

2 **Supplementary Information for**
3 **The Effects of Recursive Communication Dynamics on Belief Updating**
4 **Proceedings of the Royal Society B**

5 **Niccolò Pescetelli & Nick Yeung**

6 **Corresponding author: Niccolo Pescetelli.**

7 **E-mail: pescetelli@mpib-berlin.mpg.de**

8 **Manuscript doi: [10.1098/rspb.2020.0025](https://doi.org/10.1098/rspb.2020.0025)**

9 **This PDF file includes:**

- 10 Supplementary text
- 11 Figs. S1 to S23
- 12 Table S1
- 13 References for SI reference citations

Supporting Information Text

Experiment 1 - Supplementary text

Full methods.

Participants Twenty-four dyads ($N = 48$, 37 females, age: $M = 22.52$, $STD = 3.16$) were recruited in pairs for money or course credits, through local announcements and the University of Oxford volunteers platform. The invitation informed potential volunteers to bring a friend of the same gender. This was done to avoid confounds due to gender differences in the use of confidence scales and it represents standard practice in the literature (1, 2). All dyads responded positively to the call, apart for one whose members were gender mixed due to unforeseen circumstances. The study was approved by local ethical committee. All participants gave informed consent before taking part to the study.

Paradigm Participants sat on opposite sides of a desk divided by an occluding screen (Figure 1a), each given a separate LCD monitor, keyboard and mouse. All devices were controlled by the same computer (Dell OptiPlex 9020). All trials consisted of two stages: a private perceptual decision followed by a social exchange. During the private perceptual decision, participants performed a dot-count decision task with confidence ratings (3): Two boxes containing dots randomly arrayed on a 20×20 grid were briefly (160 ms) flashed on each participant's screen to the left and right of a central fixation cross. On each trial, one box contained $200 + d$ dots and the other $200 - d$ dots. Participants had to indicate which box contained most dots. Task difficulty was controlled by changing the d parameter and titrated to each participant independently so to reach an accuracy of around 70% (2-down 1-up procedure (4)). This ensured that independent of their individual sensitivity to the task, both experienced an equal amount of correct and error trials. Notice that, given the double staircase procedure, different dot displays were presented to the two dyad members on each trial, but the box with most dots (i.e., the correct answer) was the same for the two participants on any given trial. Thus, social information coming from the other person carried meaningful information.

Each dyad member indicated their independent response by mouse-click on a semi-continuous post-decision wagering scale (5), ranging from "100% sure LEFT" to "100% sure RIGHT", with the middle level removed to force participants commit to one or other decision. The scale had fifty levels per side. Participants were informed that each level of the scale corresponded to one token, which was awarded if the answer was correct and subtracted from their total score if the answer provided was wrong. Each token was worth £0.01, given to participants as cumulative earnings at the end of the experiment. Post-decision wagering scales are known to be dependent on the pay-off matrix used (6), which produces confounds if participants are too risk-seeking (7). As a control, participants were tested for loss aversion using the coin gamble task (8) and shown to be significantly loss averse. Furthermore, we replicated the results in Experiment 2 and 3 using different confidence scale and instructions, to avoid the limitations affecting the use of post-decision wagering scales (7). Unless stated otherwise, all reported key Experiment 1 effects were replicated in Experiments 2 and 3, as described below.

The member who responded first waited until the second had input their response. As soon as both members confirmed their answer by pressing the spacebar, the social exchange stage started, where each dyad member was informed about their partner's belief. At this point, confidence changes were recorded continuously. In contrast to the standard judge-advisor system paradigm (9), where confidence updates happen in discrete steps, here we recorded confidence judgements as they evolved over time: The mouse x-position along the scale was recorded every 200 ms and each data point so collected was treated as an individual post-decisional bet, contributing to the total amount of tokens participants were supposed to maximise. This was done to incentivize participants to update their cursor position along the scale as soon as their internal confidence changed. Furthermore, participants were explicitly told in the instructions to continuously monitor and update their decision confidence, and the incentive mechanisms was clearly explained. No clicking nor confirmation were required during the social stage to facilitate reliable and continuous tracking of confidence change. This stage expired after five seconds (26 data points). Experiments 2 and 3 shortened this time to four seconds as this was shown to be sufficient to fully capture dyadic interaction. At this point feedback was provided to both members about the tokens earned by each member, and then a new trial began.

Manipulation Our manipulation concerned only the social exchange stage. Two conditions alternated across blocks. In the Static condition, the choice and confidence level selected by each dyad member in the private phase appeared on their partner's scale as a static coloured cursor. Dyad members were at this point asked to continuously monitor and update their confidence by moving their mouse along the scale. In the Dynamic condition, the social exchange part started exactly as in the Static one, with each dyad member's cursor appearing on their partner's scale. However, and for the whole duration of the social part (five seconds), if a member updated their confidence, this would instantly appear also on their partner's scale and vice-versa. This led to a situation where participants were not only informed of their partner's original beliefs, but also how those beliefs changed in real-time as a function of their own updates (Figure 1b, main text).

The experiment began with three blocks of practice of 10 trials each (practice with the perceptual task alone, then including the social exchange stage separately with static or dynamic interaction) followed by 14 experimental blocks of 25 trials each. Each experimental block contained 2 null trials randomly interleaved, which were private decision only trials, included so that participants were motivated to report their confidence accurately also during the private decision. In null trials, earnings were calculated from the confidence expressed during the private phase only. All other trials' earnings were computed instead from the social part. Analyses were performed to assess how social exchange (interactive or static) affected dependent variables of interest: confidence, accuracy and confidence-to-accuracy calibration.

Confidence adjustments. Trial-wise transitions—or the number of times within each trial that confidence changed from time-step t to $t+1$, having remained static on the previous time-step (i.e., from $t-1$ to t)—significantly differed between conditions, but the effect size was small (Static = 1.11, Dynamic = 1.18, $t(47) = 2.95$, $p = .004$, $d = 0.12$) indicating that marginally (but consistently) more updates happened during the Dynamic condition. As the average trial consisted of a single transition in both conditions, analyses on confidence reported below are performed on last confidence points registered on each trial (unless explicitly specified).

A toy model of recursive belief update. Figure 2c in the main manuscript refers to a model of belief update that we explain in this paragraph. Consider an example trial in which a participant (P_{max}) starts off on a confidence level of $C_{max}^{pre} = 0.6(a.u.)$ while their partner (P_{min}) weakly disagrees ($C_{min}^{pre} = -0.4(a.u.)$). Here, the negative sign indicates disagreement. Suppose next that both participants use a simple update strategy, namely summing their own initial confidence with their partner's weighted confidence (here: weight = 0.80). This strategy has been shown elsewhere to be a good approximation of confidence change in joint decisions (10). In a situation where no interaction is allowed, participants can only use their partner's initial belief to get to a new confidence, thus reaching levels of $C_{max}^{post} = 0.28$ and $C_{min}^{post} = 0.08$: As might be expected intuitively, both participants reduce their confidence when learning of their partner's disagreement with their initial decision. However, now consider an interactive scenario where each participant has access to his/her partner's current confidence level at each timestep in the social phase, and uses this information to recursively update his/her initial confidence. Figure 2c in the main text shows how this simple strategy leads to an oscillatory update that stabilises for P_{max} on a higher confidence than initially held. The effect can be explained by the fact that as soon as P_{min} crosses the decision boundary 0, disagreement turns into agreement, thus supporting P_{max} 's initial belief, instead of providing contradictory evidence.

An analysis of vacillations. To test for recursive dynamics in our behavioural data we counted, for each condition, the average number of vacillations in a trial, namely the number of times the direction of the update (i.e., stationary/increase/decrease) changed in the update window. Formally: $V_i = (r_t - r_{t-1}) \neq 0$, with $r_t = \text{sign}(C_t - C_{t-1})$, for each trial i and timestamp t . Given that confidence positions were recorded every 200ms for a 5s update window, we collected 26 timestamps for every trial. Across participants, the average number of vacillations in a trial was significantly more frequent in the Dynamic than Static condition, providing some support for the intuition behind our simple model, although the effect was a small one (Static = 2.41 ± 1.26 , Dynamic = 2.55 ± 1.27 , $t(47) = 2.62$, $p = .01$, $d = 0.11$). Accordingly, participants showed a significantly higher rate of irrational increases in the Dynamic condition compared to a Static condition when they were the more confident of the two partners on the trial (0.012 vs. 0.008 of disagreement trials, $t(47) = 3.29$, $p = .001$, $d = .21$), and not when less confident (0.007 vs. 0.007 of disagreement trials, $p > .8$).

Comparing confidence updates with Bayesian aggregation. We compared participants' confidence changes in interaction with a normative Bayesian strategy. For this analysis, we treated confidence ratings as subjectively estimated probabilities that a particular decision is correct (11–13), while applying a linear transformation to prevent values of 0 and 1 and thus avoid impossible solutions: range of 0.01 = "100% Sure LEFT" to 0.99 = "100% Sure RIGHT". The probabilities p_s and p_p so obtained—representing dyad members' independent priors on RIGHT being correct (cf. 14)—can now be integrated into the normative posterior:

$$post_{norm} = \frac{p_s p_p}{p_s p_p + \neg p_s \neg p_p} \quad [1]$$

where $\neg p_s$ and $\neg p_p$ are $1 - p_s$ and $1 - p_p$ respectively, representing the subjective prior probability on LEFT. The resulting posterior confidence represents the post-exchange confidence held by a normative belief aggregation method on RIGHT being the correct answer. The normative model reproduces some patterns of belief change observed in empirical dyads (10).

Figure S3a plots the difference, Res , between the normatively prescribed confidence change, δ_{norm} , and the empirically observed value, δ_{emp} , as a function of consensus and interaction condition. A 2x2 ANOVA on participants' mean Res values revealed a significant effect of consensus ($F(1, 47) = 68.37$, $p < .001$, $\eta_G^2 = 0.46$) and interaction condition ($F(1, 47) = 4.97$, $p = .03$, $\eta_G^2 = .002$) but no significant interaction ($F < 1$). Participants were more confident than prescribed by Bayesian updating in disagreement for both conditions ($t(47) > 8.34$, $p < .001$), and less confident in agreement, significantly so in Static blocks ($t(47) = -3.06$, $p = .003$) but not in Dynamic blocks ($t = -1.56$, $p > .1$). These findings indicate that participants systematically underweighted their partner's advice, but did so particularly strongly when they disagreed, thus replicating previously observed decision biases of egocentric discounting and confirmation bias that affect the perception of social information (15–17). The main effect of interaction condition mirrored the overall effect of dynamic interaction leading to increased confidence, thus increasing the discrepancy from optimal updating in case of disagreement, but reducing it in case of agreement.

Bayes theorem can also be used to infer the participants' perception of the social information. Equation 1 is used by the optimal observer to infer the predicted posterior confidence given a prior confidence level p_s and a partner's belief p_p . However, by solving the equation for p_p (i.e., the likelihood term), we can infer the *perceived* partner's confidence \hat{p}_p , given the degree to which the participant updated their confidence from their initial judgment (the prior, p_s) to a final decision following social interaction (the posterior, $post_s$):

$$\hat{p}_p = \frac{post_s(p_s - 1)}{2p_s post_s - p_s - post_s}; \quad [2]$$

In this way we can quantify the weight that participants assigned to their partner's judgement across trials, contrasted with the weight (i.e., probability correct) that the partner themselves conveyed in their confidence judgments. Comparing the distribution of these values across trials reveals how objective social information (i.e., partner's stated support for the participant's belief) becomes distorted when perceived and acted on by the participant (i.e., actual use of partner's social information). As shown in Figure S3b, whereas partners' stated confidence ratings were fairly evenly distributed in terms of conveying social information ranging from maximal disagreement (0 on the x-axis) to maximal agreement (1), the weight that participants assigned to their partners views showed almost categorical behaviour, with separate peaks at $\hat{p}_p \sim 0.5$ (i.e., advice treated as carrying little or no information) and $\hat{p}_p \sim 1.0$ (i.e., advice treated as objectively correct) (18).

Perceptual decision task performance. A tenet in the judgment aggregation literature is that social interaction hampers collective wisdom by breaking the independence of the individual judgments (19). The traditional interpretation of the wisdom of crowds (20, 21), named here the "noise cancelling hypothesis", explains the accuracy improvement observed in belief aggregates as a statistical phenomenon where noise reduces after independent samples (here the private initial beliefs) are averaged together. This hypothesis predicts that breaking the independence between measures should have negative effects on accuracy, as errors become correlated instead of averaging out. According to this view, in the present study we should observe that (1) simple exposure to another person's belief negatively affects performance; (2) the effect of social exposure is even more damaging on performance in the Dynamic condition, as this condition affects the independence of the individual estimates more than the Static one.

A 2-way ANOVA on accuracy with factors condition (Static vs. Dynamic) and decision stage (pre- vs. post-exchange) showed a significant effect of decision stage ($F(1, 47) = 47.00, p < .001, \eta_G^2 = .16$) but no significant effect of condition ($F < 1$) nor a significant interaction ($F < 1$). Social information had a beneficial effect on average accuracy (pre-exchange accuracy = 0.72, post-exchange accuracy = 0.75). Similar null effects of interaction condition when we measured accuracy improvement as confidence changes toward or away the correct end of the scale: $\delta_{accG} = (C_{post} - C_{pre})^{Acc} * (C_{pre} - C_{post})^{1-Acc}$, where $Acc \in \{0, 1\}$. A corresponding pattern was apparent in final measure of performance we considered: the calibration of confidence relatively to objective accuracy, defined here as the type II A_{ROC} (22). The same two-way ANOVA used for choice accuracy was run on type II A_{ROC} . Results show a significant effect of decision stage ($F(1, 47) = 89.58, p < .001, \eta_G^2 = .25$), indicating calibration improvement from pre- to post-exchange phase (0.60 vs. 0.66), but no effect of condition nor interaction between the two ($F_s < 1$). Taken together, these findings indicate that, contrary to (19), exposure to another person's belief did not reduce accuracy. Furthermore, Dynamic interaction did not reduce accuracy improvement compared to Static, indicating that increased dependence between confidence updates (as indicated in the analyses of Figure 2d of the main text) had no significant damaging effect on accuracy (or accuracy improvement) either.

Reaction Times. Another indirect cue that participants could have used in the Dynamic condition, but not in the Static condition, to inform their update was the speed of their partner's confidence update. On one side, movement speed is known to be associated with confidence, with longer reaction times corresponding to lower confidence levels (23). On the other hand, in the context of observed cursor movements during the interaction phase after an initial decision, it is more plausible that *resistance* to move one's cursor is interpreted as a signal of higher confidence (stubbornness)—i.e., although participants should be faster to *express* beliefs in which they more confident, they should be slower to *change* highly confident beliefs as a consequence of interaction. Due to the fixed time allocated to the social part, a direct measure of reaction times was not available. Therefore, to test for variations in movement speed across conditions, we fitted a sigmoid curve to each trial's confidence time series, namely each member's cursor's positions over the five seconds each social exchange lasted. Slope ϕ and offset λ along the time axis were free parameters to be estimated on each trial. The offset parameter λ was used as a proxy for reaction times and entered into an ANOVA with factors trial-dominance (more confident vs. less confident), condition and consensus (agreement vs. disagreement). Results showed a significant effect for all main effects ($F(1, 47) > 7.25, p < .009$) and a significant interaction between trial-dominance and consensus ($F(1, 47) = 4.98, p = .03$). Pairwise comparisons showed that members who started off less confident than their partners were slower in updating their cursor ($t(47) = 7.91, p < .001$), making the hypothesis that cursor's stickiness was (or could be) interpreted as a cue for confidence implausible. Furthermore, estimated reaction times λ were slower in the Dynamic than Static condition ($t(47) = 2.69, p = .009$) and faster in agreement than in disagreement ($t(47) = -7.71, p < .001$), indicating that longer reaction times were associated with more uncertainty (cf. 23).

Linear-mixed effects model. A linear mixed-effects model on trial-by-trial absolute confidence change was run (a) to estimate the relative weight of different predictors on trial-level absolute confidence change; (b) to take into account the nested structure of the data. A full model (main effects and all interaction terms) was specified with participant's absolute confidence change ($|\delta_C^s|$) as dependent variable (REML fitting method). Predictors included self initial confidence C_{pre}^s as well as partner's C_{pre}^p , condition (Static vs. Dynamic), the relationship between self and partner's initial views (agreement vs. disagreement), absolute confidence change observed in the partner $|\delta_C^p|$, and fitted partner's reaction time to update rt_2 . Continuous variables were normalised within participants; condition and consensus were declared categorical predictors and entered in the model using reference coding. Trial-level data points were grouped into participants and participants into dyads. Random intercepts were declared for participants and dyads and random slopes were declared for participants (but not dyads) for each main effect. Non-significant predictors were removed and a new model so obtained was run and compared to the previous model using a likelihood ratio test (*compare* function in MATLAB). The procedure was iterated until all predictors and random effects were significant. Resulting significant coefficients for fixed effects are listed in Table S1.

Agreement showed on average smaller updates than disagreement ($\beta = -0.27, SE = 0.05, p < .001$). The interaction term between condition and consensus ($\beta = 0.11, SE = 0.02, p < .001$) confirmed that, compared to Static condition, agreement trials showed larger updates in Dynamic interaction. Importantly, this interaction was positively modulated by partner's absolute confidence change ($\beta = 0.28, SE = 0.03, p < .001$), suggesting that the more a participant's partner was willing to change their initial confidence the greater the participant's changes were in agreement trials, compared to a disagreement baseline. This evidence confirmed our hypothesis that participants made use of non-independent information. The same condition by consensus interaction was negatively modulated by participant's initial confidence ($\beta = -0.05, SE = 0.01, p = .008$) suggesting that the stronger the confidence initially held, the less the effect of Dynamic interaction differed from a Static baseline. The opposite relation was true for the participant's partner's initial confidence ($\beta = 0.05, SE = 0.02, p = .008$) indicating that in agreement, the greater a partner's initial confidence, the more a participant's confidence increases in Dynamic compared to Static condition. Partner's fitted reaction times did not interact with condition nor with condition and agreement, suggesting that, during Dynamic interaction, observed movement speed did not factor in subjective updates in obvious ways (23). Significant interaction terms between partner's fitted reaction times and condition were only four-ways interactions, which are typically difficult to interpret. Besides, the coefficients associated with fitted partner's reaction times were smaller than the ones associated with partner's update magnitude, indicating that the latter was a stronger predictor of personal updates in interaction.

Egocentric and confirmation biases. Self-serving cognitive distortions of social information can arise from a different weighting of self and others' beliefs, a phenomenon known in the judge-advisor system literature as egocentric bias (9, 17, 24). To quantify the extent of egocentric bias, we fitted a linear model on perceived evidence with only predictor partner's stated evidence (from 0=confident disagree to 1=confident agree). Regressions were fitted for each participant, dominance type and for agreement and disagreement separately, and coefficients were used as an empirical estimate of partner's belief weighting factor. Regression lines were anchored at (0.5,0.5), so to obtain a bilinear transfer function from stated to perceived evidence. This extra degree of freedom allowed different weighting factors for agreement and disagreement trials, thus allowing to account for potential asymmetries and confirmation biases (15). Fitted coefficients α (i.e., slope in agreement trials) and β (i.e., slope in disagreement trials) represent the weight that participants give to their partners' stated belief, in agreement and disagreement trials respectively. A 3-way ANOVA on weights with factors trial-dominance, condition and consensus showed an effect of consensus ($F(1, 47) = 10.97, p = .001, \eta_G^2 = 0.035$). This effect indicates that contradictory social information (i.e., disagreement) was discounted more than supporting evidence, namely agreement (0.46 vs. 0.62), thus indicating the presence of a confirmation bias (15). No significant main effects of condition or trial-dominance were found ($F < 1$) nor a significant interaction between the two ($F < 1$). Significant interactions between consensus and condition ($F(1, 47) = 10.05, p = .002, \eta_G^2 = 0.003$) and consensus and trial-dominance ($F(1, 47) = 19.79, p < .001, \eta_G^2 = 0.03$) were found, indicating that Dynamic condition tended to increase discounting in disagreement and decrease it in agreement. The result can partially be explained by the increased agreement effect and decreased disagreement effect observed in Dynamic interaction.

Loss aversion. The use of post-decision wagering as a measure of confidence can be distorted by risk or loss seeking behavior (7, 10). To this end we tested all participants after the experiment using the coin flipping gambling task (8). Participants were on average highly loss averse in both conditions ($t > 4.2, p < .001, d > .93$), thus making less likely that participants' rated confidence was inflated by risk-seeking behaviour.

Experiment 2

Here we tested an alternative hypothesis to the results found in the main text, namely that participants in the Dynamic condition simply tended to forget their initial confidence judgment, and were instead updating the *current* confidence of their partner with their own *current* confidence. Modifying the toy simulation described above easily shows that this strategy quickly leads confidence of both participants to escalate towards the maximum confidence boundary on the side of the most confident initial belief.

To test whether the effect of interaction found in Experiment 1 was due to participants' failure to remember their own initial confidence, a third experimental condition was created and compared to the previous two. In this new condition (called *Dynamic_{self}*), participants were given a static reminder of their own pre-social information confidence during the social phase. This reminder was presented on the confidence scale along with the standard personal and partner's cursors typically presented during the Dynamic condition. If the effects of interaction are only due to memory failures, then the presence of a reminder should make those same effects disappear. Failure to reduce the interaction effects should be taken as evidence that differences between dynamic and static conditions are not due to forgetfulness.

A worry from Experiment 1 was that people often used extreme values when rating their initial confidence. We thus introduced different instructions regarding the input of confidence ratings, incentivising confidence calibration over confidence magnitude. This gave us the opportunity to assess the robustness of key Experiment 1 effects with a different confidence scale.

Methods.

Participants Twenty-four dyads (14 female dyads, 1 mixed gender dyad) were tested. Mean age was 23.16 ± 3.42 . Participants were recruited online using the University volunteers platform and local advertisement websites. All participants gave informed consent before starting the experiment. The study was approved by local ethical committee.

Paradigm The experiment comprised of 432 experimental trials divided in 18 blocks and 20 practice trials divided in 4 blocks. Practice blocks were designed to practice with the first-order task, the static condition, the dynamic condition, the dynamic plus reminder condition respectively. The methods (Figure S10) were very similar to those used in Experiment 1, in regard of the dot-count task, trial stages and input modalities, with the following key differences. First, three conditions were defined by manipulating the access participants had to their own and their partner's information during the social phase: the two conditions already presented in Experiment 1 and a reminder condition. Conditions were varied within-participants across blocks (i.e. six blocks per condition). Participants experienced six identical modules, each comprising the three different conditions into three separate blocks. The order of the three conditions within a module was randomised across dyads. Second, the social part window was reduced to 4 seconds (21 data points), given that most updates in Experiment 1 occurred within 2 seconds of the social part. Third, it was decided to change the incentive system used for Experiment 1 and the instructions given to participants to use the confidence scale. This modification was motivated by two main reasons. The first reason was to make participants' confidence distributions less extreme and more uniform across the scale. Although in Experiment 1 all participants showed some evidence of loss aversion, confidence judgments were skewed toward the high end of the scale, creating potential issues in detecting small confidence changes in this direction (i.e., confidence increases) due to ceiling effects. The second reason was to check whether the effects found in the previous experiment were robust to changes in the incentive system and thus in the use of the confidence scale. Failing to reproduce Experiment 1 results when changing the incentive system would be a strong indication that they were (at least partially) dependent on the specific instructions participants received. Details about how the new incentive scheme worked and about the instructions given to participants are described below.

Manipulation Three conditions were defined that affected only the social part of the trial. A Static and a Dynamic conditions were defined as in Experiment 1, which allowed us to see if those effects replicated. A Dynamic plus self-reminder condition (Dynamic_{self}) was constructed so that participants were shown a reminder of their own initial private confidence during interaction. The reminder was presented as a static gray shaded cursor.

Notice that in all conditions the social part started exactly with the same initial configuration of objects on the screen and cursors were presented in the same position as they were left at the end of the private part. Any difference among conditions must then be attributed to the specific communication channels that each condition entails, assuming equal initial conditions of the dyad state. Conditions alternated regularly over blocks (six repetitions each) and their order was shuffled across participants.

Incentive scheme In the current experiment participants were informed that their final reward would be inversely proportional to the average absolute deviation of their accuracy from the calibration line. The calibration line was defined by the line $y = x$, i.e. where confidence expressed in percentage points is identical to the probability of a correct response. Instructions stated: "We will average all trials when you were 60% confident and see if you were indeed 60% accurate. Then we'll see if you were 70% accurate on trials where you said you were 70% confident and so on. The higher the discrepancy the less you will get.". Importantly participants were told that during the social part this measure was computed on a moment-by-moment basis and that the best strategy to maximise their gains was thus to continuously update their confidence cursor based on their internal sense of confidence.

For this calculation, at the end of each block the confidence distribution of each participant was divided into 5 bins and the weighted average absolute distance between bin accuracy and bin center was taken as a calibration error:

$$Err = \frac{\sum_{b=1}^5 |Acc_b - Conf_b| * N_b}{\sum_{b=1}^5 N_b} \quad [3]$$

where N_b is the total number of data points recorded in each bin. Err was computed for pre-social information and post-social information separately and the two were averaged together so that an equal weight was given to private and social parts. Importantly the formulation above computes the calibration error on each data point collected - i.e. 1 for pre-social information and 21 for post-social information decisions. This ensures that the error during the social part is a weighted average among bins based on the time spent in each one.

Results.

Continuous update During the social part of each trial, the x-position of the cursor along the confidence scale was recorded every 200 ms, giving 21 confidence data points over the course of 4 seconds. The absolute difference between a data point and the previous one can be used as a measure of the stability of the confidence updates over time, with smaller numbers indicating that participant's updates have stabilised. This update stability measure is shown in Figure S11 for the three different conditions separately. It can be seen that in all conditions the larger confidence update occurred around one second from the start of the social part. Both dynamic conditions showed larger updates on average around this period, followed by longer times to reach an equilibrium as suggested by the larger right tail.

Asymmetry in confidence increases Figure S12 plots confidence change distributions, divided by consensus and averaged across participants. It can be clearly seen that distributions peak at zero, suggesting that most frequent confidence update was to not update. Right-tails in disagreement and left-tails in agreement represent irrational confidence changes. A two-way repeated

measures ANOVA on the probability of an irrational change (corrected for total number of agreement and disagreement trials and trembling hand issues) showed a significant effect of consensus ($F(1, 47) = 16.98, p < .001, \eta_G^2 = .12$) but not of condition ($F < 1$) and no significant interaction between the two ($F(2, 94) = 1.25, p = .28, \eta_G^2 = .002$), suggesting that irrational increases were more frequent than irrational decreases ($M \pm \text{STD}$: irrational increases = 0.0166 ± 0.022 vs. irrational decreases = 0.0031 ± 0.004), but no consistent differences were found among conditions.

The results partly replicate what found in Experiment 1, suggesting that irrational changes are more frequent after disagreement than after agreement. Experiment 2 does not however replicate the finding that irrational increases were more frequent in the Dynamic than the Static condition, suggesting that perhaps this result was an effect of a different use of the confidence scale.

Influence in belief space Visualising confidence changes along the belief space can better represents participants' behaviour during the update window. Median confidence changes δ_C were plotted in belief space to understand which subsets of trials (i.e. which initial conditions) showed larger confidence changes and which ones showed the strongest difference between experimental conditions. Confidence changes were plotted separately for trial-dominant and trial-dominated trials. Figure S13 shows the belief spaces so obtained. Two major areas of interest were identified in Experiment 1, one corresponding to weak agreement (participants are both unsure but happen to agree) and the other corresponding to unbalanced disagreement (one participant is very confident while the other weakly disagrees). In Experiment 2, similar areas of interest emerged. In both dominant and dominated trials, participants in dynamic conditions showed larger confidence increases compared to a Static baseline after weak agreement. The magnitude of the increase in these areas, indicates that in interaction participants converged on high confidence agreement. A real-time animation of the density distribution of dyad states during the 4-second update window, as well as an animation of the contrast between conditions, can be found at osf.io/7b6py. The animated contrast plot shows that, although in the two conditions dyad states were similarly distributed along the belief surface at the beginning of the update, more trials in the dynamic conditions than in the Static one gravitated towards point (50,50). The analyses above allow us to understand which subsets of trials are similar across conditions and which ones are not, making it easier to determine what effects the manipulation has on behaviour. They can inform subsequent analyses by restricting the trials of interest to trials that are likely to generate the effects observed.

Coupling of confidence changes in interaction Experiment 1 showed that interaction produced positive correlation in dyad members confidence changes under agreement and negative under disagreement. We thus tested whether the results replicated here. Figure S14 shows the average Pearson's r coefficient, divided by condition and consensus. Coefficients were entered into an ANOVA across dyads with factors condition and consensus. Results show that both condition ($F(2, 44) = 16.35, p < .001, \eta_G^2 = .12$) and consensus ($F(1, 22) = 13.16, p = .001, \eta_G^2 = .09$) had a significant effect on the correlation observed. The interaction between the two terms was also significant ($F(2, 44) = 27.01, p < .001, \eta_G^2 = .09$). No correlation was found in any of the three conditions in disagreement ($t(23) < .8, p > .4$). On the contrary in agreement both dynamic conditions showed positive correlation coefficients ($t(23) > 4.5, p < .001$) while coefficients in the Static condition were not significantly different from zero ($p > .1$). The results partly replicate results observed in Experiment 1. Similarly to Experiment 1, Experiment 2 indicated that confidence changes of members of the same dyad remained independent from each other in the Static condition and interaction introduced positive correlations between confidence changes in agreement trials, with no difference found between dynamic conditions ($p > .2$). The negative correlation found in disagreement trials in Experiment 1 between same dyad members was however not replicated in Experiment 2.

The same mixed-effects model used in Experiment 1 was applied to Experiment 2 data and largely replicated the main effects found there. The interaction terms between both dynamic conditions and consensus were significantly above zero (Dynamic: $\beta = 0.17, SE = 0.03, p < .001$; Dynamic_{self}: $\beta = 0.20, SE = 0.03, p < .001$), indicating that during interaction agreement led to greater confidence updates compared to a Static reference. Importantly, both terms were positively modulated by partner's absolute confidence change (Dynamic: $\beta = 0.37, SE = 0.03, p < .001$; Dynamic_{self}: $\beta = 0.26, SE = 0.03, p < .001$), replicating the finding that in interaction participants tended to make use of their partner's confidence changes to inform their own confidence updates.

Performance analysis Experiment 1 had found a significant benefit of social exchange, but no significant differences between conditions according to different measures of performance, including choice accuracy, graded accuracy and confidence calibration. In Experiment 2, a two-way ANOVA on choice accuracy with factors condition and decision stage (pre-social information vs. post-social information), showed a significant effect of decision stage ($F(1, 47) = 101.66, p < .001, \eta_G^2 = .17$), replicating the finding that choice accuracy significantly improved from pre- to post-social information phase (M : 0.71 vs. 0.74), but no significant difference of condition and no significant interaction ($F < 1$). A corresponding analysis of confidence calibration (measured as type II A_{ROC}) revealed a significant main effect of decision stage ($F(1, 47) = 94.08, p < .001, \eta_G^2 = .26$), indicating that calibration significantly improved after exchanging social information (0.56 vs. 0.62), but no effect of condition ($F < 1$), indicating that neither interaction or the presence of an anchor negatively affected calibration. A significant interaction term was also found ($F(2, 94) = 3.53, p = .03, \eta_G^2 = .02$), indicating differences in improvement across conditions. Pairwise comparisons showed that calibration improved significantly more in the Dynamic compared to Static condition ($t(47) = 2.53, p = .01, d = 0.38$). No significant difference in calibration improvement was found between dynamic conditions nor between Dynamic_{self} and Static condition ($p > .1$).

Overall, decision performance improved after social exchange, and increased dependency between judgments through interaction did not hamper improvement, but instead, if anything, fostered it.

Experiment Discussion. Experiment 2 was run with two main aims in mind: (1) reproduce results from Experiment 1 when using different incentive schemes; (2) test whether those same effects were produced by failure in remembering one's own initial judgment. The results described above replicate the key findings of Experiment 1 and reproduce the overall pattern of differences between dynamic and static conditions. In particular, interaction seems to significantly increment confidence increases observed from pre- to post-social information phase after agreement, regardless of the presence of a confidence reminder. Similarly to what observed in Experiment 1, the decrease in confidence observed after a disagreeing belief is reduced in Experiment 2 in both dynamic conditions, although not significantly in the Dynamic condition. Differences in confidence change among conditions were once again driven by weak agreement trials and unbalanced disagreement trials. The findings also replicated the positive correlation emerging during real-time interaction between dyad members' confidence changes. Contrarily to Experiment 1 however, no negative correlation was found in disagreement trials, suggesting that during these trials dyad members' updates remained independent from each other irrespective of condition.

Accuracy improvements from pre- to post-social information were all positive and significantly different from zero. Conditions did not differ from each other suggesting that, notwithstanding the reduced independence of participants' judgments, performance improvements were robust. Dynamic interaction favoured greater improvements compared to the static baseline in terms of confidence calibration.

The normative framework described for Experiment 1 was also applied here to show that people adopt qualitatively different strategies in agreement and disagreement, with greater weights put on partner's beliefs in agreement trials. Confirming results found in Experiment 1, Experiment 2 provided further evidence that social information perception differs from the objective social information received. In particular, participants tend to categorise received social information into strong evidence in favour of their initial belief or null evidence.

Overall the experiment showed that most of the effects observed after the interaction manipulation are robust to changes in the use of the confidence scale, with few differences found in disagreement trials. Importantly, the introduction of a confidence reminder little affected the Dynamic condition, suggesting that differences between dynamic and static conditions were not simply due to memory failures of one's own initial confidence. The experiment thus offered a proof that dynamic and static paradigms differ not only in terms of low-level characteristics but instead differences are intrinsic to the dynamics of how information is shared and manipulated across individuals.

Experiment 3 was carried out to test whether differences between dynamic and static conditions were instead due to memory failures of one's *partner's* initial confidence.

Experiment 3

Experiment 1 showed differences in behaviour emerging from the manipulation of how social partners can exchange their independent pieces of information. The independence of confidence updates was affected by the presence of real-time interaction, suggesting that participants updated their confidence not only using the initial confidence of their partner but also their partner's updates. Experiment 2 ruled out a simple alternative explanation in terms of participants forgetting their own initial confidence. However, another explanation for confidence escalation is that participants forgot *their partner's* initial confidence and were thus incentivised to use, when available, their partner's current position as a proxy for it.

To test whether this explanation could explain the effects found in the Dynamic condition a new condition was set out and compared with the Dynamic interaction and the Dynamic_{self} conditions. In this condition, called Dynamic_{other}, the Dynamic condition is enhanced by the presence of a static reminder about one's partner's initial confidence that remains on screen for the whole duration of the social exchange. If the memory explanation is correct we expect the effects of interaction to diminish when a reminder is presented. Failure in finding such results can be taken as evidence that the effects of interaction are not due to failures in memory. In contrast to previous experiments, Experiment 3 did not include a condition with static information sharing, since the primary interest here was to replicate effects observed with dynamic interactions, and assess their sensitivity to reminders of self and partner's initial decisions.

Methods.

Participants 24 dyads (17 female dyads, 1 mixed gender) were recruited using University volunteers recruitment platform and local advertisement websites. Dyads were recruited by asking an interested volunteer to bring along a friend of the same gender. Participants (age=20.66±2.76) signed a consent form prior the beginning of the experiment. The study received ethical approval from the University ethical committee.

Paradigm Participants performed 18 blocks of 24 trials each. Perceptual task, trial sequence and response modality were kept equal to previous experiments S17. The social window was kept to 4 seconds as in Experiment 2. Given that Experiment 2 was successful in making participants less extreme in their confidence ratings, the same incentive scheme was used here. The experiment started with 4 practice blocks of 5 trials each, corresponding to practice with the perceptual task and with each condition separately. Performance was titrated to 70.7% accuracy using a 2-down 1-up procedure.

Manipulation Three experimental conditions were implemented and alternated across blocks in six identical modules of three blocks each. The order of the three conditions within a module was randomly shuffled across dyads but remained identical within the same dyad. The first two conditions were the Dynamic and Dynamic_{self} conditions, already described in Experiments 5. A third new condition, named Dynamic_{other}, was implemented by adding to the Dynamic condition a static cursor reminding

the participant of their partner's initial confidence level. A colour code was used so to avoid confusion on what each cursor meant. Participant-related cursors were represented in white (active cursor) and grey (static reminder). Partner-related cursors were represented in bright colour (active cursor) and dark colour (static reminder).

Results.

Continuous update Similarly to Experiment 1 and 2, Figure S18 shows that a sharp confidence update occurred in all conditions around the first second of the social window and settled into an equilibrium by the end of it. No significant differences were observed between conditions, indicating that the time used by dyads to reach their final decision was not reliably affected by the presence of a confidence reminder.

Asymmetry in confidence increases The confidence change distributions of Experiment 3 are shown in Figure S19 as root density plots. As in previous experiments the most common confidence change was zero, suggesting that very often participants decided not to act upon social information. To test for asymmetries in irrational confidence changes, a two-way repeated measures ANOVA on the probability of an irrational confidence change was run. Results showed only a marginal effect of consensus ($F(1, 47) = 3.08, p = .08, \eta_G^2 = .02$) and no effect of condition ($F(2, 94) = 1.03, p = .35, \eta_G^2 = .001$) nor significant interaction ($F < 1$). Experiment 3 replicates the finding found in the previous two experiments that irrational changes were more frequent in disagreement than in agreement trials ($M \pm \text{STD}$: irrational increases = 0.0124 ± 0.018 ; irrational decreases = 0.0072 ± 0.008). The lack of consistent differences among conditions suggests that the presence of a confidence anchor did not affect the presence of irrational confidence changes in the baseline Dynamic condition.

Influence in belief space Figure S20 shows that the pattern of results is very similar to those observed in the previous two experiments. No difference in weak agreement areas were found, indicating that the presence of confidence anchors did not alter median confidence change in these trials. Unbalanced disagreement (points y, in the main text) showed no difference among conditions for dominant trials but positive differences for dominated trials. The latter finding suggests that in these trials, dominated members seemed to be more swayed by dominant beliefs in both anchor conditions compared to baseline Dynamic condition, but dominant ones were not.

Coupling of confidence changes in interaction Pearson's correlation coefficients between confidence change magnitudes were compared across conditions and divided by agreement to test whether changes in one participant were correlated with changes in the other (Figure S21). Results of a repeated measure ANOVA on Pearson's coefficients showed that a significant effect of consensus was found ($F(1, 23) = 39.74, p < .001, \eta_G^2 = .24$) but not of condition ($F < 1$) and no significant interaction between the two ($F < 1$). Contrary to Experiment 2 but similarly to Experiment 1, in all conditions confidence change magnitudes in disagreement trials were marginally or significantly below zero (Dynamic: $t(23) = -1.96, p = .06, d = -0.40$; Dynamic_{self}: $t(23) = -1.87, p = .07, d = -0.38$; Dynamic_{other}: $t(23) = -3.18, p = .004, d = -0.64$), indicating that interaction produced an inverse coupling also in disagreement, with little effect of reminders. Similarly, in agreement trials, dyad members' confidence changes were positively correlated as indicated by the significantly positive correlation coefficients (Dynamic: $t(23) = 3.40, p = .002, d = 0.69$; Dynamic_{self}: $t(23) = 3.00, p = .006, d = 0.61$; Dynamic_{other}: $t(23) = 2.97, p = .006, d = 0.60$). In conclusion, irrespective of reminder presence, interaction coupled together partners' confidence changes: greater confidence changes in one dyad member produced greater partner's confidence changes in agreement but lower partner's confidence changes in disagreement.

The mixed-effects linear regression described in Experiments 1 and 2, was run here to check whether the mediating role of partner's confidence change on subjective confidence changes differed across conditions. Consensus positively interacted with partner's absolute confidence change ($\beta = 0.4318, SE = 0.0244, p < .001$), suggesting that the larger was a partner's update during interaction, the more participants tended to shift their confidence in agreement and the less they tended to shift in disagreement. Importantly, the effect was not modulated by either of the anchor conditions (Dynamic_{self}: $\beta = 0.0230, SE = 0.0338, p = .49$; Dynamic_{other}: $\beta = -0.0113, SE = 0.0339, p < .7$) suggesting that the introduction of confidence reminder did not affect the baseline Dynamic condition in the extent to which a partner's updates affected each subject's own updates.

Performance analysis Both Experiments 1 and 2 showed that performance improved after social exchange and that interaction did not negatively affect the size of the improvement. Results were replicated in Experiment 3. A two-way ANOVA on accuracy with factors condition and decision stage (pre-social information vs. post-social information) showed a significant effect of decision stage ($F(1, 47) = 103.96, p < .001, \eta_G^2 = .19$), indicating accuracy improvement due to social information exchange (M : 0.71 vs. 0.74). Importantly no effect of condition nor interaction were found (both $F < 1$), confirming that different conditions did not affect average accuracy or average accuracy improvement. A corresponding analysis of confidence calibration (type II A_{ROC}) revealed a significant main effect of decision stage ($F(1, 47) = 112.49, p < .001, \eta_G^2 = .23$), indicating that calibration improved thanks to social information exchange ($M = 0.57$ vs. 0.62), but not of condition ($F < 1$). A marginally significant interaction between the two terms was also found ($F(2, 94) = 2.95, p = .05, \eta_G^2 = .007$). Pairwise comparisons showed that both the Dynamic_{self} ($t(47) = 2.23, p = .03$) and the Dynamic_{other} ($t(47) = 1.96, p = .05$) conditions produced significant greater calibration improvement over the Dynamic baseline. The results suggest that, although not having any effect on accuracy, the presence of a confidence reminder (either own or partner's) helped participants to have a more accurate metacognitive evaluation, likely because of an increased access to independent estimates.

Experiment Discussion. Experiment 3 replicated all the key findings observed in the previous two experiments. Results show that the introduction of confidence reminders had moderate effects compared to the baseline Dynamic condition. The presence of the other person's confidence reminder made participants decrease their confidence more in disagreement, thus ending on lower absolute confidence levels. Reminders did not seem to affect the independence of the confidence updates over and above what already observed in the Dynamic condition. They did not produce differences in choice accuracy or accuracy improvements either. Only marginal differences in calibration improvements were found among conditions.

Comparison with a Bayesian belief integration strategy confirmed the observations made in the previous two experiments, suggesting that participants discounted social information received from partner. Participants tended to differently treat agreeing and disagreeing evidence and asymmetrically discount the two. Furthermore, greater social information discounting was operated by participants holding the trial-dominated belief, probably due to a general tendency to discount/ignore social information that deviated from Bayes particularly in dominated trials.

Although some effects were observed by the introduction of the other member's confidence reminder, the experiment provided little evidence that the results observed in the Dynamic condition reflected a memory failure in remembering initial beliefs (own or other's). The pattern of results observed in the Dynamic condition was nearly unaltered, suggesting that even in the presence of a constant reminder anchoring participants to their initially expressed views, phenomena of confidence escalation and updates coupling were observed. Thus it seems that confidence escalation and the correlations emerging in interaction between updates of members of a same dyad cannot be explained away by simple mechanisms specific to our paradigm. The results are so far in agreement with an explanation in terms of interaction modifying the dynamics of information exchange between two decision makers. Dynamic interaction creates a situation where both participants can not only use the independent belief of their partner to inform their post-decisional judgments but also how their partners react to the participants' belief. When interaction was allowed, seeing larger updates in their partners made participants' confidence change size increase in agreement and decrease in disagreement. The results add to a large body of evidence suggesting that confidence judgments are not only the product of a careful evaluation of decision-relevant variables, but often include several contingent cues that are not decision-relevant but flow into creating a unitary internal sense of confidence (25–27). Interestingly, interaction decreased the independence of the two members' judgments in all three experiments using the current paradigm. Contrary to a common interpretation of Wisdom-of-Crowds phenomena, in terms of noise cancellation through averaging of independent measures (19, 20), it was repeatedly shown that increased dependence did not significantly affect accuracy nor accuracy improvement. This suggests that when people are allowed to share their confidence judgments instead of their choice preferences only, individual and dyadic performance can be robust to failures (28, 29).

References

- Buchan NR, Croson RT, Solnick S (2008) Trust and gender: An examination of behavior and beliefs in the Investment Game. *Journal of Economic Behavior & Organization* 68(3-4):466–476.
- Mahmoodi A, et al. (2015) Equality bias impairs collective decision-making across cultures. *Proceedings of the National Academy of Sciences* p. 201421692.
- Boldt A, Yeung N (2015) Shared Neural Markers of Decision Confidence and Error Detection. *Journal of Neuroscience* 35(8):3478–3484.
- Treutwein B (1995) Adaptive Psychophysical Procedures. *Vision Research* 35(17):2503–2522.
- Persaud N, McLeod P, Cowey A (2007) Post-decision wagering objectively measures awareness. *Nature neuroscience* 10(2):257–61.
- Clifford CWG, Arabzadeh E, Harris JA (2008) Getting technical about awareness. *Trends in cognitive sciences* 12(2):54–8.
- Fleming SM, Dolan RJ (2010) Effects of loss aversion on post-decision wagering: implications for measures of awareness. *Consciousness and cognition* 19(1):352–63.
- De Martino B, Camerer CF, Adolphs R (2010) Amygdala damage eliminates monetary loss aversion. *Proceedings of the National Academy of Sciences of the United States of America* 107(8):3788–92.
- Bonaccio S, Dalal RS (2006) Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes* 101(2):127–151.
- Pescetelli N, Rees G, Bahrami B (2016) The perceptual and social components of metacognition. *Journal of Experimental Psychology: General* 145(8):949–965.
- Pouget A, Beck JM, Ma WJ, Latham PE (2013) Probabilistic brains: knowns and unknowns. *Nature neuroscience* 16(9).
- Meyniel F, Schlunegger D, Dehaene S (2015) The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLOS Computational Biology* 11(6):e1004305.
- Ma WJ, Jazayeri M (2014) Neural Coding of Uncertainty and Probability. *Annual review of neuroscience* 37:205–220.
- Harris AJL, Hahn U, Madsen JK, Hsu AS (2016) The Appeal to Expert Opinion: Quantitative Support for a Bayesian Network Approach. *Cognitive Science* 40(6):1496–1533.
- Nickerson R (1998) Confirmation bias: a ubiquitous phenomenon in many guises. *Review of general psychology* 2:175–220.
- Yaniv I, Milyavsky M (2007) Using advice from multiple sources to revise and improve judgments. *Organizational Behavior and Human Decision* ... 103:104–120.
- Yaniv I, Kleinberger E (2000) Advice Taking in Decision Making: Egocentric Discounting and Reputation Formation. *Organizational behavior and human decision processes* 83(2):260–281.

18. Soll JB, Larrick RP (2009) Strategies for revising judgment: how (and how well) people use others' opinions. *Journal of experimental psychology: Learning, memory, and cognition* 35(3):780–805.
19. Lorenz J, Rauhut H, Schweitzer F, Helbing D (2011) How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences of the United States of America* 108(22):9020–5.
20. Galton F (1907) Vox Populi. *Nature* 75(1949):450–451.
21. Armstrong JS (2001) *Principles of forecasting: a handbook for researchers and practitioners*. (Kluwer Academic, Norwell, MA), pp. 417–439.
22. Fleming SM, Lau HC (2014) How to measure metacognition. *Frontiers in Human Neuroscience* 8.
23. Patel D, Fleming SM, Kilner JM (2012) Inferring subjective states through the observation of actions. *Proceedings of the Royal Society B: Biological Sciences* 279(1748):4853–4860.
24. Soll JB, Mannes AE (2011) Judgmental aggregation strategies depend on whether the self is involved. *International Journal of Forecasting* 27(1):81–102.
25. Slegers DW, Brake GL, Doherty ME (2000) Probabilistic Mental Models with Continuous Predictors. *Organizational Behavior and Human Decision Processes* 81(1):98–114.
26. Gigerenzer G, Hoffrage U, Kleinbölting H (1991) Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review* 98(4):506–528.
27. Koriat A (2012) The self-consistency model of subjective confidence. *Psychological review* 119(1):80–113.
28. Navajas J, Niella T, Garbulsky G, Bahrami B, Sigman M (2018) Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour* 2(2):126–132.
29. Bahrami B, et al. (2012) What failure in collective decision-making tells us about metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367(1594):1350–1365.

549 **1. Figures**

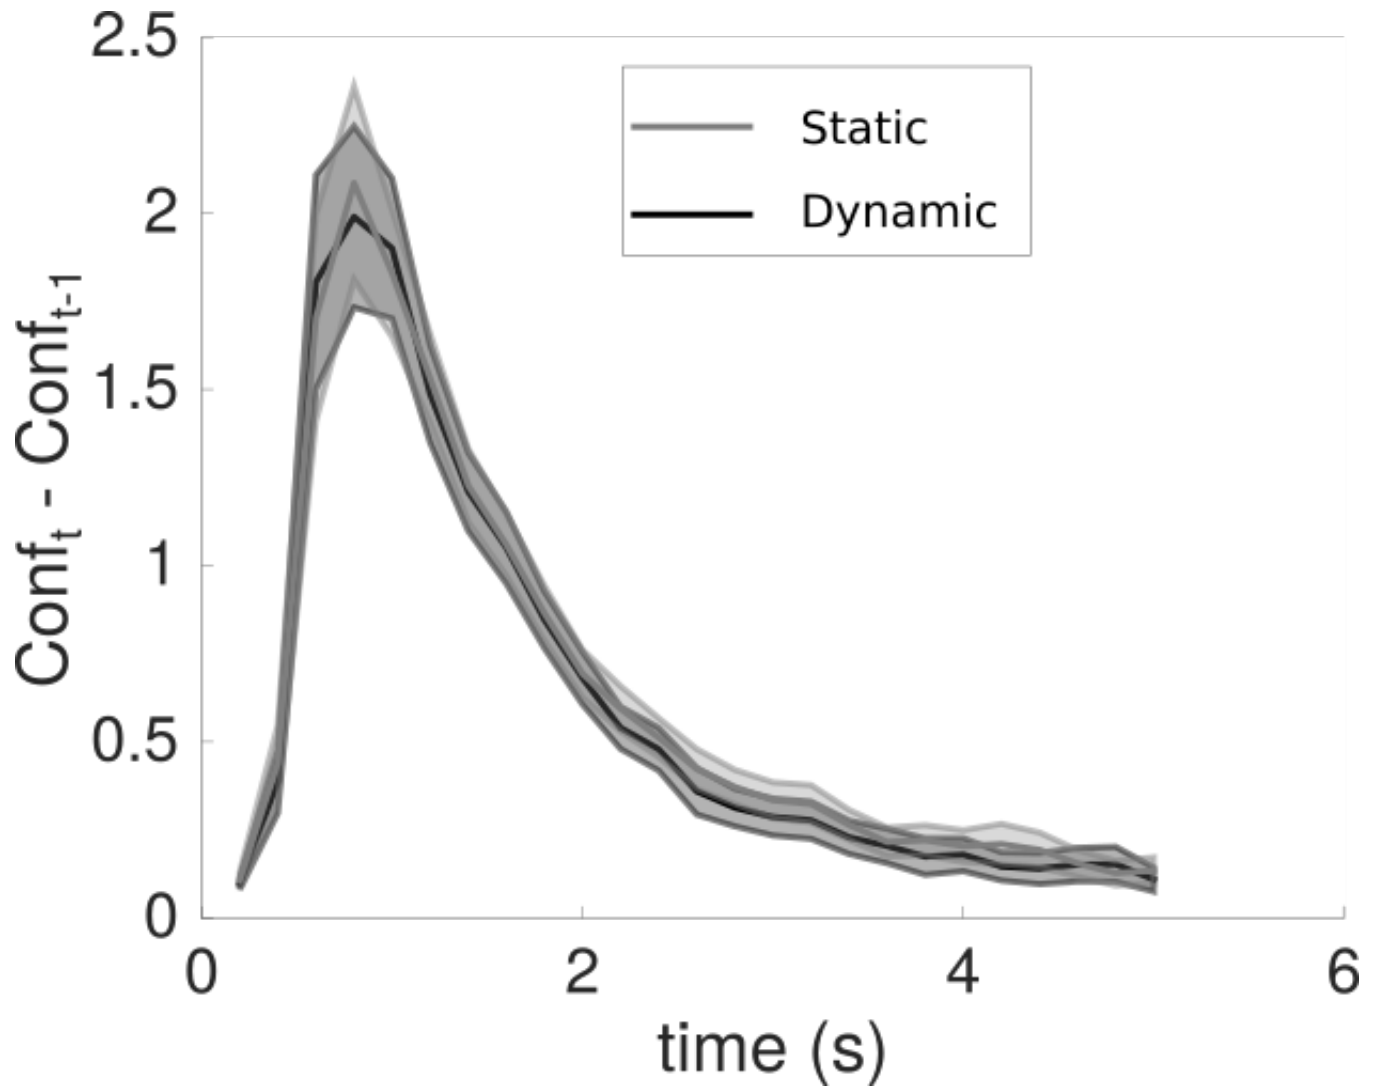


Fig. S1. Average confidence difference between two consecutive data point recorded during the social window (5 seconds). The higher the difference the bigger the update. It can be observed that in all conditions the biggest updates are observed around the first second of the social part. Both dynamic conditions show a larger update around the same time compared to the static baseline condition and a longer time to reach an equilibrium.

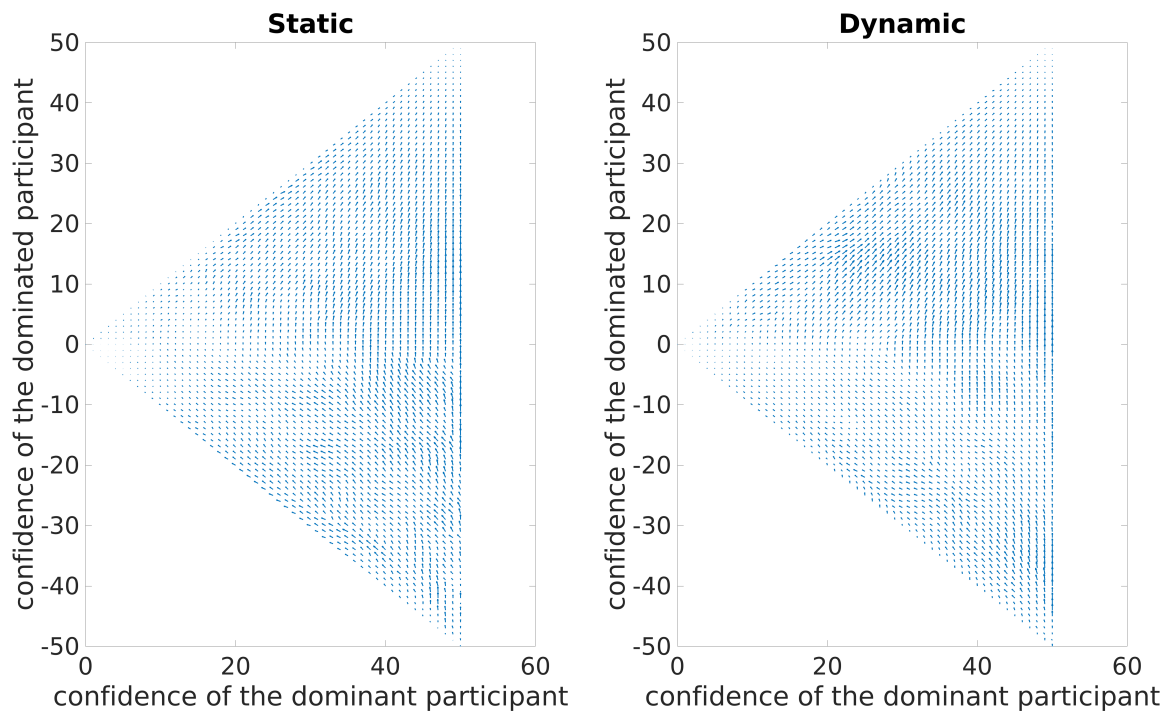


Fig. S2. Dyadic transitions in belief space. Each vector's x and y components are the trial-dominant and trial-dominated participant's confidence changes, as illustrated in Figure 3 in the main text (panels a-b for the Static condition and panels d-e for the Dynamic condition)

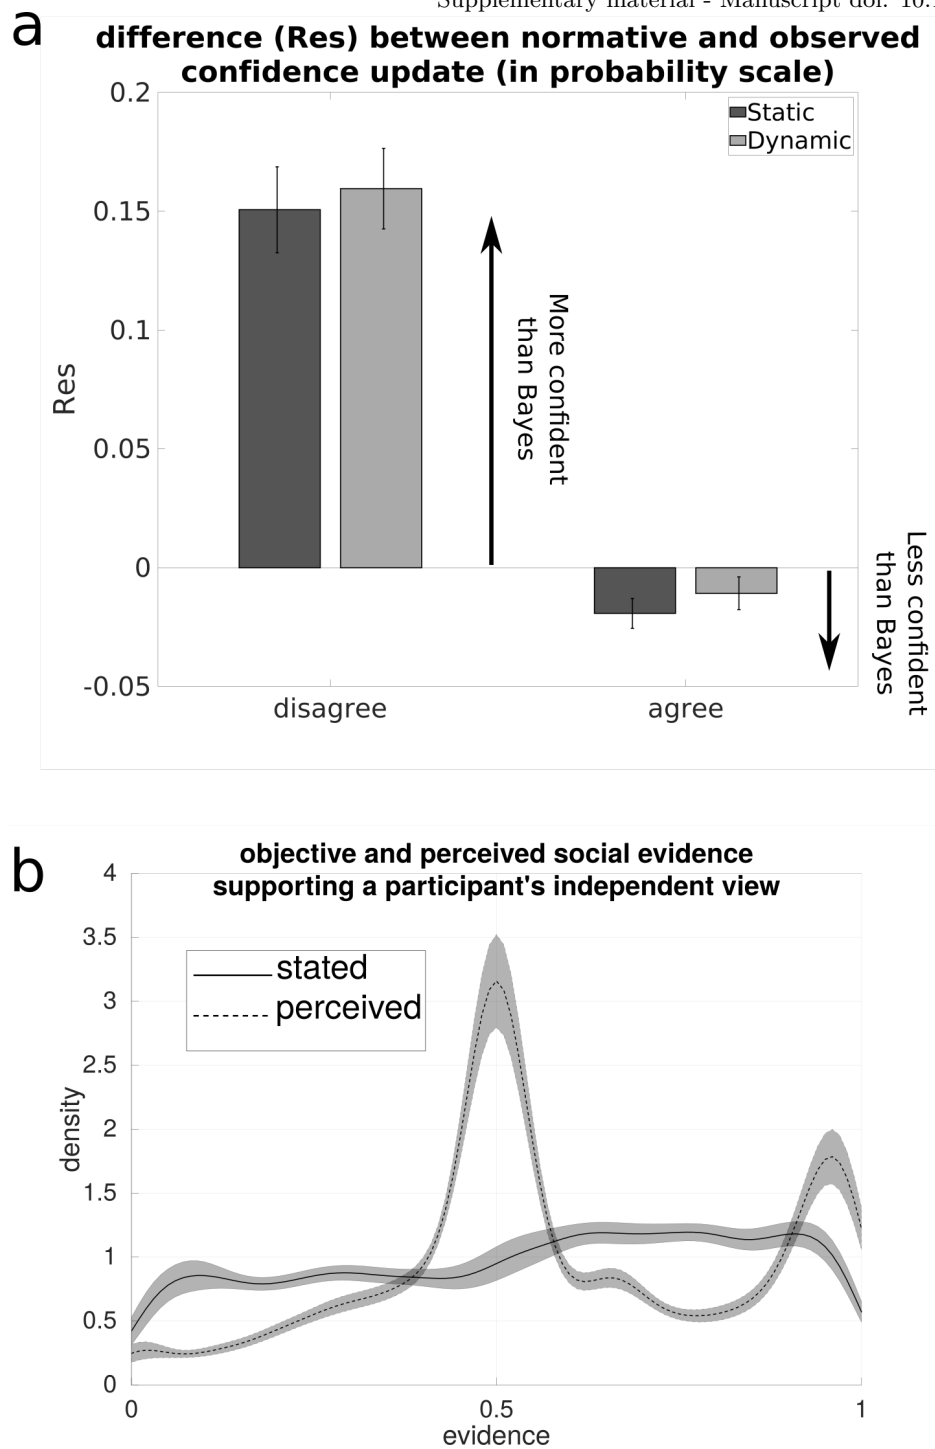


Fig. S3. (a) Difference Res between observed confidence change and normatively prescribed confidence change. **(b)** Contrast between partner's stated support for the participant's independent view (objective social evidence) and the participant's perceived support of the partner's belief (perceived social evidence), as inferred using inverted Bayes. The plots show density distributions calculated with a Gaussian kernel method (bandwidth=0.04). A value of 1 (a value of 0) corresponds to social information that maximally agrees (disagrees) with one's initially expressed belief.

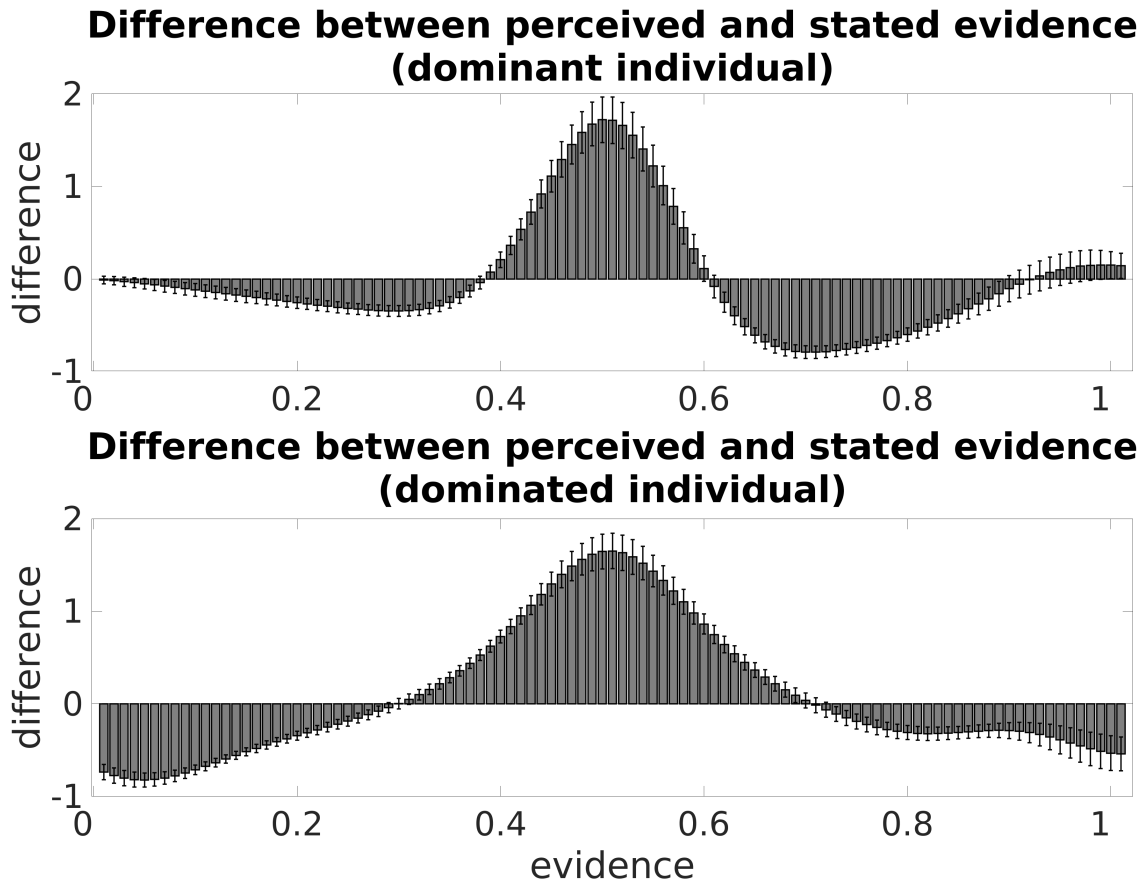


Fig. S4. The figure shows the difference between the social evidence perceived by participant and the evidence actually contained in the participant's partner stated advice. Systematic distortions occur if the participant of interest holds the trial-dominant or trial-dominated position in the trial. Positive bars indicate that the individual uses the advice more than normatively prescribed by an optimal Bayesian observer. Negative bars indicate underuse of the advice.

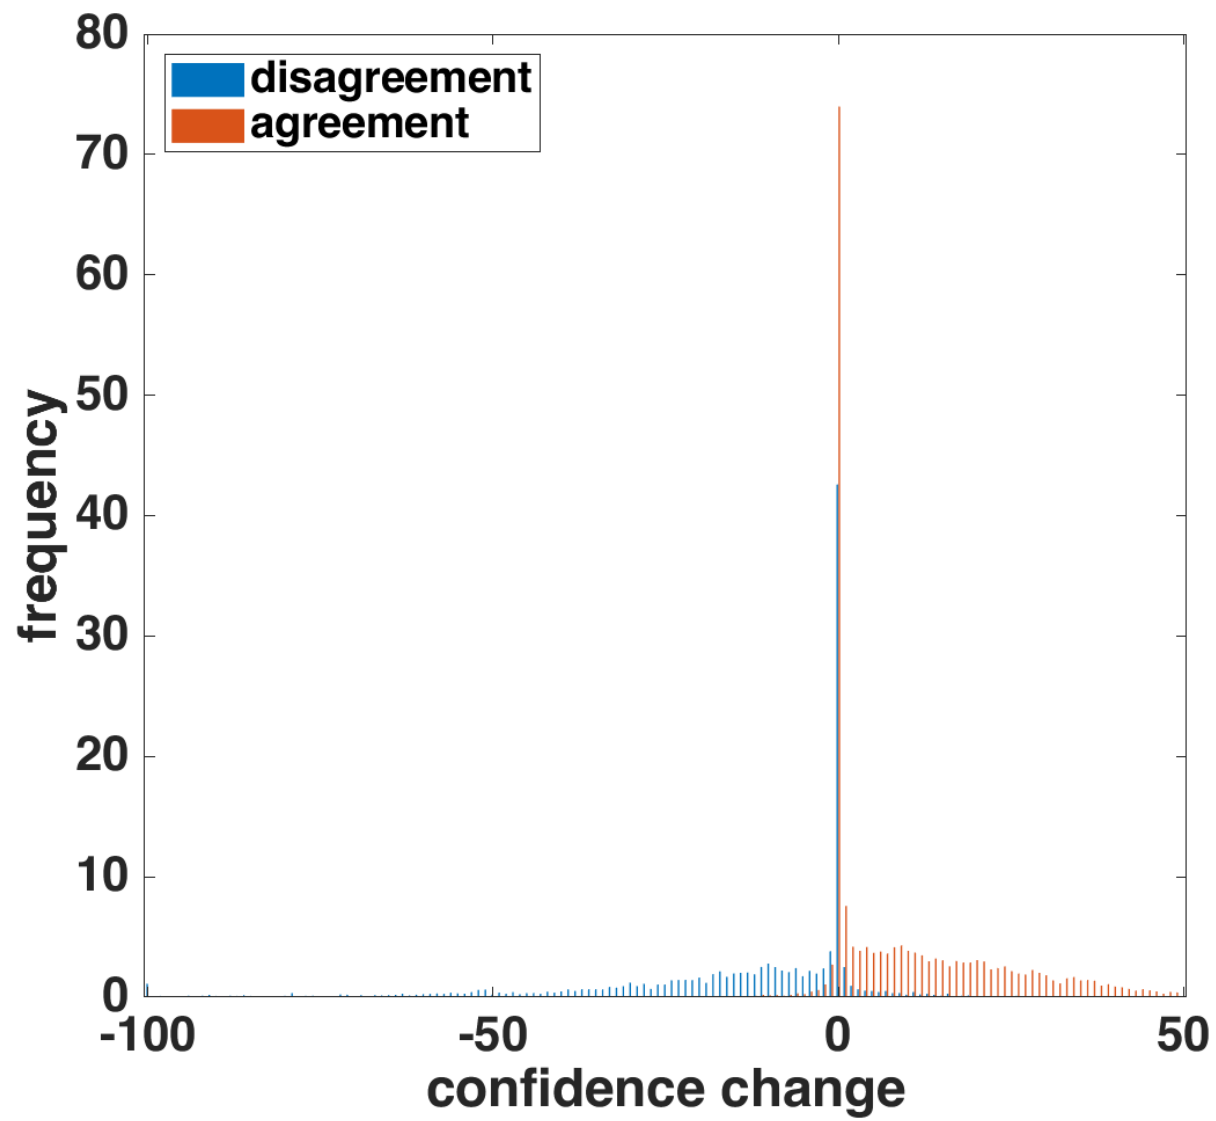


Fig. S5. The figure shows the histogram of raw confidence changes. The figure corresponds to Figure 2b in the main text without any root transformation.

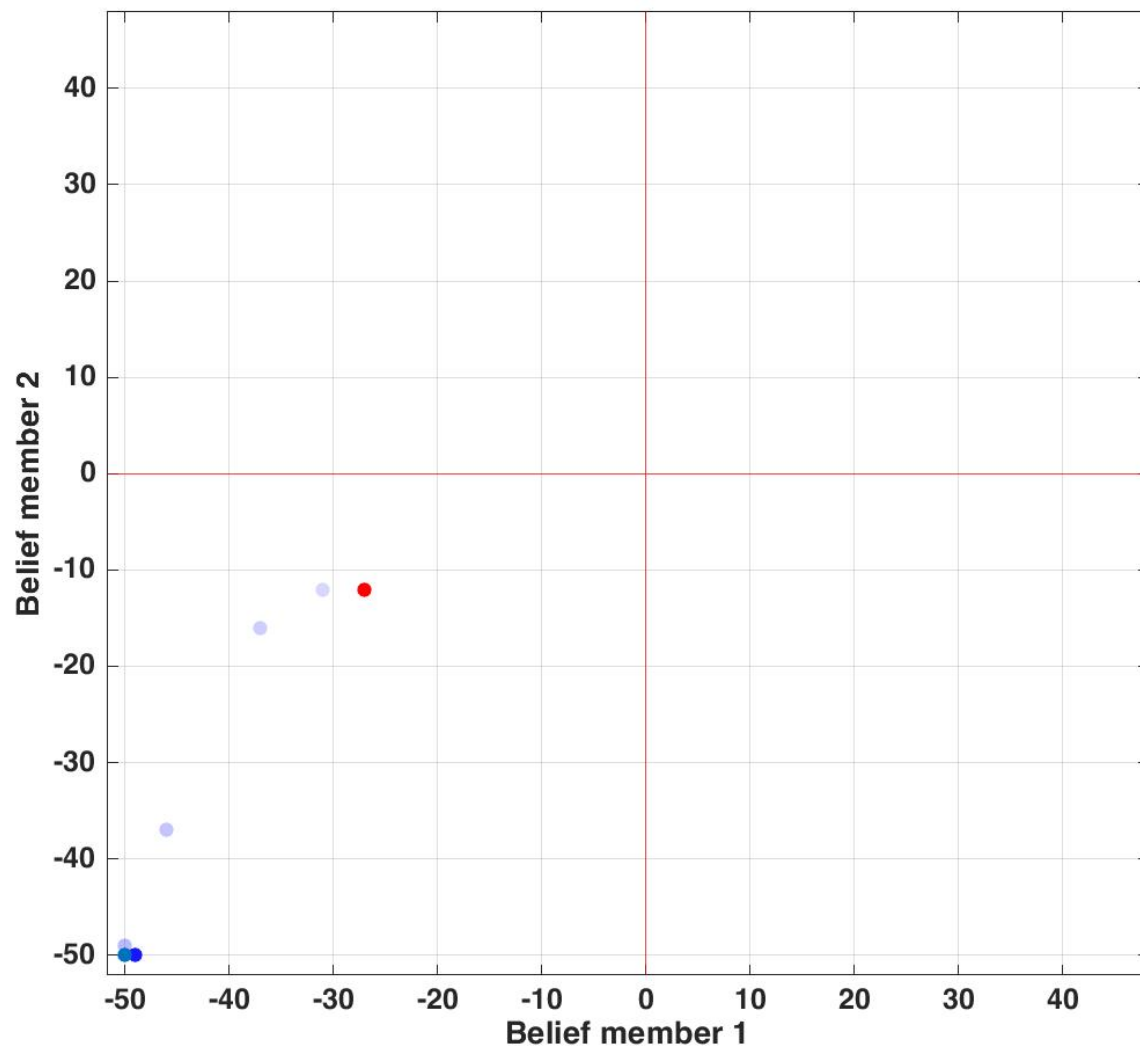


Fig. S6. A agreement trial, where both members started uncertain but ended up very confident. The full opinion space is represented (x-axis = belief of participant 1, y-axis= belief of participant 2). Transparency represent time (from transparent to saturated). The red point represent the starting value, namely dyadic initial state before social interaction.

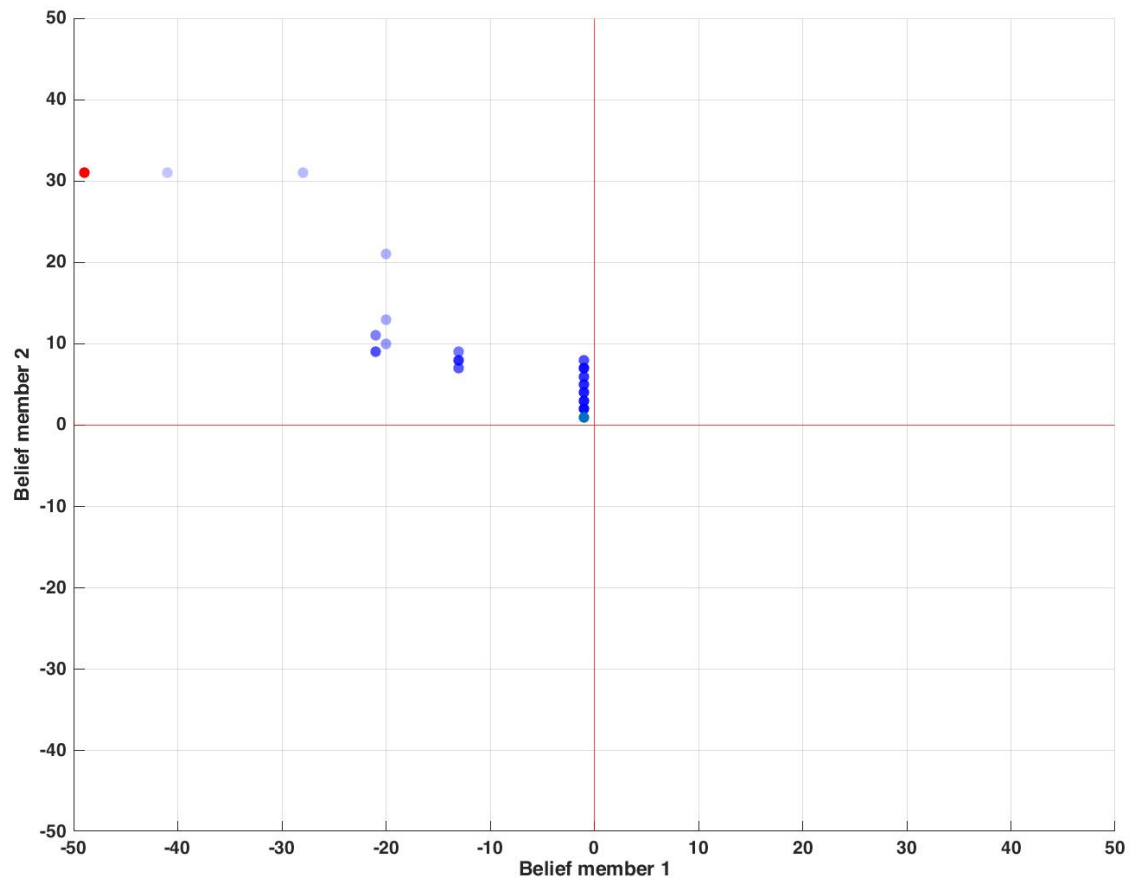


Fig. S7. A disagreement trial, where both members started very confident but ended up very uncertain. The full opinion space is represented (x-axis = belief of participant 1, y-axis= belief of participant 2). Transparency represent time (from transparent to saturated). The red point represent the starting value, namely dyadic initial state before social interaction.

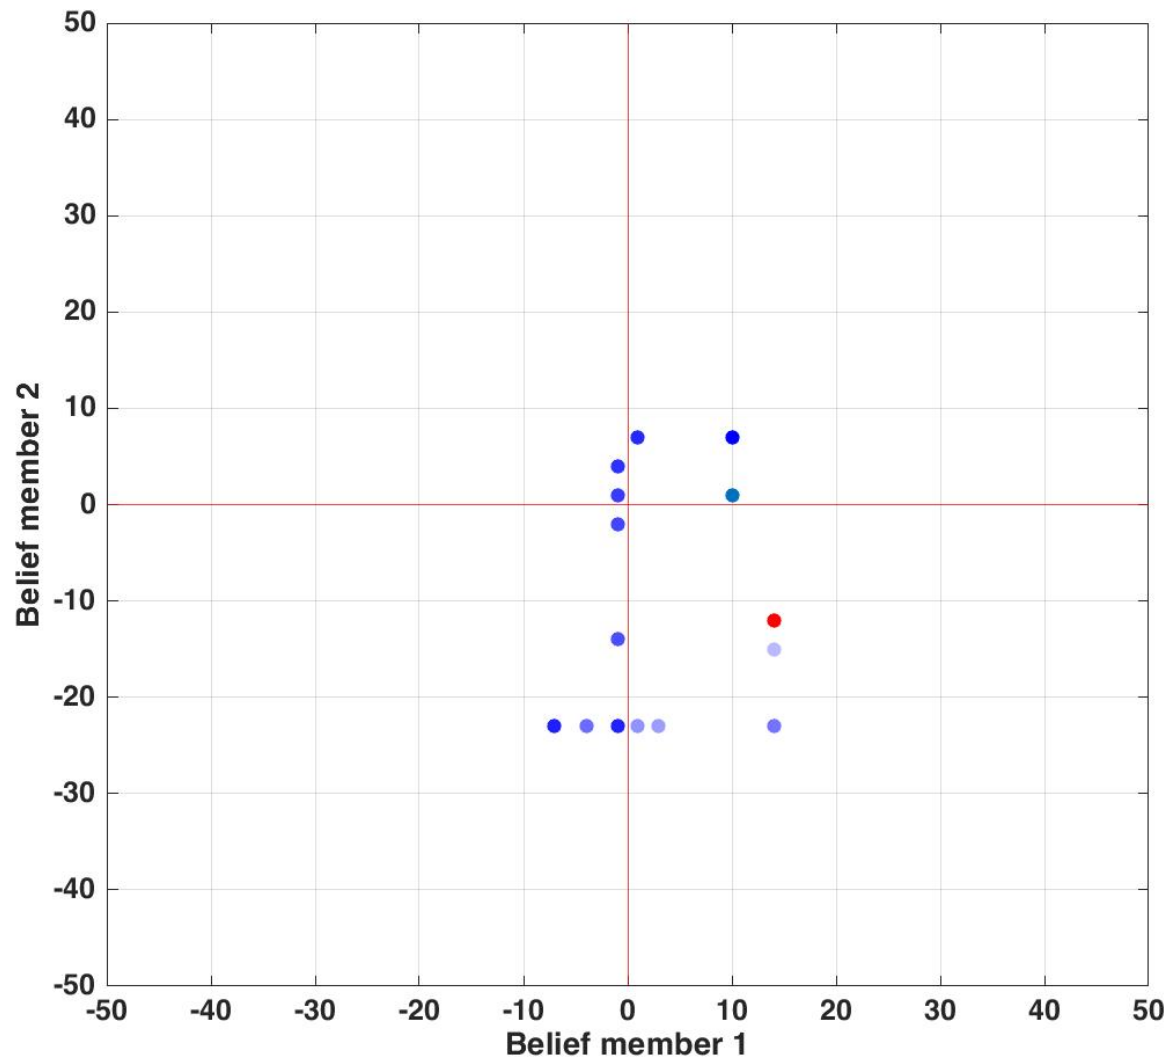


Fig. S8. A dynamic disagreement trial, where members are initially uncertain and end up vacillating between intervals. The full opinion space is represented (x-axis = belief of participant 1, y-axis= belief of participant 2). Transparency represent time (from transparent to saturated). The red point represent the starting value, namely dyadic initial state before social interaction.

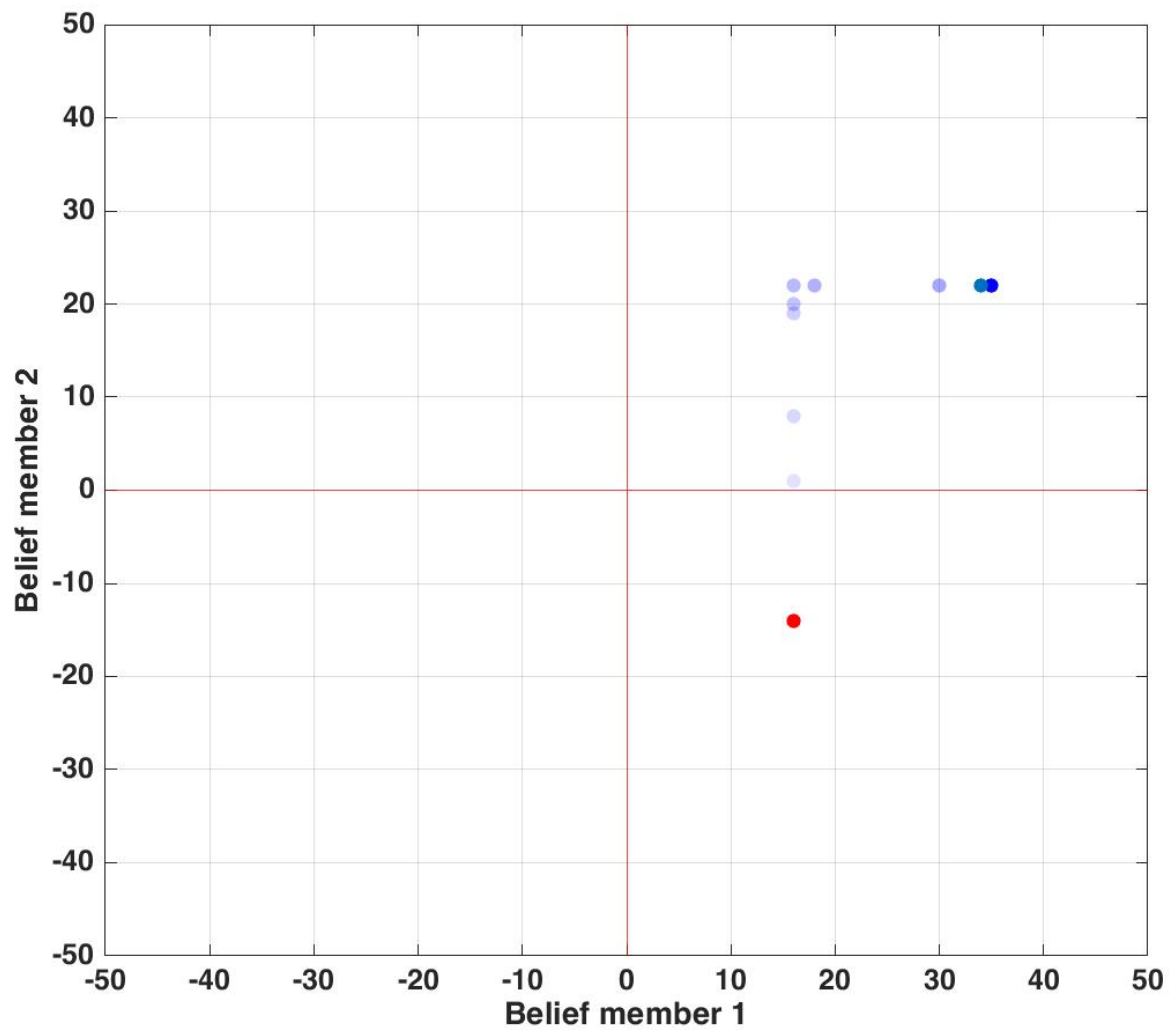


Fig. S9. A dynamic disagreement trial with irrational increase in confidence. Notice that once disagreement is resolved (member 2 changes their mind) member 1 increases their initial confidence. The full opinion space is represented (x-axis = belief of participant 1, y-axis= belief of participant 2). Transparency represent time (from transparent to saturated). The red point represent the starting value, namely dyadic initial state before social interaction.

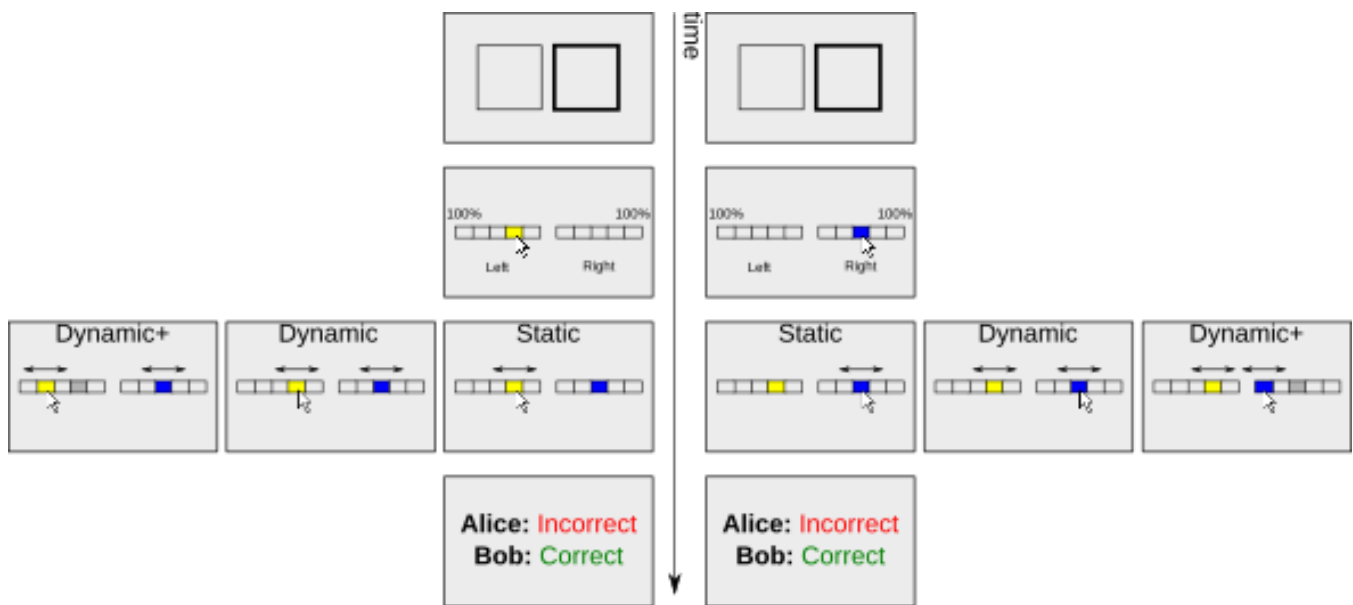


Fig. S10. Experimental paradigm implemented in Experiment 2. Three conditions are explored and compared within-participants. During the Static condition participants are shown the initial independent belief of their partner. During the Dynamic condition participants are shown the current belief of their partner in real-time. During the Dynamic plus self-reminder condition (*Dynamic_{self}*) participants are shown the current real-time belief of their partner and are at the same time reminded of their own original belief as a shaded cursor on the scale. This manipulation makes sure that if participants update their initial confidence they are constantly reminded of where along the scale they started from. In all conditions participants have four seconds when they are asked to track their confidence state in real-time. The confidence scale that was actually used had 50 levels per interval.

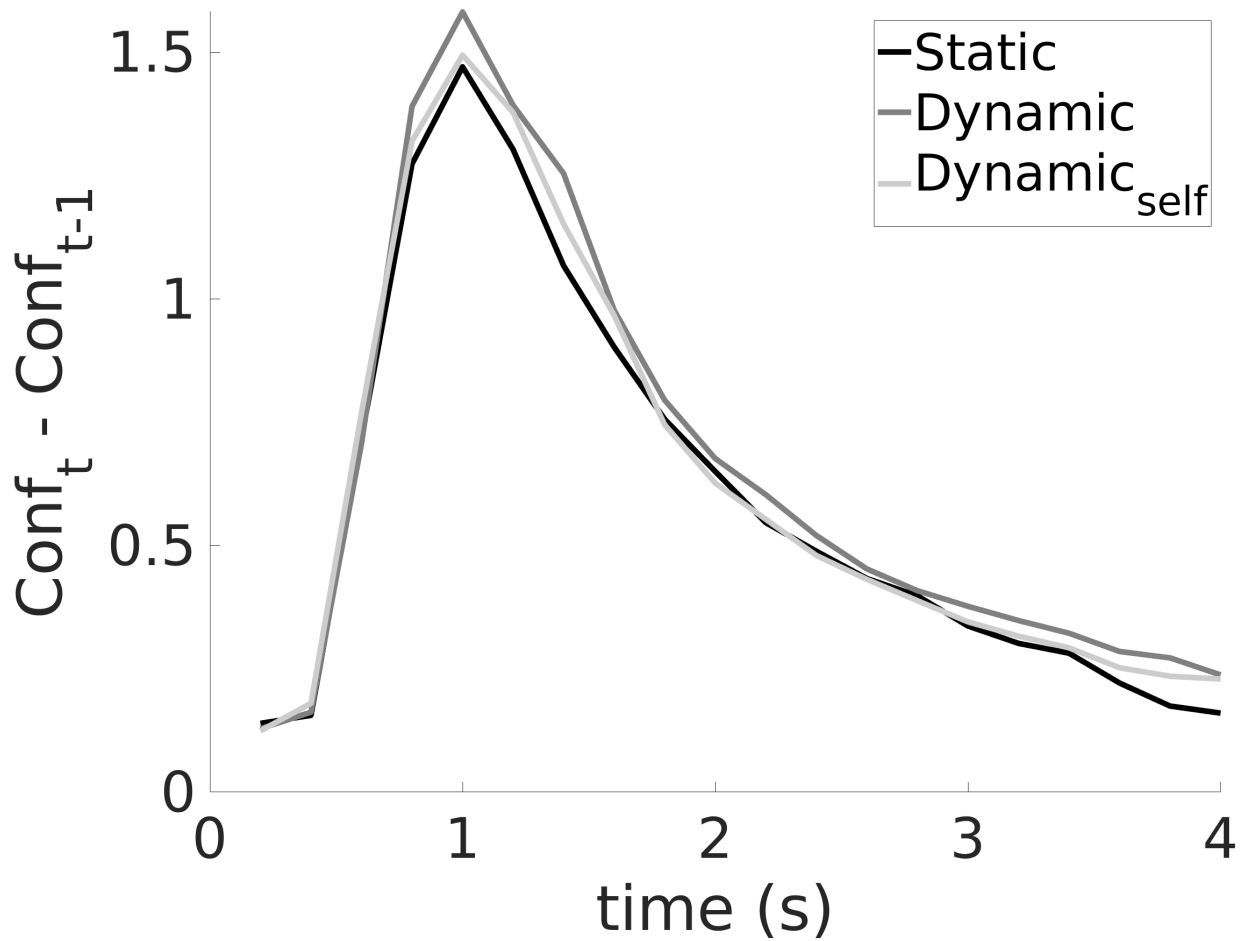


Fig. S11. Experiment 2. Average confidence difference between two consecutive data point recorded during the social window. The higher the difference the bigger the update. It can be observed that in all conditions the biggest updates are observed around the first second of the social part. Both dynamic conditions show a larger update around the same time compared to the static baseline condition and a longer time to reach an equilibrium.

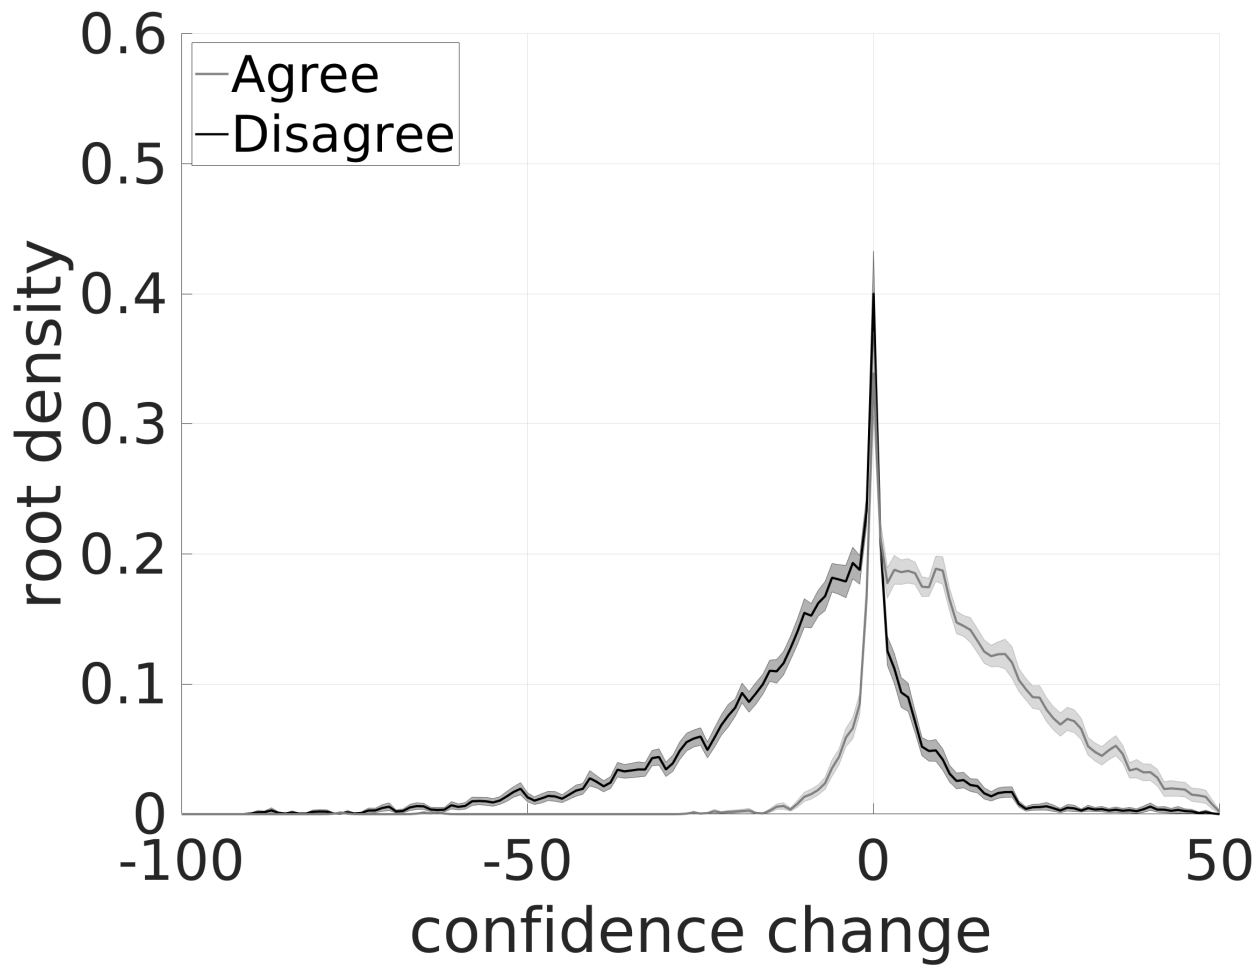


Fig. S12. Experiment 2. Confidence change distributions observed in the most confident participants divided by consensus. Plots represent estimated probability density functions using a normal kernel method (bandwidth = 0.50). Error bars represent s.e.m.

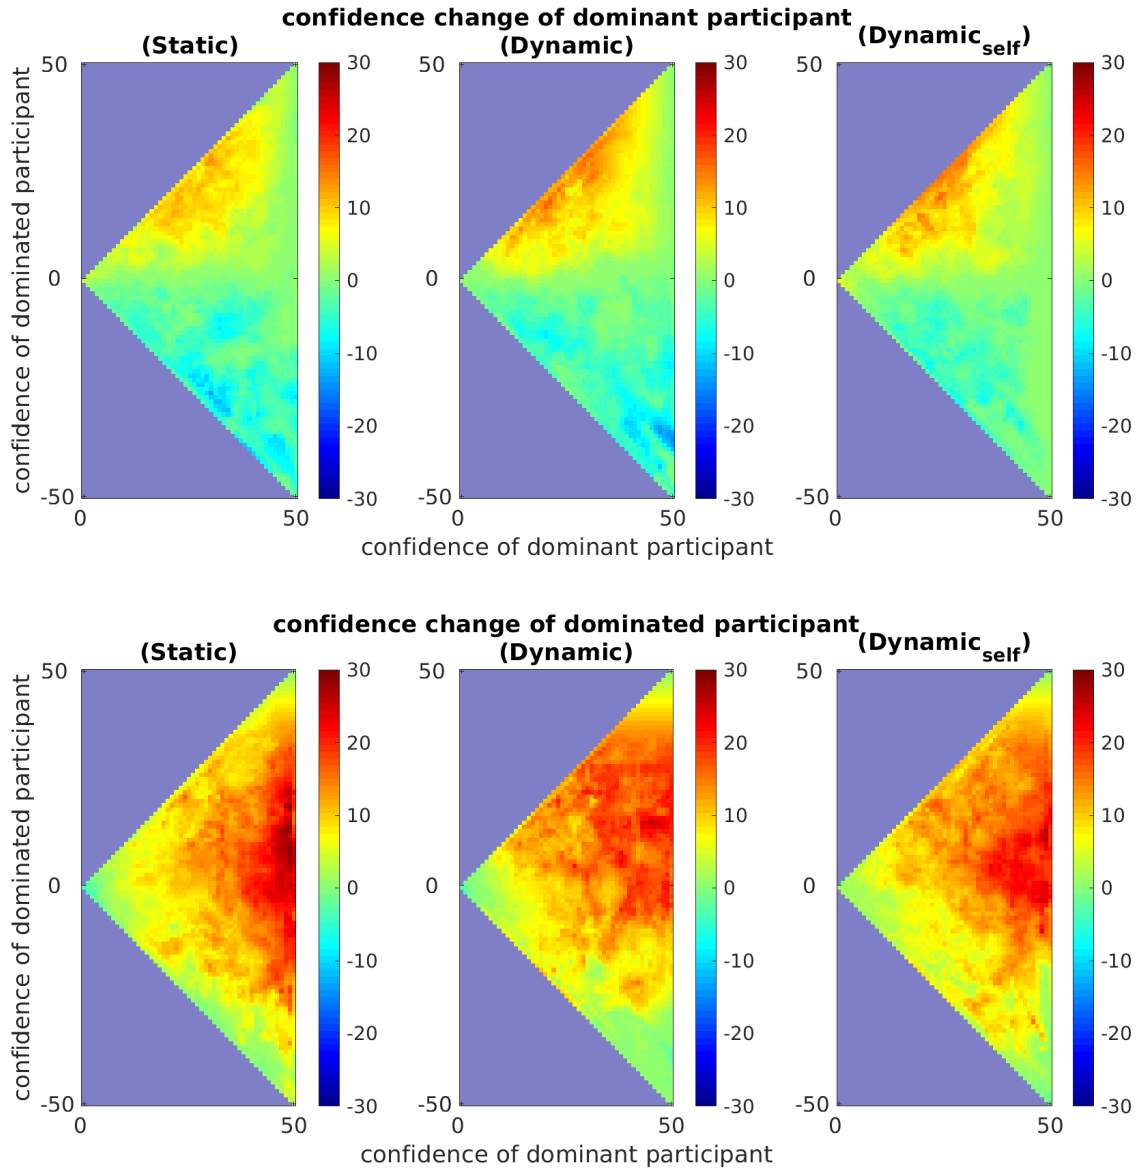


Fig. S13. Median confidence change in belief space divided by condition and trial-dominance (first three columns). Warmer colours represent confidence changes in the direction of the trial-dominant belief, while colder colours represent confidence changes further away from it.

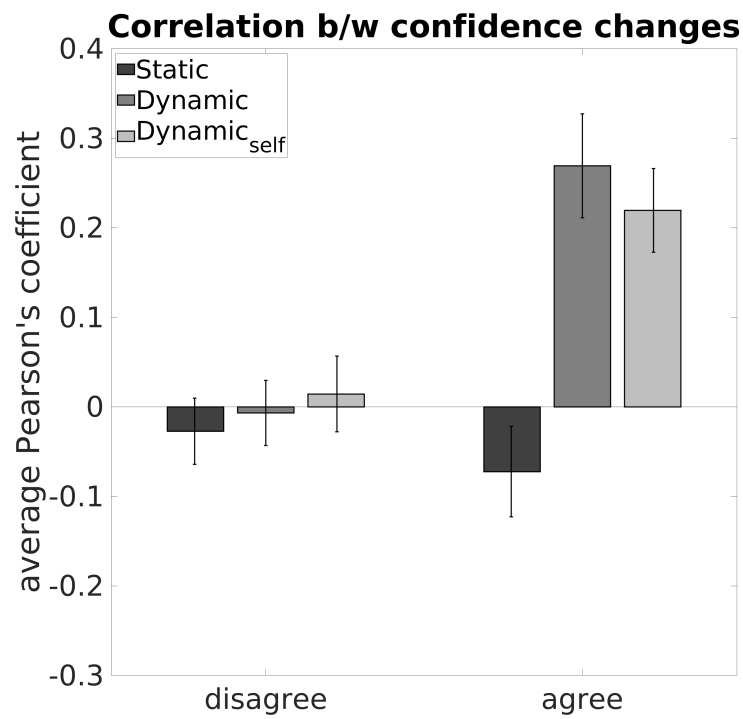


Fig. S14. Coupling between absolute confidence updates of the two participants across different conditions and divided by consensus. It can be seen that in disagreement updates of one dyad member are not correlated with updates of the other member. In agreement on the contrary a positive correlation emerges as soon as participants are allowed to interact in real-time.

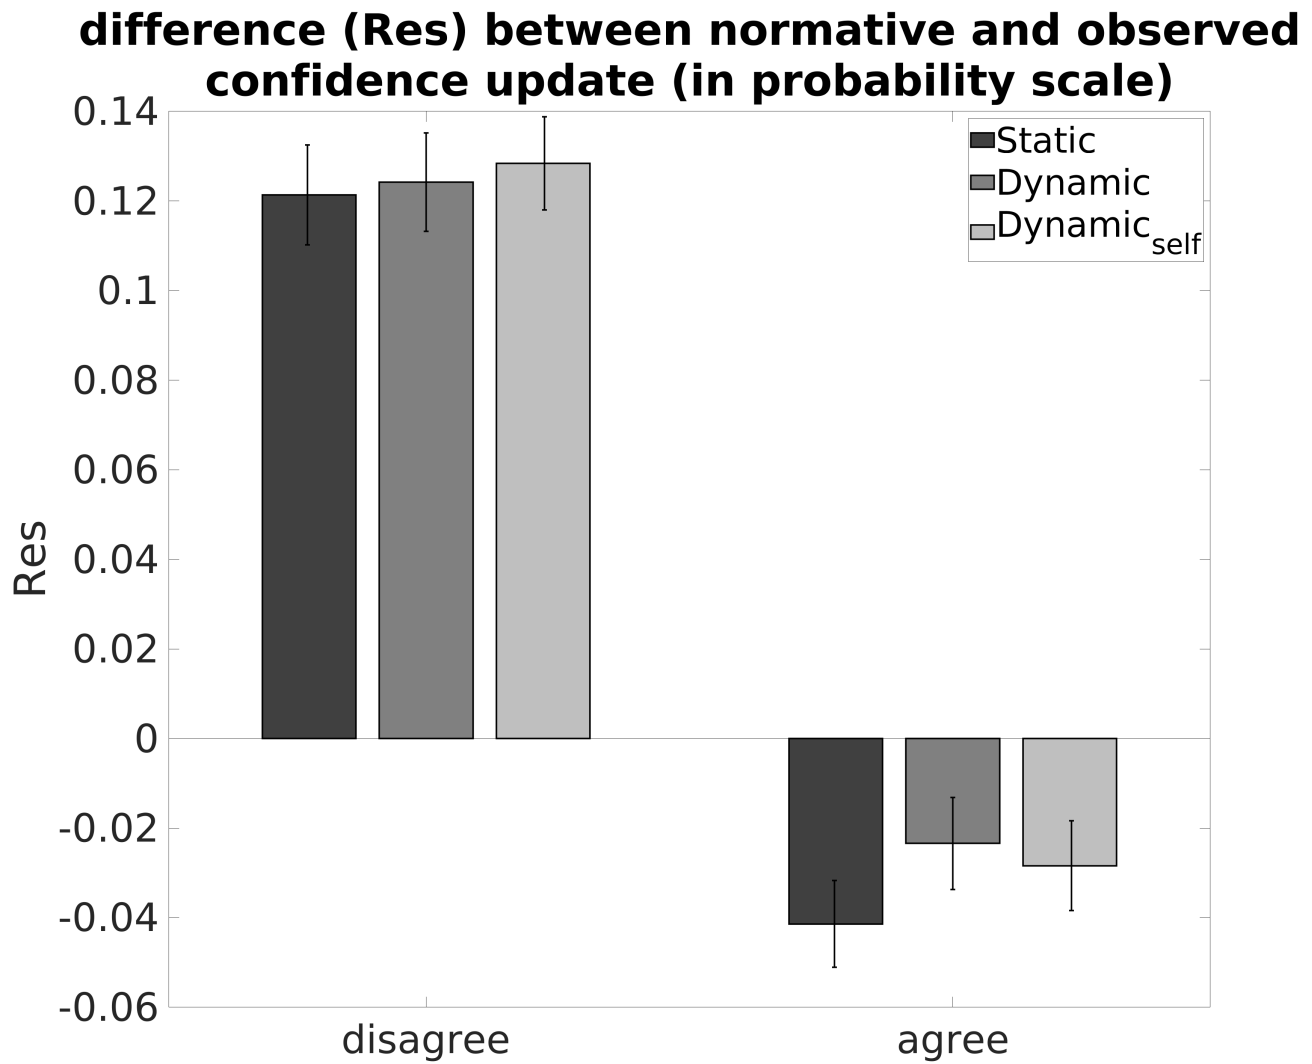


Fig. S15. Experiment 2 - Human data compared to equal-weights model. The figure shows how confidence changes observed in the data relate with the confidence changes expected by applying a normative Bayesian update rule. Participants showed a conservative bias thus decreased their confidence too little in disagreement and increased it too little in agreement trials.

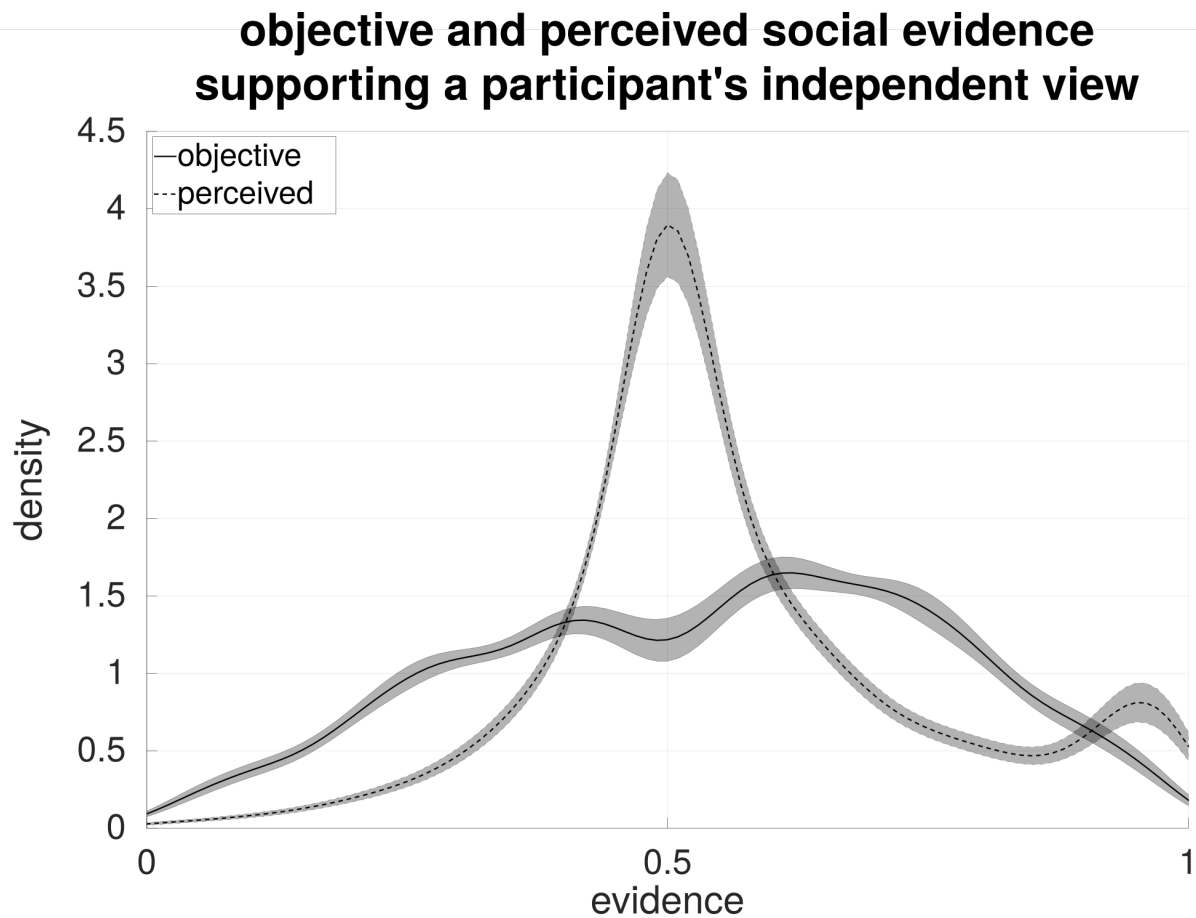


Fig. S16. How much a partner's belief is perceived to support one's own independent judgment, compared to objectively stated partner's supporting evidence. Differences between the two indicate cognitive distortions of social information.

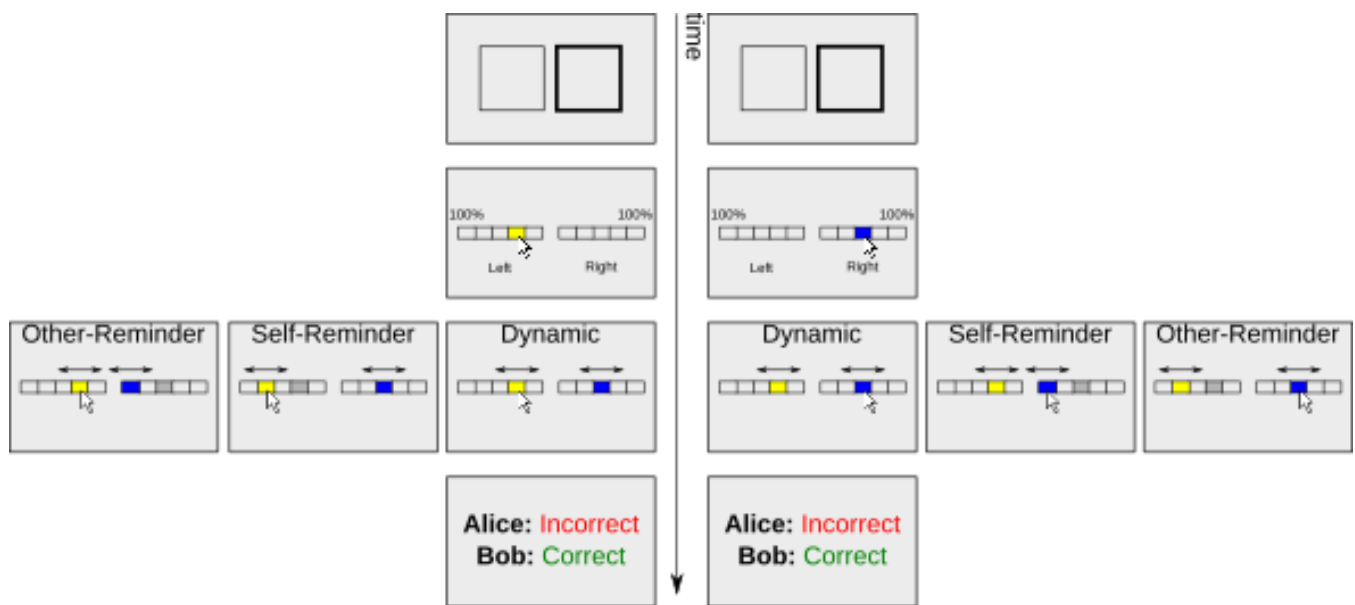


Fig. S17. Experimental paradigm implemented in Experiment 3. During the Dynamic condition participants are shown the current belief of their partner in real-time. During the $Dynamic_{self}$ participants are shown the current real-time belief of their partner and are at the same time reminded of their own original belief as a shaded cursor on the scale. During the $Dynamic_{other}$ participants are shown the current real-time belief of their partner and are at the same time reminded of their partner's original belief as a shaded cursor on the scale. This manipulation makes sure that after a change in the configuration of the elements present on screen participants are reminded of where along the scale they started from or where their partner started from. In all conditions participants have four seconds when they are asked to update their own original confidence level using post-decisional information.

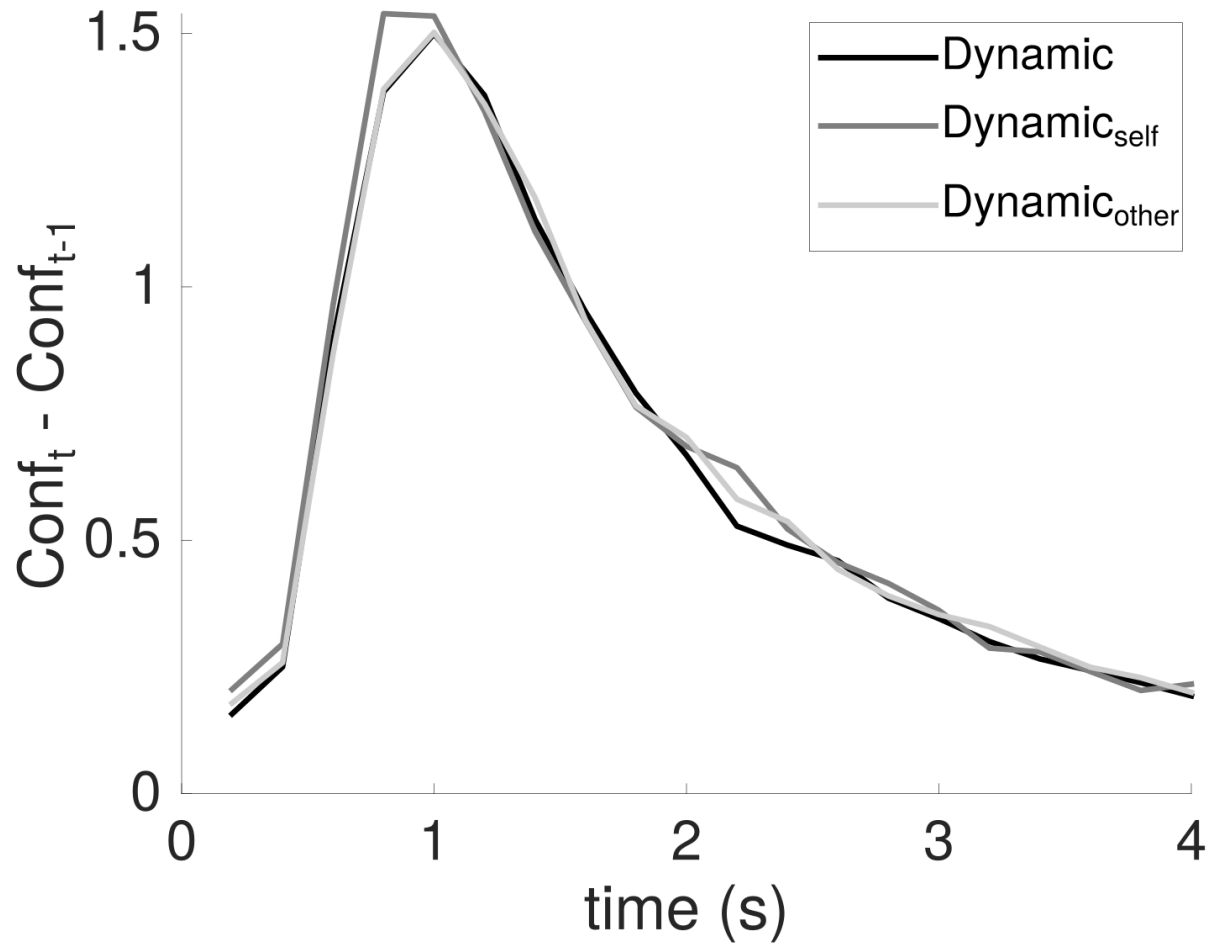


Fig. S18. Experiment 3 - Continuous update over time. Difference in recorded confidence between two subsequent data points during the social window. The measure can be used to plot how quickly participants' updates converged to a final confidence level. Right panels: within-participants point-wise difference between anchor conditions and dynamic baseline.

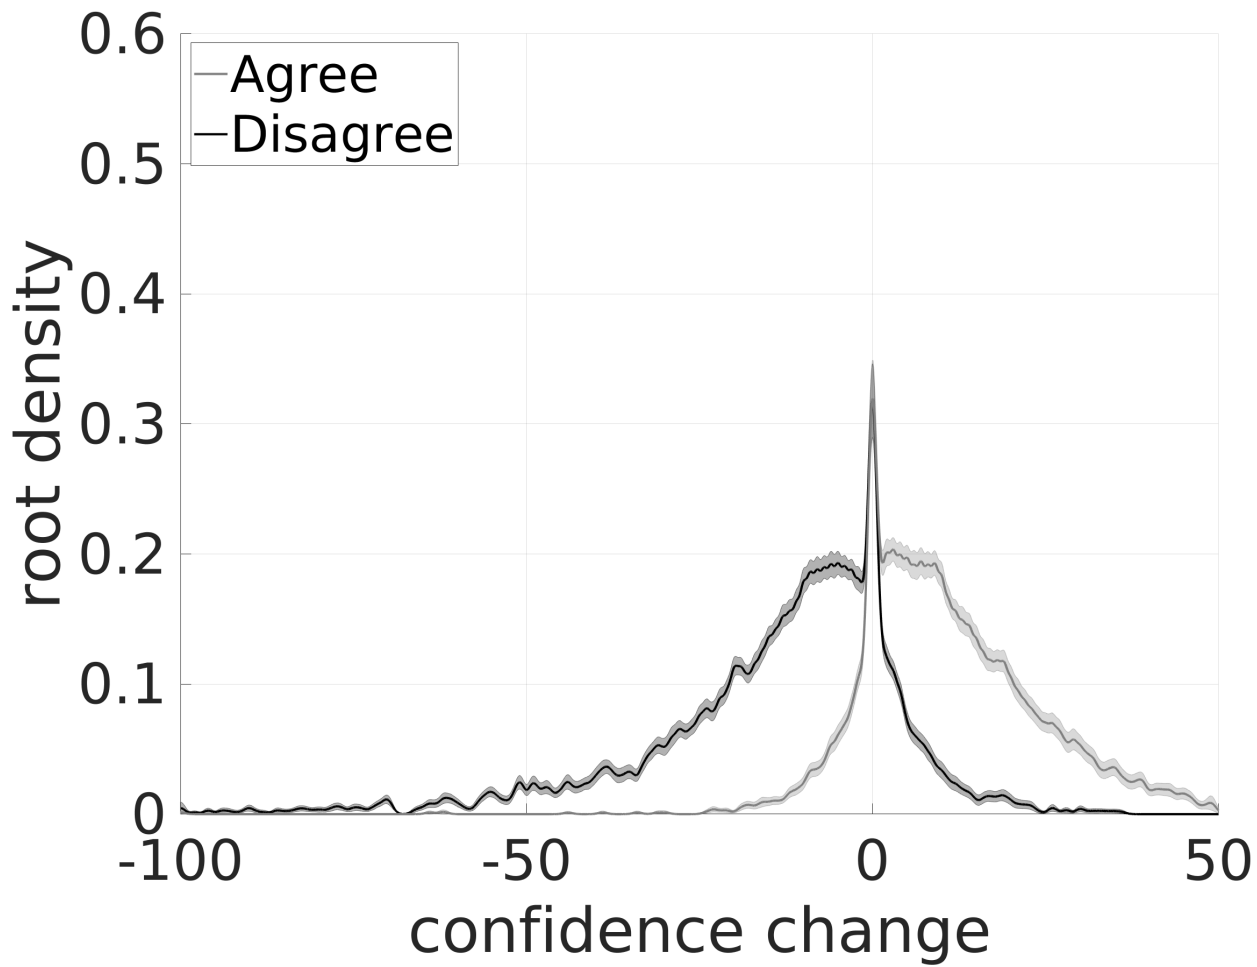


Fig. S19. Experiment 3 - Root density distributions of confidence changes divided by condition and consensus. Density plots are obtained from Gaussian kernel function with bandwidth = 0.50.

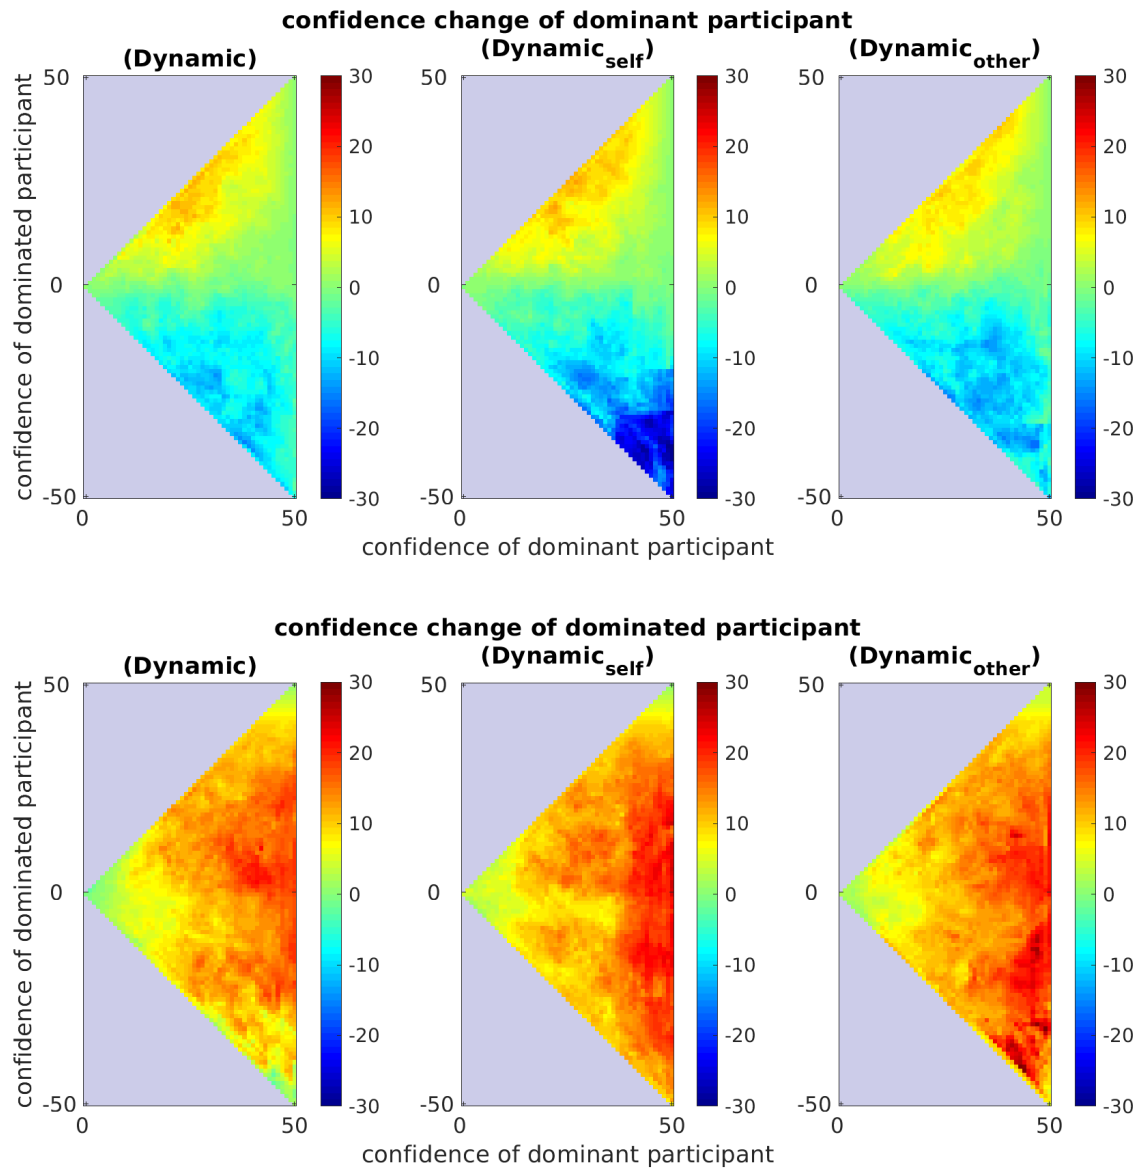


Fig. S20. Experiment 3 - Confidence change in belief space. Median confidence change in belief space for trial-dominant and trial-dominated trials and divided by condition.

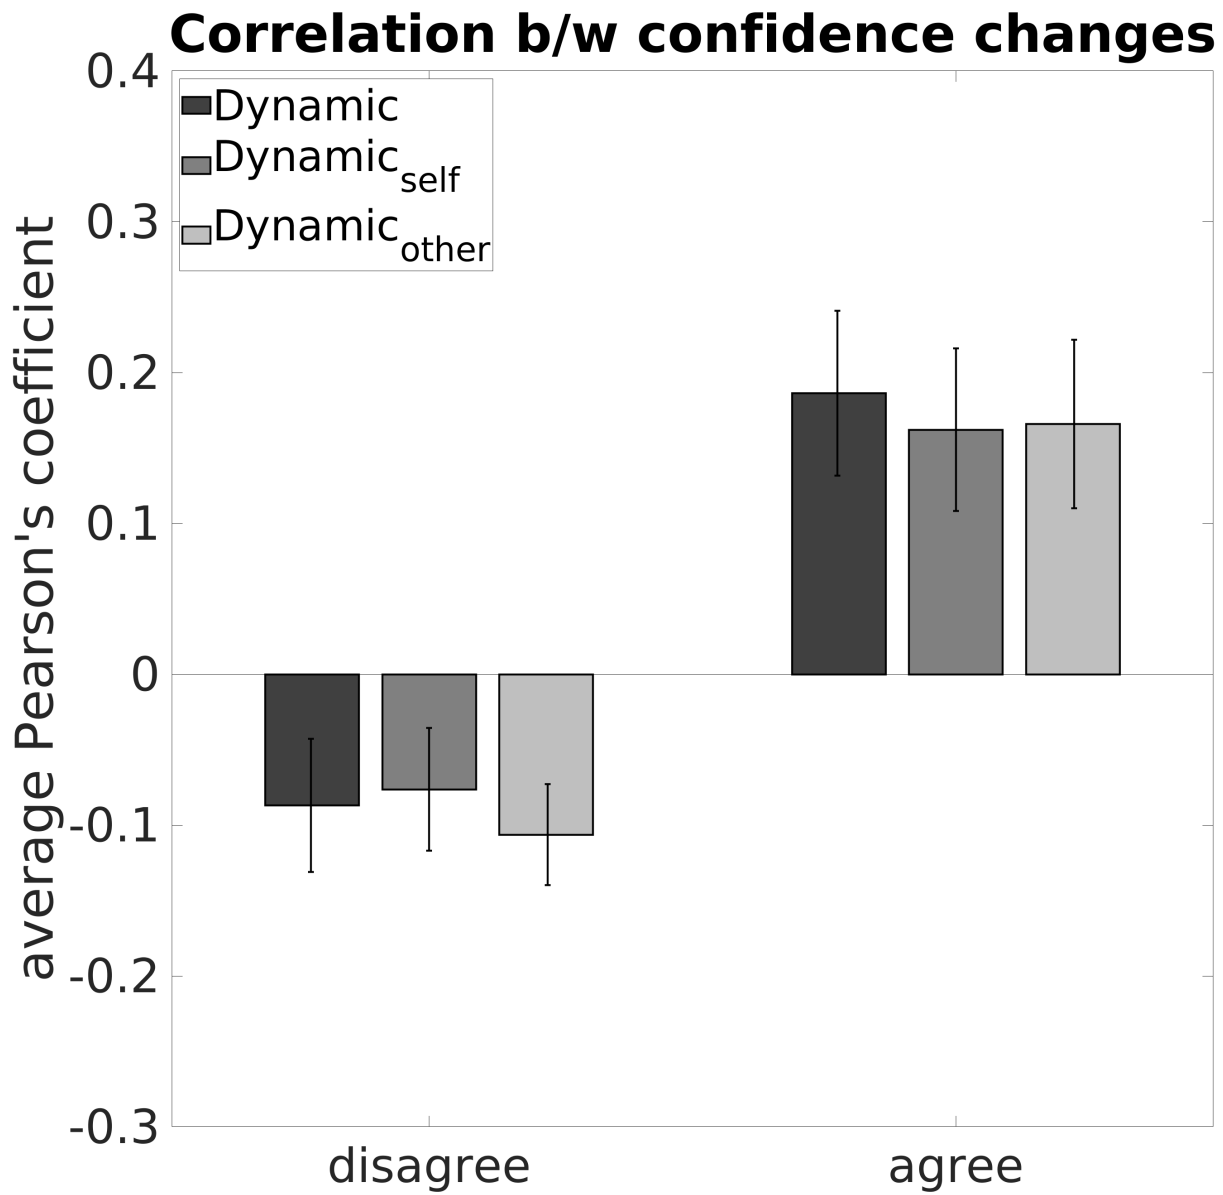


Fig. S21. Experiment 3 - Coupling (as measured by Pearson's correlation coefficient r) between absolute confidence changes of members of the same dyad. Error bars represent s.e.m.

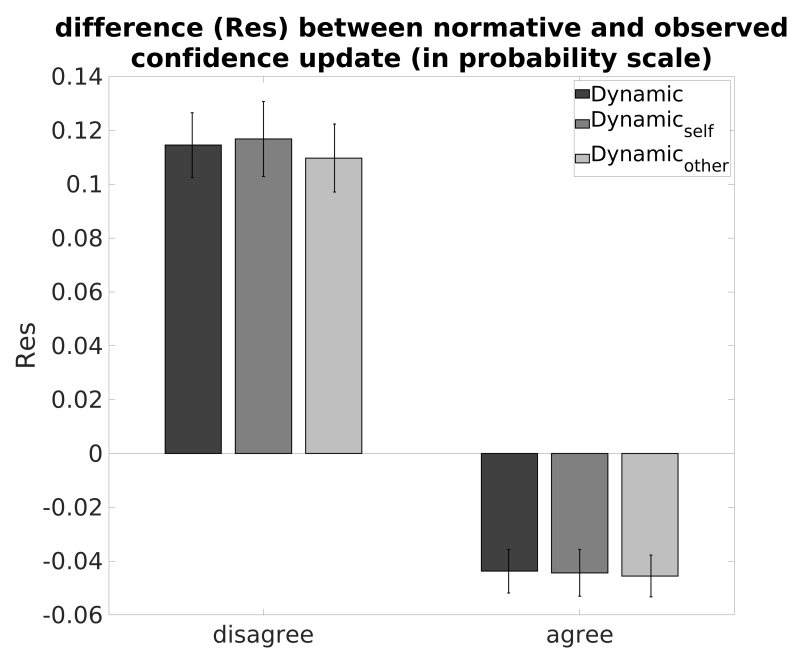


Fig. S22. Experiment 3. Residuals between human participants and a simple Bayesian model aggregating the two beliefs using equal weights. Residuals represent over- or under-confidence compared to model's predictions.

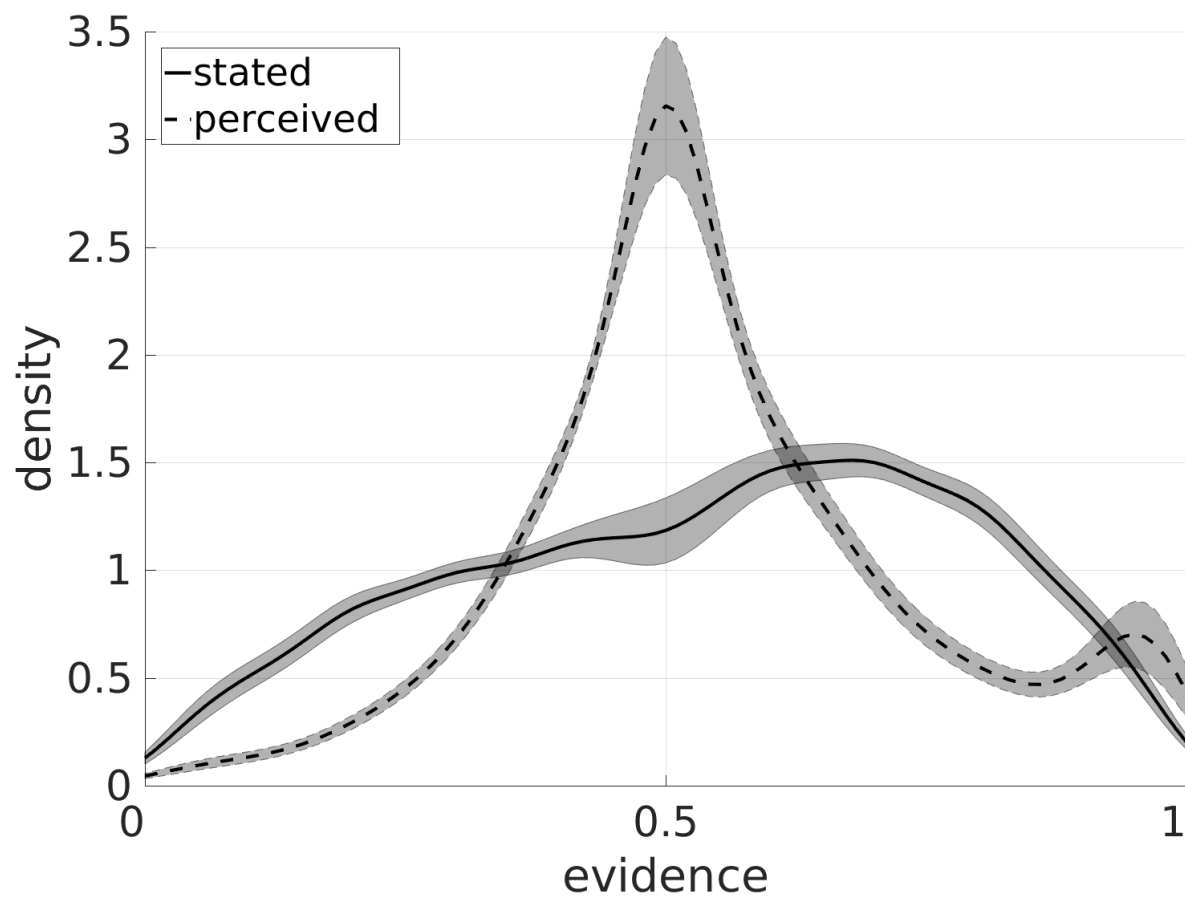


Fig. S23. Experiment 3 - Objective and perceived social support. The figure compares the distribution of the supporting evidence provided by the partner's social information (objective evidence) with the evidence estimated to be perceived by the participant. The graphs show a dissociation between the two.

	Estimate	SE	tStat	DF	p
<i>Intercept</i>	0.1712	0.0331	5.1645	15015	2.44e-07
<i>Agreement(Agr)</i>	-0.2796	0.0523	-5.3444	15015	9.20e-08
C_{pre}^p	0.2759	0.0195	14.131	15015	4.72e-45
$rt2$	0.1263	0.0149	8.4359	15015	3.58e-17
<i>Dynamic(Dyn) : Agr</i>	0.11391	0.0239	4.7536	15015	2.01e-06
<i>Dyn : δ_C^p</i>	-0.1075	0.0152	-7.0354	15015	2.07e-12
<i>Agr : δ_C^p</i>	0.0777	0.0211	3.6827	15015	.0002
<i>Agr : C_{pre}^s</i>	-0.3891	0.0172	-22.55	15015	8.05e-111
<i>$\delta_C^p : C_{pre}^s$</i>	-0.0635	0.0078	-8.1064	15015	5.61e-16
<i>Agr : C_{pre}^p</i>	-0.0581	0.0190	-3.052	15015	.0022
<i>$\delta_C^p : C_{pre}^p$</i>	-0.0238	0.0103	-2.2993	15015	0.0215
<i>$C_{pre}^s : C_{pre}^p$</i>	-0.07427	0.0118	-6.2762	15015	3.56e-10
<i>Agr : $rt2$</i>	-0.0653	0.0146	-4.4555	15015	8.43e-06
<i>$C_{pre}^s : rt2$</i>	0.0705	0.0101	6.9787	15015	3.10e-12
<i>Dyn : Agr : δ_C^p</i>	0.2808	0.0309	9.0767	15015	1.25e-19
<i>Dyn : Agr : C_{pre}^s</i>	-0.0502	0.0190	-2.6382	15015	.0083
<i>Dyn : Agr : C_{pre}^p</i>	0.0563	0.0212	2.6462	15015	.0081
<i>Dyn : $\delta_C^p : C_{pre}^p$</i>	-0.0670	0.0171	-3.9075	15015	9.36e-05
<i>Agr : $C_{pre}^s : C_{pre}^p$</i>	-0.0589	0.0150	-3.9073	15015	9.37e-05
<i>$\delta_C^p : C_{pre}^s : C_{pre}^p$</i>	0.0325	0.0097	3.3368	15015	.0008
<i>Agr : $C_{pre}^s : rt2$</i>	-0.0653	0.014	-4.6677	15015	3.07e-06
<i>Dyn : Agr : $\delta_C^p : C_{pre}^p$</i>	0.1045	0.0197	5.2838	15015	1.28e-07
<i>Dyn : $\delta_C^p : C_{pre}^s : C_{pre}^p$</i>	0.0650	0.0175	3.695	15015	.0002
<i>Dyn : $\delta_C^p : C_{pre}^s : rt2$</i>	-0.0267	0.0093	-2.8589	15015	.004
<i>Dyn : $C_{pre}^s : C_{pre}^p : rt2$</i>	0.0215	0.0095	2.5262	15015	.0240
<i>Dyn : Agr : $\delta_C^p : C_{pre}^s : C_{pre}^p$</i>	-0.0497	0.0189	-2.6244	15015	.0086

Table S1. Experiment 1 - Fixed effects of linear mixed-effect multilevel model run on trial-by-trial absolute confidence update. Main predictors are (a) condition: Static (reference), Dynamic (*Dyn*); (b) consensus: Disagreement (reference), Agreement (*Agr*); (c) partner's absolute confidence change ($|\delta_C^p|$); (d) personal initial confidence (C_{pre}^s); (e) partner's initial confidence (C_{pre}^p); (f) partner's update reaction time ($rt2$).