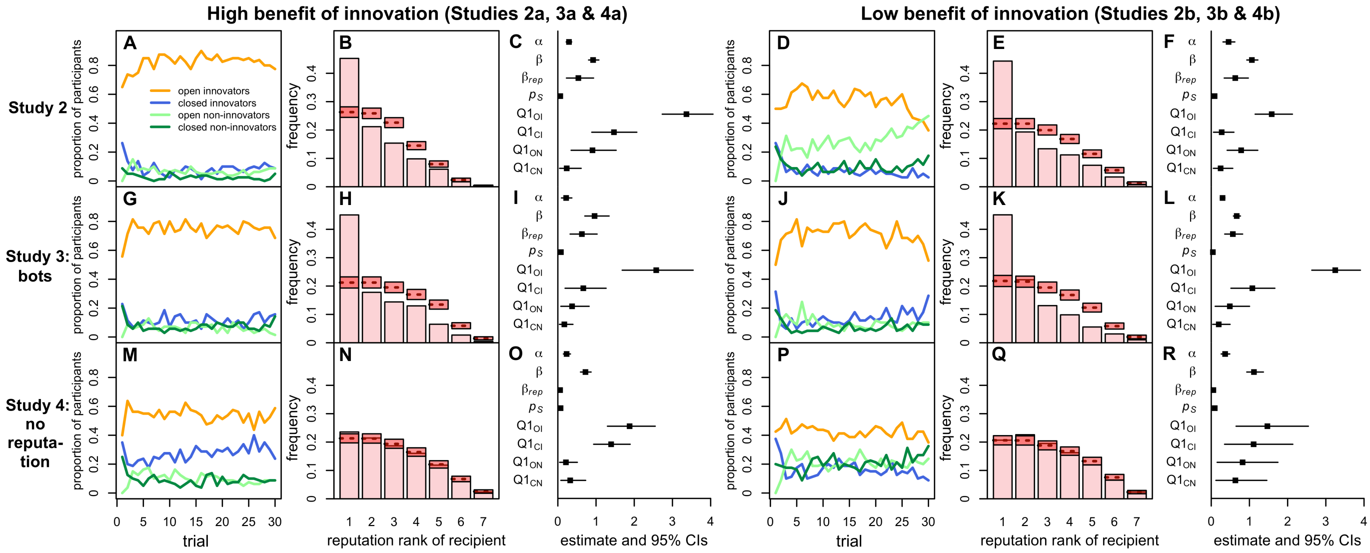

Figure 3. Strategy frequencies and reputation rank choices for Studies 2–4. See caption of figure 2 for details. Left panels show the condition with high benefit of innovation $b = 20$, right panels show low benefit of innovation $b = 5$. Study 2 (panels A–F) had several methodological improvements compared to study 1 (see text for details). Study 3 (panels G–L) had one bot in each group who always engaged in Closed Innovation (for $b = 20$) or Closed Non-innovation (for $b = 5$). Study 4 (panels M–R) removed reputation information, making partner choice strictly random. See electronic supplementary material, figures S5–S7 for strategy frequencies by group, and electronic supplementary material, tables S5–S7 for parameter estimates.

4. Study 3

(a) Methods

Study 3 addressed two related issues with studies 1 and 2. First, lack of variation in strategies: unlike in PGGs, where a minority of participants typically always free-ride [33], in studies 1 and 2 few participants engaged in Closed Innovation (or Closed Non-innovation for study 2b). Consequently, Open Innovators could not see via the leaderboard that Closed strategies could out-perform Open Innovation. Second, lack of variation in reputation: because most participants engaged in Open Innovation, most participants had similar reputation scores. Without variation in reputations, partner choice is ineffective.

Consequently, in study 3, we artificially introduced exploitative strategies via one computer-controlled bot per group. In the $b = 20$ condition (study 3a) bots always played Closed Innovation (i.e. never shared but always attempted to innovate). In the $b = 5$ condition (study 3b) bots always played Closed Non-innovation (i.e. never shared and never innovated). Model results indicate that these strategies out-perform Open Innovation when partner choice is weak (electronic supplementary material, figure S4). Participants were unaware that these were artificial bots until a post-experiment debrief. Study 3a comprised 10 groups of 7 real participants with $b = 20$ ($n = 70$; gender: 42 male, 37 female, 1 other; age: mean = 39.63, s.d. = 11.60, range = 20–67), and study 3b comprised 10 groups with 7 real participants with $b = 5$ ($n = 70$; gender: 39 male, 31 female; age: mean = 38.1, s.d. = 11.76, range = 18–68). Comprehension of task instructions was very good (study 3a mean = 4.49, s.d. = 0.81; study 3b mean = 4.50, s.d. = 0.81). Hypotheses and study design were preregistered (https://osf.io/rhefc). Hypotheses and analyses were the same as for study 2, but with bots removed from analyses.

(b) Results

Despite the addition of exploitative bots, study 3 had very similar results to study 2. Open Innovation was again the most frequent strategy for both $b = 20$ (figure 3G) and $b = 5$ (figure 3J). Model parameters were similar to study 2, and similar across conditions (figure 3I,L). Most relevantly, the strength of partner choice was no different in study 3 (study 3a: β_{rep} = 0.63[0.30,1.09]; study 3b: β_{rep} =0.56[0.31,0.90]) than study 2, indicating that the presence of exploitative bots did not increase the use of reputation. A majority of participants in study 3a (58.6%) and study 3b (61.4%) had low $\beta_{rep} < 1.25$ inconsistent with partner choice (electronic supplementary material, figure S8). The frequency of Open Innovation again did not predict the number of accumulated elements in study 3, after removing bots (electronic supplementary material, table S4).

Analysis of sharing patterns showed that the exploitative bots were mostly excluded from sharing (electronic supplementary material, figures S9 and S10). Bots were shared with less than most other real participants, with a mean sharing rank of 6.4 out of 8 in study 3a and 6.7 in study 3b. Nevertheless, this exclusion was not perfect: even though bots had a fixed reputation score of zero throughout the experiment, in 7 of the 10 study 3a groups and 6 of the 10 study 3b groups the bots were *not* the player shared with least, and in no group were bots *never* shared with, as would be expected under partner choice. In study 3b, the exclusion from sharing led to bots performing worse than most genuine participants, with an average score rank of 7.1 out of 8. Without learning about elements via sharing, Closed Non-innovators knew few elements. In study 3a bots did better, with an average score rank of 2.9 out of 8, because Closed Innovators could benefit from private innovation. Indeed, in 4 of the 10 groups they finished with higher scores than any real participant. Yet participants still failed to copy this high performing strategy, indicated by the low frequency of Closed Innovation and low probability of payoff-biased strategy copying (p_s = 0.08[0.03,0.14]) which did not differ from zero (electronic supplementary material, table S3).

Participants in studies 2 and 3 were also asked on what basis they chose partners with whom to share. In study 2 around 40%, and in study 3 around 50–55% of participants reported selectively sharing with the group member who had the highest reputation, as would be expected under reputation-based partner choice (electronic supplementary material, table S8). This proportion was higher in study 3 than study 2, likely due to the exploitative bots. Often these responses invoked partner choice, showing an awareness that sharing with high-ranked players led to being preferentially shared with. The next highest category entailed randomly or indiscriminately sharing with all other group members. Some responses in this category again invoked cooperative motivations, indicating that sharing indiscriminately is the 'fairest' approach. A smaller proportion preferentially shared with low-ranking players. Again, these often invoked cooperative motivations, such as to incentivise low-ranking sharers to share more, or to equalise traits across the group consistent with a levelling mechanism.

5. Study 4

(a) Methods

While study 3 showed that non-sharing bots were largely excluded from sharing, it is still puzzling that participants failed to adopt the high-performing Closed Innovator strategy in study 3a. In study 4, we explored the limits of participants' tendencies to share information by removing reputation information altogether. If participants chose to share, the computer picked another group member at random with whom to share. There were no reputations and sharing had no consequences beyond providing others with element information. Theoretically, there is no reason to share element information given that sharing does not return any benefits in terms of an increased chance of being the recipient of sharing. Models (electronic supplementary material, figure S4) predicted that Closed Innovation (for b = 20 in study 4a) or Closed Non-innovation (for b = 5 in study 4b) should be favoured if participants engage in payoff-biased strategy copying, or indifference between Open and Closed Innovation (for study 4a) or Open Innovation and Closed Non-innovation (for study 4b) without payoff-biased strategy copying.

We ran 10 groups of 8 participants for study 4a where b = 20 (n = 80; gender: 34 male, 45 female; age: mean = 37.57, s.d. = 14.39, range = 18–76) and 10 groups of 8 participants for study 4b where b = 5 (n = 80; gender: 32 male, 45 female; age: mean = 36.21, s.d. = 11.15, range = 19–72). The experimental design and hypotheses were preregistered (https://osf.io/79emr). Participant comprehension was tested with two questions, excluding the three original questions that were related to reputation, and showed excellent understanding (study 4a mean = 1.93, s.d. = 0.27; study 4b mean = 1.79, s.d. = 0.41). We hypothesized that given the impossibility of partner choice, in study 4a Closed Innovation should have a higher frequency over the final 50% of trials than any other strategy as indicated by non-overlapping CIs, and in study 4b Closed Non-innovation should have the highest frequency. We also hypothesized a lower frequency of Open Innovation in study 4 than for the equivalent conditions in studies 2 and 3, given the absence of partner choice.

(b) Results

By design, the strength of partner choice was virtually zero (study 4a: β_{rep} = 0.06[0.03,0.12]; figure 3O; study 4b: β_{rep} = 0.06[0.02,0.11]; figure 3R; electronic supplementary material, table S3) with sharing frequencies at random levels (figure 3N,Q). Nevertheless, and contrary to our hypothesis, Open Innovation was again the most frequent strategy in the final 50% of trials of both study 4a (figure 3M) and study 4b (figure 3P) (electronic supplementary material, table S1) However, cross-study comparisons confirmed that, as predicted, the frequency of Open Innovation was lower in study 4 when partner choice was impossible than in studies 2 and 3 when partner choice was possible, for both the b = 20 and b = 5 conditions (electronic supplementary material, table S9). Indeed, 2 of the 10 groups in study 4a ended with Closed Innovation higher than Open Innovation, while 6 of the 10 groups in study 4b ended with Closed or Open Non-innovation higher than Open Innovation. So while Open Innovation was still the most common strategy in study 4, at least some participants had learned that without reputation, Closed Innovation (in study 4a) or Closed Non-innovation (in study 4b) were at least as good. Moreover, the strong initial preference for Open Innovation seen in previous studies was no longer present in study 4, with CIs for $Q1_{OI}$ estimates overlapping with those of other strategies (figure 3O,R). Reinforcement learning and payoff-biased strategy copying were largely identical to previous studies (figure 3O,R).

Electronic supplementary material, figure S11 shows the rate at which elements accumulated over time across all eight experiments. As expected, the predominance of Open Innovation facilitated the steady accumulation of elements across all experiments. While study 4 showed a slower rate of accumulation than the other studies, consistent with the lower frequency of Open Innovation, cultural accumulation still occurred despite the lack of partner choice.

6. Discussion

Our aim here was to test whether reputation-based partner choice can maintain a combination of information sharing and costly innovation, despite temptations to let others engage in costly innovation without oneself doing so (information free-riding) and to innovate but refuse to share information with others (information hoarding). Earlier models [7] showed that the mechanism of partner choice can maintain open innovation and hence facilitate CCE in the face of information free-riding and hoarding.

Reputation-based partner choice works by conferring high reputations on information sharers, and allowing individuals to preferentially share with those who have high reputations. Hence the initial cost of innovation and cost of advantaging others with beneficial information is outweighed by preferentially receiving beneficial information from others via reputation. Yet while this mechanism works in theory [7], and there is experimental evidence for partner choice in non-informational, material cooperative dilemmas [14], it is unknown whether people actually use reputation-based partner choice to maintain open innovation in the face of informational cooperative dilemmas.

Overall, our four studies showed limited evidence that partner choice is used to facilitate information sharing, with most participants showing a strong preference for costly innovation and indiscriminate information sharing irrespective of reputation. In study 1 (both UK and China samples) and study 2 (which addressed various potential methodological artefacts), Open Innovation, i.e. costly innovation plus information sharing, was clearly the dominant strategy. While partner choice was operating as indicated by the partner choice parameter (β_{rep}) exceeding zero, it was not strong, and showed substantial variation across participants. While some participants had high β_{rep} at levels which models predicted would support information sharing in the face of information free-riding and hoarding, a majority had β_{rep} close to zero, indicative of sharing at random irrespective of recipients' sharing reputations. This was true whether the benefit of innovation was high (study 2a) making Closed Innovation, i.e. information hoarding, the relevant exploitative strategy, or low (study 2b) making Closed Non-innovation, i.e. information free-riding, the relevant exploitative strategy. In study 3, unbeknownst to participants, we introduced bots who never shared and consequently had fixed reputations of zero. Genuine participants used partner choice to mostly, but not entirely, exclude these bots from sharing. Yet Open Innovation remained the most frequent strategy despite participants being able to see, at least in the high benefit of innovation condition (study 3a), that non-sharing bots out-performed indiscriminately sharing participants. In study 4, we removed reputation entirely, precluding partner choice by design. As predicted, Open Innovation was less frequent than in studies 1–3 and no longer initially favoured over other strategies. Yet Open Innovation remained the most common strategy overall, even though there was no mechanism by which information sharing could be rewarded.

Our computational models (electronic supplementary material, figures S1 and S4) highlight exactly how our participants' behaviour deviates from expectations. In the models, Open Innovation is only favoured when partner choice is strong (high β_{rep}). Here, only Open Innovators accrue reputations for sharing their elements and consequently only they receive beneficial information from others. Yet our participants predominantly employed Open Innovation even though partner choice was weak (low β_{rep}) or, in study 4, absent ($\beta_{rep} = 0$), such that Open Innovators were no more likely to receive beneficial information than participants playing non-sharing Closed strategies. In the models, when partner choice is weak or zero, the optimal strategy depends on the strength of payoff-biased strategy copying. When payoff-biased copying is strong (high p_{S1} or p_S), then exploitative Closed strategies do best. Closed strategies benefit from both public and private knowledge whereas Open strategies benefit only from public knowledge, so Closed strategies do relatively better and are preferentially copied. Yet our participants showed virtually no payoff-biased strategy copying (low p_{S1} or p_S indistinguishable from zero). Here, the models predict that participants should be indifferent between Open and Closed strategies. This is because whether you share makes no difference to others' sharing behaviour, so in absolute terms you neither gain nor lose from sharing because you receive beneficial knowledge regardless. In contrast to this prediction, our participants predominantly engaged in Open Innovation, and seldom Closed strategies. This is the puzzling aspect of our findings: given weak or absent partner choice and weak or absent payoff-biased strategy copying, participants have no incentive to prefer Open over Closed strategies, yet they predominantly do so.

That participants failed to engage in payoff-biased strategy copying conflicts with evidence from cultural evolution experiments that payoff-biased social learning is frequently used in general [17,23], and economic experiments in which payoff-biased copying promotes the spread of exploitative strategies [31,32]. Our participants even failed to employ payoff-biased copying in study 3a when they could see Closed Innovator bots performing well, and in 4 of 10 groups out-performing all genuine participants. The lack of payoff-biased copying may be because we rewarded participants based on absolute score. If instead we had rewarded participants based on relative group rank, we may have seen more payoff-biased copying, and consequently higher frequencies of Closed strategies. This may in turn have led to the establishment of reputation-based partner choice, and the emergence of Open Innovation due to strong partner choice as predicted. We plan to test this prediction in future work. If this is the case, it would suggest that partner choice is most likely to emerge with zero-sum relative payoffs rather than non-zero-sum absolute payoffs. This could be tested historically; indeed, many of the examples noted in the Introduction where partner choice facilitated information sharing probably involved relative payoffs (e.g. firms competing for market share; scientists competing for priority).

As noted above, however, even with absolute payoffs we would still not expect the high frequency of Open Innovation that was observed. One potential proximate explanation is that people are driven to increase their reputation by indiscriminately sharing, even when this indiscriminate sharing leaves them vulnerable to exploitation by non-sharers. Psychologists have shown that the desire for status or reputation is a fundamental human motive tied to self-esteem, wellbeing and physical health outcomes [34]. Consequently, our participants could have been maximizing reputations rather than payoffs. However, while this might explain some of our participants' behaviour, at least in initial trials, it cannot explain why information sharing remained the most common strategy in study 4 when reputation was removed entirely. Moreover, very few participants (5/300) in studies 2 and 3 reported being motivated to share indiscriminately purely because they wanted to maximize their own reputation.

Another potential explanation is that our participants, and people generally, have acquired culturally evolved norms to over-share information within groups, with weak tendencies to monitor and respond to information hoarding and free-riding. These norms may have evolved via cultural group selection [7,35], where groups of sharers accumulate information faster and hence outperform groups of non-sharers. This fits with some participants' self-reported strategies of sharing indiscriminately to maximize the spread of elements across their group, or of levelling knowledge across the group by sharing with low reputation group members. There may therefore be a trade-off between individual- and group-level interests relating to information sharing. Highly discriminate sharing only with the highest reputation members of one's group (high β_{rep}) maximizes individual interests and relative payoffs by protecting against exploitation within groups, but at the expense of hindering the spread of information across the group. Indiscriminate sharing (β_{rep} = 0) maximizes group interests and absolute payoffs by rapidly spreading information across the group but at the expense of vulnerability to exploitation. The weak partner choice seen in our experiments may therefore reflect a balance between past individual- and group-level selection: too weak to maintain Open Innovation on its own, but strong enough when combined with inter-group competition to favour Open Innovation. While we cannot test this in our data given that our groups were independent with no inter-group competition, we can model this scenario to evaluate its plausibility. Electronic supplementary material, figure S12 shows that when inter-group competition is added to our existing partner choice models, Open Innovation is favoured at lower values of β_{rep} than when inter-group competition is absent. Hence, the low β_{rep} found in our experiments may be a product of past cultural group selection, with individual-level partner choice preventing it becoming entirely indiscriminate. Future empirical research might test this by incorporating inter-group competition to see whether people are more responsive to group-level incentives than they were to individual-level incentives in our current experiments. In particular, the group selection hypothesis predicts that participants should share information within groups but not between them, if groups are in competition.

The strong tendency to share information observed in our experiments fits with previous experimental findings from the advice literature, where participants generally choose to leave honest advice for others [18], albeit where there is no temptation to free-ride or hoard information, nor any mechanism to maintain information sharing in the face of such temptations. Our studies show that information sharing persists even in such contexts. Our results do not match those of [17] where participants could set informational access costs and high performing participants set excessively high access costs that prohibited any information sharing. Perhaps when participants can potentially profit from 'selling' their information, as in [17], different dynamics emerge compared to the present experiments which entailed a binary choice of whether to share or not, and the benefit of sharing was more diffuse, i.e. preferentially receiving information in the future. Alternatively, it may be because participants in [17] were not paid monetarily and so had less incentive to maximize points in the task and more incentive to maximize relative performance rank in their group, as discussed above.

Despite the improvements of study 2, several methodological limitations remain. First, our implementation of CCE as the linear accumulation of independent elements only captures one simple kind of CCE, sometimes described as 'core CCE' [3]. More elaborate tasks might implement trait interdependencies, branching lineages or recombination ('extended CEE' [3]), or qualitative shifts to new design spaces [36]. These complexities may increase the difficulty of innovation and reduce people's willingness to indiscriminately share information that is more costly to acquire. Second, the timeframe of our experiment, while increased in studies 2–4, was still relatively short. Increasing the number of timesteps further may give more opportunity for hoarding and free-riding to emerge and the partner choice system to be established. Third, innovation and information sharing were the only route to higher payoffs. In reality, costly innovation and information sharing may be compensated by receiving other kinds of benefits, such as food, mating opportunities or money. Future experiments might integrate CCE games such as ours with more traditional PGGs to explore interactions between different kinds of cooperative activities and goods.

In summary, our studies provide tentative evidence that people are over-sharers when it comes to information: they have strong preferences to share costly information with others across a range of contexts and conditions, even when exploitative information hoarding strategies give higher payoffs (study 3a), and even with no possibility of benefits being returned via reputations for sharing (study 4). It may be that this relatively indiscriminate information sharing is at least as important as high-fidelity social learning for enabling and accelerating human CCE. However, further research is needed to explore the limits of this indiscriminate information sharing, and whether other mechanisms of cooperation are better candidates for facilitating human CCE than partner choice appears to be.

Ethics. All studies were approved by the University of Exeter Faculty of Environment, Science and Economy Cornwall Ethics Committee (Application ID: 4523647).

Data accessibility. All data and code for all studies, as well as oTree code for the experimental task, is openly available at Zenodo [37].

Supplementary material is available online [38].

Declaration of AI use. AI image generators were used to create images for the materials used in the online experiments. We have not used AI-assisted technologies in any other way in creating this article.

Authors' contributions. Á.V.J.: conceptualization, data curation, formal analysis, investigation, methodology, resources, software, validation, writing—review and editing; L.C.: conceptualization, funding acquisition, methodology, supervision, writing—review and editing; K.J.: conceptualization, funding acquisition, methodology, writing—review and editing; H.J.L.: investigation, supervision; A.M.: conceptualization, data curation, formal analysis, funding acquisition, methodology, project administration, supervision, validation, visualization, writing—original draft.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest. We declare we have no competing interests.

Funding. This work was supported by the Cultural Evolution Society Transformation Fund, underwritten by the John Templeton Foundation, Grant 61913. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

Acknowledgements. We thank Oliver Hauser, Kenji Itao, Alex Thornton, Wataru Toyokawa, members of the University of Exeter's Behaviour Discussion Group and two anonymous reviewers for valuable feedback.

References

1. Boyd R, Richerson PJ. 1996 Why culture is common, but cultural evolution is rare. *Proc. Brit. Acad.* **88**, 77–93.
2. Henrich J. 2015 *The secret of our success*. Princeton, NJ: Princeton University Press.
3. Mesoudi A, Thornton A. 2018 What is cumulative cultural evolution? *Proc. R. Soc. B* **285**, 20180712. (doi:10.1098/rspb.2018.0712)
4. Lewis HM, Laland KN. 2012 Transmission fidelity is the key to the build-up of cumulative culture. *Phil. Trans. R. Soc. B* **367**, 2171–2180. (doi:10.1098/rstb.2012.0119)
5. Kempe M, Lycett SJ, Mesoudi A. 2014 From cultural traditions to cumulative culture. *J. Theor. Biol.* **359**, 29–36. (doi:10.1016/j.jtbi.2014.05.046)
6. Morgan TJH *et al*. 2015 Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nat. Commun.* **6**, 6029. (doi:10.1038/ncomms7029)
7. Mesoudi A, Jimenez AV, Jensen K, Chang L. 2025 From information free-riding to information sharing. *Phil. Trans. R. Soc. B* **380**, 20240258. (doi:10.1098/rstb.2024.0258)
8. Rogers AR. 1988 Does biology constrain culture? *Am. Anthropol.* **90**, 819–831. (doi:10.1525/aa.1988.90.4.02a00030)
9. McCray WP. 1999 Creating networks of skill. *J. Eur. Econ. Hist.* **28**, 301–333.
10. Ogilvie S. 2004 Guilds, efficiency, and social capital. *Econ. Hist. Rev.* **57**, 286–333. (doi:10.1111/j.1468-0289.2004.00279.x)
11. Ogilvie S. 2014 The economics of guilds. *J. Econ. Persp.* **28**, 169–192. (doi:10.1257/jep.28.4.169)
12. Mokyr J. 2016 *A culture of growth*. Princeton, NJ: Princeton University Press.
13. Open Science Collaboration. 2015 Estimating the reproducibility of psychological science. *Science* **349**, aac4716. (doi:10.1126/science.aac4716)
14. Roberts G *et al*. 2021 The benefits of being seen to help others. *Phil. Trans. R. Soc. B* **376**, 20200290. (doi:10.1098/rstb.2020.0290)
15. Leimar O, Hammerstein P. 2001 Evolution of cooperation through indirect reciprocity. *Proc. R. Soc. B* **268**, 745–753. (doi:10.1098/rspb.2000.1573)
16. Anthony D, Smith SW, Williamson T. 2009 Reputation and reliability in collective goods. *Ration. Soc.* **21**, 283–306. (doi:10.1177/1043463109336804)
17. Mesoudi A. 2008 An experimental simulation of the 'copy-successful-individuals' cultural learning strategy. *Evol. Hum. Behav.* **29**, 350–363. (doi:10.1016/j.evolhumbehav.2008.04.005)
18. Anlló H, Salamander G, Raihani N, Palminteri S, Hertz U. 2024 Experience and advice consequences shape information sharing strategies. *Commun. Psychol.* **2**, 1–13. (doi:10.1038/s44271-024-00175-5)
19. Fehr E, Gächter S. 2000 Cooperation and punishment in public goods experiments. *Am. Econ. Rev.* **90**, 980–994. (doi:10.1257/aer.90.4.980)
20. Chen DL, Schonger M, Wickens C. 2016 oTree—An open-source platform for laboratory, online, and field experiments. *J. Behav. Experim. Fin* **9**, 88–97. (doi:10.1016/j.jbef.2015.12.001)
21. Peer E, Rothschild D, Gordon A, Evernden Z, Damer E. 2022 Data quality of platforms and panels for online behavioral research. *Behav. Res. Methods* **54**, 1643–1662. (doi:10.3758/s13428-021-01694-3)
22. Sutton RS, Barto AG. 2018 *Reinforcement learning*. Cambridge, MA: MIT Press.
23. McElreath R, Bell AV, Efferson C, Lubell M, Richerson PJ, Waring T. 2008 Beyond existence and aiming outside the laboratory. *Phil. Trans. R. Soc. B* **363**, 3515–3528. (doi:10.1098/rstb.2008.0131)
24. Toyokawa W, Saito Y, Kameda T. 2017 Individual differences in learning behaviours in humans. *Evol. Hum. Behav.* **38**, 325–333. (doi:10.1016/j.evolhumbehav.2016.11.001)
25. Stan Development Team. 2024 Stan: Software for bayesian data analysis, version 2.36. See https://mc-stan.org.
26. Gabry J, Češnovar R, Johnson A, Bronder S. 2024 cmdstanr, version 0.8.1. See https://mc-stan.org/cmdstanr/.
27. R Core Team. 2024 R: A language and environment for statistical computing, version 4.4.2. See https://www.r-project.org.
28. Chang L, Mak M, Li T, Wu BP, Chen BB, Lu HJ. 2011 Cultural adaptations to environmental variability. *Edu. Psychol. Rev.* **23**, 99–129. (doi:10.1007/s10648-010-9149-0)
29. Mesoudi A, Chang L, Murray K, Lu HJ. 2015 Higher frequency of social learning in China than in the west shows cultural variation in the dynamics of cultural evolution. *Proc. R. Soc. B* **282**, 20142209. (doi:10.1098/rspb.2014.2209)
30. Romer PM. 1990 Endogenous technological change. *J. Polit. Econ.* **98**, S71–S102. (doi:10.1086/261725)
31. Burton-Chellew MN, El Mouden C, West SA. 2017 Social learning and the demise of costly cooperation in humans. *Proc. R. Soc. B* **284**, 20170067. (doi:10.1098/rspb.2017.0067)
32. van den Berg P, Molleman L, Weissing FJ. 2015 Focus on the success of others leads to selfish behavior. *Proc. Natl Acad. Sci. USA* **112**, 2912–2917. (doi:10.1073/pnas.1417203112)
33. Fehr E, Fischbacher U. 2003 The nature of human altruism. *Nature* **425**, 785–791. (doi:10.1038/nature02043)
34. Anderson C, Hildreth JAD, Howland L. 2015 Is the desire for status a fundamental human motive? *Psych. Bull.* **141**, 574–601. (doi:10.1037/a0038781)
35. Richerson PJ *et al*. 2016 Cultural group selection plays an essential role in explaining human cooperation. *Behav. Brain Sci.* **39**, e30. (doi:10.1017/S0140525X1400106X)
36. Derex M. 2021 Human cumulative culture and the exploitation of natural phenomena. *Phil. Trans. R. Soc. B* **377**, 20200311. (doi:10.1098/rstb.2020.0311)
37. Mesoudi A, Jiménez Á. 2026 Data and code for: Limited evidence that reputation-based partner choice facilitates information sharing in humans. Zenodo. (doi:10.5281/zenodo.19887197)
38. Jiménez ÁV, Chang L, Jensen K, Lu HJ, Mesoudi A. 2026 Supplementary material from: Limited evidence that reputation-based partner choice facilitates information sharing in humans. Figshare. (doi:10.6084/m9.figshare.c.8502444)