



Topics in Cognitive Science 14 (2022) 550–573

© 2022 The Authors. *Topics in Cognitive Science* published by Wiley Periodicals LLC on behalf of Cognitive Science Society.

ISSN: 1756-8765 online

DOI: 10.1111/tops.12596

This article is part of the topic “Extending Rationality,” Emmanuel M. Pothos and Timothy J. Pleskac (Topic Editors).

Overcoming Individual Limitations Through Distributed Computation: Rational Information Accumulation in Multigenerational Populations

Mathew D. Hardy,^a Peaks M. Krafft,^b Bill Thompson,^{a,c}
Thomas L. Griffiths^{a,c}

^a*Department of Psychology, Princeton University*

^b*Creative Computing Institute, University of the Arts London*

^c*Department of Computer Science, Princeton University*

Received 28 February 2021; received in revised form 29 November 2021; accepted 30 November 2021

Abstract

Many of the computational problems people face are difficult to solve under the limited time and cognitive resources available to them. Overcoming these limitations through social interaction is one of the most distinctive features of human intelligence. In this paper, we show that information accumulation in multigenerational social networks can be produced by a form of distributed Bayesian inference that allows individuals to benefit from the experience of previous generations while expending little cognitive effort. In doing so, we provide a criterion for assessing the rationality of a population that extends traditional analyses of the rationality of individuals. We tested the predictions of this analysis in two highly controlled behavioral experiments where the social transmission structure closely matched the assumptions of our model. Participants made decisions on simple categorization tasks that relied on and contributed to accumulated knowledge. Success required these microsocieties to accumulate information distributed across people and time. Our findings illustrate how in certain settings, distributed

Mathew D. Hardy and P. M. Krafft contributed equally.

Correspondence should be sent to Thomas L. Griffiths, Department of Psychology, Princeton University, Princeton, NJ 08540, USA. Email: tomg@princeton.edu

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

computation at the group level can pool information and resources, allowing limited individuals to perform effectively on complex tasks.

Keywords: Bayesian inference; Cultural evolution; Collective intelligence; Distributed computation; Group rationality; Social learning

1. Introduction

Humans' success and survival depends on our ability to make intelligent choices and judgments. Like other animals, we must do so under hard constraints on the resources available to make our decisions (Griffiths, 2020). Not only is our cognition constrained by our restricted brain capacity, but we also must learn from limited and often incomplete data. This is true not only in our finite childhoods and lifespans but also in the limited opportunities we have to learn many crucial facts—to survive, we need to learn about what life-threatening dangers exist in our immediate environments (from chemicals to electric outlets to poisonous insects and animals).

Given these constraints, understanding how people routinely form accurate beliefs on complex topics is a central focus of research in cognitive science, psychology, and economics (Simon, 1990). However, people rarely develop beliefs alone—rather, they learn from the knowledge, experiences, and opinions of other people. By doing so, individuals can obtain useful information while expending little physical or cognitive effort. A particularly important focus of previous research has been on the effects of social learning when repeated over successive groups of individuals, such as child–parent learning, formal education, and other domains where knowledge is transmitted from older members of a population to younger learners. In these multigenerational settings, knowledge can accumulate over time in a population, allowing individuals to extend their cognitive skills by learning from others (see, e.g., Almaatouq, Alsobay, Yin, & Watts, 2021; Almaatouq et al., 2020; Belikov, Rzhetsky, & Evans, 2020; Caldwell, Atkinson, & Renner, 2016; Frey & Goldstone, 2018; Galesic, Olsson, Dalege, van der Does, & Stein, 2021; Goldstone, Wisdom, Roberts, & Frey, 2013; Hazła, Jadbabaie, Mossel, & Rahimian, 2021; Kempe & Mesoudi, 2014b; Mesoudi, 2016; Mesoudi & Thornton, 2018; Miton & Charbonneau, 2018; Riedl, Kim, Gupta, Malone, & Woolley, 2021; Rzhetsky, Foster, Foster, & Evans, 2015; Salhab, Ajorlou, & Jadbabaie, 2020; Wisdom, Song, & Goldstone, 2013; Wojtowicz & DeDeo, 2020, for recent overviews and related studies). Crucially, the constraints and structure of interpersonal transmission often lead collective knowledge and learning to differ from individual outcomes (Kirby, Tamariz, Cornish, & Smith, 2015; Ravignani, Thompson, Grossi, Delgado, & Kirby, 2018; Silvey, Kirby, & Smith, 2019).

The accumulation of collective knowledge through sequential social learning is known as *cultural transmission* (Boyd & Richerson, 1985) and is thought to underpin cumulative cultural evolution (Mesoudi, 2011). Cultural transmission has been studied experimentally in a number of paradigms, including the evolution of simple technologies such as knots (Muthukrishna, Shulman, Vasilescu, & Henrich, 2014), virtual fishing nets (Derex, Beugin,

Godelle, & Raymond, 2013), stone tools (Morgan et al., 2015), or arrowheads (Mesoudi & O'Brien, 2008); artificial languages (Kirby, Cornish, & Smith, 2008); jigsaw puzzles (Kemp & Mesoudi, 2014a); and social phenomena like stereotypes (Martin et al., 2014).

Theories of cultural evolution have primarily been underpinned by a Darwinian framework grounded in parallels and disanalogies with biological evolution (Boyd & Richerson, 1985; Mesoudi & Whiten, 2008; Nettle, 2020; Smolla et al., 2020). These frameworks have shown that the accumulation of knowledge and technology can be understood as a form of evolution, helping to situate cultural evolution within the biological and evolutionary sciences (Laland, Sterelny, Odling-Smee, Hoppitt, & Uller, 2011). While successful, evolutionary frameworks have been difficult to connect to psychological theory (Heyes, 2018) and are therefore difficult to connect with the concepts of computation that are central to the study of cognition and intelligence (Nettle, 2020). Similarly, while research on collective intelligence has extensively documented the advantages of groups over individuals, the computational structure of the social processes that lead to long-term, open-ended collective intelligence remain unclear (Krafft et al., 2016).

One way to measure the accumulation of knowledge in populations is by viewing collective behavior as distributed computation, a process that allows groups to “store and process the cumulative innovations and collaborations of generations of individuals” (Smaldino & Richerson, 2013). From this perspective, knowledge accumulation becomes a problem of *distributed* Bayesian inference, extending probabilistic models of inference in individuals to the group setting (Chater, Tenenbaum, & Yuille, 2006; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010; Harper, 2009). This perspective establishes a formal connection between cultural evolution and statistical models of social learning used in cognitive science (Cushman & Gershman, 2019) and economics (Acemoglu, Dahleh, Lobel, & Ozdaglar, 2011).

In this paper, we build on this work and show how social learning can facilitate rational action that goes beyond the direct experience of individuals. That is, we show how social interactions allow limited individuals to improve their cognition without modifying their time and resource constraints. We do this in part by offering a formal criterion for *population rationality*: the probability that any individual in a network makes a particular decision is the same as the probability of that decision under a Bayesian posterior distribution conditioned on *all* the information observed by the population (Foster, 2018, cf.).

To develop our account, we first formulate a model of individual social learning based on a simple heuristic that requires only limited social observation. We then show that under certain conditions, groups of individuals following this heuristic will accumulate information through distributed Bayesian inference. This model offers insight into how social learning can extend the cognitive abilities of limited individuals—while conditioning on all accumulated information would be too complex for any individual, social interactions allow individuals to benefit from this information while expending little cognitive effort. That is, population rationality can be achieved by networks of highly limited individuals.

We tested our model in two large-scale experiments where participants made basic categorization decisions in simple multigenerational networks. In both experiments, the decisions of participants at one generation were propagated to those in the next generation, allowing us to study belief accumulation and transmission in a controlled laboratory setting. We found

that a substantial proportion of the data from these experiments are well-approximated by our model. Furthermore, by comparing participants' choices with the Bayesian posterior distribution conditioned on the information observed by the entire population, we were able to quantify accumulation in our experimental networks relative to the Bayesian ideal. Our findings thus offer a clear demonstration of how limited individuals can use simple social learning heuristics to make intelligent inferences.

2. Information accumulation through social sampling

Here, we describe a simple social learning heuristic we call *social sampling* that yields distributed Bayesian inference at the population level. That is, populations of individuals following social sampling will accumulate information in a way that is consistent with Bayesian inference, allowing individuals to offload computation to the group and make accurate inferences with little cognitive effort. We focus on multigenerational settings, where individuals are organized in discrete “batches” and learn from those in the previous generation. We also assume individuals observe others' true beliefs and do not need to discard or modify any observations.

2.1. Model

Information accumulation by rational agents is specified by Bayes' rule, which indicates how a probability distribution over hypotheses θ (known as the *prior* distribution) should be updated (to the *posterior* distribution) in light of evidence D :

$$p(\theta|D) = \frac{p(D|\theta) \cdot p(\theta)}{p(D)}. \quad (1)$$

In this paper, we extend this characterization of optimal belief updating from individuals to groups (see Chater, Oaksford, Hahn, & Heit, 2010; Griffiths, Kemp, & Tenenbaum, 2008; Tenenbaum, Kemp, Griffiths, & Goodman, 2011, for reviews of Bayesian models of individual cognition). Viewing a population as a single agent is a perspective with roots in many traditions, including philosophy (Easwaran, 2019), economics (Gale & Kariv, 2003), economics (Hayek, 1945), organization science (Argote, 2012), cognitive science (Krafft, Shmueli, Griffiths, Tenenbaum, & Pentland, 2021), anthropology (Hutchins, 1995), collective intelligence (Engel et al., 2015; Woolley, Chabris, Pentland, Hashmi, & Malone, 2010), ethology (Sasaki & Pratt, 2011), and computer science (Lynch, 1996; Shoham & Leyton-Brown, 2009). Crucially, as we will show, rational action at the population level requires only limited computation by individuals—while each individual agent may follow a simple heuristic, the information accumulated by the population as a whole makes it possible for those individuals to act rationally.

We consider the simple case where individuals reason about a set of that can be either true or false, represented as a binary feature vector \mathbf{x} . Individuals beliefs are shaped by pieces of evidence $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_T)$ that arrive over time. At each time t , individuals observe J

pieces of information about feature i , with $P_i = P(y_{ijt} = 1|x_i = 0) = \theta_{i0}$ if $x_i = 0$ and $P_i = P(y_{ijt} = 1|x_i = 1) = \theta_{i1}$ if $x_i = 1$. For simplicity, we let $S_{it} = \sum_{j=1}^J 1(y_{ijt} = 1)$ be the number of positive observations made about feature i at time t .

We assume that the popularity of a belief among members of the population acts as a prior distribution over beliefs for new individuals drawing inferences about the environment. Specifically, an individual a first chooses a member of the population a' to learn from uniformly at random.¹ The learner a then accepts or rejects their companion's decision $d_{a',i,t-1}$ with probability proportional to $(\theta_{i,d_{a',i,t-1}})^{S_{it}}(1 - \theta_{i,d_{a',i,t-1}})^{J-S_{it}}$, the likelihood of the evidence that a observes at time t based on the beliefs of a' .

In an infinite population of individuals following this strategy, the probability that an individual makes a categorization decision about feature i is equal to the posterior probability of that decision conditioned on the evidence observed by the entire population:

$$p_{i,t+1} = P(x_i = 1|y_{i,\cdot,\leq t}). \tag{2}$$

This result establishes that social sampling is a valid algorithm for distributed Bayesian inference in infinitely large populations, satisfying our criterion for population rationality (see Supplementary Electronic Material (SEM) for proof). The finite population case is more complex, but social sampling can be seen to be formally equivalent to a class of sequential Monte Carlo algorithms known as particle filters (Crisan & Doucet, 2002; Murphy, 2012). Social sampling can thus support Bayesian inference in expectation in finite populations, allowing groups of bounded individuals to accumulate knowledge over time.

2.2. Measuring accumulation relative to optimal inference

The model we have outlined illustrates how a simple social decision-making heuristic can lead to optimal information aggregation over generations. This suggests that in certain contexts the Bayesian posterior distribution can be used to construct a valid upper limit on information accumulation, because posterior distributions are information-theoretically and decision-theoretically optimal belief representations (Bernardo & Smith, 2000; Jaynes, 2003; Ortega, 2011). To quantify ideal information accumulation, we use the sufficient statistics of the Bayesian posterior distribution over the environment features given the evidence observed by the entire population. We assume a uniform prior, and so the posterior probability that a feature categorization is correct given the evidence received by the group up to generation t is

$$P(x_{it} = 1|y_{i,\cdot,\leq t}) = \frac{(\theta_{i1})^{S_{i,\leq t}}(1 - \theta_{i1})^{t \cdot J - S_{i,\leq t}}}{(\theta_{i1})^{S_{i,\leq t}}(1 - \theta_{i1})^{t \cdot J - S_{i,\leq t}} + (\theta_{i0})^{S_{i,\leq t}}(1 - \theta_{i0})^{t \cdot J - S_{i,\leq t}}}, \tag{3}$$

where $S_{i,\leq t} = \sum_{k=1}^t \sum_{j=1}^J \mathbf{1}(y_{ijk} = 1)$ is the total positive evidence observed for feature i up to time t .

¹ This assumption may seem implausible in large populations, but it is satisfied if each person chooses from a small number of others and those others are an unbiased sample from the full population.

The sufficient statistics of this posterior distribution are $\frac{S_{i \leq t}}{I \cdot J}$, or the proportion of pieces of evidence favoring a feature categorization. Because these statistics can be used to exactly compute the posterior distribution, we can measure accumulation in observed networks by comparing this statistic with the popularity of a belief, or proportion of people who believe it. We express this formally by letting $p_{it} = \frac{1}{N} \sum_a \mathbf{1}(d_{a,i,t-1} = 1)$ be the proportion of individuals who believe $x_i = 1$ at time $t - 1$, with d_{ait} indicating whether individual a at generation t chooses category i , and N representing the number of individuals in each generation. If the popularity is close to $\frac{S_{i \leq t}}{I \cdot J}$ at time t , we can conclude that information is being effectively accumulated in the population.

2.3. Summary

This analysis demonstrates that individuals in multigenerational networks can improve their inferences by following a simple social learning heuristic. To perform social sampling, an individual first randomly selects a person in the previous generation to learn from. They then evaluate their companion's belief against observed evidence—the more aligned the belief with the evidence, the higher the chance the individual will accept it as their own, otherwise they will continue searching and sample another person. While similar social learning models have been explored in multiarmed bandit problems (Celis, Krafft, & Vishnoi, 2017; Krafft, 2017; Krafft et al., 2021), they have not been tested in controlled laboratory settings.

Social sampling is simple to perform and only requires limited computation by individuals. More precisely, at any time the popularity of a hypothesis will approximate the true Bayesian posterior probability conditioned on all the evidence observed by the population, and so the population itself performs distributed Bayesian inference. Because this distributed inference matches the inference problems individuals face, our model allows us to be precise about *how* social learning extends people's cognitive abilities.

In our model, individuals make unbiased judgments and cannot choose which information to transmit. Instead, every person's true beliefs can be observed by those in the next generation. While this setup—and assumptions of uniform sampling of others' true beliefs—is highly simplified, it reflects a common approach to modeling individual cognition as a two-stage decision-making process (Howard & Sheth, 1969; Krumme, Cebrian, Pickard, & Pentland, 2012; Payne, 1976; Pratt, Sumpter, Mallon, & Franks, 2005; Seeley & Buhrman, 1999; Vul, Goodman, Griffiths, & Tenenbaum, 2014) that has also been applied to modeling iterated learning in populations (Bonawitz, Denison, Gopnik, & Griffiths, 2014; Kalish, Griffiths, & Lewandowsky, 2007; Mozer, Pashler, & Homaei, 2008; Sanborn & Griffiths, 2008; Vul & Pashler, 2008) and allows us to make progress in understanding knowledge accumulation in a way that can be studied empirically and extended to more naturalistic settings.

While our model shows how group-level Bayesian inference is possible, can we identify and quantify accumulation in real populations? To investigate, we ran two behavioral experiments where participants made categorization decisions in multigenerational networks. In both experiments, we assessed population rationality by comparing observed popularity with

the Bayesian ideal and evaluated our social sampling model by comparing it with several alternative models.

3. Experiment 1: Gem classification

In Experiment 1, participants completed a categorization task that was framed as gemstone classification. This experiment was designed to closely mirror an idealized setting in which optimal information aggregation is possible through our proposed Bayesian social sampling mechanism.

In this experiment, participants played the role of technicians classifying gemstones in a certain shift (i.e., generation).² Each gemstone could have up to eight potential classifications, with each classification being a randomized nonsense word (e.g., “pesho,” “ivil,” “thyun”). On each trial, participants observed social information in the form of classification judgments drawn from the previous shift of technicians, as well as non-social information in the form of a set of four new laboratory results. The laboratory results presented stochastic positive or negative evidence for each potential gemstone classification. We recruited three independent networks of participants, with each network including 10 shifts of 20 technicians (i.e., 200 participants per network). The interface for the experiment is illustrated in Fig. 1.

3.1. Methods

3.1.1. Participants

We recruited 600 participants from Amazon Mechanical Turk. Recruitment was restricted to participants in the United States with an Amazon Mechanical Turk approval rating of 95 or above. We paid participants \$1.75 as compensation, plus a performance-based bonus payment of up to \$0.50. The task typically took participants under 10 min to complete. Prior to starting the experiment, participants completed an attention and comprehension check which included questions about details of the study, including the probability of different kinds of evidence appearing in individual tests. Participants who failed to answer these questions correctly in three attempts were excluded from taking part in the study.

3.1.2. Stimuli

On each round of the experiment, participants observed a cartoon gemstone positioned at the top of the screen on a white background. Each gem had eight possible binary classifications, and participants performed lab tests to make these classifications. Each test resulted in either positive or negative evidence for each classification. Fig. 1 shows an example participant view.

Classifications were given labels from an artificial vocabulary. These labels were different on each trial a participant completed (i.e., none of the labels recurred across trials for a

² See <https://github.com/pkrafft/Overcoming-individual-limitations-through-distributed-computation> for experiment code and data.

View of Participants in **Generations 1 and 2**

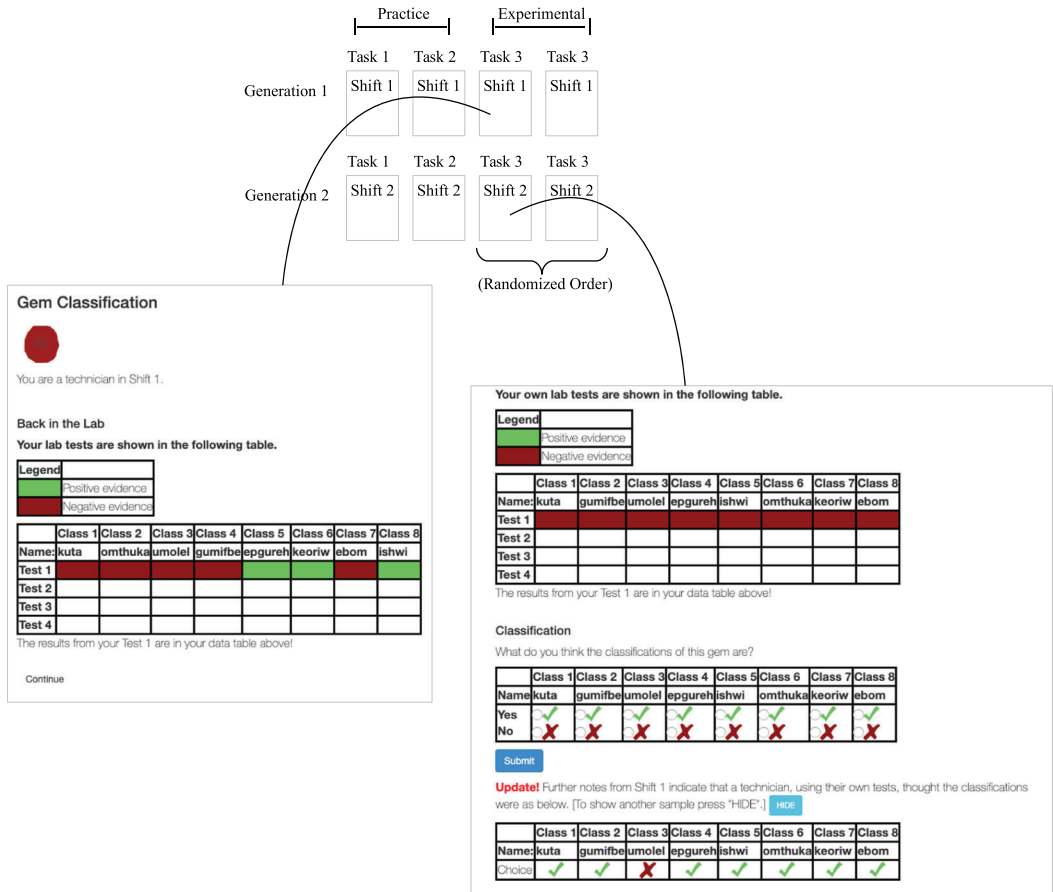


Fig. 1. Participant interface for Experiment 1. Each categorization task consisted of a different set of gem classifications that were given random nonsense words. Each participant belonged to a “generation” of participants, called a “shift” in this experiment. In each generation after the first, participants could view the gem classifications made by participants in the last generation. Each participant made their own classification decisions for a gem four times after observing four pieces of non-social evidence.

participant) and were presented in randomized order. After each test, participants indicated which of the eight classifications they thought were true for this gem and could modify these classifications throughout the trial.

Participants in generations 2–10 observed the final classification decisions of a randomly sampled participant from the previous generation of the same network. Participants could choose to resample from the previous generation at any time during the trial. When participants chose to resample, the new sample was selected at random with replacement from the 20 participants in the previous generation of the same network. Social information (an earlier participant’s classification decisions) was presented in a feedback table below the participant’s

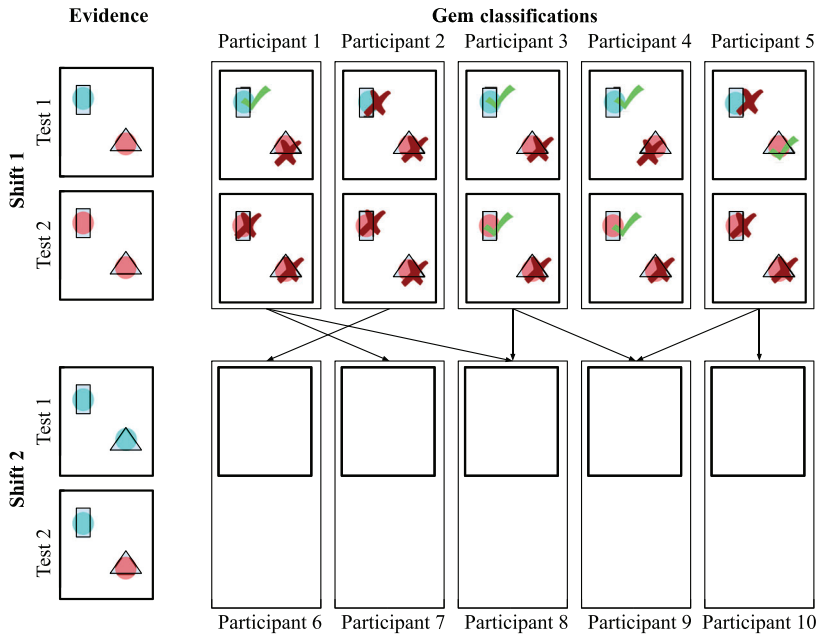


Fig. 2. Illustration of the evidence correlation structure used in Experiment 1. If a gem classification test was positive for one participant, it was positive for all participants in that generation on that classification task. Positive evidence about a classification is illustrated here using turquoise circles and negative evidence with red circles. The classification judgments made by participants are shown as checkmarks and x-marks. Participants could change their classifications after each new test.

own feedback table. After completing the fourth lab test, the results of the test were displayed and participants moved on to the next trial.

3.1.3. Procedure

Participant completed four trials (i.e., gem classifications). The first two trials were practice rounds and were presented in the same order for all participants. The last two tasks were experimental trials and were ordered using simple randomization. We limited our analysis to data from the two experimental trials.

Mirroring the social sampling model, participants were organized into discrete shifts, or “generations.” We recruited three independent networks of participants, with each network consisting of 10 generations with 20 participants per generation (see Fig. 2). Participants in generations 2–10 observed the classifications of a randomly sampled member of the previous generation and could choose to draw additional samples as many times as they wished.

At the beginning of the experiment, the true classifications for each task were chosen uniformly at random. To reduce variance between networks, the lab test results for a given classification were the same for all participants in a single generation (see Fig. 2). The probability of positive evidence in a lab test result was 0.6 for true classifications and 0.3 for false classifications (participants were informed of both probabilities).




Fig. 3. Experiment 1 results. The graph on the left shows the relationship between the total evidence available for each gem classification and the proportion of participants who selected that classification (i.e., the popularity of the classification). The graph on the right plots the relationship between the most recently observed evidence and the popularity. Each point represents a gem classification for a single shift, trial, and replication.

3.2. Results and discussion

We limited our analysis to participants in generations 2–10 and preregistered our statistical analyses and sample size before the experiment.³ These analyses, however, were adjusted and expanded over the course of revisions. This included adding model comparisons, and using all the test data in the regression described below (rather than only one decision in generations 6–10) for consistency with our analysis for Experiment 2. Limiting our regression to the preregistered subset does not change our findings.

As predicted by our social sampling model, we observed a strong correlation between the popularity of each gem classification and both the total evidence ($r(430) = 0.6, p < .0001$) and the most recent evidence ($r(430) = 0.88, p < .0001$; see Fig. 3). Furthermore, total evidence was significantly predictive of popularity in a regression that included the total evidence, most recent evidence, and fixed effects for each network ($t(424) = 4.54, p < .0001$; see Table 1 for full regression results). The correlation between popularity and total evidence shows that the proportion of social sampling was high enough to facilitate significant information aggregation over time.

We then compared the performance of our Bayesian social sampling model to three asocial models and two social models—a probability matching social sampling model and naive copying model (models are described in detail in the SEM). In the probability matching social sampling model, individuals accept candidate hypotheses with probability equal to the fraction of most recent positive evidence and in the naive copying model, individuals simply copy others' categorizations. To account for non-social learners, both social sampling models include a term capturing the proportion of social decision making. We estimated this value as the proportion of trials where participants chose to resample at least once and found partici-

³ Preregistration available at https://osf.io/7qx8e/?view_only=53781c6aa4c742e2a9843fcef533dfd7.

Table 1
Regression results for Experiments 1 and 2

Variable	Regression Model	
	Experiment 1	Experiment 2
Intercept	-0.10** (-3.15)	0.29*** (12.69)
Total evidence	0.24*** (4.54)	0.05* (2.44)
Last evidence	0.91*** (29.26)	0.25*** (18.74)
Dependent variable	Popularity	Popularity
Degrees of freedom	424	2270
Observations	432	2304

t statistics are in parentheses.

* $p < .05$, ** $p < .01$, *** $p < .001$.

Note that fixed effects for each network (Experiments 1 and 2) and each condition (Experiment 2) were included in the regressions but are not shown.

Table 2
Mean squared error of predicted and observed popularity for each model in Experiment 1

Model	Mean Squared Error
Non-social problem matching	0.0213
Non-social Bayesian problem matching	0.0220
Non-social utility maximizing	0.1220
Naive copying	0.1260
Social sampling (problem matching)	0.0181
Social sampling (Bayesian)	0.0182

The best performing model is shown in bold.

pants did so on 19% of trials. This is a lower bound on the true proportion of social sampling, since participants may have used the initial social information without resampling. However, it is close to both a separate estimate of the asocial learning proportion based on a qualitative coding of participants' strategy descriptions, and the estimated regression coefficient on total evidence available, which should correspond to the level of social sampling (see Table 1).

No free parameters are used in any of these models, and so we evaluate each model using their performance on all the experimental data. We found that the probability matching social sampling model achieved the lowest mean squared error (MSE) in predicting the popularity of each gem classification (MSE : 0.0181), followed closely by the Bayesian social sampling model (MSE : 0.0182; see Fig. 4 and Table 2 for results). Because both social sampling models predict accumulation, these results support our other analyses from Experiment 1, suggesting that networks of participants accumulated information about each categorization across time.

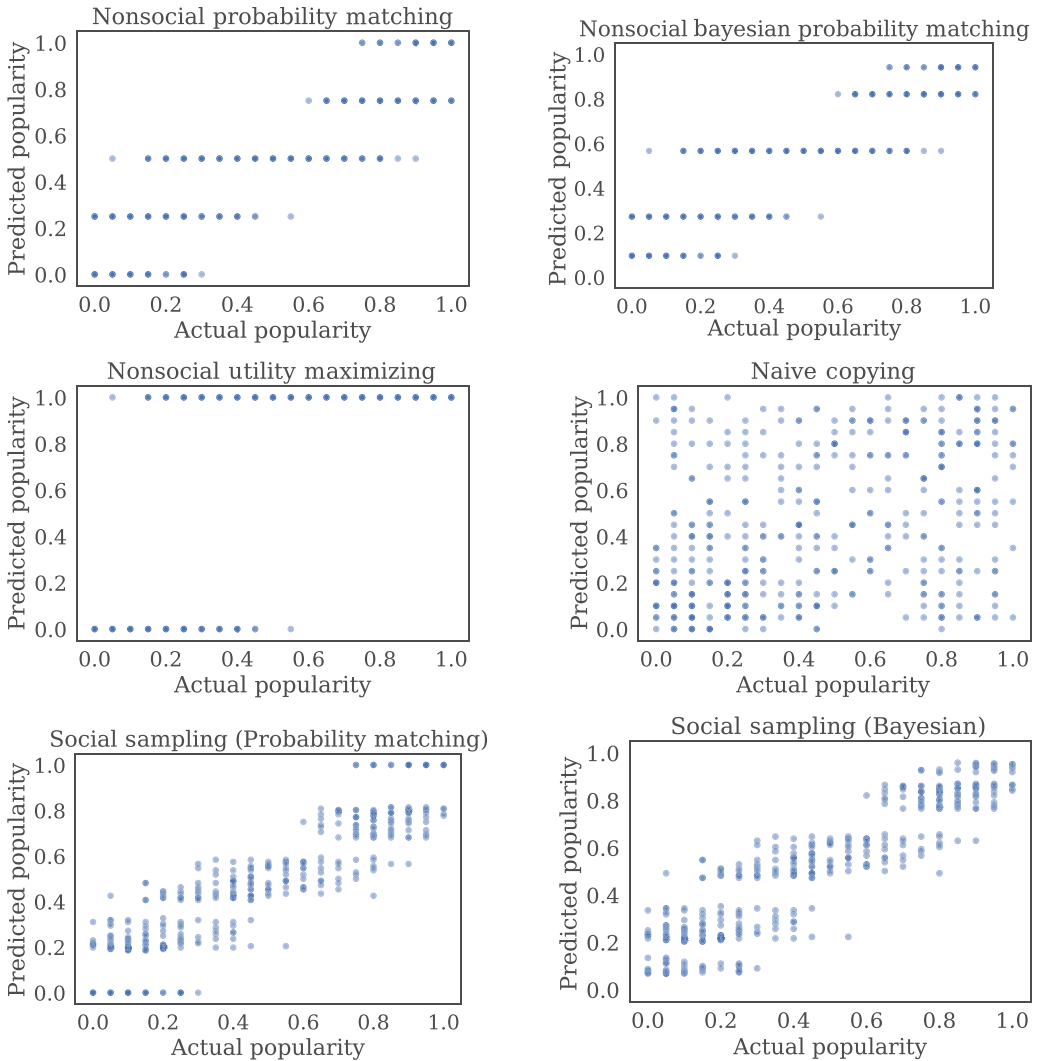



Fig. 4. Model predictions for the three non-social models and three social models we compared in Experiment 1. As in Fig. 3, 432 points are shown in each plot.

4. Experiment 2: Spaceship construction

We constructed Experiment 1 to have a high correspondence with the assumptions of our social sampling model. Experiment 2 was designed to test our model in a less idealized context. Most notably, in Experiment 2 all information was social—participants received feedback on the *previous generation's* decisions, rather than their own. Furthermore, we did not allow participants to resample from the previous generation, and feedback was censored based

View of Participant in **Generation 3**



1. View action of Generation 2 participant and their results, and choose own action

2. View results of own action for this game

Game 4 of 8

You are number 3 to play.

Below, you can view the parts *the last player* used on their ship.

You will also be shown which parts of the last player's design worked or failed on that player's flights.

Any parts highlighted in red were faulty on that flight of the last player's ship.

The ship design has been refined over the 2 generations of players before you.

Flight 1 Flight 2 Flight 3 Flight 4

Last Player's Ship: Your Ship:

Please click to select/deselect the parts you want to use on your ship:

Submit

Game 4 of 8

You are number 3 to play.

Below, you can view the parts *the last player* used on their ship.

You will also be shown which parts of the last player's design worked or failed on that player's flights.

Any parts highlighted in red were faulty on that flight of the last player's ship.

The ship design has been refined over the 2 generations of players before you.

Flight 1 Flight 2 Flight 3 Flight 4

Last Player's Ship: Your Ship:

Roll over above icons to view the flights of the last player's ship.

Please click to select/deselect the parts you want to use on your ship:

Submit

Your flights:

Flight 1 Flight 2

Next

Fig. 5. Example of the participant interface for the Experiment 2 spaceship design task. Each trial was associated with a different set of spaceship parts (using different visual icons). Participants were organized into discrete generations, or “turns.” In each generation after the first, a participant started by viewing an example spaceship design from a participant of the last generation. The participant then chose their own parts for that task and observed the part failures on each flight of that design.

on the observed participant’s choices. That is, participants did not always observe evidence on every possible categorization.

In Experiment 2, participants designed spaceships by choosing which components to include from an inventory of alternatives (see Fig. 5).⁴ Each component was either a good

Table 3

The parameter settings associated with each condition of Experiment 2

Experiment	Reps	N	J	θ
Condition 1	2	20	4	0.6
Condition 2	1	20	1	0.6
Condition 3	1	20	4	0.8
Condition 4	2	5	4	0.6
Condition 5	2	5	1	0.6

Note. The repetitions (reps) is the number of repetitions of the parameter settings we ran. N is the number of participants per generation. J is the number of flights, that is the amount of evidence, shown per spaceship design trial. θ gives the probability of success of good parts, and the probability of failure of bad parts—that is, the strength of evidence.

part that rarely failed or a bad part that failed often. On each trial, participants observed spaceships designed in the previous generation, along with the success or failure of each of the included components on one or more flights. We organized participants into five conditions, varying the number of participants in each generation (large vs. small), the amount of evidence presented to each participant (high vs. low), and the strength of that evidence (high vs. low) (see Table 3).

4.1. Methods

4.1.1. Participants

We recruited 1000 participants from Amazon Mechanical Turk. Recruitment was restricted to participants living in the United States. The task typically took participants around 5–10 min to complete, and participants earned \$1.80 as compensation.

4.1.2. Stimuli

Experiment 2 consisted of a spaceship design categorization task inspired by prior experiments on social learning (Brand, Brown, & Cross, 2018). Example stimuli from this categorization task are shown in Fig. 5. In this task, participants designed spaceships by selecting which parts to include from a set of available options. A different set of eight parts was available on each trial. Each of the eight items in a trial's set of spaceship parts could be included in that trial's spaceship design. Participants could select or remove a part in a trial's design by clicking on an icon of the part.

Participants were organized into 10 discrete generations, with each generation composed of a different set of participants. In the first generation, each participant had no information about part failures and had to simply guess which parts might be good or bad. After selecting which parts to include, a participant was given feedback from a number of “flights” of their constructed spaceship. On each flight, spaceship parts could either succeed or fail. Each

⁴ See <https://github.com/pkrafft/Overcoming-individual-limitations-through-distributed-computation> for experiment code and data.

flight was shown in sequence, and we did not allow participants to change the spaceship design between flights.

In all generations except the first, participants viewed the spaceship design and flight outcomes from a spaceship selected uniformly at random from the designs made in the previous generation. Participants viewing another participant's design could select which of the prior flights to examine, as different parts may have failed on different flights. A prior design was displayed in the same way that a participant's own spaceship design was displayed, except that parts could not be modified on a previous design and part failures were shown at the start of the trial. Unlike Experiment 1, participants could not resample a choice from the previous generation.

4.1.3. Procedure

Participants completed eight categorization tasks. The first four trials participants completed were practice trials, and the last four were test trials. Practice trials were displayed in the same order for all participants, and test trials were ordered using simple randomization. We did not inform participants of this practice-test distinction, and we limited our analysis to data from the four test trials.

In order to mirror the structure of the social sampling model, we recruited participants in discrete generations. Generation t was recruited after all participants in generation $t - 1$ completed the experiment. In the first generation, good and bad parts were chosen at random with probability .5. As in Experiment 1, spaceship failures were perfectly correlated across designs for a single flight in a given generation to reduce variance between networks. That is, if two different participants both used the same part in a certain generation, then the part would either succeed or fail for both participants on a given flight (see Fig. 6).

4.1.4. Results and discussion

As in Experiment 1, all statistical tests were two-tailed with an alpha level of 0.05 and we excluded the initial generation of participants from our analysis. We ran experiment conditions separately, and so no between-subject randomization into conditions was used. Instead, participants could only participate in the experiment once and thus could not complete multiple experimental conditions.

Replicating our findings from Experiment 1, we found that total evidence was significantly predictive of popularity, controlling for the most recent evidence observed in a regression that included the total evidence, most recent evidence, and fixed effects for each network and condition ($t(2270) = 2.44, p = .015$; full results in Table 1). While we found significant positive correlations between popularity and total evidence in each of our conditions, the strength of this correlation varied between 0.16 in Condition 5 and 0.65 in Condition 3 (see Table 4 and Fig. 7). Indeed, popularity tended to be less correlated with total evidence in networks with small generation sizes, low evidence strength, and low evidence amount.

All decisions in Experiment 2 were social, and so we limited model comparison to the three social models used in Experiment 1 and set the proportion of asocial learners to zero. As in

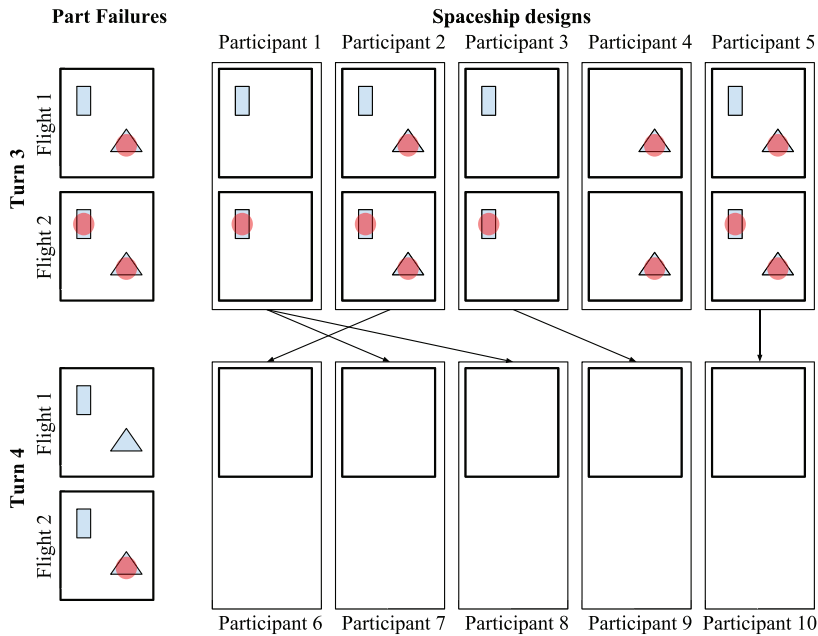


Fig. 6. Illustration of the evidence constraints in Experiment 2. As in Experiment 1, evidence was correlated across participants: If a part failed on one flight, it failed for all participants in that generation on that flight (part failures are displayed in the figure as red circles). In contrast to Experiment 1, participants could not change their spaceship designs in between flights, and evidence about a particular spaceship part was only shown if the part was chosen in the sample design shown to the participant.

Table 4

Overview of the five conditions of Experiment 2 and the correlation between popularity and total evidence within each condition

Condition	Reps	Population Size	Evidence Amount	Evidence Strength	Correlation Between Popularity and Total Evidence
Condition 1	2	Large	High	Low	.43***
Condition 2	1	Large	Low	Low	.39***
Condition 3	1	Large	High	High	.65***
Condition 4	2	Small	High	Low	.29***
Condition 5	2	Small	Low	Low	.16***

* $p < .05$, ** $p < .01$, *** $p < .001$.

Experiment 1, each model was assessed by comparing the popularity predicted by the model at each generation with the observed popularity. We found that in four of the five experimental conditions, the Bayesian social sampling model achieved the lowest MSE in predicting observed popularity. However, in the remaining condition, where the evidence

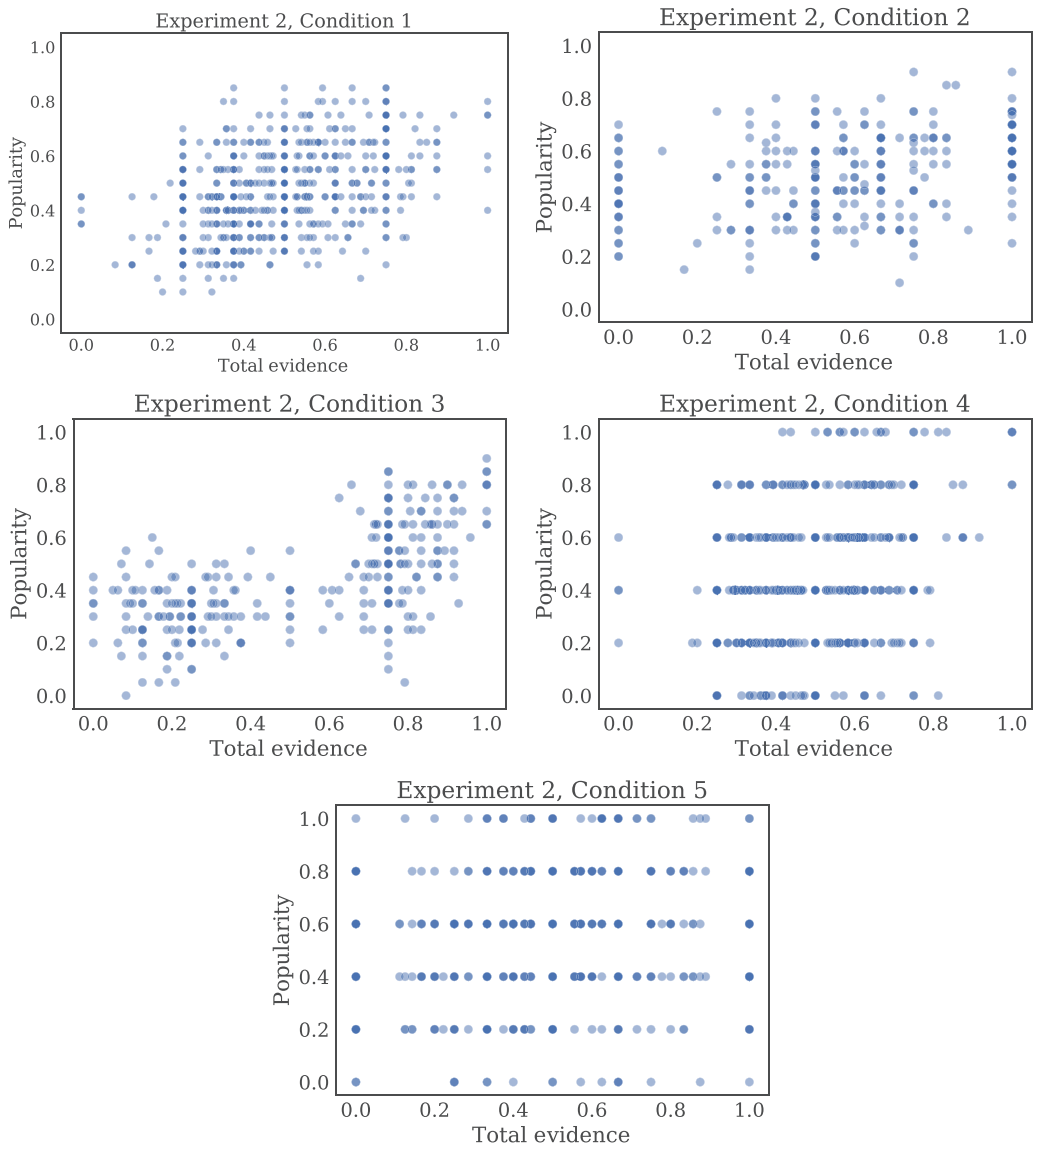


Fig. 7. Plots showing the relationship between the total evidence available for each particular spaceship part and the proportion of participants who selected that part in Experiment 2. Each point represents one part (of eight) in one social generation (of nine) on one trial (of four) in a single repetition.

strength was high, the naive copying model outperformed both the Bayesian and probability matching social sampling models (Bayesian social sampling MSE: 0.1188; probability matching social sampling MSE: 0.0825; naive copying MSE: 0.0371; see Table 5).

Table 5

Mean squared error of predicted and observed popularity for the three social models we compared for each of the conditions (C1–C5) in Experiment 2

Model	Mean Squared Error				
	C1	C2	C3	C4	C5
Naive copying	0.0433	0.0424	0.0371	0.1078	0.1030
Social sampling (matching)	0.0542	0.1703	0.0825	0.1069	0.2196
Social sampling (Bayesian)	0.0366	0.0333	0.1188	0.0969	0.0960

The best performing model in each column is shown in bold.

5. General discussion

Understanding how people make intelligent decisions under limited time and cognitive resources is a central focus of research in cognitive science, psychology, and economics. Social interactions offer a way for people to overcome these limitations by distributing computation across a group. By observing and learning from others, individuals do not need to perform complex computations or to condition on large amounts of data to make rational inferences. Instead, they can use simple heuristics that leverage accumulated social information.

To show how distributed inference can emerge in populations, we derived a social sampling model of individual decision making in multigenerational networks. We then showed that social sampling at the individual level can lead to distributed Bayesian inference at the population level. We tested the predictions of our social sampling model in two highly controlled behavioral experiments where participants made simple categorization decisions. While the transmission structure we used in both experiments was idealized and highly simplified compared to real-world social networks, this allowed us to directly quantify the degree to which information accumulated across time relative to the Bayesian ideal.

Although we observe information accumulation in both experiments, the performance of the Bayesian social sampling model varied in different conditions. Most notably, in Experiment 1 the probability social sampling model achieved a slightly lower MSE than the Bayesian social sampling model (see Table 2). Furthermore, in Experiment 2 the naive copying model had a lower MSE than both social sampling models in large networks with high levels of strong evidence (see Tables 4 and 5). These results suggest that individuals may adapt their social learning strategies to different domains and may be more likely to use social sampling in noisy, low-information environments (Toyokawa, Whalen, & Laland, 2018).

It is important to emphasize that for information aggregation to occur, the precise details of how social sampling takes place are less important than that individuals' decisions incorporate social information and new evidence using a probabilistic rule. Indeed, our social sampling models in Experiment 1 included a mixture of social and non-social decision makers, replicating related findings on underexploitation of social information in behavioral experiments (Mercier & Morin, 2019). Our analyses are thus not intended to show that social sampling is a definitive description of human behavior, but that a social decision-making model derived

from a normative Bayesian standard can help make sense of the extent of information aggregation we observe in specific experimental contexts.

While we find that information aggregation can be robust to individual differences, accumulation can fail if people's decision-making differs systematically from a probabilistic belief adoption–rejection strategy. Indeed, related research on the wisdom of crowds has shown that in certain contexts social interactions and observations can actually decrease group performance by reducing the diversity and independence of people's beliefs (Jenness, 1932; Lorenz, Rauhut, Schweitzer, & Helbing, 2011; Myers & Bishop, 1971). These dynamics—that is, whether social interactions improve or worsen group outcomes—appear to vary depending on people's social learning strategies (Toyokawa et al., 2018). For example, accumulation can be disrupted if people are utility-maximizing rather than probabilistic in their responses to evidence (Anderson & Holt, 1997; Bikhchandani, Hirshleifer, & Welch, 1992). While probability matching in individual decision-making is observed in certain domains (Shanks, Tunney, & McCarthy, 2002; Vul et al., 2014; Vulkan, 2000), our experiments show that this strategy can be extended to social settings: when a substantial proportion of individuals incorporate social information into their probability matching behavior, information accumulates across people and time. However, people may be less likely to follow a social sampling strategy in domains with low levels of probability matching, such as tasks with large financial incentives or consistent feedback (Shanks et al., 2002).

Our findings complement other approaches to understanding rationality in limited individuals. For example, previous work on adaptive heuristics (Gigerenzer & Goldstein, 1996) has shown that people use simple decision-making rules to exploit the structure of the environment and make decisions on complex tasks. Our work suggests that this framework can be extended to social decision making. Furthermore, by providing a rational motivation for people's social heuristics, our work offers a way to connect adaptive heuristics with resource-rational analysis (Griffiths, Lieder, & Goodman, 2015) in simple multigenerational populations. Integrating our model and findings with other frameworks for studying bounded rationality, such as models that utilize quantum probability theory (Pothos et al., 2021), should be addressed in future work.

We tested our model on a simple binary-choice decision-making task where the transmission structure was explicitly designed to match our model assumptions. These simplifications limit our ability to draw general conclusions about information accumulation in natural populations. For example, while enforcing uniform sampling of the previous generation gave us tight control over the transmission dynamics, people's sampling, and exploration strategies may depend on the task (Oaksford & Chater, 1994) and are likely considerably more complex in naturalistic domains and networks (Latora, Nicosia, & Russo, 2017). Similarly, people may not always transmit their true beliefs—instead, they may give noisy, incomplete, or misleading accounts of their opinions and decisions to others (Xu & Griffiths, 2010).

Despite these limitations, the general framework we have outlined and the principle of identifying the computational structure of population dynamics could therefore be extended to more complex and naturalistic domains in future work. For example, individuals could be

given greater control over which (if any) beliefs to transmit or sample from others. Models of learning in structured representational domains, such as language of thought models (Goodman, Tenenbaum, & Gerstenberg, 2015) or posterior sampling in general Markov decision processes (Agrawal & Jia, 2017; Osband, Van Roy, & Russo, 2013), could also be extended to the population setting in a way that is analogous to our extension of simple categorization decisions. Testing our model and quantifying knowledge accumulation in more naturalistic domains is an exciting challenge for future work.

Acknowledgments

This work was funded in part by NSF grant 1456709, DARPA Cooperative Agreement D17AC00004, and the support of the NOMIS Foundation.

References

- Acemoglu, D., Dahleh, M. A., Lobel, I., & Ozdaglar, A. (2011). Bayesian learning in social networks. *The Review of Economic Studies*, 78(4), 1201–1236.
- Agrawal, S., & Jia, R. (2017). Posterior sampling for reinforcement learning: Worst-case regret bounds. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems*. NeurIPS Foundation, Vol. 30 (pp. 1184–1194).
- Almaatouq, A., Alsobay, M., Yin, M., & Watts, D. J. (2021). Task complexity moderates group synergy. *Proceedings of the National Academy of Sciences*, 118(36), e2101062118.
- Almaatouq, A., Noriega-Campero, A., Alotaibi, A., Krafft, P., Moussaid, M., & Pentland, A. (2020). Adaptive social networks promote the wisdom of crowds. *Proceedings of the National Academy of Sciences*, 117(21), 11379–11386.
- Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, 87, 847–862.
- Argote, L. (2012). *Organizational learning: Creating, retaining and transferring knowledge*. Berlin, Germany: Springer Science & Business Media.
- Belikov, A. V., Rzhetsky, A., & Evans, J. (2020). Detecting signal from science: The structure of research communities and prior knowledge improves prediction of genetic regulatory experiments. arXiv preprint arXiv:2008.09985.
- Bernardo, J. M., & Smith, A. F. (2000). *Bayesian Theory*. New York: Wiley.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026.
- Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-stay, lose-sample: A simple sequential algorithm for approximating Bayesian inference. *Cognitive Psychology*, 74, 35–65.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
- Brand, C. O., Brown, G. R., & Cross, C. P. (2018). Sex differences in the use of social information emerge under conditions of risk. *PeerJ*, 6, e4190.
- Caldwell, C. A., Atkinson, M., & Renner, E. (2016). Experimental approaches to studying cumulative cultural evolution. *Current Directions in Psychological Science*, 25(3), 191–195.
- Celis, L. E., Krafft, P., & Vishnoi, N. (2017). A distributed learning dynamics in social groups. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, (pp. 441–450). New York: ACM.

- Chater, N., Oaksford, M., Hahn, U., & Heit, E. (2010). Bayesian models of cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(6), 811–823.
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287–291.
- Crisan, D., & Doucet, A. (2002). A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on Signal Processing*, 50(3), 736–746.
- Cushman, F., & Gershman, S. (2019). Editors' introduction: Computational approaches to social cognition. *Topics in Cognitive Science*, 11(2), 281–298.
- Derex, M., Beugin, M.-P., Godelle, B., & Raymond, M. (2013). Experimental evidence for the influence of group size on cultural complexity. *Nature*, 503(7476), 389–391.
- Easwaran, K. (2019). The concept of rationality for a city. *Topoi*, 1–13.
- Engel, D., Woolley, A. W., Aggarwal, I., Chabris, C. F., Takahashi, M., Nemoto, K., Kaiser, C., Kim, Y. J., & Malone, T. W. (2015). Collective intelligence in computer-mediated collaboration emerges in different contexts and cultures. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)* (pp. 3769–3778). New York: ACM.
- Foster, J. G. (2018). Culture and computation: Steps to a probably approximately correct theory of culture. *Poetics*, 68, 144–154.
- Frey, S., & Goldstone, R. L. (2018). Cognitive mechanisms for human flocking dynamics. *Journal of Computational Social Science*, 1(2), 349–375.
- Gale, D., & Kariv, S. (2003). Bayesian learning in social networks. *Games and Economic Behavior*, 45(2), 329–346.
- Galesic, M., Olsson, H., Dalege, J., van der Does, T., & Stein, D. L. (2021). Integrating social and cognitive aspects of belief dynamics: towards a unifying framework. *Journal of the Royal Society Interface*, 18(176), 20200857.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4), 650.
- Goldstone, R. L., Wisdom, T. N., Roberts, M. E., & Frey, S. (2013). Learning along with others. In B. Ross (Ed.), *Psychology of learning and motivation* (Vol. 58, pp. 1–45). Amsterdam, The Netherlands: Elsevier.
- Goodman, N. D., Tenenbaum, J. B., & Gerstenberg, T. (2015). *Concepts in a probabilistic language of thought*. Cambridge, MA: MIT Press.
- Griffiths, T. L. (2020). Understanding human intelligence through human limitations. *Trends in Cognitive Sciences*, 24(11), 873–883.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14(8), 357–364.
- Griffiths, T. L., Kemp, C., & Tenenbaum, J. B. (2008). Bayesian models of cognition. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology*. Cambridge, England: Cambridge University Press. (pp. 59–100).
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229.
- Harper, M. (2009). The replicator equation as an inference dynamic. arXiv preprint arXiv:0911.1763.
- Hayek, F. (1945). The use of knowledge in society. *The American Economic Review*, 35(4), 519–530.
- Hazła, J., Jadbabaie, A., Mossel, E., & Rahimian, M. A. (2021). Bayesian decision making in groups is hard. *Operations Research*, 69(2), 632–654.
- Heyes, C. (2018). Enquire within: Cultural evolution and cognitive science. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1743), 20170051.
- Howard, J. A., & Sheth, J. N. (1969). *The theory of buyer behavior*. Cambridge, MA: MIT Press.
- Hutchins, E. (1995). *Cognition in the wild*. Number 1995. Cambridge, MA: MIT Press.
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge, England: Cambridge University Press.
- Jenness, A. (1932). The role of discussion in changing opinion regarding a matter of fact. *The Journal of Abnormal and Social Psychology*, 27(3), 279.

- Kalish, M. L., Griffiths, T. L., & Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, *14*(2), 288–294.
- Kempe, M., & Mesoudi, A. (2014a). An experimental demonstration of the effect of group size on cultural accumulation. *Evolution and Human Behavior*, *35*(4), 285–290.
- Kempe, M., & Mesoudi, A. (2014b). Experimental and theoretical models of human cultural evolution. *WIREs Cognitive Science*, *5*, 317–326.
- Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(31), 10681–10686.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, *141*, 87–102.
- Krafft, P. (2017). *A rational choice framework for collective behavior* (doctoral thesis). Massachusetts Institute of Technology.
- Krafft, P., Shmueli, E., Griffiths, T. L., Tenenbaum, J. B., & Pentland, A. (2021). Bayesian collective learning emerges from heuristic social learning. *Cognition*, *212*, 104469.
- Krafft, P., Zheng, J., Pan, W., Della Penna, N., Altshuler, Y., Shmueli, E., Tenenbaum, J. B., & Pentland, A. (2016). Human collective intelligence as distributed Bayesian inference. arXiv preprint arXiv:1608.01987.
- Krumme, C., Cebrian, M., Pickard, G., & Pentland, A. (2012). Quantifying social influence in an online cultural market. *PLoS ONE*, *7*(5), e33785.
- Laland, K. N., Sterelny, K., Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and effect in biology revisited: Is Mayr's proximate-ultimate dichotomy still useful? *Science*, *334*(6062), 1512–1516.
- Latora, V., Nicosia, V., & Russo, G. (2017). *Complex networks: Principles, methods and applications*. Cambridge, England: Cambridge University Press.
- Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, *108*(22), 9020–9025.
- Lynch, N. A. (1996). *Distributed algorithms*. Burlington, MA: Morgan Kaufmann.
- Martin, D., Hutchison, J., Slessor, G., Urquhart, J., Cunningham, S. J., & Smith, K. (2014). The spontaneous formation of stereotypes via cumulative cultural evolution. *Psychological Science*, *25*(9), 1777–1786.
- Mercier, H., & Morin, O. (2019). Majority rules: How good are we at aggregating convergent opinions? *Evolutionary Human Sciences*, *1*, E6.
- Mesoudi, A. (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences*. Chicago, IL: University of Chicago Press.
- Mesoudi, A. (2016). Cultural evolution: Integrating psychology, evolution and culture. *Current Opinion in Psychology*, *7*, 17–22.
- Mesoudi, A., & O'Brien, M. J. (2008). The cultural transmission of great basin projectile-point technology I: An experimental simulation. *American Antiquity*, *73*(01), 3–28.
- Mesoudi, A., & Thornton, A. (2018). What is cumulative cultural evolution? *Proceedings of the Royal Society Biological Sciences*, *285*(1880), 20180712.
- Mesoudi, A., & Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3489–3501.
- Miton, H., & Charbonneau, M. (2018). Cumulative culture in the laboratory: Methodological and theoretical challenges. *Proceedings of the Royal Society B: Biological Sciences*, *285*(1879), 20180677.
- Morgan, T. J. H., Uomini, N. T., Rendell, L. E., Chouinard-Thuly, L., Street, S. E., Lewis, H. M., Cross, C. P., Evans, C., Kearney, R., de la Torre, I., Whiten, A., & Laland, K. N. (2015). Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nature Communications*, *6*(1), 6029.
- Mozer, M. C., Pashler, H., & Homaei, H. (2008). Optimal predictions in everyday cognition: The wisdom of individuals or crowds? *Cognitive Science*, *32*(7), 1133–1147.
- Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. Cambridge, MA: MIT Press.

- Muthukrishna, M., Shulman, B. W., Vasilescu, V., & Henrich, J. (2014). Sociality influences cultural complexity. *Proceedings of the Royal Society B: Biological sciences*, 281(1774), 20132511.
- Myers, D. G., & Bishop, G. D. (1971). Enhancement of dominant attitudes in group discussion. *Journal of Personality and Social Psychology*, 20(3), 386.
- Nettle, D. (2020). Selection, adaptation, inheritance and design in human culture: the view from the Price equation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1797), 20190358.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608.
- Ortega, P. A. (2011). *A unified framework for resource-bounded autonomous agents interacting with unknown environments* (doctoral thesis). University of Cambridge.
- Osband, I., Van Roy, B., & Russo, D. (2013). (More) efficient reinforcement learning via posterior sampling. In: C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*. NeurIPS Foundation, Vol. 26, (pp. 3003–3011).
- Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance*, 16(2), 366–387.
- Pothos, E. M., Lewandowsky, S., Basieva, I., Barque-Duran, A., Tapper, K., & Khrennikov, A. (2021). Information overload for (bounded) rational agents. *Proceedings of the Royal Society B: Biological Sciences*, 288(1944), 20202957.
- Pratt, S. C., Sumpter, D. J., Mallon, E. B., & Franks, N. R. (2005). An agent-based model of collective nest choice by the ant *Temnothorax albipennis*. *Animal Behaviour*, 70(5), 1023–1036.
- Ravignani, A., Thompson, B., Grossi, T., Delgado, T., & Kirby, S. (2018). Evolving building blocks of rhythm: How human cognition creates music via cultural transmission. *Annals of the New York Academy of Sciences*, 1423(1), 176–187.
- Riedel, C., Kim, Y. J., Gupta, P., Malone, T. W., & Woolley, A. W. (2021). Quantifying collective intelligence in human groups. *Proceedings of the National Academy of Sciences*, 118(21), e2005737118.
- Rzhetsky, A., Foster, J. G., Foster, I. T., & Evans, J. A. (2015). Choosing experiments to accelerate collective discovery. *Proceedings of the National Academy of Sciences*, 112(47), 14569–14574.
- Salhab, R., Ajorlou, A., & Jadbabaie, A. (2020). Social learning with sparse belief samples. In *2020 59th IEEE Conference on Decision and Control (CDC)*, (pp. 1792–1797). Piscataway, NJ: IEEE.
- Sanborn, A., & Griffiths, T. L. (2008). Markov chain Monte Carlo with people. In: *Advances in Neural Information Processing Systems*. NeurIPS Foundation, Vol. 21, (pp. 1265–1272).
- Sasaki, T., & Pratt, S. C. (2011). Emergence of group rationality from irrational individuals. *Behavioral Ecology*, 22(2), 276–281.
- Seeley, T. D., & Buhrman, S. C. (1999). Group decision making in swarms of honey bees. *Behavioral Ecology and Sociobiology*, 45(1), 19–31.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15(3), 233–250.
- Shoham, Y., & Leyton-Brown, K. (2009). *Multiagent systems*. Cambridge, England: Cambridge University Press.
- Silvey, C., Kirby, S., & Smith, K. (2019). Communication increases category structure and alignment only when combined with cultural transmission. *Journal of Memory and Language*, 109, 104051.
- Simon, H. A. (1990). Bounded rationality. In *Utility and probability* (pp. 15–18). Berlin, Germany: Springer.
- Smaldino, P. E., & Richerson, P. J. (2013). Human cumulative cultural evolution as a form of distributed computation. In *Handbook of human computation* (pp. 979–992). New York: Springer.
- Smolla, M., Jansson, F., Lehmann, L., Houkes, W., Weissing, F., Hammerstein, P., Dall, S. R., Kuijper, B., & Enquist, M. (2020). Underappreciated similarities and differences between cultural and biological evolution, and the future of cultural evolution thinking. <https://doi.org/10.31235/osf.io/3xvdw>
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285.
- Toyokawa, W., Whalen, A., & Laland, K. N. (2018). Social learning strategies regulate the wisdom and madness of interactive crowds. *bioRxiv*, 326637.

- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599–637.
- Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, 19(7), 645–647.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14(1), 101–118.
- Wisdom, T. N., Song, X., & Goldstone, R. L. (2013). Social learning strategies in networked groups. *Cognitive Science*, 37(8), 1383–1425.
- Wojtowicz, Z., & DeDeo, S. (2020). From probability to consilience: How explanatory values implement Bayesian reasoning. *Trends in Cognitive Sciences*, 4(12), 981–993.
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688.
- Xu, J., & Griffiths, T. L. (2010). A rational analysis of the effects of memory biases on serial reproduction. *Cognitive Psychology*, 60(2), 107–126.

Supporting Information

Additional Supporting Information may be found
online in the supporting information tab of this article:
Supporting Information
Supporting Information