https://ddmrmjournal.in
Academic Journal 2025, 18, 44-56

Original Research Article

Wisdom of the crowds: An overview

Surajit Barua

Department of philosophy, Seth Anandram Jaipuria College, Kolkata, West Bengal-700005, India
Email: surajitphilosophy@gmail.com

Abstract:
A number of authors have argued for the idea that crowds and groups are often, if not always, better or more accurate than individuals (including experts) at problem solving and making reliable judgments. This article gives an overview of the some of the prominent contributions made in this area by different authors in recent times. In particular, we explore the merits and demerits of different methods of aggregating individual opinions to generate group judgments. The conditions conducive to the sustenance of groups or crowds, which are capable of producing accurate judgments or predictions, are also discussed briefly.

Keywords:
Wisdom, Crowd, Deliberation, Group

Article History:
Received: 27 September 2025
Revised on: 11 November 2025
Accepted: 29 November 2025

Correspondence to:
Surajit Barua
surajitphilosophy@gmail.com

ISSN: 2331-315X

Introduction

During the last 20 years or so, a number of researchers have reported the superior decision-making and cognitive capability of the collective over the individual in various case studies. In this article, I survey the existing literature to discuss the claim that the crowd is often better and more accurate than any individual at making predictions and judgments.

I begin section 2 by introducing several examples cited by leading authors such as James Suroweiki, on the wisdom of the crowds. These examples, which are varied in nature, reveal some counterintuitive characteristics about crowds. Section 3 argues that crowds manage to achieve a high level of collective accuracy even if (i) they fail to represent the population, (ii) they are composed of complete strangers, (iii) members do not interact with each other. Section 4, 5 and 6 deal with three different methods of aggregation of individual decisions or predictions, namely deliberation with others, mathematical aggregation and prediction markets, respectively. A brief discussion of each method is followed with a description of some of their limitations. Section 7 briefly engages with the possibility of combining the three methods by means of the Delphi method. Section 8 summarizes and concludes the article by briefly suggesting how the recent research on collective wisdom can advance our philosophical understanding on other issues.

Wisdom of the crowds: some examples

James Suroweiki, in his seminal 2004 book, *Wisdom of the crowds*, made a bold claim: the collective opinion of a group of random people is often more accurate than the opinion of any individual or of a group of a like-minded people or even experts. Centuries before him, the Greek philosopher Aristotle made a similar claim. In *Politics*, he claims that a crowd of persons who are not experts, can collectively deliver better results than the experts, in their individual capacity. However, his commentators generally agree that Aristotle's claim to collective intellectual superiority is based on his conviction in the epistemic benefits of deliberation. Suroweiki's claim is radical in that it attributes superior intelligence and accuracy to the collective even in the absence of any exchange of information among group members. He describes several case studies in support of his claim, two of which are mentioned below:
Weight-judging competition: This story is about Francis Galton, who is known for his research into human intelligence. In 1906, he visited an annual regional fair in Plymouth, England, where local people gather every year to assess the quality (i.e. breed, weight etc.) of the town's livestock. During his visit, Galton observed that a competition was being held where the spectators were asked to bet on the weight of an ox on display. The person whose estimate was closest to the actual weight would win. Galton reported that about 800 people from different backgrounds participated, including non-experts who did not have specialized knowledge of cattle breeding and rearing.

After the contest was over, Galton collected the chits of papers on which the individual participants had written their estimate of the weight of the ox. Then he calculated the mean of those numbers. That number represented the collective wisdom of the Plymouth crowd. Since the participating individuals consisted of both experts and laymen, it is intuitive to conclude, at first glance, that the mean weight would be way off the mark. Unexpectedly, that was not the case. The mean weight was very close to the actual weight of the ox. This is remarkable because many of the participants in the competition had no special knowledge of cattle. Also, there is no reason to think that they were not very rational in their decisions. Yet, collectively they managed to make a wise decision. The second case reported by Suroweiki demonstrates that a group's epistemic prowess is not limited to solving a simple data problem like the one above.

The missing submarine: On May 22, 1968, the U.S. navy reported losing contact with one of its submarines, USS Scorpion, which was at sea. Its last verified radio contact, made on May 21, 1968, was approximately 50 miles, south of the Azores islands. The navy had no idea where it might have gone or what might have happened to it after that contact. One might think that asking three or four top experts in the field was the best possible way to find its location. However, a navy officer had a different idea. He scrambled a team of members having diverse knowledge. He then offered several competing hypotheses regarding the fate of the submarine to the group and asked each member to rank those hypotheses in terms of their verdict on how likely each scenario was, without consulting others in the team. He took all those individual opinions and applied probability theorems to determine the team's collective estimate of the present location of the lost submarine. This collective estimate differed from all particular locations which the individual team members had ranked the highest. And yet, the submarine was discovered within 220 yards from where the navy officer's team had said it would be. Thus, in this case too, the group, as a whole, had greater information and accuracy compared to any of the members.

The reliability of crowds is not confined to groups existing in space and time. It holds even for groups or crowds existing merely in the virtual space. Take the search engine Google, for example. Google has managed to become and remain number one search engine in terms of market share despite competition simply because it gives better (i.e. more relevant, more quickly) results in response to search queries. Out of the several methods google uses in order to deliver better results, one algorithm is called PageRank. The PageRank algorithm ranks billions of webpages in seconds according to their relevance to a particular search. Sergey Brin and Lawrence Page, who founded Google together, explain its functioning in their famous 1998 paper. According to them, the algorithm basically assigns weights to different webpages in proportion to its popularity among internet browsers. The popularity of a particular webpage is gauged on the basis of the citations (determined by means of mouse clicks, screen time and so on) it receives. So, if a particular webpage has many citations or only a few citations, but from other webpages which are themselves heavily cited, its rank goes up.

Thus, the algorithm treats the internet as a group or crowd, albeit a huge one. Google aggregates the votes (citations/links clicked) in favour of different websites and ranks them accordingly. However, not all votes have equal weights. The votes from those websites which themselves have been highly visited or cited by the internet are assigned greater weights. The votes in favour of each website are then counted and the websites are then stacked one after the other on the result page, with the website getting the maximum votes at the top.

The cognitive capacity of a crowd goes beyond its collective ability to reliably judge or predict facts. Experiments show that even its reasoning ability surpasses that of any of its individual members. Sperber and Mercier (2012) report of an experiment by Moshman and Geil, in which participants perform the Wason selection task[1] either individually or in

[1] It is a famous experiment in psychology, designed to test deductive reasoning skills using four cards. See Ragni et al. (2017) for a description of the task.

groups. It was found that 70 percent of the groups reported the correct solution but only 9 percent of the individuals succeeded in their individual capacity.

The counterintuitive features of reliable groups

The examples cited above brings out some remarkable features of crowds or groups:

(i) the epistemic superiority of groups does not necessarily require that the group in question has a particular constitution or nature. Some groups are tightly organized and the members are conscious of their belongingness to their groups. Others are not formally organized. And still others have ever-changing membership, like the cattle market;

(ii) groups could achieve a high level of accuracy without engaging in mutual exchange of ideas or deliberation;

(iii) the group sizes were often too small to represent the population about whom they were predicting;

(iv) there was lack of demographical diversity within the groups.

However, it is fact that not all groups or crowds produce accurate judgments and decisions. Mob behaviour and decisions are often irrational. For example, it is not uncommon to find instances of mob violence right after a major sports tournament between two traditionally rival teams. Similarly, there are instances where a coup is often followed by chaos and irrational mob behaviour, resulting in widespread suffering and plunder. Likewise, during stock market crash or exuberance, market participants (i.e. traders) often make collectively irrational decisions, leading to unjustifiable increase or decrease in stock prices.

Suroweiki therefore suggests that there must be a set of rules which good, epistemically superior groups or crowds follow, which helps them to maintain order and coherence. He suggests that group judgments which are based on diversity[2] and independence of individual judgments of its members are particularly reliable. Explaining the importance of the above characteristics, he argues that collective decisions which are a product of consensus or compromise amongst group members often fail to incorporate the individual wisdom of each member[3]; The most accurate collective judgments are often borne out of disagreements or divergence of opinions. A cognitively smart crowd or group does not coerce its members to abandon or modify their personal opinions in order to conform to the majority or the most influential members of the group. It does not seek to

[2] Note however, that diversity, as meant here, does not imply the stronger condition that the crowd must be a representative sample of the larger population.

[3] This point has been emphasized by other writers also, as discussed elsewhere in this article.

influence or ignore their opinions so that a happy or desirable collective decision can be generated. Instead, it utilizes neutral, mechanistic forces like market prices or technologically smart voting procedures to aggregate and generate collective decisions. Thus, the best way for a group to judge smartly or intelligently is to have conditions where each member can think and act freely, without any direct or indirect influence from others.

Several other authors too have emphasized the significance of independence and diversity in enhancing a group's cognitive performance. For example, Lorenz et al. (2010) demonstrate through several experiments that the accuracy of the median estimate of a crowd worsened when individual members were informed about the information provided by others in the group on the same question. The information about what others think about the same issue often leads to unjustified and unreasonable revision of an individual's one's own judgment towards consensus, thereby worsening the accuracy of the entire group. Thus, independence is necessary. Similarly, Hong and Page (2004) conducted a series of computational experiments wherein they divided their test subjects into two teams, one comprising of randomly selected volunteers and the other consisting of only those individuals who excelled in solo performance. Then, similar tasks were assigned to the two teams. The result corroborated the above claim- the former team outperformed the later in terms of accuracy of judgments.

Bhatt et al. (2017) reached similar conclusions about virtual crowds by analysing online social media (Twitter) communications; They found that a group which is diverse (in terms of tweet content of its members) is better at generating reliable, true judgments compared to crowds whose members are randomly selected.

Methods of aggregation of individual judgments: Deliberation

However, even if a group has the desirable characteristics noted above, another issue needs resolution: How the individual judgments of members should be aggregated (or collected) to determine what the group, as a whole, thinks? This question is crucial because a wrong method of aggregation of individual judgments may deliver a collective verdict which is unreliable despite the group having diversity and independence.

A common and popular method of aggregation is deliberation with others. Unstructured deliberation unearths hidden presuppositions, sometimes transmits additional information/evidence during the public analysis of reasons and leads to correction of errors by individual believers. For example, Lyon and Pacuit (2013) point out that unregulated discussion between random participants in the comments section of the polymathblog led to a new proof of the Hales-Jewett Theorem.

However, empirical research also shows that unstructured debate and discussion often lead to worse judgments than can be obtained without deliberation. Irving Janis has even coined a term for a malaise affecting unstructured deliberation- *groupthink*. Miriam Solomon (2006) discusses its effect in her article. When a group of individuals sit together to discuss a controversial issue with the aim of reaching consensus, its decision is susceptible to several factors. The peer pressure to arrive at a consensus can discourage a member from voicing an opinion which might be different from that of the majority; this person might fear of being identified as a spoiler(outlier) to a consensus. Thus, fear of rejection by the group may prompt this person to either remain silent or nod in affirmation to the majority opinion. Likewise, in actual deliberations, it is often seen that a few members are particularly influential and/or vocal. If any other member has a different opinion, they may resist sharing their judgment for the fear of being seen as a challenge to authority. So, the urge to be accepted by this influential few may prevent someone from voicing their own opinion. These effects ultimately lead to a group consensus which does not truly represent the collective wisdom of the group, for the opinion of the uninfluential or silent members is left out in determining the collective judgment. Even if one is aware of such biases and tries to be unbiased and independent, he/she might be unable to resist groupthink due to its pervasive influence. Consequently, such a deliberating group is prone to making poor decisions.

A real-life example of groupthink, cited by both Janis and Solomon, is the ‘Bay of Pigs’ invasion. It is a failed covert attempt in 1961 by the then U.S. government to overthrow the Fidel Castro regime in Cuba, by training and funding a force of exiled counter-revolutionary Cubans. Irving Jenis disagreed with the official explanation. According to him, the desire by members of Kennedy’s inner circle to reach a consensus through deliberation and discussion led to flawed tactical decisions and that led to the failure of the project. This happened because reasonable doubt and opinions, divergent from the one espoused by the most influential member(s) of the group, got suppressed in the urge for conformity with the views of the powers that be. Consequently, the group judgments were not really the aggregated decisions of all the members of the group but only a few influential and vocal ones. It is therefore no surprise that the project failed miserably.

Lyon and Pacuit (2013) too have identified three major issues with the method of deliberation. One is that the opinions of those members who have a lower rank or status or represent the minority are often ignored while aggregating individual opinions. Secondly, deliberation often gives disproportionately heavy weight on one piece of information (for example, the first announced judgement, the judgement of the most senior person in the group, or the judgment of the loudest person in the group). Thirdly, the weight of a particular information is often proportionate to how widely it is held. That is, if a piece of information is held by all, it is often given more weight just for that reason, compared to the weight assigned to information held by only a few.

Moreover, deliberation with others is susceptible to a bias known as ‘group polarization’. According to Cass Sunstein, this bias refers to a tendency of a deliberating group to lean heavily towards a point of view its members already believed in, before they started deliberating. He further adds that this tendency gets accentuated when the group members are like-minded people who are not open to competing views. In such situations, deliberation gets reduced to a mere opportunity to rehearse and reinforce their pre-deliberation views.

Sunstein cites results of several empirical case studies to bolster his claim. For example, in one study, it was found that test subjects who were moderately profeminist before deliberation became more strongly so after interacting with each other. Another experiment he cited found that Caucasian test subjects who were predisposed to be racially prejudiced mostly answered the following question in negative—whether white racism is responsible for conditions faced by African-Americans in American cities. On the contrary, those white subjects who were not so predisposed gave a more positive reply to the same question.

A possible solution that comes to mind immediately is heterogeneity; if you make the deliberating group heterogenous enough, perhaps the effect of group polarization can be minimized. However, even a heterogenous group is susceptible to the objections raised previously about deliberation. In such deliberative groups, views of low-status members are given least weight[4]. Sunstein too argues that in deliberating bodies, a handful of influential persons initiate and steer the conversation. Other members merely conform to their view. The less-influential members having a low status in the group do not express any divergent views they might be entertaining in their mind, for fear of retribution or lack of confidence. Empirical research[5] shows, for example, that in cultural mixed groups and mixed gender groups, cultural minorities and women respectively, have "disproportionately little influence on decisions. Thus, heterogeneity of a group, in itself, cannot guarantee protection from the ill effects of deliberation.

Methods of aggregation of individual judgments: Mathematical aggregation

Apart from deliberation, another common method of aggregation used to collate individual information in the group is *weighted averaging* of the answers of individual group members. Since experts are supposed to be more knowledgeable than the laymen, it is plausible that a policy of aggregation that assigns more weight to experts' judgments than to the novice is preferable to one which assigns equal weight to every member. Thus, everyone in the group would arrive at a certain value/content of the collective judgment, based on his own assessment of who in the group are experts. We should now average again, using the new judgments and original weights to arrive at the final collective judgment. This method seems to avoid the limitations or defects of deliberation with others.

However, this method is not foolproof either. One potential drawback, pointed out by Lyon and Pacuit, is that "people's judgements of each other's level of expertise do not track the accuracies of their judgements." (2012:605) Not only do different members assign different weight to each expert in the group, there is no correlation between the weight

[4] See Solomon (2006) and Christensen and Abbott (2000).

[5] Kirchmeyer and Cohen (1992), p-153, Christensen and Abbott (2000), p-273.

assigned to an expert and the reliability of their decisions/judgments[6]. Consequently, averaging is not always a reliable way of aggregating accurate group judgments.

When the information sought is qualitative in nature, a good alternative is to use the following majority rule: the collective judgment of the group is what more than 50% of the individual members say. In the special case where each member (a) is individually more likely than not to choose the correct verdict and (b) reaches their individual verdict independently, the *Condorcet Jury theorem* justifies the majority rule. It answers the question: what is the probability that a majority of a group of jurors will choose the right verdict, given that, each juror: The reply to the question is that the number of jurors is directly proportional to the probability of a majority of the jurors reaching the correct verdict. Suppose for example, that each of the jurors has a 60% chance of espousing the right verdict. In that case, if we have three jurors, the probability of a majority selecting the right verdict is roughly about 65%. But if the number of jurors is 26, the probability shoots to 85%.

A problem with using the majority rule as a method of aggregation is the so-called *doctrinal paradox*. This paradox occurs when the inputs and outputs are not single, all-or-nothing type of judgments but a set of logically connected propositions. Here is a hypothetical case, similar to the one suggested by Lyon and Pacuit (2013).

Consider the following three propositions: B means the individual is healthy, A means the person is tall and A → B means that if a person is tall, then that person is healthy. Assume that there are 30 people in the crowd under examination, and they make the following judgements on the three propositions, A, A → B, and B

	A	A → B	B
10 people say	True	True	True
Another 10 people say	True	False	False
The remaining 10 people say	False	True	False
So, the greater than 50% majority rule says	True	True	False

The last row in the above table shows that when we apply the majority rule to the group of 10 people, the truth value of the three propositions A, A → B, and B constitute an invalid inference; collectively speaking, the three propositions are inconsistent. However, the judgment of every individual in the crowd is consistent with respect to those three propositions. Thus, a simple majority rule may yield inconsistent results.

I think that the above anomaly is not a problem *per se* for the majority rule. Rather, it is a problem about the applicability of the rule. Notice that the subset of the group whose members believe that A, is extensionally different from that subset which believes that A

---

[6] See Burgman et al. (2011).

→ B. Similarly, the other subset that believes that G, is extensionally different from the other two. We should not ascribe different premises of the same argument, to different set of persons; if you are examining the cogency of an argument, all of its premises and conclusion should be attributed to the same person or a particular group of persons. Thus, the doctrinal paradox rests on a weak assumption. It isn't a good objection to the majority rule.

Methods of aggregation of individual judgments: Prediction markets

Yet another method of aggregation is the so-called "prediction markets". Prediction markets consist of participants who bet money on unknown future event outcomes. Information is aggregated from such individual bets and predictions of the collective about future events are produced. The current price represents the traders' collective consensus about the expected value of the contract. For example, the Iowa Electronic Markets (IEM) Presidential vote-share markets is a computerized, electronic, real-time exchange where traders buy and sell future contracts, with payoffs based on election outcomes. Thus, it is a prediction market. Berg et al. documented the performance of this market over 964 polls from the five U.S. Presidential elections since 1988. They concluded that "the market is closer to the eventual two-party vote split 74% of the time. Further, the market significantly outperforms the polls in every election when forecasting more than 100 days in advance." (Berg et al., 2008: pp. 287)

In a study conducted by Pennock et al. (2001), researchers studied the correlation between the going prices and observed outcome frequencies in two different market games which are played online with play money only; In one of the games, the traders had to bet on the possible outcomes to scientific questions which have not yet been considered settled by the scientific community. In the other, they had to bet on the outcome of Oscar, Emmy, and Grammy awards. The researchers collected historical price information for 161 expired securities, corresponding to questions that had been definitively answered "yes" or "no". The found that prices of securities in Oscar, Emmy, and Grammy awards correlate well with actual award outcome frequencies. Similarly, they discovered that market prices in the other game also strongly correlate with actual observed frequencies.

Sunstein (2006) lists several advantages of prediction markets over unstructured deliberation. In a deliberating group, members having low status or influence have little incentive to share their views. On the contrary, investments in prediction markets are not generally disclosed to the public. Thus, market participants do not have to suffer adverse consequences for expressing their own views. Sunstein also noted that the desire to make profit and avoid losses in a highly competitive market incentivizes traders to utilize whatever information they have in making trading decisions. Prediction markets, according to him, do not accentuate individual errors. The individual errors cancel each other in the fight between bulls and bears.

Like the other two methods of aggregation discussed above, prediction markets too have limitations. Sunstein lists several issues. One is that it is not clear how to use prediction markets in determining the answers to normative ('ought') questions. This is because factual questions have either a 'yes' or 'no' answer. Because the option is either positive or negative, the current market price can be deemed to be the result of the tug of war between bulls (those who answer the question in positive) and bears (those who answer in negative). With normative questions, such binary options are not available because they are often conditional in nature. Two persons may assent to a normative question, but they have different conditions in mind. It is not clear how to incorporate such nuances in the market price. Normative questions seek answers regarding what one ought to do or think. Even if prediction markets come up with an answer, objective way to test whether the settled price represents the right answer.

Another problem may arise if the prediction market in question consists of very few participants. In the absence of enough traders, the market price of the security in question may fail to price in all the relevant information about that security. For example, it is common knowledge that the price of illiquid stocks tends to experience high volatility and very often their price do not represent their underlying financial fundamentals—either the share would be overvalued or undervalued.

Which method of aggregation to choose?

In the above paragraphs, we have discussed three methods of aggregating individual judgments of a crowd or group: deliberation, mathematical aggregation (averaging and majority rule) and prediction markets. None of these methods are without a drawback. However, the problems can be minimized by either combining and complementing the different methods of aggregation or by demarcating the domain of application of the three methods. For example, the Delphi method is widely used in different sectors or industries to address problems "that could benefit from the subjective judgments of individuals on a collective basis." (Skulmoski et al. 2007, p-2)

Rowe and Wright (1999) list four key features of the Delphi method:

Questionnaires, rather than debate and discussion, are used to elicit responses from individual group members. In this manner, they are given the opportunity to express their opinions and judgments privately and protects them from undue social pressures– as from dominant or dogmatic individuals, or from a majority.

The questionnaire is administered repeatedly to give individuals the chance to change their opinions and judgments without fear of losing credibility in the group.

Between each questionnaire iteration, controlled feedback is provided through which the group members are informed of the opinions of their anonymous colleagues. Feedback comprises the opinions and judgments of all group members and not just the most vocal.

After several rounds of administration of questionnaire, the group judgment is taken as the statistical average of the volunteers' estimates on the final round.

Note that the characteristics of anonymity and iteration have the advantages of prediction markets; controlled feedback represents the advantage of deliberating with other group members and the last mentioned characteristic statistical aggregation correspond to mathematical aggregation. Thus, the Delphi attempts to incorporate the good things about all the three methods of aggregation without attracting the pitfalls of any of them.

Conclusion

Recent research demonstrates that under certain conditions, the collective verdict, decision or judgment of a group (or crowd) is generally more accurate than that made by any of its individual members. What makes this feature remarkable is that group members need not collaborate or cooperate amongst themselves to achieve this feat. Nor is it necessary that the group must consist of only carefully selected experts in the field. A bunch of complete strangers can also produce a collective output whose accuracy cannot be matched consistently by any individual expert. In fact, diversity and independence amongst group members in matters of constructing individual opinions are features which contribute positively to the cognitive prowess of the group.

The method of aggregation of individual judgments or predictions also greatly influence the collective judgment of the group as a whole. Each of the three methods of aggregation discussed in this article, namely mathematical aggregation, deliberation with other group members, and prediction markets, has certain advantages and some limitations. In order to extract maximum epistemic benefits from a group, it has to be ensured that (a) the group has the right constitutive characteristics of diversity, independence and decentralization and (b) the right method is employed to aggregate individual judgments into a collective one. Attempts have been made to devise a method which combines the good attributes of all the three methods of aggregation while overcoming their limitations. One such method, the Delphi, is briefly discussed above. Needless to say, there is room for further exploration in this area, in light of the fact that nowadays groups of different constitution have proliferated which were unheard of even a decade back; Not only are there numerous groups which exist online only, there are even groups whose members are not aware of their membership. There are groups whose members are all anonymous and there are still others whose members are all not unique or even human (think of bots). Due to the increasing complexity of groups, we have to be more diligent in extracting information from such groups. Consequently, more research is required to integrate and demarcate the different methods of aggregation, so that we can properly utilize the wisdom of the crowds

How can the philosopher benefit from the above discussion? One obvious application is in epistemology. We ought to justify our beliefs by giving reasons or evidence. Very often, the reasons we give do not adequately or genuinely support the corresponding beliefs. Yet we have several personal practical or emotional interests, which may differ from person to person. These non-epistemic factors may cloud our personal judgment about the extent to which a particular fact qualifies as a reason for a given belief. The collective judgment of one's epistemic community is better suited to this task since, in a sufficiently diverse crowd, individual emotional and practical biases get cancelled out.

Wisdom of the crowds also puts into doubt the great importance philosophers have been assigning to deliberation and discussion with others. Human emotions and selfish tendencies creep into deliberations in such a way that the method of deliberation is liable to be exploited as a means to advance personal gains, rather than being used as a means to pursuing truth. Nonetheless, the discussion above shows that it is possible to transcend or cancel individual interests and aggregate individual wisdoms into collective wisdom which is far more reliable.

References

1. Aristotle. *Politics*. Trans. C. D. C. Reeve. Indianapolis: Hackett.1998.
2. Berg, Joyce E., Forrest D. Nelson, and Thomas A. Rietz. "Prediction market accuracy in the long run." *International Journal of Forecasting* vol. 24, no.2, 2008, pp. 285-300.
3. Bhatt, S., Minnery, B., Nadella, S., Bullemer, B., Shalin, V., & Sheth, A. "Enhancing crowd wisdom using measures of diversity computed from social media data." In *Proceedings of the international conference on web intelligence,* 2017, pp. 907-913.
4. Brin, Sergey, and Lawrence Page. "The anatomy of a large-scale hypertextual web search engine." *Computer networks and ISDN systems* vol. 30. no. 1-7, 1998, pp. 107-117.
5. Burgman M, McBride M, Ashton R, Speirs-Bridge A, Flander L, Wintle B, Fidler F, Rumpff L, Twardy C. "Expert status and performance." *PLoS One,* vol. 6, no.7, 2011, pp. e22998, pp.1-7
6. Caryn. C. and A. Abbott, "Team Medical Decision Making," in *Decision Making in Health Care*, ed. Gretchen Chapman and Frank Sonnenberg (New York: Cambridge University Press, 2000), vol. 267, pp. 273–276.
7. Davis-Stober, Clintin P., David V. Budescu, Jason Dana, and Stephen B. Broomell. "When is a crowd wise? " *Decision* vol. 1, no.2, 2014.
8. Hong, Lu, and Scott E. Page. "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." *Proceedings of the National Academy of Sciences,* vol. 101, no.46, 2004, pp. 16385-16389.
9. Kirchmeyer, Catherine, and Aaron Cohen. "Multicultural groups: Their performance and reactions with constructive conflict." *Group & Organization Management,* vol. 17, no. 2, 1992, pp. 153-170.
10. Lorenz, Jan, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. "How social influence can undermine the wisdom of crowd effect." *Proceedings of the national academy of sciences*, vol. 108, no. 22, 2011, pp. 9020-9025.
11. Landemore, Hélène. "Collective wisdom: Old and new." *Collective wisdom: Principles and mechanisms,* vol. 1, 2012, pp. 1-20.

12. Lyon, Aidan, & Pacuit, Eric. “The wisdom of crowds: Methods of human judgement aggregation.” In *Handbook of human computation,* Springer, New York, NY, 2013, pp. 599-614.
13. Pennock, David M., Steve Lawrence, C. Lee Giles, and Finn Arup Nielsen. "The real power of artificial markets." *Science,* vol. 291, no. 5506, 2001, pp. 987-988.
14. Ragni, Marco, Ilir Kola, and Phil N. Johnson-Laird. "The Wason selection task: A meta-analysis." In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 39. 2017. pp. 980-985.
15. Skulmoski, Gregory. J., Hartman, Francis. T., & Krahn, Jennifer. “The Delphi method for graduate research.” *Journal of Information Technology Education: Research*, vol 6, no. 1, 2007, pp. 1-21.
16. Solomon, Miriam. "Groupthink versus the wisdom of crowds: The social epistemology of deliberation and dissent." *The Southern journal of philosophy* vol 44, 2006, pp. 28-42.
17. Sperber, Dan, and Hugo Mercier. “Reasoning as a social competence”, *Collective wisdom: Principles and mechanisms,* 2012, pp. 368-392.
18. Sunstein, Cass. R. “The Law of Group Polarization.” *Journal of Political Philosophy*, vol. 10, 2002, pp. 175–95.
19. Sunstein, Cass. R. “Deliberating groups versus prediction markets (or Hayek's challenge to Habermas).” *Episteme*, vol. *3*, no. 3, 2006. pp. 192-213.
20. Surowiecki. James. *The wisdom of crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Doubleday, New York. 2004.