Population-level amplification of perceptual bias

Mathew D. Hardy[1 +] (mdhardy@princeton.edu)
Bill Thompson[2 +] (wdt@princeton.edu)
Peaks M. Krafft[3] (p.krafft@oii.ox.ac.uk)
Thomas L. Griffiths[1,2] (tomg@princeton.edu)
[1]Department of Psychology, Princeton University
[2]Department of Computer Science, Princeton University
[3]Oxford Internet Institute, University of Oxford
[+]These authors contributed equally.

Abstract

A longstanding conjecture that has been difficult to test holds that social interactions amplify the effects of people's biases. We tested this conjecture in a perceptual decision-making paradigm. First, we formalized the algorithmic structure of decision making in networked crowds when individuals' perceptions are biased by their utilities. Our analysis predicts that even weak cognitive biases can be amplified by social interaction. We tested this prediction in a large networked behavioral experiment. Using a monetary incentive structure to induce a bias known as motivated perception, we manipulated the presence of a weak cognitive bias in social and asocial populations. Social decision making increased participants' perceptual accuracy relative to an asocial baseline. However, social decision making also led to significantly amplified rates of motivated perception, confirming the prediction that shared cognitive biases can be amplified in social networks.

Keywords: bias; decision making; perception; social networks; social modeling; Bayesian modeling

Introduction

How do you choose where to eat, who to vote for, or whether to trust the scientific consensus on global warming? Many of our decisions both large and small are influenced by interactions with other people in increasingly complex social networks. However, it remains unclear how social interactions influence the cognitive biases that also shape human decisions. On the one hand, social interactions can increase the information people have access to (Couzin, 2007), potentially allowing them to reduce their bias. On the other, the dynamics of social processes can compound small effects, potentially leading to socially-induced bias amplification (Anderson & Holt, 1997).

Understanding how social networks impact our decisions is increasingly important to both social scientists and public policy makers. Of particular interest to both groups has been motivated reasoning, a bias where people overweight information that confirms existing or positive beliefs (Kunda, 1990; Leong et al., 2019). Motivated reasoning is thought to contribute to political and climate-change belief polarization (Mutz, 2006; Hart & Nisbet, 2012). As polarization increases (Abramowitz & Saunders, 2008), understanding the impact of social interactions on this bias is an important challenge (Bail et al., 2018).

Experimentally investigating bias amplification presents two significant challenges. The first challenge is that large-scale, randomized, replicated social experiments are difficult to conduct. We used Dallinger[1] to develop a scalable experimental paradigm that facilitates recruitment of dozens of generation-structured experimental social networks (Fig. 1). The second challenge is that understanding bias amplification requires a formal framework that connects psychological theory with mathematical approaches to network dynamics. To address this difficulty, we extended a resource-rational model of utility-weighted decision making to the population setting.

We used this framework to investigate the impact of social influence on motivated reasoning in networked populations on a perceptual task (*motivated perception*). We focused on generation-structured social networks to capture the temporal dependencies and open-ended structure characteristic of many modern social interactions. Our analysis shows that utility-weighted decision making can break the connection between social decision making and distributed algorithms for Bayesian inference (Krafft et al., 2016). This leads to a potentially counter-intuitive prediction: social influence in networked groups can both increase decision-making accuracy and decrease objectivity, by amplifying the motivated perception bias. We tested these predictions in a large behavioral experiment. We manipulated the presence of an induced perceptual bias in social and asocial populations. In line with our model predictions, participants in social treatments made more accurate perceptual judgments than participants in asocial treatments. However, participants in the social bias-inducing treatment also made judgments that were more biased than an asocial control. These results confirm the prediction that social interactions can amplify people's biases, even on simple perceptual problems.

Background

Motivated reasoning & perception

People's perceptions are often colored by their desires, preferences, and expectations (Bruner & Goodman, 1947). For example, sports fans tend to interpret ambiguous fouls as reflecting positively on their preferred team, (Hastorf & Cantril, 1954) and also preferentially bet on their home team (Staněk, 2017). In more controlled experiments, people's preferences have been shown to influence their perceptions of letters and numbers (Balcetis & Dunning, 2006) as well as faces and scenes (Leong et al., 2019). This bias reflects differences

[1]Available at `https://github.com/Dallinger/Dallinger`

in actual, not just reported, perceptions (Leong et al., 2019). This bias can have substantial real-world impacts. Research has shown that motivated reasoning is an important contributor to the public divide in people's beliefs about climate change (Hart & Nisbet, 2012). Furthermore, minimizing the effects of motivated reasoning appears to be an effective method of reducing polarization (Arceneaux & Vander Wielen, 2017).

Group decision making

Models of collective intelligence (Couzin, 2007) and network dynamics (Watts & Strogatz, 1998; Kearns et al., 2006) have shown clear information processing advantages in networked groups. This social advantage has been replicated in some experiments (Mason & Watts, 2012), but refuted in others (Anderson & Holt, 1997; Toyokawa et al., 2019). These experiments typically focus on real-time interaction in *closed-group* (Mesoudi & Whiten, 2008) populations of a fixed set of individuals, rather than on open-ended, longer term social processes with evolving populations.

Generation-structured networks have been extensively studied in models of cultural evolution (Mesoudi & Whiten, 2008) and social learning (Rendell et al., 2010). These models have shown that individual biases may be amplified over generations (Boyd & Richerson, 1988). Models of cultural evolution are typically formalized using frameworks adapted from population genetics, which can be difficult to connect to psychological theory (Heyes, 2018). One way to bridge population-level processes with psychological theory is to analyze populations of Bayesian agents (Griffiths & Kalish, 2007). This approach has been successful in other cognitive domains such as language (Griffiths & Kalish, 2007) and category systems (Canini et al., 2014). Bayesian models also predict amplification of cognitive biases (Griffiths & Kalish, 2007). However, existing models focus on problems of *C-induction* (Chater & Christiansen, 2010), in which the computational problem facing individuals is to coordinate on conventions. Here, we extend this paradigm to the domain of social decision-making, and thus to problems of *N-induction*, in which populations must solve a "natural", externally defined computational problem.

Model: Utility-weighted social filtering

We analyzed social decision making in generation-structured networks in which n individuals at generation t make a decision about a unknown quantity θ after observing 1) decisions of n peers at generation $t-1$ and 2) independent evidence y^t.

Social transmission as particle filtering

Previous models (Krafft et al., 2016) introduced the idea that social observations play the role of a prior distribution on people's beliefs, connecting social learning with rational models of inference and decision making. This leads to a formal relationship between social populations and a class of algorithms known as Sequential Monte Carlo methods or *particle filtering* (Doucet et al., 2001). This relationship can be expressed

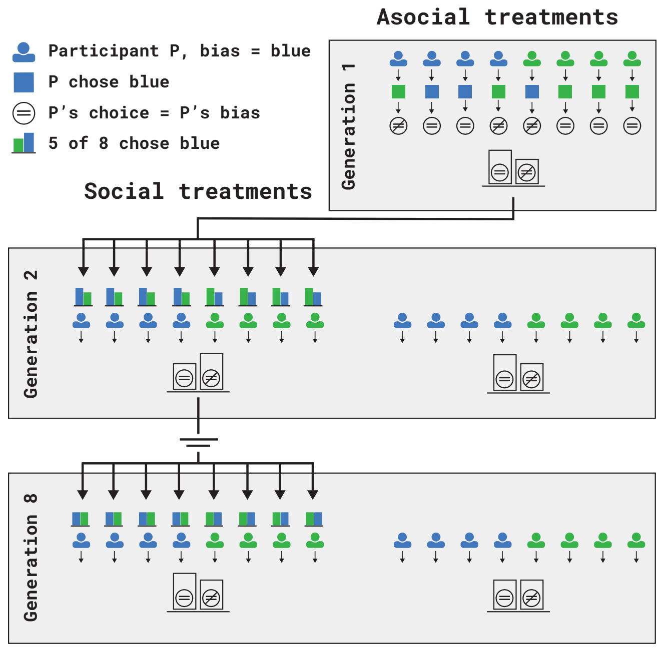

Figure 1: Transmission dynamics for a single set of treatment replications on a single trial for participants in bias-inducing treatments. Participant icon color indicates marked color. Participants in the first generation were assigned to an asocial treatment (top right). A participant's choice was coded as a binary measure indicating whether they chose their marked color. Participants in generation t were recruited after all participants in generation $t-1$ completed the experiment. Participants in social treatments (left) viewed choices made by participants in the previous generation.

by assuming individuals follow a sample-based algorithm for approximate Bayesian inference. Individual i at generation t makes a decision θ^i by first sampling a finite set of K hypotheses $\hat{\theta}^i_1, \ldots, \hat{\theta}^i_K$. These initial samples are drawn from a distribution $\hat{\theta}^i_k \sim p_{t-1}(\theta)$ that reflects the decisions made by the previous generation. Individuals re-weight these samples in light of data y^t, before sampling a single hypothesis according to the updated weights. The weight assigned to each hypothesis is $w^i_k = f(y^t|\theta^i_k)$. Under a Bayesian interpretation, $f(y|\theta)$ represents a likelihood function capturing the individual's generative model for data. The distribution over beliefs for individual i is $p\left(\theta^i = \hat{\theta}^i_k\right) = w^i_k / \sum^K_j w^j_k$. Under the simplifying assumption that individuals are exchangeable (all individuals at generation t observe the same choices from the previous generation and the same evidence y^t), the distribution of hypotheses at each generation is:

$$p_t(\theta) \propto f(y^t|\theta^i_k)p_{t-1}(\theta) = f(y^{1:t}|\theta^i_k)p_1(\theta^i_k), \qquad (1)$$

where $y^{1:t}$ denotes the evidence accumulated from generation 1 to t and $p_1(\theta)$ is the initial distribution. This is a particle filtering algorithm for approximate Bayesian inference, resulting in samples from the posterior distribution $\pi(\theta|y^{1:t})$. This connection explains why social learning can improve people's inferences in generation-structured settings.

Utility-weighted particle filter

In order to model motivated perception, we extend this framework to contexts in which people incur differential utilities $u(\theta)$. In these contexts, the computational problem faced by individuals is to estimate the *expected utility* of the unknown quantity θ. Previous research (Lieder et al., 2018) showed that this problem can be solved using a resource-rational sampling algorithm known as *utility-weighted sampling* (UWS).

The key idea behind UWS is that deliberation should account for the magnitude of a hypotheses' utilities, even if these hypotheses have low probability (such as winning the lottery). When the goal is to estimate expected utility from a limited set of samples, the optimal sampling distribution reflects a combination of utility and probability (Lieder et al., 2018). The UWS algorithm samples hypotheses from this distribution in order to calculate a weighted expectation of utility. This approach to decision making provides a resource-rational foundation for utility-based cognitive biases like motivated reasoning and perception. In the limit of a single sample, the UWS decision-making algorithm implies a small change to the hypothesis re-weighting procedure: after sampling a finite set of hypotheses from the inherited distribution $p_{t-1}(\theta)$, re-weight those samples according to $w_k^i \equiv f(y|\theta_k^i) \cdot |u(\theta_k^i)|$. Here $u(\theta)$ captures individuals' shared utilities. Under this model, the distribution of beliefs at generation t given by Equation 1 becomes:

$$p_t(\theta) \propto f(y^t|\theta)|u(\theta)|p_{t-1}(\theta) = f(y^{1:t}|\theta)|u(\theta)|^t p_0(\theta)\,. \quad (2)$$

The recursive filter defined in Equation 2 is not an algorithm for performing distributed Bayesian inference. In this algorithm, utility biases accumulate because they are reintroduced at every generation. While utility-weighted decision making is a functional solution to the computational problem faced by individuals, it breaks the connection between social decision making and distributed Bayesian inference. If the goal of collective decision making is to accumulate accurate beliefs, the *utility-weighted particle filter* (UWPF) defined by Equation 2 is not a good solution to this computational problem because the algorithm is biased (see Thrun et al. (2002) for a related analysis of biased particle filters). This analysis makes a potentially counter-intuitive prediction. On the one hand, evidence accumulates over time even if the filter is biased – being part of a social network should improve people's decisions. However, on the other hand, utility biases also accumulate over time, leading to bias amplification. This formulation connects directly to a psychological theory of decision making, allowing us to test these predictions experimentally.

Experiment

Method

We tested the predictions of the utility-weighted particle filter on a numerosity estimation task (Kao et al., 2018). Participants briefly viewed displays of blue and green dots and chose which color was more numerous. The experiment used a 2×2 factorial design that varied the presence of an induced perceptual bias and access to social information.

Participants We recruited 2,617 participants from Amazon's Mechanical Turk, limiting our study to those living in the United States. Participants received a base payment of \$0.65 for participation, plus an average bonus of \$0.65. Participants that failed to pass a short comprehension test were excluded from the experiment. Our recruitment algorithm used planned over-recruitment to accelerate data collection (over-recruited participants were excluded from analyses) enabling us to efficiently recruit our preregistered target sample size of 2,400 participants.

Treatments Participants were recruited in batches (generations). Generation t was recruited after all participants in generation $t-1$ completed the experiment. Participants were assigned to one of four treatments using block randomization - asocial without induced bias (ASO-control), asocial with induced bias (ASO-bias), social without induced bias (SOC-control), and social with induced bias (SOC-bias).

All participants received 50 points for every correct judgment. Participants in social treatments viewed the choices made by a set of participants in the previous generation (see Fig. 2). Participants in bias-inducing treatments were assigned a *marked color* of blue or green. Participants received a 1 point bonus for every dot of their marked color. Participants received this marked-color bonus regardless of whether their choice was right or wrong. This reward was included to induce a motivated perception bias towards the participant's marked color, adapting a procedure used in related work (Leong et al., 2019). Participants in non bias-inducing treatments were assigned a marked color in the same manner, but were not aware of this assignment. In every generation and in every treatment replication (described in detail below), marked color was randomly assigned via block randomization. Participants in non bias-inducing treatments were given an additional bonus for completing the experiment so that participants in all treatments earned the same unconditional reward. At the end of the experiment, each participant's points were paid as a bonus with 10 points equal to \$0.01.

Generation structure We recruited 8 generations of participants. 10 treatment replications were conducted for each treatment. Participants in generation 1 were assigned to one of the two asocial treatments (ASO-bias and ASO-control). In each generation, 8 participants were assigned to each treatment replication, corresponding to a sample size of 160 participants for generation 1 and 320 participants per generation for generations 2-8. Participants in social treatments at generation 2 observed the decisions made by generation 1 participants in same treatment replication and bias treatment. For example, a SOC-control participant in generation 2 treatment replication 5 viewed the choices of the 8 ASO-control participants in generation 1 treatment replication 5 (see Fig. 1).

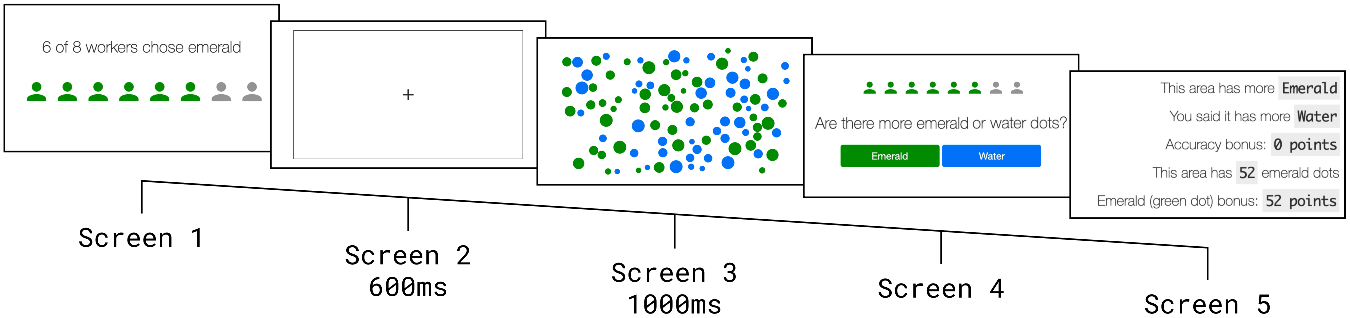

Figure 2: Experimental interface for participants in the social bias-inducing treatment with a green marked color during practice rounds. Screen 5 (feedback) is not presented after test rounds.

Procedure Participants completed 2 practice trials and 8 test trials, resulting in a dataset of 24,000 decisions. Trial order was randomized for practice and test trials. Participants received immediate feedback after practice trials but not after test trials. Test trials included two difficulty levels (medium and hard) corresponding to marked-color proportions of {0.48, 0.52}, and {0.49, 0.51}, respectively. Stimuli were preceded by a fixation cross and bounding box for 600 milliseconds. Each stimulus consisted of 100 randomly positioned and sized blue and green dots displayed for 1 second. Dot colors were matched for luminance in LAB color space. Dot positioning and sizing varied across generations and across treatment replications, but were constant across treatments.

Counterbalancing of marked color required re-coding dot color and social information to prevent unwanted confounds from color-perception biases. The color assigned to each dot was determined by the participant's marked color. For example, a trial with marked-color proportion of 0.52 included 52 green dots for participants with a marked color of green, and 52 blue dots for participants with a marked color of blue. Social information presented number of participants making the majority decision and was also recoded into the current participant's (P) marked color. For example, if 6 of 8 participants in the previous generation chose their marked color, then if P's marked color was blue, she was told that 6 of 8 participants in the previous generation chose blue. However, if only 3 of 8 participants chose their marked color and P's marked color was green, then P was told that 5 of 8 participants chose blue. Ties were broken via simple randomization.

Participants in all treatments were informed that they were working for an imaginary mining company looking for valuable gemstones. Participants whose marked color was blue in bias-inducing treatments were told that they were looking for blue sapphires in green grass, and would judge whether there were more sapphire dots or more grass dots. Participants whose marked color was green in bias-inducing treatments were told they were looking for green emeralds in blue water, and would judge whether there were more emerald or water dots. Participants in non bias-inducing treatments were told they were looking for blue sapphires and green emeralds.

Results

Analyses were performed on 19,200 test trial decisions. Accuracy and bias were binary coded for each observation. Accuracy was coded as successful when a participant chose correctly. Bias was coded as successful when a participant chose their marked color. We performed preregistered mixed effects regressions with a logit link predicting accuracy and bias.[2] All regressions included fixed effects for condition and random intercepts for treatment replication.[3] The estimated fixed effects and their standard errors are shown in Fig. 3. To determine whether bias or accuracy differed significantly between two treatments, we performed a likelihood ratio test between an unrestricted logistic regression predicting bias or accuracy with fixed effects for all treatments, and a restricted logistic regression where the two treatments were coded as a single treatment.

Bias Participants in the asocial bias-inducing treatment were significantly more biased than participants in both the asocial no-bias treatment (ASO-bias marked-color choice proportion: 0.562; ASO-control marked-color choice proportion: 0.526; $\chi^2(1) = 4.14, p = 0.042$) and the social no-bias treatment (SOC-control marked-color choice proportion: 0.513; $\chi^2(1) = 7.31, p = 0.007$). Participants in the social bias-inducing treatment were significantly more biased than participants in the asocial bias-inducing treatment (SOC-bias marked-color choice proportion: 0.618; $\chi^2(1) = 9.69, p = 0.002$), the asocial no-bias treatment ($\chi^2(1) = 21.77, p < 0.0001$), and the social-no bias treatment ($\chi^2(1) = 26.34, p < 0.0001$). There was no statistically significant difference in bias between participants in the asocial no-bias treatment and social no-bias treatment ($\chi^2(1) = 0.58, p = 0.445$).

[2]Preregistration available at `https://osf.io/yth5r`

[3]Our preregistered analysis included participant random effects which we omitted as they led to singular fits.

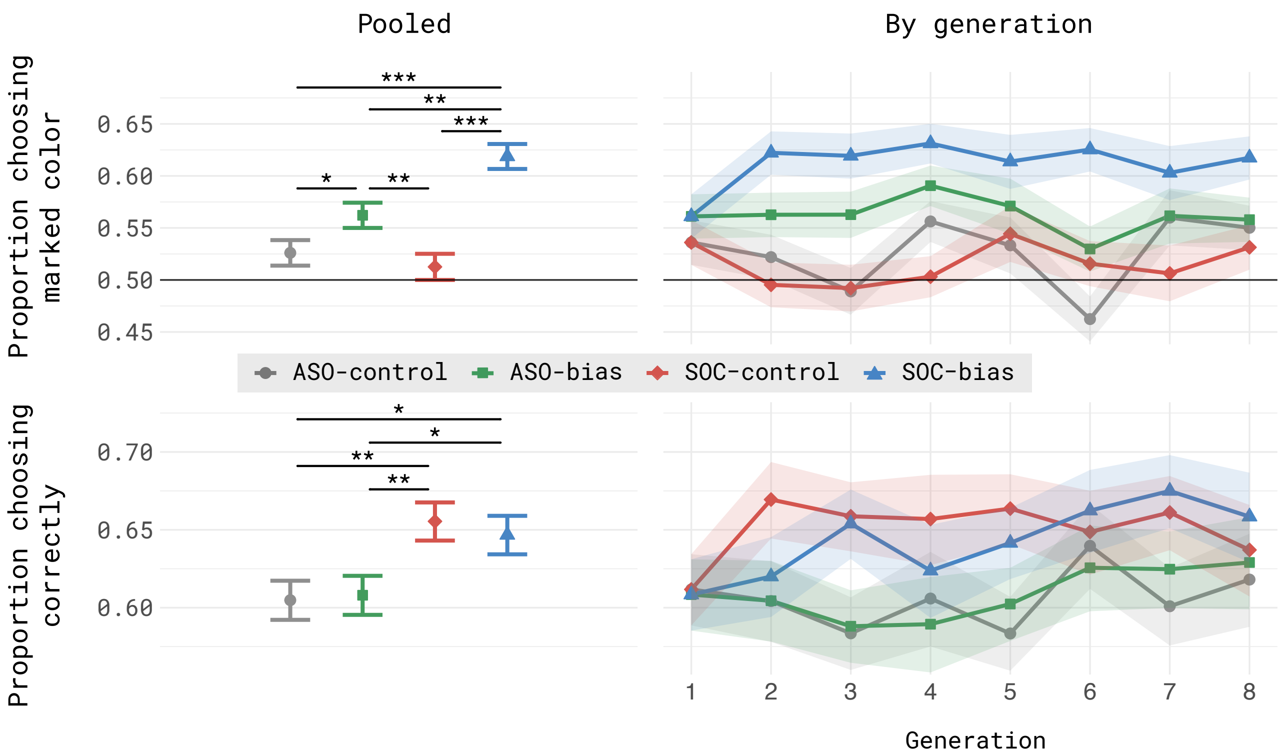

Figure 3: Experimental results. The top plots show the proportion of participants choosing their marked color. More biased participants will choose their marked color more often. The bottom plots show the participants' accuracy. Plots on the left show these estimates pooled across generations, and plots on the right show these estimates broken down by generation. Error bars on the left plots and shading on the right plots show standard error estimates derived from the mixed-effect logistic regressions. While only asocial participants were recruited in generation 1, the generation plots show generation 1 social participants to illustrate the yoking procedure described in the Methods section.

Accuracy Participants in the social no-bias treatment chose correctly significantly more often than participants in both the asocial no-bias treatment (SOC-control proportion correct: 0.654; ASO-control proportion correct: 0.604; $\chi^2(1) = 7.55, p = 0.006$) and the asocial biasing-inducing treatment (ASO-bias proportion correct: 0.607; $\chi^2(1) = 6.72, p = 0.01$). There was no statistically significant difference in accuracy between participants in the two social treatments (SOC-bias proportion correct: 0.646; $\chi^2(1) = 0.25, p = 0.617$). Participants in the social bias-inducing treatment made judgments that were significantly more accurate than participants in the asocial no-bias treatment ($\chi^2(1) = 5.27, p = 0.022$) and participants in the asocial bias-inducing treatment ($\chi^2(1) = 4.56, p = 0.033$). There was no statistically significant difference in accuracy between participants in the two asocial treatments ($\chi^2(1) = 0.03, p = 0.858$).

Exploratory analyses We investigated how SOC-bias participants were both more biased and more accurate relative to ASO-control participants. To do so, we performed similar regressions to those above predicting accuracy, but partitioning the data into trials where the true answer matched a participant's marked color and those where it did not. SOC-control participants made judgments that were not significantly more accurate than ASO-control when marked color matched the correct answer (ASO-control proportion correct: 0.630; SOC-control proportion correct: 0.667; $\chi^2(1) = 1.84, p = 0.175$). However, SOC-control participants were more accurate than ASO-control participants when their marked color did not match the correct answer (ASO-control proportion correct: 0.579; SOC-control proportion correct: 0.642; $\chi^2(1) = 7.58, p = 0.006$). SOC-bias participants made judgments that were more accurate than ASO-control participants when their marked color matched the correct answer (SOC-bias proportion correct: 0.764; $\chi^2(1) = 20.97, p < 0.0001$), but less accurate when marked color did not match the correct answer (SOC-bias proportion correct: 0.528; $\chi^2(1) = 4.64, p = 0.031$).

Conclusion

In line with our model predictions and previous research, participants in social networks made judgments that were more accurate than participants who performed the task individually. However, participants made judgments that were more biased when judgments were transmitted through social networks. Social information increased participants' performance but amplified motivated perception. These results confirm the predictions of earlier mathematical frameworks which identified the potential for bias amplification. We extended these models to the domain of social decision making to test this prediction in a large behavioural experiment. Although the perceptual bias was weak in individuals, social transmission significantly amplified its effects. Our study was limited to simple social networks with relatively small generation sizes. Priorities for future work include developing methods to combat bias amplification and applying this framework to more naturalistic domains.

References

Abramowitz, A. I., & Saunders, K. L. (2008). Is polarization a myth? *The Journal of Politics*, *70*(2), 542–555.

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, *87*(5), 847–862.

Arceneaux, K., & Vander Wielen, R. J. (2017). *Taming intuition: How reflection minimizes partisan reasoning and promotes democratic accountability*. Cambridge University Press.

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., … Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, *115*(37), 9216–9221.

Balcetis, E., & Dunning, D. (2006). See what you want to see: motivational influences on visual perception. *Journal of Personality and Social Psychology*, *91*(4).

Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process*. University of Chicago press.

Bruner, J. S., & Goodman, C. C. (1947). Value and need as organizing factors in perception. *The Journal of Abnormal and Social Psychology*, *42*(1).

Canini, K. R., Griffiths, T. L., Vanpaemel, W., & Kalish, M. L. (2014). Revealing human inductive biases for category learning by simulating cultural transmission. *Psychonomic Bulletin & Review*, *21*(3), 785–793.

Chater, N., & Christiansen, M. H. (2010). Language acquisition meets language evolution. *Cognitive Science*, *34*(7), 1131–1157.

Couzin, I. (2007). Collective minds. *Nature*, *445*(7129), 715–715.

Doucet, A., de Freitas, N., & Gordon, N. (2001). *Sequential Monte Carlo methods in practice*. Springer.

Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive science*, *31*(3), 441–480.

Hart, P. S., & Nisbet, E. C. (2012). Boomerang effects in science communication: How motivated reasoning and identity cues amplify opinion polarization about climate mitigation policies. *Communication Research*, *39*(6), 701-723.

Hastorf, A. H., & Cantril, H. (1954). They saw a game; a case study. *The Journal of Abnormal and Social Psychology*, *49*(1).

Heyes, C. (2018). Enquire within: Cultural evolution and cognitive science. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1743), 20170051.

Kao, A. B., Berdahl, A. M., Hartnett, A. T., Lutz, M. J., Bak-Coleman, J. B., Ioannou, C. C., … Couzin, I. D. (2018). Counteracting estimation bias and social influence to improve the wisdom of crowds. *Journal of The Royal Society Interface*, *15*(141), 20180130.

Kearns, M., Suri, S., & Montfort, N. (2006). An experimental study of the coloring problem on human subject networks. *Science*, *313*(5788), 824–827.

Krafft, P. M., Zheng, J., Pan, W., Della Penna, N., Altshuler, Y., Shmueli, E., … Pentland, A. (2016). Human collective intelligence as distributed Bayesian inference. *arXiv preprint arXiv:1608.01987*.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*(3).

Leong, Y. C., Hughes, B. L., Wang, Y., & Zaki, J. (2019). Neurocomputational mechanisms underlying motivated seeing. *bioRxiv*, 364836.

Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, *125*(1).

Mason, W., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, *109*(3), 764–769.

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.

Mutz, D. C. (2006). How the mass media divide us. In P. S. Nivola & D. W. Brady (Eds.), *Red and Blue Nation? Characteristics and Causes of America's Polarized Politics* (p. 223 - 248). Washington, D.C.: Brookings Institution Press.

Rendell, L., Fogarty, L., & Laland, K. N. (2010). Rogers'paradox recast and resolved: Population structure and the evolution of social learning strategies. *Evolution: International Journal of Organic Evolution*, *64*(2), 534–548.

Staněk, R. (2017). Home bias in sport betting: Evidence from czech betting market. *Judgment and Decision Making*, *12*(2).

Thrun, S., Langford, J., & Verma, V. (2002). Risk sensitive particle filters. In *Advances in neural information processing systems* (pp. 961–968).

Toyokawa, W., Whalen, A., & Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, *3*(2), 183–193.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world'networks. *Nature*, *393*(6684).