Social learning strategies regulate the wisdom and madness of interactive crowds

Wataru Toyokawa[1,2,3,*], Andrew Whalen[1], and Kevin N. Laland[1]

[1]*School of Biology, University of St Andrews, Harold Mitchel Building, St Andrews, Fife, KY16 9TH, Scotland*

[2]*Japan Society for the Promotion of Science, Kojimachi, Chiyoda-ku, Tokyo, 102-0083, Japan*

[3]*Department of Evolutionary Studies of Biosystems, School of Advanced Sciences, SOKENDAI (The Graduate University for Advanced Studies), Shonan Village, Hayama, Kanagawa 240-0193, Japan*

[*]*corresponding author: wt25@st-andrews.ac.uk*

Abstract

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-ftting to identify the individual learning strategies exhibited by participants. Challenging tasks elicit greater conformity amongst individuals, with rates of copying increasing with group size, leading to high probabilities of herding amongst 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 longstanding puzzle in the literature.

Keywords:

collective intelligence, herding, social learning, computational modelling, web-based experiment, hierarchical Bayesian approach

Understanding the mechanisms that account for accurate collective decision-making amongst 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 a high performance in solving cognitive problems under uncertainty[6–10]. Although these findings suggest fundamental cognitive benefits of grouping[11], 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 (i.e. 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-maker[16]. 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 crowds effect[21] (but see[22]), whilst potential disadvantages of information transfer are well-recognised 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 collecitve intelligence, but at the same time, increase the risk of an informational cascade[16]. 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, disproportionally strong positive responses to recruitment signals in social insects have been shown to trap the whole colony to exploit a suboptimal, out-dated 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]. On the other hand, 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, contingent upon 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], whilst 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 crowds theory, increasing group size may increase collective accuracy[45–48]. The relative advantage of the 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 to 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 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 in quantitative studies of animal social learning strategies[35;49–51]. In particular, computational modelling allowed us to conduct a parametric description of different information-gathering processes and to estimate the parameter values at an individual-level resolution. This approach allows us to characterise the complex relationship between individual-level decision, learning strategies and collective-level behavioural dynamics.

Second, we added resolution to our analyses by manipulating both task uncertainty and group size in our web-based experiments with adult human subjects, predicting that these factors would induce heavier use of social information in humans, and thereby alter the balance between collective intelligence and the risk of inflexible herding. To do this, we focused on human groups exposed to a simple gambling task called a multi-player 'multi-armed bandit', where both asocial and social sources of information were available[35;51;52]. Through development of an interactive, web-based collective decision-making task, and use of hierarchical Bayesian statistical methods in fitting our computational model to the experimental data, we identify the individual-level learning strategies of participants as well as quantify variation in different learning parameters, allowing us to conduct an informed exploration of the population-level outcomes. The results provide clear evidence that the collective behavioural dynamics can be predicted with knowledge of human social learning strategies.

Below, we firstly deploy agent-based simulation to illustrate how the model parameters relating to social learning can in principle affect the collective-level behavioural dynamics. The simulation provides us with precise, quantitative predictions concerning the complex relationship between individual behaviour and group dynamics. Second, we present the findings of a multi-player web-based experiment with human participants that utilises the gambling task framework. Applying a hierarchical Bayesian statistical method, we estimated the model's parameters for each of 699 different individuals, allowing us to (*i*) examine whether and, if so, how social information use is affected by different group size and task uncertainty, and (*ii*) whether and how social-information use affects both collective intelligence and the risk of maladaptive herding.

1 Results

1.1 The relationship between social learning and the collective behaviour

Figure 1 shows the relationship between the average decision accuracy and individual-level social information use obtained from our individual-based model simulations, highlighting the trade-off between accuracy and flexibility of collective decision-making. When the mean *conformity exponent* is small (i.e. $\bar{\theta} = (\sum_i \theta_i)/individuals = 1$), large groups are able to recover the decision accuracy quickly as do small groups after the location of the optimal option has been switched, whereas overall improvement by increasing group size in decision accuracy is subtle when the average *social learning weight* is also small (i.e. $\bar{\sigma} = (\sum_i \sum_t \sigma_{i,t})/(individuals \times rounds) = 0.3$; Figure 1A and 1C). On the other hand, when both the conformity exponent $\bar{\theta}$ and the social learning weight $\bar{\sigma}$ are large, average performance is no longer monotonically improving with increasing group size, and it is under these circumstances that the strong herding effect becomes prominent (Figure 1D). Although the high conformity bias with low social learning weight makes large groups more accurate before the environment changes, larger groups are less flexible in performance recovery (Figure 1C). The patten is robust for other parameter regions (Supplementary Supplementary Figure 2).

Figure 2C and 2D indicate that when both $\bar{\theta}$ and $\bar{\sigma}$ are large the collective choices converged either on the good option or on one of the poor options almost randomly, regardless of the option's quality, and that once individuals start converging on an option the population gets stuck. As a result, the distribution of the groups' average performance over the replications becomes a bimodal 'U-shape'. Interestingly, however, the maladaptive herding effect remains relatively weak in smaller groups (see Figure 1D; the dotted line). This is because the majority of individuals in smaller groups (i.e. 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 (e.g. six out of ten).

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

1.2 Collective performance of human participants

Figure 3A shows behavioural dynamics of human participants in different group sizes and different task uncertainty conditions (see Supplementary Supplementary Figure 3 for each group's behaviour). The average decision performance of collectives (i.e. group size ≥ 2) exceeded that of solitary individuals (i.e. group size $= 1$) in the Moderate-uncertainty condition (i.e. the 95% Bayesian CI of ξ_t exceeds 0 at regions $t \in [9, 40]$ and $[67, 70]$; Figure 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 the 95% Bayesian CI of $\omega_2 = 0.67$ [0.44, 0.91]; Table 1), suggesting that the average performance of large groups (e.g. $12 \leq$ group size ≤ 16) were better, and hence more flexible, than smaller groups and solitaries (Figure 3A). On the other hand, in the Moderate-uncertainty condition, the average performance of the collectives dropped below that of the solitaries after the environmental change (i.e. $\xi_t < 0$ at a region $t \in [42, 45]$; Figure 3B). Also, the main effect of group size was negative in the post-change period (mean and the 95% Bayesian CI of ω_2 = -0.26 [-0.44, -0.11]; Table 1), suggesting that larger groups were more likely to get stuck in the out-dated option in the Moderate-uncertainty condition. In the High-uncertainty condition, the main effect of

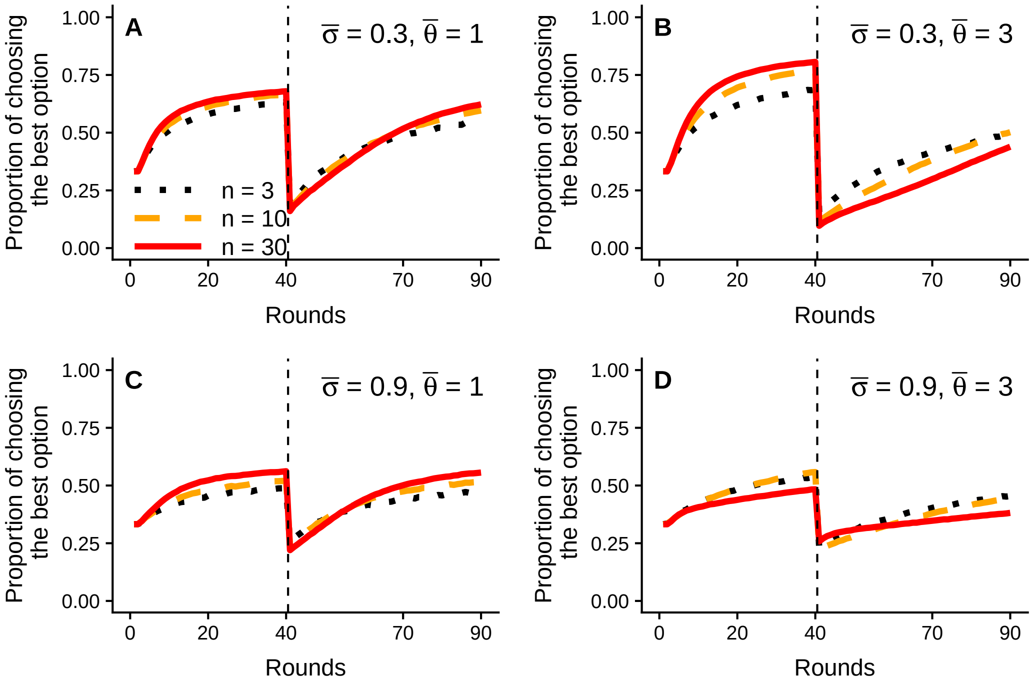

Figure 1: Findings of the individual-based model showing the effects of social information use on the average decision accuracy over replications. The x-axis gives the round and y-axis gives the proportion of individuals expected to choose the optimal slot (i.e. decision accuracy) averaged over all replications. The vertical dashed line indicates the timing of environmental (i.e. payoff) change (at $t = 41$). Different group sizes are shown by different styles (black (dotted): $n = 3$, orange (dashed): $n = 10$, red (solid): $n = 30$). We set the average slopes for the *social learning weight* to be equal to zero for the sake of simplicity; namely, $\mu_\delta = 0$. Other free parameter values (i.e. μ_α, $\mu_{\beta_0^*}$, μ_ϵ,ν_α, $\nu_{\beta_0^*}$, ν_ϵ, ν_σ, ν_δ and ν_θ) are best approximates to the experimental fitted values (see Table 2 and Supplementary Table 1).

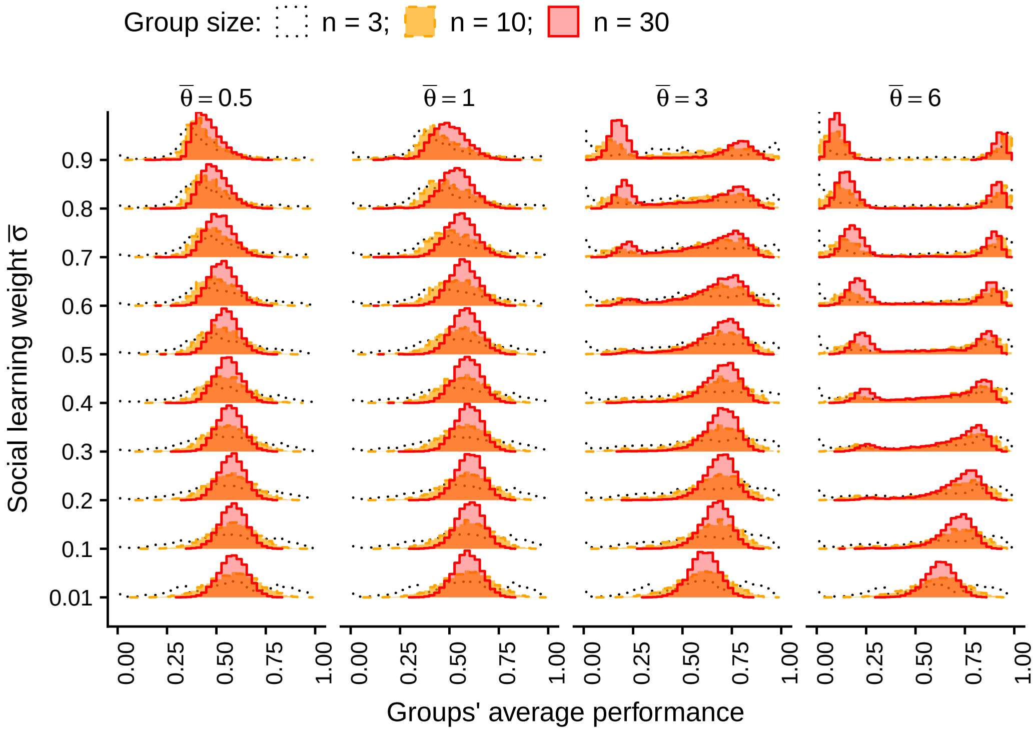

Figure 2: Results from the individual-based model simulations showing the distribution of each group's mean accuracy before environmental change ($t \leq 40$). The x-axis gives the mean decision accuracy over the first 40 rounds (i.e. the environment 1) for each replication. Different group sizes are shown by different styles (black (dotted): $n = 3$, orange (dashed): $n = 10$, red (solid): $n = 30$). The other free parameter values are the same as in Figure 1.

group size was positive in the prior-change period and negative post-change (mean and the 95% Bayesian CIs are $\omega_1 = 0.07$ [0.00, 0.15] and $\omega_2 = -0.10$ [-0.17, -0.02]; Table 1), although the effect size was too small to differentiate performances of different group sizes visually (Figure 3A). Using monetary earnings as an outcome variable of decision performance did not change our conclusions qualitatively (supporting Supplementary Figure 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.

1.3 Estimation of human social information use

Using posterior estimation values obtained by the hierarchical Bayesian model fitting method (Table 2), we were able to categorise the participants as deploying one of three different learning strategies based on their fitted conformity exponent values; namely, the 'positive frequency-dependent copying' strategy ($\theta_i \gg 0$), the 'negative-frequency dependent copying' strategy ($\theta_i \ll 0$) and the 'random choice' strategy ($\theta_i \approx 0$). 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 Figure 5). Some individuals whose 'true' conformity exponent fell between zero and one would have been categorised as exhibiting a random choice strategy (Supplementary Figure 7). Individuals identified as exhibiting a positive frequency-dependent copiers were mainly those whose conformity exponent was larger than one ($\theta_i > 1$).

Figure 4A show the estimated frequencies of different learning strategies. Generally speak-

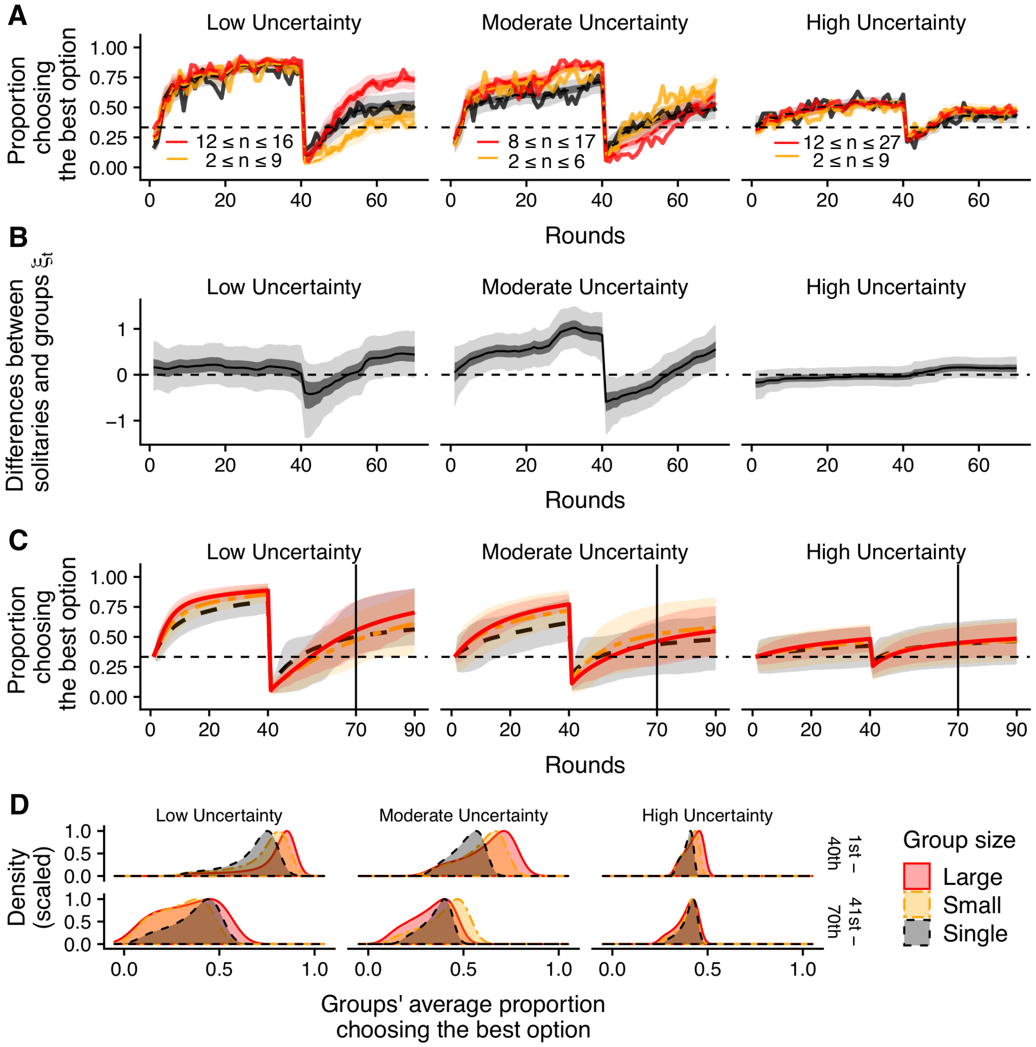

Figure 3: Time evolutions and distributions of decision performance for each condition. A: The average decision accuracies of the experimental participants (red: large groups, orange: small groups, dark grey: lone individuals). All individual performances were averaged within the same size category (solid lines). The light-shaded areas, dark-shaded areas, and dashed curves show the 95%, 50%, and median Bayesian credible intervals of the phenomenological, time-series logistic regression. Sample sizes for large, small, and lone groups are: $N = 43$, $N = 44$ and $N = 38$ for the Low-uncertainty condition; $N = 52$, $N = 56$ and $N = 37$ for the Moderate-uncertainty condition; and $N = 259$, $N = 168$ and $N = 58$ 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 shaded areas are the Bayesian CIs and solid curves are the median. **C, D:** Change and distribution in average decision accuracy of the individual-based post-hoc simulations of the learning process model using the experimentally fit parameter values. **C:** All replications were averaged within the same size category (solid lines). The shaded areas give the 50% quantiles. The experimental horizon (i.e. $t = 70$) is indicated by the vertical line. **D:** Performance was averaged within prior- and post-change periods for each replication for each group sizes category.

ing, 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 GLMM predicting the probability to use the positive frequency-dependent copying strategy is above zero, [1.05, 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 is below zero, [-1.88, -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 × uncertainty interaction = [-1.46, -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 (i.e. $\bar{\sigma}_i = (\sum_t \sigma_{i,t})/(\text{total rounds})$) increased with group size (the 95% Bayesian CI = [0.15, 0.93]; Figure 4B; Supplementary Table 5), and decreased with uncertainty (the 95% Bayesian CI = [-0.98, -0.22]), and age of subject (the 95% Bayesian CI = [-0.36, -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 confirmed (Supplementary Table 6; Supplementary Figure 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 (the 95%

Bayesian CI = [0.38, 1.41]), but we found no significant effects of group size, age, gender or interactions (Figure 4C; Supplementary Table 7). These results were qualitatively unchanged when we focused only on the positive frequency-dependent copying individuals (Supplementary Table 8; Supplementary Figure 6).

We observed extensive individual variation in social information use. The greater the task's uncertainty, the larger were individual variances 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, 1.62] (Supplementary Table 5) and for θ_i was [1.07, 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, 1.80] (Supplementary Table 6) and [0.71, 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 visualised in Figure 5. The social learning weights generally decreased with experimental round. However, some individuals in the Moderate- and the High-uncertain conditions accelerated rather than decreased their reliance on social learning over time. Interestingly, those accelerating individuals tended to have a larger conformity exponent (Supplementary Figure 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 Figure 10), although the fitting of $\theta_{i,t}$ in the alternative model was relatively unreliable (Supplementary Figure 9). These findings suggest that conformists tended to use asocial learning at the outset (i.e. exploration asocially) but increasingly started to conform as the task proceeded (i.e. exploitation socially).

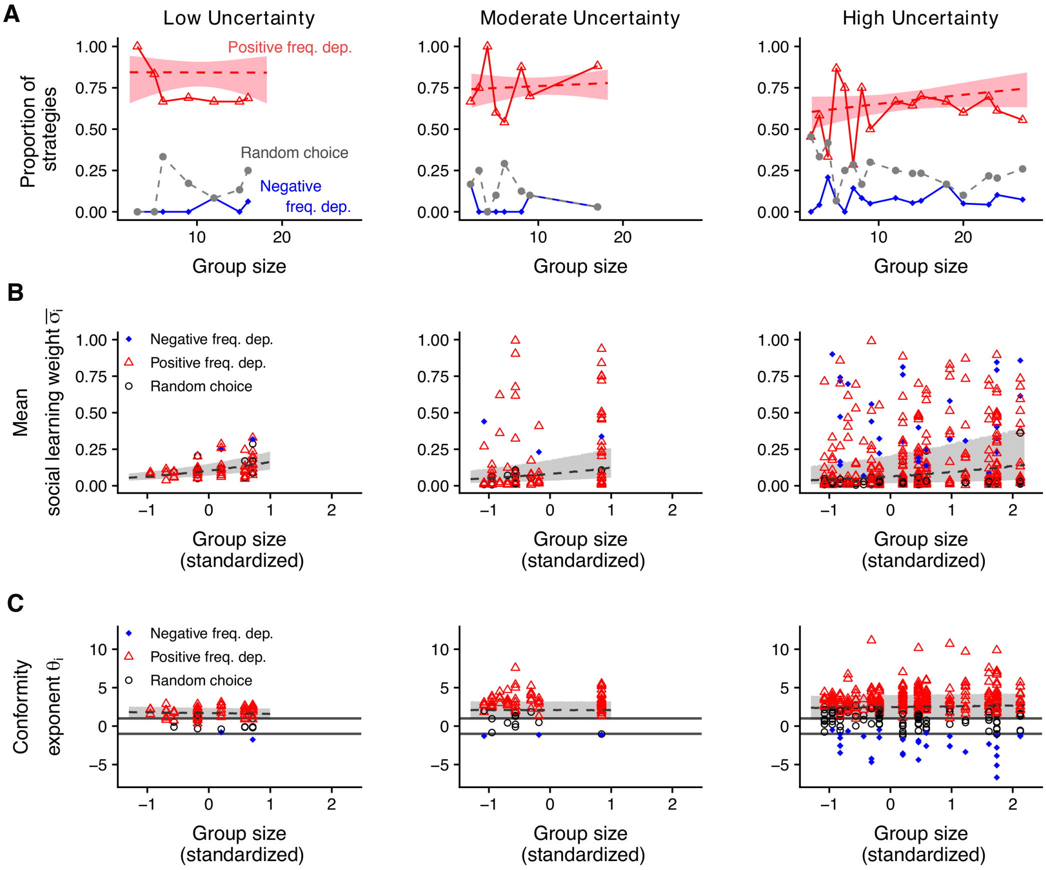

Figure 4: Model fitting for the three different task's uncertain conditions (the Low-, Moderate- and High-uncertainty) and the different group size. Three different learning strategies are shown in different styles (red-triangle: positive frequency-dependent learning, blue-circle: negative frequency-dependent 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) Estimated social learning weight, and (C) estimated conformity exponent, 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 a region $-1 < \theta_i < 1$. Sample sizes for Negative Frequency Dependent, Positive Frequency Dependent, and Random Choice strategies are: $N = 2$, $N = 61$ and $N = 14$ for the Low-uncertainty condition; $N = 3$, $N = 80$ and $N = 15$ for the Moderate-uncertainty condition; and $N = 32$, $N = 260$ and $N = 106$ for the High-uncertainty condition, respectively.

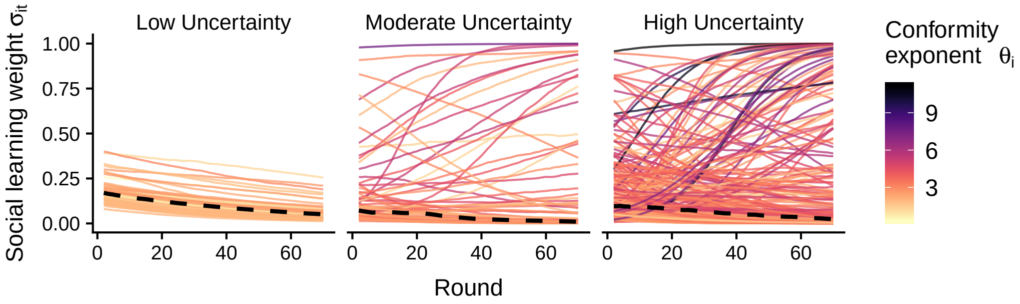

Figure 5: Change in fitted values (i.e. median of the Bayesian posterior distribution) of the social learning weight $\sigma_{i,t}$ with time for each Positive Frequency Dependent individual, for each level of task uncertainty. Thick dashed lines are the median values of $\sigma_{i,t}$ across the subjects for each uncertainty condition. Individual conformity exponent values θ_i are shown in different colours (higher θ_i is darker). Sample size for each task uncertainty condition is: $N = 61$ (Low-uncertainty), $N = 80$ (Moderate-uncertainty) and $N = 260$ (High-uncertainty).

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 found for random choice individuals (Supplementary Figure 5). Individuals deploying a random choice strategy exhibited a $\sigma_{i,t}$ that approached to 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 Table 9, Supplementary Table 10 and Supplementary Figure 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-uncertain condition. Conversely, group size had a positive effect on the mean social learning weight, but did not affect conformity.

1.4 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 Figure 3). The trajectories of the simulated dynamics recovered nicely the pattern observed in the experiment (Figure 3A and 3C), suggesting that the strategic changes in the individual-level social information use (Figure 4) could explain the collective-level behavioural pattern.

Figure 3D shows that larger groups are more likely to make accurate decisions than are both small groups and solitaries in the period prior to change across all uncertainty conditions, suggesting collective intelligence was operating. In the post-change period, however, 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 (i.e. $\bar{\theta} = 1.65$), large groups performed better than did small groups over 59.5% of total 10,000 repetitions. However, in the Moderate-uncertainty condition, where we found that participants were most likely to have strong positive frequency dependence ($\bar{\theta} = 3.00$, c.f. 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 did the large groups over 69.5% of repetitions (Figure 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 did 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 lefthand side to the peak of the performance distribution in the post-change period; Figure 3D). This might be due to the lower softmax exploration rates among social learners in the Low-uncertainty condition (i.e. both $\mu_{\beta_0^*}$ and μ_ϵ were large; Table 2): the whole population gets stuck because all individuals are very exploitative on their past experience.

2 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 a 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 (Figure 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 (i.e. $\bar{\theta} \approx 1.65$), larger groups were able to be more accurate than smaller groups while retaining flexibility in their decision-making[9], 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 (i.e. $\bar{\theta} \approx 3.0$ and 2.7, respectively), larger groups performed better prior to environmental change but were vulnerable to getting stuck with an out-dated 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 $\bar{\sigma}_i \approx 0.3$ (Figure 4B; Supplementary Figure 5; Supplementary Figure 6) and it decreased on average as the task proceeded (Figure 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 (Figure 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 (Figure 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_i, irrespective of the options' quality. Negative-frequency dependent copying individuals could also be highly exploratory. These individuals tended to avoid choosing an option upon which other people had converged and would explore the other two 'unpopular' options. Interestingly, in the High-uncertain condition the mean social learning weights of the negative-frequency dependent copying individuals ($\bar{\sigma}_i \approx 0.5$) were larger than that of the other two strategies ($\bar{\sigma}_i \approx 0.1$, Supplementary Figure 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 (Figure 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 different learning strategies[63–65]. For example, very exploitative asocial reinforcement learners (i.e. for whom 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-uncertain 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, are 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 to theory[57;58]. Indeed, the post-hoc simulation sometimes failed to recover the observed behavioural trajectories. In particular, experimental groups with $n = 12$, $n = 16$, and one group in $n = 9$, in the Low-uncertainty condition performed very well, exceeding the 95% CIs of the post-hoc simulation after the environmental change (Supplementary Figure 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, e.g. ‘copy-rapidly-increasing-option’ strategy[66] 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 processes[68] 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 suggest 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 performance[74]. In such temporally dynamic environments, decision-making and learning are coordinated to affect future behavioural outcomes recursively[75]. 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 new light on when groups of individuals will exhibit the wisdom of the crowds and when inflexible herding. Our analysis implies that herding is most likely amongst individuals in large groups exposed to challenging tasks. That is because challenging tasks lead to greater uncertainty and thereby elicit greater conformist learning amongst individuals, whilst rates of copying increase with group size. Difficult tasks, by definition, render identification of the optimal behavior harder, allowing groups sometimes to 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.

3 Material and methods

3.1 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 SD.) = 34.33 (10.9)) 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 to 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, resulted in 699 subjects (573 subjects entered the group (i.e. $N \geq 2$) and 126 entered the solitary (i.e. $N = 1$) condition). The task was only available for individuals who had greater than or equal to 90% HIT approval rate and who accessed 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.

3.2 Design of the experimental manipulations

The three uncertainty conditions were: Low-uncertainty condition (differences between mean payoffs were 1.264), Moderate-uncertainty condition (differences between mean payoffs were 0.742) and High-uncertainty condition (differences between mean payoffs were 0.3). The mean payoff associated with the ‘excellent’ slot in all three conditions was fixed to 3.1 cents (Supplementary Figure 1). Each task uncertainty condition was randomly assigned for each different HIT session, and participants were allowed to participate in one HIT only. Sample size after the data exclusion for each uncertainty condition was: $N = 113$ (Low-uncertainty condition), $N = 132$ (Moderate-Uncertain condition), and $N = 454$ (High-uncertain condition). We assigned more sessions to 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.

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 HITs 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 (Supplementary Figure 3, Supplementary Table 1). Data collection and analysis were not performed blind to the conditions of the experiments.

3.3 The multi-player 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 Figure 1, Video 1, for detail see Supplementary Method). Participants played the task for 70 rounds. The slots paid off money noisily (in the US cents), varying around two different means during the first 40 rounds such that there was one 'good' slot and two other options giving poorer average returns. From the round 41st, however, one of the 'poor' slots abruptly increased its mean payoff to become 'excellent' (i.e. 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 asocial 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, whilst in the task with the greatest uncertainty the distribution of payoffs overlapped considerably (Supplementary Figure 1).

3.4 The 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 $P_{i,t}(m)$ that is a weighted average of social and asocial influences:

$$P_{i,t}(m) = \sigma_{i,t} \times \mathit{Social\ influence}_{i,m,t} + (1 - \sigma_{i,t}) \times \mathit{Asocial\ influence}_{i,m,t}, \tag{1}$$

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]:

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

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, higher social influence is given to an option which was chosen by more individuals (i.e. positive frequency bias). When this exponent is larger than zero ($\theta_i > 0$), higher social influence is afforded to an option chosen by more individuals (i.e. positive frequency bias), with conformity bias arising when $\theta_i > 1$, such that disproportionally more social influence is given to the most common option[28]. When $\theta_i < 0$, on the other hand, higher social influence is afforded to the option that fewest individuals chose in the preceding round $t-1$ (i.e. 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 (i.e. $F_{t-1} = 0$) could provide the highest social influence when $\theta_i < 0$. Note, there is no social influence when $\theta_i = 0$ because in this case the 'social influence' favours an uniformly random choice, i.e., $S_{i,t}(m) = f_m^0/(f_1^0+f_2^0+f_3^0) = 1/3$, independent of the social frequency distribution. Note also that, in the first round $t = 1$, we assumed that the choice is only determined by the asocial softmax function because there is no social information available.

For the asocial influence, we used a standard reinforcement learning with 'softmax' choice rule[75], widely applied in human social learning studies e.g.[35;51;55]. An individual i updates the estimated average reward associated with an option m at round t, namely Q-value ($Q_{i,t}(m)$), according to the Rescorla-Wagner rule as follows:

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

where α_i ($0 \le \alpha_i \le 1$) 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. $\mathbb{1}(m, m_{i,t})$ is the binary action-indicator function of individual i, given by

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

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

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

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

where $\beta_{i,t}$, called *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 (i.e. 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 (i.e. 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-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 $\epsilon_i t$ 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:

$$\sigma_{i,t} = \frac{1}{1 + \exp(-(\sigma^*_{i,0} + \delta_i t/70))}. \tag{6}$$

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 group size is one (i.e. when an individual participated in the task alone and/or when $\sum_{k \in options} F_{t-1}(k) = 0$).

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 different patterns of learning in this social learning task[63]. The parameter recovery test confirmed that we were able to differentiate such processes under these assumptions (Supplementary Figure 7, Supplementary Figure 8). While we also considered the possibility of the conformity exponent being time-dependent (i.e. $\theta_{i,t} = \theta^{*}_{i,0} + \gamma_i t/70$), the parameter recovery test suggested that the individual slope parameter γ_i was not reliably recovered (Supplementary Figure 9), and hence we concentrated our analysis on the time-independent θ_i model. We confirmed that instead using the alternative model where both social learning parameters were time-dependent (i.e. $\sigma_{i,t}$ and $\theta_{i,t}$) did not qualitatively change our results (Supplementary Figure 10).

One concern might be the asymmetry between the asocial softmax influence which takes many prior experiences into account (depending upon a learning rate) and the social influence referring only the most recent frequency information F_{t-1}. The choice frequency appeared at round t is the most reliable social information, compared to the past frequencies, because it could be the most ‘updated’ information as long as the other individuals have made informed decisions in their best knowledge. In contrast, option’s reward just obtained at $t - 1$, which was independently and randomly drawn from a probability distribution, is less reliable than accumulated Q-values taking past rewards into account. 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_i, $\beta^*_{i,0}$, ϵ_i, $\sigma^*_{i,0}$, δ_i, and θ_i. To fit the model, we used a hierarchical Bayesian method (HBM), estimating the global means (μ_α, $\mu_{\beta^*_0}$, μ_ϵ, $\mu_{\sigma^*_0}$, μ_δ, and μ_θ) and the global variations (v_α, $v_{\beta^*_0}$, v_ϵ, $v_{\sigma^*_0}$, v_δ, and v_θ) for each of the three experimental conditions (i.e. the Low-, Moderate- and High-uncertain condition), which govern overall distributions of individual parameter values. It has become recognised that the HBM can provide more robust and reliable parameter estimation than conventional maximum likelihood point estimation in complex cognitive models[76], a conclusion with which our parameter recovery test agreed (Supplementary Figure 7, Supplementary Figure 8).

3.5 Agent-based model simulation

We ran a series of individual-based model simulations assuming that a group of individuals play our three-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 ($n \in \{3, 10, 30\}$), the mean social learning weight ($\bar{\sigma} \in \{0.01, 0.1, 0.2, 0.3, ..., 0.9\}$) and the mean conformity exponent ($\bar{\theta} \in \{0.5, 1, 3, 6\}$), running 10,000 replications for each of the possible parameter × group size combinations. As for the other parameter values (e.g. the asocial reinforcement learning parameters; α, β^*_0, ϵ), here we used the experimentally fitted global means (Table 2 and Supplementary Table 3). Relaxation of this assumption (i.e. using a different set of asocial learning parameters) does not qualitatively change our story (Supplementary Figure 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 were different between individuals even within the same social group.

3.6 Generalised 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 valuable was whether the participant chose the best option (1) or not (0). The model includes fixed effects of grouping ξ, standardised group size ω, and an intercept with a random effect of individuals $\mu + \rho_i$. We assumed that the intercept and the effect of grouping change from round to round, as a random walk process. For the effect of group size we considered the effect of the 1st environment $1 \leq t \leq 40$ and that of the 2nd environment, namely, ω_1 and ω_2, separately.

To examine whether increasing group size and increasing 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 groups. The dependent valuable was whether the participant used the positive frequency-dependent copying (1) or not (0). The model includes fixed effects of group size (standardised), task uncertainty (0: Low, 0.5: Moderate, 1: High), age (standardised), 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 (standardised), task uncertainty (0: Low, 0.5: Moderate, 1: High), age (standardised), 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 upon task uncertainty because we had found in the computational model-fitting result that the fitted global variance parameters (i.e. $\nu_{\sigma_0^*}$, ν_δ and ν_θ) were larger in more uncertain conditions (Supplementary Table 2).

3.7 Statistical analysis

We used a hierarchical Bayesian method (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 ensures 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 4 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.

3.7.1 Parameter recovery test

To check the validity of our model-fitting method, we conducted a ‘parameter recovery test’ so as 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, and then re-fitted the model with this synthetic data using the same procedure. The parameter recovery test showed that the all true global parameter values were fallen into the 95% Bayesian credible interval (Supplementary Figure 7), and at least 93% of the true individual parameter values were correctly recovered (i.e. 96% of α_i, 93% of $\beta^*_{i,0}$, 95% of ϵ_i, 97% of $\sigma^*_{i,0}$, 96% of δ_i and 97% of θ_i values were fallen into the 95% Bayesian CI. Supplementary Figure 7).

3.7.2 Categorisation 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 ($\theta_{lower} > 0$), below zero ($\theta_{upper} < 0$) or including zero ($\theta_{lower} \leq 0 \leq \theta_{upper}$), she/he was categorised 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 categorisation instead of using the more conservative 95% CI because the latter would cause much higher rates of ‘false negatives’, by which an individual who applied either a positive frequency-dependent copying or a negative-frequency dependent copying strategy was falsely labelled as an asocial random choice individual (Supplementary Figure 7). Four hundred agents out of 572 ($\approx$ 70%) were falsely categorised as a random choice learner in the recovery test when we used the 95% criterion (Supplementary Figure 7). On the other hand, the 50% CI criterion seemed to be much better in terms of the false negative rate which was only 18.5% (i.e. 106 agents), although it might be slightly worse in terms of ‘false positives’: Thirty-seven agents (6.5%) were falsely labelled as either a positive frequency-dependent copier or a negative-frequency dependent copier by the 50% CI, whereas the false positive rate of the 95% CI was only 0.2% (Supplementary Figure 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.

3.7.3 The post-hoc model simulation

So as to evaluate how accurately our model can generate observed decision pattern in our task setting, we ran a series of individual-based model simulation using the fitted individual parameter values (i.e. 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 uncertain condition to an simulated agent, until the number of agents reaches the simulated group size. We allowed duplicate choice of experimental subject 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, task uncertainty.

3.8 Data availability

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

3.9 Code availability

The browser based online task was built by Node.js (`https://nodejs.org/en/`) and socket.io (`https://socket.io`), and the code are available on a GitHub repository (`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`), both are available on an online repository (`https://github.com/WataruToyokawa/ToyokawaWhalenLaland2018`).

4 Competing interest

The authors declare no competing interests.

5 Authors' contributions

WT, AW and KNL planned the study and built the computational model. WT ran simulations. WT and AW made the experimental material, ran the web-base experiment, and collected the experimental data. WT, AW and KNL analysed the data and wrote the manuscript.

6 Acknowledgements

The experiment was supported by The John Templeton Foundation (KNL; 40128) and Suntory Foundation Research Support (WT; 2015-311). The computer simulations and computational model analyses were supported by JSPS Overseas Research Fellowships (WT; H27-11). 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.

Table 1: The mean and the 95% Bayesian credible intervals of the posterior for the group size effect at the phenomenological logistic model

	Low Uncertainty		Moderate Uncertainty		High Uncertainty	
ω_1	0.08	[-0.15, 0.33]	0.10	[-0.06, 0.26]	0.07	[0.00, 0.15]
ω_2	0.67	[0.44, 0.91]	-0.26	[-0.44, -0.11]	-0.10	[-0.17, -0.02]

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

Table 2: The mean and the 95% Bayesian credible intervals of the posterior global means for the parameter values. The number of participants (N) for each experimental condition are also shown.

	Groups			Solitary individuals		
Uncertainty:	Low	Moderate	High	Low	Moderate	High
μ_{α^*} (learning rate)	0.99	0.90	0.61	0.85	-0.17	0.46
	[0.34, 1.73]	[0.43, 1.44]	[0.21, 1.03]	[-0.07, 1.95]	[-1.27, 0.89]	[-0.39, 1.36]
$\mu_{\beta_0^*}$ (inv. temp.)	1.84	1.68	1.38	1.10	1.44	0.85
	[1.15, 2.70]	[1.25, 2.18]	[1.16, 1.62]	[0.69, 1.54]	[0.80, 2.07]	[0.46, 1.22]
μ_{ϵ} (inv. temp.)	3.70	3.01	2.97	2.39	2.81	2.27
	[1.98, 5.71]	[1.88, 4.27]	[2.37, 3.60]	[1.46, 3.53]	[1.64, 4.07]	[1.40, 3.31]
$\mu_{\sigma_0^*}$ (soc. wight)	-1.55	-2.37	-2.16	–	–	–
	[-2.71, -0.71]	[-4.12, -1.01]	[-2.81, -1.63]	–	–	–
μ_{δ} (soc. wight)	-1.39	-1.55	-1.87	–	–	–
	[-2.66, -0.03]	[-4.29, 0.91]	[-3.04, -0.81]	–	–	–
μ_{θ} (conformity coeff.)	1.65	3.00	2.67	–	–	–
	[0.83, 2.82]	[1.57, 4.85]	[1.80, 3.73]	–	–	–
N	77	98	398	36	34	56

References

[1] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY, 1999. ISBN 0-19-513159-2.

[2] Scott Camazine, Jean-Louis Deneubourg, Nigel R. Franks, James Sneyd, Guy Theraulaz, and Eric Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001. ISBN 9780691116242.

[3] Jens Krause, Graeme D Ruxton, and Stefan Krause. Swarm intelligence in animals and humans. *Trends in ecology & evolution*, 25(1):28–34, jan 2010. ISSN 0169-5347. doi: 10.1016/j.tree.2009.06.016. URL `http://www.ncbi.nlm.nih.gov/pubmed/19735961`.

[4] Thomas D. Seeley. *The Wisdom of the Hive*. Harvard University Press, Cambridge, MA, 1995. ISBN 9780674953765.

[5] David J. T. Sumpter. *Collective Animal Behavior*. Princeton University Press, Princeton, NJ, 2010. ISBN 9781400837106.

[6] Andrew J. King and Cédric Sueur. Where Next? Group Coordination and Collective Decision Making by Primates. *International Journal of Primatology*, 32(6):1245–1267, jun 2011. ISSN 0164-0291. doi: 10.1007/s10764-011-9526-7. URL `http://www.springerlink.com/index/10.1007/s10764-011-9526-7`.

[7] J. Morand-Ferron and J. L. Quinn. Larger groups of passerines are more efficient problem solvers in the wild. *Proceedings of the National Academy of Sciences*, 108(38):15898–15903, 2011. ISSN 0027-8424. doi: 10.1073/pnas.1111560108.

[8] Takao Sasaki and Dora Biro. Cumulative culture can emerge from collective intelligence in animal groups. *Nature Communications*, 8:1–6, 2017. ISSN 20411723. doi: 10.1038/ncomms15049. URL `http://dx.doi.org/10.1038/ncomms15049`.

[9] Zachary Shaffer, Takao Sasaki, and Stephen C. Pratt. Linear recruitment leads to allocation and flexibility in collective foraging by ants. *Animal Behaviour*, 86(5):967–75, sep 2013. ISSN 00033472. doi: 10.1016/j.anbehav.2013.08.014. URL `http://linkinghub.elsevier.com/retrieve/pii/S0003347213003825`.

[10] Chris R. Reid and Tanya Latty. Collective behaviour and swarm intelligence in slime moulds. *FEMS Microbiology Reviews*, 40(6):798–806, 2016. ISSN 15746976. doi: 10.1093/femsre/fuw033.

[11] Jens Krause and Graeme D. Ruxton. *Living in groups*. Oxford University Press, Oxford, New York, 2002. ISBN 978-0-19-850818-2.

[12] Charles Mackay. *Extraordinary Popular Delusions and the Madness of Crowds*. Richard Bentley, London, UK, 1841.

[13] Tatsuya Kameda and Reid Hastie. Herd Behavior. In *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*, pages 1–14. Wiley Online Library, 2015. ISBN 9781118900772.

[14] Gustav Le Bon. *The crowd: study of the popular mind*. Unwin, London, UK, 4th edition, 1896.

[15] Ramsey M Raafat, Nick Chater, and Chris Frith. Herding in humans. *Trends in cognitive*

sciences, 13(10):420–8, oct 2009. ISSN 1364-6613. doi: 10.1016/j.tics.2009.08.002. URL http://www.ncbi.nlm.nih.gov/pubmed/19748818.

[16] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100 (5):992, 1992. ISSN 0022-3808. doi: 10.1086/261849.

[17] V. V. Chari and Patrick J. Kehoe. Financial crises as herds: Overturning the critiques. *Journal of Economic Theory*, 119(1 SPEC. ISS.):128–150, 2004. ISSN 00220531. doi: 10.1016/S0022-0531(03)00225-4.

[18] Irving L. Janis. *Victims of groupthink: A psychological study of foreign policy*. Houghton Mifflin Company, Boston, NY, 1972.

[19] Lev Muchnik, Sinan Aral, and Sean J Taylor. Social influence bias: a randomized experiment. *Science (New York, N.Y.)*, 341(6146):647–51, aug 2013. ISSN 1095-9203. doi: 10.1126/science.1240466. URL http://www.ncbi.nlm.nih.gov/pubmed/23929980.

[20] Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. Experimental Study of Inequality and Cultural Market. *Science*, 311(February), 2006.

[21] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences of the United States of America*, 108(22):9020–5, may 2011. ISSN 1091-6490. doi: 10.1073/pnas.1008636108. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3107299{&}tool=pmcentrez{&}rendertype=abstract.

[22] Bertrand Jayles, Hye-rin Kim, Ramón Escobedo, Stéphane Cezera, Adrien Blanchet, Tatsuya Kameda, Clément Sire, and Guy Theraulaz. How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):201703695, 2017. ISSN 0027-8424. doi: 10.1073/pnas.1703695114. URL http://www.pnas.org/lookup/doi/10.1073/pnas.1703695114.

[23] Luc-Alain Giraldeau, Thomas J Valone, and Jennifer J Templeton. Potential disadvantages of using socially acquired information. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 357(1427):1559–66, nov 2002. ISSN 0962-8436. doi: 10.1098/rstb.2002.1065. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1693065{&}tool=pmcentrez{&}rendertype=abstract.

[24] Claire Detrain and Jean Louis Deneubourg. *Collective Decision-Making and Foraging Patterns in Ants and Honeybees*, volume 35. Elsevier, 2008. ISBN 9780123743299. doi: 10.1016/S0065-2806(08)00002-7. URL http://dx.doi.org/10.1016/S0065-2806(08)00002-7.

[25] Christian List, Christian Elsholtz, and Thomas D Seeley. Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 364(1518):755–762, 2009. ISSN 0962-8436. doi: 10.1098/rstb.2008.0277.

[26] J. L. Deneubourg, S. Aron, S. Goss, and J. M. Pasteels. The self-organizing exploratory pat-

tern of the argentine ant. *Journal of Insect Behavior*, 3(2):159–168, 1990. ISSN 08927553. doi: 10.1007/BF01417909.

[27] R Beckers, J L D Deneubourg, S Goss, and J M Pasteels. Collective decision making through food recruitment. *Insectes Sociaux*, 37:258–267, 1990.

[28] Robert Boyd and Peter J. Richerson. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago, IL, 1985.

[29] Peter J. Richerson and Robert Boyd. *Not By Genes Alone*. University of Chicago Press, Chicago, IL, 2005.

[30] Marcus W. Feldman, Kenichi Aoki, and Jochen Kumm. Individual Versus Social Learning: Evolutionary Analysis in a Fluctuating Environment. *Anthropological Science*, 104(3):209–231, 1996. ISSN 0918-7960. doi: 10.1537/ase.104.209. URL `http://www.santafe.edu/media/workingpapers/96-05-030.pdf{%}5Cnpapers2://publication/uuid/2408B672-F4CE-4EE0-A7B0-AFAC56879407`.

[31] Kevin N. Laland. Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14, 2004. ISSN 0090-4996. doi: 10.3758/BF03196002.

[32] Tatsuya Kameda and Daisuke Nakanishi. CostâĂŞbenefit analysis of social/cultural learning in a nonstationary uncertain environment. *Evolution and Human Behavior*, 23(5): 373–393, 2002. ISSN 10905138. doi: 10.1016/S1090-5138(02)00101-0. URL `http://www.sciencedirect.com/science/article/pii/S1090513802001010`.

[33] Rachel L. Kendal, Isabelle Coolen, and Kevin N. Laland. The role of conformity in foraging

when personal and social information conflict. *Behavioral Ecology*, 15(2):269–277, 2004. ISSN 10452249. doi: 10.1093/beheco/arh008.

[34] T J H Morgan, L E Rendell, M Ehn, W Hoppitt, and K N Laland. The evolutionary basis of human social learning. *Proceedings. Biological sciences / The Royal Society*, 279(1729): 653–62, feb 2012. ISSN 1471-2954. doi: 10.1098/rspb.2011.1172. URL `http://www.ncbi.nlm.nih.gov/pubmed/21795267`.

[35] Wataru Toyokawa, Yoshimatsu Saito, and Tatsuya Kameda. Individual differences in learning behaviours in humans: asocial exploration tendency does not predict reliance on social learning. *Evolution and Human Behavior*, 38(3):325–333, 2017. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2016.11.001. URL `http://linkinghub.elsevier.com/retrieve/pii/S1090513816301180`.

[36] M.M Webster and K.N Laland. Social learning strategies and predation risk: minnows copy only when using private information would be costly. *Proceedings of the Royal Society B: Biological Sciences*, 275(1653):2869–2876, 2008. ISSN 0962-8452. doi: 10.1098/rspb.2008.0817. URL `http://rspb.royalsocietypublishing.org/cgi/doi/10.1098/rspb.2008.0817`.

[37] M M Webster and K N Laland. Reproductive state affects reliance on public information in sticklebacks. *Proceedings. Biological sciences / The Royal Society*, 278(1705):619–627, 2011. ISSN 0962-8452. doi: 10.1098/rspb.2010.1562.

[38] Robert Boyd and Peter J. Richerson. Social learning as an adaptation. *Lectures on Mathematics in the Life Sciences*, 20:1–26, 1989.

[39] R. Bond. Group Size and Conformity. *Group Processes & Intergroup Relations*, 8(4): 331–354, oct 2005. ISSN 1368-4302. doi: 10.1177/1368430205056464. URL `http://gpi.sagepub.com/cgi/doi/10.1177/1368430205056464`.

[40] M. A. Kline and R. Boyd. Population size predicts technological complexity in Oceania. *Proceedings of the Royal Society B: Biological Sciences*, 277(1693):2559–2564, 2010. ISSN 0962-8452. doi: 10.1098/rspb.2010.0452. URL `http://rspb.royalsocietypublishing.org/cgi/doi/10.1098/rspb.2010.0452`.

[41] Michael Muthukrishna, Ben W Shulman, Vlad Vasilescu, and Joseph Henrich. Sociality influences cultural complexity. *Proceedings of Royal Society B*, 281(November 2013): 20132511, 2014.

[42] Sally E Street, Ana F Navarrete, Simon M Reader, and Kevin N Laland. Coevolution of cultural intelligence , extended life history , sociality , and brain size in primates. *Proceedings of the National Academy of Sciences*, 114(30):1–7, 2017. ISSN 0027-8424. doi: 10.1073/pnas.1620734114.

[43] Sc Nicolis and Jl Deneubourg. Emerging patterns and food recruitment in ants: an analytical study. *Journal of theoretical biology*, 198(4):575–92, jul 1999. ISSN 1095-8541. doi: 10.1006/jtbi.1999.0934. URL `http://www.ncbi.nlm.nih.gov/pubmed/10373356`.

[44] Stephen C Pratt and David J T Sumpter. A tunable algorithm for collective decision-making. *Proceedings of the National Academy of Sciences of the United States of America*, 103(43):15906–10, oct 2006. ISSN 0027-8424. doi: 10.1073/pnas.

0604801103. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1635101{&}tool=pmcentrez{&}rendertype=abstract`.

[45] Christian List. Democracy in animal groups: a political science perspective. *Trends in Ecology and Evolution*, 19(4):166–168, 2004. ISSN 01695347. doi: 10.1016/j.tree.2004.02.001.

[46] Andrew J King and Guy Cowlishaw. When to use social information: the advantage of large group size in individual decision making. *Biology letters*, 3(2):137–139, 2007. ISSN 1744-9561. doi: 10.1098/rsbl.2007.0017.

[47] Max Wolf, Ralf H J M Kurvers, Ashley J W Ward, Stefan Krause, and Jens Krause. Accurate decisions in an uncertain world: collective cognition increases true positives while decreasing false positives. *Proceedings. Biological sciences / The Royal Society*, 280(1756):20122777, apr 2013. ISSN 1471-2954. doi: 10.1098/rspb.2012.2777. URL `http://www.ncbi.nlm.nih.gov/pubmed/23407830`.

[48] Andres Laan, Gabriel Madirolas, and Gonzalo G De Polavieja. Rescuing collective wisdom when the average group opinion is wrong. *Frontiers in Robotics and AI*, 4(November):1–28, 2017. ISSN 2296-9144. doi: 10.3389/frobt.2017.00056.

[49] Lucy M. Aplin, Ben C. Sheldon, and Richard McElreath. Conformity does not perpetuate suboptimal traditions in a wild population of songbirds. *Proceedings of the National Academy of Sciences*, page 201621067, 2017. ISSN 0027-8424. doi: 10.1073/PNAS.1621067114. URL `http://www.pnas.org/content/early/2017/07/18/1621067114.full`.

[50] Brendan J Barrett, Richard L Mcelreath, Susan E Perry, and Brendan J Barrett. Pay-off-biased social learning underlies the diffusion of novel extractive foraging traditions in a wild primate. *Proceedings of Royal Society B*, 284, 2017.

[51] Richard McElreath, Adrian V Bell, Charles Efferson, Mark Lubell, Peter J Richerson, and Timothy Waring. Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 363(1509):3515–28, nov 2008. ISSN 1471-2970. doi: 10.1098/rstb.2008.0131. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2607339{&}tool=pmcentrez{&}rendertype=abstract`.

[52] Wataru Toyokawa, Hye-Rin Kim, and Tatsuya Kameda. Human Collective Intelligence under Dual Exploration-Exploitation Dilemmas. *PloS one*, 9(4): e95789, jan 2014. ISSN 1932-6203. doi: 10.1371/journal.pone.0095789. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3995913{&}tool=pmcentrez{&}rendertype=abstract`.

[53] Anne Kandler and Kevin N Laland. Tradeoffs between the strength of conformity and number of conformists in variable environments. *Journal of theoretical biology*, 332:191–202, sep 2013. ISSN 1095-8541. doi: 10.1016/j.jtbi.2013.04.023. URL `http://www.ncbi.nlm.nih.gov/pubmed/23643630`.

[54] Charles Efferson, Rafael Lalive, Peter J. Richerson, Richard McElreath, and Mark Lubell. Conformists and mavericks: the empirics of frequency-dependent cultural transmission.

Evolution and Human Behavior, 29:56–64, 2008. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2007.08.003.

[55] Richard McElreath, Mark Lubell, Peter J. Richerson, Timothy M. Waring, William Baum, Edward Edsten, Charles Efferson, and Brian Paciotti. Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, 26(6):483–508, nov 2005. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2005.04.003. URL `http://linkinghub.elsevier.com/retrieve/pii/S1090513805000231`.

[56] Alex Mesoudi. An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior*, 32(5): 334–342, sep 2011. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2010.12.001. URL `http://linkinghub.elsevier.com/retrieve/pii/S1090513810001443`.

[57] Charles Perreault, Cristina Moya, and Robert Boyd. A Bayesian approach to the evolution of social learning. *Evolution and Human Behavior*, 33(5):449–459, sep 2012. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2011.12.007. URL `http://linkinghub.elsevier.com/retrieve/pii/S1090513811001437`.

[58] L Rendell, R Boyd, D Cownden, M Enquist, K Eriksson, M W Feldman, L Fogarty, S Ghirlanda, T Lillicrap, and K N Laland. Why copy others? Insights from the social learning strategies tournament. *Science (New York, N.Y.)*, 328 (5975):208–13, apr 2010. ISSN 1095-9203. doi: 10.1126/science.1184719. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2989663{&}tool=pmcentrez{&}rendertype=abstract`.

[59] Jolle W Jolles, Kate L Laskowski, Neeltje J Boogert, and Andrea Manica. Repeatable group differences in the collective behaviour of stickleback shoals across ecological contexts. *Proceedings of Royal Society B*, pages 13–16, 2018.

[60] Pablo Michelena, Raphaël Jeanson, Jean-Louis Deneubourg, and Angela M Sibbald. Personality and collective decision-making in foraging herbivores. *Proceedings. Biological sciences / The Royal Society*, 277(1684):1093–9, apr 2010. ISSN 1471-2954. doi: 10.1098/rspb.2009.1926. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2842771{&}tool=pmcentrez{&}rendertype=abstract`.

[61] Isaac Planas-sitja, Jean-louis Deneubourg, Celine Gibon, and Gregory Sempo. Group personality during collective decision-making : a multi-level approach. *Proceedings of the Royal Society B: Biological Sciences*, 282:20142515, 2015.

[62] Alex Mesoudi, Lei Chang, Sasha R.X. Dall, and Alex Thornton. The Evolution of Individual and Cultural Variation in Social Learning. *Trends in Ecology & Evolution*, xx:1–11, 2016. ISSN 01695347. doi: 10.1016/j.tree.2015.12.012. URL `http://linkinghub.elsevier.com/retrieve/pii/S0169534715003237`.

[63] Brendan J Barrett. Equifinality in empirical studies of cultural transmission. *Behavioural Processes*, 2018. ISSN 03766357. doi: 10.1016/j.beproc.2018.01.011. URL `http://dx.doi.org/10.1016/j.beproc.2018.01.011`.

[64] Dora Biro, Takao Sasaki, and Steven J. Portugal. Bringing a Time-Depth Perspective to Collective Animal Behaviour. *Trends in Ecology and Evolution*, 31(7):550–562, 2016.

ISSN 01695347. doi: 10.1016/j.tree.2016.03.018. URL http://dx.doi.org/10.1016/j.tree.2016.03.018.

[65] William Hoppitt and Kevin N. Laland. *Social Learning: An Introduction to Mechanisms, Methods, and Models*. Princeton University Press, 2013. ISBN 0691150702, 9780691150703.

[66] Ulf Toelch, Matthew J. Bruce, Marius T.H. Meeus, and Simon M. Reader. Humans copy rapidly increasing choices in a multiarmed bandit problem. *Evolution and Human Behavior*, 31(5):326–333, sep 2010. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2010.03.002. URL http://linkinghub.elsevier.com/retrieve/pii/S1090513810000292.

[67] Bahador Bahrami, Karsten Olsen, Peter E. Latham, Andreas Roepstorff, Geraint Rees, and Chris D. Frith. Optimally intercting minds. *Science*, 329:1081–1085, 2010. doi: 10.5061/dryad.5t110.Supplementary.

[68] Nathaniel D Daw, John P O'Doherty, Peter Dayan, Ben Seymour, and Raymond J Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876–9, jun 2006. ISSN 1476-4687. doi: 10.1038/nature04766. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2635947{&}tool=pmcentrez{&}rendertype=abstract.

[69] Frédéric Dandurand, Thomas R. ShultzEmail, and Kristine H. Onishi. Comparing online and lab methods in a problem-solving experiment. *Behavior Research Methods*, 40(2):428–434, 2008. doi: 10.3758/BRM.40.2.428.

[70] Jérôme Hergueux and Nicolas Jacquemet. Social preferences in the online laboratory: a randomized experiment. *Experimental Economics*, 18(2):251–283, 2015. ISSN 13864157. doi: 10.1007/s10683-014-9400-5.

[71] Joshua Becker, Devon Brackbill, and Damon Centola. Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, page 201615978, 2017. ISSN 0027-8424. doi: 10.1073/pnas.1615978114. URL `http://www.pnas.org/lookup/doi/10.1073/pnas.1615978114`.

[72] I. Lorge, D. Fox, J. Davitz, and M. Brenner. A survey of studies contrating the quality of group performance and individual performance, 1920-1957. *Psychological Bulletin*, 55(6): 337–372, 1958. ISSN 0033-2909. doi: 10.1037/h0042344.

[73] Reid Hastie and Tatsuya Kameda. The robust beauty of majority rules in group decisions. *Psychological review*, 112(2):494–508, apr 2005. ISSN 0033-295X. doi: 10.1037/0033-295X.112.2.494. URL `http://www.ncbi.nlm.nih.gov/pubmed/15783295`.

[74] Elise Payzan-Lenestour and Peter Bossaerts. Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Computational Biology*, 7(1), 2011. ISSN 1553734X. doi: 10.1371/journal.pcbi.1001048.

[75] R S Sutton and a G Barto. *Reinforcement learning: an introduction.* MIT Press, Cambridge, MA, 1998. ISBN 0262193981. doi: 10.1109/TNN.1998.712192.

[76] Woo Young Ahn, Georgi Vasilev, Sung Ha Lee, Jerome R. Busemeyer, John K. Kruschke, Antoine Bechara, and Jasmin Vassileva. Decision-making in stimulant and opiate addicts in

protracted abstinence: Evidence from computational modeling with pure users. *Frontiers in Psychology*, 5(AUG):1–15, 2014. ISSN 16641078. doi: 10.3389/fpsyg.2014.00849.