scientific reports

OPEN

Cultural diversity and wisdom of crowds are mutually beneficial and evolutionarily stable

Benoît de Courson[1,2], Léo Fitouchi[3], Jean-Philippe Bouchaud[2,4] & Michael Benzaquen[1,2,4✉]

The ability to learn from others (social learning) is often deemed a cause of human species success. But if social learning is indeed more efficient (whether less costly or more accurate) than individual learning, it raises the question of why would anyone engage in individual information seeking, which is a necessary condition for social learning's efficacy. We propose an evolutionary model solving this paradox, provided agents (i) aim not only at information quality but also vie for audience and prestige, and (ii) do not only value accuracy but also reward originality—allowing them to alleviate herding effects. We find that under some conditions (large enough success rate of informed agents and intermediate taste for popularity), both social learning's higher accuracy and the taste for original opinions are evolutionarily-stable, within a mutually beneficial *division of labour*-like equilibrium. When such conditions are not met, the system most often converges towards mutually detrimental equilibria.

In 1785, Condorcet proved that aggregating predictions can be much more accurate than individuals' predictions[1]. As the group size increases, the average prediction can even approach perfect accuracy, provided that individual predictions are unbiased. Empirically, this phenomenon was famously illustrated by Galton's "Vox populi" paper[2]: when he gathered the estimates of an ox weight from 787 participants, the average opinion was just one pound away from reality.

Yet, the very fact that the aggregate prediction is better than individual ones creates a tension: why would anyone make up his/her mind independently when he/she would be better off sampling around and simply endorsing the majority opinion? A very close analogy can be found in the Grossman-Stiglitz paradox in financial economics[3]: if, as often assumed, the price of an asset aggregates all available information, then there is no reason to gather information, which makes it absurd that the price reflects anything to begin with.

From an evolutionary perspective, this parallels the question of the evolution of social and individual learning strategies. Social learning refers to learning in interaction with another individual, whereas individual learning (which we employ here as synonymous to *asocial learning*) refers to learning by oneself, usually by trial and error, without any input from others[4,5]. In this context, the aforementioned efficiency of socially adopting the majority opinion implies that a population where all agents learn independently would swiftly be invaded by a "conformist" strategy (i.e. social learning), as long as the sample size is greater than one, or if there is some information cost. This creates what Boyd, Richerson and Henrich[6,7] called a conformist bias (technically defined by the fact that the probability to adopt the belief held by the majority is greater than the frequency of the majority opinion in the population, see Fig. 1). Importantly, this idea has been invoked to explain group-level cultural differences: as Henrich[8] puts it, "without a conformist component to create 'cultural clumps', social learning models predict (incorrectly) that populations should be a smear of ideas, beliefs, values and behaviours, and that group differences should only reflect local environmental differences."

This rationale brings however an internal tension within Condorcet's jury theorem, which rests on a crucial assumption: that opinion are somewhat independent in the probabilistic sense. But when allowing imitation, the theorem jeopardises the very independence it is assuming: if most of the population imitates others, opinions are no longer independent, but converge towards one another (an effect that has been found empirically[9]). This independence breakdown has consequences—at some point, opinions are so tightly correlated that the accuracy achieved by imitation collapses, as individuals agree more with one another than with reality. The snake bites its own tail: the increase of accuracy achieved by sampling the opinion of others triggers a spread of imitators,

[1]LadHyX, UMR CNRS 7646, Ecole Polytechnique, 91128 Palaiseau Cedex, France. [2]Chair of Econophysics & Complex Systems, Ecole Polytechnique, 91128 Palaiseau Cedex, France. [3]Institut Jean Nicod, Département d'études cognitives, ENS, EHESS, PSL Research University, CNRS, Paris, France. [4]Capital Fund Management, 23-25, Rue de l'Université, 75007 Paris, France. ✉email: michael.benzaquen@polytechnique.edu
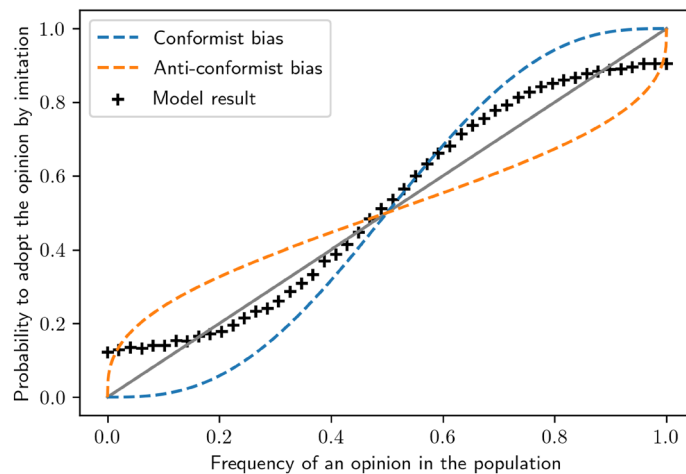
Figure 1. Definition of a (anti-) conformist bias (dashed orange and blue lines) in the cultural evolution literature[6,7], compared with the outcome of our model (explained in Sect. 2.5) where some endogenous heterogeneities are generated.

which continues until all anchor to "truth" has disappeared. Social learners are 'information scroungers', whose spread is reminiscent of Rogers' paradox[10]. However, the paradox we study is different: In Rogers', social learners spread because they do not pay information search cost, whereas this rationale holds even if information is free, and comes from the fact that social learners are more accurate.

This intuition, elegantly modeled by Curty & Marsili[11] (see Sect. 2.1), sheds light on herding phenomena. However, it leads to paradoxical predictions. In particular, at equilibrium,

(i) the vast majority of agents imitate, so that virtually every agent holds the same belief (i.e. full in-group cultural homogeneity, which is also a prediction of the conformist transmission literature[7]). This conclusion can be understood from Fig. 1: if we think of imitation as a dynamic process, then frequencies 0 and 1 are the only stable fixed points of the system.

(ii) hence, the "wisdom of crowds" cannot work, since the aggregation of similar opinion does not lead to more accurate predictions. This result is reminiscent of the "Rogers' paradox"[10] (see Sect. 3.1).

Both conclusions are unrealistic. First, there are countless examples of the wisdom of crowds in action[12–14], including evidence that access to other people's opinion improves collective accuracy[9,15–17].

Second, there is usually a substantial in-group heterogeneity in people's beliefs (see the conclusion of Ref.[18]). A study trying to measure within- and between-group cultural variation incidentally found the first component to be the larger than the second[19]. Furthermore, humans are found to use individual learning quite often[20–22] and to under-use social information[15,23], a phenomenon known as "egocentric discounting"[24], which leads to a sub-optimal situation in terms of belief accuracy, both at the individual[15,22] and collective[25] levels. Also, the very tendency to conformism is disputed[18]: minorities seem to have a literally disproportionate social influence, greater than their share in the population[26].

Anyway, as Eriksson & Coultas[18] argue in response to Henrich's aforementioned statement, "we can turn the argument around: with a conformist component, social learning models predict, incorrectly, that populations should be culturally homogeneous. In this light, it is not surprising that the evidence from social psychology does not support the Conformist Hypothesis." However, an anti-conformist bias does not seem credible either: it would imply evenly split populations (in Fig. 1, 1/2 is the only stable fixed point of the orange curve), which is unrealistic.

Hence, we need a model which, while allowing individuals to exploit the wisdom of crowds, accounts for its effectiveness at equilibrium and for persistent within-group heterogeneity. These two phenomena are actually intertwined. The wisdom of crowds effect is contingent on some opinion diversity in the group, otherwise the opinion of the crowd is the same as that of any individual. In a theoretical model, Scott E. Page analogously showed that *the amount by which the crowd out-predicts its average member increases as the crowd becomes more diverse*[27]. In other words, in order to benefit from the wisdom of crowds effect, one should aggregate *diverse* opinions.

This is a feature that should be incorporated in theoretical models. Previous studies implicitly assume that agents listen to randomly chosen other agents. Instead, we propose that agents can develop a "taste for originality", implemented as a preferential attachment to individuals displaying deviant opinions, in a dynamic network representing who listens to whom. This taste for originality is not only normatively warranted, but also empirically reasonable. Several studies have found that individuals having atypical opinions were rated as more likeable, courageous or admirable[28–30]. Other studies have found that individuals put more weight on advice from people who form their opinions independently[31–33] (see[34] for a review). Similarly, financial analysts who produce

p	Informed accuracy
q	Followers accuracy
s	Fraction of links pointing toward informed agents
γ	Memory decay rate
α	Taste for originality
Ω_i	Network of agent i
S_i^j	Score attributed by i to j
ω	Incentive to gather in-degrees
c	Information searching fitness cost
σ	Mutation rate

Table 1. Notations.

"bold forecasts" tend to be preferentially heeded[35]. Agents seem to be intuitively aware of the danger of blindly following the "crowd".

Being independent-minded generates *interesting*–because unexpected–opinions, likely to be listened to. This could be the primary goal of certain individuals[36]. In current evolutionary models, fitness is proxied by belief accuracy, which means that humans are assumed to be accuracy-maximizers. Here, we propose that humans also care about their social influence (which is empirically related to reproductive success[37]). The dynamic network representing who listens to whom allows one to do so in a simple way: agents simply maximize a combination of their opinion accuracy and the number of people who listen to them.

To sum up, we introduce two novel assumptions: (i) agents have an incentive to be listened to, and (ii) agents can choose to listen to deviant opinions in order to benefit from the diversity effect. This framework allows us to shed a new light on individual learning: although it is sub-optimal in terms of accuracy, it is a way to signal independent mindedness, which makes one worthy to be listened to. It leads to a kind of division of labour equilibrium: as they do not benefit from the wisdom of crowds effect, individual learners achieve lower accuracy, which is offset by the reward they get from being listened to. Such reward may be prestige, influence, or more concretely university positions, consulting fees, copyright, etc. Our scenario is consistent with the empirical findings we mentioned: (i) the effectiveness of the wisdom of crowds effect despite widespread social learning, (ii) a stable non-zero degree of within-group cultural diversity, and (iii) an over-reliance on individual learning, with respect to what accuracy-maximisation would require.

The paper is organised as follows. In Sect. 2.1 we recall the seminal model of Curty & Marsili, which serves as groundwork for our paper. Then, we proceed by making the model gradually more complex. In 2.2, we introduce an adaptive network of imitation and assess the consequences on the model's predictions. In 2.3, we relax the assumption that agents imitate randomly, and we allow them to listen preferentially to original opinions. In 2.5, we introduce an evolutionary framework to endogenize the strategy (individual or social learning) and the taste for originality. In 2.6, we explore the possibility of an anti-conformist bias to "mimic originality".

Methods and results

The Curty-Marsili forecasting game. *The classical setting.* The seminal model of Curty & Marsili[11] considers a population of agents making a binary forecast for the next time step. Table 1 gathers the notations used for the different variables and parameters used throughout the paper. Agents have two possible strategies: looking for information (*I*) or being "follower" (*F*) and relying on others. These two strategies correspond to what is usually called respectively individual and social learning in cultural evolution. An agent playing the *I* strategy has a probability $p > 1/2$ of being right (a quantity that we will, henceforth, call the "accuracy" of the strategy). The follower strategy consists in an iterative imitation process, specified as follows:

(i) All agents playing *F* randomly initialize their opinion, with a 1/2 probability of being right: they have no private information, and thus a neutral prior.
(ii) One of them, randomly chosen, observes a random sample of *m* individuals, and endorses the majority opinion among them (for simplicity, *m* is assumed to be odd to avoid ties).
(iii) Step (ii) is repeated a large number of times until the proportion of followers holding the right opinion converges to a value $\hat{q}$.

Unlike the *I* strategy, the *F* accuracy depends on the fraction *z* of followers in a non-trivial way, see Fig. 2. For a small fraction *z*, there is only one attractive fixed point $q_>$, larger than *p*, which means that herding is efficient. This is simply the Condorcet theorem mechanism at play: aggregating information yields more accurate predictions. Furthermore, the performance of followers increases at first with their number, as the probability to sample from an accurate follower increases, thereby improving sample quality. Notwithstanding, above a critical value z_c, the system becomes bi-stable: $\hat{q}$ has two attractive fixed points (solid red line), and an unstable one in the middle (dashed red line), which marks the frontier between the two basins of attraction. Thus, the population has two possible end states, one where most people are right, and another where most people are wrong.
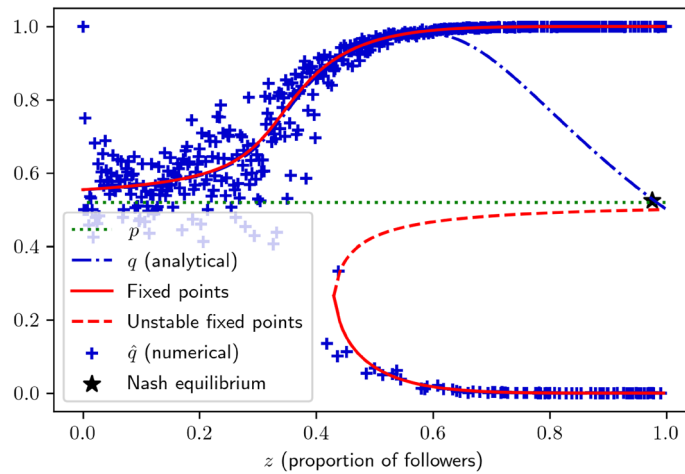
Figure 2. Impact of followers prevalence z on their accuracy q, for $m = 11$, $N = 2000$, and $p = .52$.

More intuitively, there is for $z > z_c$ a substantial probability (increasing with z) that, due to initial conditions, followers trap themselves in the wrong self-reinforcing forecast. When z becomes very large, these two fixed point are close to 0 and 1, the followers basically reach a consensus. This is a herding phenomenon: when imitation is too frequent, followers are mostly imitating one another and no longer aggregating genuine information. It is possible to compute analytically $q = \mathbb{E}[\hat{q}]$[11], which gives the bell-shaped blue dotted curve shown in Fig. 2.

Under the hood, this is due to the fact that the more numerous the followers, the less independent their opinions. As we mentioned in the Introduction, the super-accuracy of followers relies on independence: if all individuals agree, aggregating opinions has no effect. Worse still: as followers have no information of their own but rely on others to make up their mind, a population made of 100% followers would decide purely by chance, which means that q has to fall towards 1/2 as $z \rightarrow 1$. Furthermore, there is a point $z^{\dagger} \lesssim 1$ (i.e. the vast majority of agents are followers) where $q = p$. If the agents are rational accuracy-maximizers, this point is the only Nash equilibrium.

The message of the model is twofold. First, the equilibrium of the game is the point where, even though most people learn from others, learning offers no longer any gain of accuracy ($p = q$, reminiscent of Rogers' paradox[10]–see Sect. 3.1 in the Discussion for more details). In other words, imitation behaviour spreads until the snake bites his own tail. There is "consensus"—nearly all the population adopts either the right or the wrong forecast. This model can thus explain, e.g., herding effects among financial analysts[38].

This model shares the basic features and conclusions of the conformist transmission literature. In Fig. 1, the dotted-blue curve represents the probability that an imitator acquires the right opinion by using the F strategy, depending on its prevalence in the population. As we already noted, this shape matches Boyd & Richerson's definition[6] of having conformist bias (the probability to adopt the belief held by the majority is greater than the frequency of the majority opinion in the population).

An adaptive social network. *Assumptions.* One of the limitations of the Curty-Marsili framework is the assumption that followers sample randomly their consociates to make up their mind, thereby neglecting all possible effects of reputation, trust, and epistemic links between agents. In order to represent the fact that agents often rely on a network of people they know and trust, we embed the game in a dynamic directed graph, where each agent i is a node. A link $i \rightarrow j$ signifies that i listens to j when he/she makes up his/her mind. The network follows the following rules.

- Each agent i initializes its network Ω_i of nearest neighbours by picking uniformly m agents in the population.
- Each member $j \in \Omega_i$ is endowed with a personal score S_i^j, initialized at 0. It is updated following:

$$\forall j \in \Omega_i, \quad S_i^j(t + 1) = (1 - \gamma)S_i^j(t) + R^j(t), \tag{1}$$

  with $0 < \gamma \leq 1$ a memory decay rate, and a reward $R^j(t) = 1$ if j was right at time t, 0 otherwise. The score thus reflects the accuracy record of an agent, with an emphasis on the recent past. $\gamma = 0$ corresponds to infinite memory, whereas $\gamma = 1$ is such that only the latest reward matters.
- At each time step, i drops the lowest performer $j^\star = \operatorname{argmin}(S_i^j, j \in \Omega_i)$ with probability γ To reduce the number of parameters, we constrained the memory decay rate and the probability to update the network to have the same value γ. (which controls the speed of the network evolution), and picks a substitute. The probability for a "target" agent j^{target} to be chosen as a substitute is proportional to $a + k(j^{\text{target}})$, with $a > 0$ a small incompressible weight that avoids diversity depletion (when $a = 0$ an agent who has lost all its audience has no possibility to ever re-enter the network), and $k(j^{\text{target}})$ is j^{target}'s in-degree (measuring his/her "reputation"). This picking process excludes oneself and all agents already present in Ω_i.
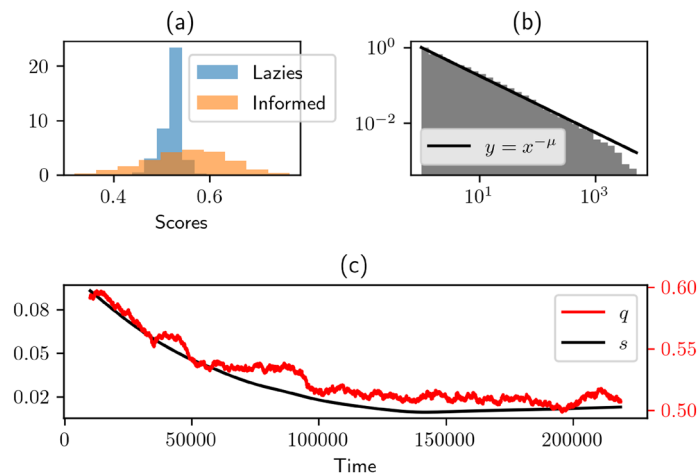
Figure 3. (**a**) Distribution of scores, (**b**) complementary CDF of in-degrees k, and (**c**) dynamics of the model, all for $z = .98$, $p = .52$. Panel (**b**) reveals that the distribution of in-degrees decays as a power-law $k^{-1-\mu}$ with $\mu \approx .75$. A value of $\mu \leq 1$ means that the first moment of k formally diverges, symptomatic of a "winner-takes-all" effect.

For the moment, we keep the strategies (I and F) fixed and exogenous, and observe the phenomenology of the model; we shall in a second stage let the strategies evolve depending on their respective payoffs. Henceforth, we choose z (the fraction of agents playing F) close to 1, and $p = 0.52$.

Collapse of the informed audience share. With this modification of the Curty-Marsili model we find that, perhaps surprisingly, q now drops below p, and settles around 1/2 (Fig. 3c).

This is caused by a subtle phenomenon. In the network, the scores S_i^t are supposed to estimate the accuracy of an agent. So, we should expect that when $q < p$, the followers tend to fill their network with I agents, who are on average more accurate. Paradoxically, the opposite happens. Even though the scores of the informed agents are on average higher than those of the followers, the latter have a much lower variance (see Fig. 3a). Indeed, as followers tend to make the same prediction, their scores fluctuate less than those of I agents who make independent predictions. As the network is updated by withdrawing worst performing agents it is often, counter-intuitively, an I agent who is jettisoned. Hence a collapse of the informed audience share, defined as the proportion of links directed towards an I agent.

Therefore, the systems ends up in a situation where the tiny minority of agents who have genuine information are not heeded. In our model, as often in real life, being the only one to be wrong has much stronger consequences than being the only one to be right. Being wrong means being wiped out the network, while being right does not provide any particular advantage in terms of audience share. The informed agents suffer most from this asymmetry. This echoes some empirical results: Yaniv & Kleinburger[24] found that it was *easier for advisors to lose a good reputation than to gain one*. This is also reminiscent of Kahneman & Tversky's prospect theory[39]: if average performance is the point of reference, then being the only one to be wrong is a loss, which is psychologically over-weighted compared with being the only one to be right. Anecdotal evidence suggests that this also true for asset managers, who get badly punished when they are alone in a drawdown.

Broad in-degree distribution. In our network, the out-degrees (the number of agents a given agent is listening to) are all fixed to m. The dynamics of in-degrees (the number of agents listening to someone) evolves dynamically. Quite interestingly, we find that the in-degrees distribution develops a power law tail, see Fig. 3b. This can be interpreted as the spontaneous emergence of "opinion leaders", i.e. agents whose opinion has a systemic impact on the population. This scale-free topology is due to the assumption that the probability to pick a given node is an affine function of its in-degree. Nodes with in-degree $k \gg a$ therefore grow exponentially with time. On the other hand, the probability that the score of an "opinion leader" remains by chance above a certain threshold that shields him/her from losing some followers decays exponentially with time. This "battle of exponentials" naturally leads to power-law distribution for k, with an exponent that depends on the parameters of the model (see e.g.[40]).

The value of originality. The results of the previous section are somewhat unsatisfying. In the situation where q is close to 1/2, any follower would benefit from listening to an informed instead of a follower. This is an assortment failure, and the situation described above is only an equilibrium because our assumptions prevent efficient assortment.

While it makes sense to assume that the strategies themselves (I or F) are not directly observable, opinions are public and may be used to infer the underlying strategies. But as we repeatedly pointed out, at equilibrium informed agents cannot be identified by their accuracy (in the Curty & Marsili's equilibrium, $p = q$).
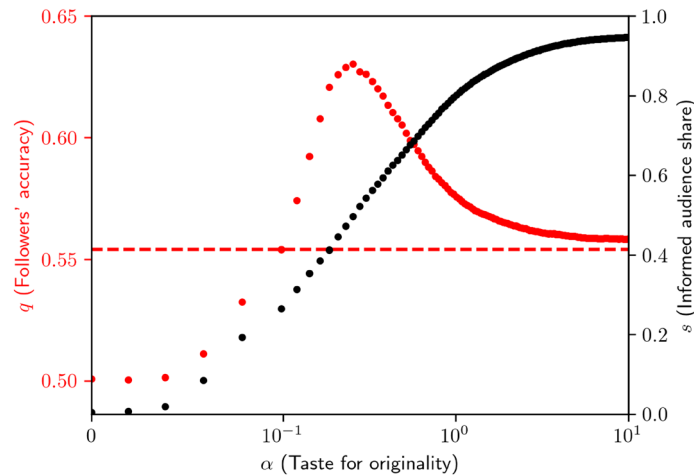
Figure 4. Impact of α on I audience share s and followers accuracy q (with the asymptotic value of q as a dashed line when $s \to 1$).

Instead, the most distinctive feature of I agents is their originality. Since their forecast is not influenced by anyone, the correlation between their opinion and the majority opinion is zero (or slightly above zero if the agent has a substantial influence in the network), whereas this quantity is close to one for followers. This distinctive feature can plausibly be observed by followers, and be used to identify potentially informed sources and listen preferentially to them. More mundanely, agents with original opinions are often deemed competent (they are said to "think outside the box", or hold "bold" opinions[35]).

This intuition can be formalized by introducing a proxy of originality, which we define to be the distance observed by agent i between the forecast of an agent j in Ω_i and the average forecast in Ω_i. Formally, the originality of an agent is $\left| R^j(t) - \overline{R}_i(t) \right|$, with $\overline{R}_i(t) := \frac{1}{m} \sum_{k=1}^{m} R^k(t)$. Such distance quantifies how deviant a given forecaster is, from the point of view of a follower. Including it in the scores update rule in the simplest way amounts to writing:

$$S_i^j(t+1) = (1 - \gamma) S_i^j(t) + R^j(t) + \alpha \left| R^j(t) - \overline{R}_i(t) \right| \tag{2}$$

where α is a weighting parameter, measuring by how much agents value *originality* over *accuracy*.

In information-theoretic terms, given the opinion of an original agent, the opinion of another randomly drawn unoriginal one has a null Shannon entropy. Conversely, the opinion of an original agent has a constant positive Shannon entropy, irrespective of the opinions that have been already voiced, since the random variable representing his/her opinion is independent from all other opinions. In other words, listening to an original agent is a way to diversify information sources, whereas listening to a consensus-follower is a waste of time. Hence, listening to original agents can actually be seen as a form of protection against group-think, which agents are arguably wary of. We start by fixing α as an exogenous parameter before letting it evolve endogenously to an equilibrium value.

Figure 4 reveals that the effect of the originality term in (2) is twofold. First, as α increases, predictably, the taste for originality increases the audience of informed agents in the network (red markers, left axis). This means that our modified score update rule is successful in implementing the idea of beneficial originality. Perhaps surprisingly though, the effect of α on q is *non-monotonic*. As α increases further, q reaches a maximum beyond which it decays.

In order to understand why this is the case, recall that in Fig. 2, z was found to have a bell-shaped effect on q. In the present framework, the informed audience share s in the network (black markers, right axis) somehow plays the role of z in Curty & Marsili's model. There needs to be an optimal fraction of informed agents in the network in order to maximize the wisdom of crowds. It is due to the fact that when $q > p$, followers are better forecasters than informed, leading to a mean-variance trade-off between diversifying information sources and listening to good forecasters.

This result is striking. Intuitively, social learners should preferentially seek to copy individual learners, who supposedly possess high quality information. Here, at some point, social learners are so good at detecting individual learners that listening to social learners instead becomes a better option. Analogously, reading a meta-analysis instead of primary literature is a good idea as long at the meta-analysis is filled with primary literature.

From the group point of view, the optimal situation should correspond to a "pyramidal" organisation with some individual learners at its base, originality-oriented social learners in the middle (who aggregate the information produced by the base) and pure social learners at the top. But if agents aimed to be right as often as possible, such a pyramid would be doubly unstable: the base would have an incentive to switch to social learning, and the middle would gain from listening to social learners instead. However, this organisation can be stable when agents do not have accuracy as only goal. This is what we explore in the following section.

Finally, increasing α concentrates the audience on the very small number of informed agents. Thus, we expect that some I agents have massive in-degrees, but numerical results are, again, counter-intuitive. The higher the α, the lower the maximum in-degrees (not shown). This makes sense if one accounts for the fact that informed draw their high audience share from their originality. Hence, as soon as an informed agent is in the network of a significant share of the population, it acquires a systemic impact which undermines its very originality. The effect of originality on scores therefore leads, self-consistently, to a limitation of the audience share of opinion leaders. Individual learners gets much more listened than average, without reaching the extremes allowed by the fat-tailed in-degree distribution obtained with $\alpha = 0$ (which gets gaussian as α increases).

Preliminary summary and open issues. Our exploratory results so far have shown that there exists an optimal positive value of the taste for originality α. One would however prefer that individuals endogenously develop such a feature. But this feature actually generates externalities: in a population with a large α (say, where followers listen only to informed), an agent who sets its personal α to 0 in order to listen mostly to followers would have a better accuracy than others. We might thus still be in a producer-scrounger problem, where free-riding would destroy the additional accuracy brought by the taste for originality.

In addition, we have so far worked with an exogenous proportion of informed agents. But if agents try to maximize their accuracy by changing strategies, this proportion should spontaneously decrease as soon as $q > p$. The foregone conclusion is that imitation produces no accuracy gain ($p = q$ is a necessary condition to have a Nash equilibrium) or, worse, that informed agents go extinct (which implies $q = 1/2$). This is precisely what happens in a naive evolutionary simulation where fitness depends only on accuracy.

Within our framework, the only way to produce a situation where $q > p$ is to modify the objective of our agents. Note that, in some contexts like financial markets, another realistic possibility is that informed agents can benefit from time priority. Take that the profit of an informed agent is given by $\beta p - c$ where $\beta > 1$ measures time priority and c are information costs, and that the profit of followers is simply q. Then one can have $q > p$ at equilibrium provided $(\beta - 1)p > c$. Concretely, if agents compete not only for accuracy, but also for in-degrees (i.e. for audience), then we might establish an efficient "division of labour". The lower accuracy of the informed would be compensated by their large in-degrees, due to the taste for originality that followers develop.

Thus, neither the taste for originality nor the survival of informed agents can be taken for granted at this point. We need to investigate these issues in a dynamical setting, i.e. by making both z and α endogenous. The next section introduces an evolutionary framework precisely to address this issue.

An evolutionary framework. At this point, analytical computations of an evolutionary equilibrium appear out of reach. We thus choose a dynamical mutation-selection algorithm, which allows us to find the equilibrium numerically, and a possible description of how the psychological traits we study evolved biologically. To do so, we need two ingredients: genes (the endogenous part of the model) and a fitness formula (maximized by evolution).

The genotype. We assume that agents have two genes:

  (i)   The strategy gene, which has two alleles: informed (I) and follower (F)
  (ii) The originality gene, with also two alleles: $\alpha = 0$ or $\alpha = \alpha^\star > 0$.

Note that henceforth the taste for originality is no longer continuous, but instead binary ($\alpha = 0$ or $\alpha^\star$), which helps for interpreting the results. Strictly speaking, what is endogenous here is not the value of α but rather the fraction of agents having $\alpha = \alpha^\star > 0$.

Three phenotypes can then be distinguished:

  (i)   Followers preferring originality, thereafter called "dandies". (We chose this word to signify that these agents also imitate, but do so in a way that make them deviant.)
  (ii) Followers without this taste, or "conformists"
  (iii) Informed (for them, the second gene makes no difference, as they do not use their social network)

The natural selection dynamics is implemented in a Wright-Fisher fashion[41]. At each time step, a uniformly drawn individual is killed, and replaced by a clone of some agent chosen with a probability proportional to fitness, defined below. At each time step, each gene has also a small probability σ to mutate to the alternative allele (typically, $\sigma = 10^{-8}$ in the simulations).

Costs and benefits. In addition to the obvious aim of agents, which is to make the right forecast as often as possible, we add two further sources of costs and benefits (otherwise, the informed inevitably go extinct and the model becomes trivial). First, information seeking has a cost. Second, we assume that being listened to has benefits. It can represent an analyst fee, university wage or simply the prestige or the power reaped from being influential.

Formally, we introduce the fitness of an individual $F_i(t)$, which is a weighted sum of three ingredients:

  (i)   A measure of an agent's forecast performance, $S_i(t)$, recorded with an exponential decay memory kernel (i.e. $S_i(t) = (1 - \gamma)S_i(t - 1) + R_i(t)$, with γ the memory decay rate and $R_i(t) = 1$ if i was right at time t and 0 otherwise.
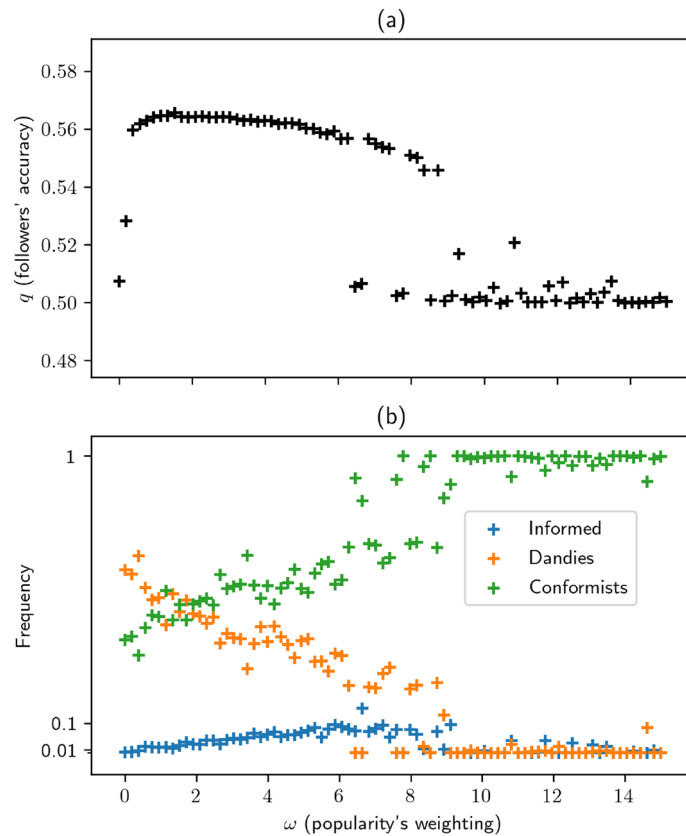
Figure 5. Effect of ω on accuracy q and phenotype prevalence, with p=.52, c=.01, $\alpha^{\star} = 1$.

(ii) The cost of information ($c \geqslant 0$), incurred by the informed

(iii) A measure of an agent's audience/prestige proxied by its in-degree $k_i := \#(j/i \in \Omega_j)$, weighted by a coefficient ω quantifying how rewarding it is to be listened to, relatively to being right.

The fitness formula, in its simplest linear form, writes:

$$F_i(t) = S_i(t) - c \cdot \mathbb{1}_{i \in I} + \omega \cdot k_i. \tag{3}$$

The simpler case where fitness depends only on accuracy corresponds to $c = 0$ and $\omega = 0$.

A mutually beneficial division of labour.   Figure 5a shows the average accuracy attained by followers as a function of ω. Quite remarkably, we find that q increases rapidly with ω for small ω, reaches a maximum and then decreases. For intermediate values of ω the system appears to be bi-stable, with two equilibria: $q_1 \approx 1/2$ and $q_2 > 1/2$, the latter becoming gradually unreachable as ω increases.

In order to explain this result, we plot the frequencies of types in the population (Fig. 5b). We see that the higher the ω, the rarer "dandies" are, whereas the effect on the fraction of informed agents is bell-shaped. The first observation comes from the phenomenon discussed in Sect. 2.2 and illustrated in Fig. 3a. Like the informed (but to a lesser extent), dandies decouple their opinion from the rest of the population. Their scores in the network are therefore more dispersed. Thus, exactly as for the informed, they tend to have fewer in-degrees than conformists. Hence, the higher the ω, the more natural selection will disfavour the dandy allele.

The bell-shaped effect on informed agents then follows naturally. We suggested in Sect. 2.4 that the survival of informed agents was contingent on two features: the reward brought by in-degrees (that compensates their lower accuracy) and the presence of dandies (who constitute the lion's share of the audience of informed agents). Here, these conditions translate into an ω value high enough for agents to have an incentive to be followers, but not too high—otherwise, for the reason stated in the previous paragraph, dandies go extinct, preventing informed agents from having an audience.

To sum up, we find three distinct regimes:

(i) Small ω: informed go extinct as they have no reason not to; they are less accurate than dandies, and accuracy is the main asset. Their extinction causes q to fall to 1/2, as no one in the network has any information any longer.

(ii) High ω: dandies cannot spread, as they lack in-degrees (which is here crucial). Their absence causes informed agents to go extinct, as the high reward for in-degrees does not favour them anymore.

(iii) Intermediate ω (from 1 to 5 in Fig. 5a): the three types can coexist, with increased accuracy at the global level ($q > p$).

The last regime is especially interesting and is that which we are looking for. We can describe it as a mutually beneficial "division of labour": informed agents do the hard work and collect reliable information, dandies listen to them to form more accurate opinions (meta-analysis), whereas conformists listen to a mix of conformists and dandies (so dandies achieve an indirect transmission of information between informed and conformists, otherwise impossible). The taste for originality is now stable, as the lower in-degrees of dandies is offset by their higher accuracy. The presence of an informed minority is also stable, as their lower accuracy is offset by their high in-degrees.

Opinion diversity. One can also analyse these results in terms of opinion diversity. Such diversity can be thought of as the average size of the minority group. One finds (not shown) that it follows very closely the prevalence of dandies. Thus, allowing for originality somehow solves the paradox we started from: the wisdom of crowds can work while maintaining opinion diversity.

In the labour division equilibrium, the population is neither a "smear of ideas"[8] nor homogenous. Instead, there is a clear majority, but also a substantial residual heterogeneity, which is endogenous and evolutionary stable.

So, how can we position the agents in terms of conformism? Similarly to Fig. 1, we can use our model (with $\omega = 1$ so we fall in the labour division equilibrium) to plot the probability for a random individual to acquire an opinion by imitation, as function of opinion's prevalence in the population—see Fig. 1. We see that individuals are generally "conformists" in the sense of Boyd & Richerson, unless the population is very close to full consensus, in which case the residual heterogeneity makes the curve saturate to a value < 1. Thus, the population is bound to stabilize in one of the two stable fixed points (around 0.1 and 0.9), which means some opinion diversity is stable. Interestingly, Fig. 1 is very close to the type of curve postulated by Schelling or Granovetter[42,43] in their models of social imitation—models that posit some inherent heterogeneities, for a review see[44].

Intuitively, agents are only cautiously influenced by the majority: they aggregate information while remaining wary of herding. Also, we can note that q (the followers' accuracy, not shown) is a strictly increasing function of opinion diversity, in line with Pages's "Diversity Prediction Theorem"[27]: *the amount by which the crowd outpredicts its average member increases as the crowd becomes more diverse*. As noted in[44,45] such a situation also prevents large opinion swings (or "crashes") when external conditions slowly evolve.

Note however that although the average value of q is larger than p, we are still in the bi-stable region of Fig. 2: in some instances, followers self-trap in the wrong belief. In our scenario, wisdom of crowd indeed exists but can sometimes badly fail. When neighbourhoods are not drawn at random in the population but according to static criteria (location, political opinion, social class, etc.) one should expect the formation of "echo chambers" with some clusters adopting the wrong belief, while others fully benefit from information aggregation.

What about anti-conformism?

Is originality per se a credible signal? The line of arguments developed above still has a loophole. If originality in itself is used by agents to detect the sources of genuine information, then original agents benefit from this signal by getting prestige in return. Thus, a social learner could benefit from being artificially original: he/she would be considered by others as a good source without paying the price of information search—just like in Batesian mimicry, where the palatable species gains protection from predators without paying the cost of being toxic[46]. Furthermore, originality is simple enough to produce: an agent can choose opinions randomly, or sample around and adopt the *minority* opinion. Therefore, the above scenario is possibly vulnerable to an invasion by parasitic behaviour mimicking originality, thus scrounging the prestige of genuinely informed agents.

It thus makes sense to include a fourth strategy: *anti-conformism*, consisting in sampling around, and endorsing the *minority* opinion. If such a behaviour spreads in the population, it challenges the very credibility of our scenario, as an original agent is no longer necessarily a source of reliable information. To do so, we include a third binary gene, whose possession of the positive allele coupled with the follower one leads to anti-conformism. Like the other genes, it evolves by mutation and selection. Now, if we re-run the simulations of Fig. 5a allowing for this new strategy (see orange curve in Fig. 6a), the $q > p$ phase disappears (in other words, the wisdom of crowds does not work anymore). The originality is at first used by dandies to hedge against herding, but is then exploited by anti-conformists, who become parasitic to the prestige of the informed.

The accuracy/prestige sweet-spot. At this point, one could conclude that our scenario producing $q > p$ is flawed, and that it can only be saved if there exists an authentication method allowing truly informed agents to stand out. The situation is in fact more complex, in an interesting way.

Actually, anti-conformism yields originality at the expense of accuracy: if the population is better than random ($q > 1/2$), then an anti-conformist mutant would be *worse than random* ($1 - q < 1/2$). More mundanely, an anti-conformist's opinion is most certainly original, but is often balderdash. Making wrong forecasts decreases fitness in two ways: directly (Eq. 3) and through prestige, as supplying good information is also key to obtaining prestige (recall the update rule 2).

The magnitude of this effect is controlled by the probability p that an informed agent makes the right forecast (which is, contrarily to q, exogenous). We can think of p as the easiness to forecast the topic in question. For instance, p would be higher in meteorology than in finance, as the predictive power is (nowadays) stronger in the
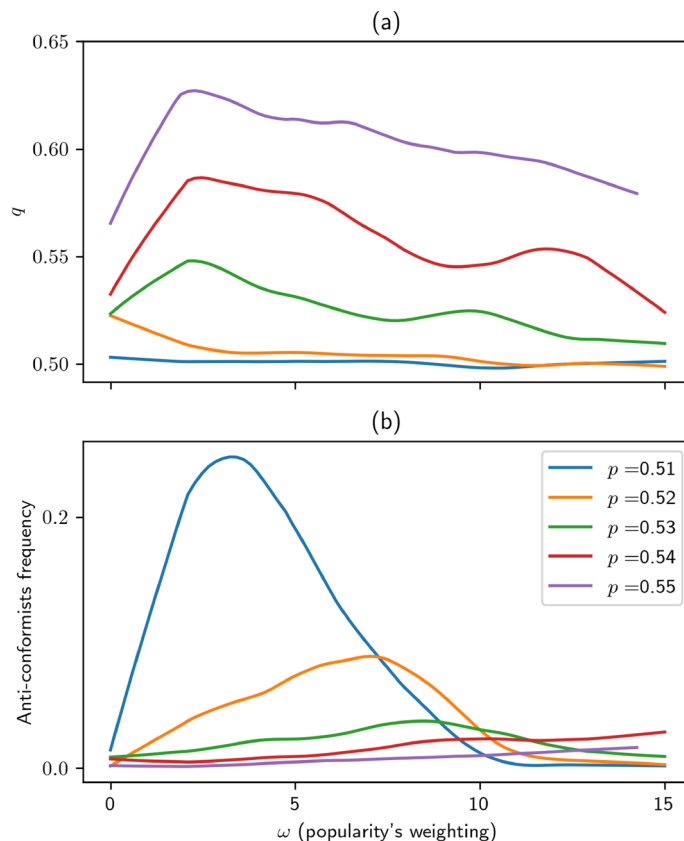
Figure 6. Effect of ω on F success rate q (**a**), and on anti-conformists prevalence (**b**), with $c = .01$ and various values of p. Lines are local regressions through data points (not shown for clarity).

former case. So, it would make sense that the lower the p, the more anti-conformism can spread in the population. Put differently, the more unpredictable a subject is, the harder it is to detect con artists. Thus, the virtuous equilibrium described in Sect. 2.5 is reachable when two conditions are met: (i) p is sufficiently high to detect counterfeits; (ii) ω takes intermediate values, as in the previous section.

This scenario is indeed confirmed numerically (cf. Fig. 6a): the $q > p$ phase reappears for higher values of p. As shown in Fig. 6b, anti-conformists only spread for low p and intermediate ω, in which case the division of labour fails.

In terms of opinion diversity (not shown), the phase where anti-conformists are present not only produces heterogeneity, but also *polarisation*: the average size of the minority group is close to 50%, i.e. the population is evenly split, as would indeed be expected for an anti-conformist bias (see Fig. 1): if the probability to follow the minority is greater than the size of such minority, then the only stable fixed point corresponds to 50% of the population holding one opinion, and 50% holding the other.

Discussion

Paradoxes of social transmission. We started this paper with an apparent paradox. Condorcet's theorem[1], together with empirical data[15,22], shows that aggregating others' opinions often allows one to form more accurate opinions than relying on individual learning. If this is the case, individuals should rely on "wisdom of crowds" heuristics, and there is indeed evidence that they do[34]. However, as agents form their opinions based on the views of others, the diversity of opinions in the group diminishes. This undermines a necessary condition ensuring the very effectiveness of the wisdom of crowds in the first place, namely the independent formation of aggregated opinions.

As a result, and as previous models suggest, we should end up in an apparently paradoxical equilibrium in which (i) the vast majority of agents adopt such a social learning strategy, leading to a large within-group homogeneity of opinions[7,11], even though (ii) this strategy, based on the wisdom of crowds effect, is not more effective than individual learning[11].

Our model offers a way to solve this apparent paradox and generate endogenous heterogeneity. First, because listening to diverse opinions increases the efficiency of information aggregation[27], it is beneficial for individuals to listen preferentially to *original* opinions. Second, by assuming that individuals benefit not only from being well informed, but also from being listened to, the preference of social learners for original opinions endogenously creates an incentive to be independent-minded.

Thus, as individuals have an extra-epistemic incentive to be listened to, some individual learning can endure despite the more efficient character of information aggregation. Importantly, the very persistence of individual learning guarantees the efficiency of social learning, as it introduces some diversity and independence in the opinions that the social learners gather—a necessary condition for the wisdom of crowds to work.

Rogers' paradox The paradox of the wisdom of crowd shares some similarities with Rogers' paradox[10], and can thus be relevant to its resolution as well. In Rogers' influential model, social learning is initially more efficient because it allows, in a varying environment, to avoid to cost of trial-and-error exploration. However, as social learners massively spread in the population, genuine information search get rarer. Social learning thus becomes less and less efficient, as the probability to copy an outdated belief or behaviour increases. Because of this negative prevalence dependence, social learning's fitness then collapses until being, at equilibrium, equal to the one of individual learning. Consequently, the mean fitness in such a group ends up being equal to what it would be in an "acultural" population, where everyone learns individually. In other words, culture (defined here as the ability to socially learn) would be, in the end, inevitable but useless. When compared to apparently massive contribution of social transmission to the ecological success of the human species[47,48], this result seemed somewhat paradoxical.

Several solutions were proposed to solve Rogers' paradox, mostly involving hybrid strategies. Boyd & Richerson's[49] proposed "selective learning", where individuals try to learn individually first, then learn socially if the first trial was inconclusive. It allows to imitate opportunistically, only when needed, and predicts a fitness gain for the group at equilibrium. Enquist et al.[50] proposed the reverse: "critical learning", consisting in learning socially first, then individually if the socially acquired strategy proved unsatisfactory.

Our model provides an alternative to solve the conundrum. Under the hood, Rogers' result comes from the prevalence-independence of the fitness of individual learning strategies. Since by definition strategies have identical fitness at equilibrium, the system is bound to have the same fitness as an asocial learners population. In Boyd & Richerson's model, this issue is bypassed by the assumption that the strategy consists in setting a threshold of self-confidence under which one learns socially, which means that the less one learns a-socially, the more one is demanding, hence the better its outcome. In our model, the positive prevalence-dependence of individual learners' payoff is quite natural: their opinions are sought for, so the rarer they are, the more they attract an audience.

It would be dubious to argue that our scenario is simpler than the existing ones. However, it has two merits. First, it is robust to a situation where a belief's quality can only be assessed *ex post* (as it is the case for forecasts, see *supra*). Second, and more importantly, it solves an extended version of the paradox: if we allow individuals to aggregate information, then social learning is not only less costly but also more accurate than individual learning, because of a "wisdom of crowds" effect. This setting precludes Boyd & Richerson's explanation, where individual learners are the ones who achieved the highest accuracy.

Predictions.
In addition to offering a solution to the above paradoxes, our model generates several predictions allowing one to explain some already known phenomena and possibly stimulate further empirical research.

Egocentric discounting and the under-use of social information. Social learning seems to contribute to our species ecological success, and many models of its evolution showed its great adaptive benefits[6,51]. However, somewhat unexpectedly, experimental evidence from social psychology, cultural evolution and experimental economics consistently shows that humans typically *under-use* social information, by placing an unjustified weight on their own opinions relative to others' (see[23], for a comprehensive review).

This sub-optimal under-reliance on social information—often called "egocentric discounting"[24]—is not easily explained within existing frameworks (see[23]). By contrast, it is a natural consequence of our model. As social learners need to diversify their pool of opinion, individual learners, although informationally worse off, get social benefits from being listened to, that out-weight the costs of belief inaccuracy. In other words, a sub-optimal reliance on social learning with respect to belief accuracy is compensated by extra-epistemic benefits brought by social influence.

In this perspective, egocentric discounting is explained by a mutually beneficial equilibrium, in which individual learners exchange, with social learners, original information for status. Such a division of labour equilibrium is consistent with the documented substantial inter-individual variation in the propensity to individually vs. socially learn[23,52,53].

The same model also predict that (i) social learning is assorted to a taste for originality rewarding the independent-minded individual learners, and that (ii) individual learners, in exchange, gain social status or reputation benefits from their independent-mindedness. The next sub-sections successively examine these two points.

Minority influence and the diversification of informational sources. A taste for originality, by which people would put more weight on independent-minded individuals' opinions, has, to our knowledge, never been tested—and could form the basis of future studies.

A general line of evidence on "minority influence"[54] might however be a manifestation of such a phenomenon. In opposition to what a pure conformist bias would predict[7], minorities have indeed frequently been found to have a disproportionate social influence relative to their size in the population[18,26,55]. Again, this kind of phenomena emerge straightforwardly from the model: the taste for originality makes agents holding frequently deviant opinions more influential, and, as a result, the influence of a rare opinion is higher than its share in the population (cf Fig. 1). In this perspective, the disproportionate social influence of minorities would emerge from social learners' disposition to diversify their informational source to increase social learning's efficiency—an adaptive form of bet-hedging.

Reputational benefits from information transmission.   Moreover, for the "division of labour" to be mutually beneficial, the informational influence of independent minded people has to translate in terms of fitness benefits.

Although there is to our knowledge no study directly investigating the link between informational originality and reputational benefits, several lines of reasoning converge to suggest its plausibility.

First, researchers long argued that prestige might be gained by providing others with social information[56,57]. This argument posits that, when some individuals have better-than-average skills in some domain (hunting, for instance), others individuals will pay deference to gain access to them in order to imitate the best available models. It is akin to the way informed agents get prestige in our model, by centralizing attention in the network. Note however that this analogy is somewhat misleading. Indeed, at variance with our model, the informed agents *are* indeed informationally better off in such a story. By contrast, in our virtuous equilibrium, informed agents are in fact the worst forecasters as they do not benefit from the Condorcet effect. But although their opinion is less accurate, it is (paradoxically) more informative about the state of the world, as it is not correlated to the opinion of the crowd. This extends Henrich & Gil-White's logic: when sampling several models, it can make sense to pay a deferential cost even to someone who is not particularly competent, as long as his/her opinion is original.

Second, some dimensions of communication suggest that it involves a form of reciprocal trade of information for status benefits. As argued by Dessalles, learning individually is costly and sharing the obtained information benefits its receiver, so doing it without compensation would amount to altruism[58]. Gaining extra social status could be the compensation that allows individual learning to endure, or as Dessalles puts it: *communication of information is part of a kind of unconscious trade in which status is the payment*. Furthermore, he remarks that people often have a hard time being heard and making their point instead of being asked for information, which would be expected for an altruistic act. Humans seem to be willing not only to share information, but to strive to make others hear their point of view and recognize their authorship of new ideas[59], which suggests that being heard is indeed profitable. Also in line with this idea, prosocial individuals tend to spend more effort on individual learning[21].

Relevance to efficient market theory.
Finally, as noted in the introduction, our story shares some similarities with, and offers some original insights on, the "impossibility of efficient markets" paradox discussed by Grossman & Stiglitz in their seminal work[3]. In particular, our distinction between information seekers (analysts), followers and anti-conformists is a relatively faithful description of the ecology of financial markets. (Note that some analysts may in fact themselves follow the consensus, see[38]).

Financial markets are notoriously difficult to predict, meaning, in our language, that the probability of success p of the informed is very close to 1/2. As we have pointed out in Sect. 2.6, this is a situation where "quacks" easily proliferate, since it is hard to distinguish them from genuine experts. In this case, the aggregation mechanism by which financial markets are supposed to reveal the "true" fundamental price of firms may completely breakdown. Furthermore, a prevalence of followers that trade chronologically *after* informed traders can give rise to bubbles that further decouple market prices from ground truth.

Hence we are less optimistic than Grossman & Stiglitz who propose that financial markets are close to efficiency, with a small residual error allowing information seekers to cover their costs while allowing the aggregation process to perform its task. We rather believe in Black's picture, where market prices are *within a factor of 2* of fundamental prices[60,61]: only when mispricing becomes large enough will p differ sufficiently from 1/2, allowing true information to reveal itself and reinstate the power of Condorcet's aggregation theorem. In fact, since p becomes itself endogenous in this case, financial markets are presumably prone to predictivity cycles: p close to 1/2 disfavours informed trading, which generates mispricing and, in turn, larger values of p, favouring informed trading, etc. For some empirical signatures of such cycles, see[62,63].

Limitations.
We should nevertheless highlight a number of limitations of our model, which might help to delineate its domain of validity. For example, our model treats beliefs as perishable instead of cumulative: agents start searching from scratch at each iteration instead of taking their previous opinion as a starting point, and search to improve it. Also, the model implicitly assumes that agents can observe the opinions of their neighbours, but can only assess directly the *quality* of these opinions *ex post*, when it is too late to change opinion. In other words, agents can measure the past accuracy of an information source, but know nothing else of the current advice he is proposing. Thus, our model was not designed to study a cultural item like canoe building techniques, where the quality of a technique can be inferred through the quality of the produced canoes. This allows more elaborate strategies, such as the *critical learning*[50] we mentioned, or a *prestige bias*[56] (agents learning the best technique by imitating the most successful individual in the group). As these were not central to the contribution of the model, we chose not to include competence heterogeneity and prestige-biased transmission, although these are clearly major forces in cultural evolution. Our model seems at first sight more suited to analyse beliefs such as *forecasts*. Indeed, it is impossible to compare the qualities of different forecasts before the event realization, and too late to change them afterward, which makes the two aforementioned solutions unworkable. We see however no fundamental reason why our model would not apply to technological cultural items: individual learning (in this case, *innovations*) is a *sine qua non* condition for the cumulative process to take place, and should at first sight be counter-selected if the aggregate opinion is of high quality. The endogenous generation of cultural diversity through individual learning may also have interesting implications regarding the dynamics of cumulative technological evolution. Researchers have indeed suggested that increasing the independence of individuals' exploration strategies can decrease the probability of populations getting "stuck" in local optima, by making them explore different cultural possibilities[64].

Other mechanisms could also be invoked to justify the persistence of some individual learning (and the resulting preservation of some opinion diversity), despite social learning's efficiency. A possibility would be for more

competent or knowledgeable individuals to discount the opinion of others in forming their own (see, e.g.,[65]). In line with this idea, studies suggests that IQ and self-confidence negatively predict reliance on social learning[51,65]. This idea, however, assumes that it is because competent individuals are *informationally better off* that they rely on individual learning (because they are "wiser" than the crowd). This, however, seems not to be the case. Throughout the experimental literature in social psychology, cultural evolution and experimental economics, individual learners' independent mindedness appears as informationally sub-optimal: they are found to *under-use* social information relative to what an accuracy-maximization imperative would warrant[23]. Furthermore, Mesoudi[22] found that the subjects who used social learning the most were also above-average individual learners.

Finally, our model neglects possible conflicts of interest or strategic interaction: agents ignore the influence they have on others. In reality, information transmission has no reason to always be benevolent, but can lead to manipulation, against which individual learning is a way to hedge oneself.

Code availibility

The Python code of the model is available here.

References

1. de Condorcet, M. *Essay on the application of analysis to the probability of majority decisions* (Imprimerie Royale, 1785).
2. Galton, F. Vox populi. *Nature* **75**, 450–451. https://doi.org/10.1038/075450a0 (1907).
3. Grossman, S. J. & Stiglitz, J. E. On the impossibility of informationally efficient markets. *Am. Econ. Rev.* **70**, 393–408 (1980).
4. Kendal, R. L. *et al.* Social learning strategies: bridge-building between fields. *Trends Cognit. Sci.* **22**, 651–665 (2018).
5. Osiurak, F. & Reynaud, E. The elephant in the room: What matters cognitively in cumulative technological culture. *Behav. Brain Sci.* **43**, (2020).
6. Boyd, R. & Richerson, P. J. *Culture and the Evolutionary Process* (Culture and the Evolutionary Process, University of Chicago Press, 1985).
7. Henrich, J. & Boyd, R. The evolution of conformist transmission and the emergence of between-group differences. *Evol. Hum. Behav.* **19**, 215–241. https://doi.org/10.1016/S1090-5138(98)00018-X (1998).
8. Henrich, J. Cultural group selection, coevolutionary processes and large-scale cooperation. *J. Econ. Behav. Organ.* **53**, 3–35. https://doi.org/10.1016/S0167-2681(03)00094-5 (2004).
9. Lorenz, J., Rauhut, H., Schweitzer, F. & Helbing, D. How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci.* **108**, 9020–9025. https://doi.org/10.1073/pnas.1008636108 (2011).
10. Rogers, A. R. Does biology constrain culture. *Am. Anthropol.* **90**, 819–831. https://doi.org/10.1525/aa.1988.90.4.02a00030 (1988).
11. Curty, P. & Marsili, M. Phase coexistence in a forecasting game. *J. Stat. Mech. Theory Exp.* **2006**, P03013–P03013. https://doi.org/10.1088/1742-5468/2006/03/P03013 (2006).
12. Surowiecki, J. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations* (Doubleday, 2004).
13. Senninger, L. K. *Wisdom of the Crowd in Experiments.* (BestMasters Gabler Verlag, 2018).
14. Kurvers, R. H. J. M. *et al.* Boosting medical diagnostics by pooling independent judgments. *Proc. Natl. Acad. Sci.* **113**, 8777–8782. https://doi.org/10.1073/pnas.1601827113 (2016).
15. Novaes, T. A., Wolf, M., Krause, J. & Kurvers, R. H. J. M. Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information. *J. R. Soc. Interface* **15**, https://doi.org/10.1098/rsif.2018.0155 (2018).
16. Wolf, M., Kurvers, R. H. J. M., Ward, A. J. W., Krause, S. & Krause, J. Accurate decisions in an uncertain world: collective cognition increases true positives while decreasing false positives. *Proc. R. Soc. B Biol. Sci.* **280**, 20122777. https://doi.org/10.1098/rspb.2012.2777 (2013).
17. Clément, R. J. G., Wolf, M., Snijders, L., Krause, J. & Kurvers, R. H. J. M. Information transmission via movement behaviour improves decision accuracy in human groups. *Anim. Behav.* **105**, 85–93. https://doi.org/10.1016/j.anbehav.2015.04.004 (2015).
18. Eriksson, K. & Coultas, J. Are people really conformist-biased? An empirical test and a new mathematical model. *J. Evol. Psychol.* **7**, 5–21. https://doi.org/10.1556/JEP.7.2009.1.3 (2009).
19. Bell, A. V., Richerson, P. J. & McElreath, R. Culture rather than genes provides greater scope for the evolution of large-scale human prosociality. *Proc. Natl. Acad. Sci.* **106**, 17671–17674. https://doi.org/10.1073/pnas.0903232106 (2009).
20. Efferson, C., Lalive, R., Richerson, P. J., McElreath, R. & Lubell, M. Conformists and mavericks: the empirics of frequency-dependent cultural transmission. *Evol. Hum. Behav.* **29**, 56–64. https://doi.org/10.1016/j.evolhumbehav.2007.08.003 (2008).
21. Eriksson, K. & Strimling, P. Biases for acquiring information individually rather than socially. *J. Evol. Psychol.* **7**, 309–329. https://doi.org/10.1556/jep.7.2009.4.4 (2009).
22. Mesoudi, A. An experimental comparison of human social learning strategies: Payoff-biased social learning is adaptive but under-used. *Evol. Hum. Behav.* **32**, 334–342. https://doi.org/10.1016/j.evolhumbehav.2010.12.001 (2011).
23. Morin, O., Jacquet, P. O., Vaesen, K. & Acerbi, A. Social information use and social information waste. *SocArXiv* **10.31235/osf.io/rqcdf**, (2020).
24. Yaniv, I. & Kleinberger, E. Advice taking in decision making: egocentric discounting and reputation formation. *Organ. Behav. Hum. Decis. Process.* **83**, 260–281. https://doi.org/10.1006/obhd.2000.2909 (2000).
25. Miu, E. & Morgan, T. J. H. Cultural adaptation is maximised when intelligent individuals rarely think for themselves. *Evol. Hum. Sci.* **2**, https://doi.org/10.1017/ehs.2020.42 ((2020/ed).).
26. Latané, B. & Wolf, S. The social impact of majorities and minorities. *Psychol. Rev.* **88**, 438–453. https://doi.org/10.1037/0033-295X.88.5.438 (1981).
27. Page, S. E. *The Difference–How the Power of Diversity Creates Better Groups, Firms, Schools and Societies* (Princeton University Press, 2008).
28. Dyne, L. V. & Saavedra, R. A naturalistic minority influence experiment: effects on divergent thinking, conflict and originality in work-groups. *Br. J. Soc. Psychol.* **35**, 151–167. https://doi.org/10.1111/j.2044-8309.1996.tb01089.x (1996).
29. Jetten, J. & Hornsey, M. J. Deviance and dissent in groups. *Annu. Rev. Psychol.* **65**, 461–485. https://doi.org/10.1146/annurev-psych-010213-115151 (2014).
30. Nemeth, C. & Chiles, C. Modelling courage: the role of dissent in fostering independence. *Eur. J. Soc. Psychol.* **18**, 275–280. https://doi.org/10.1002/ejsp.2420180306 (1988).
31. Bloomfield, R. & Hales, J. An experimental investigation of the positive and negative effects of mutual observation. *Account. Rev.* **84**, 331–354. https://doi.org/10.2308/accr.2009.84.2.331 (2009).

32. Hess, N. H. & Hagen, E. H. Psychological adaptations for assessing gossip veracity. *Hum. Nat. (Hawthorne, N.Y.)* **17**, 337–354, https://doi.org/10.1007/s12110-006-1013-z (2006).
33. Mercier, H. & Miton, H. Utilizing simple cues to informational dependency. *Evol. Hum. Behav.* **40**, 301–314. https://doi.org/10.1016/j.evolhumbehav.2019.01.001 (2019).
34. Mercier, H. & Morin, O. Majority rules: How good are we at aggregating convergent opinions?. *Evol. Hum. Behav.* **1**, https://doi.org/10.1017/ehs.2019.6 (2019).
35. Hong, H., Kubik, J. D. & Solomon, A. Security analysts' career concerns and herding of earnings forecasts. *RAND J. Econ.* **31**, 121–144 (2000).
36. Bardoscia, M., De Luca, G., Livan, G., Marsili, M. & Tessone, C. J. The social climbing game. *J. Stat. Phys.J. Stat. Phys.* **151**, 440–457 (2013).
37. von Rueden, C., Gurven, M. & Kaplan, H. Why do men seek status? fitness payoffs to dominance and prestige. *Proc. R. Soc. B Biol. Sci.* **278**, 2223–2232. https://doi.org/10.1098/rspb.2010.2145 (2011).
38. Guedj, O. & Bouchaud, J.-P. Experts' earning forecasts: Bias, herding and gossamer information (2004). arXiv:cond-mat/0410079.
39. Kahneman, D. & Tversky, A. Prospect Theory: An Analysis of Decision Under Risk. In *Handbook of the Fundamentals of Financial Decision Making*, vol. Volume 4 of *World Scientific Handbook in Financial Economics Series*, 99–127, https://doi.org/10.1142/9789814417358_0006 (WORLD SCIENTIFIC, 2012).
40. Bouchaud, J. P. More lévy distributions in physics. In Shlesinger, M. F., Zaslavsky, G. M. & Frisch, U. (eds.) *Lévy Flights and Related Topics in Physics*, 237–250 (Springer, 1995).
41. Imhof, L. A. & Nowak, M. A. Evolutionary game dynamics in a wright-fisher process. *J. Math. Biol.* **52**, 667–681 (2006).
42. Schelling, T. C. Dynamic models of segregation. *J. Math. Biol.* **1**, 143–186 (1971).
43. Granovetter, M. Threshold models of collective behavior. *Am. J. Sociol.Am. J. Sociol.* **83**, 1420–1443. https://doi.org/10.1086/226707 (1978).
44. Bouchaud, J.-P. Crises and collective socio-economic phenomena: Simple models and challenges. *J. Stat. Phys.* **151**, 567–606, https://doi.org/10.1007/s10955-012-0687-3 (2013).
45. Michard, Q. & Bouchaud, J.-P. Theory of collective opinion shifts: from smooth trends to abrupt swings. *Eur. Phys. J. B Condens. Matter Complex Syst.* **47**, 151–159 (2005).
46. Wickler, W. Mimicry and the evolution of animal communication. *Nature* **208**, 519–521. https://doi.org/10.1038/208519a0 (1965).
47. Henrich, J. *The Secret of Our Success–How Culture is Driving Human Evolution, Domesticating Our Species, and Making us Smarter* (Princeton University Press, 2015).
48. Boyd, R., Richerson, P. J. & Henrich, J. The cultural niche: why social learning is essential for human adaptation. *Proc. Natl. Acad. Sci.* **108**, 10918–10925. https://doi.org/10.1073/pnas.1100290108 (2011).
49. Boyd, R. & Richerson, P. J. Why does culture increase human adaptability?. *Ethol. Sociobiol.* **16**, 125–143. https://doi.org/10.1016/0162-3095(94)00073-G (1995).
50. Enquist, M., Eriksson, K. & Ghirlanda, S. Critical social learning: a solution to Rogers's paradox of nonadaptive culture. *Am. Anthropol.* **109**, 727–734. https://doi.org/10.1525/aa.2007.109.4.727 (2007).
51. Muthukrishna, M., Morgan, T. J. & Henrich, J. The when and who of social learning and conformist transmission. *Evol. Hum. Behav.* **37**, 10–20 (2016).
52. Olsen, K., Roepstorff, A. & Bang, D. Knowing whom to learn from: individual differences in metacognition and weighting of social information. *PsyArXiv* (2019).
53. Toelch, U., Bruce, M. J., Newson, L., Richerson, P. J. & Reader, S. M. Individual consistency and flexibility in human social information use. *Proc. R. Soc. B Biol. Sci.* **281**, 20132864 (2014).
54. Moscovici, S. *Innovation and minority influence*, 9–52. European Studies in Social Psychology (Cambridge University Press, 1985).
55. Tanford, S. & Penrod, S. Social influence model: a formal integration of research on majority and minority influence processes. *Psychol. Bull.* **95**, 189 (1984).
56. Henrich, J. & Gil-White, F. J. The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evol. Hum. Behav.* **22**, 165–196. https://doi.org/10.1016/S1090-5138(00)00071-4 (2001).
57. Jiménez, Á. V. & Mesoudi, A. Prestige-biased social learning: current evidence and outstanding questions. *Palgrave Commun.* **5**, 1–12 (2019).
58. Dessalles, J. -L. *Altruism, Status, and the Origin of Relevance*. **20** (Cambridge University Press, 1998).
59. Altay, S., Majima, Y. & Mercier, H. It's my idea! reputation management and idea appropriation. *Evol. Hum. Behav.* **41**, 235–243 (2020).
60. Black, F. Noise. *J. Finance* **41**, 528–543. https://doi.org/10.1111/j.1540-6261.1986.tb04513.x (1986).
61. Bouchaud, J.-P. *et al.* Black was right: Price is within a factor 2 of value. *Risk.net* (2018).
62. Schmitt, N. & Westerhoff, F. On the bimodality of the distribution of the s&p 500's distortion: empirical evidence and theoretical explanations. *J. Econ. Dyn. Control* **80**, 34–53. https://doi.org/10.1016/j.jedc.2017.05.002 (2017).
63. Majewski, A. A., Ciliberti, S. & Bouchaud, J.-P. Co-existence of trend and value in financial markets: estimating an extended chiarella model. *J. Econ. Dyn. Control* **112**, 103791. https://doi.org/10.1016/j.jedc.2019.103791 (2020).
64. Derex, M. & Boyd, R. Partial connectivity increases cultural accumulation within groups. *Proc. Natl. Acad. Sci.Proc. Natl. Acad. Sci.* **113**, 2982–2987 (2016).
65. Hawthorne-Madell, D. & Goodman, N. D. Reasoning about social sources to learn from actions and outcomes. *Decision* **6**, 17 (2019).

Acknowledgements

Author contributions

M.B., J.P.B. and B.C. designed research; B.C. performed research, L.F. and B.C. contextualised the model in the evolutionary social sciences literature, all the authors wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.B.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.