From information free-riding to information sharing: how have humans solved the cooperative dilemma at the heart of cumulative cultural evolution?

Alex Mesoudi[1*], Ángel Jiménez[1], Keith Jensen[2] & Lei Chang[3]

[1]Human Behaviour and Cultural Evolution Group, Centre for Ecology and Conservation, Penryn Campus, University of Exeter, UK

[2]Division of Human Communication, Development & Hearing, University of Manchester, UK

[3]Department of Psychology, University of Macau, Macau

[*]Corresponding author, a.mesoudi@exeter.ac.uk

10 June, 2024

Abstract

Cumulative cultural evolution, where populations accumulate ever-improving knowledge, technologies and social customs, is arguably a unique feature of human sociality and responsible for our species' ecological dominance of the planet. However, at the heart of cumulative cultural evolution is a cooperative dilemma. Assuming asocial learning is more costly than social learning, social learners can act as 'information free-riders' by copying innovations from asocial learners without paying the cost. This cost asymmetry will reduce innovation, inhibiting cumulative culture. Innovators might respond by protecting their knowledge and keeping the benefits to themselves – 'information hoarding' - but then others cannot build on their discoveries and again cumulative culture is inhibited. Here we formally model information free-riding and information hoarding within a cumulative cultural evolution framework using both analytical and agent-based models. Model 1 identifies the restrictive conditions under which information sharing can evolve in the face of information free-riding and hoarding. Models 2-4 then show how three mechanisms known to favour cooperation in non-informational contexts - kin selection, reputation-based partner choice and cultural group selection – can also solve the informational cooperative dilemma and facilitate cumulative cultural evolution, each with distinct signatures potentially detectable in historical, ethnographic and other empirical data.

Keywords: cooperation; cultural evolution; cumulative culture; innovation; social learning.

Introduction

Cumulative cultural evolution (CCE), where populations accumulate ever-improving knowledge, technologies and social customs, is arguably a unique feature of human sociality and responsible for our species' ecological dominance of the planet [1–3]. However, it is often overlooked that at the heart of CCE lies a cooperative dilemma. CCE requires both asocial learning (aka 'individual learning' or 'innovation') to create new knowledge, and social learning (aka 'cultural transmission' or 'cultural learning') to preserve and accumulate knowledge across generations [4]. Asocial learning is assumed to be more costly than social learning; it takes more time and effort to invent something new than copy it from someone else. This cost asymmetry can generate an informational tragedy of the commons [5]. If innovators freely share their knowledge, then social learners can become 'information free-riders' by copying beneficial knowledge from innovators without bearing the costs of innovation. All else being equal, social learners will therefore outperform and replace all innovators. However, if everyone copies, then no-one is innovating, and CCE stops. Innovators might respond by protecting their knowledge and keeping the benefits to themselves – 'information hoarding' – although then others cannot build on their discoveries, and again CCE stops.

While humans have clearly to some extent solved this cooperative dilemma given that we exhibit CCE, a glance through history and across societies illustrates how pervasive and challenging information free-riding and information hoarding have been. Amongst West Papuan hunter-gatherers, expert male adze-makers transmit their skills exclusively to their sons or nephews via apprenticeships [6]. In medieval Venice, expert glassmakers were legally prohibited from leaving the city to prevent their skills spreading to rival states [7]. In such cases, while specific actors benefit (adze-makers' kin; the Venetian Republic), overall CCE is inhibited due to the reduced pool of innovators [8,9].

Sometimes, however, widespread information-sharing emerges. Mokyr [10] attributes the exponential accumulation of knowledge during the 17th–18th century Enlightenment to a 'Republic of Letters', a network of innovators such as Francis Bacon and Isaac Newton who openly shared ideas, data and methods. Mokyr attributes this to a process of 'competitive patronage', where powerful families or governments protected and rewarded innovators in exchange for reputational benefits. The recent 'open science' movement also represents a shift from scientists hoarding information for their own benefit to the open release of data and methods, aiding the identification of replicable results that can be more reliably built upon by others [11]. Patent and copyright systems provide financial benefits to innovators to offset the costs of innovation [12], although there is little consensus on their effectiveness [13], and practices such as patent thicketing or patent hoarding can block innovation, illustrating the fragility of information sharing.

Our aim here is to formally model information free-riding and information hoarding in a CCE context, to (i) explore when and why these phenomena inhibit CCE, and (ii) whether solutions to free-riding from the evolution of cooperation literature [14], originally designed for material cooperative dilemmas, also apply to informational cooperative dilemmas. Formal models can serve to clarify verbal arguments and historical case studies such as those of Mokyr [10], highlighting often hidden underlying assumptions, and generating unexpected insights or predictions that are not apparent to the unaided mind [15].

While previous theoretical and empirical studies have touched on the cooperative dilemma at the heart of CCE, none have directly addressed it, nor fully examined potential solutions to it. Rogers [16] modelled the evolution of social learning as a producer-scrounger dilemma where innovators generate knowledge at higher cost than social learners can copy that knowledge. This generates a dilemma where, assuming constant environments, social learners entirely replace innovators because they bear lower costs and gain equal benefits. When environments change, a stable equilibrium exists between social learners and innovators due to the added advantage to innovators of discovering newly adaptive knowledge when previous knowledge becomes out-dated. In this model and its extensions [17–19], social learning is a form of information free-riding. However, these models are not directly relevant to the context described above. Rogers' model does not allow the accumulation of knowledge, only alternation between one of two behaviours. Environmental change renders all existing knowledge useless, again preventing CCE. Furthermore, innovators have no control over whether others can copy them, thus excluding the possibility of information hoarding, nor solutions to the cooperative dilemma.

One study that incorporated both CCE and information hoarding modelled the accumulation of technology in a producer-scrounger game, with a parameter controlling the excludability of the technology produced by innovators [20]. Generally, social learners facilitated CCE by acquiring multiple beneficial technologies from different innovators and passing them all to the next generation. Excludability acted against this: when innovators could prevent social learners copying them, fewer social learners persisted, inhibiting CCE. While this model demonstrates the negative consequence of information hoarding for CCE, it does not model cooperative solutions to information hoarding.

A more recent model allowed innovators to either teach, i.e. improve the chances of social learning, or mask, i.e. decrease the chances of social learning, equivalent to information hoarding [21]. This and other models of the evolution of teaching [22] (and language/communication more generally [23,24]) find that teaching evolves due to kin selection, as both genes for teaching and adaptive information are transmitted together from parent to offspring. Yet these models are designed to explore the evolution of teaching across different species rather than the emergence of information sharing within human CCE, and make assumptions (e.g. that teaching and masking are genetically transmitted) that are inconsistent with the diversity and sporadic emergence and loss of information sharing through human history.

The following models build up in complexity to first formally analyse the informational dilemma at the heart of CCE (Model 1) before exploring how kin selection (Model 2), partner choice (Model 3) and cultural group selection (Model 4) might solve this dilemma and facilitate CCE. Figure S1 presents a schematic of the four models.

Model 1: The dilemma

The aim of Model 1 is to formalise the informational public goods dilemma described above. This is intended as a clarification of the logic using many simplifying assumptions, not a realistic simulation of specific historical processes. Model 1a is an analytic population-level model. Model 1b introduces an agent-based version of Model 1a with explicit agents, traits and transmission events which recreates and confirms the findings of Model 1a and serves as a basis for Models 2-4.

Model 1a formalises two processes: information free-riding and information hoarding. Assume three types of individual. Open Innovators innovate at cost c to generate a cultural trait that yields a benefit b (see Table S1 for a full list of parameters). Open Innovators then freely and unconditionally share their trait with others. They receive a benefit x for sharing their trait, via a mechanism which for now remains unspecified. Later we replace x with explicit mechanisms of cooperation (kin selection, Model 2; partner choice, Model 3; cultural group selection, Model 4), but for now we leave it as a placeholder in order to clarify the general logic of informational cooperative dilemmas. Non-Innovators never innovate, are content to let others undergo costly innovation, and benefit from others' (open) innovation. They can be seen as information free-riders. Closed Innovators innovate at cost c with benefit b, and then bear an additional cost d to hoard their trait from others. With probability p_h this hoarding is successful, otherwise the innovation is released to the community just like those of Open Innovators. If their innovation is released, they also get the benefits of sharing x. Closed Innovators can be seen as information hoarders. Our guiding question is under what conditions do Open Innovators displace Closed Innovators and Non-Innovators.

Assume the number of Open, Closed and Non-Innovators in the population are X, Y and Z respectively. The payoff to an Open Innovator, $W(OI)$, is then:

$$W(OI) = w_0 + bX + bY(1 - p_h) + x - c \tag{1}$$

In addition to baseline payoff w_0, each Open Innovator receives a payoff b from each of the innovations generated by the X Open Innovators in the population (including themselves), a payoff b from the $Y(1 - p_h)$ Closed Innovators who unsuccessfully protect their innovation such that it becomes openly known, a benefit x for releasing their innovation, and they bear a cost of innovation c.

The payoff to a Closed Innovator, $W(CI)$, is:

$$W(CI) = w_0 + bX + bY(1 - p_h) + bp_h + x(1 - p_h) - c - d \tag{2}$$

Each Closed Innovator receives the same benefits as Open Innovators from the open knowledge generated by Open Innovators and unsuccessful Closed Innovators, $bX + bY(1 - p_h)$. Closed Innovators also receive, with probability p_h, another benefit b from their successfully protected innovation, and with probability $1 - p_h$ a benefit x from sharing their unsuccessfully protected innovation. Finally, each Closed Innovator bears a cost of innovation c and a cost of attempting to protect their knowledge d.

The payoff to a Non-Innovator, $W(NI)$, is:

$$W(NI) = w_0 + bX + bY(1 - p_h) \tag{3}$$

Non-Innovators pay no costs of innovating, and receive the benefits of open innovation from Open Innovators and unsuccessful Closed Innovators

Note that in this formulation, innovation is a perfectly *non-rival* good via its benefit b. Every Open Innovator and unsuccessful Closed Innovator generates a payoff b which is received by (not divided amongst) every individual in the population. This captures the notion that knowledge, unlike material goods, can be shared via cultural transmission without loss; one person can transmit knowledge to another without themselves losing that knowledge [25]. Innovation is *excludable* for Closed Innovators via parameter p_h: with probability p_h, Closed Innovators generate private knowledge, and with probability $1 - p_h$ it is non-excludable. For Open Innovators, innovation is always non-excludable.

Open Innovators can invade Non-Innovators when $x > c$. This is when the benefits to Open Innovators of releasing their innovation outweigh the cost of producing that innovation. Note that b is not present in this inequality. It does not matter how big the benefit is of the innovation, because both Open Innovators and Non-Innovators receive that benefit.

Closed Innovators can invade Non-Innovators when $bp_h + x(1-p_h) > c+d$. This is when the benefit to Closed Innovators of private innovation, bp_h, plus the reputational benefit of unsuccessfully protected innovation, $x(1-p_h)$, outweigh the costs to Closed Innovators of innovation and of protecting their innovation. Here the benefit of innovation, b, does matter, because when Closed Innovators are successful then only they receive their private benefit.

Open Innovators can invade Closed Innovators when $d > p_h(b-x)$. This is fulfilled when d is large (imposing high costs on Closed Innovators), when p_h is small (low probability of Closed Innovators generating private knowledge not available to Open Innovators) and when x is large relative to b (because Open Innovators benefit uniquely from x and not b). When $x > b$ then, assuming that $d > 0$, Open Innovators will always outperform Closed Innovators. Note that c is not present in this inequality, as both Open Innovators and Closed Innovators pay innovation costs.

Consider how to get to a population of Open Innovators from a mix of Non-Innovators and Closed Innovators. Consider first a situation when $p_h = 1$, i.e. Closed Innovators are always successful in protecting their knowledge. This is perhaps due to poor communication and limited social networks. This means that Open Innovators will do better than Non-Innovators when $x > c$ and better than Closed Innovators when $x > b - d$. Only when both of these conditions are met will Open Innovators spread. Collectively, then, for Open Innovators to spread, x and/or d need to increase relative to c and/or b (figure 1).

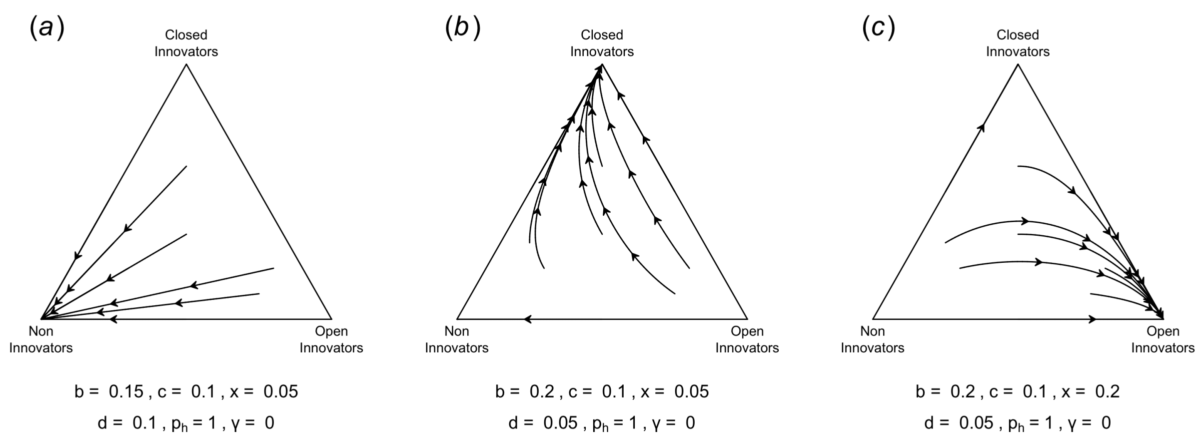

Figure 1. Three conditions in which Non Innovators, Closed Innovators and Open Innovators are respectively favoured. (a) $x < c$ and $b < c + d$ so Non-Innovators spread at the expense of Open Innovators and Closed Innovators. (b) b is increased and d is decreased to now favor Closed Innovators. (c) x is increased such that $x > c$ and $x > b - d$, favoring Open Innovators. Plots created using a modified version of package *baryplot* [26].

Note the apparent paradox here related to b. Open Innovators are more likely to emerge if the benefit from innovation b is small. Yet CCE by definition surely *increases* b in absolute terms, as technology becomes more effective and knowledge more accurate. So the consequence of open innovation - larger benefits of innovation - paradoxically make it harder for open innovation to emerge.

Consider now the case when $p_h < 1$, i.e. when Closed Innovators sometimes fail to protect their knowledge. This might be due to inventions or institutions such as the printing press, postal services, cheaper transportation or the internet making it harder to keep innovations secret. Reducing p_h does not affect whether Open Innovators outcompete Non-Innovators, which remains when $x > c$. Benefits of sharing still need to outweigh the costs of innovation. However, it does reduce the parameter space within which Closed Innovators can outcompete Open Innovators, assuming $x > c$. At the extreme when $p_h = 0$, Open Innovators will replace Closed Innovators whenever $d > 0$ (figure S2). This confirms the notion that making it harder to protect or conceal knowledge favours open innovation.

Finally, we implement a crude form of CCE by assuming that the benefits to innovation, b, increase with accumulating knowledge. A steel axe is more effective and durable than a stone axe, but only appeared after the accumulation of prior axe manufacturing and steelworking knowledge. Quantum physics is more accurate than Newtonian physics, but only appeared after the latter was established.

Accumulated knowledge must be openly available in order to accumulate, otherwise it dies with the innovator. We therefore assume that every timestep the parameter b increases by an amount $\gamma(bX + bY(1-p_h))$. Here γ is a constant controlling the extent to which b increases, and the term in brackets is the benefit accrued from open knowledge generated by Open Innovators and unsuccessful Closed Innovators. When $\gamma = 0$ we retrieve the results shown in figure 1; as γ increases, so does b in each timestep.

Figure 2 shows the same parameter values as figure 1c where Open Innovators were previously favoured, but with $\gamma = 1$. Now Closed Innovators are favoured. This is because increasing b favours Closed Innovators by increasing their benefit from private innovation (bp_h). The increased benefits generated by Open Innovators and unsuccessful Closed Innovators is shared by all agents and therefore does not affect the dynamics.

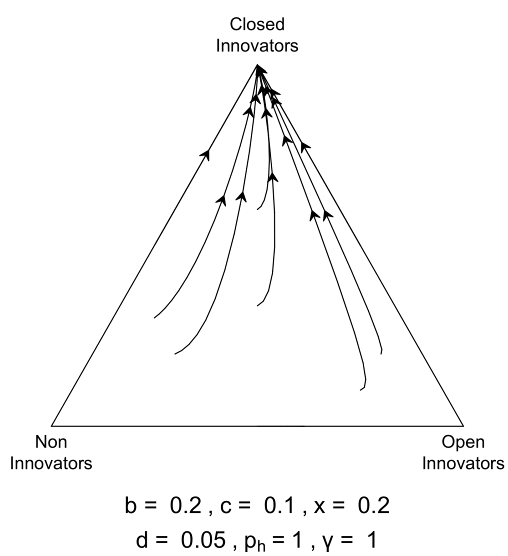

Figure 2. Allowing benefits from innovation to accumulate by setting $\gamma = 1$ favours Closed Innovators where Open Innovators would otherwise be favoured (as in figure 1c).

Model 1b replicates and extends Model 1a to incorporate explicit agents, interactions and traits. This added complexity and need to explicitly model social interactions favours an agent-based modelling approach [27,28]. Model 1b assumes N agents each of whom can learn any number of L traits. Traits are represented as 'bit strings', sequences of L ones and zeroes where a 1 in position l indicates that the lth trait has been learned and a 0 indicates a lack of knowledge of the trait. For example, 00110 is a set of $L = 5$ traits of which this agent only knows the third and fourth. Each learned trait gives the agent a benefit b each timestep; this agent would receive $2b$ from their two learned traits.

As in Model 1a, agents are Open Innovators, Closed Innovators or Non-Innovators. In each of t_{max} timesteps, each Open Innovator and Closed Innovator engages in innovation by picking one of its L traits at random and, if that trait is 0, switching it to 1 with probability p_i. Innovation costs c whether successful or not. Following innovation, there is sharing/copying. As before, Closed Innovators successfully hoard their traits with probability p_h. Each agent copies each trait known by at least one non-hoarding agent (all Open Innovators, all Non-Innovators, and unsuccessful Closed Innovators) with probability p_c per trait. When $p_c = 1$, each agent acquires all the traits known by all non-hoarding agents in the population. This extreme case matches Model 1a, with $p_c < 1$ the more realistic case when only some innovations are acquired due to time constraints or copying error. As before, Open Innovators and unsuccessfully-hoarding Closed Innovators who innovated on that timestep receive a benefit x for openly releasing their innovation. Closed Innovators pay a hoarding cost d whether hoarding is successful or unsuccessful. However, Closed Innovators who fail to innovate (because $p_i < 1$) do not pay the cost d as they have nothing to hoard.

After fitnesses are calculated, there is payoff-biased copying of strategies (Open Innovation, Closed Innovation or Non-Innovation). Each agent picks another agent in the population at random. If the chosen agent has higher fitness than the focal agent, then the focal agent adopts the strategy of the chosen agent with probability p_s. Then there is agent turnover. With probability p_d, each agent 'dies' and is replaced with a naive, unknowledgeable agent whose traits are all zero. The new agent keeps the same strategy as its 'parent'. The parameter p_d therefore controls how much timesteps overlap. When $p_d = 1$, all agents die each timestep and are replaced with unknowledgeable agents, recreating Model 1a. When $p_d = 0$, agents live forever, or at least until t_{max} timesteps. This resembles a fixed population of agents engaging in repeated cycles of innovation and copying, allowing the accumulation of learned traits over time.

Finally, it is unrealistic to assume that if a population learns all L possible traits then cultural evolution stops. Consequently we assume that when the mean proportion of traits known across all agents in the population reaches 0.9, then a new set of L unknown traits are added to the set of possible traits. This can be seen as the opening up of a new space of possibilities once knowledge in one domain has reached a certain point. For example, automobile tyre designs increased rapidly in diversity and effectiveness following the invention of the automobile in the late 19th century before stabilising and converging on a single design containing a specific set of innovations [29]. This generates patterns of punctuated equilibria with rapid increases in cultural knowledge followed by stasis, then rapid increases again [30,31]. This makes Model 1b resemble genuine sequential cumulative cultural evolution, rather than the crude implementation in Model 1a.

Model 1b replicated the findings of Model 1a with equivalent parameter values (figures S3-S5), supporting the conclusions drawn by the analytical model. We additionally find that reducing the probability of innovation p_i favours Non-Innovators given that Open and Closed Innovators pay a cost for innovation that increasingly yields no benefits (figure S6). We also find that when the information dilemma is eliminated by making innovation easier than social learning ($p_c < p_i$), then

Open Innovators do best, as expected (figure S6). Finally, Model 1b shows a clear link between scenarios in which parameter values favour open innovation and CCE (figure S7): more open innovation means more traits are accumulated, and when combined with overlapping timesteps ($p_d < 1$) generates open-ended CCE.

Summary of Model 1

The aim of Model 1 was to formalise the informational dilemmas at the heart of CCE. Using both analytical and agent-based models, we showed two ways in which Open Innovators who freely share information can be exploited, resulting in the disruption of CCE. First, information free-riding occurs when Non-Innovators can copy cultural traits from Open Innovators without paying the cost of innovation, and when any benefits accrued to Open Innovators for information sharing fail to compensate for this (i.e. when $x < c$). The actual benefit gained from the innovation, b, does not matter here, given that both Open Innovators and Non-Innovators receive it. Second, and perhaps less commonly appreciated, information hoarding occurs when Closed Innovators can generate and hoard their own private knowledge, plus receive public knowledge from Open Innovators. Open Innovators only have public knowledge. Unless the benefits of information sharing to Open Innovators outweigh the benefits to Closed Innovators of private knowledge (i.e. when $x > b$, assuming non-overlapping timesteps), then Closed Innovators will displace Open Innovators. Interestingly, when cultural traits can accumulate, in Model 1a by increasing the benefit b in proportion to openly released knowledge, and in Model 1b by making timesteps overlap such that traits are inherited across generations, then information sharing is even less likely. This is because private knowledge hoarded by Closed Innovators is increasingly beneficial, and increasingly likely to outweigh the benefits of information sharing. Finally, our agent-based Model 1b demonstrated a clear link between the presence of information sharing and CCE: when Open Innovators dominate, then cultural traits accumulate, and when combined with overlapping timesteps generates open-ended CCE. Overall, crucial to the emergence of CCE is the currently-unspecified benefits of openly releasing information, x. In Models 2-4 we replace this parameter with three mechanisms from the evolution of cooperation literature which endogenously direct benefits back to Open Innovators: kin selection (Model 2), partner choice (Model 3) and cultural group selection (Model 4).

Model 2: Kin selection

A widespread solution to cooperative dilemmas in nature is kin selection [32]. Individuals direct helping behaviour towards kin who share genes with the helper due to common ancestry. On average, relatives will share the genes underlying kin-directed cooperation, and so such genes will increase in frequency as a result of the cooperation. Kin selection explains the vast majority of cooperative behaviour in non-human species [33], as well as various forms of human cooperation [34–36].

Here we do not model the evolution of cooperation via kin selection, which has been extensively modelled in evolutionary biology [32,37]. We assume kin selection has already evolved in our species, and provides motivation to preferentially help genetic relatives. In the context of information sharing, individuals should preferentially share information with relatives. We focus on parent-to-offspring transmission (often labelled ‘vertical cultural transmission’: [38]), given that parents will typically have greater knowledge than, and spend considerable time with, their children. This resembles the case of father-son adze making apprenticeships amongst West Papuans [6]. (Note however that there is extensive evidence of learning from non-parents across societies [39–43], both obliquely from unrelated elders and horizontally from unrelated peers.) However, while kin selection may be one solution to the informational dilemma, it has downsides. Learning from just two parents - or one, in cases of sex-specific skills and knowledge - provides a far smaller pool of demonstrators and teachers than learning from any member of society. Fewer demonstrators means slower CCE [8,9].

The aim of Model 2 is to formalise and test the notion that kin-directed information sharing can avoid exploitation by information free-riding and information hoarding, as well as its limitations in terms of fewer demonstrators. Model 2 is identical to Model 1b except for the social learning stage. We no longer assume social learning from all non-hoarders, i.e. horizontal cultural transmission, nor Open Innovators who indiscriminately share with all other agents. Instead, social learning takes the form of vertical transmission and occurs at the start of each timestep (except the first, when there are no parents from whom to learn). Open Innovators are replaced with Kin-directed Innovators who share traits exclusively with their offspring. Each agent produces one offspring, i.e. asexual reproduction. While humans are obviously not asexual reproducers, this might resemble sex-specific transmission of skills, as well as providing a simple test case. Parents transmit their strategies (Kin-directed Innovator, Closed Innovator or Non-Innovator) to their one offspring. This could be via either genetic or cultural inheritance. Non-hoarding agents (Kin-directed Innovators, Non-Innovators and unsuccessful Closed Innovators) also transmit each of their known cultural traits to that offspring with probability p_c per trait. There is then innovation amongst Kin-directed Innovators and Closed Innovators, fitness calculations, and payoff-biased copying of strategies, all as in Model 1b. There is now no agent mortality (i.e. p_d is absent) because all agents are assumed to die each timestep and are replaced with their offspring. Unlike Model 1, there are now no direct benefits of openness ($x = 0$) given that kin selection is explicitly modelled as a mechanism providing those benefits. Initially we assume copying and innovation are both errorless ($p_c = p_i = 1$) and hoarding always successful ($p_h = 1$).

Figure 3 shows that under parameter values that would not have favoured Open Innovators in Model 1 (figure 3a), Kin-directed Innovators go to fixation (figure 3b). Kin selection favours the emergence of agents who preferentially share traits with offspring. This is because both strategies and cultural traits are inherited vertically together. Only Kin-directed Innovator parents will innovate and pass on those beneficial innovations to their offspring. Closed Innovators innovate but do not pass on that knowledge to their offspring, while Non-Innovators do not innovate so have nothing to pass on. This is the same mechanism that explains the evolution of teaching in other models [21,22], where both genes for teaching and the superior taught knowledge that results from teaching are passed on together from parents to offspring.

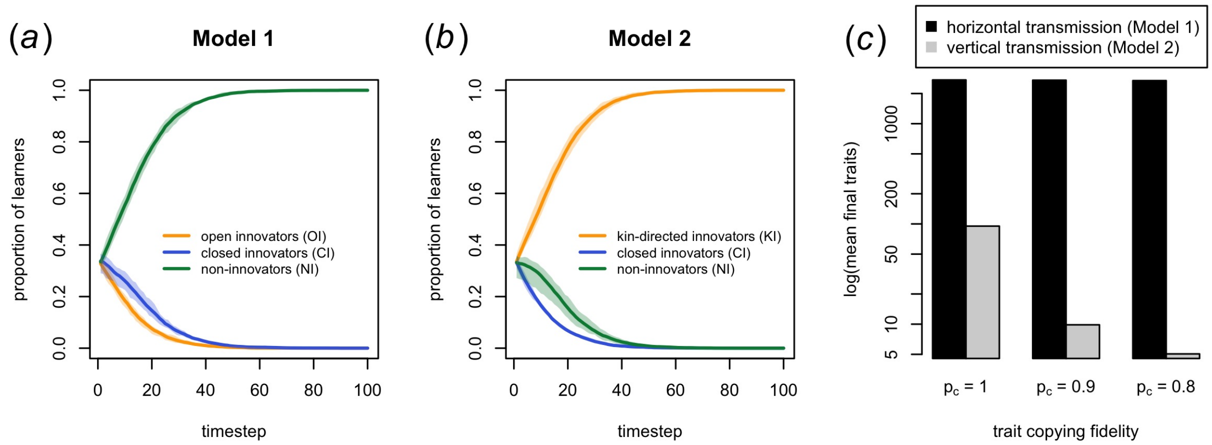

Figure 3. (a) Time series for the agent-based version of Model 1 with no benefits to Open Innovators for sharing knowledge ($x = 0$), favouring Non-Innovators (other parameter values: $b = 0.15$, $c = 0.1$, $d = 0.1$, $p_h = p_i = p_c = p_d = 1$; thick lines are means of 10 independent runs with shaded areas showing the range). (b) Time series with the same parameter values for Model 2, now favouring Kin-directed Innovators who share traits exclusively with kin (offspring). (c) Logged mean final number of traits accumulated at $t_{max} = 100$ for Model 1b and Model 2 in populations entirely composed of Open Innovators or Kin-directed Innovators respectively, showing that horizontal transmission in Model 1 supports orders of magnitude more traits than vertical transmission in Model 2, and is far less vulnerable to reduced trait copying fidelity ($p_c < 1$).

Figure 3c illustrates the downside of Kin-directed Innovation. Compared to the horizontal cultural transmission of Model 1b with 100% Open Innovators, Model 2 with 100% Kin-directed Innovators accumulates far fewer cultural traits. Learning from only a single parent creates a bottleneck that inhibits CCE compared to horizontal cultural transmission from all N individuals. Another notable difference between Models 1 and 2 is the effect of reducing the probability of copying, p_c. Figure 3c shows that reducing the probability of copying by a small amount severely inhibits CCE in Model 2, while having a negligible effect in Model 1. Vertical cultural transmission is far more vulnerable to copying error given that there is only a single demonstrator, compared to horizontal cultural transmission from the entire population. This effect of demonstrator number on CCE is well known from other models [8,9]. This vulnerability to copying error might explain why real world cases of kin-directed information sharing involve lengthy apprenticeships: the West Papuan adze making example from [6] involved apprenticeships of five years or more. Such lengthy apprenticeships may be one way to reduce copying error and maximise the probability of successfully copying cultural traits.

Figure S8 shows two alternative versions of Model 2. Model 2b assumes sexual reproduction where agents mate and produce offspring who inherit the strategy and traits of one randomly chosen parent. This yields almost identical results as the main Model 2 with Kin-directed Innovators favoured (figure S8a), again because traits and strategies are inherited together. Model 2c assumes that offspring combine the traits of both parents, while inheriting their strategy from just one. This reduces the success of Kin-directed Innovators, who now co-exist with Non-Innovators (figure S8b). This is because the inheritance link between strategy and cultural traits has been partially broken. If a Kin-directed Innovator and a Non-Innovator mate and produce an Non-Innovator offspring, this offspring will inherit the cultural traits of its Kin-directed Innovator parent (and potentially also its Non-Innovator parent), without paying the cost of innovation that Kin-directed Innovator offspring do. Figure S8*c* further confirms the trade-off between selective information sharing and the speed of CCE: when Kin-directed Innovators are less favoured in Model 2*c*, traits accumulate slower, especially when copying fidelity is less than perfect.

Summary of Model 2

Model 2 shows that kin selection is a viable solution to information free-riding and information hoarding, on the assumption that parents exclusively share cultural traits with their kin, and learning strategies are also passed on to those kin. Because both strategies and cultural traits are inherited together, this favours Kin-directed Innovator lineages due to their accumulating beneficial knowledge. This mechanism is similar to that present in models of the evolution of teaching [21,22], which is similarly argued to have evolved due to kin selection. However, Kin-directed Innovation suffers from having fewer demonstrators, which severely slows CCE relative to horizontal cultural transmission in Model 1, especially in the uniparental case where cultural traits are transmitted from only one parent. Biparental transmission (Model 2c) increases the rate of CCE, but at the expense of allowing Non-Innovators to exploit Kin-directed Innovators. In all cases where Kin-directed Innovators are favoured, reducing the fidelity of copying even slightly severely slows CCE. These patterns pleasingly match the empirical example of West Papuan adze making [6], which is strictly uniparental along male lineages and involves lengthy apprenticeships that potentially serve to maximise transmission fidelity.

Model 3: Reputation-based partner choice

Partner choice, aka competitive altruism [44,45], involves individuals selecting interaction partners based on the partners' past history of cooperation (their 'reputation'). If individuals who cooperate are subsequently more likely to be selected as recipients of cooperation, then cooperators are paired with cooperators and free-riders are excluded. This mechanism can favour cooperation even amongst non-kin as individuals compete to be the most cooperative partner and thus benefit by being chosen to receive help.

In Model 3 we implement partner choice as a potential mechanism for maintaining open innovation and facilitating CCE. Model 3 is identical to Model 1b (indiscriminate horizontal cultural transmission from all N agents) except that we relabel the previously indiscriminate Open Innovators as Reputational Innovators, and modify the social learning phase. In Model 3, each agent now selects n_c candidates to be learners at random from the $N-1$ other agents in the population. When $n_c = 1$ then partner choice is random, and should not lead to increased information sharing. When $n_c > 1$, the candidate with the highest reputation is chosen to be the learner for that agent, who acts as the demonstrator in their interaction. The larger n_c, the stronger the partner choice. If $n_c = N-1$, then the agent with the highest reputation in the entire population is guaranteed to be picked.

Once demonstrators and learners have been paired up, if the demonstrator is not successfully hoarding their knowledge, the chosen learner copies each known trait of the demonstrator with probability p_c per trait. Demonstrators receive a reputation increase of r_s per trait shared. In theory, this should lead to sharing demonstrators acquiring higher reputations than non-sharing demonstrators, and so in subsequent timesteps information sharers are more likely to be chosen as learners, receive more knowledge, and have higher payoffs due to their greater knowledge. In our model, information sharers are Reputational Innovators and Non-Innovators, and sometimes Closed Innovators when $p_h < 1$.

In addition, if the demonstrator was the first agent to innovate and share that particular trait - i.e. they were its inventor - then the demonstrator gets an additional reputation increase of r_i per invented trait. This is intended to represent the 'priority advantage' commonly seen for scientific or technological discovery, where originators of ideas (e.g. Darwin, Newton or Edison) receive higher reputation increases for sharing what they have originated than people who share others' inventions or discoveries. The frequent fights over priority (e.g. Darwin vs Wallace; Newton vs Leibniz; Edison vs Tesla) indicate the potential importance of priority advantage [46]. We added this assumption because this should lead to Reputational Innovators receiving higher reputations than Non-Innovators, given that only Reputational Innovators can originate (and then share) ideas.

Finally, we model two different ways in which reputations are acquired, local and global. For global partner choice, every agent has a single reputation known by every other agent in the population which is updated whenever they share knowledge with any other agent. For example, if agent i shares knowledge with agent j, then agent i's reputation increases in the eyes of all N agents in the population. This resembles reputation-based models of indirect reciprocity [47,48] where reputations (e.g. images scores or standings) are globally known, and might represent a modern situation where mass communication or the internet allows the rapid spread of reputational information. For local partner choice, every agent has a specific and potentially different reputation score for each other agent in the population, which is only updated when knowledge is shared with that specific agent. For example, if agent i shares knowledge with agent j, then only agent j increases their reputation score associated with agent i. This resembles previous agent-based models of partner choice [49], and might represent a fragmented community with infrequent, face-to-face communication.

Figure 4 shows the likelihood of Reputational Innovators going to fixation as a function of the number of candidate learners (n_c), whether reputations increase when any trait is shared ($r_s = 1$) or only traits that that agent invented are shared ($r_i = 1$), and for local vs global reputations, under parameter values that in Model 1 would have favoured Non-Innovators. As expected, when partner choice is random ($n_c = 1$), Reputational Innovators are not favoured under any condition. As n_c increases and partner choice gets stronger, Reputational Innovators are increasingly favoured. At the maximum ($n_c = N - 1$), Reputational Innovators are virtually guaranteed to go to fixation when reputations are global, and likely (but not guaranteed) to go to fixation when reputations are local.

As expected, increasing reputations only when agents share traits that they themselves invented ($r_i = 1$) is more effective at promoting information sharing than increasing reputations due to sharing any trait ($r_s = 1$) (figure 4). However, even the latter still favours Reputational Innovators at high values of n_c especially when reputations are global, even though Non-Innovators can receive reputation increases for sharing Reputational Innovators' invented traits without bearing the costs of innovation. This is because while Non-Innovators initially increase in frequency due to this advantage, few new traits are then being innovated, and rare Reputational Innovators gain an advantage by sharing traits that only they innovate (figure S9).

Figure 4 also shows that it takes much higher values of n_c, i.e. much stronger partner choice, for Reputational Innovators to be favoured under local partner choice than for global partner choice. This is as expected, given that localised, agent-specific knowledge of reputations acquired through direct interactions is inevitably less reliable than universally-known reputations acquired indirectly via every interaction in the population.

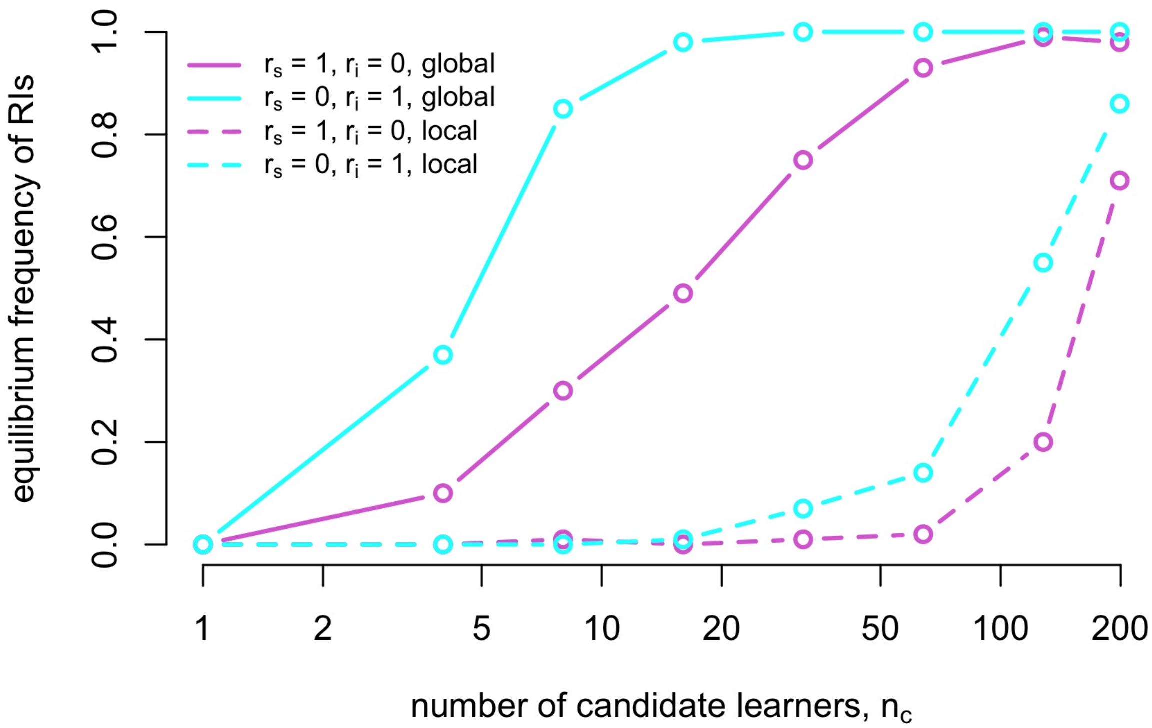

Figure 4. The proportion of 100 simulation runs in which Reputational Innovators go to fixation, as a function of the number of candidate learners (n_c), for the case when sharing any trait yields reputation rewards ($r_s = 1$; purple lines) or when only sharing traits that that agent invented yields reputation rewards ($r_i = 1$), and for global (solid lines) or local (dotted lines) reputations. Other parameters are set to values that ordinarily favour Non-Innovators: $N = 200$, $L = 1000$, $p_c = 0.8$, $p_i = 0.1$, $p_h = 0.9$, $p_s = 1$, $b = 0.2$, $c = 0.1$, $d = 0.05$, $p_d = 0.001$.

Summary of Model 3

Model 3 shows that reputation-based partner choice is a viable solution to information free-riding and information hoarding. If sharing information with others leads to higher reputations, and individuals with higher reputations for sharing are themselves subsequently more likely to be chosen to receive information from others, then those individuals will benefit from the received information and sharers will do better than non-sharers. The effectiveness of partner choice depends on being able to select potential partners (learners) from a large proportion of the population, i.e. a high n_c. In the real world, n_c might be seen as a measure of interconnectedness or interaction probability between community members that allows learning from others. More interconnected communities are therefore more likely to favour information sharing via partner choice. This supports Mokyr's [10] suggestion that letter writing, postal services, the printing press and international travel led to the formation of the Republic of Letters in which open innovation bloomed during the Enlightenment. Partner choice was also more effective when reputations accrue for sharing traits that the agent themself invented ($r_i = 1$) than for sharing any trait irrespective of provenance ($r_s = 1$). Therefore, we might expect that cultural systems that reward inventors of traits would be more likely to favour and sustain open innovation than systems that reward the sharing of any traits. Interestingly, however, Reputational Innovators were favoured when partner choice was sufficiently strong even when reputations increase for sharing any trait. This unanticipated robustness suggests that partner choice is particularly likely to lead to the spread of open innovation. Finally, partner choice is more effective with global reputations that are known by all agents and updated when any agent shares information with any other agent, compared to local reputations which are agent-specific and only updated when that agent has information shared with them. This again speaks to the importance of social interconnectedness and mass communication for partner choice to work, which would allow the widespread communication of reputations across the entire population.

Model 4: Cultural group selection

Another hypothesis for human cooperation amongst non-kin is cultural group selection [50,51]. Generally, group selection is the idea that groups of cooperators out-compete groups of non-cooperators. Genetic group selection is an unlikely explanation for cooperation in nature, given that migration rapidly breaks down the between-group genetic variation that is required for selection to operate at the level of the group, and group-wide cooperation will be undermined by selection within groups for selfish free-riders. However, *cultural* group selection rests on the assumption that cultural evolution generates better conditions for group-level selection. Cultural processes like conformity [52], punishment [53] or reciprocity [54] can maintain group-wide norms of cooperation and thus between-group variation in cooperation that is then subject to selection via direct (e.g. warfare) or indirect (e.g. economic) inter-group competition. In the context of information sharing and CCE, we might imagine groups of open innovators who freely share information exclusively within their group to accumulate more beneficial knowledge, and consequently out-compete, both groups of Non-Innovators who have nothing to share and accumulate, or groups of Closed Innovators who never share information and thus fail to accumulate information.

Here we model this scenario, adapting a previous agent-based simulation of cultural group selection [53]. We assume g groups each containing n agents, giving $ng = N$ agents in total. Agents can be Non-Innovators, Closed Innovators, or Group Innovators. The first two are identical to previous models. The latter are innovators who only share traits with other members of their own group. To test the above logic, and assuming that Non-Innovators are the 'ancestral' state, we start with one group of Group Innovators, one group of Closed Innovators, and the rest (i.e. the majority) Non-Innovators.

Each timestep there is first innovation for Closed Innovators and Group Innovators with probability p_i as before, and then social learning where each agent acquires each trait known by every non-hoarding agent in their group with probability p_c per trait. Then, following [53], there is costly punishment, an empirically supported within-group mechanism for maintaining cooperation [55,56]. Each Group Innovator reduces each Non-Innovator and hoarding Closed Innovator's payoff by u/N, and bears a cost of k/N per punished agent. Then there is payoff-biased copying of strategies, as before, followed by migration. With probability p_m per agent, each agent moves to a randomly chosen group, taking their traits and strategy with them, as per [57]. Finally there is group selection. Groups are paired up at random, and with probability p_g enter into a contest. The group with more traits wins the contest. The losing group's strategies and traits are replaced with those of the winning group. Note that unlike previous models where intergroup conflict success was determined by frequencies of agent types (e.g. groups with more defectors more likely to lose: [53]), here we make the more plausible assumption that group success is determined by number of cultural traits; hence group selection does not act directly on the Group Innovator phenotype, it acts only via the Group Innovators' ability to generate and accumulate cultural knowledge.

Figure 5 shows how Group Innovators are favoured as the probability of intergroup competition, p_g, increases, and the fitness cost of punishment imposed by Group Innovators on the other two types within groups, u, also increases, given a relatively low migration rate of $p_m = 0.01$. As in previous models [53], cultural group selection alone is not sufficient to spread group-beneficial behaviour in the face of migration. Punishment within groups is also required. Figure S10 shows how intergroup conflict is vulnerable to migration, as expected. Figure S11 shows the effect of population structure, with information sharing (like cooperation in general) less likely to emerge in larger populations, either via the number (g) or size (n) of groups.

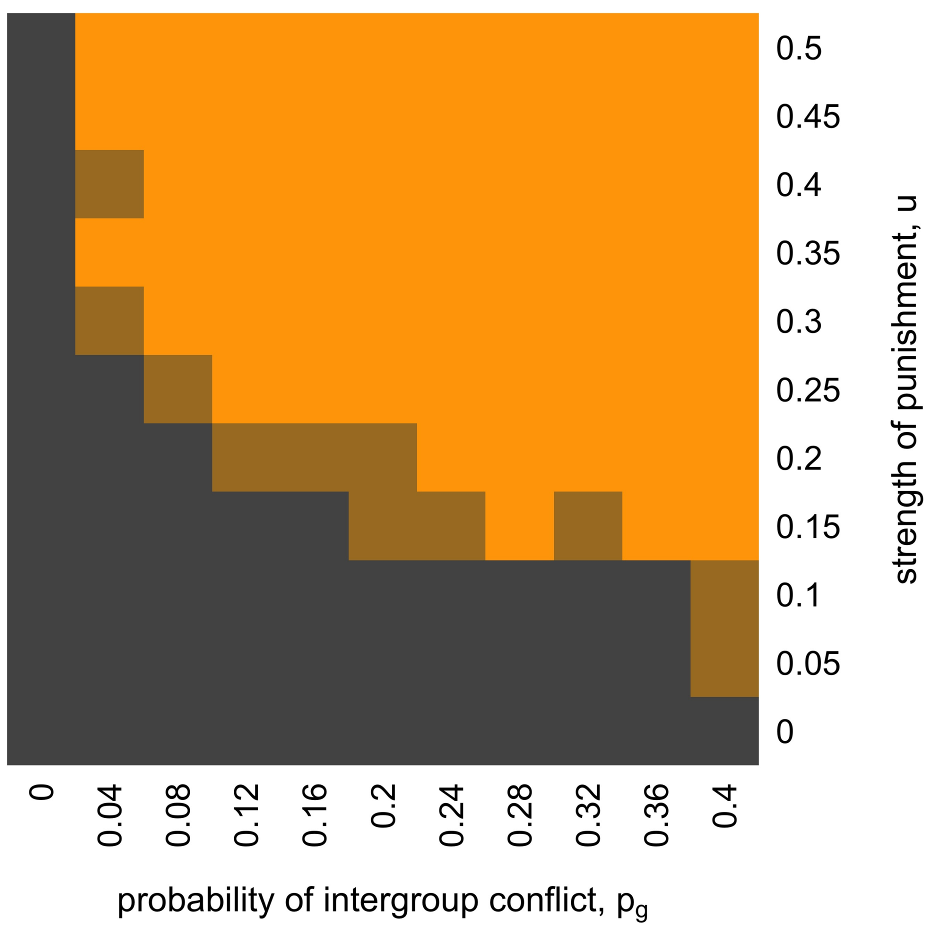

Figure 5. The frequency at equilibrium of Group Innovators in Model 4 (cultural group selection), for different values of the probability of intergroup conflict p_g, and the within-group punishment penalty u of agents with other strategies (Non-Innovators and Closed Innovators). Colours indicate

the mean frequency at equilibrium of Group Innovator agents across 20 independent simulation runs; orange = frequency greater than 2/3, black = frequency < 1/3, grey = intermediate frequencies. Other parameter values: $n = 50$, $g = 20$, $L = 1000$, $p_c = 0.8$, $p_i = 0.2$, $p_h = 0.9$, $p_s = 0.2$, $p_d = 0.1$, $b = 0.2$, $c = 0.1$, $d = 0.05$, $p_m = 0.01$, $k = 0.2$.

Summary of Model 4

Model 4 confirms that cultural group selection is a potential mechanism for the spread of group-based open innovation, where innovators share knowledge exclusively with other members of their groups. However, Model 4 also replicated previous model findings that the conditions for cultural group selection to act alone are rather restrictive: the rate of intergroup conflict must be high, migration rates low, and the total population size small. Also as in previous models [53], adding punishment as a within-group mechanism for maintaining cooperation boosts the power of cultural group selection such that it can favour open innovation at lower rates of intergroup conflict, higher rates of migration and larger population sizes. Historically, we might seek signatures of cultural group selection alongside within-group punishment, or indeed other within-group mechanisms such as reciprocity [54] or partner choice (Model 3). Such a scenario might resemble the example of glassmaking in the medieval Venetian Republic. Here the intergroup conflict was economic rather than military, with the Republic seeking to maintain its technological advantage over other states. There is also evidence of punishment, albeit punishment of glassmakers who sought to benefit individually by taking their knowledge to other states.

Discussion

Here we attempted to formally explore the factors that affect the cooperative dilemma at the heart of cumulative cultural evolution, as well as some solutions to the dilemma borrowed from evolutionary biology. Model 1 formalised the essential cooperative dilemma of CCE: if innovation (i.e. asocial learning) is more costly than copying (i.e. social learning), then it is optimal to be the latter, an informational free-rider acquiring costly information from others at no cost. Yet if everyone is copying, no-one is innovating, and CCE halts. Alternatively, innovators might try to protect their knowledge from being copied ('information hoarding'), but this again prevents the spread and accumulation of information, and CCE stops.

The notion that social learning can be seen as information free-riding has been modelled previously in the cultural evolution literature [16], yet not within the context of CCE where beneficial traits can accumulate over time. In fact, Model 1 found that making culture cumulative makes the informational cooperative dilemma even more pronounced. As culture accumulates, its benefits increase. As culture accumulates, therefore, information hoarders (called Closed Innovators in our models) can increasingly benefit from their private knowledge, out-competing Open Innovators who freely release their knowledge to others and have no private knowledge. Model 1 also shows that information hoarding is more likely to emerge when the cost of hoarding is low or non-existent, and when hoarders have a high probability of successfully hoarding their knowledge. Information free-riding (Non-Innovators in our models) is more likely to emerge when innovation is costly or difficult, which reduces the fitness of both Open and Closed Innovators.

In Model 1, Open Innovators only emerge when unspecified benefits to openness are large enough to outweigh the benefits of information free-riding and information hoarding. Models 2, 3 and 4 unpacked this unspecified benefit, replacing it with mechanisms known in evolutionary biology to promote cooperation more generally. Model 2 showed that kin selection can favour information sharing when that sharing is preferentially directed to kin, in this context offspring. This occurs because strategies and traits are inherited together: individuals who share knowledge with their kin produce offspring who inherit a propensity to share with kin and also inherit their parent's accumulated cultural traits. This resembles models of the evolution of teaching, where costly teaching is inherited together with cultural traits whose acquisition is facilitated by that teaching [21,22]. However, Model 2 also showed that the downside of kin-directed innovation is a notable decrease in the rate of CCE due to the reduction in effective sample size of demonstrators from potentially the entire population down to one or two parents, reflecting general findings related to population size and CCE [8,9]. This downside is particularly pronounced when copying fidelity is even slightly less than perfect, and when transmission is biparental rather than uniparental, which partially unlinks the inheritance of strategies and traits. All this suggests that kin-directed information sharing is viable, but more likely for hard-to-acquire traits that are uniparentally transmitted. This fits with empirical cases such as male-kin adze-making apprenticeships in West Papuan hunter-gatherers [6].

Model 3 showed that reputation-based partner choice is another viable solution to the informational dilemma. If sharers acquire reputations by sharing, and others preferentially direct their sharing to those with high reputations, then the benefits of sharing are returned back to the sharer. Partner choice increases in effectiveness when high reputation individuals can be more easily identified and when reputations are global (known by everyone) rather than local (acquired via direct interaction). Such factors are more likely as populations become more interconnected and communication improves. Partner choice is also more likely to favour information sharing when there is a 'priority advantage' such that the inventor of an idea receives a higher reputation increase for sharing that idea than someone who shared it second-hand and did not invent it. This sheds light on frequent and sometimes fierce priority battles in the history of science and technology [46]. Nevertheless, this was not an absolute condition, and even when reputations increased for sharing information irrespective of provenance, partner choice was still effective. Overall, the robustness of partner choice supports suggestions from historians pointing to the importance of reputation and partner choice in the emergence of open knowledge sharing in periods such as the Enlightenment [10].

Finally, Model 4 showed that cultural group selection is another viable solution to the informational dilemma, where individuals preferentially share costly innovations exclusively with other members of their group, and groups of such group-biased innovators out-compete groups of Non-Innovators or Closed Innovators due to the former's superior culturally accumulated knowledge. However, as in previous non-informational models, cultural group selection acting alone is highly vulnerable to migration breaking down between-group variation and preventing group-based information sharers from associating with each other. Adding punishment greatly increases the range of conditions under which cultural group selection favoured information sharing.

These findings suggest several qualitative signatures to look for in historical and ethnographic data which may be indicative of the different cooperation mechanisms. Model 2 suggests that hard-to-acquire traits should co-occur with uniparental kin-directed cultural transmission, which resembles the aforementioned adze-making apprenticeships in West Papuan hunter-gatherers, but may also be observed in other small-scale societies where vertical transmission is common. Model 3 suggests that innovators should be concerned with and foster their reputations for sharing, and priority advantage is important but not crucial for partner choice to work. Model 4 suggests that any group-wide advantage due to group-directed sharing is likely to covary with within-group mechanisms of cooperation such as punishment. Potentially the latter is needed before the former, which could be tested historically. Given that migration is a major obstacle to cultural group selection working, we might expect specific mechanisms to have evolved to deal with this, such as the aforementioned ban on glassmakers leaving the Venetian Republic, which could be seen as migration-specific punishment.

There are several limitations of our models that warrant caution in our conclusions and should stimulate further modelling work. First, we assumed for simplicity discrete behavioural strategies (Open Innovators, Non-Innovators, etc.). It is clearly unrealistic to assume that individuals either always innovate or never innovate, and always share or never share. Allowing innovation or sharing to be a continuous probability may well change the dynamics of our models, as has been found for non-informational models of cooperation [54]. Second, in both our analytic (Model 1a) and agent-based (Models 1b-5) models we implemented a crude form of cumulative culture. In the analytic model, the benefits of knowledge increase in proportion to the amount of openly-shared knowledge. In the agent-based models, traits accumulate within a large trait-space that expands once 90% of traits have been discovered. More realistic implementations of CCE might allow for path-dependence and parallel lineages [4]. We might also assume that the benefit of traits reduces over time according to a discount factor such that knowledge gradually becomes out-of-date, akin to Schumpeterian 'creative destruction' [58]. Alternatively, the benefits of knowledge might increase with the number of users of that knowledge (i.e. network effects: [59]). Both of these factors might favour more open innovation than in our current models. Conversely, innovation might become harder as knowledge accumulates [60], reducing the benefits of innovation over time and favouring less open innovation. Third, the cooperation mechanisms implemented in Models 2-4 are all first steps and could be extended in multiple ways. Kin selection in Model 2 is implemented as vertical cultural transmission, but this could be extended by examining phenomena such as kin competition, dispersal, assortative mating, or paternity certainty [37]. Model 3 could incorporate the explicit transmission and perhaps dishonest faking of reputations, while Model 4 could be extended by incorporating preferential rather than random migration [61], and other within-group mechanisms such as reciprocity or conformity [54]. Fourth, while these models are rather generic, intended to clarify the logic of how cooperation mechanisms might apply to informational dilemmas, models might also be developed for specific historical scenarios, such as Mokyr's [10] competitive patronage hypothesis for the Enlightenment, which resembles a bidirectional version of our partner choice in Model 3. Finally, an additional non-cooperative mechanism that we did not implement at all is information theft, where individuals might pay an initial cost, and also risk costly punishment, for accessing another individual's knowledge against their will.

The questions raised here have parallels in other literatures. In economics, endogenous growth theory conceptualises technological change as a cumulative process [25], sharing our assumptions here that knowledge is nonrivalrous such that its benefits replicate rather than deplete when it is shared, that people behave rationally to maximise benefits relative to costs, and that knowledge is partially excludable such that innovation occurs when the benefits returned to innovators exceed the costs of innovation. However, endogenous growth theory is typically concerned with economic growth in contemporary industrialised societies rather than the broader process of CCE itself across societies and throughout history, as we are here, and seldom considers more universal and concrete mechanisms of cooperation drawn from evolutionary biology and cultural evolution theory, as we do. Elsewhere, scholars have transferred theory designed for managing material commons such as irrigation systems or fisheries to 'knowledge commons' [5], with parallels to the models presented here. Again, however, there is little engagement with mechanisms of cooperation from evolutionary biology (although see [62] for links with non-knowledge commons), and little formal modelling.

Analysing patent systems through the lens of cultural evolution would be a fruitful avenue for future research. Patent systems can be seen as culturally evolved, legally enforced institutions for solving the problem of information free-riding in modern capitalist societies. They work by giving innovators exclusive rights to the innovation's use, or a share in the profits from that innovation, for a fixed time period. Yet patents come with a fundamental trade-off [12]: longer patents provide higher incentives to innovators so make innovation more likely, facilitating CCE, yet longer patents also reduce the ability of other innovators to build on that patented knowledge, inhibiting CCE. There is little theoretical or empirical consensus on how this trade-off can be optimally balanced [13]. In our models, Closed Innovators are essentially permanent patent-holders, so time-limiting their ability to hoard information could provide a means of modelling the patent system within our broad framework. The priority advantage implemented in Model 3 whereby only the first inventor of a trait gets a reputational reward also resembles the patent system, where likewise only the first to patent an idea or discovery receives benefits [63]. Unlike patents, however, there is no trade-off inherent in reputation-based priority advantage systems, where ideas or discoveries can be immediately used or learned by others [63]. Technological evolution may be characterised by patents, while scientific evolution by reputation [63].

In conclusion, our models have clarified the exact nature of the informational dilemma at the heart of CCE - informational free-riding and informational hoarding - and presented potential mechanisms known from evolutionary biology that might solve this dilemma in different ways. We hope that these simple, initial models will stimulate further modelling and empirical work to further identify the conditions under which information is likely to be shared or hoarded, and the consequences for cumulative cultural evolution.

Acknowledgements

We thank Joanna Bryson, Marco Campenni, Daniel Sankey and Cameron Turner for valuable comments on the models and manuscript.

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."

References

1. Henrich J. 2015 *The secret of our success.* Princeton University Press.

2. Hill KR, Barton M, Hurtado AM. 2009 The emergence of human uniqueness. *Evolutionary Anthropology* **18**, 187–200.

3. Tomasello M. 1999 *The cultural origins of human cognition.* Cambridge, MA: Harvard Univ. Press.

4. Mesoudi A, Thornton A. 2018 What is cumulative cultural evolution? *Proceedings of the Royal Society B* **285**, 20180712. (doi:10.1098/rspb.2018.0712)

5. Hess C, Ostrom E. 2011 *Understanding knowledge as a commons.* MIT Press.

6. Stout D. 2002 Skill and cognition in stone tool production. *Current Anthropology* **43**, 693–722.

7. Macfarlane A, Martin G. 2002 *Glass: A world history.* University of Chicago Press.

8. Henrich J. 2004 Demography and cultural evolution. *American Antiquity* **69**, 197–214.

9. Derex M, Mesoudi A. 2020 Cumulative cultural evolution within evolving population structures. *Trends in Cognitive Sciences* **24**, 654–667. (doi:10.1016/j.tics.2020.04.005)

10. Mokyr J. 2016 *A culture of growth: The origins of the modern economy.* Princeton University Press.

11. Open Science Collaboration. 2015 Estimating the reproducibility of psychological science. *Science* **349**, aac4716. (doi:10.1126/science.aac4716)

12. Nordhaus WD. 1969 *Invention growth, and welfare: A theoretical treatment of technological change.* MIT Press. See https://cir.nii.ac.jp/crid/1130282269289605888.

13. Budish E, Roin BN, Williams H. 2016 Patents and research investments: Assessing the empirical evidence. *American Economic Review* **106**, 183–187. (doi:10.1257/aer.p20161091)

14. Apicella CL, Silk JB. 2019 The evolution of human cooperation. *Current Biology* **29**, R447–R450. (doi:10.1016/j.cub.2019.03.036)

15. Smaldino PE. 2016 Models are stupid, and we need more of them. In *Computational Models in Social Psychology* (eds RR Vallacher, A Nowak, SJ Read), Psychology Press.

16. Rogers AR. 1988 Does biology constrain culture? *American Anthropologist* **90**, 819–831.

17. Enquist M, Eriksson K, Ghirlanda S. 2007 Critical social learning: A solution to Rogers' paradox of nonadaptive culture. *American Anthropologist* **109**, 727–734.

18. Rendell L, Fogarty L, Laland KN. 2010 Rogers' paradox recast and resolved: Population structure and the evolution of social learning strategies. *Evolution* **64**, 534–548. (doi:10.1111/j.1558-5646.2009.00817.x)

19. Kameda T, Nakanishi D. 2003 Does social/cultural learning increase human adaptability? Rogers' question revisited. *Evolution and Human Behavior* **24**, 242–260.

20. Lehmann L, Feldman MW. 2009 Coevolution of adaptive technology, maladaptive culture and population size in a producer-scrounger game. *Proceedings of the Royal Society B* **276**, 3853–3862.

21. Turner CR, Mann SF, Spike M, Magrath RD, Sterelny K. 2023 Joint evolution of traits for social learning. *Behavioral Ecology and Sociobiology* **77**, 47. (doi:10.1007/s00265-023-03314-w)

22. Fogarty L, Strimling P, Laland KN. 2011 The evolution of teaching. *Evolution* **65**, 2760–2770. (doi:10.1111/j.1558-5646.2011.01370.x)

23. Cavalli-Sforza LL, Feldman MW. 1983 Paradox of the evolution of communication and of social interactivity. *Proceedings of the National Academy of Sciences* **80**, 2017–2021. (doi:10.1073/pnas.80.7.2017)

24. Čače I, Bryson JJ. 2007 Agent Based Modelling of Communication Costs: Why Information Can Be Free. In *Emergence of Communication and Language* (eds C Lyon, CL Nehaniv, A Cangelosi), pp. 305–321. London: Springer. (doi:10.1007/978-1-84628-779-4_16)

25. Romer PM. 1990 Endogenous technological change. *Journal of Political Economy* **98**, S71–S102. (doi:10.1086/261725)

26. McElreath R. 2012 Baryplot: Barycentric plots. R package version 1.1.

27. Acerbi A, Mesoudi A, Smolla M. 2022 *Individual-based models of cultural evolution: A step-by-step guide using R.* London: Routledge. (doi:10.4324/9781003282068)

28. Smaldino P. 2023 *Modeling Social Behavior: Mathematical and Agent-Based Models of Social Dynamics and Cultural Evolution.* Princeton University Press.

29. Klepper S, Simons KL. 2000 The making of an oligopoly. *Journal of Political Economy* **108**, 728–760.

30. Valverde S, Solé RV. 2015 Punctuated equilibrium in the large-scale evolution of programming languages. *Journal of The Royal Society Interface* **12**, 20150249. (doi:10.1098/rsif.2015.0249)

31. Derex M. 2021 Human cumulative culture and the exploitation of natural phenomena. *Philosophical Transactions of the Royal Society B: Biological Sciences* **377**, 20200311. (doi:10.1098/rstb.2020.0311)

32. Hamilton WD. 1964 The genetical evolution of social behaviour I and II. *Journal of Theoretical Biology* **7**, 1–52.

33. Clutton-Brock T. 2009 Cooperation between non-kin in animal societies. *Nature* **462**, 51–57. (doi:10.1038/nature08366)

34. Kurzban R, Burton-Chellew MN, West SA. 2015 The evolution of altruism in humans. *Annual Review of Psychology* **66**, 575–599. (doi:10.1146/annurev-psych-010814-015355)

35. Silk JB. 1980 Adoption and kinship in Oceania. *American Anthropologist* **82**, 799–820.

36. Daly M, Wilson M. 1988 Evolutionary social psychology and family homicide. *Science* **242**, 519–524. (doi:10.1126/science.3175672)

37. Marshall JAR. 2015 *Social evolution and inclusive fitness theory.* Princeton University Press. (doi:10.1515/9781400866564)

38. Cavalli-Sforza LL, Feldman MW. 1981 *Cultural transmission and evolution.* Princeton: Princeton Univ. Press.

39. Harris JR. 1995 Where is the child's environment? A group socialization theory of development. *Psychological Review* **102**, 458–489.

40. Aunger R. 2000 The life history of culture learning in a face-to-face society. *Ethos* **28**, 1–38.

41. Henrich J, Broesch J. 2011 On the nature of cultural transmission networks: Evidence from Fijian villages for adaptive learning biases. *Philosophical Transactions of the Royal Society of London B* **366**, 1139–1148. (doi:10.1098/rstb.2010.0323)

42. Demps K, Zorondo-Rodríguez F, García C, Reyes-García V. 2012 Social learning across the life cycle: Cultural knowledge acquisition for honey collection among the Jenu Kuruba, India. *Evolution and Human Behavior* **33**, 460–470. (doi:10.1016/j.evolhumbehav.2011.12.008)

43. Lew-Levy S, Kissler SM, Boyette AH, Crittenden AN, Mabulla IA, Hewlett BS. 2020 Who teaches children to forage? Exploring the primacy of child-to-child teaching among Hadza and BaYaka Hunter-Gatherers of Tanzania and Congo. *Evolution and Human Behavior* **41**, 12–22. (doi:10.1016/j.evolhumbehav.2019.07.003)

44. Roberts G. 1998 Competitive altruism: From reciprocity to the handicap principle. *Proceedings of the Royal Society of London B: Biological Sciences* **265**, 427–431. (doi:10.1098/rspb.1998.0312)

45. Roberts G, Raihani N, Bshary R, Manrique HM, Farina A, Samu F, Barclay P. 2021 The benefits of being seen to help others: Indirect reciprocity and reputation-based partner choice. *Philosophical Transactions of the Royal Society B: Biological Sciences* **376**, 20200290. (doi:10.1098/rstb.2020.0290)

46. Merton RK. 1957 Priorities in scientific discovery: A chapter in the sociology of science. *American Sociological Review* **22**, 635–659. (doi:10.2307/2089193)

47. Nowak MA, Sigmund K. 1998 Evolution of indirect reciprocity by image scoring. *Nature* **393**, 573–577.

48. Leimar O, Hammerstein P. 2001 Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **268**, 745–753. (doi:10.1098/rspb.2000.1573)

49. Campenni M, Schino G. 2014 Partner choice promotes cooperation: The two faces of testing with agent-based models. *Journal of Theoretical Biology* **344**, 49–55. (doi:10.1016/j.jtbi.2013.11.019)

50. Richerson PJ *et al.* 2016 Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences* **39**. (doi:10.1017/S0140525X1400106X)

51. Smith D. 2020 Cultural group selection and human cooperation: A conceptual and empirical review. *Evolutionary Human Sciences* **2**, e2. (doi:10.1017/ehs.2020.2)

52. Boyd R, Richerson PJ. 1985 *Culture and the evolutionary process.* Chicago, IL: Univ. Chicago Press.

53. Boyd R, Gintis H, Bowles S, Richerson PJ. 2003 The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences* **100**, 3531–3535. (doi:10.1073/pnas.0630443100)

54. Efferson C, Bernhard H, Fischbacher U, Fehr E. 2024 Super-additive cooperation. *Nature* **626**, 1034–1041. (doi:10.1038/s41586-024-07077-w)

55. Fehr E, Gächter S. 2002 Altruistic punishment in humans. *Nature* **415**, 137–140. (doi:10.1038/415137a)

56. Mathew S, Boyd R. 2011 Punishment sustains large-scale cooperation in prestate warfare. *Proceedings of the National Academy of Sciences* **108**, 11375–11380. (doi:10.1073/pnas.1105604108)

57. Mesoudi A. 2018 Migration, acculturation, and the maintenance of between-group cultural variation. *PLOS ONE* **13**, e0205573. (doi:10.1371/journal.pone.0205573)

58. Schumpeter JA. 1942 *Capitalism, socialism and democracy.* 5th edn. New York: Harper. (doi:10.4324/9780203202050)

59. Swann GMP. 2002 The functional form of network effects. *Information Economics and Policy* **14**, 417–429. (doi:10.1016/S0167-6245(02)00051-3)

60. Mesoudi A. 2011 Variable cultural acquisition costs constrain cumulative cultural evolution. *PLOS One* **6**, e18239. (doi:10.1371/journal.pone.0018239)

61. Boyd R, Richerson PJ. 2009 Voting with your feet: Payoff biased migration and the evolution of group beneficial behavior. *Journal of Theoretical Biology* **257**, 331–339. (doi:10.1016/j.jtbi.2008.12.007)

62. Wilson DS, Ostrom E, Cox ME. 2013 Generalizing the core design principles for the efficacy of groups. *Journal of Economic Behavior & Organization* **90**, S21–S32. (doi:10.1016/j.jebo.2012.12.010)

63. Stephan PE. 1996 The economics of science. *Journal of Economic Literature* **34**, 1199–1235.

Supplementary material for 'From information free-riding to information sharing: how have humans solved the cooperative dilemma at the heart of cumulative cultural evolution?'

Alex Mesoudi[1], Ángel Jiménez[1], Keith Jensen[2] & Lei Chang[3]

[1]Human Behaviour and Cultural Evolution Group, Centre for Ecology and Conservation, Penryn Campus, University of Exeter, UK

[2]Division of Human Communication, Development & Hearing, University of Manchester, UK

[3]Department of Psychology, University of Macau, Macau

27 May, 2024

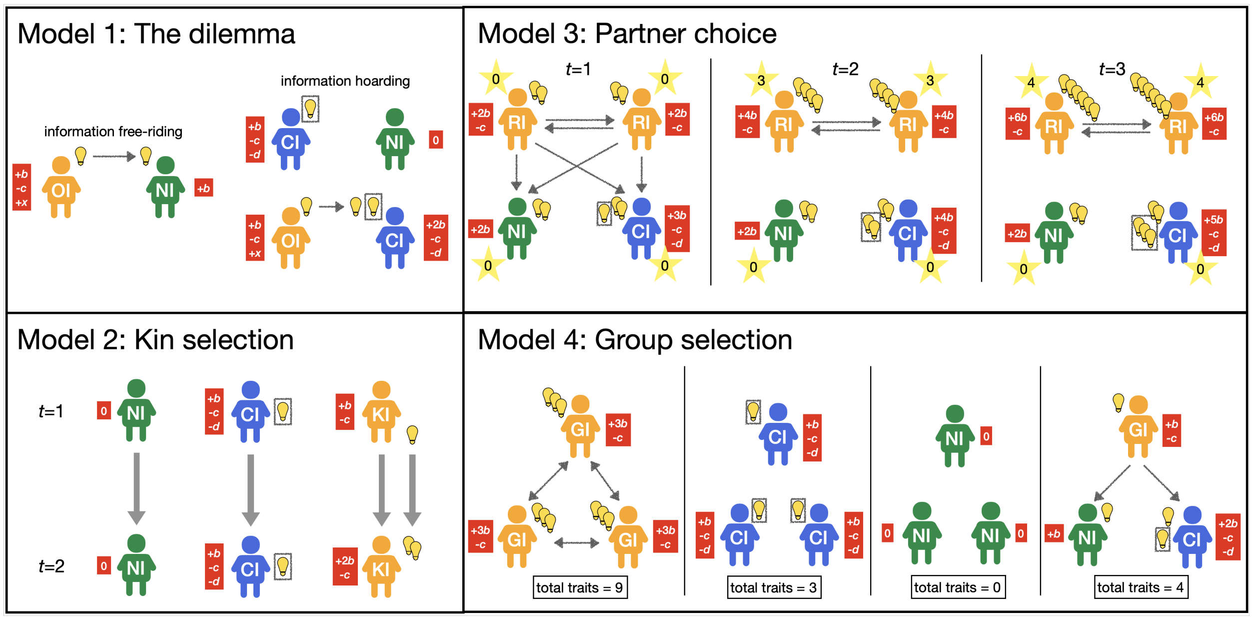

Figure S1. A schematic of the logic underlying the models. Agent types are denoted by colour and initials (OI = Open Innovator; NI = Non-Innovator; CI = Closed Innovator; KI = Kin-based Innovator; RI = Reputation-based Innovator; GI = Group Innovator); traits as lightbulbs, with lightbulbs enclosed in boxes denoting hoarded knowledge; costs and benefits in red boxes (see Table S1 for parameter definitions); information flows by grey arrows; and reputations in Model 3 in stars. Model 1 explores the basic dilemma of costly innovation. Information free-riding occurs when Open Innovators (OI; orange) produce knowledge at cost c which benefits them b, while Non-Innovators (NI; green) copy that information and receive benefit b but without paying any cost. Information hoarding occurs when Closed Innovators (CI; blue) produce knowledge that is protected from being copied by others; when interacting with OIs, CIs then have both their private knowledge (enclosed in a grey box) plus OI's shared public knowledge. In Model 2 (kin selection), strategies are passed from parents in $t = 1$ to offspring in $t = 2$. Lineages of Kin-based Innovators (KI; orange) also culturally transmit their traits, unlike lineages of NIs who have no traits to transmit, and CIs who protect their traits. In Model 4 (partner choice), agents have reputations for sharing shown in stars, and Reputation-based Innovators (RI; orange) share only with others who have the highest reputations. Initially, the lack of reputations causes the two RIs to share with every other agent. From $t = 2$, RIs share only with each other, as they have higher reputations than both NIs and CIs who do not share. Eventually this results in RIs accumulating more traits than the other types. In Model 5 (group selection), Group Innovators (GI; orange) only share with other members of their group. Here there are three groups, one all GIs, one all CIs, one all NIs, and one mixed. Cultural group selection favours the group with the most number of traits, which here is the group of GIs. Within the mixed group, however, GIs do poorly relative to the other types.

Table S1. Definitions of model parameters. Paramters are common to all models unless specified. ABMs = Agent-Based Models (Models 1b, 2, 3 and 4). All probabilities are denoted p, non-probabilities with other symbols.

Parameter	Definition
c	The cost to an Open or Closed Innovator of innovation
b	The benefit of innovation to the innovator and, if released, to every other individual
d	The cost to Closed Innovators of attempting to protect their innovation
p_h	The probability of a Closed Innovator successfully protecting their innovation
x	The unspecified benefit to individuals for releasing their innovation to others (Model 1 only)
γ	The proportion of open knowledge that b increases by in each timestep (Model 1a only)
N	The total number of agents in the population (ABMs only)
L	The total number of learnable traits (ABMs only)
p_i	The probability of innovation (ABMs only)
p_c	The probability of copying a trait (ABMs only)
p_s	The probability of payoff-biased copying of strategies (ABMs only)
p_d	The probability of each agent in each timestep dying and being replaced with a new unknowledgable agent (ABMs only)
n_c	The number of candidate learners from whom agents select the one with the highest reputation with whom to share (Model 3 only)
r_s	The reputational increase from sharing any trait with another agent (Model 3 only)
r_i	The reputational increase from sharing a trait that this agent has innovated with another agent (Model 3 only)
n	The number of agents per group (Model 4 only)
g	The number of groups (Model 4 only)
p_g	The probability of intergroup conflict (Model 4 only)
k	The fitness cost to the punishing Group Innovator agent per punished agent (Model 4 only)
u	The fitness cost to the punished agent per punishing agent (Model 4 only)

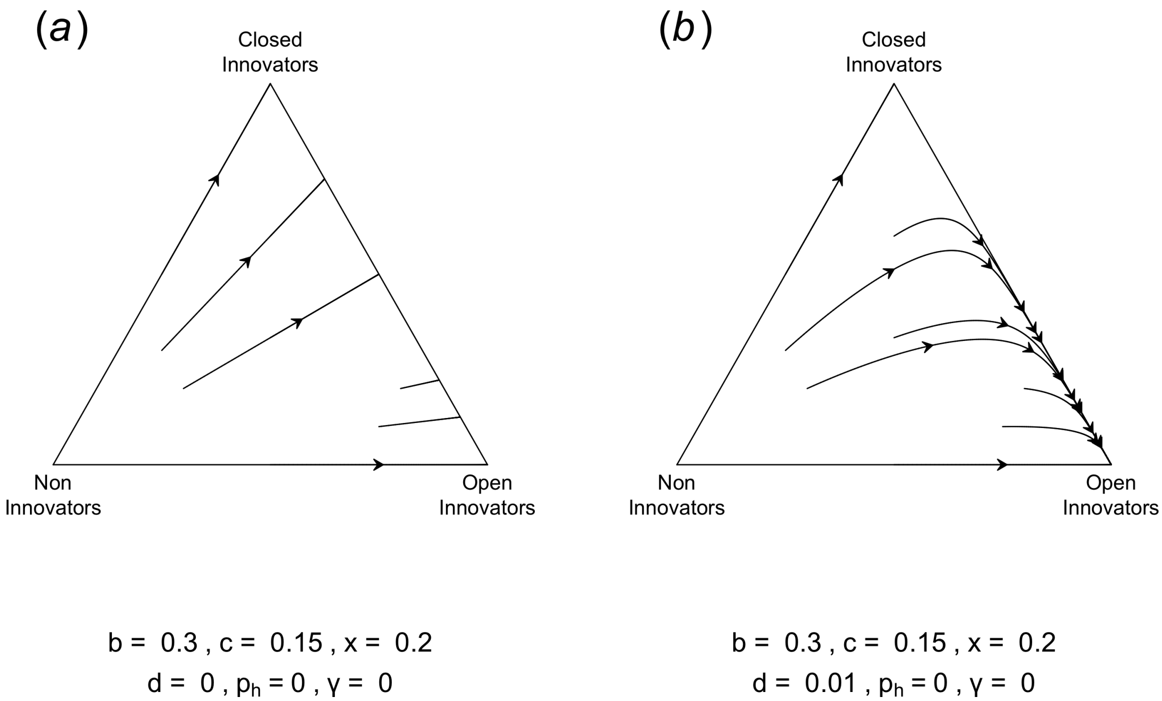

Figure S2. Reducing the probability of Closed Innovator agents successfully hoarding their knowledge (p_h) favours Open Innovators in Model 1a. (a) Setting $p_h = 0$, meaning that Closed Innovators can never protect their knowledge, makes it easier for Open Innovators to replace Closed Innovators, assuming $x > c$. When $d = 0$, Open Innovators and Closed Innovators have the same payoffs. (b) Increasing d any amount with otherwise identical parameter values as in (a) favors Open Innovators.

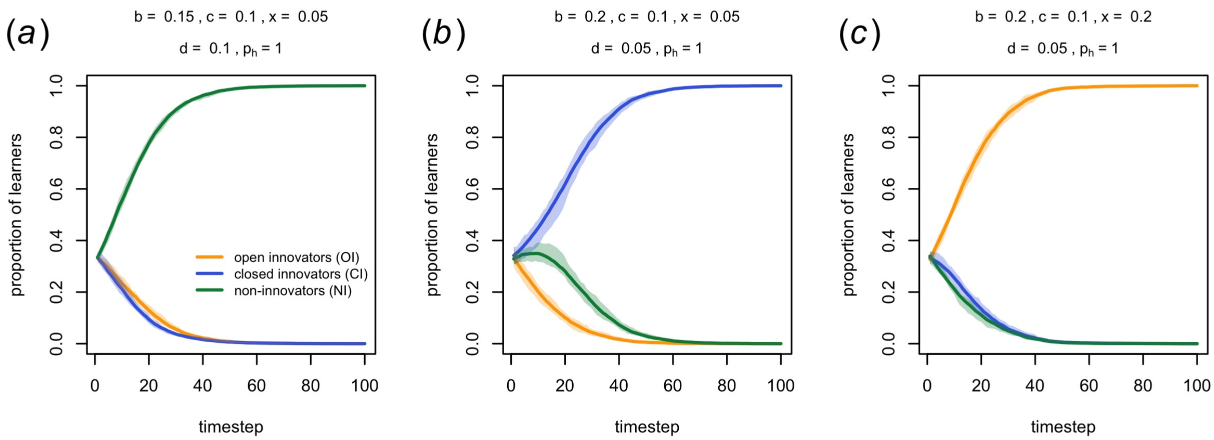

Figure S3. Time series from Model 1b for the same parameter values as shown in Figure 1 for Model 1a, which favour (a) Non-Innovators, (b) Closed Innovators, and (c) Open Innovators, respectively. To match Model 1a in all other respects, we assume perfect innovation ($p_i = 1$) and copying ($p_c = 1$), and non-overlapping timesteps ($p_d = 1$). This replicates and confirms the analytical Model 1a but within the agent-based modelling framework of Model 1b. (Other parameters: $N = 1000$, $L = 1000$, $p_s = 0.1$. Lines shown are means of 10 independent runs, with shaded areas showing the range across all runs.)

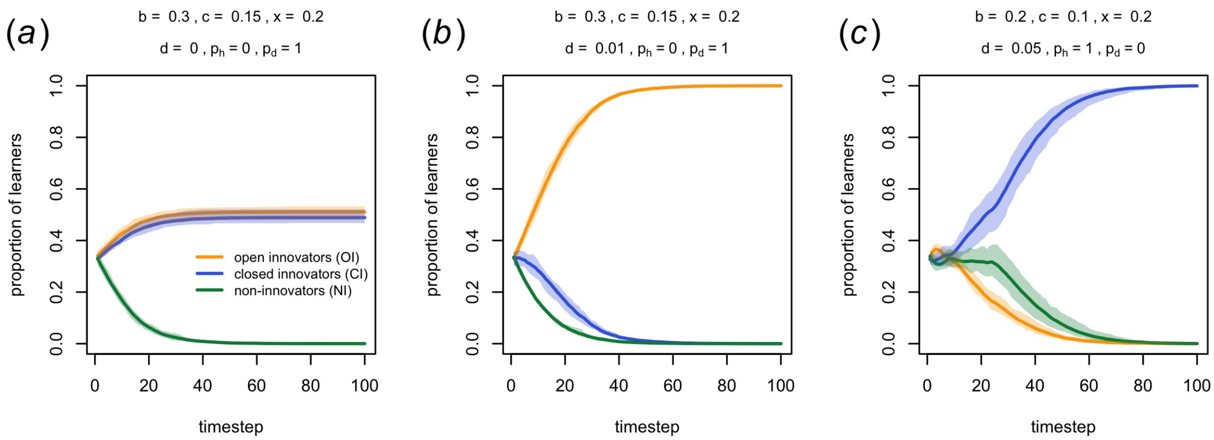

Figure S4. Time series from Model 1b for the same parameter values as shown in figure S2 and figure 2 for Model 1a, replicating the results for Model 1a using the agent-based Model 1b. (a) When Closed Innovators cannot hoard their information ($p_h = 0$), they co-exist with Open Innovators when hoarding is costless ($d = 0$), as in figure S2(a). (b) When hoarding is costly ($d > 0$), Open Innovators outperform Closed Innovators, as in figure S2(b). (c) When agents never die ($p_d = 0$), allowing the accumulation of knowledge over timesteps, then Closed Innovators who can hoard outperform other agents as they can accumulate private knowledge, as well as benefit from accumulated public knowledge generated by Open Innovators, as in figure 2. (Other parameters: $N = 1000$, $L = 1000$, $p_c = 1$, $p_i = 1$, $p_s = 0.1$. Lines shown are means of 10 independent runs, with shaded areas showing the range across all runs.)

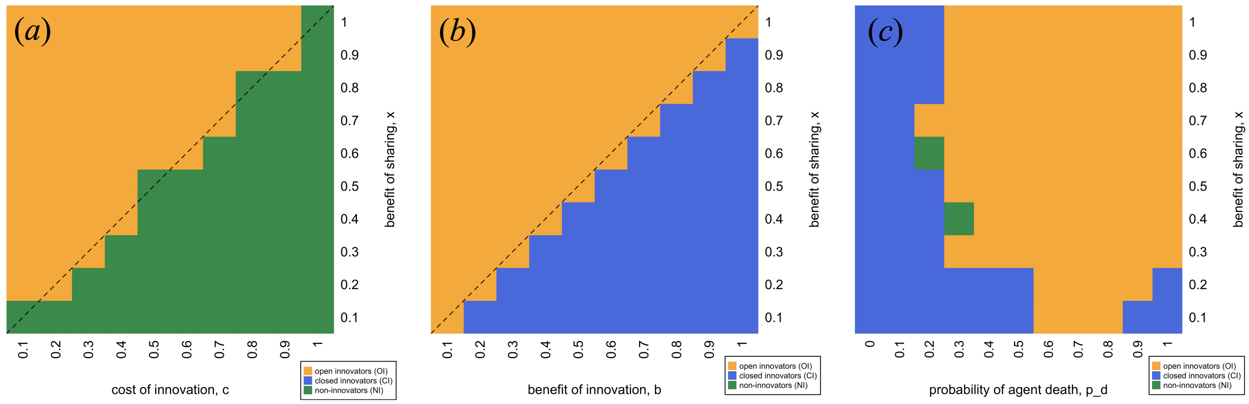

Figure S5. Heatmaps showing the most frequent strategy (Open Innovation, orange; Closed Innovation, blue; Non-Innovation, green) across different parameter values in the agent-based Model 1b. (*a*) Open Innovation is favoured over Non-Innovation when the benefit of sharing is greater than the cost of innovation (i.e. $x > c$), as predicted by the analytic model. (*b*) Open Innovation is favoured over Closed Innovation when the benefit of sharing is greater than the benefit of innovation (i.e. $x > b$), as predicted by the analytic model. (*c*) When timesteps overlap ($p_d << 1$), then Closed Innovators are favoured, as they can accumulate private beneficial knowledge.

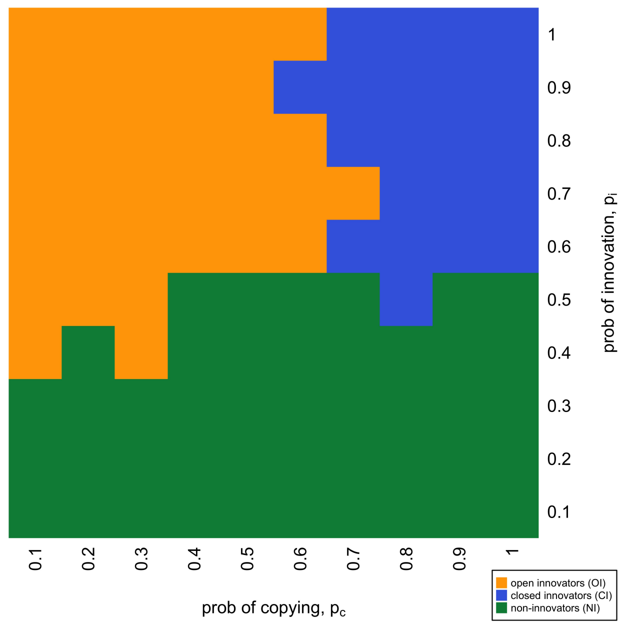

Figure S6. Heatmap showing the most frequent strategy (Open Innovation, orange; Closed Innovation, blue; Non-Innovation, green) across different probabilities of successful innovation p_i and different probabilities of copying traits when shared p_c, for Model 1b. When the probability of innovation p_i is low, Non-Innovators perform better given that Open Innovators and Closed Innovators both bear a cost for innovation that it increasingly likely to yield no benefits. When the probability of innovation p_i is high and the probability of copying p_c is high, then Closed Innovators do well because they can easily copy beneficial traits innovated by Open Innovators, as well as keep their own privately innovated traits. When the probability of innovation p_i is high and the probability of copying p_c is low, then Open Innovators do well because most knowledge is now private but they do not bear the cost of hoarding that Closed Innovators do. However, when p_c is lower than p_i then the information dilemma does not apply, as individual learning is effectively less costly than social learning. Other parameter values: $b = 0.2$, $c = 0.1$, $x = 0.1$, $d = 0.05$, $p_h = 1$, $p_d = 1$, $p_s = 0.1$, $t_{max} = 200$, $N = 100$, results are means of 10 independent runs.

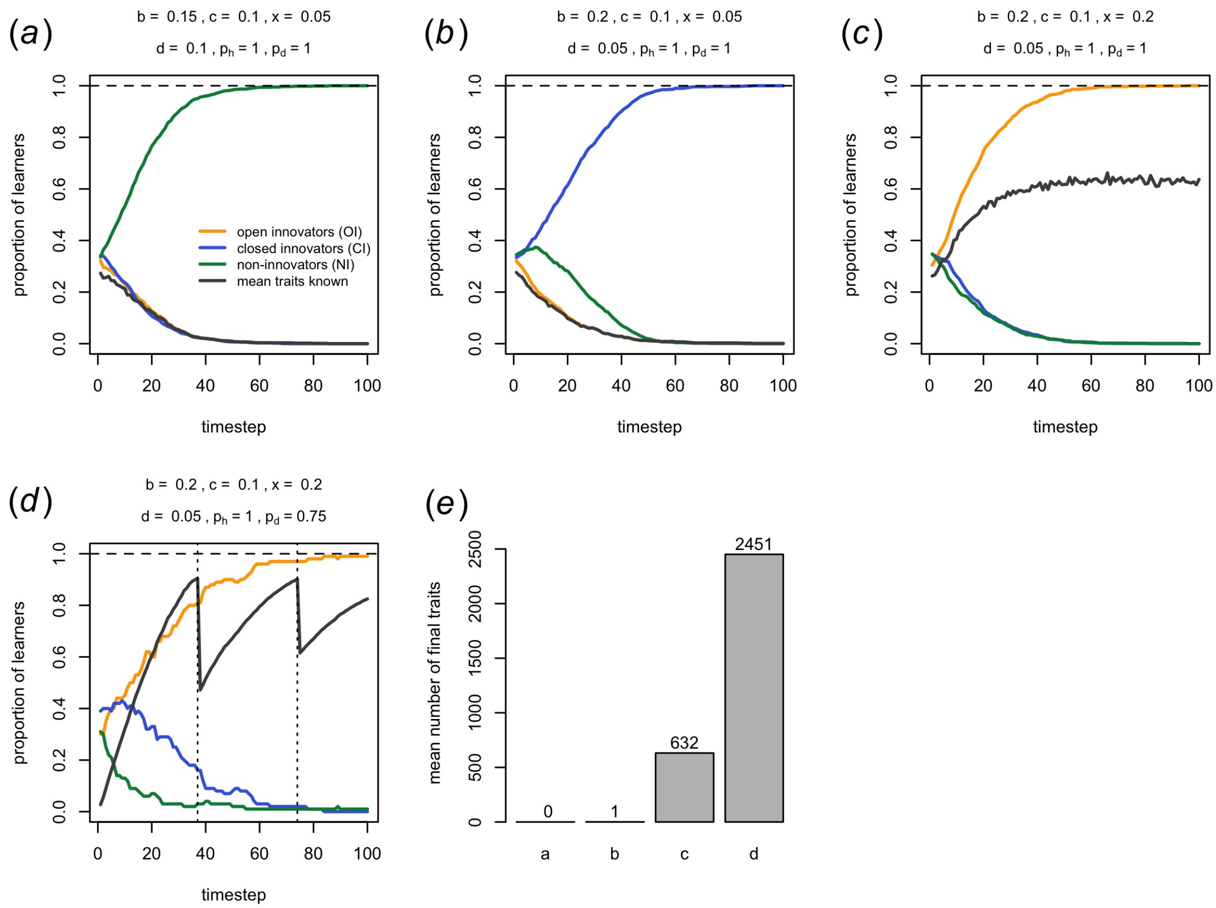

Figure S7. Open Innovation favours the accumulation of cultural traits in Model 1b, especially when there are overlapping timesteps. (a-c) show the equivalent panels from figure S3 with non-overlapping timesteps, for a single run, and with the mean proportion of the L traits learned by all agents shown in dark grey. (a) When Non-Innovators dominate, then the mean proportion of traits known drops to zero given that Non-Innovators do not innovate, and no traits are passed down generations. (b) When Closed Innovators dominate, then the mean proportion of traits known drops to one, given that in each timestep Closed Innovators discover one new trait (because $p_i = 1$) which is not shared within a generation nor passed to subsequent generations. (c) When Open Innovators dominate, then the mean proportion of traits known stabilises at around 0.6. This is because Open Innovators innovate and share their traits with all other agents in their generation, but these are wiped in the next timestep so the next generation starts from scratch. A single generation only learns 60% rather than 100% of the L traits even though $p_i = p_c = 1$ because some agents innovate the same traits as other agents, so not all L traits are available to share. (d) When timesteps overlap, traits can accumulate over time in an open-ended fashion. Here, vertical dotted lines indicate instances when 90% of all L traits are known and another L traits are added to the trait space. This allows continual accumulation of traits amongst the Open Innovators. (e) The mean absolute number of traits known in the final timestep of the scenarios depicted in panels a-d of this figure, averaged across 10 independent runs. Numbers above bars show the mean value depicted in the bar.

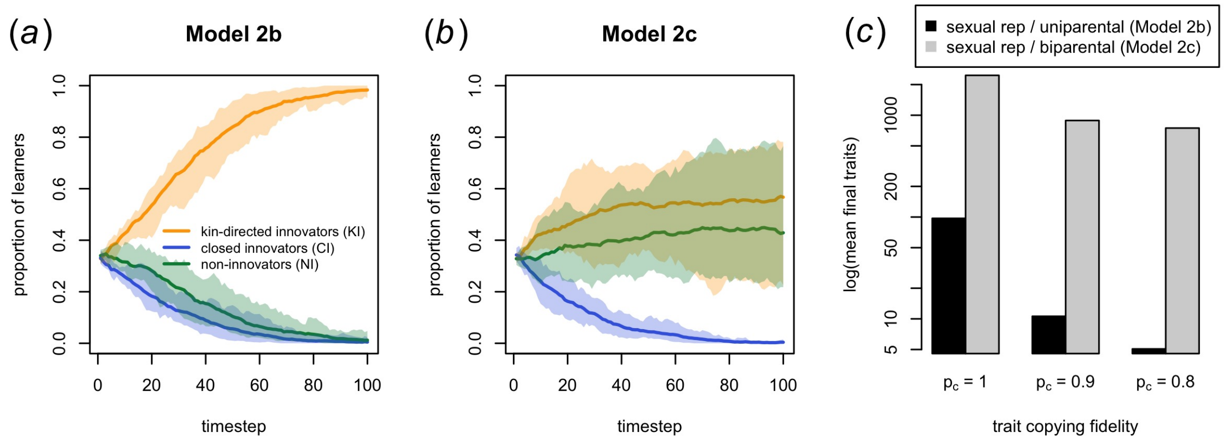

Figure S8. Two extensions of the asexual reproduction kin selection Model 2 reported in the main text, both featuring sexual reproduction. During the social learning phase, each agent picks two other agents at random as potential mates. With probability p_s, they select the potential mate with the higher fitness to be their actual mate, otherwise they pick one of the two at random. This mate selection via p_s replaces payoff-biased copying of strategies. Each pair then produces one offspring (agents can appear in more than one mating pair, so the number of agents remains N). The offspring inherits the strategy of one randomly chosen parent. In Model 2b there is uniparental vertical transmission of traits such that the offspring copies all of the traits of the same parent from whom they inherited their strategy with probability p_c per trait, unless the parent is a hoarder in which case no copying occurs. Model 2c is identical to Model 2b except that traits are learned additively from both non-hoarding parents. Offspring acquire every trait known by either parent with probability p_c per trait. For example, if one parent has trait string 0100101 and the other parent has 1001001, the offspring will attempt to learn each of the five traits known by either parent (the 1s in the combined string 1101101) with probability p_c per trait. Strategies are still inherited from a single randomly chosen parent, as these discrete traits cannot be combined or merged. Results: (a) Time series for Model 2b with sexual reproduction of agents but uniparental transmission of traits, favouring Kin-directed Innovators as in Model 2 (see figure 3b). (b) Time series for Model 2c with sexual reproduction of agents and biparental (blending) transmission of traits, now showing co-existence of Kin-directed Innovators and Non-Innovators. (c) With populations entirely composed of Kin-directed Innovators, biparental transmission (Model 2c) supports more traits at timestep $t = 100$ than does uniparental transmission (Model 2b), and is less vulnerable to copying error, much like horizontal transmission shown in figure 3(c).

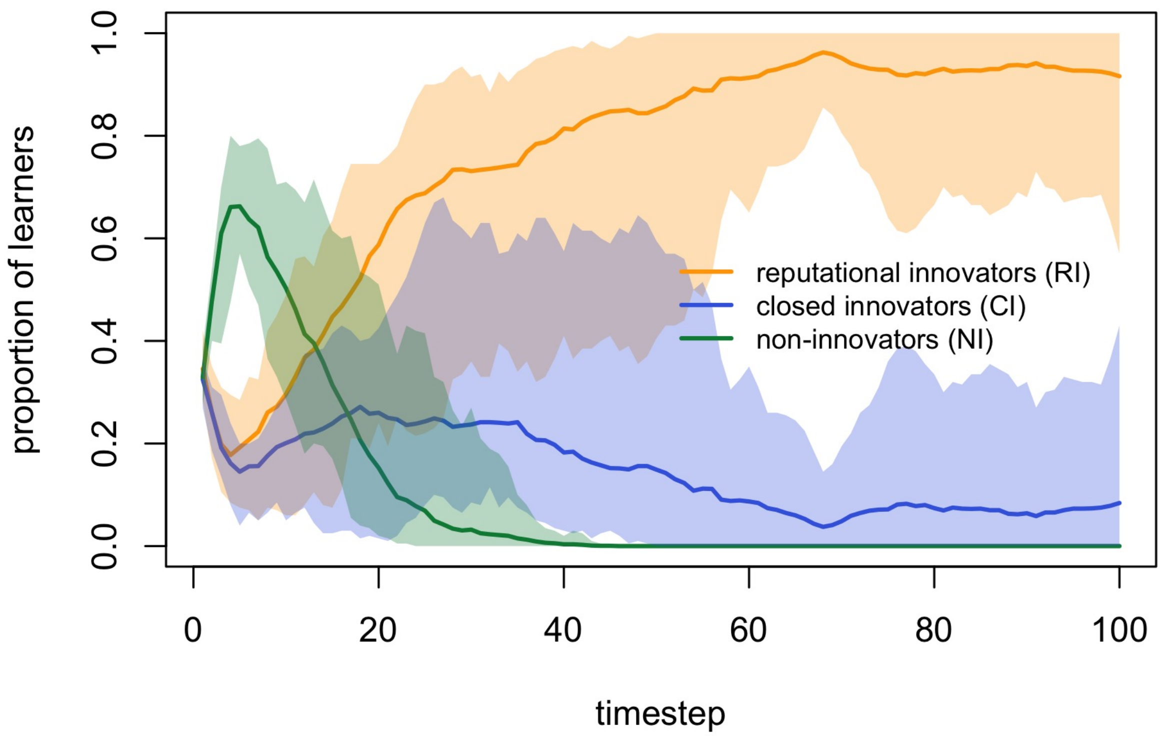

Figure S9. Time dynamics for Model 3 (partner choice) where Reputational Innovators are favoured due to strong partner choice ($n_c = N - 1$), and when reputations are global and increase when any traits are shared ($r_s = 1$, $r_i = 0$). Solid lines show means of 10 independent runs, with shading showing the range across runs. Initially, Non-Innovators do better than Reputational Innovators and increase in frequency due to their increasing reputations for sharing others' traits plus not paying any cost of innovation. However, once Non-Innovators reach a high enough frequency, few new traits are being innovated, and every agent comes to know every known trait. Non-Innovators therefore no longer gain reputation increases because there is nothing new to share. At this point Reputational Innovators have an advantage by innovating new traits and getting a reputation increase from sharing those new traits. Reputational Innovators start to accumulate higher reputations and gradually spread. Parameters: $n_c = 199$, $N = 200$, otherwise as in figure 4.

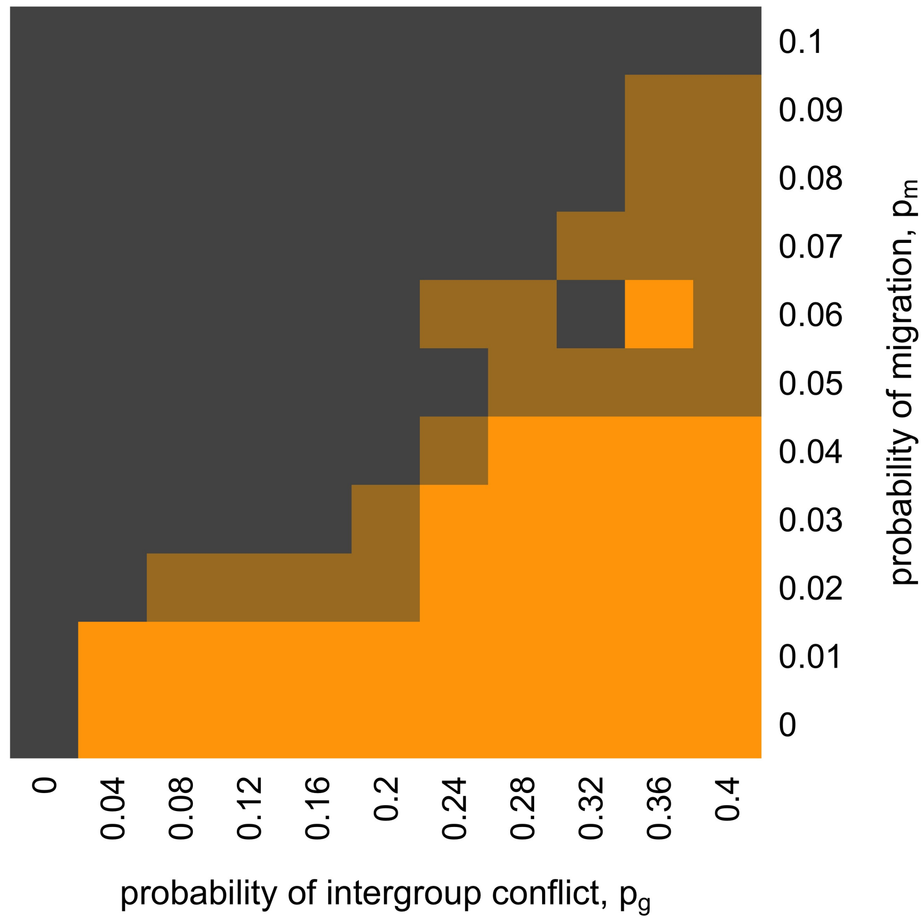

Figure S10. The frequency at equilibrium of Group Innovators in Model 4 (cultural group selection), for different values of the probability of intergroup conflict p_g, and the probability of migration between groups p_m. Colours indicate the mean frequency at equilibrium of Group Innovators across 20 independent simulation runs; orange = frequency greater than 2/3, black = frequency $< 1/3$, grey = intermediate frequencies. Other parameter values: $n = 50$, $g = 20$, $L = 1000$, $p_c = 0.8$, $p_i = 0.2$, $p_h = 0.9$, $p_s = 0.2$, $p_d = 0.1$, $b = 0.2$, $c = 0.1$, $d = 0.05$, $k = 0.2$, $u = 0.4$.

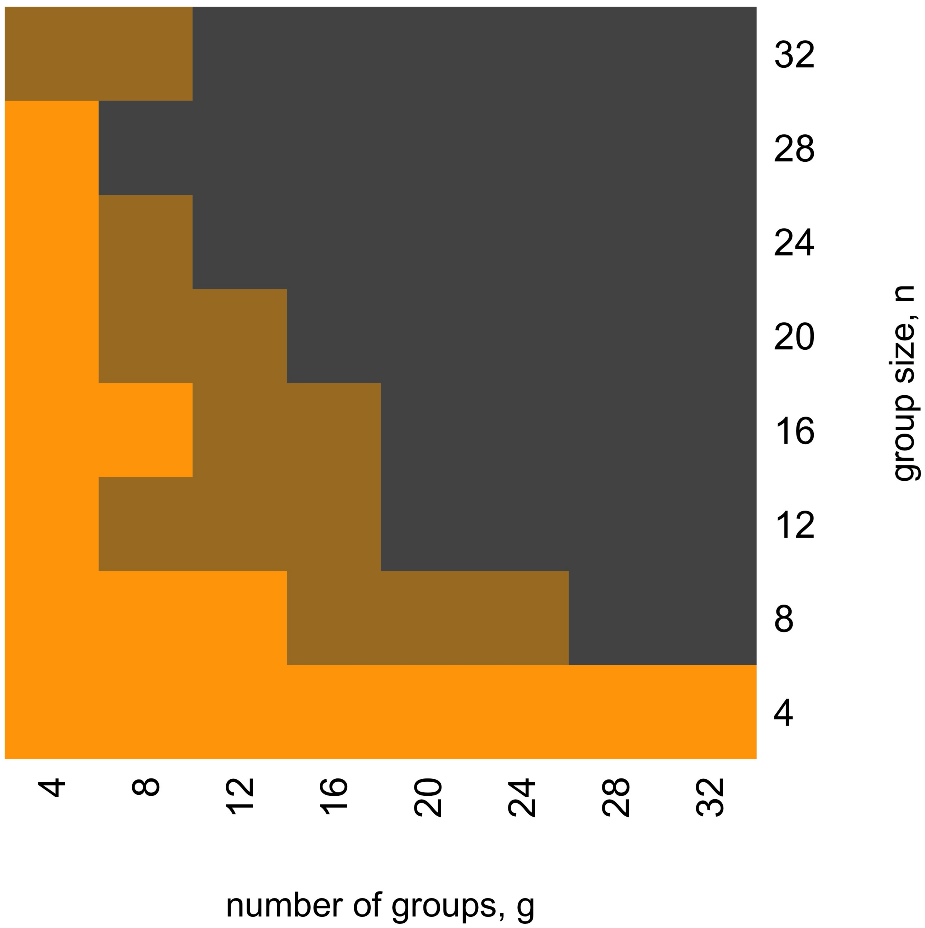

Figure S11. The frequency at equilibrium of Group Innovators in Model 4 (cultural group selection), for different numbers of groups g and different group sizes n. Colours indicate the mean frequency at equilibrium of Group Innovators across 20 independent simulation runs; orange = frequency greater than 2/3, black = frequency $< 1/3$, grey = intermediate frequencies. Other parameter values: $L = 1000$, $p_c = 0.8$, $p_i = 0.2$, $p_h = 0.9$, $p_s = 0.2$, $p_d = 0.1$, $p_m = 0.01$, $b = 0.2$, $c = 0.1$, $d = 0.05$, $k = 0.2$, $u = 0.4$.