FRONT MATTER

Title

A cognitive mechanism for the persistence of widespread false beliefs

Authors

Vincent V. Frigo^{1*}, Lang Chen², and Timothy T. Rogers^{1*}

Affiliations

¹Department of Psychology, University of Wisconsin—Madison

²Department of Psychology, Santa Clara University

* vfrigo@wisc.edu

* ttrogers@wisc.edu

Abstract

Cognitive science commonly assumes human learners combine prior knowledge with new information to form statistically optimal beliefs. Why then are false beliefs so resilient and widespread in society? We develop and test a hypothesis stemming from computational models of human category learning. Over four experiments, we find learners weight conflicting information sources pursuant to a nonlinear function dependent on the learners' initial beliefs. Highly agreeable sources receive great weight, preventing belief change, but this deference drops sharply and then flattens with increasing disparity, allowing even very distal sources to effect belief change. To understand the implications of this bias, we simulate groups of agents learning from one another and from an oracle. The learning bias often leads agents to fractionate into mutually distrustful groups or converge on a false belief. Opposing belief bubbles in contemporary society may reflect stable states in this dynamic system.

Teaser

A new cognitive learning bias explains why our society often fractionates into groups possessing stable, mutually-incompatible beliefs.

MAIN TEXT

Introduction

In the late 1840s, innovations in handwashing radically decreased the death rate in a Vienna birthing ward. The technique's champion was promptly fired, and the death rate rebounded¹. In 1990's England, MMR vaccinations dropped following ill-founded claims linking the treatment to autism²; incidence of measles rose by 780% over four years^{2,3}. In the US in 2020, a prominent TV personality contended COVID19 was no more dangerous than the flu, against all medical evidence⁴. A poll that month found 20% of Americans agreed⁵; within a year the US led the world in pandemic deaths⁶. Cognitive science commonly presumes human learning is rational: prior knowledge coupled with experience begets optimal probabilistic inference⁷. Why then do widespread false beliefs persist in daily life?

We develop and test a hypothesis that arises from studies of human category learning. This task has long provided an empirical testbed for computational models of belief formation,

with an extensive literature establishing that people behave rationally when learning from items that are correctly labeled by a teacher or “oracle.”^{8–10} Yet real-world experience rarely provides such unambiguous feedback, instead serving a hodge-podge of conflicting information from sources varying in their knowledge, accuracy, and trustworthiness. Learners must weight these contradictory sources when updating beliefs, but little is known about the factors governing such weighting.

We conducted four experiments assessing how people weight conflicting labels in a simple category-learning task where factors commonly considered critical to false belief maintenance—motivated reasoning¹⁷, cognitive dissonance^{18,19}, heuristic reasoning biases^{20,21}, domain-specific prior knowledge^{22,23}, memory coherence^{24,25}, and social pressures²⁶—play little part. The results revealed a learning bias that can lead people to converge on differing beliefs despite overwhelmingly similar experiences, and that differs from classic biases identified in other literatures. We further show that the bias is not explained by either generative or discriminative computational models that have successfully accounted for human behavior in standard supervised and semi-supervised contexts, but can be described mathematically using a simple heuristic model. To understand the implications of this bias for the propagation of beliefs in social groups, we conducted simulations with agent-based models^{11,12} whose nodes behave in accordance with the learning bias observed in experiments. The simulations suggest how this bias supports the emergence of mutually distrustful groups sustaining distinct beliefs, and how such fractionation might be remediated by the algorithms that curate news feeds on social media platforms.

Approach. All behavioral experiments followed the same general design; full details appear in Methods and Materials. Participants imagined they were shipwrecked on a deserted island together with “veteran sailors who know the lands intimately, mates who have sailed there for a few years, novices, and criminals who lie habitually.” Stimuli were 301 supershapes varying along a one-dimensional perceptual continuum from “spiky” to “smooth,”²⁷ and introduced as edible native fruits some of which taste better when cooked and others when eaten raw. The participant was to gather fruit with two unfamiliar companions rendered as faces on the computer screen. The participant decided “cooked” or “raw” for each fruit, then viewed each companion’s opinion. Across studies we assessed how learners weighted the feedback from the two companions when updating their beliefs. Each study proceeded in five phases as shown in Figure 1A:

(1) Participants completed twenty *supervised learning trials* involving one example from each category repeated 10 times in permuted order, with correct feedback provided by a “knowledgeable and trustworthy” captain, to establish an initial belief about the boundary between categories.

(2) Participants then classified, without feedback, 38 equally-spaced items centered on the midpoint of the stimulus range, allowing us to estimate the learner’s initial belief l_i about the boundary location.

(3) Next participants classified all 301 stimuli presented in random order, receiving feedback from each companion displayed as a written label (RAW or COOKED) below the corresponding face. Each companion’s opinion was based on a fixed category boundary unknown to the participant.

(4) Participants then classified the same 38 items from Step 2 without feedback, allowing estimation of the *final belief* about boundary location l_f . Belief change was measured as the boundary shift $\Delta l = l_f - l_i$.

(5) Participants used a 100-point scale to rate each companion’s knowledge, trustworthiness, accuracy, and, as a control, attractiveness.

94 Some elements of this design are worth emphasizing. Because the items were novel,
95 participant behavior could not be strongly constrained by prior real-world beliefs. The
96 categorization task—should this item be eaten raw or cooked—was purposely mundane so that
97 motivational factors did not favor one answer over another. The distribution of observed stimuli
98 was flat so that probability density did not provide a cue as to the location of the category
99 boundary. As each participant viewed the opinions of just two sources plus the “captain,” they
100 received little information about the degree of consensus or majority opinion in the population.
101 The scenario described made it clear that the two companions may or may not be knowledgeable
102 or trustworthy, to reduce reliance on prior beliefs about the likelihood of encountering a
103 knowledgeable, trustworthy source. Finally, we counterbalanced all nuisance factors across
104 conditions—which face image was paired with which companion boundary, whether the “raw”
105 label encompassed the spikier or smoother end of the stimulus manifold, and the alignment of
106 companion boundaries toward/away from each end of the stimulus manifold—so that changes in
107 boundary location must be due to the experimental factors in each study.

108 Across experiments we manipulated the boundaries used by the companions and the
109 location of the participant’s initial boundary to study how learners weight conflicting input from
110 uncertain information sources, both when updating their beliefs and when forming explicit
111 opinions about the sources. To assess whether the observed behaviors are consistent with
112 common discriminative and generative models of category learning, we simulated each study
113 using variants of these models adapted to learn from multiple potentially conflicting sources.
114 Finally, we used simulation and analysis of agent-based models to understand what the behavioral
115 results suggest about the propagation of beliefs in social groups.

116 **Results.**

117 **Experiment 1** assessed how learners weight sources differing in their agreement with the
118 learner’s initial belief. The stimuli can be characterized by integers ranging from 0 (spikiest) to
119 300 (smoothest) with the numeric distance between items proportional to their perceived
120 similarity. The supervised phase employed category exemplars at 100 and 200 to generate an
121 initial boundary near the range midpoint (150). To validate our procedure, control conditions
122 assessed how beliefs change when (1) only one companion provides feedback from a boundary
123 either 15 or 25 units from the midpoint, and (2) neither source provides feedback. In the former
124 case, learner boundaries should shift toward the single-source boundary with a larger shift for the
125 farther source, while in the latter they should not shift at all.

126 In the experimental condition (Figure 1B), one companion adopted a “close” boundary 15
127 units from the midpoint (at 135 or 165) while the other adopted a “far” boundary 100 units from
128 the midpoint in the opposite direction (at 250 or 50, respectively). Companions thus provided
129 conflicting feedback for many items, and the far source provided some labels that directly
130 contradicted the initial supervised learning phase. Learners might respond to this feedback in at
131 least three ways:

132 *(H1) They may weight sources equally so that the learner’s boundary shifts close to the*
133 *midpoint between source boundaries.*

134 *(H2) They may ignore the far source since this feedback conflicts with both the*
135 *supervised trials and the learner’s initial belief, in which case the learner’s boundary should shift*
136 *toward the close companion’s boundary.*

137 *(H3) They may ignore feedback entirely, in which case the boundary should not shift.*

138 These hypotheses represent different parameterizations of a heuristic model under which
139 the learner’s mental boundary shifts as a weighted sum of the distance between the initial

boundary and each source's boundary. Let m and n be the boundaries used by the farther and the closer source, respectively. The heuristic model stipulates:

$$\Delta l = \theta(m - l_i) + (1 - \theta)(n - l_i) \quad \text{Eqn. 1}$$

...where $\theta \in [0,1]$ is a free parameter specifying how the two sources are weighted. When $\theta = 0.5$ sources are weighted equally and the boundary shifts to the midpoint between them (H1). When $\theta = 0$ the far source is ignored and the learner's boundary shifts to match the closer source n 's boundary (H2). When $\theta = |m_i - l_i|/(|m - l_i| + |m - n|)$, the product of the weight and distance for each opposing source cancel out, so Δl will be 0 (H3) and both sources are ignored. This leads to a fourth hypothesis:

(H4) Learners may assign some other weighting θ , which can be computed from empirical estimates of Δl , and the known location of m , and n by rearranging Equation 1:

$$\theta = \frac{l_f - n}{m - n} \quad \text{Eqn. 2}$$

If the confidence limits on this estimate exclude the values for H1-H3 above, participants must weight the two sources according to some heuristic beyond the noted hypotheses.

Figure 1 A. Task procedure for all studies. **B. Design for Experiment 1.** Dots show locations of two supervised training items, one from each category (raw/cooked, shown as red/blue). The dashed line indicates the midpoint between these. Yellow shading shows the items used to estimate initial and final boundaries. Gray shading shows the items viewed during the two-source phase. Red and green vertical lines show the far and close source boundaries while red/blue lines above show which labels each source provided. The line labelled "agree" indicates where the two sources agree (blue for one label, red for the other) and where they disagree (purple). **C. Mean and 95% confidence intervals** for the boundary shift in each condition of Experiment 1. Positive numbers indicate shift toward the close source. Crosses show the location of the close source boundary; the red dash shows the midpoint between sources in the two-source condition. **D. Mean and standard error of source ratings** for knowledge, trustworthiness, accuracy, and attractiveness in Experiment 1. Asterisks indicate reliable differences at $p < 0.005$ based on a paired-samples t-test. **E. Estimate of weighting parameter θ** for each participant in Experiment 2 plotted against the absolute distance v between the participant's initial boundary and the varying source boundary. The red curve shows the least-squares estimate of the heuristic evidence-weighting model.

Figure 1C shows how boundaries shifted in each condition, with positive numbers indicating a shift toward the close source boundary. In the one-source control conditions, learner

boundaries moved toward the companion's boundary; in the no-source control they did not shift at all. In the experimental condition, boundaries shifted reliably toward the *far* source, but not as far as the midpoint between sources ($M = -14$, $p < .001$ for both contrasts). The weighting parameter θ had a mean estimate of 0.23 across subjects, suggesting learners gave less weight to the far source when updating their beliefs. Participants also rated the far source as less knowledgeable, accurate, and trustworthy ($p < 0.005$ for all contrasts), but not less attractive, than the close source (Figure 1D). Thus category beliefs shifted toward the far source boundary despite greater skepticism about this companion and a lesser weight given to his feedback.

Experiment 2 considers whether the estimated source weighting varies systematically with distance from the learner's initial belief to each source's boundary. Each participant viewed feedback from a *fixed source* 15 units from the midpoint (at 135 or 165) and a second *variable source* whose boundary was sampled independently for each subject with uniform probability from the half of the stimulus range opposite the fixed source. For each participant we estimated the boundary shift and computed the weight θ given to the variable source using Equation 2. We then assessed how this weight varied with the distance v between the learner's initial boundary and the variable-source boundary.

Results are shown in Figure 1E: the variable source received less weight the further its boundary lay from the learner's initial boundary. This decline in weight was, however, nonlinear, initially dropping steeply but then flattening out, and never decaying to zero so that even when a source provided highly implausible labels it had some influence on the learner's boundary. We fit this relationship with the following curve, shown as the red line in Figure 2E:

$$\theta = 1 - \frac{[v - \alpha]^+}{[v - \alpha]^+ + \beta}$$

In this *heuristic evidence-weighting (HEW)* model, v is the distance from the learner's initial boundary l_i to the variable-source boundary m , $\{\beta \in \mathbb{R} \mid \beta > 0\}$ is a parameter controlling the rate at which θ decays toward zero, and α is an offset constant that shifts the curve left or right. The notation $[\]^+$ indicates that the difference $v - \alpha$ is rectified to 0 when negative so that θ is properly bounded at 1 and 0. When $v \leq \alpha$ (so that the varying source is within some radius α of the learner's initial boundary) then $\theta=1$, the varying source gets all the learner's weight, and the close source is ignored. As v grows, θ diminishes but never decays to zero so that the varying source pulls the learner's boundary toward it to some degree even when very distal. The parameters β and α can be estimated empirically via least squares. The best-fitting curve shown in Figure 1E explains 41% of the variance in the data, better than linear (35%) and sigmoidal (34%) models, which have the same number of free parameters.

Importantly, the HEW model fit to Experiment 2 data suggests that a source receives the *entirety* of a learner's weight when its boundary lies within about 5 units of the learner's initial boundary, in which case the learner will *ignore* feedback from the other source. This extreme weighting drops steeply with distance from initial belief, but the rate of change declines so that moderately and extremely distal sources receive somewhat similar weightings—thus learners will be influenced by distant sources as long as no other source agrees perfectly with their initial belief.

Figure 2 A. Design for Experiment 3. Dots show labelled items for each group in the supervised phase; vertical dotted lines show the expected initial boundary for each group. Blue and orange vertical bars show source boundaries in the two-source phase. **B. Mean and 95% confidence interval** of the boundary shift in each group; positive numbers indicate shift toward the midline. **C. Mean and standard error of source ratings** for each group.

Experiment 3. The HEW model predicts that beliefs will not change if one source strongly concurs with the learner, since the affirming voice receives all the weight and merely echoes the learner's opinion. Experiment 3 tested this prediction in two groups who initially learned boundaries near either 100 or 200 from the sparse initial supervision (Figure 2A), then received identical feedback on 300 trials from two sources employing boundaries at 100 (agreeing with Group 1) and 200 (agreeing with Group 2). If learner boundaries always move toward the farther source (as in Experiments 1 & 2), the two groups should grow more similar in their beliefs. Under the HEW model, however, no learning should occur and the two groups should retain discrepant beliefs despite overwhelmingly similar learning experiences.

Consistent with this prediction, initial boundaries did not shift ($M=0.43$, $t(29)=0.11$, $p<0.92$; Figure 2B) and final category boundaries differed reliably between groups ($M=142.58$ and 174.73 , $t(20.95)=-4.22$, $p<.001$). The groups also showed opposing judgments about the knowledge ($F(2, 6)=6.69$, $p<.01$ for interaction) and trustworthiness ($F(2,6)=11.42$, $p<.001$) of each source, preferring whichever provided labels more consistent with their own initial beliefs.

Experiment 4. The two groups in Experiment 3 retained discrepant beliefs even though (1) there existed a category boundary consistent with the initial supervised learning experiences of both groups and (2) participants in both groups received exactly the same feedback on the vast majority of trials (300 of the 320 total). Thus groups of learners can differ in their beliefs despite highly similar learning experiences, if each learner's belief is strongly affirmed by one source. What might cause greater agreement between dissenting groups? Obviously this should occur if both sources provide similar information—that is, if all sources provide the same feedback—but the HEW model suggests an additional possibility. Figure 3A shows the predicted boundary shift in the Experiment 3 study design (x-axis) when the two source boundaries lie at varying distances from the range midpoint, the learner, and each other (y-axis) for various values of the slope

parameter β (different colors). The learner boundary shifts toward the midline (and hence toward the other group) when both source boundaries are near the midpoint between groups and thus pulling in the same direction—but also when they are near the extremes, pulling in opposite directions. When at the extremes, both source boundaries are distal to the group’s initial boundary, and the product of source weight and distance leads the farther source to have a greater impact on belief change than the closer source. Thus, the HEW model makes the counter-intuitive prediction that two discrepant groups should also come into greater agreement when the two source boundaries lie at opposite extremes, both far from either group’s initial boundary. As in Experiment 3, it also suggests that no learning occurs in the “bubble” where one source boundary lies near the initial boundary.

Experiment 4 tests this prediction. Participants again learned an initial boundary near 100 or 200, then were randomly assigned to one of four experimental conditions (Figure 3B). In the *between-initial* condition, source boundaries lay at 125 and 175, both *between initial* expected group boundaries and so pulling both groups toward the midline. Other conditions positioned the two sources at increasing distances from the range midpoint (at 100/200 for *at-initial* condition, replicating Experiment 3; 75/225 for the *near-initial* condition; and 50/250 for the *at-extreme* condition). A no-source control condition allowed us to assess how group boundaries change without feedback in the two-source phase.

Figure 3 A. Predicted boundary shift toward midline (x-axis) as source boundaries move from near the midline to the poles (y-axis). Colors show predictions for different model slope parameters; black shows the slope estimated in Experiment 2. **B. Design for experiment 4** using same conventions as Figure 3. **C. Mean and 95% confidence interval** of shift toward the midline in each condition. **D. Vertical lines/shading show the mean and 95% confidence interval** for the initial boundary location in each group. Dots and horizontal bars show the mean/standard error of the final boundary location in each condition and group. Stars indicate reliably different final boundary locations in the two groups (* $p < 0.05$, ** $p < 0.01$).

Consistent with the HEW model, boundaries for each group of learners shifted toward the midpoint when both source boundaries lay between group initial boundaries ($t(28)=6.53$, $p<0.001$) and when they lay at extremes distal to group initial boundaries ($t(23)=3.16$, $p<0.005$). As in Experiment 3, the boundaries did not shift reliably when one source boundary was identical to each group’s expected initial boundary ($t(29)=1.25$, $p=n.s.$), nor did they shift when sources

were situated beyond, but still close to, each group's initial expected boundary ($t(30)=0.24$, $p=n.s.$). The two groups agreed on the final boundary location (i.e., no reliable difference) for both the *between-initial* ($t(27)=1.53$, $p=n.s.$) and *at-extreme* conditions ($t(22)=0.23$, $p=n.s.$), but retained reliably different beliefs in the *at-initial* ($t(28)=-2.71$, $p<0.02$) and *near-initial* ($t(29)=-2.68$, $p<0.02$). In other words, the two groups continues to disagree so long as they received approximate "validation" of their belief from at least one source.

Simulation study 1: Discriminative and generative models. Together the behavioral studies suggest a heuristic model in which learners weight conflicting sources as a function of their respective proximities to the learner's own belief. We describe this as a heuristic model because the work does not reveal any principled reason why learners should behave in this way—it only describes behavior, albeit in a manner that allowed us to make and test some counter-intuitive predictions. A key question is whether existing learning models from cognitive science can offer a more principled account of the behavior. To answer this question, we simulated behavior in experiments 1-4 using variants of both *discriminative* and *generative* category learning models^{37,51}. Both frameworks are common in cognitive science, accounting well for human behavior in both fully-supervised and semi-supervised scenarios, including for the same binary classification task and stimuli used in the current work. Both frameworks are also easily adapted to make use of weighted labels (see Methods).

We simulated behavior for all four experiments and each model type in the same way. Each stimulus was represented with a real-valued coordinate $x \in [0,300]$ indicating the shape's "spikiness," perturbed by a small amount of Gaussian noise to capture the perceptual confusability of neighboring stimuli. The raw/cooked category labels were represented with a binary integer. We first fit an *initial model* to the two labeled items (each repeated 10 times) and the 38 unlabeled items from the experiment's supervised and initial-boundary phases; from this we computed the initial category boundary b_i as the coordinate x yielding equal probability of each category. We then fit a *final model* using the same initial labeled and unlabeled items plus the 301 stimuli from the experiment's two-source phase, each paired with a label from source 1 and source 2 and with a parameter θ specifying how the sources were to be weighted in the model fit. From the final model we computed a final boundary location b_f , then calculated the boundary shift $b_f - b_i$ exactly as in the experiment.

The simulations require specification of the source-weighting parameter θ , representing some hypothesis about how learners weight the different sources of information. We considered three different hypotheses. The first is simply that learners weight both sources equally, as would be the case for standard supervised learning. In this case $\theta=0.5$ for all training items. The remaining two hypotheses both assume that θ depends on an *evidence-weighting* model E that, given an example x_i labeled by source s , returns the probability that the label is correct:

$$\theta = p(y_i = y_{s_i} | x_i, y_{s_i}, E)$$

Hypothesis 2 is that, from the supervised examples, the learner forms a mental model of the *oracle*, identical to the learner's initial model. When the two sources provide conflicting labels, the learner uses the oracle model to infer the label probabilities the oracle *would have* provided. The learner then updates her own model using these inferred probabilities as the weighting parameter θ . When the two sources agree, the label they provide is treated as the correct label. Hypothesis 3 is that the learner uses her own current beliefs as the evidence weighting model. When the two sources disagree, the learner estimates the posterior probability of the label given her current model, then updates parameter estimates using the inferred probability as the weight θ . Further simulation details appear in Methods and Materials.

Results.

Figure 4. Simulations with discriminative models. Panels show results for discriminative models in which both sources receive equal weight (top row) or are weighted by a model of the oracle (middle row) or by the learner's current model (bottom). Columns show simulated and real data for the primary result from each study. Study 1 (left) shows the initial (green) and final (red) model probability curves with implied boundaries (dashed lines) and the predicted (yellow) and confidence intervals for observed (blue) boundary shift (right). Study 2 shows the estimated weight given to the variable source based on model-predicted boundary shift (black curve/dots) and the fitted curve from real data (blue). Study 3 (left) shows initial and final model probability curves and (right) predicted/observed boundary shift (yellow/blue). Study 4 shows predicted (rainbow) and confidence intervals for observed (light blue) boundary shift for the four study conditions; positive numbers indicate shift toward the midline. BI: between initial; AI: at initial; PI: past initial; AE: at extreme.

Figure 4 shows results for discriminative models. None showed a pattern of predicted behaviors commensurate with any of the four studies. All models showed a much larger shift of the boundary toward the far source than was every observed in any study; no simulation yielded a result in which the boundary failed to change (as was observed empirically in studies 3 and 4); and the magnitude of the boundary shift varied monotonically with the distance between sources in Experiment 4, rather than exhibiting the U-shaped pattern in our studies. Thus weighted discriminative models do not provide a good account of the behavioral phenomena.

337

338 **Figure 5 . Simulations with generative models.** Panels show results for generative models in which both sources
339 receive equal weight (top row) or are weighted by a model of the oracle (middle row) or by the learner's current
340 model (bottom). Columns show simulated and real data for the primary result from each study. Study 1 shows the
341 estimated probability density functions for each category (left) for initial (top) and final (bottom) models, together with the
342 model category boundary (green/red dashed lines) and predicted shift (yellow bar) and observed confidence interval
343 (blue box). Study 2 shows the estimated weight given to the variable source based on model-predicted boundary shift
344 (black curve/dots) and the fitted curve from real data (blue). Study 3 shows initial and final probability densities
345 (right), predicted shift (yellow bar) and observed confidence interval (blue box). Study 4 shows predicted shift
346 (rainbow) and observed confidence intervals (blue) for the four study conditions; positive numbers indicate shift
347 toward the midline. BI: between initial; AI: at initial; PI: past initial; AE: at extreme.

348 Figure 5 shows results for generative models. Columns 1 and 3 show the estimated
349 probability densities on x for each category (pink and blue) for the initial (top) and final (bottom)
350 models, as well as the optimal category boundary (dashed lines). Some models provided a good
351 quantitative and qualitative match to some experimental results, but no model produced results
352 commensurate with all four, and some results were not explained by any model. When source
353 labels were weighted equally, simulations of Study 1 predicted a boundary-shift toward the far
354 source of a magnitude comparable to that observed in experiment, but also a shallower/more
355 linear decay of the weighting parameter θ with source distance in Study 2, and much larger
356 boundary shifts than were observed for Studies 3 and 4. Oracle- and learner-weighted models
357 provided a poor fit to Study 1, predicting no boundary shift, but yielded evidence-weighting
358 curves qualitatively similar to those observed empirically in simulations of Study 2. For Study 3
359 they both predicted quite small boundary shifts not too different from the observed shift of zero,
360 but neither model accounted for the U-shaped pattern of boundary shift observed across
361 conditions in Study 4, instead predicting either small (oracle-weighted) or large (learner-
362 weighted) shifts away from the midline in the at-extreme condition. Thus the model that explains
363 the shift in Study 1 cannot explain the evidence-weighting curve from Study 2 or the absence of
364 boundary-shift in Study 3 and the moderate conditions of Study 4. Models that do account for
365 these latter phenomena do not explain the boundary shift in Study 1, and no model explains the
366 U-shaped pattern of boundary shift observed in Study 4.

In sum, discriminative and generative models that have proven useful for understanding human category learning in fully-supervised and semi-supervised settings, when adapted to learning from multiple conflicting sources, do not clearly explain our findings. Note, we do *not* contend that no variant of these approaches is capable of explaining our results—only that the variants we have considered, which correspond to fairly straightforward hypotheses about how learners might weigh questionable evidence, do not transparently account for our findings.

Simulation study 2: Agent-based models. The HEW model suggests a mechanism for the persistence of false beliefs in social groups. In everyday life we typically learn from other individuals trying to form their own opinions from uncertain evidence, which may or may not be correct and which can change over time. We may also receive information from true "oracles"—people who know and tell the truth, in which case the feedback they provide will not change. Often a challenge for the learner is that she does not know which sources are oracles and which not. Simulation study 2 was designed to understand how the learning bias demonstrated in our experiments impacts formation of beliefs across groups in this setting, where a learner receives information from other learners and oracles, without knowing which is which.

We used an *agent-based* modeling scenario—a formalism in which individuals are represented as *agents* possessing internal states that govern their overt behaviors, interacting with one another via connections represented as a graph^{44,45}. Agents independently update their internal states and/or alter their overt behaviors based on the inputs they receive from other agents and the environment. All elements of the system—the agents’ internal states, the behaviors they generate, how these change through interactions, and the structures governing interactions—are described mathematically, allowing the theorist to study, either analytically or via simulation, the group-level behaviors that emerge from interactions among otherwise independent actors.

Consider the case where two agents, Alice and Bob, observe an object, and Alice labels it for Bob. This interaction can be represented as a directed edge from A to B in a social graph. The environment provides a sample to be classified, and Alice uses her own current belief to generate a label—in this sense the environment and the agent's internal state jointly determine Alice's observable behavior (the label she produces). In this interaction, Bob observes the same item together with the label provided by Alice, and, if he is not an oracle, updates his own internal state (his belief about the boundary location) according to some learning rule. An oracle can label items for other agents but does not update its internal state since this is wedded to the ground truth. This interaction, replicated many times over groups of learners connected in a social graph, constitutes an agent-based model of the category learning task at a social level. Clearly the distribution of beliefs across individuals will change over time, in ways that are hard to foresee but depend on both the behaviors of individual agents and the structure of the graph governing their interactions.

Our simulations simplified the above scenario in two ways. First, since our experiments involve learning from two sources, we consider a model scenario in which, on each simulated learning trial, the learning agent updates its current beliefs based on inputs from two source agents (either of which may be another learner or an oracle). Second, rather than sampling individual items to be labeled, we instead suppose that, on each interaction, source agents communicate their current belief about the category boundary directly to the learning agent, approximating a scenario where the learner sees many samples labeled by the source agents on each round. Formal model definition and analysis is provided in Materials and Methods.

Figure 6. **A. Left:** Two simulated agents (green circles) learning from each other and from a single oracle (blue circle). **Right:** Trajectories show how beliefs change from different starting configurations (green dots) when sources are weighed equally or via the HEW function. Gray lines trace how initial beliefs change over time; red dots show final beliefs after 200 updates; blue lines show the true boundary. **B. Left:** Ten simulated agents learning from each other and from two oracles under guidance from a feed curation algorithm that chooses which two sources an agent learns from on each round. Lines show possible communication pathways from agents (gray) and oracles (blue). The feed curator chooses which two paths are used for a given agent on each update (red lines). **Right:** Trajectories showing how the 10 initial beliefs change over time for different feed-curation algorithms, when beliefs are updated using the HEW model. The blue line shows the ground truth. When the feed curator chooses sources at random with uniform probability, agents end up agreeing with one another but not with the ground truth (conformity). When it selects the two closest sources for each agent, fractionated subgroups with different beliefs form (echo chamber). The same result obtains when the curator chooses the closest and farthest beliefs as sources for each agent (both sides). When the curator chooses sources that are similar to the learner but outside the “bubble” that produces complete discounting of other voices, all agents converge on the ground truth.

We first considered the minimal case of two individuals learning a category boundary, each receiving feedback from two sources. Unbeknownst to either learner, one source is a truth-telling oracle while the other is the second learner. Each learner updates her beliefs by weighting the two sources according to some policy. Learner 1’s updates depend on learner 2’s beliefs and vice-versa so that the two learners jointly constitute a dynamic system in which belief-updating is coupled. Figure 4A shows how learner beliefs evolve over time from different initial states under two weighting policies. When learners weight sources equally they converge on the oracle’s boundary. When they weight sources according to the HEW function, they only converge on the true boundary when it lies between their initial beliefs. When these reside on the same “side” of the true boundary the learners come into perfect agreement before either agrees with the oracle

and learning stops. Thus three groups emerge across all starting beliefs: one that agrees on the truth, a second that is too extreme in one direction, and a third too extreme in the other.

The remaining simulations considered how learning evolves with larger numbers of agents. A broad literature has considered how the structure of a social graph influences the propagation of beliefs and related emergent phenomena including belief polarization^{36,46–48}. We view this work as highly relevant and important but leave its development in the current context to future research and instead consider a fully connected graph where, on each interaction, a learning agent may receive information from any two other sources (Figure 6B, left). Within this setting we considered how beliefs evolve depending upon the algorithm determining which two agents serve as the sources for each learning agent on each round. The models thus capture a scenario in which a social-media platform has a “bird’s-eye view” of the current beliefs of all users and uses this information to determine which opinions get shared amongst which subsets of users at any given time.

As previously, each learning agent viewed the opinions of two source agents once in every epoch and updated its current belief according to the HEW function parameterized as determined in Experiment 2. Source agents were either other learners or oracles that report a common ground truth. We simulated the evolution of beliefs over time for 10 learners and 2 oracles under four simple information-sharing algorithms (see Methods). In the *random* condition, two sources were sampled with equal probability from among all agents excluding the learner, independently for each learning agent on each round. In this case learning agents all converged to a final belief near the mean of the initial belief distribution, but still distal to the ground truth reported by the oracle. That is, the system exhibited a kind of “conformity” effect, in which learners all adopted the majoritarian view, even though it differed from the lone voices providing a ground truth (Figure 4C, “Conformity”). The reason is clear from the preceding analysis of the minimal case: learning agents do not update their beliefs when at least one source is in strong agreement. With random source sampling, beliefs of learning agents all grow more similar to one another, leading to higher and higher probability of each “freezing” in place. When all learners agree, none of them will adjust their beliefs except in the rare learning trials when both randomly selected sources are oracles. Even in that case, the belief will snap back to the majority view on future learning rounds.

In the *echo chamber* condition, each learning agent received input from the two sources with beliefs most similar to the learner on each round. In this case, near neighbors were drawn toward one another until they fell into agreement, at which point learning stopped. As a result the learners fractionated into groups with different beliefs (Figure 4C, “Echo chamber”).

In the *both sides* condition, each learning agent received input from the source with the most similar and the most dissimilar belief on every round, providing a model of the “both sides” approach to persuasion. Prior work has sometimes assumed that exposure to diverse views can help prevent polarization or disinformation bubbles^{49,50}. When learning was weighted by the HEW function, however, the learners again fractionated into subgroups possessing different stable beliefs. The reason is clear from the minimal-case analysis: once the nearest neighbor expresses a belief identical to the learner, the second source is ignored—so exposure to the “other side” of the argument has no impact.

The preceding conditions make it clear that, if learners weight information as suggested by our experiments, then the group will fail to converge on a ground truth whenever all learners hear from a source with a very similar belief. The final condition considered what happens when sources are selected to have beliefs similar to the learning agent, but *outside the bubble* that creates complete discounting of the second source. On each round for each learning agent, we first discarded from consideration all sources whose beliefs were sufficiently similar to “freeze”

the agent's learning (ie, within α of learner's boundary), then selected the two most similar sources from among the remaining agents. Under this rule, when learners agree with one another but not with the two oracles, they hear only from the two oracles. Thus learning will only stop completely when all agents possess very similar beliefs—ie, when they have converged on the ground truth (Figure 4C, “Outside the bubble”). This simulation suggests that algorithms for sharing information across social networks may help to remediate false-belief bubbles.

Discussion

We have documented a bias in human category learning that suggests one reason why false beliefs persist in society. In a learning task requiring no prior knowledge and with little social or affective stakes, participants weighted conflicting information pursuant to a nonlinear function dependent upon their own initial belief. Under this *heuristic evidence-weighting model*, sources highly agreeable with a learner's initial belief are strongly weighted and drown out other voices, preventing belief change. The weighting drops sharply for more discrepant sources, but the rate of change slows so that even very distant sources can effect significant belief change. Simulated agents adopting this model often fractionated into mutually distrustful groups or converged on incorrect beliefs. Differing belief bubbles in contemporary society may therefore reflect stable states in a dynamic system composed of individual learners communicating their beliefs to one another and updating those beliefs according to an evidence-weighting function that strongly favors close agreement but, in the absence of such agreement, never completely discounts extreme voices. This proposed mechanism relates to, but differs from, prior accounts offered across diverse literatures. Here we briefly review several alternative hypotheses and consider how they relate to the current work.

Motivated reasoning encompasses any situation in which a person's thoughts or beliefs are in some manner influenced by a desire to arrive at a conclusion with favorable implications⁵⁵. If belief formation optimizes affective cost jointly with fidelity to the truth, this may cause false beliefs—for instance, when a participant discounts evidence suggesting that she is at risk of cancer^{25,56–58}. While an extensive literature suggests that motivational factors can indeed lead to false beliefs, motivated reasoning does not explain the current results since, in these studies, a given belief (e.g. that spiky fruits should be cooked) levies no emotional cost. Neither conclusion (cooked or raw) is in any way more desirable than the alternative. The sole motivational pressure that exists is a desire to be accurate and/or to please the experimenter, neither of which would explain the selective weighting of information we have documented.

Confirmation and disconfirmation bias. The finding that learners more strongly weight concurring evidence might be viewed as a manifestation of confirmation or disconfirmation bias^{24,59}—the ubiquitous tendency for people to seek out evidence consistent with their existing beliefs (confirmation bias) while actively avoiding, discounting or scrutinizing contradictory evidence (disconfirmation bias). Yet the current work also differs from classic examples of these biases in several respects. First, whereas confirmation bias typically stresses that people seek out information confirming their preferred conclusion,⁶⁰ our studies offered no opportunity for explicit information selection, and learning was clearly influenced by both sources. Second, participant beliefs in our studies most often shifted toward the *disconfirming* source, a finding difficult to reconcile with confirmation/ disconfirmation bias^{17,59}. Third, many studies use explicit ratings of source reliability to measure these biases^{21,22}, but our studies show that beliefs can shift toward disconfirming sources even when explicit measures (e.g. ratings of knowledge, trust, and accuracy) favor confirming sources. Thus the explicit ratings of evidence common in prior work may provide an incomplete picture of how a source actually influences belief formation. Fourth, to our knowledge, accounts of confirmation/disconfirmation bias have not been formulated quantitatively in a manner that permits computational exploration of the consequences of such bias in groups of learners. For these reasons, we believe the learning bias we have documented,

534 though clearly related to the idea that people strongly weight confirming information, is best
535 viewed as distinct from the classical confirmation or disconfirmation biases in the reasoning
536 literature.

537 *Cognitive dissonance.* A related idea is that the holding of mutually incompatible beliefs
538 or “cognitions” causes an uncomfortable state known as *cognitive dissonance*, and that beliefs
539 change in order to reduce such dissonance—a mechanism that can cause or maintain false
540 beliefs¹⁸. In a broad sense, cognitive dissonance arises in any learning task where feedback
541 conflicts with the learner’s expectation⁵⁷—the expectation and observed evidence can be viewed
542 as mutually incompatible mental states. Dissonance theory otherwise offers little scope for
543 understanding the behaviors documented here. For instance, when confronted with one close and
544 one distal source both differing from the current belief, the learner might reduce dissonance by (1)
545 adjusting her learning boundary to the midpoint between sources, thus minimizing prediction
546 error as in standard learning models, (2) deciding that the far source is incorrect, thus discounting
547 this feedback and moving toward the closer source or (3) deciding that neither source is correct,
548 leading to no change in the current belief. Because all three possibilities are consistent with
549 dissonance theory generally, it is difficult to attribute the specific change observed to cognitive
550 dissonance.

551 *Selective memory.* People better remember stimuli and events that cohere with their
552 existing long-term knowledge²⁵ or implicit theories²⁴, a phenomenon that can produce false
553 beliefs^{24,61}. For instance, when asked to memorize and later recall a list of semantically related
554 words (e.g. “sleep”, “night,” “pillow,” “blanket”), people will often falsely remember other
555 related words that did not appear on the list (e.g. “bed”)⁶². Such a phenomenon could potentially
556 influence new category learning if people better remember or even falsely remember experiences
557 that cohere with their existing store of knowledge. In general it has been suggested that people
558 may defend a belief against inconsistent evidence through “some combination of avoiding,
559 disbelieving, misperceiving, forgetting, or misremembering [the discordant information]”²⁴.
560 Coherence with existing knowledge is unlikely to explain the current results, however, because
561 our studies employed novel stimuli in a task for which participants had little relevant prior
562 knowledge.

563 *Social identity theory.* A broad literature in social cognition suggests that learners give
564 greater weight or credence to information provided by sources perceived as belonging to a similar
565 social group^{35,63,64}. For instance, political partisans rate their agreement with a particular policy
566 position differently depending on whether the policy is attributed to a Democrat or Republican
567 politician⁶⁵. Simulations with agent-based models have illustrated how such effects can lead to
568 polarized beliefs in social groups under certain circumstances^{66,67}. Such phenomena cannot
569 explain the current results, however, because the two sources were identified only by the face
570 images appearing with the labels on each trials, and these were matched for relevant social
571 characteristics (including gender, age, attractiveness, race, and trustworthiness) and also counter-
572 balanced across participants.

573 *Bounded rationality.* In many cases, truly rational behaviors can only arise from
574 computations that are prohibitively costly in terms of time or computational resources. As a
575 consequence, human cognition and behavior may exploit resource-efficient approximations of
576 rational models, which can produce suboptimal behaviors in certain circumstances. This
577 “bounded rationality” perspective has had a large impact on behavioral economics⁶⁸ and more
578 recently on Bayesian approaches to cognition⁶⁹. The approach does not help to explain the current
579 phenomena, however, because human behavior is in fact well-approximated by rational models in
580 both fully-supervised and semi-supervised variants of the same task, using the same stimuli⁴³. If
581 cognitive resources are sufficient to support rational learning in these cases, the failure of

analogous models to capture the current phenomena cannot then be attributed to poor approximations generated by a boundedly rational learning system.

Other cognitive agent-based models. Agent-based models have been used for many years to understand emergent phenomena in groups of interacting individuals, and have become increasingly influential in understanding political polarization^{11,66,67}, influence maximization on social networks⁷⁰, and many other phenomena related to mass human behavior⁷¹. The current project was motivated by this work, with the aim of filling an important gap: much of the research we have encountered adopts assumptions about learners that render formal analysis of emergent behaviors more tractable but may not accurately capture human behavior. We have instead developed an empirically-motivated heuristic model of the individual learner, then considered the group behaviors that emerge from such a model in very simple agent-based settings.

Our approach is similar in spirit to other recent work in the cognitive science literature. For instance, Mäs and colleagues have considered how belief polarization can emerge in social networks if agents give greater weight to information sources perceived as being similar along a set of social dimensions¹¹. Hahn and colleagues use agent-based models to show that, even with arguably rational learners, groups can arrive at incorrect beliefs depending upon the networks that structure their social interactions³⁶. The current research offers a complementary perspective: rather than assuming rationality, our agents behave according to the heuristic model derived from experiment, and we then consider the group behaviors that emerge under this heuristic model. The work contributes to a growing literature suggesting that cognitively-motivated agent-based models can elucidate specific mechanisms by which false beliefs persist in groups of learners.

Broader implications for other theories. The experiments from the main paper show that false beliefs can arise in a learning scenario where factors highlighted in the broader literature play no significant role. In this sense the work suggests a hypothesis about the curious stability of false beliefs in real-world scenarios that differs from canonical proposals in the literature^{31,36–38}. This does not invalidate the import of these other factors; indeed, the current work offers a new empirical framework for evaluating their impact both on learning and on explicit beliefs about information sources. For instance, motivational factors can be introduced by employing a categorization task with higher emotional stakes; social factors such as group identity, conformity or authority may be introduced by varying the number of sources or the information the participant receives about them; memory coherence effects may be studied by using real-world stimuli for which participants have prior knowledge; and so on. Such research may help to connect existing hypotheses about false belief formation more directly to computational theories of human learning and to adjudicate the relative contributions of different factors within a common task—goals shared by other recent work in this area^{36,39–41}.

Finally, although the HEW model describes and predicts behavior in our studies, it does not explain why human learning behaves in this way, nor do our results transparently arise from normative category-learning models. A more principled account will depend on a better understanding of the inferences we make about the information sources with which we interact. This is a topic of considerable research in cognitive pedagogy^{28–35}, and we hope insights derived there may also shed light on phenomena like those we have documented.

Still, the heuristic model is useful in suggesting new avenues for remediating the spread and persistence of false beliefs in social networks. If belief bubbles arise partly from the distribution of opinions to which a learner is exposed, there may exist algorithms for sharing information in social networks that promote widespread adherence to the truth rather than total discord.

Materials and Methods

Participants.

Together the studies included 404 participants, all undergraduate students at the University of Wisconsin-Madison between the ages of 18 and 23. All participants were volunteers recruited from the Psychology Human Participants Pool through SONA Systems Ltd. online postings of the experiments. Students received course credit for their participation. Studies were approved by the Internal Review Board for Social and Behavioral Science at UW-Madison.

Experimental setup.

Studies were run in groups of 8 or fewer participants, all completing the study in the same room, each at a stand-alone workstation with the experimental program loaded and running, and with workstations separated by vertical dividers so participants could not see other screens. After giving written informed consent, all participants read the same instructions, then completed the study at their own pace while the experimenter was in the room. Prior to their arrival in the lab, participants were randomly assigned to experimental conditions in cohorts to ensure all conditions completed at an equal pace. For instance, in a study with 8 conditions, the first 8 sign-ups were each assigned at random to one of the available conditions; the next 8 were again assigned randomly to conditions, etc. While each study opportunity was posted in sessions that included 8 slots, the number of participants run in each session varied. We collected data in this manner until all conditions had a pre-specified minimum number of participants. This procedure led to slightly different numbers of participants in each condition depending on the number of sign-ups and no-shows. Rather than discarding participants to ensure equal numbers, we instead include all participants in our analyses.

Software.

All experiments were conducted using the E-Prime 2.0 software suite run on Intel computers under the Windows 2000 operating system.

Stimuli.

Stimuli were 301 rendered visual images of 3-dimensional “supershapes” varying along a perceptual manifold from spiky to smooth as shown in Figure 1 of the main paper⁴². The shapes are useful for studying human category learning for several reasons. They are complex and organic-looking shapes unfamiliar to participants, and the single dimension along which they vary does not conform to an already-familiar perceptual feature such as color or orientation. Thus participants have no prior expectations about which stimulus features “matter” for categorization⁴³. Because they vary along a unidimensional perceptual manifold, no questions arise pertaining to feature salience or attentional selection. The stimuli have been used in several prior experiments involving fully-supervised, semi-supervised, and active learning^{27,42,43}, with results that match the predictions of statistically rational models of probabilistic inference^{34,43}.

In each study, feedback from the two sources was provided by displaying two face images on the screen and providing the written label (RAW or COOKED) for each source below the corresponding face. The two face images were matched in a prior norming study for age, gender, race, and attractiveness, to minimize social cues the learner might use to weight the feedback each source provided. Consistent with these prior norms, participants in the no-source control condition of Experiment 1—who viewed the faces but did not receive feedback associated with them—rated the two images as equivalent in their knowledge, trustworthiness, likely accuracy in guessing category labels, and attractiveness (paired t-test $|t| < .98$, $p = \text{n.s.}$ for all comparisons).

General procedure for all studies.

All experiments required participants to learn a binary classification of the stimuli, which can be represented as points ranging from 0-300 along a continuous one-dimensional perceptual manifold. Each experiment followed the same procedural template involving 6 steps:

1. Instructions. Participants read the following instructions on a computer screen:

You are hired as a shipmate on a voyage to lands you have never before visited. Your crew includes veteran sailors who know these lands intimately, mates who have sailed there for a few years, novices like yourself, and criminals who lie habitually. Along the way a storm sinks the ship. You and a handful of others are washed up on the shore of a deserted island. Your first task is to gather food for the group. A variety of strange fruits grow on the island, some of which taste better after being cooked and others that taste better when eaten raw. Luckily, your captain, a kind and knowledgeable sailor, has survived, and can show you which fruits are which.

The captain wants you to learn which fruits taste good after being cooked and which taste good when eaten raw so you can identify them on your own. By chance, he found one example of the fruit that needs to be cooked and one example of the fruit that needs to be eaten raw on the shore. He will show you each of them several times. After you guess, the captain will let you know whether or not you were correct.

These instructions aimed to create a scenario in which participants had little relevant prior knowledge, no affective stake in the response beyond accuracy, and no strong beliefs about the knowledge or trustworthiness of other information sources (i.e. fruit gathering companions).

2. Supervised learning. Participants completed twenty supervised learning trials involving ten repetitions of a single example from each category, presented in a permuted order determined independently for each participant. On each trial the participant viewed one stimulus on the screen, guessed the category (cooked/raw) by pressing a button, and received feedback from the knowledgeable and trustworthy “captain”. The stimulus remained on the screen until a button press was detected. Feedback showed the image of the item with the correct label printed underneath and remained onscreen for 1s. Trials were separated by a blank screen appearing for 300ms.

3. Initial boundary test. Participants next read the following instructions:

Now it is time to practice. Based upon what you have learned from the captain, try to decide whether the following fruits are better cooked or eaten raw by yourself. No feedback will be provided by your captain.

They then classified, without feedback, 19 stimuli ranging from 60 to 240 in intervals of 10, each appearing twice for 38 total trials, with stimulus ordering permuted independently for each participant. On each trial participants viewed a single item and pressed a button to indicate their category decision; stimuli remained onscreen until a button-press was detected. No feedback was provided. Trials were separated by a blank screen appearing for 1s. Data from this phase were used to estimate the participant’s initial category boundary location.

Following this task participants were shown a random stimulus that they could adapt to be more/less spiky using the wheel of the mouse. They were asked to adjust the shape until it showed the item for which they were most uncertain about category membership. This was intended to provide a measure of their explicit belief about the category boundary, but faulty logging by the E-prime software produced many missing trials, so these data were not analyzed.

4. Learning from sources. In this phase participants first read the following instructions:

Now that you have seen some fruits that need to be cooked and some that can be eaten raw, you are sent out in the forest with two other survivors to gather fruits. The captain puts you in charge---for each fruit you find, you must decide whether it goes into the "Cooked" basket or the "Raw" basket. After you make your decision, your companions will give you their own opinions, which may or may not be helpful. Use their input as best you can to help figure out how to tell fruit that needs to be cooked from fruit that can be eaten raw.

They then classified all 301 stimuli presented in permuted order generated independently for each participant. On each trial a single item appeared on the screen together with two face images depicting the two companions, and participants again indicated their categorization decision by button press. In experimental conditions, participants then received feedback from both companions in the form of written text (RAW or COOKED) below each face which remained onscreen until the participant pressed the spacebar. No other indication of the correct choice was given. The labels provided by each source were generated deterministically according to a fixed category boundary, with the location of the boundaries being the primary independent variable manipulated across studies. In one-source control conditions, all aspects of the design were identical except that only one companion provided feedback (ie, both companions were shown on every trial but category labels only appeared beneath one, the same companion on every trial). In the no-source control conditions, both faces appeared on every trial but no category labels were shown. Trials were separated by a blank screen appearing for 300ms.

5. Final boundary test. Participants classified the same 38 items used in the initial boundary test without feedback, following the identical procedure.

6. Source ratings. In the final phase, participants rated each companion's level of knowledge ("How well can this person tell the difference between raw and cooked fruits?"), trustworthiness ("How much do you trust this person to tell the truth?"), accuracy ("Out of 100 guesses, how often will this person get the right answer?"), and, as a control, attractiveness ("How attractive is this person?"). On each trial, participants viewed a single face with the corresponding question written below, and indicated their response using a slider on a scale ranging from 0 to 100. Participants judged both faces for each question in the same fixed order. Judgments were collected for both experimental and control conditions.

All experimental studies fully counterbalanced three factors: (1) which end of the stimulus continuum (spiky/smooth) was labeled "raw" vs "cooked", (2) which face image was paired with which feedback boundary, and (3) whether the farther source boundary was toward the spiky or smooth side of the continuum relative to the expected initial learner boundary. Control conditions balanced the same factors where these were relevant—for instance, in "no-source" control conditions, neither source provides feedback to the learner so there was no need to counterbalance factors related to source feedback, but the remaining nuisance factor—which end of the continuum was labeled spiky/smooth—was counterbalanced.

Procedural details for each study.

Experiment 1. 146 participants completed this study. In all conditions the supervised phase used a labeled stimulus at 100 and 200, with an expected initial boundary near 150. In the experimental conditions, a "close" source boundary was 15 units from the expected initial boundary at either 165 or 135, while a "far" source boundary was 100 units from the initial boundary on the opposite side of the range (at 50 or 250 respectively). In the no-source control condition, participants viewed the faces of the two companions with no feedback from either source. In the one-source control conditions, participants viewed both faces with feedback appearing under only one face. This feedback was generated according to a category boundary

either 15 units or 25 units from the expected initial boundary (ie at 135 or 165 in the one-source-15 control condition, and at 125 or 175 in the one-source-25 condition).

Experiment 2. 64 participants completed Experiment 2. As with Experiment 1, the supervised phase used labeled items at 100 and 200 to teach an initial boundary near 150. In the *learning from sources* phase a “fixed” source provided feedback based on a boundary at either 135 or 165, while a second “variable” source used a boundary that was sampled at random, independently for each participant, with uniform probability from the opposite side of the stimulus range (e.g. from 151-300 when the fixed source was at 135, or from 0-149 when the fixed source was at 165).

Experiment 3. 35 participants completed Experiment 3. Participants were split at random into 2 groups, A and B. During the *supervised learning* phase group A observed labeled items at 30 and 170 while group B observed labeled items at 130 and 270. Note that both supervised conditions are consistent with a “true” category boundary somewhere between 130-170, but the tendency to place initial boundaries at the midpoint between labeled examples leads the two groups to form differing initial boundaries. During the *learning from sources* stage, all participants observed feedback from two companions based on boundaries at 100 and 200.

Experiment 4. 159 participants completed experiment 4. All participants were assigned at random to one of two groups, A and B. In the supervised phase, group A again observed labelled items at 70 and 170 while group B observed labelled items at 130 and 230. Participants in each group were then assigned randomly to one of 5 conditions: strong-agree, moderate-agree, moderate-disagree, strong-disagree, or no-source-control. In the four experimental conditions, all participants viewed feedback from two sources employing decision boundaries at a fixed distance on either side of the range midpoint of 150. Across these conditions, the distance from source boundary to range midpoint varied: from 25 units away (ie at 125 and 175) in the strong-agree condition, to 50 units away (100 and 200) for moderate-agree, to 75 units (75 and 225) for moderate-disagree, to 100 units (50 and 250) for strong-disagree. In the no-source control condition, participants received the same supervised learning and viewed the faces of the two companions on each trial but did not receive feedback from either source.

Statistical Analysis.

All analyses were carried out in R version 3.6.1 and are documented in JuPyTeR notebooks, together with de-identified data, at github.com/ttrogers/frigo-chen-rogers.

All studies involve measuring how an initial category boundary formed from sparse supervised learning changes following learning from two conflicting sources. This requires estimating each participant’s initial and final category boundaries and taking the difference between these as a measure of change. Initial and final boundary locations were estimated for each participant from the unsupervised category decisions collected during the initial-boundary-test and final-boundary-test phases of the study. For each participant we fit a logistic regression model to predict binary category decisions from the coordinate of the corresponding item along the stimulus range, then from the fitted model parameters computed the expected stimulus location where the participant showed equal probability of each category label. This procedure was carried out separately for data from the initial and final boundary-test phases for each subject.

To estimate the boundary-shift for each subject, we took the simple difference (final location – initial location), then multiplied this by -1 in conditions where the farther source boundary was situated toward the low end of the stimulus range so that positive shift numbers always indicate a shift toward the *far* source boundary (or, in Experiment 2, toward the variable source boundary).

Outliers were identified and removed from the dataset as follows. First, a small number of boundary estimates (n=4) yielded impossible values (ie, boundaries lying outside the full stimulus range), indicating subjects who were confused or guessing randomly. These participants were removed from the sample. Next we Z-scored the initial boundary estimates and removed participants with estimates greater than two standard deviations from the mean to ensure that all initial boundaries were representative of the sample. For all experiments this resulted in normally-distributed initial boundary estimates, with the estimate for every participant lying between the two labeled items from the supervised phase. Finally, we Z-scored estimates of the boundary shift and again discarded participants where the shift was greater than 2 standard deviations from the mean. The resulting boundary shift estimates were again normally distributed in all experiments. The same procedure was applied to all studies, removing ~7% of the 404 participants across all studies.

Detailed Methods for Simulation Study 1.

Discriminative models learn a function that maps from a continuous feature space to a probability distribution over discrete, mutually exclusive categories. In our studies stimuli are characterized by a single continuous feature dimension representing the item's position along the continuum from spiky to smooth, and each item belongs to one of just two mutually-exclusive categories—so binomial regression provides a class of discriminative models to simulate category learning in this case. Let x_i indicate the continuous-valued coordinate of a stimulus i along the 1D manifold and let $y_i \in \{0,1\}$ indicate the binary category label provided by an oracle for item i . The discriminative model specifies the probability that i belongs to the category labeled 1 via the sigmoid function:

$$p(y_i = 1 | x_i, \alpha, \beta) = \frac{1}{1 + e^{-(\alpha + \beta x_i)}}$$

...where α and β are free parameters controlling the offset and slope of the fitted curve, respectively. For a sample of data consisting of n (x,y) pairs one can analytically solve for the parameters α and β that minimize the discrepancy between the output of this function and the true binary category labels, formulated as the logistic loss:

$$\frac{1}{n} \sum_{i=1}^n -y_i(\log \hat{y}_i) - (1 - y_i) \log(1 - \hat{y}_i)$$

...where $\hat{y}_i$ is the probability output by the sigmoid function for x_i .

A fitted model then provides, for any stimulus x_i , an estimated probability p that it should have category label 1 (and probability $1-p$ that it should have category label 0). Learning can be simulated by estimating model parameters from an ever-expanding sample of data. The boundary is the coordinate x_b that gives an equal probability of 0.5 to both categories, which can be computed directly from the model parameters

$$x_b = \frac{-\alpha}{\beta}$$

This standard formulation assumes an oracle providing a single true label for all examples. In our studies, two labels are provided for each item. These sometimes conflict, and the learner may weight input from each source differently. To accommodate this setting we apply weighted logistic regression. Let $L_s(x_i)$ be the logistic loss for x_i on the label provided by source s , let $\theta \in [0,1]$ be a weight given to source 1, and $1-\theta$ be a weight given to source 2. We then find model parameters to minimize:

$$\frac{1}{n} \sum_{i=1}^n \theta L_1(x_i) + (1 - \theta) L_2(x_i)$$

So the loss for each data point x_i is just a weighted sum of the losses for the labels provided by each source, with source weighting specified by θ .

Parametric generative models provide one way of implementing *prototype* theories in cognitive science. They view participants as learning, for each category, the parameters of a distribution that characterizes the probability density of category exemplars in the stimulus feature space. Gaussian mixture models (GMMs) provide the paradigm example. They assume that items from a given category are sampled from a Gaussian in the stimulus feature space. Different categories are characterized by Gaussians with different parameters and different base frequencies or prior probabilities. The task of the learner is to accurately estimate the base frequency and parameters of the distribution for each category given samples of labeled data drawn from the mixture. New items can then be classified by using the parameter estimates to compute the posterior probability that the item was drawn from each component of the mixture.

For the current studies, because stimuli vary along a single feature dimension, supervised learning in a GMM framework involves estimating the mean and variance of the stimulus coordinates for all items receiving a given category label, as well as the prior probability of the label. As before let x_i denote the coordinate of stimulus i along the continuous feature-dimension x , and let y_i denote the true binary category label for x_i . Additionally, let $z \in [0,1]$ be the prior probability $y_i=1$ and $N(\mu_{y=k}, \sigma_{y=k}^2)$ be a Gaussian distribution on x with mean μ and variance σ^2 from which category k items are sampled. For n labeled data points, we estimate:

$$z = \frac{1}{n} \sum_{i=1}^n y_i$$

... and estimate $\mu_{y=k}, \sigma_{y=k}^2$ as the mean and variance on x across all labeled items for which $y=k$. Since there are only two mutually exclusive categories, these parameter estimates for category $y=1$ can be written:

$$\mu_{y=1} = \frac{\sum_{i=1}^n y_i x_i}{\sum_{i=1}^n y_i}, \sigma_{y=1}^2 = \frac{\sum_{i=1}^n y_i (x_i - \mu_{y=1})^2}{\sum_{i=1}^n y_i}$$

...and for $y=0$:

$$\mu_{y=0} = \frac{\sum_{i=1}^n (1 - y_i) x_i}{\sum_{i=1}^n (1 - y_i)}, \sigma_{y=0}^2 = \frac{\sum_{i=1}^n (1 - y_i) (x_i - \mu_{y=0})^2}{\sum_{i=1}^n (1 - y_i)}$$

Learning is simulated by estimating these parameters from ever-increasing samples of labeled data.

To categorize a new, unlabeled item x_j , we use current parameter estimates to compute, for each category k , the posterior probability that the item was sampled from k , that is, $p(y_j = k | x_j)$. Bayes' rule specifies that:

$$p(y_j = k | x_j) \propto p(y_j = k) p(x_j | y_j = k)$$

Here $p(y_j = 1)$ is the prior estimated as z above, while the prior on the category $y=0$ is just $1-z$ in the binomial case. $p(x_j | y_j = k)$ is the probability of sampling x_j from $N(\mu_{y=k}, \sigma_{y=k}^2)$ which we can denote as $p(x_j | \mu_k, \sigma_k^2)$. So:

$$p(y_j = 1 | x_j) \propto z \times p(x_j | \mu_1, \sigma_1^2)$$

$$p(y_j = 0|x_j) \propto (1 - z) \times p(x_j|\mu_0, \sigma_0^2)$$

...and since these are the only two categories, these proportional relations can be standardized to probabilities by taking:

$$p(y_j = 1|x_j) = \frac{z \times p(x_j|\mu_1, \sigma_1^2)}{z \times p(x_j|\mu_1, \sigma_1^2) + (1 - z) \times p(x_j|\mu_0, \sigma_0^2)}$$

$$p(y_j = 0|x_j) = 1 - p(y_j = 1|x_j)$$

The category boundary is again defined as the coordinate x_b where $p(y_j = 1|x_b) = p(y_j = 0|x_b) = 0.5$.

Learning in the GMM framework involves updating estimates of the parameters (mean and variance) associated with each category distribution. For supervised learning all data points receiving a given label contribute equally to these estimates. How can this approach be adapted for the current case, where learners can receive conflicting information about feedback and may weight sources? One approach known as *soft supervision* weights the data provided by each source, similar to weighted logistic regression. Let $\theta \in [0,1]$ be the weight given to source 1, $1-\theta$ be the weight given to source 2, and $y_{s,i} \in \{0,1\}$ be the binary label given by source s to item i . Weighted estimates of the mean and variance for the category labeled with $y=1$ can then be written:

$$\mu_{y=1} = \frac{\sum_{i=1}^n (\theta y_{1,i} x_i + (1-\theta) y_{2,i} x_i)}{\sum_{i=1}^n (\theta y_{1,i} + (1-\theta) y_{2,i})}$$

$$\sigma_{y=1}^2 = \frac{\sum_{i=1}^n (\theta y_{1,i} (x_i - \mu_{y=1})^2 + (1-\theta) y_{2,i} (x_i - \mu_{y=1})^2)}{\sum_{i=1}^n (\theta y_{1,i} + (1-\theta) y_{2,i})}$$

...and for the category $y=0$:

$$\mu_{y=0} = \frac{\sum_{i=1}^n (\theta (1-y_{1,i}) x_i + (1-\theta) (1-y_{2,i}) x_i)}{\sum_{i=1}^n (\theta (1-y_{1,i}) + (1-\theta) (1-y_{2,i}))}$$

$$\sigma_{y=0}^2 = \frac{\sum_{i=1}^n (\theta (1-y_{1,i}) (x_i - \mu_{y=0})^2 + (1-\theta) (1-y_{2,i}) (x_i - \mu_{y=0})^2)}{\sum_{i=1}^n (\theta (1-y_{1,i}) + (1-\theta) (1-y_{2,i}))}$$

...and for the prior $p(y = 1)$:

$$\frac{1}{n} \sum_{i=1}^n (\theta y_{1,i} + (1-\theta) y_{2,i})$$

When the two sources provide the same label, these reduce to standard supervised learning estimates. Likewise, when $\theta=1$ they reduce to standard supervised learning considering only labels from source 1, and when $\theta=0$ they reflect supervised learning for labels from source 2 only. When the two sources provide different labels and $0 < \theta < 1$, estimates of the mean and variance for each category are influenced by both sources in proportion to the weighting parameter.

It is worth noting that this approach is related to models developed for semi-supervised GMMs⁴³. In the semi-supervised setting, training data include a mix of fully-labeled and completely unlabeled examples, which are jointly used to estimate the parameters of the distributions on x that characterize each category. It is clear how labeled examples contribute to these estimates, but how should unlabeled samples be used? Because we don't know from which category a sample was drawn, it is unclear which distribution's parameters should be updated. One common approach called *boosting* weights the contribution of an unlabeled item to a given

category's parameters proportional to the posterior probability that the item was sampled from the category. For instance, if the current model suggests .8 probability that the label is 1 and a .2 probability that the label is 0, the parameters for both categories are updated, but the unlabeled item has a larger impact on the estimates for category 1 than for category 0. Boosted semi-supervised GMMs of this kind have accounted for several aspects of human semi-supervised learning using the same task and stimuli in prior work⁴³, so we viewed the soft-supervised variant just described as a promising candidate for understanding learning with multiple conflicting labels.

Detailed Methods for Simulation Study 2

We first define the agent-based model for the minimal case of two learners and one oracle. Let l_t be a learning agent's belief at time t about the location of the category boundary along the stimulus manifold, $\{l \in \mathbb{R} \mid 0 \leq l \leq 300\}$. After receiving input from two sources possessing category boundary beliefs m and n , the agent updates its belief proportional to a weighted sum of the difference between its current belief and each source's belief:

$$\Delta l_t = \theta(m - l_t) + (1 - \theta)(n - l_t) \quad \text{Eqn. S1}$$

...where $\theta \in [0,1]$ is the proportional weight given to source m . The magnitude of the update at time t is controlled by a rate constant $\gamma \in [0,1]$ so that $l_{t+1} = l_t + \gamma \Delta l_t$.

A source agent can be either another learner or an oracle—the sole difference is that oracles do not update their beliefs (since they know and report the ground truth). Consider the minimal case scenario involving two learners and one oracle. On each epoch, each learning agent takes one turn “listening” to the other two agents and updating its current belief according to equation S1. Epochs iterate until updates for both agents stop, that is, applying the update rule leads to a belief change of zero for both. Let l_t and m_t denote the beliefs of the two learning agents at time t and n denote the ground truth reported by the oracle. Then the updates are:

$$\Delta l_t = \theta_t(m_t - l_t) + (1 - \theta_t)(n - l_t)$$

$$\Delta m_t = \theta_t(l_t - m_t) + (1 - \theta_t)(n - m_t)$$

Updates to l depend on m , while the updates to m depend on l , so the two learners are *coupled* and form a dynamical system. We can find the steady states of this system analytically, by solving for values of l and m where the update rule produces no change to either, that is where $\Delta l = \Delta m = 0$. The solution will depend on how the weighting parameter θ_t is set.

Consider the simple case where all sources are always given equal weighting, $\theta_t = 1 - \theta_t = \frac{1}{2}$ for all t . Learning will stop when Δl and Δm are both zero. In that case they are equal, so

$$\frac{(m - l) + (n - l)}{2} = \frac{(l - m) + (n - m)}{2}$$

Simplifying this equality shows that $\Delta l = \Delta m$ when $m = l$; that is, belief updates for the two learners will be identical only when their beliefs are identical. Further, equation S1 indicates that $\Delta l = 0$ only when $2l - m = n$. Putting the two observations together we can see that belief updates stop for both learners (ie $\Delta l = \Delta m = 0$) only when $l = m = n$, that is, when both learners agree with the oracle and each other.

That is the result shown through simulation in Figure 4A (left) of the Results section. In each simulation, the two learning agents were each given an initial belief selected from a grid in the range $[0, 300]$, while the oracle reported a ground-truth boundary of 150. On each round the learning agents updated their beliefs according to equation S1 with $\theta = 0.5$ and $\gamma = 0.1$. The

figure shows the results of 200 such updates, after which both agents converged on the oracle's belief regardless of their starting beliefs.

The more interesting analysis considers what happens when each learner updates its belief in the manner suggested by our experiments. Let l_t again be the belief of a learning agent at time t , and let m_t be the belief of the *closer* of two source agents at time t , while n_t is the belief of the other source agent. The weighting parameter for a given learning agent at time t is given by the heuristic evidence-weighting function derived in Experiment 2:

$$\theta_t = 1 - \frac{m_t - l_t}{m_t - l_t + \beta} \quad \text{Eqn. S2}$$

...where $\{\beta \in \mathbb{R} \mid \beta > 0\}$ is a real-valued positive number controlling the slope of the HEW function. (Here we have set the offset term from Equation 2 to zero to simplify analysis but the generalization to offsets with positive values is straightforward.) Note that $\theta \rightarrow 0$ only as $|m - l| \rightarrow \infty$, so the weight on the closest source m will never be zero. When $m = l$, however, the closest source and the learner have identical beliefs so $m - l = 0$ and $\theta = 1$. Plugging these values into Equation S1 shows that $\Delta l = 0$ whenever $m = l$. In other words, the learning agent stops updating its beliefs when it agrees perfectly with its nearest source, regardless of the input from a second source. It should be clear that the second learning agent m will also stop updating its beliefs in this case, so that the coupled system of learners has stable states lying along the identity line $m = l$ in the space of learning-agent beliefs.

This is the behavior shown by simulation in Figure 4A (right). We replicated the preceding simulation exactly, but using the HEW function, parameterized as indicated by Experiment 2 (including the offset α), to set θ_t independently on each round. As expected from the analysis, learning always stopped within α of the identity line $m = l$ indicating states where the two learners have beliefs similar enough they give each other full weight. When the oracle was the closest source to both learning agents, the learners converged on the ground truth, that is $l = m = n$, which transparently lies on the identity line. When the closest source for one learning agent was the other learning agent, they came into perfect agreement before either reached the ground truth, and learning stopped for both—producing considerable disagreement about the boundary location across all pairs of simulated learners.

Multiple learners and oracles. As previously, each learning agent viewed the opinions of two source agents once in every epoch and updated its current belief according to the HEW function parameterized as determined in Experiment 2. Source agents were either other learners or oracles that report a common ground truth. We simulated the evolution of beliefs over time by updating each learning agent once in each epoch using Equation S1 for the update, the HEW function to set the weighting parameter independently for each learner on each round, and a fixed rate constant of $\gamma = 0.05$. We again ran 200 updates at which point the mean absolute change across learning agents fell close to zero within a tolerance of 0.03, then considered how the resulting distribution of beliefs relate to one another and to the ground truth. Initial beliefs for 10 learners were sampled randomly with uniform probability from the range [140-240] for each learning agent, while 2 oracles reported a ground-truth boundary at 150. This corresponds to a situation in which the initial group consensus (expressed as the mean or median of the initial beliefs across all agents) is incorrect. We then considered four simple information-sharing algorithms as described in Results. Note that, although the figure shows the results of a single model run, the same pattern of results was observed across many runs varying the initial distribution of beliefs and the rate constant, for each information-sharing algorithm.

References

1. Hanninen, O., Farago, M. & Monos, E. Ignaz Philipp Semmelweis, The Prophet of Bacteriology. *Infect. Control* **4**, 367–370 (1983).
2. McIntyre, P. & Leask, J. Improving Uptake of MMR Vaccine. *BMJ* **336**, 729–30 (2008).
3. Cockman, P., Dawson, L., Mathur, R. & Hull, S. Improving MMR vaccination rates: herd immunity is a realistic goal. *Bmj* **343**, d5703–d5703 (2011).
4. Bursztyn, L., Rao, A., Roth, C. & Yanagizawa-Drott, D. Misinformation During a Pandemic. *SSRN Electron. J.* (2020) doi:10.2139/ssrn.3580487.
5. Pennycook, G., Mcphetres, J., Bago, B. & Rand, D. G. Attitudes about COVID-19 in Canada, the UK, and the USA: A novel test of political polarization and motivated reasoning. *[Preprint]* (2020).
6. Reported cases and deaths by country, territory, or conveyance. *WorldoMeter* (2020).
7. Griffiths, T. L., Kemp, C. & Tenenbaum, J. B. Bayesian models of cognition. in *The Cambridge handbook of computational psychology*. 59–100 (Cambridge University Press, 2008). doi:10.1017/CBO9780511816772.006.
8. Anderson, J. R. The adaptive nature of human categorization. *Psychol. Rev.* **98**, 409–429 (1991).
9. Sanborn, A.N., Griffiths, T.L., & Navarro, D. J. A more rational model of categorization. in 726–731 (2006).
10. Griffiths, T. L., Sanborn, A. N., Canini, K. R. & Navarro, D. J. Categorization as nonparametric Bayesian density estimation. *Probabilistic Mind Prospect. Bayesian Cogn. Sci.* (2012) doi:10.1093/acprof:oso/9780199216093.003.0014.
11. Mäs, M. & Flache, A. Differentiation without distancing. explaining bi-polarization of opinions without negative influence. *PLoS One* **8**, (2013).
12. Lorenz, J. M. N. T. S. Individual attitude change and societal dyamics: Computational experiments with psychological theories. *J. Chem. Inf. Model.* **53**, 1689–1699 (2020).
13. Perfors, A. Why do echo chambers form ? The role of trust , population heterogeneity , and objective truth. 918–923 (2018).
14. Madsen, J. K., Bailey, R. M. & Pilditch, T. D. Large networks of rational agents form persistent echo chambers. *Sci. Rep.* **8**, 1–8 (2018).
15. Lorenz, J. & Urbig, D. About the power to enforce and prevent consensus by manipulating communication rules. *Adv. Complex Syst.* **10**, 251–269 (2007).
16. Flache, A. *et al.* Models of social influence: Towards the next frontiers. *Jasss* **20**, (2017).
17. Kunda, Z. The Case for Motivated Reasoning. *Psychol. Bull.* **108**, 480–98 (1990).
18. Festinger, L. *A Theory of Cognitive Dissonance*. (Stanford University Press, 1957).
19. Devine, P. G., Tauer, J. M., Barron, K. E., Elliot, A. J. & Vance, K. M. Moving beyond attitude change in the study of dissonance-related processes. *Cogn. dissonance Prog. a pivotal theory Soc. Psychol.* 297–323 (1999) doi:Doi 10.1037/10318-012.
20. Dawson, E., Gilovich, T. & Regan, D. T. Motivated Reasoning and Performance on the Wason Selection Task. doi:10.1177/014616702236869.
21. Lord, C. G., Ross, L. & Lepper, M. R. Biased assimilation and attitude polarization: The

- 1057 effects of prior theories on subsequently considered evidence. *J. Pers. Soc. Psychol.* (1979)
1058 doi:10.1037/0022-3514.37.11.2098.
- 1059 22. Taber, C. S., Lodge, M. & Brook, S. Motivated Skepticism in the Evaluation of Political
1060 Beliefs Physicists do (Glanz. *Am. J. Pol. Sci.* **50**, 755–769 (2006).
- 1061 23. Kim, K.-H., Payne, G. T. & Tan, J. A. An examination of cognition and affect in strategic
1062 decision making. *Int. J. Organ. Anal.* **14**, 227–294 (2006).
- 1063 24. Jacobson, G. C. Perception, Memory, and Partisan Polarization on the Iraq War. *Polit. Sci.*
1064 *Q.* **125**, 31–56 (2010).
- 1065 25. Yeo, S. K., Cacciatore, M. A. & Scheufele, D. A. News Selectivity and Beyond: Motivated
1066 Reasoning in a Changing Media Environment. (2015).
- 1067 26. Agrawal, N. & Maheswaran, D. Motivated Reasoning in Outcome-Bias Effects Outcome
1068 Bias: The Case for Motivated Reasoning. *J. Consum. Res.* (2005).
- 1069 27. Zhu, X., Rogers, T., Qian, R. & Kalish, C. Humans perform semi-supervised classification
1070 too. *Proc. Natl. Conf. Artif. Intell.* **1**, 864–869 (2007).
- 1071 28. Shafto, P., Goodman, N. D. & Frank, M. C. Learning From Others: The Consequences of
1072 Psychological Reasoning for Human Learning. *Perspect. Psychol. Sci. a J. Assoc.*
1073 *Psychol. Sci.* **7**, 341–351 (2012).
- 1074 29. Shafto, P., Goodman, N. D. & Griffiths, T. L. A rational account of pedagogical reasoning:
1075 Teaching by, and learning from, examples. *Cogn. Psychol.* **71**, 55–89 (2014).
- 1076 30. Goodman, N. D. & Stuhlmüller, A. Knowledge and implicature: modeling language
1077 understanding as social cognition. *Top. Cogn. Sci.* **5**, 173–184 (2013).
- 1078 31. Hahn, U. & Warren, P. A. Perceptions of Randomness: Why Three Heads Are Better Than
1079 Four. *Psychol. Rev.* **116**, 454–461 (2009).
- 1080 32. Landrum, A. R., Eaves, B. S. J. & Shafto, P. Learning to trust and trusting to learn: a
1081 theoretical framework. *Trends Cogn. Sci.* **19**, 109–111 (2015).
- 1082 33. Ho, M. K., Cushman, F., Littman, M. L. & Austerweil, J. L. People teach with rewards and
1083 punishments as communication, not reinforcements. *J. Exp. Psychol. Gen.* **148**, 520–549
1084 (2019).
- 1085 34. Young, A. G., Alibali, M. W. & Kalish, C. W. Disagreement and causal learning: others’
1086 hypotheses affect children’s evaluations of evidence. *Dev. Psychol.* **48**, 1242–1253 (2012).
- 1087 35. Kahan, D. M. Ideology, Motivated Reasoning, and Cognitive Reflection: An Experimental
1088 Study. *SSRN Electron. J.* (2012) doi:10.2139/ssrn.2182588.
- 1089 36. Hahn, U., Hansen, J. U. & Olsson, E. J. Truth tracking performance of social networks:
1090 how connectivity and clustering can make groups less competent. *Synthese* (2018)
1091 doi:10.1007/s11229-018-01936-6.
- 1092 37. Love, B. C., Ramscar, M., Griffiths, T., Ramscar, M. & Jones, M. Generative and
1093 Discriminative Models in Cognitive Science. 25–26 (2001).
- 1094 38. Jern, A., Chang, K. -m. K. & Kemp, C. Bayesian belief polarization. *Adv. Neural Inf.*
1095 *Process. Syst.* 853–861 (2009).
- 1096 39. Jern, A., Chang, K. K. & Kemp, C. Belief Polarization Is Not Always Irrational. *Psychol.*
1097 *Rev.* **121**, 206–224 (2014).
- 1098 40. Cook, J. & Lewandowsky, S. Rational Irrationality: Modeling Climate Change Belief

- 1099 Polarization Using Bayesian Networks. *Top. Cogn. Sci.* **8**, 160–179 (2016).
- 1100 41. Sherman J.W., Kruschke J.K., Sherman S.J., Percy E.J., Petrocelli J.V., C. F. R. Attentional
1101 processes in stereotype formation: a common model for category accentuation and illusory
1102 correlation. *J. Pers. Soc. Psychol.* **96**, 305–323 (2009).
- 1103 42. Zhu, J., Rogers, T., Qian, R. & Kalish, C. Humans perform semi-supervised learning too.
1104 *Proc. Natl. Conf. Artif. Intell.* **22**, 864–869 (2009).
- 1105 43. Gibson, B. R., Rogers, T. T. & Zhu, X. Human Semi-Supervised Learning. *Top. Cogn. Sci.*
1106 **5**, 132–172 (2013).
- 1107 44. Gilbert, N. & Banks, S. Platforms and methods for agent-based modeling. *Proc. Natl.*
1108 *Acad. Sci. U. S. A.* **99**, 7197–7198 (2002).
- 1109 45. Berry, B. J. L., Kiel, L. D. & Elliott, E. Adaptive agents, intelligence, and emergent human
1110 organization: Capturing complexity through agent-based modeling. *Proc. Natl. Acad. Sci.*
1111 *U. S. A.* **99**, 7187–7188 (2002).
- 1112 46. Bienenstock, E. J., Bonacich, P., Oliver, M. The effect of network density and
1113 homogeneity on attitude polarization. *Soc. Networks* **12**, 153–172 (1990).
- 1114 47. Flache, A. & Macy, M. W. Small worlds and cultural polarization. *J. Math. Sociol.* **35**,
1115 146–176 (2011).
- 1116 48. Kook Lee, J., Choi, J., Kim, Cheonsoo, K., Yonghwan, K. Social Media, Network
1117 Heterogeneity, and Opinion Polarization. *J. Commun.* **64**, 702–722 (2014).
- 1118 49. Garimella, K., Gionis, A., Parotsidis, N. & Tatti, N. Balancing information exposure in
1119 social networks. *Adv. Neural Inf. Process. Syst.* **2017-Decem**, 4664–4672 (2017).
- 1120 50. Aslay, C., Matakos, A., Galbrun, E., Gionis, A. Maximizing the diversity of exposure in a
1121 social network. *IEEE Int. Conf. Data Min.* 863–868 (2018).
- 1122 51. Hsu, A. & Griffiths, T. Effects of generative and discriminative learning on use of category
1123 variability. *32nd Annu. Conf. Cogn. Sci. Soc.* 242–247 (2010).
- 1124 52. Griffiths, T. L., Canini, K. R., Sanborn, A. N. & Navarro, D. J. Unifying rational models of
1125 categorization via the hierarchical Dirichlet process. *Proc. 29th Annu. Conf. Cogn. Sci.*
1126 *Soc.* 323328 (2007).
- 1127 53. Klayman, J. *Varieties of confirmation bias*. (New York: Academic Press, 1995).
- 1128 54. Nickerson, R. S. Confirmation Bias : A Ubiquitous Phenomenon in Many Guises. *Rev.*
1129 *Gen. Psychol.* **2**, 175–220 (1998).
- 1130 55. Sobel, D. M. & Munro, S. E. Domain Generality and Specificity in Children ’ s Causal
1131 Inference About Ambiguous Data. **45**, 511–524 (2009).
- 1132 56. Shiloh, S., Drori, E., Orr-Urtreger, A. & Friedman, E. Being ‘at-risk’ for developing
1133 cancer: Cognitive representations and psychological outcomes. *J. Behav. Med.* **32**, 197–208
1134 (2009).
- 1135 57. Wyer, R. S. & Albarracín, D. *Belief Formation, Organization, and Change: Cognitive and*
1136 *Motivational Influences*. (2004).
- 1137 58. Kahneman, D. *Thinking, Fast and Slow*. (Farrar, Straus and Giroux, 2011).
- 1138 59. Taber, C. S. & Lodge, M. The Illusion of Choice in Democratic Politics: The Unconscious
1139 Impact of Motivated Political Reasoning. *Polit. Psychol.* **37**, (2016).
- 1140 60. Jonas, E., Schulz-hardt, S., Frey, D. & Thelen, N. Confirmation Bias in Sequential

- Information Search After Preliminary Decisions : An Expansion of Dissonance Theoretical Research on Selective Exposure to Information. *J. Pers. Soc. Psychol.* **80**, 557–571 (2001).
61. Chapman, L. J. & Chapman, J. P. Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of abnormal psychology* vol. 74 271–280 (1969).
 62. Stadler, MA; Roediger, HL; McDermott, K. Norms for word lists that create false memories. *Mem Cogn.* **27**, 494–500 (1999).
 63. Kahan, D. M. *Misconceptions, misinformation, and the logic of identity-protective cognition. Cultural Cognition Project Working Paper Series No. 164* <https://ssrn.com/abstract=2973067> (2017).
 64. Pennycook, G. & Rand, D. G. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci.* **116**, 2521 LP – 2526 (2019).
 65. Druckman, James N, Bolsen, T. Framing, motivated reasoning, and opinions about emergent technologies. *J. Commun.* **61**, 659–688 (2011).
 66. Hogg, M. A., Turner, J. C. & Davidson, B. Polarized Norms and Social Frames of Reference: A Test of the Self-Categorization Theory of Group Polarization. *Basic Appl. Soc. Psych.* **11**, 77–100 (1990).
 67. Mummendey, A., Kessler, T., Klink, A. & Mielke, R. Strategies to cope with negative social identity : Predictions by Social Identity Theory and Relative Deprivation Theory. *J. Pers. Soc. Psychol.* **76**, 1–5 (1999).
 68. Simon, H. A. *Models of Bounded Rationality, Volume 3.* (MIT Press, 1997).
 69. Griffiths, T. L., Lieder, F. & Goodman, N. D. Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* **7**, 217–229 (2015).
 70. Li, Y., Fan, J., Wang, Y., Tan, K. Influence maximization on social graphs: a survey. *IEEE Trans. Knowl. Data Eng.* **30**, 1852–1872 (2018).
 71. Bonabeau, E. Agent-based modeling: Methods and techniques for simulating human systems. *Proc. Natl. Acad. Sci.* **99**, 7280–7287 (2002).

Acknowledgments

Funding:

This work was supported by:

NSF 154548: DESE LUCID: A project-focused cross-disciplinary graduate training program for data-enabled research in human and machine learning and teaching

DOD, Army W911NF2110317: Cohesive and Robust Human-Bot Cybersecurity Teams

Author contributions:

Conceptualization: VF, LC, TTR

Methodology: VF, LC, TTR

Investigation: VF, LC

Simulation: TR

Visualization: VF, TTR

Supervision: TTR

Writing—original draft: VF, TTR

Writing—review & editing: VF, LC, TTR

Competing interests: The authors declare no competing interests.

Data and materials availability: All data and code for replicating simulations are available at <https://github.com/ttrogers/frigo-chen-rogers>