

Social learning strategies regulate the wisdom and madness of interactive crowds

Wataru Toyokawa ^{1,2,3*}, Andrew Whalen¹ and Kevin N. Laland¹

Why groups of individuals sometimes exhibit collective ‘wisdom’ and other times maladaptive ‘herding’ is an enduring conundrum. Here we show that this apparent conflict is regulated by the social learning strategies deployed. We examined the patterns of human social learning through an interactive online experiment with 699 participants, varying both task uncertainty and group size, then used hierarchical Bayesian model fitting to identify the individual learning strategies exhibited by participants. Challenging tasks elicit greater conformity among individuals, with rates of copying increasing with group size, leading to high probabilities of herding among large groups confronted with uncertainty. Conversely, the reduced social learning of small groups, and the greater probability that social information would be accurate for less-challenging tasks, generated ‘wisdom of the crowd’ effects in other circumstances. Our model-based approach provides evidence that the likelihood of collective intelligence versus herding can be predicted, resolving a long-standing puzzle in the literature.

Understanding the mechanisms that account for accurate collective decision-making among groups of animals—‘collective intelligence’—has been a central focus of animal behaviour research^{1–5}. There are a large number of biological examples showing that collectives of poorly informed individuals can achieve high performance in solving cognitive problems under uncertainty^{6–10}. Although these findings suggest fundamental cognitive benefits of grouping¹¹, there is also a long-standing recognition, especially for humans, that interacting individuals may sometimes be overwhelmed by the ‘extraordinary popular delusions and madness of crowds’¹². Herd behaviour (that is, an alignment of thoughts or behaviours of individuals in a group) occurs because individuals imitate each other^{13–15}, even if each is a rational decision-maker¹⁶. Imitation is thought to be a cause of financial bubbles^{12,17}, ‘groupthink’¹⁸ and volatility in cultural markets^{19,20}. More generally, interdependence between individual decisions may undermine the wisdom of the crowds effect²¹ (but see ref. 22), while potential disadvantages of information transfer are well-recognized in the biological literature^{23,24}. It seems that information transmission among individuals, and making decisions collectively, is a double-edged sword: combining decisions may provide the benefits of collective intelligence, but at the same time increase the risk of an informational cascade¹⁶. Collectively, an understanding of whether, and if so how, it is possible to prevent or reduce the risk of maladaptive herding while concurrently keeping or enhancing collective intelligence is largely lacking.

A balance between using individual and social information may play a key role in determining the trade-off between collective wisdom and ‘madness’²⁵. If individuals are too reliant on copying others’ behaviour, any idea—even a maladaptive one—can propagate in the social group through positive feedbacks^{2,26}. For instance, disproportionately strong positive responses to recruitment signals in social insects have been shown to trap the whole colony to exploit a suboptimal, outdated resource^{24,27}. Likewise, conformity-biased transmission in humans and other animals can potentially lead groups to converge on a maladaptive behaviour^{16,23,28,29}. In contrast, however, if individuals completely ignore social information so as to

be independent, they will fail to exploit the benefits of aggregating information through social interactions. The extent to which individuals should use social information should fall between these two extremes. Evolutionary models predict that the balance between independence and interdependence in collective decision-making may be changeable and contingent on the individual-level flexibility and inter-individual variability associated with the social learning strategies deployed in diverse environmental states^{28,30,31}.

Experimental studies report that animals (including humans) increase their use of social information as the returns from asocial learning become more unreliable^{32–37}, while theory and data suggest that the benefits to individuals of social learning increase with group size^{34,38–42}. Selectivity in the predicted use of social information may impact on collective decision-making because slight differences in the parameter values of social information use are known to be able to alter qualitatively the collective behavioural dynamics^{1,2,5,43,44}. Therefore, researchers should expect populations to exhibit a higher risk of being trapped with maladaptive behaviour with increasing group size and decreasing reliability of asocial learning (and concomitant increased reliance on social learning).

From the viewpoint of the classic wisdom of the crowds theory, increasing group size may increase collective accuracy^{45–48}. The relative advantage of collective over solitary individuals may also be highlighted by increased task difficulty, because there would be more room for the performance of difficult tasks to be improved compared with easier tasks in which high accuracy can be achieved by asocial learning only. To understand the collective decision performance of social learners fully requires fine-grained quantitative studies of social learning strategies and their relations to collective dynamics, linked to sophisticated computational analysis.

The aims of this study were twofold. First, we set out to test the hypothesis that the circumstances under which collective decision-making will generate ‘wisdom’ can be predicted with knowledge of the precise learning strategies individuals deploy, through a combination of experimentation and theoretical modelling. The choice of an abstract decision-making task allowed us to implement a computational modelling approach, which has been increasingly deployed

¹School of Biology, University of St Andrews, St Andrews, UK. ²Japan Society for the Promotion of Science, Tokyo, Japan. ³Department of Evolutionary Studies of Biosystems, School of Advanced Sciences, The Graduate University for Advanced Studies, Hayama, Japan. *e-mail: wf25@st-andrews.ac.uk

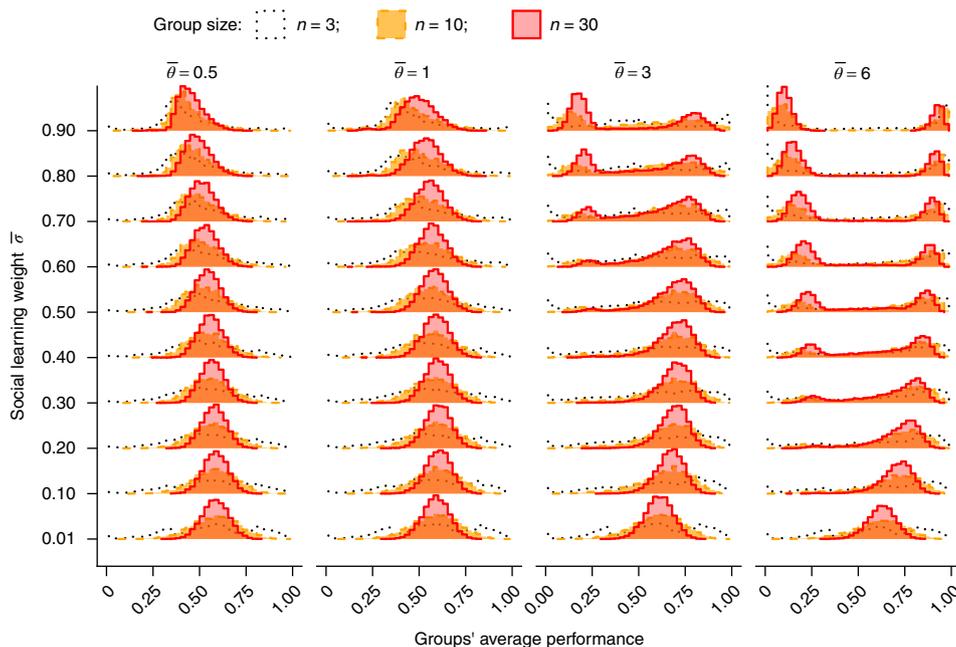


Fig. 2 | Results from the individual-based model simulations showing the distribution of each group's mean accuracy before environmental change ($t \leq 40$). The x axis gives the mean decision accuracy over the first 40 rounds (that is, environment 1) for each replication. Black dotted, orange dashed and red solid lines represent group sized of $n=3$, 10 and 30, respectively. The other free parameter values are the same as in Fig. 1.

relatively weak in smaller groups (see dotted line in Fig. 1d). This is because the majority of individuals in smaller groups (that is two individuals out of three) are more likely to break the cultural inertia by simultaneously exploring another option by chance than are the majority in larger groups (for example, six out of ten).

In summary, the model simulation suggests an interaction between the social learning weight $\bar{\omega}$ and conformity exponent $\bar{\theta}$ on decision accuracy and the risk of inflexible herding. When the conformity exponent is not too large, increasing the group size can increase the decision accuracy while concurrently retaining the decision flexibility across a broad range of mean social learning weights. However, when the conformity bias becomes large, the risk of inflexible herding arises, and when both social learning parameters are large, collective intelligence is rare and inflexible herd behaviour dominates.

Collective performance of human participants. Figure 3a shows the behavioural dynamics of human participants in different group sizes and different task uncertainty conditions (see Supplementary Fig. 3 for each group's behaviour). The average decision performance of collectives (that is, group size ≥ 2) exceeded that of solitary individuals (that is, group size = 1) in the moderate-uncertainty condition (that is, the 95% Bayesian credible interval (CI) of ξ_t exceeds 0 at regions $t \in 9-40$ and $67-70$; Fig. 3b). In other uncertainty conditions, no global positive effect of grouping was observed, suggesting that collective intelligence was prominent only in the moderate-uncertainty condition. However, the main effect of group size was positive in the post-change period of the low-uncertainty condition (mean and 95% Bayesian CI of $\omega_2 = 0.67$ (0.44 to 0.91); Table 1), suggesting that the average performances of large groups (for example, $12 \leq \text{group size} \leq 16$) were better, and hence more flexible, than smaller groups and solitaries (Fig. 3a). In contrast, in the moderate-uncertainty condition, the average performance of the collectives dropped below that of the solitaries after the environmental change (that is, $\xi_t < 0$ at a region $t \in 42-45$; Fig. 3b). Also, the main effect of group size was negative in the post-change period (mean and 95% Bayesian CI of $\omega_2 = -0.26$ (-0.44 to -0.11);

Table 1), suggesting that larger groups were more likely to get stuck in the outdated option in the moderate-uncertainty condition. In the high-uncertainty condition, the main effect of group size was positive in the pre-change period and negative post-change (mean and 95% Bayesian CIs: $\omega_1 = 0.07$ (0.00 to 0.15); $\omega_2 = -0.10$ (-0.17 to -0.02); Table 1), although the effect size was too small to differentiate between the performances of different group sizes visually (Fig. 3a). Using monetary earnings as an outcome variable of decision performance did not change our conclusions qualitatively (supporting Supplementary Fig. 4 and Supplementary Table 2).

Our phenomenological model regression established that manipulating both task uncertainty and group size indeed affected the collective decision dynamics. Below, we address whether or not the pattern could be explained with knowledge of human social learning strategies estimated through our learning-and-decision-making computational model.

Estimation of human social information use. Using posterior estimation values obtained by the hierarchical Bayesian model-fitting method (Table 2), we were able to categorize the participants as deploying one of three different learning strategies based on their fitted conformity exponent values; namely, the 'positive frequency-dependent copying' strategy ($\theta_i \gg 0$), 'negative frequency-dependent copying' strategy ($\theta_i \ll 0$) and 'random choice' strategy ($\theta_i \sim 0$). Note that we could not reliably detect the 'weak positive' frequency-dependent strategy ($0 < \theta_i \leq 1$) due to the limitation of statistical power (Supplementary Fig. 5). Some individuals whose 'true' conformity exponent fell between zero and one would have been categorized as exhibiting a random choice strategy (Supplementary Fig. 7). Individuals identified as exhibiting the positive frequency-dependent copying strategy were mainly those whose conformity exponent was larger than 1 ($\theta_i > 1$).

Figure 4a shows the estimated frequencies of different learning strategies. Generally speaking, participants were more likely to utilize a positive frequency-dependent copying strategy than the other two strategies (the 95% Bayesian CI of the intercept of the generalized linear mixed model (GLMM) predicting the probability to use

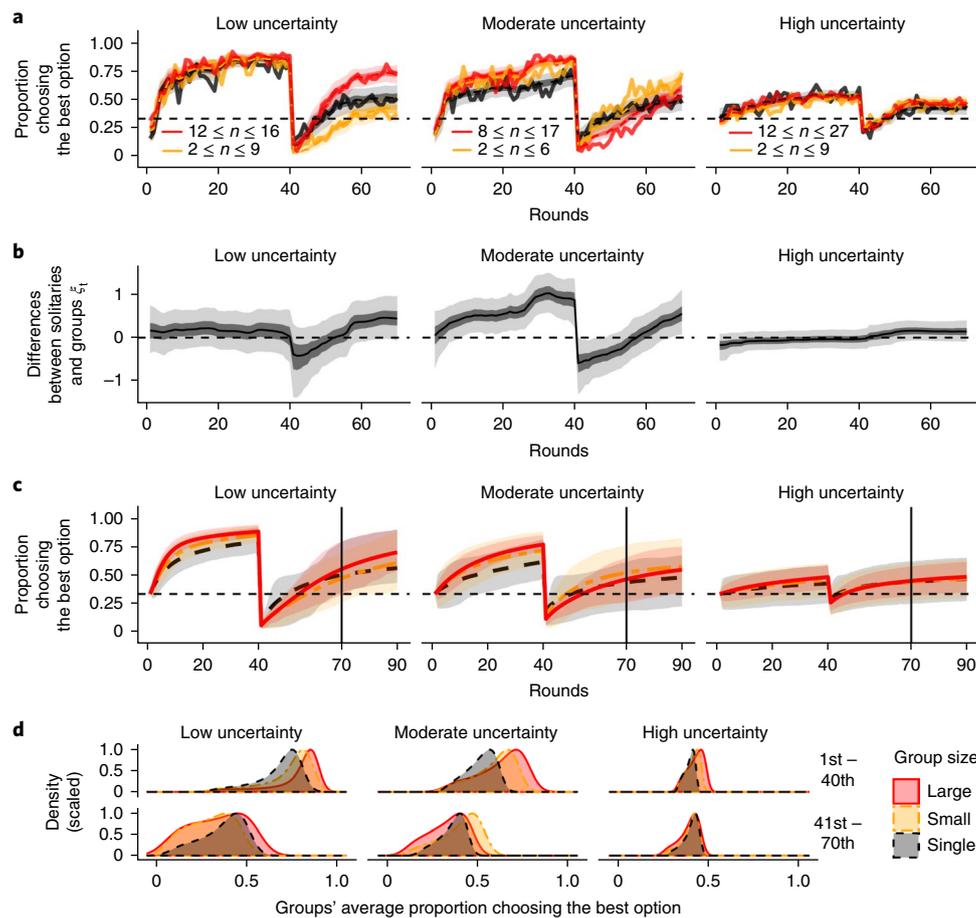


Fig. 3 | Time evolutions and distributions of decision performance for each condition. **a**, Average decision accuracies of experimental participants (red, large groups; orange, small groups; dark grey, lone individuals). All individual performances were averaged within the same size category (solid lines). The light-shaded areas, dark-shaded areas and dashed curves show the 95%, 50% and median Bayesian CIs of the phenomenological, time-series logistic regression. Sample sizes for large, small and lone groups are: $n = 43, 44$ and 38 for the low-uncertainty condition; $n = 52, 56$ and 37 for the moderate-uncertainty condition; and $n = 259, 168$ and 58 for the high-uncertainty condition, respectively. **b**, Change in the main effect of the dummy variable of grouping on the decision accuracy at the phenomenological regression model. The light- and dark-shaded areas are the 95 and 50% Bayesian CIs and the solid curves are the medians. **c,d**, Change and distribution of average decision accuracy of the individual-based post hoc simulations of the learning process model using the experimentally fit parameter values. In **c**, all replications were averaged within the same size category (solid lines); the shaded areas give the 50% quantiles; and the experimental horizon (that is, $t = 70$) is indicated by the vertical line. In **d**, the performance was averaged within pre- (top, 1st–40th) and post-change periods (bottom, 41st–70th) for each replication for each group-size category.

Table 1 | Mean and 95% Bayesian CIs of the posterior for the group size effect in the phenomenological logistic model

	Low uncertainty		Moderate uncertainty		High uncertainty	
ω_1	0.08	−0.15 to 0.33	0.10	−0.06 to 0.26	0.07	0.00 to 0.15
ω_2	0.67	0.44 to 0.91	−0.26	−0.44 to −0.11	−0.10	−0.17 to −0.02

All \hat{R} values are 1.0 and the effective sample sizes are larger than 837.

the positive frequency-dependent copying strategy was above zero (1.05 to 2.50); Supplementary Table 4). We found that positive frequency-dependent copying decreased with increasing task uncertainty (the 95% Bayesian CI of task uncertainty effect was below zero (−1.88 to −0.25); Supplementary Table 4). We found no clear effects of either the group size, age or gender on adoption of the positive frequency-dependent copying strategy, except for the negative interaction effect between age and task uncertainty (the 95% Bayesian CI of the age \times uncertainty interaction was −1.46 to −0.15; Supplementary Table 4).

We also investigated the effects of group size and task uncertainty on the fitted individual parameter values. We found that the individual mean social learning weight parameter (that is, $\bar{\sigma}_i = (\sum_t \sigma_{i,t}) / (\text{total rounds})$) increased with group size (95% Bayesian CI = 0.15 to 0.93; Fig. 4b and Supplementary Table 5), and decreased with uncertainty (95% Bayesian CI = −0.98 to −0.22) and the age of the subject (95% Bayesian CI = −0.36 to −0.02). However, the negative effects of task uncertainty and age disappeared when we focused only on $\bar{\sigma}_i$ of the positive frequency-dependent copying individuals, and only the positive effect of the group size was

Table 2 | Mean and the 95% Bayesian CIs of the posterior global means for the parameter values.

Uncertainty	Groups			Solitary individuals		
	Low (<i>n</i> = 77)	Moderate (<i>n</i> = 98)	High (<i>n</i> = 398)	Low (<i>n</i> = 36)	Moderate (<i>n</i> = 34)	High (<i>n</i> = 56)
μ_{α} (Learning rate)	0.99 0.34 to 1.73	0.90 0.43 to 1.44	0.61 0.21 to 1.03	0.85 −0.07 to 1.95	−0.17 −1.27 to 0.89	0.46 −0.39 to 1.36
μ_{β_0} (Inverse temperature)	1.84 1.15 to 2.70	1.68 1.25 to 2.18	1.38 1.16 to 1.62	1.10 0.69 to 1.54	1.44 0.80 to 2.07	0.85 0.46 to 1.22
μ_{ϵ} (Inverse temperature)	3.70 1.98 to 5.71	3.01 1.88 to 4.27	2.97 2.37 to 3.60	2.39 1.46 to 3.53	2.81 1.64 to 4.07	2.27 1.40 to 3.31
μ_{σ_0} (Social learning weight)	−1.55 −2.71 to −0.71	−2.37 −4.12 to −1.01	−2.16 −2.81 to −1.63	—	—	—
μ_{δ} (Social learning weight)	−1.39 −2.66 to −0.03	−1.55 −4.29 to 0.91	−1.87 −3.04 to −0.81	—	—	—
μ_{θ} (Conformity coefficient)	1.65 0.83 to 2.82	3.00 1.57 to 4.85	2.67 1.80 to 3.73	—	—	—

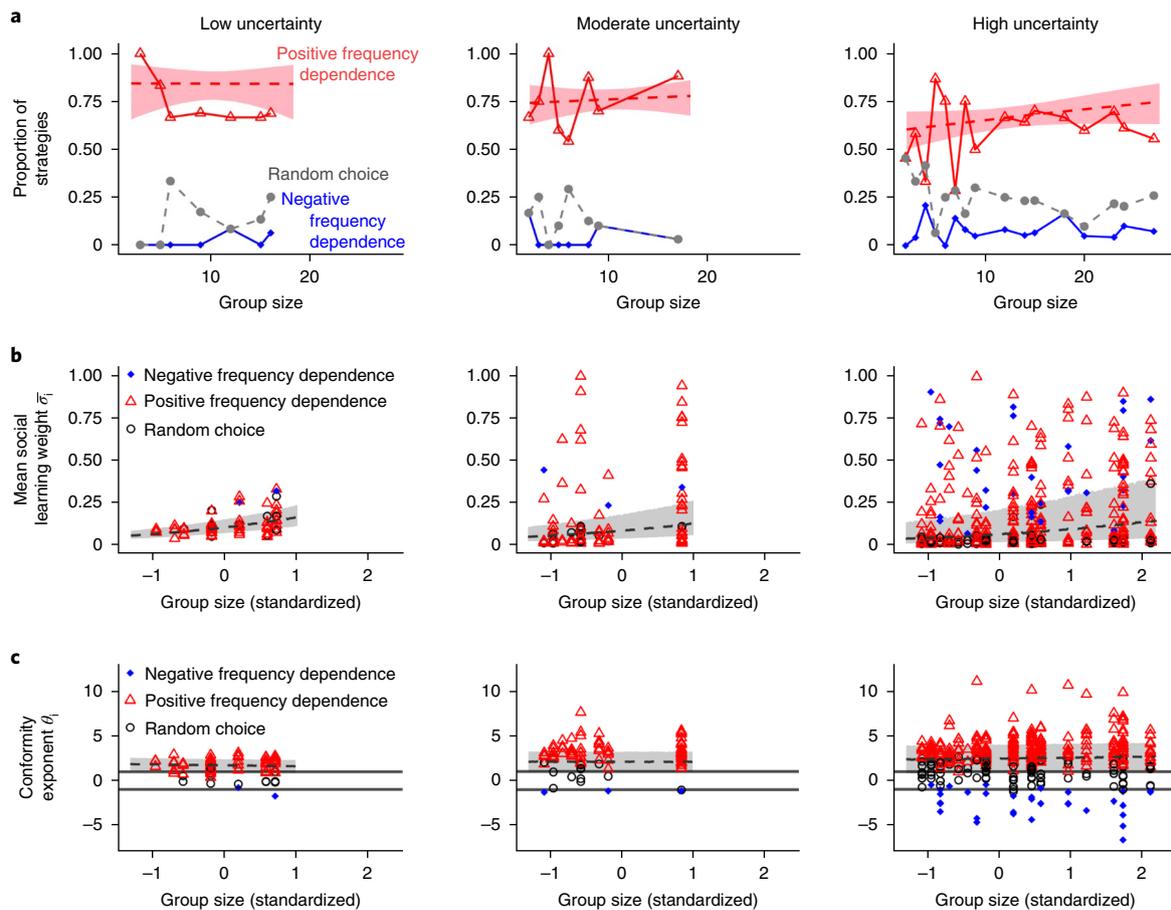


Fig. 4 | Model fitting for the three different task's uncertain conditions (low-, moderate- and high-uncertainty) and the different group sizes.

Three different learning strategies are shown (red triangle, positive frequency-dependent learning; blue circle, negative frequency-dependent learning; grey circle, nearly random choice strategy). **a**, Frequencies of three different learning strategies. Note that a sum of the frequencies of these three strategies in the same group size does not necessarily equal to 1, because there are a small number of individuals eliminated from this analysis due to insufficient data. **b,c**, Estimated social learning weight (**b**) estimated conformity exponent (**c**) for each individual shown for each learning strategy. The 50% Bayesian CIs of the fitted GLMMs are shown by dashed lines and shaded areas. The horizontal lines in **c** show the region $-1 < \theta_i < 1$. The sample sizes for negative frequency-dependent, positive frequency-dependent and random choice strategies were $n = 2, 61$ and 14 , respectively, for the low-uncertainty condition; $n = 3, 80$ and 15 , respectively, for the moderate-uncertainty condition; and $n = 32, 260$ and 106 , respectively, for the high-uncertainty condition.

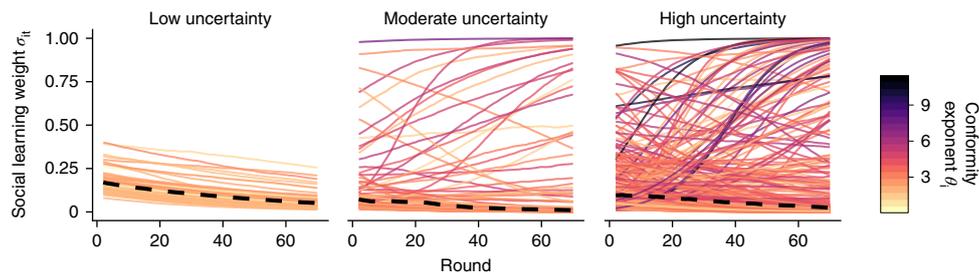


Fig. 5 | Change in fitted values (that is, the median of the Bayesian posterior distribution) of the social learning weight σ_{it} with time for each positive frequency-dependent individual for each level of task uncertainty. Thick dashed lines are the median values of σ_{it} across the subjects for each uncertainty condition. Individual conformity exponent values θ_i are shown in different colours (darker = higher θ_i). Sample sizes for each task uncertainty condition are $n = 61$ (low uncertainty), $n = 80$ (moderate uncertainty) and $n = 260$ (high uncertainty).

confirmed (Supplementary Table 6 and Supplementary Fig. 6). It is worth noting that the meaning of the social learning weight is different between these three different strategies. The social learning weight regulates positive reactions to the majorities' behaviour for positive frequency-dependent copiers, whereas it regulates avoidance of the majority for negative frequency-dependent copiers, and determines the probability of random decision-making for the random choice strategists.

The individual conformity exponent parameter θ_i increased with task uncertainty (95% Bayesian CI = 0.38 to 1.41), but we found no significant effects of group size, age, gender or interactions (Fig. 4c and Supplementary Table 7). These results were qualitatively unchanged when we focused only on the positive frequency-dependent copying individuals (Supplementary Table 8 and Supplementary Fig. 6).

We observed extensive individual variation in social information use. The greater the task's uncertainty, the larger the individual variations in both the mean social learning weight and the conformity exponent (the 95% Bayesian CI of the GLMM's variation parameter for $\bar{\sigma}_i$ was 1.11 to 1.62 (Supplementary Table 5) and for θ_i it was 1.07 to 1.54 (Supplementary Table 7)). This was confirmed when focusing only on the positive frequency-dependent copying individuals (the Bayesian 95% CIs were 1.14 to 1.80 (Supplementary Table 6) and 0.71 to 1.10 (Supplementary Table 8), respectively).

The manner in which individual variation in social information use of positive frequency-dependent copying individuals changes over time is visualized in Fig. 5. The social learning weights generally decreased with experimental round. However, some individuals in the moderate- and high-uncertainty conditions accelerated rather than decreased their reliance on social learning over time. Interestingly, those accelerating individuals tended to have a larger conformity exponent (Supplementary Fig. 5). In addition, the time-dependent $\theta_{i,t}$ in our alternative model generally increased with experimental round in the moderate- and the high-uncertainty conditions (Supplementary Fig. 10), although the fitting of $\theta_{i,t}$ in the alternative model was relatively unreliable (Supplementary Fig. 9). These findings suggest that conformists tended to use asocial learning at the outset (that is, exploration asocially), but increasingly started to conform as the task proceeded (that is, exploitation socially).

Extensive variation in the temporal dynamics of the social learning weight $\sigma_{i,t}$ was also found for the negative frequency-dependent copying individuals, but not for random choice individuals (Supplementary Fig. 5). Individuals deploying a random choice strategy exhibited a $\sigma_{i,t}$ that approached zero, indicating that their decision-making increasingly relied exclusively on the softmax choice rule, rather than unguided random choices, as the task proceeded.

No significant fixed effects were found in other asocial learning parameters, such as the learning rate α_i and the mean inverse

temperature $\bar{\beta}_i = (\sum_t \beta_{i,t}) / (\text{total rounds})$ (Supplementary Tables 9 and 10 and Supplementary Fig. 6).

In summary, our experiments on adult humans revealed asymmetric influences of increasing task uncertainty and increasing group size on the social learning parameters. The conformity exponent increased with task uncertainty on average, but the proportion of positive frequency-dependent copying individuals showed a corresponding decrease, due to the extensive individual variation emerging in the high-uncertainty condition. Conversely, group size had a positive effect on the mean social learning weight, but did not affect conformity.

Social learning strategies explain the collective dynamics. The post-hoc simulation provides statistical predictions on how likely it is, given the fitted learning model parameters, that groups of individuals make accurate decisions and that they exhibit inflexible herding. Figure 3c shows the change over time in performance with different group sizes and different uncertainty conditions, generated by the post-hoc simulation (see also Supplementary Fig. 3). The trajectories of the simulated dynamics recover nicely the pattern observed in the experiment (Fig. 3a,c), suggesting that the strategic changes in individual-level social information use (Fig. 4) could explain the collective-level behavioural pattern.

Figure 3d shows that larger groups are more likely to make accurate decisions than both small groups and solitaires in the period before change across all uncertainty conditions, suggesting that collective intelligence is operating. However, in the post-change period, performance differed between the conditions. In the low-uncertainty condition, where we found that the participants were most likely to have a relatively weak positive frequency-dependence (that is, $\bar{\theta} = 1.65$), large groups performed better than small groups over 59.5% of a total of 10,000 repetitions. However, in the moderate-uncertainty condition, where we found that participants were most likely to have strong positive frequency dependence ($\bar{\theta} = 3.00$ compared with 1.65 in the low-uncertainty condition), the large groups were more likely to get stuck on the suboptimal option, and hence the small groups performed better than the large groups over 69.5% of repetitions (Fig. 3d). The decision accuracy did not substantially differ with group size in the post-change period in the high-uncertainty condition, although the large groups performed slightly better than the small groups (50.8% of the repetitions).

Interestingly, although their relatively low conformity biases, there were some groups in the low-uncertainty condition that seemed to exhibit herding (the 'humped' area at the left-hand side to the peak of the performance distribution in the post-change period; Fig. 3d). This might be due to the lower softmax exploration rates among social learners in the low-uncertainty condition (that is, both μ_{β_i} and μ_ϵ were large; Table 2): the whole population gets stuck because all individuals are very exploitative on their past experience.

Discussion

We investigated whether and how human social learning strategies regulate the trade-off between collective intelligence and inflexible herding behaviour using a collective learning-and-decision-making task combined with simulation and model fitting. We examined whether manipulating the reliability of asocial learning and group size would affect the use of social information, and thereby alter the collective human decision dynamics, as suggested by our computational model simulation. Although a theoretical study has suggested that reliance on social learning and conformity bias would play a role in collective dynamics^{2,5,53}, thus far no empirical studies have quantitatively investigated the population-level consequences of these two different social learning processes. Our high-resolution, model-based behavioural analysis using hierarchical Bayesian statistics enabled us to identify individual-level patterns and variation of different learning parameters, and to explore their population-level outcomes. The results provide quantitative support for our hypothesis that the collective decision performance can be predicted with quantitative knowledge of social learning strategies.

Overall, our individual-based computational model recovered the behavioural pattern suggested by the phenomenological regression (Fig. 3). Using the post-hoc simulation with individually fit model parameters, we confirmed that, in the low-uncertainty condition, where individuals had weaker positive frequency bias (that is, $\bar{\theta} \approx 1.65$), larger groups were able to be more accurate than smaller groups while retaining flexibility in their decision-making⁹, although their low asocial exploration rates seemed to undermine the potential flexibility. However, in the moderate- and the high-uncertain conditions, where individuals had the higher conformity exponent parameters (that is, $\bar{\theta} \approx 3.0$ and 2.7 , respectively), larger groups performed better before environmental change but were vulnerable to getting stuck with an outdated maladaptive option post-change. Therefore, the changes in the level of conformity in human individuals^{34,41} indeed incurred a trade-off between the collective intelligence effect and the risk of inflexible herding.

Although the social learning weight increased with increasing group size, the overall mean value was $\bar{\sigma}_i \approx 0.3$ (Fig. 4b and Supplementary Figs. 5 and 6), and it decreased on average as the task proceeded (Fig. 5). This implies a weaker social than asocial influence on decision-making, as reported in several other experimental studies^{35,54–56}, although evolutionary models tend to predict heavier reliance on social learning than experimental studies report^{57,58}. Thanks to this relatively weak reliance of social learning, the kind of extreme herding that would have blindly led a group to any option regardless of its quality, such as the ‘symmetry breaking’ known in trail-laying ant collective foraging systems^{2,5,26}, did not occur (Fig. 2).

Individual differences in rates of exploration might also help to mitigate potential herding. Although a majority of participants adopted a positive frequency-dependent copying strategy, some individuals exhibited negative frequency dependence or random decision-making (Fig. 4a). The random choice strategy was associated with more exploration than the other strategies, because it led to an almost random choice at a rate σ_r , irrespective of the options’ quality. Negative frequency-dependent copying individuals could also be highly exploratory. These individuals tended to avoid choosing an option on which other people had converged and would explore the other two ‘unpopular’ options. Interestingly, in the high-uncertainty condition, the mean social learning weights of the negative frequency-dependent copying individuals ($\bar{\sigma}_i \approx 0.5$) were larger than those of the other two strategies ($\bar{\sigma}_i \approx 0.1$; Supplementary Fig. 5), indicating that these individuals engaged in such majority-avoiding exploration relatively frequently. Such a high variety in social information use^{59–62} and exploratory tendencies would prevent individuals from converging on a single option, leading to a mitigation of herding

but concurrently diminishing the decision accuracy in high-uncertainty circumstances (Fig. 3).

A methodological advantage of using computational models to study social learning strategies is its explicitness of assumptions about the temporal dynamics of behaviour, which enabled us to distinguish between different learning strategies^{63–65}. For example, very exploitative asocial reinforcement learners (that is, those for whom the exploitation parameter $\beta_{i,t}$ is large and the social learning weight $\sigma_{i,t}$ is nearly zero, as seen in the low-uncertainty condition) and conformity-biased social learners (where the conformity exponent θ_i is large and $\sigma_{i,t}$ is positive, as seen in the moderate-uncertainty condition) would eventually converge on the same option, resulting in the same final behavioural steady state. However, how they explored the environment, as well as how they reacted to the other individuals in the same group, was significantly different, and they could produce qualitatively different collective temporal dynamics.

However, our computational model could not fully capture other, potentially more sophisticated forms of social learning strategies that participants might deploy, which might be a reason for the seemingly low rate of social learning observed in the experiment compared with theory^{57,58}. Indeed, the post-hoc simulation sometimes failed to recover the observed behavioural trajectories. In particular, experimental groups with $n=12$ or 16 , and 1 group with $n=9$ in the low-uncertainty condition performed very well, exceeding the 95% CIs of the post-hoc simulation after the environmental change (Supplementary Fig. 3). This indicates that collective behaviour in these groups was more flexible than our model predicted. Further empirical studies that consider a wider range of possible social learning strategies (for example, the ‘copy-rapidly-increasing-option’ strategy⁶⁶ or Bayesian updating^{57,67}) are needed to explore computational underpinnings of social learning and collective behaviour.

The Internet-based experimentation allowed us to conduct a real-time interactive behavioural task with larger subject pools than a conventional laboratory-based experiment. This enabled us not only to quantify the individual-level learning-and-decision processes⁶⁸ but also to map these individual-level processes on to the larger-scale collective behaviour^{5,15,20}. Although there are always questions about the validity of participants’ behaviour when recruited via web-based tools, we believe that the computational modelling approach, coupled with higher statistical power due to the large sample size, compensates for any drawbacks. The fact that our learning model could approximate the participants’ decision trajectories effectively suggests that most of the participants engaged seriously with solving the task. An increasing body of evidence supports the argument that web-based behavioural experiments are as reliable as results from the laboratory^{69,70}.

The diverse effects of social influence on the collective wisdom of a group has been drawing substantial attention^{19,21,22,71,72}. The bulk of this literature, including many jury models and election models^{45,73}, has focused primarily on the static estimation problem, where the ‘truth’ is fixed from the outset. However, in reality, there are many situations under which the state of the true value is changing over time so that monitoring and tracking the pattern of change is a crucial determinant of decision performance⁷⁴. In such temporally dynamic environments, decision-making and learning are coordinated to affect future behavioural outcomes recursively⁷⁵. Our experimental task provides a simple vehicle for exploring collective intelligence in a dynamic situation, which encompasses this learning-and-decision-making feedback loop. Potentially, integrating the wisdom of crowds with social learning and collective dynamics research will facilitate the more tractable use of collective intelligence in a temporary changing world.

In summary, a combination of experimentation and theoretical modelling sheds light on when groups of individuals will exhibit the wisdom of the crowds and when they will choose inflexible

herding. Our analysis implies that herding is most likely among individuals in large groups exposed to challenging tasks. This is because challenging tasks lead to greater uncertainty and thereby elicit greater conformist learning among individuals, while rates of copying increase with group size. Difficult tasks, by definition, render identification of the optimal behaviour harder, allowing groups to sometimes converge on maladaptive outcomes. Conversely, the reduced conformity levels of individuals in small groups, and the greater probability that social information would be accurate for less challenging tasks, generated ‘wisdom of the crowd’ effects in most other circumstances. Our findings provide clear evidence that the conflict between collective intelligence and maladaptive herding can be predicted with knowledge of human social learning strategies.

Methods

Participants. The experimental procedure was approved by the Ethics Committee at the University of St Andrews (BL10808). A total of 755 subjects (354 females, 377 males, 2 others and 22 unspecified; mean age (1 s.d.) = 34.33 years (10.9 years)) participated through Amazon’s Mechanical Turk. All participants consented to participation through an online consent form at the beginning of the task. We excluded subjects who disconnected from the online task before completing at least the first 30 rounds from our computational model-fitting analysis due to unreliability of the model-parameter estimation, resulting in 699 subjects (573 subjects entered the group condition (that is, $n \geq 2$) and 126 entered the solitary condition (that is, $n = 1$)). The task was only available for individuals who had a $\geq 90\%$ Human Intelligence Task (HIT) approval rate and who accessed the task from the United States. Although no sample-size calculation was performed in advance, our parameter recovery test confirmed that the sample size was sufficient for estimation of individual parameters using a hierarchical Bayesian method (HBM).

Design of the experimental manipulations. The 3 uncertainty conditions were: low uncertainty (differences between mean payoffs were 1.264), moderate uncertainty (differences between mean payoffs were 0.742) and high uncertainty (differences between mean payoffs were 0.3). The mean payoff associated with the ‘excellent’ slot in all 3 conditions was fixed at 3.1 cents (Supplementary Fig. 1). Each task uncertainty condition was randomly assigned for each different HIT session, and participants were allowed to participate in one HIT only. The sample sizes after data exclusion for each uncertainty condition were: $n = 113$ (low uncertainty), $n = 132$ (moderate uncertainty) and $n = 454$ (high uncertainty). We assigned more sessions to the high-uncertainty condition compared with the other two conditions because we expected that larger group sizes would be needed to generate the collective wisdom in noisier environments.

To manipulate the size of each group, we varied the capacity of the waiting room from 10 to 30. Because the task was being advertised on the Worker website at Amazon’s Mechanical Turk for approximately 2 h, some participants occasionally arrived after the earlier groups had already started. In these cases, the participant entered a newly opened waiting room, which was open for the next 5 min. The number of participants arriving declined with time because newly posted alternative HITs were advertised at the top of the task list, which decreased our task’s visibility. This meant that later-starting sessions tended to begin before reaching maximum room capacity, resulting in smaller group sizes. Therefore, the actual sizes differed between groups (Supplementary Fig. 3 and Supplementary Table 1). Data collection and analysis were not performed blind to the conditions of the experiments.

Multiplayer three-armed bandit task. To study the relationship between social information use and collective behavioural dynamics, we focused on a well-established learning-and-decision problem called a ‘multi-armed bandit’ task, represented here as repeated choices between three slot machines (Supplementary Fig. 1 and Supplementary Video 1; for details, see the Supplementary Methods). Participants played the task for 70 rounds. The slots paid off money noisily (in US cents), varying around 2 different means during the first 40 rounds such that there was one ‘good’ slot and two other options giving poorer average returns. However, from the 41st round, 1 of the ‘poor’ slots abruptly increased its mean payoff to become ‘excellent’ (that is, superior to ‘good’). The purpose of this environmental change was to observe the effects of maladaptive herding by potentially trapping groups in the out-of-date suboptimal (good) slot, as individuals did not know whether or how an environmental change would occur. Through making choices and earning a reward from each choice, individuals could gradually learn which slot generated the highest rewards.

In addition to this social learning, we provided social information for each member of the group specifying the frequency with which group members chose each slot. All group members played the same task with the same conditions simultaneously, and all individuals had been instructed that this was the case, and hence understood that the social information would be informative.

Task uncertainty was experimentally manipulated by changing the difference between the mean payoffs for the slot machines. In the task with the least uncertainty, the distribution of payoffs barely overlapped, while in the task with the greatest uncertainty the distribution of payoffs overlapped considerably (Supplementary Fig. 1).

Computational learning-and-decision-making model. We modelled individual behavioural processes by assuming that individual i makes a choice for option m at round t , in accordance with the choice probability $p_{i,t}(m)$, which is a weighted average of social and asocial influences:

$$P_{i,t}(m) = \sigma_{i,t} \times \text{social influence}_{i,m,t} + (1 - \sigma_{i,t}) \times \text{asocial influence}_{i,m,t} \quad (1)$$

where $\sigma_{i,t}$ is the social learning weight ($0 \leq \sigma_{i,t} \leq 1$).

For the social influence, we assumed a frequency-dependent copying strategy by which an individual copies others’ behaviour in accordance with the distribution of social frequency information^{49–51,55}:

$$\text{Social influence}_{i,m,t} = \frac{(F_{t-1}(m) + 0.1)^{\theta_i}}{\sum_{k \in \text{options}} (F_{t-1}(k) + 0.1)^{\theta_i}} \quad (2)$$

where $F_{t-1}(m)$ is a number of choices made by other individuals (excluding her/his own choice) for the option m in the preceding round $t - 1$ ($t \geq 2$). θ_i is individual i ’s conformity exponent, $-\infty \leq \theta_i \leq +\infty$. When this exponent is larger than zero ($\theta_i > 0$), higher social influence is given to an option that was chosen by more individuals (that is, positive frequency bias), with conformity bias arising when $\theta_i > 1$, such that disproportionately more social influence is given to the most common option²⁸. In contrast, when $\theta_i < 0$, higher social influence is afforded to the option that fewest individuals chose in the preceding round $t - 1$ (that is, negative frequency bias). To implement the negative frequency dependence, we added a small number 0.1 to F so that an option chosen by no one (that is, $F_{t-1} = 0$) could provide the highest social influence when $\theta_i < 0$. Note that there is no social influence when $\theta_i = 0$ because in this case the ‘social influence’ favours a uniformly random choice (that is, $S_{i,t}(m) = f_m^0 / (f_1^0 + f_2^0 + f_3^0) = 1/3$), independent of the social frequency distribution. Note also that, in the first round $t = 1$, we assumed that the choice was only determined by the asocial softmax function because there was no social information available.

For the asocial influence, we used a standard reinforcement learning with ‘softmax’ choice rule⁷⁵ that is widely applied in human social learning studies (for example, refs. ^{35,51,55}). An individual i updates the estimated average reward associated with an option m at round t ; namely, the Q value ($Q_{i,t}(m)$), according to the Rescorla–Wagner rule as follows:

$$Q_{i,t+1}(m) = Q_{i,t}(m) + \alpha_i (1 - m_{i,t}) (r_{i,t}(m) - Q_{i,t}(m)) \quad (3)$$

where α_i ($0 \leq \alpha_i \leq 1$) is a learning rate parameter of individual i determining the weight given to new experience and $r_{i,t}(m)$ is the amount of monetary reward obtained from choosing the option m in round t . $1(m, m_{i,t})$ is the binary action-indicator function of individual i , given by:

$$1(m, m_{i,t}) = \begin{cases} 1, & \text{if } m_{i,t} = m \text{ or } t = 1 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

Therefore, $Q_{i,t}(m)$ was updated only when the option m was chosen; when the option m was not chosen, $Q_{i,t}(m)$ was not updated (that is, $Q_{i,t+1}(m) = Q_{i,t}(m)$). Note that, only in the first round $t = 1$ were all Q values updated using the chosen option’s reward $r_{i,t}(m)$, so that the individual could set a naive ‘intuition’ about the magnitude of reward values she/he would expect to earn from a choice in the task; namely, $Q_{i,t=1}(1) = Q_{i,t=2}(2) = Q_{i,t=3}(3) = \alpha_i r_{i,t=1}(m)$. In practical terms, this prevents the model from being overly sensitive to the first experience. Before the first choice, individuals had no previous preference for either option (that is, $Q_{i,1}(1) = Q_{i,1}(2) = Q_{i,1}(3) = 0$).

The Q value is then translated into the asocial influence through the softmax (or logit choice) function:

$$A_{i,t}(m) = \frac{\exp(\beta_{i,t} Q_{i,t}(m))}{\sum_{k \in \text{options}} \exp(\beta_{i,t} Q_{i,t}(k))} \quad (5)$$

where $\beta_{i,t}$, called the inverse temperature, manipulates individual i ’s sensitivity to the Q values (in other words, controlling the proneness to explore). As $\beta_{i,t}$ goes to zero, asocial influence approximates to a random choice (that is, highly explorative). Conversely, if $\beta_{i,t} \rightarrow +\infty$, the asocial influence leads to a deterministic choice in favour of the option with the highest Q value (that is, highly exploitative). For intermediate values of $\beta_{i,t}$, individual i exhibits a balance between exploration and exploitation^{35,68}. We allowed for the possibility that the balance between exploration and exploitation could change as the task proceeds. To depict such time dependence in exploration, we used the equation $\beta_{i,t} = \beta_{i,0}^* + \epsilon_i t / 70$. If the slope ϵ_i is positive (negative), asocial influence $A_{i,t}$ becomes more and more exploitative

(explorative) as round t increases. For a model-fitting purpose, the time-dependent term $\varepsilon_{i,t}$ is scaled by the total round number 70.

We allowed that the social learning weight $\sigma_{i,t}$ could also change over time as assumed in the inverse temperature $\beta_{i,t}$. To let $\sigma_{i,t}$ satisfy the constraint $0 \leq \sigma_{i,t} \leq 1$, we used the following sigmoidal function:

$$\sigma_{i,t} = \frac{1}{1 + \exp(-(\alpha_{i,0}^* + \delta_i t/70))} \quad (6)$$

If the slope δ_i is positive (negative), the social influence increases (decreases) over time. We set the social learning weight equal to zero when the group size was one (that is, when an individual participated in the task alone and/or when $\sum_{k \in \text{options}} F_{i,t}(k) = 0$).

We modelled both the inverse temperature $\beta_{i,t}$ and the social learning weight $\sigma_{i,t}$ as a time function since otherwise it would be challenging to distinguish between different patterns of learning in this social learning task⁶⁵. The parameter recovery test confirmed that we were able to differentiate such processes under these assumptions (Supplementary Figs. 7 and 8). While we also considered the possibility of the conformity exponent being time dependent (that is $\theta_{i,t} = \theta_{i,0}^* + \gamma t/70$), the parameter recovery test suggested that the individual slope parameter γ was not reliably recovered (Supplementary Fig. 9), and hence we concentrated our analysis on the time-independent θ_i model. We confirmed that using the alternative model where both social learning parameters were time dependent (that is, $\sigma_{i,t}$ and $\theta_{i,t}$) did not qualitatively change our results (Supplementary Fig. 10).

One concern might be the asymmetry between the asocial softmax influence, which takes many previous experiences into account (depending on a learning rate), and the social influence referring only the most recent frequency information $F_{i,t}$. The choice frequency appearing at round t is the most reliable social information, compared with past frequencies, because it could be the most 'updated' information as long as the other individuals have made informed decisions to their best knowledge. On the other hand, the accumulated Q values are the most reliable asocial information, compared with an option's reward earned at the most recent round, which was independently and randomly drawn from a probability distribution. Therefore, although many other formulations for asocial and social learning processes were possible, we believe that our current choice—time-depth asocial reinforcement learning with the most updated frequency-dependent copying—was a reasonable first step.

In summary, the model has six free parameters that were estimated for each individual human participant; namely, α , $\beta_{i,0}^*$, ε , $\sigma_{i,0}^*$, δ , and θ_i . To fit the model, we used a HBM, estimating the global means (μ_α , $\mu_{\beta_{i,0}^*}$, μ_ε , $\mu_{\sigma_{i,0}^*}$, μ_δ and μ_θ) and global variations (ν_α , $\nu_{\beta_{i,0}^*}$, ν_ε , $\nu_{\sigma_{i,0}^*}$, ν_δ and ν_θ) for each of the three experimental conditions (that is, low, moderate and high uncertainty), which govern the overall distributions of individual parameter values. It has become recognized that the HBM can provide more robust and reliable parameter estimation than conventional maximum-likelihood point estimation in complex cognitive models⁶⁶—a conclusion with which our parameter recovery test agreed (Supplementary Figs. 7 and 8).

Agent-based model simulation. We ran a series of individual-based model simulations assuming that a group of individuals play our 3-armed bandit task for 90 rounds (under the moderate-uncertainty condition) and that individuals behave in accordance with the computational learning-and-decision model. We varied the group size ($n \in \{3, 10, 30\}$), mean social learning weight ($\bar{\sigma} \in \{0.01, 0.1, 0.2, 0.3, \dots, 0.9\}$) and mean conformity exponent ($\bar{\theta} \in \{0.5, 1, 3, 6\}$), running 10,000 replications for each of the possible parameter \times group size combinations. As for the other parameter values (for example, the asocial reinforcement learning parameters α , $\beta_{i,0}^*$, ε), here we used the experimentally fitted global means (Table 2 and Supplementary Table 3). Relaxation of this assumption (that is, using a different set of asocial learning parameters) does not qualitatively change our story (Supplementary Fig. 2). Note that each individual's parameter values were randomly drawn from the distributions centred by the global mean parameter values fixed to each simulation run. Therefore, the actual composition of individual parameter values was different between individuals even within the same social group.

Generalized linear mixed models. To directly analyse the effects of group size and task uncertainty on the time evolution of decision performance, we conducted a statistical analysis using a phenomenological model; namely, a hidden Markov process logistic regression without assuming any specific learning-and-decision-making processes. The dependent variable was whether the participant chose the best option (1) or not (0). The model includes fixed effects of grouping ξ , standardized group size ω and an intercept with a random effect of individuals $\mu + \rho_i$. We assumed that the intercept and effect of grouping changed from round to round, as a random walk process. For the effect of group size, we considered the effect of the first environment $1 \leq t \leq 40$ and that of the second environment (namely, ω_1 and ω_2), separately.

To examine whether increasing the group size and task uncertainty affected individual use of the positive frequency-dependent copying strategy, we used a

hierarchical Bayesian logistic regression model with a random effect of group.

The dependent variable was whether the participant used the positive frequency-dependent copying (1) or not (0). The model includes fixed effects of group size (standardized), task uncertainty (0, low; 0.5, moderate; 1, high), age (standardized), gender (0, male; 1, female; NA, others or unspecified), and possible two-way interactions between these fixed effects.

We also investigated the effects of both group size and the task's uncertainty on the fitted values of the learning parameters. We used a hierarchical Bayesian Gaussian regression model predicting the individual fitted parameter values. The model includes effects of group size (standardized), task uncertainty (0, low; 0.5, moderate; 1, high), age (standardized), gender (0, male; 1, female; NA, others or unspecified) and two-way interactions between these fixed effects. We assumed that the variance of the individual parameter values might be contingent on task uncertainty because we had found in the computational model-fitting result that the fitted global variance parameters (that is, $\nu_{\sigma_{i,0}^*}$, ν_δ and ν_θ) were larger in more uncertain conditions (Supplementary Table 2).

Statistical analysis. We used a HBM to estimate the free parameters of our statistical models, including both the phenomenological regression model and the computational learning-and-decision-making model. The HBM allows us to estimate individual differences, while ensuring that these individual variations are bounded by the group-level global parameters. The HBM was performed under Stan 2.16.2 (<http://mc-stan.org>) in R 3.4.1 (<https://www.r-project.org>) software. The models contained at least four parallel chains and we confirmed convergence of the MCMC using both the Gelman–Rubin statistics and the effective sample sizes. Full details of the model-fitting procedure and prior assumptions are shown in the appendix.

Parameter recovery test. To check the validity of our model-fitting method, we conducted a 'parameter recovery test' to examine how well our model-fitting procedure had been able to reveal true individual parameter values. To do this, we generated synthetic data by running a simulation with the empirically fitted global parameter values, then re-fitted the model with these synthetic data using the same procedure. The parameter recovery test showed that the all-true global parameter values fell within the 95% Bayesian credible interval (Supplementary Fig. 7), and at least 93% of the true individual parameter values were correctly recovered (that is, 96% of α , 93% of $\beta_{i,0}^*$, 95% of ε , 97% of $\sigma_{i,0}^*$, 96% of δ and 97% of θ_i values fell within into the 95% Bayesian CI; Supplementary Fig. 7).

Categorization of individual learning strategies. Based on the 50% CI of the individual conformity exponent parameter values θ_i , we divided the participants into the following three different social learning strategies. If her/his 50% CI of θ_i fell above zero ($\theta_{\text{lower}} > 0$), below zero ($\theta_{\text{upper}} < 0$) or including zero ($\theta_{\text{lower}} \leq 0 \leq \theta_{\text{upper}}$), she/he was categorized as a 'positive frequency-dependent copier', a 'negative frequency-dependent copier' or a 'random choice individual', respectively. We used the 50% Bayesian CI to conduct this categorization instead of using the more conservative 95% CI because the 95% CI would have caused much higher rates of 'false negatives', by which an individual who applied either a positive frequency-dependent copying strategy or negative frequency-dependent copying strategy would have been falsely labelled as an asocial random choice individual (Supplementary Fig. 7). Of 572 agents, 400 (~70%) were falsely categorized as random choice learners in the recovery test when we used the 95% criterion (Supplementary Fig. 7). In contrast, the 50% CI criterion seemed to be much better in terms of the false negative rate, which was only 18.5% (that is, 106 agents), although it might be slightly worse in terms of 'false positives': 37 agents (6.5%) were falsely labelled as either a positive frequency-dependent copier or negative frequency-dependent copier by the 50% CI, whereas the false positive rate of the 95% CI was only 0.2% (Supplementary Fig. 7). To balance the risk of false positives and false negatives, we decided to use the 50% CI, which seemed to have more strategy-detecting power.

The post-hoc model simulation. To evaluate how accurately our model can generate observed decision patterns in our task setting, we ran a series of individual-based model simulations using the fitted individual parameter values (that is, means of the individual posterior distributions) for each group size for each uncertainty condition. At the first step of the simulation, we assigned a set of fitted parameters of a randomly chosen experimental subject from the same group size and the same uncertainty condition to a simulated agent, until the number of agents reached the simulated group size. We allowed duplicate choices of experimental subjects in this parameter assignment. At the second step, we let this synthetic group of agents play the bandit task for 90 rounds. We repeated these steps 10,000 times for each group size and task uncertainty.

Reporting Summary. Further information on experimental design is available in the Nature Research Reporting Summary linked to this article.

Code availability

The browser-based online task was built by Node.js (<https://nodejs.org/en/>) and socket.io (<https://socket.io>), and the codes are available from GitHub

(<https://github.com/WataruToyokawa/MultiPlayerThreeArmedBanditGame>). Analyses were conducted in R (<https://www.r-project.org>) and simulations of individual-based models were conducted in Mathematica (<https://www.wolfram.com>), and the codes for both of these are available from GitHub (<https://github.com/WataruToyokawa/ToyokawaWhalenLaland2018>).

Data availability

Both experimental and simulation data are available in an online repository (<https://github.com/WataruToyokawa/ToyokawaWhalenLaland2018>).

Received: 8 June 2018; Accepted: 12 December 2018;

Published online: 21 January 2019

References

- Bonabeau, E., Dorigo, M. & Theraulaz, G. *Swarm Intelligence: From Natural to Artificial Systems* (Oxford Univ. Press, New York, 1999).
- Camazine, S. et al. *Self-Organization in Biological Systems* (Princeton Univ. Press, Princeton, 2001).
- Krause, J., Ruxton, G. D. & Krause, S. Swarm intelligence in animals and humans. *Trends Ecol. Evol.* **25**, 28–34 (2010).
- Seeley, T. D. *The Wisdom of the Hive* (Harvard Univ. Press, Cambridge, MA, 1995).
- Sumpter, D. J. T. *Collective Animal Behavior* (Princeton Univ. Press, Princeton, 2010).
- King, A. J. & Sueur, C. Where next? Group coordination and collective decision making by primates. *Int. J. Primatol.* **32**, 1245–1267 (2011).
- Morand-Ferron, J. & Quinn, J. L. Larger groups of passerines are more efficient problem solvers in the wild. *Proc. Natl Acad. Sci. USA* **108**, 15898–15903 (2011).
- Sasaki, T. & Biro, D. Cumulative culture can emerge from collective intelligence in animal groups. *Nat. Commun.* **8**, 1–6 (2017).
- Shaffer, Z., Sasaki, T. & Pratt, S. C. Linear recruitment leads to allocation and flexibility in collective foraging by ants. *Anim. Behav.* **86**, 967–975 (2013).
- Reid, C. R. & Latty, T. Collective behaviour and swarm intelligence in slime moulds. *FEMS Microbiol. Rev.* **40**, 798–806 (2016).
- Krause, J. & Ruxton, G. D. *Living in Groups* (Oxford Univ. Press, Oxford & New York, 2002).
- Mackay, C. *Extraordinary Popular Delusions and the Madness of Crowds* (Richard Bentley, London, 1841).
- Kameda, T. & Hastie, R. in *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource* 1–14 (Wiley, Hoboken, NJ, USA, 2015).
- Le Bon, G. *The Crowd: A Study of the Popular Mind* 4th edn (Unwin, London, 1896).
- Raafat, R. M., Chater, N. & Frith, C. Herding in humans. *Trends Cogn. Sci.* **13**, 420–428 (2009).
- Bikhchandani, S., Hirshleifer, D. & Welch, I. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100**, 992–1026 (1992).
- Chari, V. V. & Kehoe, P. J. Financial crises as herds: overturning the critiques. *J. Econ. Theory* **119**, 128–150 (2004).
- Janis, I. L. *Victims of Groupthink: A Psychological Study of Foreign Policy* (Houghton Mifflin Company, Boston, 1972).
- Muchnik, L., Aral, S. & Taylor, S. J. Social influence bias: a randomized experiment. *Science* **341**, 647–651 (2013).
- Salganik, M. J., Dodds, P. S. & Watts, D. J. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**, 854–856 (2006).
- Lorenz, J., Rauhut, H., Schweitzer, F. & Helbing, D. How social influence can undermine the wisdom of crowd effect. *Proc. Natl Acad. Sci. USA* **108**, 9020–9025 (2011).
- Jayles, B. et al. How social information can improve estimation accuracy in human groups. *Proc. Natl Acad. Sci. USA* **114**, 201703695 (2017).
- Giraldeau, L.-A., Valone, T. J. & Templeton, J. J. Potential disadvantages of using socially acquired information. *Phil. Trans. R. Soc. Lond. B* **357**, 1559–1566 (2002).
- Detrain, C. & Deneubourg, J. L. Collective decision-making and foraging patterns in ants and honeybees. *Adv. Insect Physiol.* **35**, 123–173 (2008).
- List, C., Elsholtz, C. & Seeley, T. D. Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Phil. Trans. R. Soc. Lond. B* **364**, 755–762 (2009).
- Deneubourg, J. L., Aron, S., Goss, S. & Pasteels, J. M. The self-organizing exploratory pattern of the Argentine ant. *J. Insect Behav.* **30**, 159–168 (1990).
- Beckers, R., Deneubourg, J. L. D., Goss, S. & Pasteels, J. M. Collective decision making through food recruitment. *Insectes Soc.* **37**, 258–267 (1990).
- Boyd, R. & Richerson, P. J. *Culture and the Evolutionary Process* (Univ. Chicago Press, Chicago, 1985).
- Richerson, P. J. & Boyd, R. *Not by Genes Alone* (Univ. Chicago Press, Chicago, 2005).
- Feldman, M. W., Aoki, K. & Kumm, J. Individual versus social learning: evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* **104**, 209–231 (1996).
- Laland, K. N. Social learning strategies. *Anim. Learn. Behav.* **32**, 4–14 (2004).
- Kameda, T. & Nakanishi, D. Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment. *Evol. Hum. Behav.* **23**, 373–393 (2002).
- Kendal, R. L., Coolen, I. & Laland, K. N. The role of conformity in foraging when personal and social information conflict. *Behav. Ecol.* **15**, 269–277 (2004).
- Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W. & Laland, K. N. The evolutionary basis of human social learning. *Proc. Biol. Sci.* **B 279**, 653–662 (2012).
- Toyokawa, W., Saito, Y. & Kameda, T. Individual differences in learning behaviours in humans: asocial exploration tendency does not predict reliance on social learning. *Evol. Hum. Behav.* **38**, 325–333 (2017).
- Webster, M. M. & Laland, K. N. Social learning strategies and predation risk: minnows copy only when using private information would be costly. *Proc. R. Soc. B* **275**, 2869–2876 (2008).
- Webster, M. M. & Laland, K. N. Reproductive state affects reliance on public information in sticklebacks. *Proc. Biol. Sci.* **B 278**, 619–627 (2011).
- Boyd, R. & Richerson, P. J. Social learning as an adaptation. *Lect. Math. Life Sci.* **20**, 1–26 (1989).
- Bond, R. Group size and conformity. *Group Process. Intergroup Relat.* **8**, 331–354 (2005).
- Kline, M. A. & Boyd, R. Population size predicts technological complexity in Oceania. *Proc. R. Soc. B* **277**, 2559–2564 (2010).
- Muthukrishna, M., Shulman, B. W., Vasilescu, V. & Henrich, J. Sociality influences cultural complexity. *Proc. R. Soc. B* **281**, 20132511 (2014).
- Street, S. E., Navarrete, A. F., Reader, S. M. & Laland, K. N. Coevolution of cultural intelligence, extended life history, sociality, and brain size in primates. *Proc. Natl Acad. Sci. USA* **114**, 1–7 (2017).
- Nicolis, S. & Deneubourg, J. Emerging patterns and food recruitment in ants: an analytical study. *J. Theor. Biol.* **198**, 575–592 (1999).
- Pratt, S. C. & Sumpter, D. J. T. A tunable algorithm for collective decision-making. *Proc. Natl Acad. Sci. USA* **103**, 15906–15910 (2006).
- List, C. Democracy in animal groups: a political science perspective. *Trends Ecol. Evol.* **19**, 166–168 (2004).
- King, A. J. & Cowlshaw, G. When to use social information: the advantage of large group size in individual decision making. *Biol. Lett.* **3**, 137–139 (2007).
- Wolf, M. et al. Accurate decisions in an uncertain world: collective cognition increases true positives while decreasing false positives. *Proc. R. Soc. B Biol. Sci.* **280**, 20122777 (2013).
- Laan, A., Madirolas, G. & Polavieja, G. G. D. Rescuing collective wisdom when the average group opinion is wrong. *Front. Robot. AI* **4**, 1–28 (2017).
- Aplin, L. M., Sheldon, B. C. & McElreath, R. Conformity does not perpetuate suboptimal traditions in a wild population of songbirds. *Proc. Natl Acad. Sci. USA* **114**, 7830–7837 (2017).
- Barrett, B. J., McElreath, R. L., Perry, S. E. & Barrett, B. J. Pay-off-biased social learning underlies the diffusion of novel extractive foraging traditions in a wild primate. *Proc. R. Soc. B* **284**, 20170358 (2017).
- McElreath, R. et al. Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Phil. Trans. R. Soc. Lond. B* **363**, 3515–3528 (2008).
- Toyokawa, W., Kim, H.-R. & Kameda, T. Human collective intelligence under dual exploration-exploitation dilemmas. *PLoS One* **9**, e95789 (2014).
- Kandler, A. & Laland, K. N. Tradeoffs between the strength of conformity and number of conformists in variable environments. *J. Theor. Biol.* **332**, 191–202 (2013).
- 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 (2008).
- McElreath, R. et al. Applying evolutionary models to the laboratory study of social learning. *Evol. Hum. Behav.* **26**, 483–508 (2005).
- Mesoudi, A. An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evol. Hum. Behav.* **32**, 334–342 (2011).
- Perreault, C., Moya, C. & Boyd, R. A Bayesian approach to the evolution of social learning. *Evol. Hum. Behav.* **33**, 449–459 (2012).
- Rendell, L. et al. Why copy others? Insights from the social learning strategies tournament. *Science* **328**, 208–213 (2010).
- Jolles, J. W., Laskowski, K. L., Boogert, N. J. & Manica, A. Repeatable group differences in the collective behaviour of stickleback shoals across ecological contexts. *Proc. R. Soc. B* **285**, 13–16 (2018).
- Michelena, P., Jeanson, R., Deneubourg, J.-L. & Sibbald, A. M. Personality and collective decision-making in foraging herbivores. *Proc. R. Soc. B* **277**, 1093–1099 (2010).

61. Planas-Sitjà, I., Deneubourg, J.-L., Gibon, C. & Sempo, G. Group personality during collective decision-making: a multi-level approach. *Proc. R. Soc. B* **282**, 20142515 (2015).
62. Mesoudi, A., Chang, L., Dall, S. R. X. & Thornton, A. The evolution of individual and cultural variation in social learning. *Trends Ecol. Evol.* **31**, 215–225 (2016).
63. Barrett, B. J. Equifinality in empirical studies of cultural transmission. *Behav. Process.* (in the press).
64. Biro, D., Sasaki, T. & Portugal, S. J. Bringing a time-depth perspective to collective animal behaviour. *Trends Ecol. Evol.* **31**, 550–562 (2016).
65. Hoppitt W. & Laland K. N. *Social Learning: An Introduction to Mechanisms, Methods, and Models* (Princeton Univ. Press, Princeton, NJ, USA, 2013).
66. Toelch, U., Bruce, M. J., Meeus, M. T. H. & Reader, S. M. Humans copy rapidly increasing choices in a multiarmed bandit problem. *Evol. Hum. Behav.* **31**, 326–333 (2010).
67. Bahrami, B. et al. Optimally interacting minds. *Science* **329**, 1081–1085 (2010).
68. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
69. Hergueux, J. & Jacquemet, N. Social preferences in the online laboratory: a randomized experiment. *Exp. Econ.* **18**, 251–283 (2015).
70. Dandurand, F., ShultzEmail, T. R. & Onishi, K. H. Comparing online and lab methods in a problem-solving experiment. *Behav. Res. Methods* **40**, 428–434 (2008).
71. Becker, J., Brackbill, D. & Centola, D. Network dynamics of social influence in the wisdom of crowds. *Proc. Natl Acad. Sci. USA* **114**, E5070–E5076 (2017).
72. Lorge, I., Fox, D., Davitz, J. & Brenner, M. A survey of studies contrasting the quality of group performance and individual performance, 1920–1957. *Psychol. Bull.* **55**, 337–372 (1958).
73. Hastie, R. & Kameda, T. The robust beauty of majority rules in group decisions. *Psychol. Rev.* **112**, 494–508 (2005).
74. Payzan-Lenestour, E. & Bossaerts, P. Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput. Biol.* **7**, e1001048 (2011).
75. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: an Introduction* (MIT Press, Cambridge, MA, 1998).
76. Ahn, W. Y. et al. Decision-making in stimulant and opiate addicts in protracted abstinence: evidence from computational modeling with pure users. *Front. Psychol.* **5**, 849 (2014).

Acknowledgements

This experiment was supported by The John Templeton Foundation (40128 to K.N.L.) and Suntory Foundation research support (2015-311 to W.T.). The computer simulations and computational model analyses were supported by JSPS overseas research fellowships (H27-11 to W.T.). The phenomenological model analyses were supported by JSPS KAKENHI (grant number 17J01559). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

W.T., A.W. and K.N.L. planned the study and built the computational model. W.T. ran the simulations. W.T. and A.W. made the experimental material, ran the web-based experiment and collected the experimental data. W.T., A.W. and K.N.L. analysed the data and wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41562-018-0518-x>.

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

Correspondence and requests for materials should be addressed to W.T.

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

© The Author(s), under exclusive licence to Springer Nature Limited 2019

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistical parameters

When statistical analyses are reported, confirm that the following items are present in the relevant location (e.g. figure legend, table legend, main text, or Methods section).

n/a | Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- An indication of whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistics including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated
- Clearly defined error bars
State explicitly what error bars represent (e.g. SD, SE, CI)

Our web collection on [statistics for biologists](#) may be useful.

Software and code

Policy information about [availability of computer code](#)

Data collection

Simulations of individual-based models were conducted in Mathematica (version: 10.0.0.0). Experimental data were collected through a browser based online task built by Node.js and socket.io. Code are available on a GitHub repository (<https://github.com/WataruToyokawa/MultiPlayerThreeArmedBanditGame>).

Data analysis

Analyses were conducted in R (3.4.1) and Stan (2.16.2). Code are available on a GitHub repository (<https://github.com/WataruToyokawa/ToyokawaWhalenLaland2018>).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

We included "Code Availability" and "Data Availability" subsections.

Field-specific reporting

Please select the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/authors/policies/ReportingSummary-flat.pdf](https://www.nature.com/authors/policies/ReportingSummary-flat.pdf)

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Computer-based experimental study using a web-based multi-player task. Data are quantitative.
Research sample	A total of 755 subjects (354 females, 377 males, 2 others and 22 unspecified; mean age (1 SD.) = 34.33 (10.9)) participated through Amazon's Mechanical Turk. The online task was only available for individuals whose 'HIT Approval Rate' was greater than or equal to 90% and who live in the US. A rationale behind the US-only sampling was that social learning strategies might vary across different cultures (e.g. America and India), which would confound our focal effects of group size and task uncertainty on social learning.
Sampling strategy	Upon connecting to the experimental game web page, participants were randomly assigned for one of three different conditions (i.e. Low-, Moderate-, and High-uncertainty condition). Sample size for each condition were: N = 113 (Low-uncertainty condition), N = 132 (Moderate-Uncertain condition), and N = 454 (High-uncertain condition). We recruited more participants in the High-uncertainty condition compared to the other two because we expected that larger group sizes would be needed to generate the collective wisdom in noisier environments. Although no sample-size calculation was performed in advance, our parameter recovery test confirmed that the sample size was sufficient for estimation of individual parameters using a hierarchical Bayesian method.
Data collection	<p>Experimental data were collected through a computer-based online task that can be played in a web browser. To minimise the risk of multiple accesses from the same person, we introduced the restriction that a single 'worker ID' associated with participants' AMT accounts, could participate only once in the experiment. We rejected access from the same IP address: If a participant's IP address had already been stored in our database, the participant directly proceeded from the instruction page to the questionnaire page. In that case, 25 cents show-up fee was still paid because it was possible that different persons might use the same IP address.</p> <p>To minimise the risk of opening other browser windows during the task (for example, browsing other websites), we used 'Page Visibility API' (https://developer.mozilla.org/en-US/docs/Web/API/Page_Visibility_API) to track whether the experimental browser window was always active and not hidden by other browser windows or tabs for more than 1 second. If it was detected that the experimental window was in a hidden state, the participant was automatically redirected to the questionnaire page. In that case, 25 cents show-up fee plus a waiting-bonus (if applicable) and a game-bonus earned so far were paid. In the instruction, participants were warned not to open any other browser windows/tabs during the task and were informed that they would not be able to participate in the task if they do so.</p>
Timing	The experimental sessions were conducted between the 11th of November 2015 and the 12th of January 2016.
Data exclusions	We excluded subjects who disconnected to the online task before completing at least the first 30 rounds from our learning model fitting analysis, resulted in 699 subjects (573 subjects entered the group (i.e. $N \geq 2$) condition and 126 entered the solitary (i.e. $N = 1$) condition). A rationale behind the 30-round criterion was that we could not reliably estimate parameter values of our computational model for a subject who played less than 30 rounds.
Non-participation	Fifty-six out of 755 participants disconnected to the online task before completing at least the first 30 rounds. The reasons of these declines might vary, e.g. opening other browser window during the task or bad Internet connection. However, we could not fully detect a reason for each dropping.
Randomization	Participants were allocated randomly to one of three uncertainty conditions. To manipulate the size of each group, we varied the capacity of the "waiting room" from 10 to 30. Because the task was being advertised on the Worker website at AMT for approximately 2 hours, some participants occasionally arrived after the earlier groups had already started. In that case the participant entered the newly opened waiting room which was open for the next 5 minutes. The number of participants arriving declined with time because newly posted alternative tasks were advertised on the top of the task list, which decreased our task's visibility. This meant that a later-starting session tended to begin before reaching maximum room capacity, resulting in the smaller group size. Therefore, the actual size differed between groups.

Reporting for specific materials, systems and methods

Materials & experimental systems

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Unique biological materials
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input type="checkbox"/>	<input checked="" type="checkbox"/> Human research participants

Methods

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics

See above.

Recruitment

Participants were recruited through Amazon's Mechanical Turk. The online task was only available for individuals whose 'HIT Approval Rate' was greater than or equal to 90%. This selection criterion might elicit a self-selection bias toward individuals who have already participated in other behavioural experiments in AMT. Although it is difficult to rule out the possibility that prior experiences in similar tasks might potentially have an influence on subjects' behavioural strategies, that concern is no greater than for any behavioural experiments on humans, particularly as the description of the task given to potential recruits was kept deliberately vague. We believe that the computational modelling approach coupled with higher statistical power due to the large sample size that AMT recruitment affords, compensates for any drawbacks, and that our conclusions were not likely to be altered qualitatively by any such bias.