Prioritizing Policy Objectives in Polarized Groups using Artificial Swarm Intelligence

Gregg Willcox
Unanimous AI
San Francisco, CA
Gregg@unanimous.ai

Louis Rosenberg
Unanimous AI
San Francisco, CA
Louis@unanimous.ai

Mark Burgman
Imperial College London
London, UK
M.burgman@imperial.ac.uk

Alexandru Marcoci
University of North Carolina, Chapel Hill
Chapel Hill, NC
Alexandru.marcoci@gmail.com

Abstract— **Groups often struggle to reach decisions, especially when populations are strongly divided by conflicting views. Traditional methods for collective decision-making involve polling individuals and aggregating results. In recent years, a new method called Artificial Swarm Intelligence (ASI) has been developed that enables networked human groups to deliberate in real-time systems, moderated by artificial intelligence algorithms. While traditional voting methods aggregate input provided by isolated participants, Swarm-based methods enable participants to influence each other and converge on solutions together. In this study we compare the output of traditional methods such as Majority vote and Borda count to the Swarm method on a set of divisive policy issues. We find that the rankings generated using ASI and the Borda Count methods are often rated as significantly more satisfactory than those generated by the Majority vote system (p<0.05). This result held for both the population that generated the rankings (the "in-group") and the population that did not (the "out-group"): the in-group ranked the Swarm prioritizations as 9.6% more satisfactory than the Majority prioritizations, while the out-group ranked the Swarm prioritizations as 6.5% more satisfactory than the Majority prioritizations. This effect also held even when the out-group was subject to a demographic sampling bias of 10% (i.e. the out-group was composed of 10% more Labour voters than the in-group). The Swarm method was the only method to be perceived as more satisfactory to the "out-group" than the voting group.**

Keywords—Artificial Swarm Intelligence, Human Swarming, Artificial Intelligence, Voting Methods, Borda Count, Majority Voting, Brexit.

I. Introduction

Groups often struggle to reach satisfactory decisions, in the sense that most participants approve of the decision, especially when the population is divided by conflicting views. This is particularly true in the realm of governmental decision making, as deeply held political and ideological opinions often prevent groups from reaching decisions that satisfy the whole, or even most of, the population.

Traditional approaches to group decision-making solicit isolated opinions. The results are then aggregated using a voting algorithm. The Majority Voting algorithm, in which each polled respondent votes for only one candidate, and the candidate with the most votes wins, is the most widely voting algorithm in the Anglo-Saxon world [20].

Other widely studied systems include ranked voting methods, such as the Condorcet [22] and Borda Count [23] methods, both developed in the 1700's as alternatives to Majority Voting. In these methods, participants rank all candidates from the most to least preferable, and the candidates that are ranked with the highest average preference, or that are ranked higher than other candidates most often, win the vote.

These traditional voting systems have often faced criticism because all fail to pass simple tests of "fair" aggregation algorithms [26]. As one example, the Borda and Majority algorithms are very open to manipulation in real voting systems [21, 27, 28]. Research that tries to distinguish between these and other methods in real-world practice and to find the best voting method for a given context, often tries to calculate the utilitarian value for each voting method's outcome, which is unrealistic in real-world scenarios [25].

In recent years, a new method called Artificial Swarm Intelligence (ASI) has been developed that enables networked human groups to deliberate in real-time systems moderated by artificial intelligence algorithms. Whereas most existing voting methods focus on collecting isolated input from a group of participants in parallel, ASI collects input from participants in real-time: the group engages in a real-time deliberation to vote on a set of alternatives, and participants can switch their response at any point.

A recent study [19] attempting to measure the utilitarian optimality of groups with conflicting opinions found that ASI reaches decisions that are significantly better, as measured by the monetary amount won by the group, than the Borda and Majority methods. In another study [24], groups made political prioritizations using both ASI and Majority Voting protocols. The group later rated it's own prioritizations made via ASI as a more accurate representation of the group's opinions (74% of responses) and a more accurate representation of individual priorities (66% of responses) than the Majority Voting protocol.

This study aims to identify the conditions under which groups benefit most from using ASI to facilitate decision-making (when compared to the Borda and Majority methods).

II. Human Swarming

In traditional voting schemes, participants provide input in isolation. In swarm-based methods, groups think together in systems modeled after biological swarms and converge on collective solutions. As shown in Figure 1, a typical ASI system includes a group of users connected in real-time over a network. Each computer, which may be a desktop, tablet, or phone, runs a unique software interface designed to capture and stream the user's real-time input to a cloud-based processing engine. The engine runs swarming algorithms and sends back real-time output to each user, creating a closed-loop.

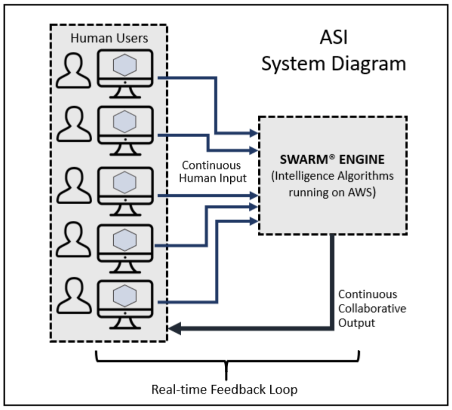

Fig.1. System Diagram for a real-time ASI System

The present study uses Swarm AI technology, which is modeled largely on the dynamic behaviors of honeybee swarms. The decision-making process that governs honeybee swarms has been researched since the 1950s and has been shown at a high level to be quite similar to decision-making in neurological brains [13,14]. Both employ populations of simple excitable units (i.e., neurons and bees) that work in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in real-time. In both brains and swarms, outcomes are arrived at through competition among sub-populations of excitable units. When one sub-population exceeds a threshold level of support, the corresponding alternative is chosen. In honeybees, this enables hundreds of scout bees to work in parallel, collecting information about their local environment, and then to converge together on a single optimal decision, frequently picking the best solution to complex multi-variable problems [15-17].

The similarity between "brains" and "swarms" is apparent when comparing the decision-making models that represent each. The decision process in primate brains is often modeled as mutually inhibitory leaky integrators that aggregate incoming evidence from competing neural populations [18]. A common framework for primate decision-making is the Usher-McClelland model in Figure 2 below.

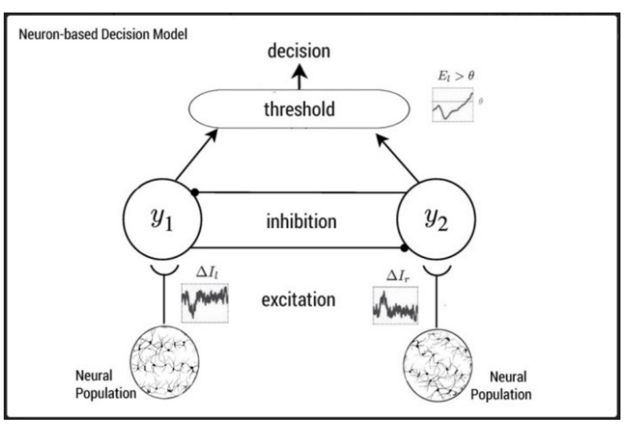

Fig. 2. Usher-McClelland model of neurological decision-making

This neurological decision model can be compared to swarm-based decision models, for example the honey-bee model represented in Figure 3. As shown below, swarm-based decisions follow a very similar process, aggregating input from sub-populations of swarm members through mutual excitation and inhibition, until a threshold is exceeded.

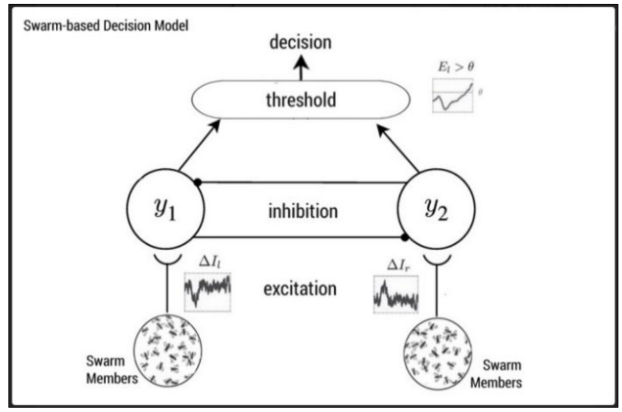

Fig. 3. Mutually inhibitory decision-making model in bee swarms

Thus, while brains and swarms are very different forms of intelligence, both enable decisions to emerge from the interactions among collections of processing units. The goal of the present study is to apply this decision-making model to human groups deliberating on divisive political issues and investigate the satisfaction of group members with the collective output.

III. SWARMING SOFTWARE

To enable swarming among groups of networked humans, ASI technology allows distributed groups of users to form closed-loop systems [5-7] and (a) integrate noisy evidence, (b) weigh competing alternatives, and (c) converge on decisions in synchrony, while also allowing all participants to perceive and react to the changing system in real-time.

As shown in Figure 4, networked human groups can answer questions as a "swarming system" by collaboratively moving a graphical puck to select among a set of alternatives. Each participant uses a mouse or touchscreen to manipulate a graphical magnet. By positioning their magnet with respect to the moving puck, participants impart their personal intent on the system as a whole. The input from each user is not a discrete vote, but a stream of real-time vectors that varies freely. Because all users can adjust their intent continuously in real-time, the puck and the participant swarm around it move, based on the dynamics of the full system. This enables a complex negotiation among all members simultaneously, empowering the group to collectively explore the decision-space and converge on the most agreeable solution in synchrony.

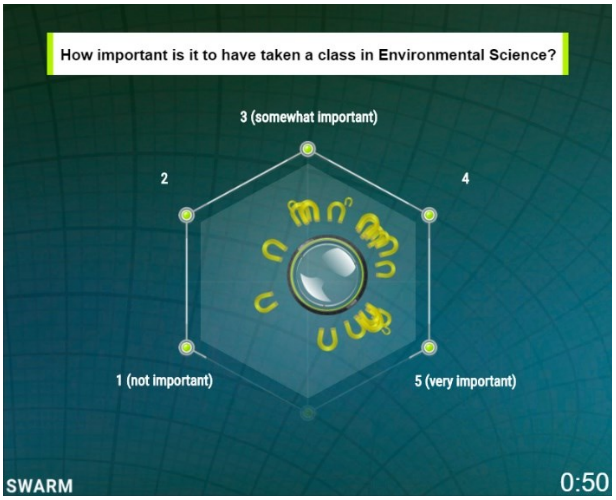

Fig. 4. A human swarm answering a question in real-time

It is important to note that participants freely modulate both the direction and magnitude of their intent by adjusting the distance between their magnet and the puck. Because the puck is in continuous motion across the decision-space, users need to continually adjust their magnet so that it stays near the puck's outer rim. This is significant, for it requires participants to remain continuously engaged throughout the decision process, evaluating and re-evaluating the strength of their opinions as they convey their contribution. If they stop adjusting their magnet with respect to the changing position of the puck, the distance grows and their imparted sentiment wanes. A more complete description of the algorithm can be found in [31, 32].

Thus, like bees vibrating their bodies to express sentiment in a biological swarm, or neurons firing activation signals to express conviction levels within a biological neural-network, the participants in an artificial swarm must continuously update and express their changing preferences during the decision process, or lose their influence over the collective outcome. In addition, algorithms monitor the behaviors of all swarm members in real-time, inferring their implied conviction based upon their relative motions over time. This reveals a range of behavioral characteristics within the swarm population and weights their contributions accordingly, from entrenched participants to flexible participants to fickle participants.

IV. PRIORITIZING POLICY OBJECTIVE IN THE AGE OF BREXIT

A study was conducted to evaluate a political constituency's satisfaction with the output of three prioritization methods—Majority voting, Borda Count voting, and ASI.

A. Pilot Study: Question Divisiveness Measurement

The questions in this study needed to be highly politically divisive for British participants, so a pilot study was conducted to evaluate the divisiveness of six policy questions among Labour and Conservative members of the UK public. With help from political science experts, six policy ranking questions were designed that were likely to be divisive. These six questions were sent to 42 UK citizens: 22 Conservative and 20 Labour voters.

The divisiveness of each question was measured as the average difference between the Labour and Conservative rankings of the question's political objectives. The three questions that were the most divisive were selected for inclusion in the following experiment, and are described in Appendix A.

B. Participants

To evaluate the satisfaction of a political constituency with various prioritization methods we recruited N=237 participants using a market research company. 119 participants were female. Participants were compensated for the 45 minutes or less session in line with the practices of the market research company at a rate of approximately £36 / hour. All participants signed an informed consent form.

C. Materials and Methods

Participants were randomly assigned into two conditions: the In-Group and the Out-Group. Participants in the former condition were randomly assigned to four groups of between 8 and 20 individuals, and were tasked with prioritizing three sets of policy objectives from least to most important and then ranking their satisfaction with each prioritization method. These four groups were considered the "in-group", as they contributed to the prioritizations that they later scored for satisfaction. One "out-group" of 170 participants was also convened that did not contribute to any prioritization, and that only ranked their satisfaction with each prioritization from the in-groups.

To give a more complete overview of the demographics of participants in this study, table 1 lists the demographic composition of each group, split by Political Affiliation, Brexit Stance, and Gender.

	Number of Participants	**Political Affiliation:** Labour / Conservative	**Brexit Stance:** Leave / Remain / Undecided	**Gender:** Male / Female
In-Group 1	8	4 / 4	3 / 5 / 0	4 / 4
In-Group 2	20	10 / 10	9 / 11 / 0	10 / 10
In-Group 3	20	12 / 8	8 / 12 / 0	6 / 14
In-Group 4	19	7 / 12	10 / 8 / 1	12 / 7
Out-Group	170	87 / 83	70 / 88 / 12	86 / 84

Table 1: Demographic Breakdown of Groups

Participants in the in-groups first provided their answers independently using a standard online survey to prioritize each set of objectives. Upon completion, the groups congregated on the Swarm AI platform (an online tool, purposefully built to facilitate ASI decision-making) to answer the same set of questions.

While the participants were completing the ordering tasks using ASI, their survey results were analyzed using the Majority and Borda Count algorithms to generate two ordered lists. For the Majority algorithm, the objectives were ordered by the number of participants that ranked each objective as the "most important", and ties were broken randomly. In the Borda Count algorithm, each participant's ranking was converted into a score for each objective: 1 point for the "most important" objective, 2 points for the second most important objective, 3 points for the third most important, etc., and the sum of these points across all participants in the group was calculated for each objective. The objectives were ordered from least points (most important) to most points (least important), with ties broken randomly.

When moving to the Swarm AI platform, the groups prioritized the objectives using an iterative elimination approach: the groups started by selecting the LEAST important objective out of the 6 objectives listed, then this objective was eliminated from consideration, and the group repeated the process, until there were two objectives left. For the final elimination, the question was flipped, and the group was asked which of the remaining objectives was the MOST important. The ranking generated in this way using the swarm platform was considered the group's ranking of the objectives.

After completing all questions on the Swarm AI platform, participants were redirected to a follow-up survey, where they were presented with the three questions they just answered, along with three anonymized rankings for each question—the Majority, Borda Count, and ASI rankings. For each question, participants were asked to rank each of the three lists based on their level of satisfaction with the list, from 1-most satisfied to 3-least satisfied. The three anonymized lists were presented in a different order for each question to eliminate ordering bias.

Finally, an out-group of participants was assembled to represent a public constituency that did not directly vote on the policy objectives. Members of the out-group were not a part of any of the four group ranking exercises. The satisfaction of the out-group with each of the in-group's rankings was measured using a standard survey. This survey contained 12 questions: each responding to a set of rankings created by one of the four in-groups in response to one of the three questions. Each question required the participant to rank their satisfaction with each of the three prioritizations that the in-group made using the Majority, Borda Count, and ASI methods.

As an example, one question in this survey asked out-group participants to rank their satisfaction with each of in-group 1's question 1 Majority, Borda Count, and ASI prioritizations.

V. Analysis and Results

All participants completed the survey fully before joining the groups on the Swarm AI platform; no survey data were missing at the time of analysis. All questions were answered in between 10 and 60 seconds.

A. *Question Divisiveness*

Overall, each of the three questions proved to be highly divisive when segmenting by the Political Affiliation of participants (Conservative or Labour). Significant and important differences in ranking ($p<0.05$) were observed in each of the three questions when segmenting this way, as shown in Table 2.

Table 2 summarizes the number of statistically significant differences between the Labour and Conservative average ranking of each of the 6 items in each question. These differences were calculated using a 2-sample t-test for each item. Notably, the Political Affiliation of participants significantly impacted their answers to all three questions.

The average effect size for all significant differences, calculated as the average difference between Conservative and Labour rankings of each of the significantly different items, is reported. Significant differences were observed when the average ranking difference was at minimum 0.83 ranks, or 16.6% of the maximum observable difference, which is a considerable difference in ranking.

Question Number	Number of Significant Differences ($p<0.05$) by Political Affiliation	Average Effect Size (Conservative minus Labour Ranking)
1: Objectives	3	1.25
2: Issues	4	1.12
3: Immigration	3	0.93

Table 2: Number of Significant Differences Observed by demographic when ranking priorities as individuals via online survey.

The average ranking difference of each item, and the confidence interval of each difference, is shown in Figures 1-3.

The fact that half of the six items in each question were ranked differently by Labour and Conservative voters in this pool indicates that the questions formulated by the research team were meaningfully divisive and therefore were suitable for use in this experiment.

B. *In-Group Satisfaction*

Next, each in-group's satisfaction with each ranking was calculated, as shown in Appendix B as the average ranking that was given to each list by participants in that group. A score of 1.0 indicates the most preferred list, while a score of 3.0 indicates the least preferred list. In Question 2, Group 1's Majority and ASI lists were the same, so only two independent lists were ranked in this instance.

Analyzing the average satisfaction ranking for each method across all groups, we find that the Swarm method was preferred to the Majority Voting method for all three questions and was significantly preferred to the Majority Voting method on the second question alone ($p<0.01$). ASI also received a higher average satisfaction ranking than the Borda Count method on two of the three questions, though there were no significant differences between the two methods on any of the three questions.

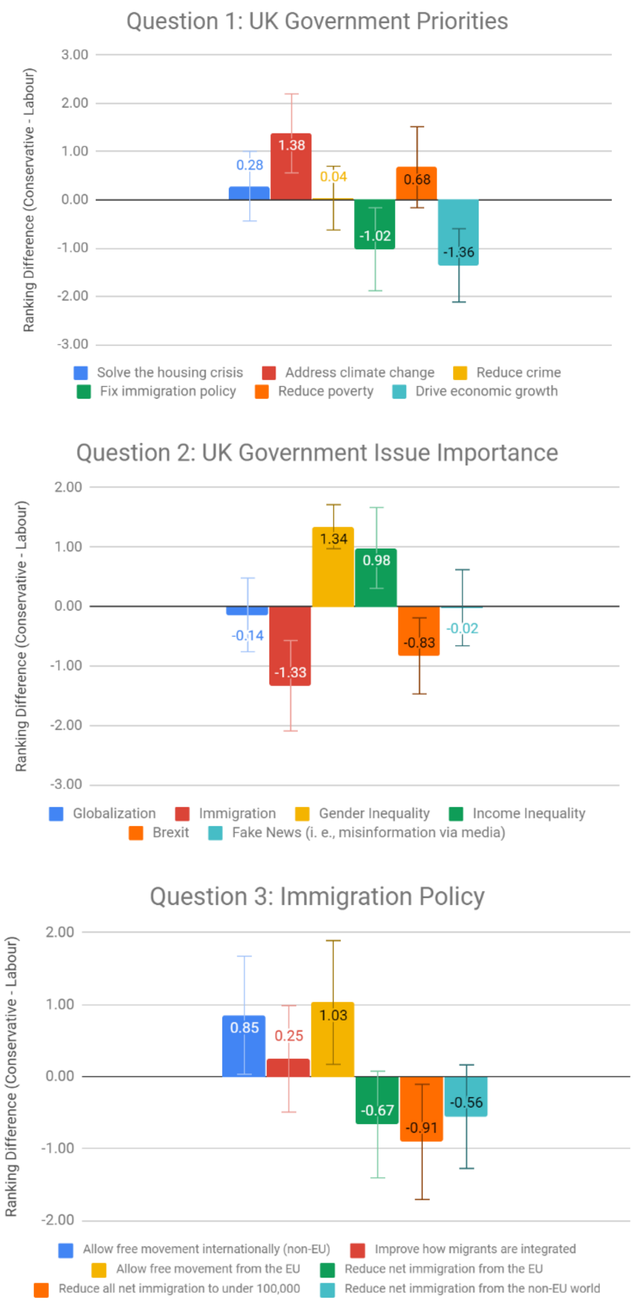

Figs. 4-6: Average Ranking Differences between Conservative and Labour Voters across each question in the study. 95% confidence intervals are shown.

C. Out-Group Satisfaction

Next, the average out-group satisfaction ranking with each of the three decision-making methods was calculated for each question. Then, the average rank of each of the three decision-making methods was calculated across questions. As shown in Appendix C, the ASI method outperformed the Majority method, resulting in superior satisfaction ranking for three of the four questions.

Interestingly, the ASI and Borda Count methods were the most-favored method in two of the groups considered. In the first group, which consisted of 4 Labour and 4 Conservative members, the ASI method outperformed the majority and Borda count methods, though the difference between the satisfaction rankings was not significant when measured with a paired t-test. In the second group, which consisted of 10 Labour and 10 Conservative members, the Borda count significantly outperformed the ASI and Majority methods ($p=0.015$, $p=5.9E-5$ respectively). In the third group, which consisted of 8 Conservative and 12 Labour members, the Borda count significantly outperformed the ASI and Majority methods ($p=0.0026$, $p=0.0040$ respectively). In the fourth group, which consisted of 12 Conservative and 8 Labour members, the ASI method significantly outperformed the Borda Count and Majority methods ($p=0.047$, $p=0.0087$ respectively).

Over all groups, the rankings generated by the ASI and Borda count methods both significantly outperformed those arising from the Majority method ($p=2.28E-5$, $p=1.66E-7$ respectively). When comparing the ASI and the Borda Count methods, the Borda Count method marginally outperformed the Swarm method, though this effect was not significant ($p=0.233$).

D. In-Group vs Out-Group Satisfaction

It is instructive to compare the satisfaction of the in-group with that of the out-group: the participants who took part in ranking the policies (the in-group) may be more satisfied with the rankings they generated than the participants who had no say (the out-group). The in-group and out-group satisfaction rankings of each decision method can be compared by subtracting each decision method's average ranking in the out-group from the same ranking in the in-group, as shown in Appendix D. The ASI was the only decision-method that was on average ranked as more preferable by the out-group than the in-group; however, the size of this effect was small.

E. Sampling Bias

Often, policy decisions are not made with a representative group of decision-makers: the elected officials that make policy decisions may be composed of different demographics than the electorate whom the officials represent. The effect where one group is sampled from a wider population, but ends up with a different demographic composition, may be referred to as sampling bias.

To investigate the impact of sampling bias of the in-group on these results, the satisfaction of the out-group was measured when composed of varying ratios of Labour and Conservative participants. Ratios between 25% Labour / 75% Conservative and 75% Labour / 25% Conservative were tested in 5% increments. To create a robust estimate of the out-group's satisfaction with a particular political composition, participants in the out-group were resampled with replacement 1000 times.

The average ranking of each method's lists across all questions is shown for groups 1-4 in Figures 4-7. The 60% confidence intervals of the average ranking of each method were calculated to give a sense of the uncertainty of the result, and is shown as the shaded interval around each line.

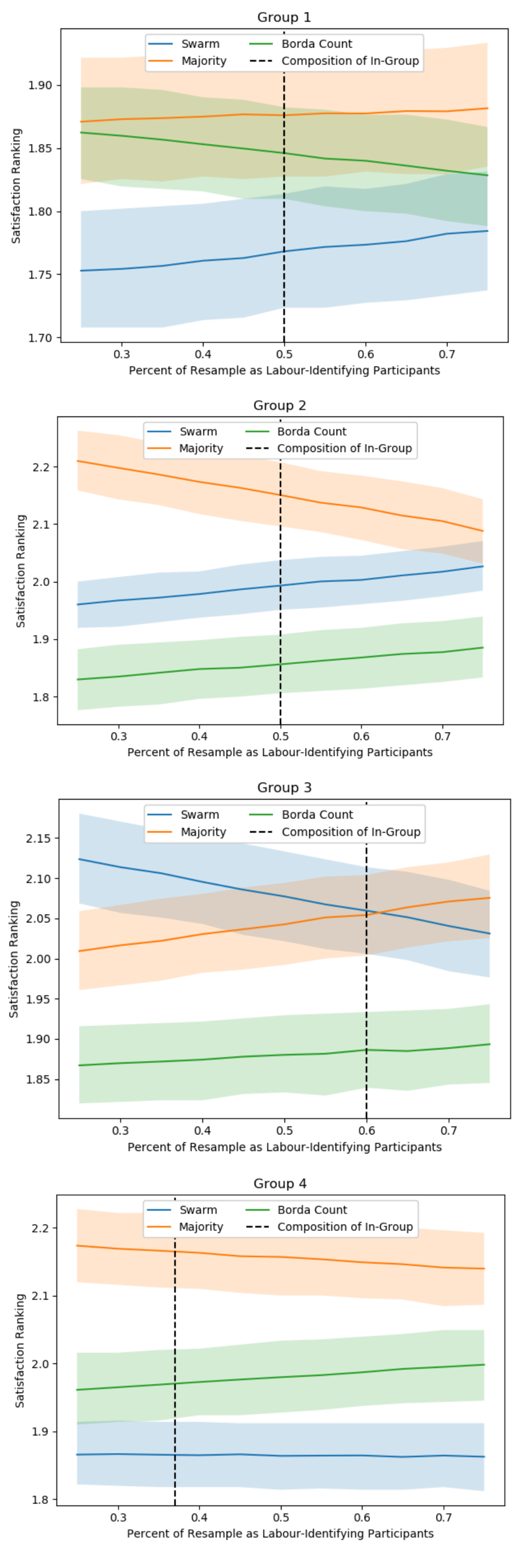

Figures 7-10: Bootstrapped Average Satisfaction Ranking of Out-Group, by the Percent of the Resample Identifying as Labour Voters. Shaded areas indicate 60% confidence intervals.

The 90% confidence intervals of these same graphs are included in Appendix E.

In these graphs, a decision method with a positive slope indicates that Labour voters favored that method's lists more than Conservative members - so including more Labour voters as part of the Out-Group sample led to a higher average group satisfaction with that method's rating. There's no clear trend from these charts that any decision method was favored by either Labour or Conservative members across the board.

There is, however, weak evidence that the preference of the out-group with the results of the ASI and Borda Count methods (as compared to the Majority method) are stable in the face of out-group sampling bias: the optimal decision method for each question is unchanged even when the Out-group's composition is changed by 25% or more. In addition, in only one case--when group 3's resampled out-group had fewer Labour voters than the in-group--did the Majority method outperform the ASI method.

VI. Other Relevant Analyses

A two-sample t-test was used to measure the effect that these demographic labels had on the in-group's ranking of each of the items in each of the three questions in this study. Table 3 below shows a breakdown of the number of items which were ranked significantly differently ($p<0.05$) for each question. For the Brexit Stance column, only the Leave and Remain demographic labels were used, since there was only one person in the in-group that self-reported a Brexit Stance of "Undecided".

Question	Political Affiliation	Brexit Stance	Gender
1: UK Government Objectives	3	2	0
2: Government Issues	4	3	1
3: Immigration Policy	3	5	0

Table 3: Number of Significant Differences Observed by demographic when ranking priorities as individuals via online survey.

Although Political Affiliation was the demographic that showed the largest number of significant ranking differences overall, Brexit Stance showed the most significant differences on a single question: the Immigration Policy question found 5 of 6 (83%) of items were ranked differently by the Leave and Remain members of the in-group. Gender was not meaningfully related to the ranking of items on this test. Only one item ("Globalization" in the Government Issues questions) was ranked significantly differently by men and women in this test.

A Chi-squared analysis was then conducted to measure whether each self-reported demographic was indeed independent of one another. Over the 67 members of the in-group, Gender was not significantly correlated with either Political Affiliation or Brexit Stance. Political Affiliation, however, was significantly correlated with the Brexit Stance of participants ($p=0.018$): Conservative participants made up 66% of participants who identified as wanting to Leave the EU, while Labour participants made up 75% of participants who identified as wanting to Remain in the EU.

VII. CONCLUSIONS

Over the three politically divisive questions in this study, and over the four groups that answered these questions, the Majority voting method was regularly the least preferred voting method. This was a surprising finding, as the Majority voting method is the most common method of aggregating voters' preferences in modern democracies. The Borda Count and ASI methods were similarly preferable in most groups, though the ASI method had slightly higher satisfaction levels among the out-group. This result suggests that, of these three methods, the Borda Count and ASI methods may be promising candidates for general public votes and for representative voting structures, which by their nature are subject to a degree of sampling bias.

Future work could collect more data with greater power to detect differences between the methods. Other studies could take steps beyond this study, investigating the group-level and individual psychological effects of using ASI and survey-based voting methods to make group decisions or prioritizations: does ASI enable individuals to feel a higher level of buy-in to the group's decisions than survey methods? Are the decisions executed satisfactorily more frequently when decisions are made on the Swarm AI platform? Do people feel like their views are more represented in the group's decision when using ASI vs traditional methods? Anecdotal evidence suggests this may be the case, but a rigorous study has yet to be conducted.

Another interesting avenue for future work includes comparing Swarm AI to a real-time iterative Majority voting framework [26]—since Swarm AI can be seen as a continuous voting framework. It would also be interesting to consider the effect on the results of deliberation before voting, since deliberation would open the doors up to both collective reasoning and strategic voting of many kinds. Other voting systems that take account of the full ordering of preferences, such as Preferential Voting, may generate equivalent levels of satisfaction.

Finally, future work may compare ASI to a Delphi method. Delphi methods use iterated rounds of deliberation and voting to refine group consensus [29]. While on the surface Delphi and ASI methods may sound similar, there are deep structural differences in the two methods that make the comparison interesting for future research: Delphi methods often take weeks to perform, but enable groups to directly deliberate and discuss problems [30], while ASI requires less than a minute for a group of any size to reach an answer, and features less language-based deliberation.

ACKNOWLEDGMENT

Thanks to Chris Hornbostel for his efforts in coordinating the swarms. Also, thanks to Unanimous AI for the use of the Swarm platform for this ongoing work. Contact Unanimous AI or the authors for more information on Swarm AI and to use this technology in a research capacity. This work was funded by NESTA.

REFERENCES

[1] Galton, F. (1907). Vox Populi. Nature, 75, 450-451.

[2] Steyvers, M., Lee, M.D., Miller, B., & Hemmer, P. (2009). The Wisdom of Crowds in the Recollection of Order Information. In Y. Bengio and D. Schuurmans and J. Lafferty and C. K. I. Williams

[3] Philip E. Tetlock and Dan Gardner. 2015. Superforecasting: The Art and Science of Prediction. Crown Publishing Group, New York, NY, USA.

[4] J Dana, P Atanasov, P Tetlock, B Mellers (2019), Are markets more accurate than polls. The surprising informational value of "just asking." Judgment and Decision Making 14 (2), 135-147

[5] Rosenberg, L.B., "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[6] Rosenberg, Louis. "Artificial Swarm Intelligence vs Human Experts," Neural Networks (IJCNN), 2016 International Joint Conference on. IEEE. J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.

[7] Rosenberg, Louis. Baltaxe, David and Pescetelli, Nicollo. "Crowds vs Swarms, a Comparison of Intelligence," IEEE 2016 Swarm/Human Blended Intelligence (SHBI), Cleveland, OH, 2016, pp. 1-4.

[8] Baltaxe, David, Rosenberg, Louis and N. Pescetelli, "Amplifying Prediction Accuracy using Human Swarms", Collective Intelligence 2017. New York, NY ; 2017.

[9] Willcox G., Rosenberg L., Askay D., Metcalf L., Harris E., Domnauer C. (2020) Artificial Swarming Shown to Amplify Accuracy of Group Decisions in Subjective Judgment Tasks. In: Arai K., Bhatia R. (eds) Advances in Information and Communication. FICC 2019. Lecture Notes in Networks and Systems, vol 70. Springer, Cham

[10] L. Rosenberg, N. Pescetelli and G. Willcox, "Artificial Swarm Intelligence amplifies accuracy when predicting financial markets," 2017 IEEE 8th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON), New York City, NY, 2017, pp. 58-62.

[11] L. Rosenberg and G. Willcox, "Artificial Swarm Intelligence vs Vegas Betting Markets," 2018 11th International Conference on Developments in eSystems Engineering (DeSE), Cambridge, United Kingdom, 2018, pp. 36-39

[12] L. Rosenberg, M. Lungren, S. Halabi, G. Willcox, D. Baltaxe and M. Lyons, "Artificial Swarm Intelligence employed to Amplify Diagnostic Accuracy in Radiology," 2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), Vancouver, BC, 2018, pp. 1186-1191.

[13] Seeley T.D, Buhrman S.C 2001 "Nest-site selection in honey bees: how well do swarms implement the 'best-of-N' decision rule?" Behav. Ecol. Sociobiol. 49, 416–427

[14] Marshall, James. Bogacz, Rafal. Dornhaus, Anna. Planqué, Robert. Kovacs, Tim. Franks, Nigel. "On optimal decision-making in brains and social insect colonies." Soc. Interface 2009.

[15] Seeley, Thomas D., et al. "Stop signals provide cross inhibition in collective decision-making by honeybee swarms." Science 335.6064 (2012): 108-111.

[16] Seeley, Thomas D. Honeybee Democracy. Princeton Univ. Press, 2010.

[17] Seeley, Thomas D., Visscher, P. Kirk. "Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making." Behavioural Ecology and Sociobiology 54 (5) 511-520.

[18] Usher, M. McClelland J.L 2001 "The time course of perceptual choice: the leaky, competing accumulator model." Psychol. Rev. 108, 550–592

[19] L. Rosenberg, G. Willcox, "Artificial Swarms find Social Optima: (Late Breaking Report)", 2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA), pp. 174-178, 2018.

[20] Laslier JF. (2012) And the Loser Is… Plurality Voting. In: Felsenthal D., Machover M. (eds) Electoral Systems. Studies in Choice and Welfare. Springer, Berlin, Heidelberg

[21] Poundstone, W. (2009): "Gaming the Vote: Why Elections Aren't Fair (and What We Can Do About It)" pp. 79, 206, 230.

[22] Condorcet, Jean-Antoine-Nicolas de Caritat, marquis de, 1743-1794 “Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix", Pre-1801 Imprint Collection (Library of Congress) DLC

[23] Merijn Van Erp and Lambert Schomaker (2000) Variants Of The Borda Count Method For Combining Ranked Classifier Hypotheses. In the seventh Intenational workshop on frontiers in handwriting recognition 443–452

[24] L. Rosenberg, D. Baltaxe, “Setting Group Priorities - Swarms vs Votes” (2016) IEEE Conference on Swarm Human Blending

[25] Kolm SC. (1993) The Impossibility of Utilitarianism. In: Koslowski P., Shionoya Y. (eds) The Good and the Economical. Studies in Economic Ethics and Philosophy. Springer, Berlin, Heidelberg

[26] Burgman, M. A., Regan, H. M., Maguire, L. A., Colyvan, M., Justus, J., Martin, T. G., & Rothley, K. (2014). Voting systems for environmental decisions. Conservation biology, 28(2), 322-332.

[27] Gibbard, Allan (1973). "Manipulation of voting schemes: A general result". Econometrica. 41 (4): 587–601. doi:10.2307/1914083. JSTOR 1914083.

[28] Satterthwaite, Mark Allen (April 1975). "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions". Journal of Economic Theory. 10 (2): 187–217. CiteSeerX 10.1.1.471.9842. doi:10.1016/0022-0531(75)90050-2.

[29] Pill, Juri (Febuary 1971). “The Delphi method: Substance, context, a critique and an annotated bibliography”. Socio-Economic Planning Sciences. 5 (1): 57-71.

[30] Boulkedid, R., Abdoul, H., Loustau, M., Sibony, O., & Alberti, C. (2011). Using and reporting the Delphi method for selecting healthcare quality indicators: a systematic review. PloS one, 6(6), e20476. doi:10.1371/journal.pone.0020476

[31] Rosenberg, L., Willcox, G. (2019). Artificial Swarm Intelligence. Unanimous AI, https://www.unanimous.ai/whitepaper

[32] How does Swarm work? -. (n.d.). Retrieved from https://unanimous.ai/what-is-si/

APPENDIX A: PRIORITIZATION QUESTIONS

1) Rank the following UK Government objectives in order of their importance (1=most important, 2=next most important, etc.).
*Address climate change
*Drive economic growth
*Fix immigration policy
*Reduce crime
*Reduce poverty
*Solve the housing crisis.

2) Rank the following issues in order of the priority the UK government should give them. (1=highest priority, 2=second highest priority, etc.)
*Brexit
*Fake News (misinformation in media)
*Gender inequality
*Globalization
*Immigration
*Income inequality

3) Rank the following proposed [immigration] policies in order of your preference. (1=most preferred, 2=next most preferred, etc.)
*Allow free movement from the EU
*Allow free movement internationally (non EU)
*Improve how migrants are integrated
*Reduce all net immigration to under 100,000
*Reduce net immigration from the EU
*Reduce net immigration from the non-EU world

APPENDIX B: IN-GROUP SATISFACTION RANKING BY QUESTION NUMBER

	In-Group Satisfaction Ranking by Question Number								
	Question 1 (UK Govt. Objectives)			Question 2 (Government Issues)			Question 3 (Immigration)		
	Swarm	Borda	Majority	Swarm	Borda	Majority	Swarm	Borda	Majority
Group 1	1.75	2.13	2.13	1.25	1.75	1.25	1.88	1.88	2.25
Group 2	1.95	1.74	2.32	2.05	1.74	2.21	2.05	1.95	2.00
Group 3	1.90	2.15	1.95	1.65**	1.85*	2.50	1.95	1.95	2.10
Group 4	2.00	2.11	1.89	1.68*	1.89	2.42	2.11	1.74	2.16
All Groups	**1.92**	2.02	2.06	**1.73****	1.82*	2.24	2.02	**1.88**	2.11

*=(p<0.05) as compared to the Majority Rating, **=(p<0.01) as compared to the Majority Rating

Appendix C: Average Ranking of Decision Methods in Out-Group Satisfaction Survey

	# Conservative	# Labour	Swarm	Majority	Borda
Group 1	4	4	**1.77**	1.87	1.85
Group 2	10	10	1.99	$2.15^{\Psi,**}$	$\mathbf{1.87}^{\Omega}$
Group 3	8	12	2.07	2.04^{**}	$\mathbf{1.89}^{\Psi}$
Group 4	12	7	$\mathbf{1.87}^{*}$	$2.14^{\Psi,*}$	1.99

$^{\Omega}$ = significant at the 0.05 level relative to the Swarm
$^{\Psi}$ = significant at the 0.01 level relative to the Swarm
* = significant at the 0.05 level relative to the Borda Count
**=significant at the 0.01 level relative to the Borda Count

Appendix D: Difference of Out-Group and In-Group Average Satisfaction Rankings

	Swarm	Majority	Borda
Group 1	0.14	0.00	-0.07
Group 2	-0.02	-0.03	0.05
Group 3	0.15	-0.14	-0.17
Group 4	-0.19	0.00	0.07
Average	0.02	-0.04	-0.03

Appendix E: Bootstrapped Average Satisfaction Ranking of Out-Group, by the Percent of the Resample Identifying as Labour Voters. Shaded areas indicate the 90% confidence interval around each method's satisfaction.