Averaging versus Sampling in Collaborative Judgement

Rebecca Floyd¹, David Leslie², Roland Baddeley¹, & Simon Farrell³

¹ University of Bristol

² University of Lancaster

³ University of Western Australia

Author Note

Rebecca Floyd and Roland Baddeley: School of Experimental Psychology, University of Bristol; David Leslie: Mathematics and Statistics Department, University of Lancaster; Simon Farrell: School of Psychology, University of Western Australia.

This research was supported by EPSRC grant EP/I032622/1 awarded to Gilchrist, Baddeley, Bogacz, Farrell, Leslie, Ludwig & McNamara. Farrell was additionally supported by an Australian Research Council Future Fellowship (FT130100149).

The stimulus materials for Experiment 3, and the primary data files and analysis script for all experiments, are available on the Open Science Framework at <https://osf.io/jex8h/>.

Correspondence concerning this article should be addressed to Simon Farrell, School of Psychological Science, University of Western Australia, 35 Stirling Hwy, Crawley WA 6009.

E-mail: simon.farrell@uwa.edu.au

Abstract

How does a dyad combine information from different members in order to arrive at a consensus judgement? One suggestion is that groups combine information in a Bayes optimal fashion: the group calculates a weighted average of individuals' estimates, with the weightings being proportional to the quality of the information each individual possesses. Alternatively, the dyad may seek to identify which member's estimate is the best, and return that as a joint judgement. These models were tested by asking members of a dyad to make private estimates of a continuous quantity (the direction of movement of a coherent motion stimulus), and to then make a joint judgement. Joint judgements were more accurate than individual judgements, but were only partly based on optimal integration. Rather, the joint judgements were often in the neighbourhood of one of the individual judgements, or an uninformed average of the two judgements. Regression analyses suggest that dyads sampled from the alternative responses according to their initial disagreement, and their relative accuracy, on a trial-by-trial basis. Rather than learning about each others' ability, dyad members appear to rely on communication of estimated precision when forming judgements, and often resolve discrepancies by taking the best guess.

Keywords: collective judgement; perceptual judgement; sampling; optimal integration; group decision-making

Averaging versus Sampling in Collaborative Judgement

Final year undergraduate projects in the UK and Australia are often assessed by two independent markers, who then meet to agree a joint final mark. Assuming both markers are estimating some common attribute of the project, how do the markers arrive at an agreed mark? If the two individual marks are 81 and 83, any mark in the range 81–83 would represent a reasonably compromise to most people. However, what strategies are available to markers when there is a wide disparity in the marks? Should the markers average the two marks (perhaps giving greater weight to the marker who is more expert in the area)? Or might the markers recognise some fundamental incompatibility in the marks, and instead return the mark from the marker who is most trusted?

This is one example of a more general problem: how do dyads and small groups resolve discrepancies in their individual judgements so as to arrive at a single joint judgement? Many of our everyday activities are carried out in the presence of others, or are the output of groups working together (see Stasser & Dietz-Uhler, 2001 for an overview). In many cases, groups can be treated as information processors that store, retrieve and combine information to make choices or judgements (Hinsz, Tindale, & Vollrath, 1997). The focus of this paper is on how dyads arrive at joint judgements on continuous dimensions, a question that has been given relatively little empirical attention (Davis, 2014). We first review the general literature on joint judgement, before considering how dyads and small groups could, and do, resolve discrepancies in their individual judgements so as to arrive at a consensus.

Judgements from staticized and interacting groups

Understanding how groups combine information to make judgements first requires understanding how such information could be combined. A common focus in group decision-making has been on the comparative accuracy of staticized and real, interacting groups (e.g., Davis, 2014; Einhorn, Hogarth, & Klempner, 1977; Gigone & Hastie, 1997; Libby, Trotman, & Zimmer, 1987; Rowe & Wright, 1999; Stasser & Dietz-Uhler, 2001). A staticized (or statisticized)

group is composed of individuals who never interact, and who make independent judgements or choices on a common stimulus. Those independent responses can then be combined—most typically, by averaging—to obtain an overall estimate from the group. A common finding is that the average response is more accurate than the individual responses feeding in to that average; the Wisdom of the Crowd (WoC) phenomenon (e.g. Galton, 1907; Surowiecki, 2004). This WoC effect has been observed in judgement tasks such as estimating the number of objects in a container (Dolder & Assem, 2018; Krause, James, Faria, Ruxton, & Krause, 2011) or the weight of an ox (Galton, 1907), as well as discrete choice tasks such as perceptual (signal detection) tasks (Bahrami et al., 2010; Barr & Gold, 2014; Sorkin, Hays, & West, 2001) and face identification (e.g., White, Burton, Kemp, & Jenkins, 2013). While a WoC obtains when the responses of group members are given equal weighting, more accurate group estimates can be obtained by giving greater weighting to members who are relatively more accurate (e.g., Davis-Stober, Budescu, Dana, & Broomell, 2014; Einhorn et al., 1977). The WoC critically relies on the crowd being, on average, well-calibrated, and that members provide independent judgements; if members of the crowd show systematic biases, or are systemically incorrect, the crowd as a whole will show those errors (e.g., Koriati, 2015; Simmons, Nelson, Galak, & Frederick, 2010; for a discussion of moderating factors, see Kramer, MacCoun, & Kerr, 2014).

Against the baseline of staticised groups, a key question is about the judgement and decision-making of real, interacting groups. One of the preconditions for the success of the WoC effect is that judgements are independent (Surowiecki, 2004). This is unrepresentative of most human collaboration—including juries, product development teams, government committees, and academics marking theses—where people communicate with each other to combine their individual judgements and reach a consensus judgement or decision. One general outcome from studies on interacting groups is that real groups rarely outperform staticized groups, and can perform worse than staticized groups (Davis, 1992; Gigone & Hastie, 1997). A number of pressures may work to inhibit performance, including social inhibition (Karau & Williams, 1993), amplification of cognitive biases (Kramer et al., 2014), informational and availability cascades

(Anderson & Holt, 1997; Diehl & Stroebe, 1987), and restricted sharing of information (Stasser & Titus, 1985). Even in cases where individuals make their own judgements, sharing of information can potentially undermine the WoC effect by introducing correlations into individual judgements (Lorenz, Rauhut, Schweitzer, & Helbing, 2011), although individual's estimates may improve in such cases (Farrell, 2011).

While the relative performance of staticized and interacting groups is interesting, here we are primarily concerned with understanding the process by which dyads arrive at a joint judgement, and the characteristics of those joint judgements. Hinsz et al. (1997) highlighted that groups are increasingly required to perform cognitive tasks, and from that perspective can be treated as information processors. Accordingly, a separate body of research has asked how groups encode, store, and retrieve information, combine information to decide and choose, and how groups learn from feedback, with a major focus being on how individual judgements are aggregated into a single choice. A useful example was provided by Sorkin et al. (2001), who examined how groups of people perform a joint signal detection task (see also, e.g., Bahrami et al., 2010; Barr & Gold, 2014). In the individual signal detection task, participants must indicate whether a stimulus was drawn from a signal distribution, or was generated from a noise distribution. Individuals will generally not all observe the same percept, even when presented with the same stimulus, and the stimulus itself might potentially vary across observers. Accordingly, if a group of observers is then required to return a single “signal” versus “noise” response, that group will often need to aggregate across the preferred response of individual members. One way in which a group might do this is to simply take the majority of “signal” versus “noise” votes, termed a Condorcet group (Sorkin et al., 2001; Sorkin, West, & Robinson, 1998). Alternatively, group members might average across some underlying quantity (in the case of signal detection, evidence) to arrive at a decision. Sorkin et al. (2001) and others (e.g., Libby et al., 1987; Snizek & Henry, 1989; Yetton & Bottger, 1982) distinguish between two forms of averaging. *Unweighted* averaging essentially implements the WoC rule. In contrast, groups might give more (or all; Einhorn et al., 1977) weighting to group members with greater expertise—a

weighted average.¹ Sorkin et al. (2001) found evidence of weighted averaging: performance increased as group size increased, and the judgements of more accurate individuals were given higher weighting. However, Sorkin et al. (2001) also found that larger groups were not as effective as expected from an optimal weighting model, and in a second experiment showed that this was attributable to a decrease in the quality of individual judgements with group size (i.e., social loafing), rather than a decrease in the efficiency of aggregating of those individual judgements into a single response.

More recent work, spurred in particular by Bahrami et al. (2010), has analysed the performance of interacting groups (usually dyads) in finer detail than earlier studies. As in Sorkin et al.'s (2001) work, Bahrami et al. (2010) defined optimality as a weighted average model, weighting the judgements from different individuals according to the precision (inverse variance) of each individual. Such optimal weighting has been found in the perceptual decision-making of individual observers combining information from different sources (Alais & Burr, 2004; Pouget, Beck, Ma, & Latham, 2013). Bahrami et al. (2010) presented dyads with a simple two-alternative forced choice (2AFC) perceptual decision-making task, and found that although dyads often performed better than the best member of the dyad, dyad performance was suboptimal when the members of the dyad differed substantially in their sensitivity. To make this inference, Bahrami et al. (2010) fit four models (from which psychophysical functions could be obtained and compared to the data): a coin flip model (i.e., random choice in the case of discrepancy), a model in which the best member's response is always returned as the joint response, an optimal weighting model, and a confidence sharing model. The relatively better fits of this last model implied that participants did not fully communicate information about their perception of the stimulus (as is specified by the optimal model), and instead only provided an estimate of the probability that their binary judgement was correct (i.e., a pronouncement of confidence).

While communication seems to be a key aspect of dyadic judgement, Koriati (2012) suggested that dyads could use simple strategies to resolve disagreements and return a joint

¹ These weights may also be determined by social factors (see Graesser, 1991); we focus on informational factors here.

judgement. Koriat found that a collective benefit could be obtained in absence of dyad interaction by taking the judgement of the most confident dyad member on a trial as the dyad's joint judgement. Koriat (2012) argued that much of the benefit from collaboration comes from the sharing of confidence, and using a "take the best" heuristic to select the response on that trial. Bang et al. (2014) similarly found that this "Maximum Confidence Slating" (MCS) algorithm performed well, and also outperformed models using reaction time as a proxy for confidence. Bang et al. (2014) found that real interacting dyads outperformed the MCS algorithm when the dyad members differed in their perceptual sensitivity, and suggested that interaction mitigates the differences between members in the reliability of judgements. Such results imply that collaborative benefit requires an element of confidence "tuning" via interpersonal interaction, despite an equality bias in opinions of group member's competence (Mahmoodi et al., 2015).

In summary, previous work on binary judgements in dyads and small groups suggests that much of the advantage from forming judgements in groups comes from the information that groups provide, rather than from the interaction within groups (Davis, 1992; Gigone & Hastie, 1997; Koriat, 2012). While some work had suggested that groups average the judgements of their members (Sorkin et al., 2001), other work suggests that discrepancies are resolved by taking the judgement of the most able or most confidence member (Bahrami et al., 2010; Koriat, 2012).

Joint judgements on continuous quantities

Much of the work looking at the construction and optimality of consensus judgements has examined discrete—typically binary—responses. However, many group activities require forming quantitative judgements on a continuous scale, such as making a sales forecast, estimating probability of an event, assessing value, or assigning damages (Ariely et al., 2000; Davis, 2014; Einhorn et al., 1977; Gigone & Hastie, 1997; Hinsz, 1999). Forming such judgements prohibits the use of simple consensus rules, with any number of factions and groupings into factions possible (Davis, 2014), meaning that the models applied to understand groups' discrete choices may not apply to continuous judgements (Hinsz, 1999). Some authors have applied formal descriptive

models to this problem—primarily the Social Decision Scheme model (Davis, 2014; Hinsz, 1999)—but this arguably fails to give any explanation for how people reach consensus, as it simply specifies transition probabilities from distributions of judgements prior to group interaction to the final group judgement.

In considering how groups make quantitative judgements, it is important to consider how groups form their judgements, both to specify alternative theoretical models, and also to implement staticized baselines against which actual performance can be compared (Einhorn et al., 1977; Gigone & Hastie, 1997). Einhorn et al. (1977) identified four such models for consensus formation: a) randomly picking the response of a single individual, irrespective of their expertise; b) unweighted (WoC) averaging; c) returning the judgement of the best member where the best is known with certainty; and d) returning the judgement of the best member even in cases where the best member is only known probabilistically. The random responding can be considered equivalent to the “coin flip” model used by Bahrami et al. (2010), and the last two models map onto the “take the best” heuristic suggested by Koriati (2012), differing in how diagnostic confidence is of superior performance. When calculating the predictions of staticized groups under these different models, Einhorn et al. (1977) found that the “best” models often outperformed the WoC average, partly because the “take the best” strategies are more robust to bias. Einhorn et al. (1977) also found that groups’ actual performance was often best fit by the “best” model, indicating that (at least for small groups) the best member of the group is identified, and the judgement of that member is returned as the groups’ estimate. Snizek and Henry (1989) similarly found that the “best” model performed better than averaging models in predicting joint judgement.

Rather than examining the relative accuracy of small groups, we are interested in characterising the judgements made by groups, and consequently making inferences about the process by which those judgements were made. As mentioned, the work on consensus choice may not generalise to quantitative judgements. In addition, one opportunity offered by quantitative judgements is that the effect of different strategies can be observed at the level of individual trials. For binary choices, if the pair in a dyad disagrees, the returned discrete response

from a dyad tells us little about how the dyad resolved their discrepancy on that particular trial. In contrast, joint judgements lie on the real line, and so the answers returned can be compared to the individual estimates, and probabilistically assigned to different strategies.

We report three experiments in which the responses of individuals and the dyad are on a continuous scale, allowing us to determine to what extent individual dyad responses are obtained as a weighted combination of individual judgements, versus discrete selection from dyad members' individual judgements. Dyad members were presented with random dot kinematograms. Participants were required to individually provide an estimate of direction of coherent movement before being asked to provide a joint response obtained by collaboration. By expressing the jointly agreed answer relative to each individual's answer, we had a measure of the influence of each member of the dyad on the joint decision for each trial. A key difference between this study and previous studies on joint performance of small groups (particularly dyads) is that we used the distribution of judgements across trials, and the relationship of that distribution to different theoretically informed predictors, to answer questions about the process by which dyads come to consensus. Previous work has been limited to inferring the underlying mechanism from the relative accuracy of the group compared to potential models; our approach allows us to more clearly determine how individual estimates are combined and weighted to generate the joint judgement on each trial.

In the experiments, we ask two key questions about how people reach and report a consensus judgement on continuous quantities. The first question is whether dyads' behaviour is best characterised by an averaging strategy, or whether the estimate of only one dyad member is returned as the joint response. On the one hand, the work of Sorkin et al. (2001), Bahrami et al. (2010) and others in the discrete domain suggests that the dyads should perform some form of averaging. On the other hand, the results just reviewed from the discrete domain (Koriat, 2012) and continuous domains of judgement (Einhorn et al., 1977) suggest that dyads may instead go with the estimate of a single member.

The second question is: what determines the weighting given to the dyad members in

determining the consensus judgement? One obvious candidate is the relative ability of the members, akin to member expertise (Einhorn et al., 1977; Gigone & Hastie, 1997; Sorkin et al., 2001; Yetton & Bottger, 1982). Under averaging (and under the assumption that the percepts of the two members are independent draws from a common latent state), the optimal solution is to take a weighted average of the individual members' estimates, the weighting being determined by the relative precision (inverse variance) of the two members (Alais & Burr, 2004). This means that the judgement of the more accurate member should carry a heavier weighting in the final judgement. Such considerations also apply to the case where the judgement of a single member is returned as the joint response. Is the choice between members a coin flip (Bahrami et al., 2010; Einhorn et al., 1977), or is the judgement of the more able and accurate member returned as the joint judgement (Einhorn et al., 1977)? Previous work suggests that dyads and small groups both should (Einhorn et al., 1977) and do (Libby et al., 1987; Yetton & Bottger, 1982) identify their best members and give those members more weighting in the group consensus. Another key factor is the relative disparity between members in their ability, though the evidence here is more mixed: while some work (expertise of Australian loan officers; Libby et al., 1987) suggests that groups will perform better when there is a disparity in ability (such that interaction allows groups to identify their best members and give them greater weighting), other work suggests that variance in ability is, if anything, detrimental to performance of a dyad (Bahrami et al., 2010; Bang et al., 2014). Here, we address how dyad members take their respective ability into account in whatever process they use to form joint continuous judgements.

In our experiments we asked each dyad to complete a number of judgement trials, meaning that we can also ask about any sequential effects or strategies in the joint judgement, and more clearly attribute judgements to random error versus systematic influences (Gigone & Hastie, 1997). If a dyad is learning about the relative ability of its members, the running average of the performance of the two members up to trial t should be predictive of the relative weighting of the two members' judgements on trial t . However, it may be that other information is used to determine weighting. Members may simply communicate their confidence on each trial, and use

their relative confidence to determine weighting (Bang et al., 2014; Koriat, 2012). While we did not record confidence judgements, we asked how well the accuracy of the two members on each trial—the relative error—predicted the weighting. Dyads may also use a strategy somewhere in-between, and track the relative error over the past several trials.

Experiment 1

Method

Participants. 56 participants (38 female, mean age 19.32 years, SD 1.32) were recruited from the University of Bristol undergraduate population, and earned credits as a course requirement. Participants were assigned to 28 dyads by self-selecting a timeslot on an experiment scheduling website. Any time slots with a single participant signed up were cancelled and the participant offered a choice of any other times with a single participant already signed up, or asked to select an alternative time.

Procedure. Participants were asked to complete a two-part personality inventory and a 10 minute practice session with the stimuli before going on to the main part of the study, the joint judgement task. Given the social nature of collaboration and the potential of personality factors to influence it (Graesser, 1991; Wolf & Krause, 2014), a number of personal factors, such as age, gender and personality traits were recorded to explore whether these could impact the decision process. For example, agreeableness has been shown to promote cooperative behaviour (Zhao & Smillie, 2015) and could result in individuals with high scores in this trait subsuming accuracy in favour of cooperation and gender, specifically levels of testosterone, could have the opposite consequence (Wright et al., 2012). We considered that these analyses are likely to be uninformative given the low power to examine individual differences, and so treated them as exploratory. In unreported analyses these variables were included as predictors on the elastic net regression (see below), and none of them predicted dyad performance.

Personality inventory.

This was a two part, self-reported questionnaire administered on a PC. The measures taken

were the “Big Five” attributes of Agreeableness, Extraversion, Conscientiousness, Intellect/Imagination and Emotional Stability (Goldberg, 1990) and a Prestige/Dominance Inventory (Cheng, Tracy, & Henrich, 2010).

Calibration phase.

Motion coherence stimuli were used as stimuli in the judgement task. As participants were assumed to be naïve in their exposure to this type of stimuli, they were given a calibration session lasting approximately 10 minutes. This phase served to give participants some experience with the task, and to obtain an a priori measure of performance. Participants were asked to try to maintain a constant distance from the monitor. Before starting the practice, they were asked to hold a piece of string, 65 cm in length, attached to the top right corner of the monitor taut to the corner of their right eye and position themselves comfortably so as to be able to maintain that distance for the duration of the practice.

Stimuli were generated in MATLAB. One hundred moving black dots were displayed in a circular aperture on a white background. Coherence was manipulated by varying the number of dots moving in a coherent direction toward the outer edge of the aperture. The coherences used were 5%, 10%, 15%, 20%, 30%, 35%, 40% and 50%. The direction of coherent dot motion could be in any of 360 degrees and was selected uniformly at random for each trial.

At the start of each trial, participants were presented with a fixation cross in the centre of the screen for 1 second. The stimulus was then presented for 1 second, followed by presentation of a circle and a pointer (as in Panel 2 of Figure 1). The initial placement of the pointer was pseudo-randomly determined for each trial. Participants used the left and right cursor key to move the pointer clockwise and anticlockwise to indicate the direction that they believed the coherently moving dots to be moving. Pressing the space bar confirmed the selection and moved on to the next trial.

Joint Judgement Task.

Participants were positioned side by side with a partition positioned to ensure that dyad members could not see each others' monitors, but such that they could make eye contact with

each other. Each monitor was attached to the same PC, which controlled the experiment.

Figure 1 shows the flow of a single trial in the joint judgement task. As in the practice phase, on each trial participants were shown a motion coherence stimulus (Figure 1, Panel 1) and individually required to input the direction of coherent dot motion without any discussion (Figure 1, Panel 2). The coherence for all trials was set at 35% based on a pilot of the stimuli, so that most participants would be able to perceive some coherent motion. In response to the stimulus, participants first provided an individual response; the pointer starting position was randomly selected, but the same for both participants. One participant provided their answer using the left and right cursor key on the keyboard, as in the practice. The other participant used the left and right mouse button to move the dial anti-clockwise and clockwise respectively, and the central mouse button to confirm their selection.

Once both participants had returned an individual response, participants were asked to return a joint judgement (Figure 1, Panel 3). Both individuals were shown the same answer dial; this dial displayed the participants' individual answers, along with an adjustable dial to be used to return the dyad's joint judgement. The starting position of the pointer for the joint judgement was pseudo-randomly determined by the computer. Participants were allowed free discussion, and both participants had control over the response dial via their respective input devices. When dyad members had arrived at a decision on which they were both agreed, they submitted their joint judgement by pressing their individual "confirm" buttons (in either order) with no other intervening button or key presses. If any additional key presses were detected, the confirmation process was reset. After confirming the joint judgement, the dyad was given feedback: The correct direction was indicated by a green line on the answer dial and displayed for 1 s before moving on to the next trial (Figure 1).

Each dyad completed 2 blocks of 20 trials (40 in total), swapping input devices between blocks. The procedure took roughly 45 minutes. At the end of the session they were debriefed and thanked for their participation.

Results

We first examined the accuracy of the individual and joint responses. The top panel of Figure 2 shows observed accuracy, and predicted accuracy under several possible models to be described shortly. The left two bars of Figure 2 respectively show the performance of the individual participants relative to the dyad (“Individual Responses”, the mean of the individual answers’ distance from the actual answer, the latter being the direction used to generate the coherent motion for that trial) and the jointly agreed answer’s distance from the actual answer (“Joint Responses”). Figure 2 shows a benefit to collaboration, with the joint estimates being more accurate on average.

Three models of potential strategies that participants might have adopted were also applied to the data on a trial by trial basis. “Take the Best” (TTB) takes the most accurate individual response on each trial as the dyad response. This is an optimistic model, as it assumes that dyad members have perfect knowledge of which response was the most accurate on each trial. This model this is analogous to Bahrami et al.’s (2010) confidence sharing model (see also Koriat, 2012), which assumes that participants communicate an answer and an estimate of a probability of being correct and that one of the answers is selected in a binary fashion based on stated confidence. Here, given that confidence cannot be measured directly, accuracy is used as a proxy and the answer from the participant closest to the truth is taken. The paradigm also allowed participants to combine their individual judgements in a Bayesian manner by weighting by their relative accuracy in performing the task (Sorkin et al., 2001). Specifically, Bayesian Weighting (BW) takes the weighted mean of the two individual judgements on a given trial, where the weights are determined by the precision (1/variance), of each individual’s overall performance:

$$w_A = \frac{\frac{1}{\sigma_A^2}}{\frac{1}{\sigma_A^2} + \frac{1}{\sigma_B^2}}$$

A third possible strategy assumed that people naively averaged their individual responses without reference to their relative accuracy (Wisdom of Crowds, or WoC).

For Experiment 1, Panel A of Figure 2 shows that the joint responses achieved a level of

accuracy comparable to a BW strategy. The dyads did not appear to be performing as well as the TTB strategy, but did clearly outperform the WoC strategy, which assumes naïve averaging of responses.

To characterise the joint estimates, a measure of the relative influence of the two dyad members' judgements on the joint response was determined. For each trial, a joint proportion (JP) was calculated as the angular distance from Participant A's answer to the joint answer divided by the angular distance from Participant A's answer to Participant B's answer. In most cases (i.e. where the joint answer lies between the individual responses), this value ranges between 0 and 1. In cases where the joint answer does not lie on the shortest arc between the participant's individual answers, the JP will be either negative or larger than 1 (see Supplemental Materials for details on these calculations). Figure 3 (Panel A) shows histograms of Joint Proportions, collapsing across all trials and all participants. The histogram shows evidence of discretization in responses, with participants generally opting to provide one of the individual judgements as the dyad's joint response, or an answer between the two individual responses (middle peak).

Discussion

On aggregate, the dyad's answers were more accurate than the individual answers. The distributions of joint responses (Figure 3) suggests a mixture of a bimodal distribution consistent with Koriatis (2012) results, and weighted averaging (Bahrami et al., 2010; Sorkin et al., 2001). Formalisation of these distributions through Gaussian mixture modelling is covered in the Results of Experiment 3. First, we address a concern that the distinct peaks are due to the dyad members anchoring (e.g., Tversky & Kahneman, 1974) on the lines provided as indicators of the individual responses. To address this concern, the experiment was repeated, but without showing the participants their initial judgements in the joint judgement phase.

Experiment 2

Method

Participants. Thirty-eight participants (22 female, mean age 26.29 years, SD 10.01) were recruited using the University of Bristol's pool of paid participants and an advertisement in a local paper, and reimbursed £7 for their time.

Procedure. This was as for Experiment 1, with the exception that the calibration phase was excluded so as to shorten the experiment the practice phase was omitted. Additionally, the individual judgements were not shown on the dial at any point during the joint decision-making phase (cf Panel 3 of Figure 1).

Results

Panel B of Figure 2 shows the mean accuracy of the individual and joint responses, along with model predictions. As in Experiment 1, the Joint Responses are on average more accurate than those individually given, showing again a benefit of collaboration. Although the experimental and modelled results show the same pattern, overall performance appeared to be poorer than in Experiment 1.

The distribution of joint responses is plotted in Panel B of Figure 3. The distribution shows pronounced peaks at 0 and 1. Those peaks—and the peak at a joint proportion of 0.5—are less obvious than in Experiment 1.

Discussion

There are two obvious reasons for the less pronounced peaks in Experiment 2. First, by removing the practice phase with the stimuli, the opportunity to acclimatise to the task was also removed. Second, by no longer displaying the participants' initial answers, we forced participants to rely on verbal descriptions or positioning the pointer from memory, meaning that the remembered individual estimates—and thus the joint estimates—were subject to greater variability.

One further issue is that in the first two experiments, participants may have been simply using a heuristic that prioritised speed or effort to arrive at a joint judgement as there was no

incentive for accuracy. Accordingly, Experiment 3 introduced a monetary bonus contingent on the accuracy of the joint judgements.

Experiment 3

Method

Participants. 58 participants (29 dyads) were recruited using the University of Bristol's pool of paid participants and an advertisement in a local paper. However, due to a computer problem during the experiment, one dyad was unable to complete all 40 trials, so was excluded. The remaining 56 participants are included in the analysis (34 female, mean age 24.17 years, SD 10.64). Participants were reimbursed £7 for their time and could also gain a performance related bonus. Participants were informed that five of their joint answers would be randomly selected and used to determine a bonus of up to £5 depending on how close their answers were from correct. The reward was calculated as the sum of $\pounds 1 - (\text{distance from actual answer} / 2)^2$ per selected trial. Only 2 participants specifically asked how the reward was derived, so they were shown this calculation during the debriefing.

Procedure. The procedure was as for Experiment 2, but with the reintroduction of a practice phase to ensure participants performed suitably well on the task. For the practice, the participants were first given 8 high coherence trials (100%, 95%, 90%, 85%, 80%, 75%, 70%, 65%) to ensure they understood the task before starting the main practice. The rest of the session was as described in Experiment 1. The practice trials were not included in any analyses.

The second change from Experiment 2 was that participants were informed of the incentive to perform well, and were shown their earned reward on screen at the end of the experiment.

Results

Consistent with the first two experiments, mean distance from actual answer (Figure 2, Panel C) showed a marked improvement for the joint responses over the individual answers. There was also a similar pattern to the three alternative models as seen in the previous two

studies.

Panel C of Figure 3 shows a histogram of joint responses. The histogram is similar to that from Experiment 2, with peaks around 0 and 1 in addition to a range of answers between the individual answers. Despite having removed the possibility of anchoring and by incentivising the dyad to perform as well as possible, participants still showed a pattern of often returning (an approximation of) one individual response as the joint response.

To characterise the apparent clusters in Figure 3, and to determine predictors cluster membership, a Gaussian mixture model (GMM) was fit to the joint proportions for each experiment using the Netlab toolbox (Nabney, 2002). For each of the three experiments, no assumptions about the centres of the Gaussians were made. Instead, a value between zero and one was randomly selected as a centre to seed the GMM algorithm (which initially assumes a unit variance Gaussian). To avoid local minima, the routine was run 100 times (each with its own starting point) and the replication with the lowest deviance was taken as the solution. This process was repeated assuming between 2 and 8 components of the mixture. In an initial phase, we used the Bayesian Information Criterion (BIC; for application to mixture models, see Keribin, 2000; Steele & Raftery, 2010) to select the number of components in the mixture. However, this procedure was found to select multiple distributions with the same mean at the end endpoints for all three experiments. We instead adopted a conceptually simpler model with four components in Experiment 1, and three in Experiments 2 and 3 (as described below). It should be noted that the BIC solution suggests that the distributions may be heavier-tailed than the Gaussians we used here.

Four Gaussians were identified for Experiment 1, and three for Experiments 2 and 3. Figure 3 (right hand column) shows the fitted GMM for each experiment. We interpreted the clusters as representing: returning a joint response anchored on Participant A's initial answer (Participant A); returning a joint answer around Participant B's initial answer (Participant B); and selecting a point somewhere in between (Middle). A fourth distribution was estimated for Experiment 1, centered on the mid-point and being less variable than the "middle" distribution;

henceforth this is referred to as the “WoC” (naïve averaging) distribution.

One output of the GMM is the estimated posterior probability that each joint proportion (JP) belonged to each of the categories. These scores were discretized by assigning each JP to the category under which it had the highest probability of membership. This nominal category membership was used as the dependent variable in a logistic regression. Numerous predictors were included, such as distance of each individual response from the true direction of the presented stimulus, and reaction time of each individual to submit their initial response. Also included were measures capturing a reliance on information from previous trials, for example category chosen by the dyad on the previous trial and cumulative mean accuracy of each participant. Interpersonal measures were derived from the demographic data supplied by participants, their scores in the Personality Inventory and in the Prestige/Dominance Inventory and from their performance in the joint decision making task. The complete list is given in Table 1 in the Supplemental Materials.

The predictors, along with the outcome variable of JP category, were entered into a multinomial logistic regression (Friedman, Hastie, & Tibshirani, 2001), using the glmnet package (Friedman, Hastie, & Tibshirani, 2010; Simon, Friedman, Hastie, & Tibshirani, 2011). This technique entails fitting a generalized linear model via penalized maximum likelihood. A penalty is imposed on the regression coefficients, resulting in efficient variable selection alongside shrinkage of the retained regression coefficients (i.e., the estimates are pulled towards zero). We ran this procedure using ridge regression (which tends to extract multiple correlated predictors) and lasso regression (which tends to select only one of a number of correlated predictors) and report both regressions below.

To calculate confidence intervals, we used a bootstrapping procedure (Efron & Tibshirani, 1994) with 1000 bootstrap samples. For each bootstrap sample, the regression was applied, then the glmnet cross validation routine was used to identify the best lambda value (the value determining the size of the penalty that in turn down weights coefficient estimates) and this was used to extract the coefficients for predictors for that bootstrap sample. To minimise any

incidental confounds relating to which participants happened to be designated as “A” and which as “B”, the designation was flipped independently for each dyad with a probability of 0.5 for each of the bootstrap samples.

The resulting sampling distributions were used to construct 95% confidence intervals for both ridge regression and lasso regression. Table 1 to Table 3 show the coefficients of the predictors whose CI did not include 0 for the ridge regression. In all three experiments, the distance between the participants’ initial answers was a predictor of whether a joint judgement lying somewhere between the two individual judgements (the “middle” category) was returned. The sign of the coefficient indicates that as the size of the difference became larger, the dyads were less likely to return a joint answer from this distribution.

In addition, Relative Error was identified as a predictor. Relative error is a measure of the relative accuracy of the two dyad members in their individual judgements on a given trial, calculated as

$$RE = \log \frac{\Delta A^2}{\Delta B^2}$$

where ΔA^2 and ΔB^2 are the distance of the participant A and B’s answer from the actual answer squared. Negative relative error values indicate that the participant designated as “A” is the more accurate, while positive values indicate participant “B” is the more accurate. The sign of the coefficient for relative error shows that as the relative error became more negative (i.e., A was more accurate), the likelihood increased that A’s individual judgement was returned as the joint judgement. Conversely, as relative error became more positive, the likelihood increased that B’s individual judgement was used as a basis for the joint response.

Several other predictors came out in the ridge regression, including several historical factors (predictors relating to judgements on previous trials), and the time taken to make a joint judgement. However, these other predictors dropped out when the lasso regression—which favours selection of only one of a set of correlated predictors—was applied. Under lasso regression (Table 4), only the predictors just discussed—relative error and distance between initial

responses—were reliably selected. In addition, for Experiment 1 the joint response time was also selected (95% CI: 0.055-0.137).

In summary, the dyad was more likely to select a joint answer in between the individual estimates when the discrepancy between those estimates is small; but when the discrepancy was large, the dyad was more likely to respond with an answer approximating the individual judgement of the participant who was more accurate on that trial. Of course, who is more accurate is not directly available to the dyad members at the point of making the decision (the correct answer feedback is presented briefly at the end of the trial), so we infer there is a relationship between accuracy, awareness of that accuracy and what is communicated between the pair as their confidence in their individual judgements.

The possibility still remains that participants could be using weighting of confidence in the “middle category”. Figure 4 shows the how joint proportion varies with relative error for those points classified as belonging to the middle category. There was a positive relationship in all experiments (linear robust regression using the `robustbase` package, (Maechler et al., 2018; Todorov & Filzmoser, 2009); E1: $\beta = 0.05$, $R^2 = 0.29$, $p < .001$; E2: $\beta = 0.01$, $R^2 = 0.04$, $p = 0.001$; E3: $\beta = 0.02$, $R^2 = 0.07$, $p < .001$), consistent with weighting by relative performance. This suggests that on some trials, participants are aware of their accuracy on that trial and arrive at an average response weighted by relative accuracy. However, the regression fits indicate only a weak relationship, especially for Experiments 2 and 3.

Discussion

Even when incentivised, participants appeared not to primarily rely on weighted averaging to arrive at joint decisions. Rather, dyads appear to use simpler rules: a) if the difference between the individual answers is relatively small, choose a point between the two that is partially weighted by the relative accuracy of each dyad member; and b) if the difference is larger, then choose a point near the more accurate participant. This leads to the question of how the dyad determines who is the more accurate participant. Although some historical factors were

suggested by the ridge regression, the more selective lasso regression primarily indicated the involvement of factors specific to a trial.

General Discussion

Consistent with previous studies of discrete joint judgement (e.g., Bahrami et al., 2010; Sorkin et al., 2001), in a continuous judgement task we observed a clear benefit of judgement in groups compared to those made individually. While the dyad's collaborative judgements clearly outperformed naïve WoC averaging, more stringent baselines are provided by allowing greater weighting of the judgement of more able members (Einhorn et al., 1977). Joint judgements were, on average, as accurate as those predicted from a Bayesian model in which each member's judgement was weighted by their relative precision. However, modelling of individual trial responses suggests that Bayesian averaging was only employed in cases where the individual judgements were relatively close, and this relationship was only weak across the three experiments.

Examination of the distribution of the joint judgements—scaled to an interval defined by the two members' individual estimates on each trial—revealed clusters of responses indicating that the dyads often returned the estimate of one of the members as the joint responses (or at least, an approximation thereof). The regression (using penalised maximum likelihood estimation) revealed that the relative accuracy of the two members on a trial predicted which of the two individual estimates was returned. This is consistent with suggestions that groups follow a “best member” strategy when making judgements (e.g., Einhorn et al., 1977; Koriat, 2012; Snizek & Henry, 1989; Yetton & Bottger, 1982). An advantage of our study is that joint response for each trial could be assigned to one member versus the other (using mixture modelling). Rather than relying solely on the comparison of the average accuracy between real and staticized dyads, we explicitly show the sampling of individual estimates by examining the distribution of trial-level joint estimates.

The regression also suggested that the only predictor related to accuracy was the relative

accuracy on that trial; predictors relating to historical accuracy—whether measured by the influence of the last few trials preceding a judgement, or the cumulative accuracy up to a trial—were not identified as substantial and unique in the penalised regression. This result rules out one possible model of repeated interaction, which is that dyad members built up a model of each others' relative accuracy, and used this to identify the “best” member. Instead, dyads relied on trial-by-trial information about their relative accuracy to identify the best member. Since the joint judgement was given prior to any feedback about the accuracy of individual or joint estimates, this implies that the dyad members had access to metacognitive information on their precision on that trial, and were able to communicate this to their partner so as to reliably choose the best member. This is consistent with suggestions that participants communicate their confidence on each trial and use this as a basis for arbitrating between discrepant judgements (e.g., Bang et al., 2014; Koriat, 2012, 2015), and the greater convergence on a common terminology for confidence seen in more successful dyads (Fusaroli et al., 2012). In particular, the results are consistent with Bang et al.'s (2017) evidence that dyad members calibrate their confidence statements to each other. Bang et al. (2017) asked their participants to make private judgements and confidence ratings in a binary perceptual choice task, with the response of the most confident member on a trial automatically being returned as the joint response for that trial. Bang et al. (2017) found that the average difference in confidence ratings between the two dyad members was not related to the difference in average accuracy between the members: the more accurate member did not tend to return higher confidence ratings. Instead, dyad members returned statements of high and low confidence with roughly equal frequency. Such a strategy is consistent with the results from our judgement task, where trial-by-trial (rather than average) accuracy predicted whose response was returned as the best. An interesting question for future research is how confidence ratings relate to different aspects of continuous judgements, and how those confidence judgements are used to weight or select response.

While the joint response was often drawn from a cluster corresponding to one of the individual member's judgement, Figure 3 shows a group of responses somewhere in-between the

two outermost clusters. This was most pronounced in Experiment 1—where the individual responses were displayed on the screen during the joint judgement formation—but was also seen in Experiments 2 and 3. Regression of these responses against relative member ability revealed a tendency for the more accurate member to be given greater weighting in this cluster of judgements, consistent with an optimal weighting strategy. However, this relationship was quite weak, and it is not clear that a weighted strategy was always, or even primarily, used. One explanation for this weak effect is that dyad members were able to identify the most accurate member on a trial, but have difficulty quantifying the relative accuracy (presumably communicated via a confidence statement) needed to take a weighted average.

The penalised regression also identified situations in which the dyads returned the response of a single member, versus returning a point somewhere in between. When the individual judgements were relatively close, the dyad was likely to take a (weighted) average, while more discrepant initial judgements predicted a response being returned from just one of the members. This is a sensible strategy under the assumption that members' individual percepts might be sampled from different latent states. Consider the case where two individuals living in Perth are asked to estimate the highest temperature the previous day, and the two estimates are 6°C and 36°C. A naïve WoC aggregator might simply average these estimates, but a more likely explanation is that these observations have not been sampled from the same city (e.g., the estimates respectively come from Perth, Scotland and Perth, Western Australia). In the context of ventriloquist effect, for which a Bayesian optimal weighting account has been offered (Alais & Burr, 2004), Körding et al. (2007) and Sato, Toyoizumi, and Aihara (2007) suggested that when perceiving a multidimensional cue, observers use the discrepancy between the cues to infer whether the individual sensory cues were generated by the same cause, or by different causes. If cues are perceived as coming from a common cause, they are integrated (averaged), while an inference of separate causes leads to segregation of the cues. Our dyads appeared to segregate (vs average) on the basis of the discrepancy in individual judgements, and thus appear to make predictions based on a causal understanding similar to that identified by Körding et al. (2007) and

Sato et al. (2007). Differences in causes might arise from the random dot kinematograms we used: an observer may not perceive any coherent motion, or simply not pay attention on a trial.

However, if one is willing to make some additional assumptions about how the response from a dyad is sampled, the data may also be accounted for by another Bayesian model that assumes continuous—rather than discrete—uncertainty gives rise to bimodality in responses. Lindley (1983) asked how an individual might combine estimates from different sources (i.e., other individuals, particularly experts) where each source communicates a measure of central tendency (e.g., mean) and spread (e.g., variance, or confidence). If there is uncertainty in how well each source estimates its spread, then under reasonable assumptions the distributions of the estimate for each source follow a t -distribution. In the case of two sources—corresponding to each individual in the dyad—and assuming the sources are independent, the posterior distribution will be bimodal as long as the difference in initial mean estimates is sufficiently large.

Figure 5 illustrates this graphically. The vertical lines in each panel show the two individual estimates, and the line plots the density across the dimension to be judged. In all cases, the density of the joint judgement is simply obtained by multiplying the densities of the two individual estimates (Lindley, 1983). The top row shows cases where the two individual estimates are assumed to be normally distributed (the individual distributions are not shown beyond the lines indicating central tendency), such that each observer has some uncertainty about their estimate, but knows exactly their best estimate (the mean) and the precision on their estimate (inverse variance). The top row shows that irrespective of the initial discrepancy in the estimates, the posterior distribution is unimodal. In contrast, the bottom row shows what happens when the individual estimates are described by a t distribution. A t distribution results when observers have uncertainty about their estimate—described by a normal distribution—but also have uncertainty about their precision. In our experiments, the uncertainty about precision might arise from metacognitive uncertainty, or the low fidelity of confidence statements. When the two initial estimates are close together, a unimodal distribution results. However, as the discrepancy increases (left to right column), the resulting posterior distribution becomes more clearly bimodal.

Lindley's (1983) model would explain the clustering seen in distributions seen in Figure 3, along with the finding that bimodality is more likely to be seen when the initial estimates are discrepant. However, to actually observe the posterior in our recorded judgements requires the additional assumption that the dyad's response is a single sample from the distribution; reporting the mean of the posterior distribution will not produce the observed patterns. This interpretation of collaborative decision making as sampling from a joint posterior is in keeping with a growing body of evidence that supports sampling as a common principle of human judgement and decision making (e.g., Lloyd & Leslie, 2013; Sanborn, Griffiths, & Navarro, 2010; Vul, Goodman, Griffiths, & Tenenbaum, 2014). May, Korda, Lee, and Leslie (2012), for example, showed that a sampling strategy can result in optimal long term behaviour when an individual must explore their environment while learning to act optimally. While compelling, a key question is whether members can (and do) estimate and communicate their estimates and estimated precision in a way that would allow the dyad to collectively construct a posterior and sample from it.

One question that is left for future research is how these results scale up to larger groups. For example, measures of central tendency other than the average may be used, such as the median or the middle of the range of scores (Davis, 2014; Hinsz, 1999). Larger groups would also give the opportunity to observe potential formation of coalitions and factions (Graesser, 1991), and a question in that case is whether discrepancies between factions are resolved by averaging, or by taking the estimate of a single faction (cf Sorkin et al., 2001). We examined dyads partly for reasons of practicality, but also because interacting groups making important decisions—especially where the decision-makers are stakeholders in the discussion—are dyads. The opening example of academic markers is one example, and our results accord with a protocol implicitly (and sometimes explicitly) adopted by markers: if the discrepancy is small, average the marks, and if the discrepancy is large, return only one of the marks (the marker having greater expertise in the topic), or solicit the input of a third party.

References

- Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*, 14, 257–262.
- Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, 847–862.
- Ariely, D., Tung Au, W., Bender, R. H., Budescu, D. V., Dietz, C. B., Gu, H., ... Zauberman, G. (2000). The effects of averaging subjective probability estimates between and within judges. *Journal of Experimental Psychology: Applied*, 6(2), 130.
- Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science*, 329(5995), 1081–1085. doi:10.1126/science.1185718
- Bang, D., Aitchison, L., Moran, R., Castanon, S. H., Rafiee, B., Mahmoodi, A., ... Summerfield, C. (2017). Confidence matching in group decision-making. *Nature Human Behaviour*, 1(6), 0117.
- Bang, D., Fusaroli, R., Tylén, K., Olsen, K., Latham, P. E., Lau, J. Y., ... Bahrami, B. (2014). Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and Cognition*, 26, 13–23. doi:10.1016/j.concog.2014.02.002
- Barr, S., & Gold, J. M. (2014). Redundant visual information enhances group decisions. *Journal of Experimental Psychology: Human Perception and Performance*, 40(6), 2124.
- Cheng, J. T., Tracy, J. L., & Henrich, J. (2010). Pride, personality, and the evolutionary foundations of human social status. *Evolution and Human Behavior*, 31(5), 334–347.
- Davis, J. H. (1992). Some compelling intuitions about group consensus decisions, theoretical and empirical research, and interpersonal aggregation phenomena: Selected examples 1950–1990. *Organizational Behavior and Human Decision Processes*, 52(1), 3–38.
- Davis, J. H. (2014). Group decision making and quantitative judgments: A consensus model. In *Understanding group behavior* (pp. 43–68). Psychology Press.

Davis-Stober, C. P., Budescu, D. V., Dana, J., & Broomell, S. B. (2014). When is a crowd wise?

Decision, 1(2), 79.

Diehl, M., & Stroebe, W. (1987). Productivity loss in brainstorming groups: Toward the solution of

a riddle. *Journal of Personality and Social Psychology, 53*(3), 497.

Dolder, D. van, & Assem, M. J. van den. (2018). The wisdom of the inner crowd in three large

natural experiments. *Nature Human Behaviour, 2*(1), 21.

Efron, B., & Tibshirani, R. (1994). *Introduction to the bootstrap*. New York: Chapman & Hall.

Einhorn, H. J., Hogarth, R. M., & Klempner, E. (1977). Quality of group judgment. *Psychological*

Bulletin, 84(1), 158.

Farrell, S. (2011). Social influence benefits the wisdom of individuals in the crowd [Letter to the

editor]. *Proceedings of the National Academy of Sciences, 108*, E625.

Friedman, J., Hastie, T., & Tibshirani, R. (2001). *The elements of statistical learning*. Springer: New York.

Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models

via coordinate descent. *Journal of Statistical Software, 33*(1), 1–22. Retrieved from

<http://www.jstatsoft.org/v33/i01/>

Fusaroli, R., Bahrami, B., Olsen, K., Roepstorff, A., Rees, G., Frith, C., & Tylen, K. (2012). Coming

to terms: Quantifying the benefits of linguistic coordination. *Psychological Science, 23*(8),

931–939. doi:10.1177/0956797612436816

Galton, F. (1907). Vox populi (the wisdom of crowds). *Nature, 75*(7), 450–451.

Gigone, D., & Hastie, R. (1997). Proper analysis of the accuracy of group judgments. *Psychological*

Bulletin, 121(1), 149.

Goldberg, L. R. (1990). An alternative" description of personality": The big-five factor structure.

Journal of Personality and Social Psychology, 59(6), 1216.

Graesser, C. C. (1991). A social averaging theorem for group decision making. In N. H. Anderson

(Ed.), *Contributions to information integration theory* (Vol. Volume II: Social). Hillsdale, N.

J.: Lawrence Erlbaum Associates.

- Hinsz, V. B. (1999). Group decision making with responses of a quantitative nature: The theory of social decision schemes for quantities. *Organizational Behavior and Human Decision Processes*, 80(1), 28–49.
- Hinsz, V. B., Tindale, R. S., & Vollrath, D. A. (1997). The emerging conceptualization of groups as information processors. *Psychological Bulletin*, 121(1), 43.
- Karau, S. J., & Williams, K. D. (1993). Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, 65(4), 681.
- Keribin, C. (2000). Consistent estimation of the order of mixture models. *Sankhyā: The Indian Journal of Statistics, Series A*, 49–66.
- Koriat, A. (2012). When are two heads better than one and why? *Science*, 336(6079), 360–362.
doi:10.1126/science.1216549
- Koriat, A. (2015). When two heads are better than one and when they can be worse: The amplification hypothesis. *Journal of Experimental Psychology: General*, 144(5), 934.
- Körding, K. P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., & Shams, L. (2007). Causal inference in multisensory perception. *PLoS One*, 2(9), e943.
- Kramer, G. P., MacCoun, R. J., & Kerr, N. L. (2014). “When are n heads better (or worse) than one?”: Biased judgment in individuals versus groups. In *Understanding group behavior* (pp. 113–144). Psychology Press.
- Krause, S., James, R., Faria, J. J., Ruxton, G. D., & Krause, J. (2011). Swarm intelligence in humans: Diversity can trump ability. *Animal Behaviour*, 81(5), 941–948.
- Libby, R., Trotman, K. T., & Zimmer, I. (1987). Member variation, recognition of expertise, and group performance. *Journal of Applied Psychology*, 72(1), 81.
- Lindley, D. (1983). Reconciliation of probability distributions. *Operations Research*, 31(5), 866–880.
doi:10.1287/opre.31.5.866
- Lloyd, K., & Leslie, D. S. (2013). Context-dependent decision-making: A simple bayesian model. *Journal of the Royal Society Interface*, 10(82), 20130069.
- 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, 9020–9025.
- Maechler, M., Rousseeuw, P., Croux, C., Todorov, V., Ruckstuhl, A., Salibian-Barrera, M., ... Anna di Palma, M. (2018). *Robustbase: Basic robust statistics*. Retrieved from <http://robustbase.r-forge.r-project.org/>
- Mahmoodi, A., Bang, D., Olsen, K., Zhao, Y. A., Shi, Z., Broberg, K., ... Bahrami, B. (2015). Equality bias impairs collective decision-making across cultures. *Proc Natl Acad Sci USA*, 201421692. doi:10.1073/pnas.1421692112
- May, B. C., Korda, N., Lee, A., & Leslie, D. S. (2012). Optimistic bayesian sampling in contextual-bandit problems. *Journal of Machine Learning Research*, 13(Jun), 2069–2106.
- Nabney, I. (2002). *NETLAB: Algorithms for pattern recognition*. Springer Science & Business Media.
- Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, 16(9), 1170.
- Rowe, G., & Wright, G. (1999). The delphi technique as a forecasting tool: Issues and analysis. *International Journal of Forecasting*, 15(4), 353–375. doi:[https://doi.org/10.1016/S0169-2070\(99\)00018-7](https://doi.org/10.1016/S0169-2070(99)00018-7)
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, 117(4), 1144.
- Sato, Y., Toyoizumi, T., & Aihara, K. (2007). Bayesian inference explains perception of unity and ventriloquism aftereffect: Identification of common sources of audiovisual stimuli. *Neural Computation*, 19(12), 3335–3355.
- Simmons, J. P., Nelson, L. D., Galak, J., & Frederick, S. (2010). Intuitive biases in choice versus estimation: Implications for the wisdom of crowds. *Journal of Consumer Research*, 38(1), 1–15.
- Simon, N., Friedman, J., Hastie, T., & Tibshirani, R. (2011). Regularization paths for cox's proportional hazards model via coordinate descent. *Journal of Statistical Software*, 39(5), 1–13. Retrieved from <http://www.jstatsoft.org/v39/i05/>

- Snizek, J. A., & Henry, R. A. (1989). Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes*, 43(1), 1–28.
doi:[https://doi.org/10.1016/0749-5978\(89\)90055-1](https://doi.org/10.1016/0749-5978(89)90055-1)
- Sorkin, R. D., Hays, C. J., & West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review*, 108(1), 183–203. doi:10.1037/0033-295x.108.1.183
- Sorkin, R. D., West, R., & Robinson, D. E. (1998). Group performance depends on the majority rule. *Psychological Science*, 9(6), 456–463.
- Stasser, G., & Dietz-Uhler, B. (2001). Collective choice, judgment, and problem solving. In M. A. Hogg & R. S. Tindale (Eds.), *Blackwell handbook of social psychology: Group processes* (pp. 31–55). Wiley-Blackwell. doi:10.1002/9780470998458.ch2
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48(6), 1467.
- Steele, R. J., & Raftery, A. (2010). Performance of bayesian model selection criteria for gaussian mixture models. *Frontiers of Statistical Decision Making and Bayesian Analysis*, 113–130.
- Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. New York: Random House.
- Todorov, V., & Filzmoser, P. (2009). An object-oriented framework for robust multivariate analysis. *Journal of Statistical Software*, 32(3), 1–47. Retrieved from <http://www.jstatsoft.org/v32/i03/>
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cogn Sci*, 38(4), 599–637. doi:10.1111/cogs.12101
- White, D., Burton, A. M., Kemp, R. I., & Jenkins, R. (2013). Crowd effects in unfamiliar face matching. *Applied Cognitive Psychology*, 27(6), 769–777.

- Wolf, M., & Krause, J. (2014). Why personality differences matter for social functioning and social structure. *Trends in Ecology & Evolution*, 29(6), 306–308.
- Wright, N. D., Bahrami, B., Johnson, E., Di Malta, G., Rees, G., Frith, C. D., & Dolan, R. J. (2012). Testosterone disrupts human collaboration by increasing egocentric choices. *Proceedings of the Royal Society B: Biological Sciences*, 279, 2275–2280.
- Yetton, P. W., & Bottger, P. C. (1982). Individual versus group problem solving: An empirical test of a best-member strategy. *Organizational Behavior and Human Performance*, 29(3), 307–321.
- Zhao, K., & Smillie, L. D. (2015). The role of interpersonal traits in social decision making: Exploring sources of behavioral heterogeneity in economic games. *Personality and Social Psychology Review*, 19(3), 277–302.

Table 1

Experiment 1: 95% confidence intervals for primary predictors from multinomial ridge regression (intervals excluding 0 for individual categories are shown in bold).

Predictor	Middle (Wide)		Middle (Narrow)		A		B	
	2.5	97.5	2.5	97.5	2.5	97.5	2.5	97.5
Rel Err ^a	-0.04	0.04	-0.05	0.05	-0.12	-0.04	0.04	0.12
AB distance ^b	-0.36	-0.12	-0.99	-0.73	0.35	0.77	0.32	0.77
AB history ^c	-0.20	0.42	0.05	0.76	-0.64	0.16	-0.66	0.16
Rel Err (-3) ^d	-0.04	0.04	-0.05	0.05	-0.09	-0.01	0.00	0.09
joint RT ^e	0.05	0.10	-0.06	0.00	-0.07	0.02	-0.07	0.02

^aRelative Error. ^bAngular distance between individual responses of participants A and B. ^cAverage angular distance between individual responses of participants A and B since beginning of experiment. ^dRelative error (3 trials back). ^eTime taken to enter final joint response.

Table 2

Experiment 2: 95% confidence intervals for primary predictors from multinomial ridge regression (intervals excluding 0 for individual categories are shown in bold).

Predictor	Middle		A		B	
	2.5	97.5	2.5	97.5	2.5	97.5
Rel Err ^a	-0.05	0.04	-0.11	-0.04	0.04	0.11
Rel Err history ^b	-0.16	0.17	-0.28	-0.04	0.03	0.27
AB distance ^c	-0.84	-0.54	0.15	0.55	0.15	0.57
AB Rel Err history ^d	-0.34	-0.02	-0.11	0.30	-0.13	0.29
joint RT ^e	0.00	0.04	-0.04	0.02	-0.04	0.02

^aRelative Error. ^bAverage relative error across all preceding trials.

^cAngular distance between individual responses of participants A and

B. ^dAverage absolute value of relative error of all preceding trials.

^eTime taken to enter final joint response.

Table 3

Experiment 3: 95% confidence intervals for primary predictors from multinomial ridge regression (intervals excluding 0 for individual categories are shown in bold).

Predictor	Middle		A		B	
	2.5	97.5	2.5	97.5	2.5	97.5
Rel Err ^a	-0.04	0.03	-0.12	-0.06	0.06	0.12
Rel Err history ^b	-0.07	0.07	-0.13	-0.03	0.03	0.13
AB distance ^c	-1.02	-0.80	0.30	0.61	0.30	0.61
AB Rel Err ^d	-0.09	0.00	-0.04	0.08	-0.04	0.09

^aRelative Error. ^bAverage relative error across all preceding trials.

^cAngular distance between individual responses of participants A and B. ^dAbsolute value of relative error.

Table 4

95% confidence intervals for primary predictors from multinomial lasso regression for Experiments 1–3 (intervals excluding 0 for individual categories are shown in bold).

Predictor	Middle (Wide)		Middle (Narrow)		A		B	
	2.5	97.5	2.5	97.5	2.5	97.5	2.5	97.5
Experiment 1								
Rel Err ^a	-0.03	0.03	-0.04	0.04	-0.18	-0.01	0.01	0.18
AB distance ^b	-0.60	-0.25	-1.51	-0.94	0.29	1.04	0.28	1.04
Experiment 2								
Rel Err ^a			0.00	0.00	-0.16	0.00	0.00	0.15
AB distance ^b			-1.35	-0.76	0.00	0.43	0.00	0.44
Experiment 3								
Rel Err ^a			0.00	0.00	-0.18	-0.01	0.01	0.17
AB distance ^b			-1.91	-1.31	0.00	0.32	0.00	0.32

^aRelative Error. ^bAngular distance between individual responses of participants A and B.

Figure 1. Structure of a single trial in the joint judgement task.

Figure 2. Mean error (distance from actual answer) for participants' individual answers and the dyads' answers (leftmost bars). Also shown are the predictions of three models (TakeBest: Take the Best algorithm; Bayes: optimal weighted average; WOC: Wisdom of crowds algorithm). Error bars show ± 1 between-subjects standard error. Note that the error bars are conservative as they do not account for the within-subjects nature of comparisons between different bars.

Figure 3. Normalised histograms of joint proportions (the placement of the joint answer with respect to the two individual answers, represented by 0 and 1 on the abscissa), overlaid with fitted Gaussian mixture model (GMM; lines). See text for details of the GMM modelling.

Figure 4. Joint Proportion score as a function of Relative Error. Each point represents an individual trial, and linear regression fit also shown.

Figure 5. Illustration of the consequences of assuming normal distributions (top row) vs t -distributions (bottom row) for integrated joint judgement. The vertical lines indicate example initial estimates, and the densities are obtained by taking the product of individual normal or t distributions around those estimates. The product of two normal distributions is always normally distributed, and thus unimodal. A similar unimodal distribution is also observed when individual distributions are described by t -distributions and the discrepancy is small (bottom left panel), but as the discrepancy gets larger, the distribution becomes more distinctly bimodal.