

Belief revision in a micro-social network: Modeling sensitivity to statistical dependencies in social learning

Jan-Philipp Fränken, Nikolaos C. Theodoropoulos, Adam B. Moore, Neil R. Bramley

{jp.franken,nik.theodoropoulos,adam.moore,neil.bramley}@ed.ac.uk

Department of Psychology, 7 George Square

Edinburgh EH89JZ, United Kingdom

Abstract

In both professional domains and everyday life, people must integrate their own experience with reports from social network peers to form and update their beliefs. It is therefore important to understand to what extent people accommodate the statistical dependencies that give rise to correlated belief reports in social networks. We investigate adults' ability to integrate social evidence appropriately in a political scenario, varying the dependence between the sources of network peers' beliefs. Using a novel interface that allows participants to express their probabilistic beliefs visually, we compare participants against a normative Bayesian standard. We find that they distinguish the value of evidence from dependent versus independent sources, but that they also treated social sources as substantially weaker evidence than direct experience. The value of our elicitation methodology and the implications of our results for modeling human-like belief revision in social networks are discussed.

Keywords: social networks; probabilistic beliefs; sequential belief updating; information cascades; Bayesian modeling

Introduction

We live and learn in a “society of minds” (Minsky, 1988). This means that we form beliefs not just on the basis of our own observations (and prior expectations) but also based on the beliefs communicated by our neighbors in our social network. For instance, interview panel members will typically discuss job applicants even after having seen mostly the same application materials and interviews, making it difficult to distinguish individual panel members' (prior) judgments of candidates' abilities from collective judgments formed on the basis of the shared evidence. Similarly, imagine you have read about two political campaign strategies, each proposed by a different candidate. If initially you find both strategies equally compelling, resulting in uncertainty about which of the two candidates you would like to support, you might well seek out new information about the candidates by talking to friends. If your friends base their beliefs partially on reading the same articles, how should you weigh their opinions?

The above examples illustrate one source of statistical dependency between opinions in a social network: shared information originating from the same source. If we are to understand learning in a social world, we must understand how people deal with such statistical dependencies while integrating their direct observations from the environment with the communicated beliefs of their social network peers. Investigating how information spreads through social networks and how statistical dependencies affect the formation of people's beliefs is thus a key issue for cognitive science with




Figure 1: Illustration of network conditions. t_0 : neighbors A-C form beliefs given data. t_1 : neighbors update beliefs based on interaction (sequential case only). t_2 : target updates belief.

implications for e.g., the study of misinformation and echo chambers (Bikhchandani, Hirshleifer, & Welch, 1992; Watts, 2002; Whalen, Griffiths, & Buchsbaum, 2018; Madsen, Bailey, & Pilditch, 2018), the dynamics of micro-targeting (Madsen & Pilditch, 2018), or advocacy organisations' attempts to shape public debate (Bail, 2016).

Here, we investigate how people integrate information based on statistical (in)dependencies underlying the beliefs of three social network peers in a fictitious political context. We first introduce a simple Bayesian model of belief revision to account for the normative case. Building on previous work on sensitivity to shared information in social learning (Whalen et al., 2018), we compare three different conditions (see Fig. 1). The first (Fig. 1a) serves as baseline in a sense that statistical *independence* between the beliefs of network peers is induced by the cover story. In the second condition, independence is violated as participants are told that the three network neighbors form their beliefs sequentially, meaning that the belief of the neighbor that formed their belief last (neighbor C) contains all information gathered by the others (Fig. 1b). In a third condition, social network peers form beliefs on the basis of shared evidence (Fig. 1c). Dependencies between beliefs in condition 2 differ from dependencies in condition 3 in the sense that neighbor C is the most relevant source since their belief already incorporates the beliefs of the other neighbors. We report on a behavioral experiment that investigated how subjects update prior beliefs under these conditions.

Information Cascades and Probabilistic Beliefs Information cascades—spreading of beliefs through networks—can produce maladaptive collective outcomes even as agents incorporate information from network neighbors in individually rational ways (Bikhchandani et al., 1992). This inherent sus-

ceptibility of social networks towards information cascades is supported by research showing that information cascades occur in simulated social networks where agents are furnished with an individually rational cognitive architecture (Fränken & Pilditch, 2020). Similar results have been obtained from empirical analyses of social media data, which showed that users cluster into communities dominated by like-minded others, resulting in proliferation of unsubstantiated beliefs or conspiratorial thinking (Del Vicario et al., 2016).

Previous models of information cascades assumed that people settle on particular beliefs by maximizing over subjective probabilities, thus leading to degradation of the information transmitted (Bikhchandani et al., 1992; Pilditch, 2017). Additionally, simulation-based work by Pilditch (2017) (see also Madsen et al., 2018) did not account for potential dependencies underlying the beliefs of network peers. The assumption of independent beliefs among network peers is frequently used in general models of opinion dynamics and consensus generation, where people’s belief revision processes have been modeled by combining their initial beliefs with the weighted average of neighboring beliefs (Hegselmann, Krause, et al., 2002; Lorenz, 2006, 2007; Becker, Brackbill, & Centola, 2017). Expanding this body of literature, recent empirical results have shown that people’s social learning strategies are adaptive - accounting for statistical dependencies underlying the beliefs in their social networks (Whalen et al., 2018). In addition to maximizing, Whalen et al. (2018) tested the assumption of “probability matching”, which assumes that people settle on beliefs stochastically, drawing particular conclusions proportional to the posterior probability of a belief (Shanks, Tunney, & McCarthy, 2002).

In the present work, we make neither assumption (matching or maximizing), instead empirically exploring a setting in which agents communicate their full probabilistic beliefs. Using probabilistic beliefs allows us to explore the influence of communicated certainty—defined here as the precision of the belief distribution—and its related probabilistic quantity confidence—the probability that a particular choice is correct. These are at the core notion of an rational agent (Pouget, Drugowitsch, & Kepecs, 2016; Fleming & Daw, 2017) and thus play a crucial role during the integration of social information to update prior beliefs (see e.g., De Martino, Bobadilla-Suarez, Nouguchi, Sharot, & Love, 2017).

Normative Framework

We explore a general sequential belief updating setting in which people first gather evidence by themselves (i.e., asocial information) before reporting their initial belief about the relative competence of two fictitious competing political candidates. Evidence comes in the form of binomial “performance tests” that result in either a 0 = loss or 1 = win for each candidate. We thus model evolving beliefs about the relative competence of candidate A over B using the beta probability distribution $X \sim \text{Beta}(\alpha, \beta)$. Following Bayes’ rule, the initial posterior probability of a belief or hypothesis $p(h)$, given

asocial information, d , thus corresponds to the normalized product of the likelihood $p(d|h)$ and the prior $p(h)$:

$$p(h|d) \propto p(d|h)p(h) \quad (1)$$

In our computational model, we use an initially uniform $\text{Beta}(1,1)$ prior which is conjugate to the binomial likelihood $\binom{n}{k} p^k (1-p)^{n-k}$, allowing us to model belief updating straightforwardly using the analytical posterior $\text{Beta}(1+k, 1+n-k)$. For example, observing data $D = \{0, 1, 1\}$ where $k = 2$ successes in $n = 3$ binomial trials, the posterior is $X \sim \text{Beta}(3, 2)$ with a mean of $\frac{3}{5}$. This reflects the nature of subjects’ beliefs about the two candidates, which include an overall preference for a candidate (if mean $\frac{\alpha}{\alpha+\beta}$ is $<$ or $>$.5) and a measurement of certainty (precision of the beta distribution, given by $\frac{(\alpha+\beta)^2}{\alpha\beta}$). The model then also gives clear qualitative (directional) and quantitative predictions for how learners should update their belief upon observing the beliefs of other network neighbors depending on the condition. As in Equation 1, we use Bayes’ theorem to model how people should integrate the beliefs (i.e., social information) from their network neighbors s :

$$p(h|s_1, \dots, s_n) \propto p(s_1, \dots, s_n|h)p(h) \quad (2)$$

Assuming that the beliefs of neighbors are perfectly independent, transparent and trustworthy (Fig. 1a), the target’s posterior after incorporating the beliefs of their peers should simply be a new beta distribution with the parameters $\text{Beta}(n_0 + k_A + k_B + k_C, n_0 + n_A + n_B + n_C - k_A - k_B - k_C)$ where n_0 and k_0 are from their prior. If neighbors are sequentially dependent, in the sense that A communicated their belief to B who then saw more data and communicated to C, the aggregated parameters of the posterior distribution should be based on the neighbor that formed their belief last (i.e., neighbor C in Fig. 1b). The normative posterior for the sequential case is thus equal to $\text{Beta}(n_0 + k_C, n_0 + n_C - k_C)$. Finally, if sources are dependent in the sense that their beliefs are based on at least partially shared information (Fig. 1c), the normative model provides an upper- and a lower bound for the revised posterior. The upper bound is equal to the independent case, and it assumes that none of the neighbors’ beliefs were influenced by the shared data (i.e., $D = \{\}$). Conceptually, this can be compared to a scenario in which panel members ignore all shared application materials and interviews, evaluating the candidate’s performance ability entirely based on their prior beliefs.

As we do not vary the parameters of peers between conditions in our experiment, the only source of variation in updating can be attributed to manipulating the dependencies between neighbors. Thus, the model lower bound is equal to subtracting the lowest α and β parameters from the aggregate parameters. For the present experiment, we do not specify on how much neighbors were influenced by the shared data. Thus, assuming all possible combinations of overlap equi-probable, we model the normative impact of shared information on belief

updating as having a magnitude intermediate between strictly independent information (higher magnitude) and sequentially updated beliefs (lower magnitude). Based on this framework (Fig. 1), we derived the following qualitative (directional) predictions: The difference between subjects initial- and revised posterior probabilistic beliefs will be smaller when the beliefs of social network peers are dependent as compared to the independent condition (1). The dependent case of sequentially updated beliefs will result in a smaller update of prior beliefs as compared to shared information (2).

Experiment

Participants Participants (N = 79, range: 21 - 69 years, mean = 39.89, SD = 12.97, 35 female) were recruited and tested through Amazon’s Mechanical Turk. Participants were native English speakers based in the United States. They were paid \$1.75 for their time (mean = 17.49 min, SD = 6.91 min).

Task Description and Measures¹ Participants imagined being a political consultant travelling around the US to help local branches of their political party decide between two fictitious competing candidates most suitable for public office. To do so, they imagined that they were travelling to three different cities, with two different candidates competing in each city. Prior to the main task, participants completed a short training phase and comprehension quiz to ensure that they understood how to provide their beliefs using the interface shown in Fig. 2. Specifically, participants used two response sliders to provide their full probabilistic beliefs, one controlling the mean of the density (belief slider) and one controlling the log precision (certainty slider). The response sliders ranged from 1-99; where a belief of 1 means full support for the left candidate, 50 is neutral, and 99 full support for the right candidate. A certainty of 1 is the lowest possible certainty and a certainty of 99 is the highest possible certainty. The resulting density was dynamically displayed to participants as they selected their response. The range of allowable values was restricted to ensure the belief function was concave (i.e., $\alpha \geq 1$ or $\beta \geq 1$).




Figure 2: Interface for rating belief and certainty.

Within each city, the order of steps was: (1) obtaining prior belief based on observing asocial evidence → (2) learning about the beliefs of social network neighbors under consideration of statistical (in)dependence → (3) providing final

posterior belief. We set the scene using an uniform prior belief $X \sim \text{Beta}(1, 1)$ telling subjects that the two competing candidates were tested on two initial tests (each winning one of them) prior to the arrival of participants. Participants then observed the performance of the two competing candidates across two additional independent selection tests assessing different qualifications not covered in the initial tests.

Following observation of asocial information, participants rated their prior belief and certainty in the relative suitability of the two competing political candidates for public office. The procedure was identical for each condition. After the initial assessment phase, subjects were shown the belief and certainty ratings of three social network neighbors (i.e., social information; see Fig. 3). The network neighbors were described as three locals that were likely voters from a subject’s political party who had learned about the candidates during debates in their local town-hall. Each city included a different cover story about the relationship between the three locals matching either statistical independence, sequential dependence, or shared information. After learning about the beliefs of locals and their relationship, subjects provided their final belief.




Figure 3: Example beliefs of locals.

Design and Procedure We employed a within-subjects design with three levels (variations of network setups implied by differing cover stories). The three levels of our independent variable were: independent information, sequential dependence, and shared information. The two data points (i.e., test outcomes) used to parameterize subjects’ initial (neutral) prior beliefs, the resulting normative prior, and the parameters of the three locals are shown in Table 1. Parameter settings were constant across conditions, with the only source of variation being our independent variable. Resulting model predictions and sufficient statistics are summarised in Fig. 4 (columns 1-2). The order of conditions and the position of the candidate supported by locals (left/right) was randomized between participants. After completion of the main task, participants provided basic demographics (e.g., gender).

Table 1: Fixed parameters used across conditions.

Data	Normative Prior	Parameters of Locals
{1,0}	B(2,2)	B(46,9.5), B(49,10), B(50,45)

Analysis Our analysis has two parts. First, we compared the aggregate parameters of subjects’ posterior judgments between conditions to evaluate qualitative (i.e., directional) alignment with our model predictions. Thus, we first contrasted subjects’ posterior means (model 1) and variances (model 2) between

¹OSF; demo video; GitHub




Figure 4: Summary of model predictions (columns 1-2) and behavioural results (columns 3-4) for each condition (rows 1-3). For our main analysis, we used the aggregate μ and σ^2 parameters (plotted below each distribution) to make the interpretation of our predictions and results more intuitive. ** = $p < 0.01$ refers to the comparison of the (log) Jensen-Shannon Divergence from priors to posteriors between conditions (see model 4 in Table 2 and Fig. 5d for details).

conditions using two linear mixed-effects models with condition (i.e., social network set-up) as fixed effect and subject as a random intercept. However, evaluating these separately may miss dependencies that exist in how participants updated these components of their beliefs. To address this, we computed the Jensen-Shannon Divergence (D_{JS}) between priors and posteriors for each subject to determine whether the magnitude of updating prior beliefs differed by condition. D_{JS} allows measurement of changes both in mean and variance between distributions through a single symmetric distance measure given by:

$$D_{JS}(P||Q) = \alpha D_{KL}(P||Q) + (1 - \alpha) D_{KL}(Q||P) \quad (3)$$

where D_{KL} is the Kullback-Leibler Divergence, a standard asymmetric measure in information theory for measuring how much a probability density P has moved compared to a reference distribution Q . By definition, $D_{KL} \geq 0$, being equal to 0 if and only if P and Q are identical. A limitation of D_{KL} is its nonsymmetry, which is resolved by D_{JS} if $\alpha = 0.5$. Having computed each subject's D_{JS} , we fitted two additional mixed-effects models with condition as fixed effect and subject as random intercept to compare mean differences in D_{JS} (model 3) and the log-transform of D_{JS} (model 4). The reason for using $\log D_{JS}$ in model 4 is that the distribution of $\log D_{JS}$ residuals was closer to a normal distribution than the distribution of D_{JS} residuals (which showed a skew to the right). All models were compared to a reduced model including only an intercept as predictor variable and subject as random effect. Models were implemented in R using the function `lmer()` from the package `lme4` (Bates, Mächler, Bolker, & Walker, 2015).

Following tests of our qualitative (directional) comparison between conditions, we compared subject and model performances across conditions to check in how far subjects aligned

quantitatively with our normative framework. Therefore, we first computed D_{JS} between subjects' prior beliefs (Fig. 4, column 3) and the normative prior (Fig. 4, column 1) across conditions. Due to the skewed distribution of D_{JS} for this comparison, we conducted a Wilcoxon signed rank test (non-parametric t-test; alternative hypothesis > 0) to check if subjects integrated the asocial information as predicted by our model. We also compared the difference between subjects' prior mean and model prior mean and subjects' prior variance and model prior variance (using two-sided Wilcoxon signed rank tests because the dependent variables did not follow a normal distribution). The three comparisons were repeated to contrast subjects' posteriors with model posteriors.

Results

Sanity checks Five subjects were removed from our analysis because they did not change the positions of sliders between their prior and posterior judgments (i.e., $\log D_{JS} = -\text{Inf}$), resulting in a final sample size of 74. Levene's test revealed that the homogeneity of variance assumption was maintained for all four dependent measures (all $ps > 0.05$) used between models 1-4. Inspection of residual plots confirmed that the residual posterior means, posterior variances and $\log D_{JS}$ residuals were normally distributed. For D_{JS} , residuals showed skew to the right. Correlations between the three levels of our fixed effect (i.e., social network set-up) were moderate, ranging from ± 0.470 to ± 0.551 . Comparing each model to its reduced version revealed that inclusion of social network set-up only contributed significantly to the proportion of explained variance in $\log D_{JS}$ (see Table 2)². The qualitative comparison in the remainder of this paper will thus focus on interpreting the results of model 4 (for completeness, regression coefficients

² $\text{BIC}_{\text{diff}} = \text{BIC}_{\text{full}} - \text{BIC}_{\text{intercept-only}}$; R_m^2 = proportion of variance explained by the fixed effect (i.e., social network set-up).

for models 1-3 are reported in the next section and in Fig. 5).

Table 2: Model fits for each dependent variable (DV).

Model	DV	BIC _{diff}	R_m^2	χ^2	p -value
1	μ	6.14	0.009	4.69	0.096
2	σ^2	9.10	0.004	1.65	0.439
3	D_{JS}	6.40	0.012	4.52	0.105
4	$\log D_{JS}$	-1.43	0.032	12.27	0.002

Qualitative comparison Fig. 4 summarises model predictions and aggregate parameters across subjects for our three experimental conditions. The directional shift from subjects' prior distributions to their posteriors and the fact that subjects' posterior distributions were more compressed than their priors suggested that social information resulted in increased belief and certainty ratings. As expected by our normative model, the results of model 4 showed that subjects changed their prior beliefs significantly less in the two dependent conditions as compared to the independent case (Fig. 5d; $b = -0.639$, $t(148) = -3.07$, $p = 0.003$ for sequentially updated beliefs and $b = -0.641$, $t(148) = -3.08$, $p = 0.003$ for beliefs based on shared information). This means that the magnitude to which people updated their beliefs (i.e., changed both the belief and certainty sliders) was in line with the predicted (directional) magnitude of our normative model. In other words, people were sensitive to differences in the statistical power of the information between the independent condition (larger statistical power implying stronger updating of prior beliefs) and the two dependent conditions (smaller statistical power implying smaller updating; see Fig. 4).




Figure 5: Coefficients of the two dependent conditions compared to the independent case (intercept in each model). Error bars correspond to standard errors of the mean.

For the comparison between the two dependent conditions, no significant effects emerged (all $ps > 0.05$; see Fig. 4 and Fig. 5d). The results of models 1 (μ) and 3 (D_{JS}) showed that the comparison between sequentially updated beliefs and independent beliefs was significant ($b = -0.028$, $t(148) = -2.16$,

$p = 0.032$ for model 1; Fig. 5a and $b = -1.703$, $t(148) = -2.02$, $p = 0.045$ for model 3; Fig. 5c). For model 2 (σ^2), no significant differences emerged (all $ps > 0.05$; Fig. 5b).

Quantitative comparison Quantitative comparisons revealed that subjects' prior distributions differed significantly from model priors in terms of D_{JS} ($V = 25425$, $p < 0.001$). Inspection of Fig. 4 suggests that this difference might be driven by dissimilar variances, as subjects' priors were less diffused than model priors. The results of further comparisons confirmed this observation, showing that the mean of subjects' prior variance was significantly smaller than model prior variance ($V = 1607$, $p < 0.001$), despite finding no significant difference between prior means ($V = 13150$, $p = 0.652$, see Table 3). These findings might be attributed to an initial overestimation of certainty upon observing the outcomes of the candidates' test trials. Specifically, the average parameters of subjects' initial estimates of the candidates were equal to $B(19.2, 19.1)$, which was 9.58 times higher in magnitude than the simple Bayesian model's $B(2,2)$.

Table 3: Means and SDs of the dependent variables (DV) used for quantitative comparison.

Measure	DV	Mean	SD
prior	D_{JS}	2.547	3.782
prior	$\mu_{\text{subj}} - \mu_{\text{model}}$	0.001	0.085
prior	$\sigma_{\text{subj}}^2 - \sigma_{\text{model}}^2$	-0.027	0.023
posterior	D_{JS}	6.270	8.491
posterior	$\mu_{\text{subj}} - \mu_{\text{model}}$	0.033	0.132
posterior	$\sigma_{\text{subj}}^2 - \sigma_{\text{model}}^2$	0.008	0.012

Posterior contrasts revealed that model and subject distributions were significantly different from each other as measured by D_{JS} ($V = 25425$, $p < 0.001$). For posteriors, this difference was driven by a mismatch both in terms of posterior means ($V = 17196$, $p < 0.001$) and variances ($V = 25183$, $p < 0.001$). These findings demonstrate that, overall, subjects changed their posterior means more than expected by the normative prediction (mainly due to a strong mismatch between posterior means in the sequential belief updating condition, see Fig. 4, row 2). Subjects' average posterior variance was significantly larger than model posterior variance ($\sigma_{\text{subj}}^2 = 0.01$; $\sigma_{\text{model}}^2 = 0.002$). Compared to the average prior comparison, this finding might be attributed to the fact that subjects downweighted the evidential value of social information obtained from their peers. The average model posterior parameters across conditions were equal to $B(107.7, 58.4)$, which was 3.55 times the magnitude of subjects average posterior parameters $B(31.8, 15.0)$.

Discussion and Further Work

We modeled a sequential belief updating process including a target agent (i.e., the participant) and three social network neighbors. The cover story describing the relationship between network peers was varied in three within-subject conditions

to investigate the effects of three statistical (in)dependencies summarised in Fig. 1. Extending the findings of Whalen et al. (2018), our behavioural results confirmed our prediction that people update their beliefs significantly less when the provided social information was coming from dependent sources (as compared to the independent case). Thus, our result shows that people are not simply combining their own beliefs with the communicated beliefs of their network neighbors. Rather, they are additionally sensitive to the origin of those beliefs and to what extent they are redundant. This contributes another empirical piece to the puzzle of how to counteract the spread of false consensus effects and information cascades, which have been suggested to occur in networks of agents forming their beliefs in individually rational ways (Bikhchandani et al., 1992; Pilditch, 2017).

We could not confirm whether people differentiated between the evidential signals of shared information and sequentially updated beliefs while revising prior beliefs (despite the trend matching the predicted pattern; see Fig. 4). This might be attributed to the context of the task: we assumed that subjects would learn about political candidates based on binomial “performance tests”; and we operationalized network peers as locals being likely voters from a subject’s political party that formed their beliefs based on attending debates in their local town-hall. A more abstract experimental paradigm, such as learning coordinating with others to estimate the proportion of blue vs. red marbles in an urn (a common paradigm used to study information cascades and sequential belief updating; see e.g., Anderson & Holt, 1997) which does not require such context specific assumptions might have resulted in a measurable difference between the two dependent conditions. Despite being unable to differentiate between the two dependent cases, our task provides a valuable contribution to the field of social learning in the context of (online) political belief formation (see e.g., Bond et al., 2012).

To address the limitation of context, further experiments need to test the ecological validity of the presented normative framework across a variety of scenarios. These might involve emotional decisions (e.g., in the context of moral dilemmas), rational tasks, such as business decisions, and more general scenarios that are abstract (e.g., urn-based tasks). An additional limitation of our work is that we assumed additive communication of evidence in the case of sequentially updated beliefs (see Fig. 1). Generally, a more formal explanation of how sources form their beliefs and a precise description of the computational processes underlying evidence accumulation between conditions are important issues that need to be addressed in further work.

Relative to our simple Bayesian account, quantitative comparison between model predictions and subjects suggested that subjects over-weighted the influence of asocial information while they under-weighted the influence of social information. This finding is in line with previous theoretical (Schöbel, Rieskamp, & Huber, 2016) and empirical (Nöth & Weber, 2003) work demonstrating that people are more influenced by

their own private information as compared to social information, though in our case this might be an artifact of simulated social information (i.e., information coming from hypothetical social network members rather than actual ones). Moreover, we assumed that the shapes of the reported distributions reflect how subjects represent the distributions of their actual beliefs. If this assumption is not met, preferential weighting of asocial over social information might also be attributed to an initial misrepresentation (i.e., overestimation) of certainty. Thus, further work using the proposed interface could include a control measure of subjects’ ability to accurately report their certainty. This might involve an initial assessment of beliefs prior to engaging with any form of evidence.

To address some of the present limitations, we plan to replicate the current experiment across a variety of contexts where actual subjects come in pairs/triples to do the same task, enabling them to update beliefs dynamically. To generate a better understanding of whether people appropriately down-weight social evidence, further empirical work could also incorporate agent-based simulations contrasting normative accounts with competing models of social influence (e.g., Schöbel et al., 2016). This approach might enable measuring the empirical degree of information degradation due to correlated sources in realistic information networks. At the present stage, our finding suggests that people are able to understand and report probabilistic beliefs, which might be useful for calibrating belief parameters in related agent-based models (ABMs) of echo chambers (Madsen et al., 2018) and scientific belief formation (Lewandowsky, Pilditch, Madsen, Oreskes, & Risbey, 2019).

In summary, our results suggest that while a Bayesian framework provides a good qualitative account of how people update their beliefs based on social information coming from sources with different levels of independence, it cannot, in the current form, account for the relative weights that people assign to private and social information while updating their beliefs. We acknowledge that these findings might depend on the specific context of our experimental paradigm and plan further experimental validation of model predictions across alternative scenarios.

Acknowledgments

JPF is supported by the German Academic Scholarship Foundation. NT is funded by Università degli Studi Suor Orsola Benincasa. Participants were paid through a Student Grant from the University of Edinburgh awarded to JPF.

References

- Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American economic review*, 847–862.
- Bail, C. A. (2016). Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media. *Proceedings of the National Academy of Sciences*, 113(42), 11823–11828.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi: 10.18637/jss.v067.i01

- Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences*, 114(26), E5070–E5076.
- 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.
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415), 295–298.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., ... Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3), 554–559.
- De Martino, B., Bobadilla-Suarez, S., Nouguchi, T., Sharot, T., & Love, B. C. (2017). Social information is integrated into value and confidence judgments according to its reliability. *Journal of Neuroscience*, 37(25), 6066–6074.
- Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-making: A general bayesian framework for metacognitive computation. *Psychological review*, 124(1), 91.
- Fränken, J.-P., & Pilditch, T. D. (2020). Cascades across networks are sufficient for the formation of echo chambers: An agent-based model. Retrieved from psyarxiv.com/8rgkc
- Hegselmann, R., Krause, U., et al. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of artificial societies and social simulation*, 5(3).
- Lewandowsky, S., Pilditch, T. D., Madsen, J. K., Oreskes, N., & Risbey, J. S. (2019). Influence and seepage: An evidence-resistant minority can affect public opinion and scientific belief formation. *Cognition*, 188, 124–139.
- Lorenz, J. (2006). Consensus strikes back in the hegselmann-krause model of continuous opinion dynamics under bounded confidence. *Journal of Artificial Societies and Social Simulation*, 9(1).
- Lorenz, J. (2007). Continuous opinion dynamics under bounded confidence: A survey. *International Journal of Modern Physics C*, 18(12), 1819–1838.
- Madsen, J. K., Bailey, R. M., & Pilditch, T. D. (2018). Large networks of rational agents form persistent echo chambers. *Scientific reports*, 8(1), 12391.
- Madsen, J. K., & Pilditch, T. D. (2018). A method for evaluating cognitively informed micro-targeted campaign strategies: An agent-based model proof of principle. *PloS one*, 13(4), e0193909.
- Minsky, M. (1988). *Society of mind*. Simon and Schuster.
- Nöth, M., & Weber, M. (2003). Information aggregation with random ordering: Cascades and overconfidence. *The Economic Journal*, 113(484), 166–189.
- Pilditch, T. D. (2017). Opinion cascades and echo-chambers in online networks: A proof of concept agent-based model. *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature neuroscience*, 19(3), 366.
- Schöbel, M., Rieskamp, J., & Huber, R. (2016). Social influences in sequential decision making. *PloS one*, 11(1).
- 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.
- Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9), 5766–5771.
- Whalen, A., Griffiths, T. L., & Buchsbaum, D. (2018). Sensitivity to shared information in social learning. *Cognitive science*, 42(1), 168–187.