



Wise groups and humble persons: the best of both worlds?

Mattias Skipper¹

Received: 30 September 2022 / Accepted: 20 December 2022
© The Author(s) 2023

Abstract

This paper is about a problem that can arise when we try to harness the “wisdom of the crowd” from groups comprised of individuals who exhibit a certain kind of epistemic humility in the way they respond to testimonial evidence. I begin by setting out the problem and then make some initial steps toward solving it. The solution I develop is tentative and may not apply in all circumstances, but it promises to alleviate what seems to me to be a problem of both theoretical interest and practical importance.

Keywords Wisdom of the crowd · epistemic humility · testimony · expert deference

1 The GroupWise Show

Every Saturday night, the popular game show “GroupWise” is broadcasted on the Veritas Channel. The rules of the game are simple: there are two teams, each with 99 members, one of whom serves as Team Captain (the teams are advised to pick as their Team Captain someone with excellent general knowledge). The game show host asks a series of yes–no questions one at a time. Once a question has been asked, each Team Captain gets sixty seconds to think about the question before they must write down an answer on a piece of paper and place it in their team’s ballot box. The rest of the team members then get a chance to see their Team Captain’s answer before they themselves must write down an answer and place it in their team’s ballot box. The answer with the most votes counts as the team’s answer. And the team with the most correct answers in the end wins the much-coveted GroupWise Trophy.

This week’s contestants call themselves the *Humbleteers* and the *Intransigents*. The two teams happen to be remarkably similar: not only is the Team Captain on each team considerably more reliable than the rest of the team members when it comes to answering trivia questions; the Team Captains are also *equally* reliable, and the remaining 196 ordinary team members are all equally reliable as well. The

✉ Mattias Skipper
skipper.mattias@gmail.com

¹ Department of Philosophy, Inland Norway University of Applied Sciences, Lillehammer, Norway

only relevant difference between the two teams concerns their attitude towards testimonial evidence. The Humbleteers see testimony as a valuable source of epistemic improvement. When they learn what someone else believes about a given matter, they are prepared to revise their own beliefs on the matter, at least to some extent, depending on how reliable they expect that other person to be, and depending on how much that other person's belief differs from their own. By contrast, the Intransigents are adamant in wanting to figure things out for themselves. When they learn what other people believe, they do not let it affect their own beliefs at all, not even if they have every reason to expect those other people to be highly reliable, and not even if those other people's beliefs differ greatly from their own.

Clearly, the Humbleteers have a much more reasonable attitude towards testimonial evidence than do the Intransigents. It is indisputable, I take it, that a substantial part of our evidence comes from testimony, and that each of us would be in a much worse epistemic situation if we were to cut ourselves off from it. There are, of course, difficult questions about how one should respond to *certain kinds* of testimony, such as testimony about moral matters or testimony from people who are (or who you consider to be) your "epistemic peers" on a given issue.¹ But regardless of one's stance on such controversial matters, it should be uncontroversial that a blanket rejection of testimony as a source of evidence is doomed to make us epistemically worse off.

Nonetheless, if each team member goes ahead and votes for whichever answer they consider more likely to be correct at the time of casting their vote, the Intransigents will stand a better chance than the Humbleteers of winning the GroupWise Trophy. The reason for this will be evident to readers who are familiar with Condorcet's jury theorem and related results from the judgment aggregation literature, but for those who are not, the crux of the matter is that the Intransigents effectively vote *independently* of each other, whereas the Humbleteers do not. When the ordinary members of the Humbleteers learn what their Team Captain voted, they each come to place more credence in their Team Captain's answer than in whatever answer they themselves might have been initially inclined towards, since they each trust the Team Captain's judgment more than their own. As a result, they all end up voting for whatever the Team Captain voted, which means that their reliability as a group will be identical to that of the Team Captain. By contrast, when the ordinary members of the Intransigents learn what their Team Captain voted, they do not let it affect their beliefs at all, with the result that they all vote for whatever answer they were initially inclined towards. And when a large number of people vote independently of each other, it is well known that the majority will often end up being more reliable than even the most reliable individual. This is one manifestation of what has become known as the "wisdom of the crowds."

Here is one example of how the Intransigents can end up more reliable than the Humbleteers:

Example 1.1. Suppose that the two team captains are 80% reliable and that the remaining 196 ordinary team members are 60% reliable. Then, the Humbleteers

¹ For discussion of moral deference, see Enoch (2014), Howell (2014), and Skarsaune (2016). For discussion of peer disagreement, see Kelly (2005; 2010), Christensen (2007), and Elga (2007).

have a collective reliability of 80%, whereas the Intransigents have a collective reliability of approximately 98%.²

For this to happen—that is, for the Intransigents to end up more reliable than the Humbleteers—a few background conditions must be satisfied. First, and most obviously, the ordinary team members cannot be *anti-reliable*: they must be more than 50% likely to vote for the correct answer, or else the majority will invariably fail to be more reliable than the Team Captain.³ Second, even if the ordinary team members are more than 50% reliable, the Team Captain cannot be *too much* more reliable than the ordinary team members, or else the “wisdom of the crowd” effect will not be strong enough to render the majority more reliable than the Team Captain. The following example illustrates what happens if this second condition is not satisfied:

Example 1.2. Suppose that the two team captains are 90% reliable and that the ordinary team members are 55% reliable. Then, the Humbleteers have a collective reliability of 90%, whereas the Intransigents have a collective reliability of only approximately 86%.⁴

So, there is no guarantee that the “wisdom of the crowd” effect will be strong enough to render the Intransigents more reliable than the Humbleteers. Nonetheless, under a wide range of conditions, the “wisdom of the crowd” effect will indeed be strong enough to render the Intransigents more reliable than the Humbleteers.

There is something a bit tragic about this result. One might have hoped that the Humbleteers would be rewarded for exhibiting what seems like an eminently reasonable kind of epistemic humility in their response to testimonial evidence, whereas the Intransigents would be punished for not doing so. Yet, the exact opposite turns out to be the case: the Intransigents are the ones who get to harness the “wisdom of the crowd” in a wide range of cases, whereas the Humbleteers do not.

The example of the GroupWise Show is supposed to illustrate a more general problem that can arise when a group of people need to form a collective judgment about some factual matter. On the one hand, it seems clear that we as individuals have much to gain from treating testimony as a valuable source of evidence. On other hand, if we do so, we risk compromising the independence between our judgments, and hence risk compromising our collective ability to take advantage of the “wisdom of the crowd.” So, we sometimes seem forced to choose between two evils: being intransigent in our response to testimonial evidence or being unable to harness the “wisdom of the crowd.”

² The collective reliability of the Intransigents can be calculated by multiplying the probability that the Team Captain votes correctly by the probability that at least 49 ordinary members vote correctly, and then adding the probability that the Team Captain votes incorrectly multiplied by the probability that at least 50 ordinary members vote correctly:

$$.8 \cdot \sum_{i=49}^{98} \frac{98!}{i!(98-i)!} \cdot .6^i (1-.6)^{98-i} + (1-.8) \cdot \sum_{i=50}^{98} \frac{98!}{i!(98-i)!} \cdot .6^i (1-.6)^{98-i} \approx .98.$$

³ More generally, in cases where the team members are allowed to have different levels of reliability, their *average* reliability must exceed 50% (cf. Owen et al., 1989).

⁴ The collective reliability of the Intransigents is calculated in the same way as in footnote 2.

Of course, in ordinary life, we rarely find ourselves in situations as extreme as that of the Humbleteers and the Intransigents. Most of us are not completely intransigent or completely humble in our response to testimonial evidence. But the same problem can arise in more realistic versions of the case. So long as the Intransigents are at least slightly less willing than the Humbleteers to revise their beliefs in response to testimonial evidence, they will exhibit a higher degree of independence and hence be in a better position to harness the “wisdom of the crowd.” As the difference between the two teams becomes smaller, the range of circumstances in which the Intransigents outperform the Humbleteers obviously narrows to some extent. But even if most real-world instances of the problem are less pronounced than in the GroupWise Show, it hardly seems safe to dismiss the problem as a mere theoretical curiosity.

The question, then, is whether there is anything we can do to solve or alleviate the problem. I will argue that there is. But before doing so, I want to situate the problem within a slightly broader context.

Several epistemologists and philosophers of science have argued, in different ways, that the epistemic performance of a group can come systematically apart from the epistemic performance of the individuals that comprise it. For example, some have argued that scientists who irrationally stick to their own theories beyond what the evidence warrants may thereby help to prevent the scientific community at large from abandoning those theories prematurely (Feyerabend, 1965; Hull, 1988; Zollman, 2010). Others have argued that biased individuals may under the right conditions combine their judgments in such a way as to create unbiased group judgments (Goodin, 2006). And yet, others have argued that diverse groups of relatively unreliable problem solvers may outperform less diverse groups of comparatively reliable problem solvers (Hong & Page, 2001, 2004).⁵

The phenomena described by these authors are, at least in some respects, more subtle and unobvious than the one exemplified by the GroupWise Show. Nonetheless, the problem as I have described it seems to have gone largely unnoticed in the literature. The closest discussion, to my knowledge, is due to Hazlett (2016), who argues that intransigent (or “non-deferential”) belief is *pro tanto* socially valuable because it facilitates the kind of independence required for “wisdom of the crowd” effects to arise.⁶ This claim of Hazlett’s is very much congenial to what has been said so far. But the point remains that just as non-deferential belief can have beneficial effects at the group level, deferential belief can have beneficial effects at the individual level. This is the tension which interests me here, and which the GroupWise Show is supposed to illustrate.

2 A solution

One reaction to what has been said so far is to concede that we sometimes have to make a trade-off between being humble in our response to testimonial evidence and being able to harness “wisdom of the crowd” effects, and then go on to

⁵ Further examples along similar lines can be found in Mayo-Wilson et al. (2011), Skipper & Steglich-Petersen (2021), and Daoust (forthcoming).

⁶ A similar point can be found in Dellsén (2020), who argues that a certain kind of autonomy among experts in a given domain is *pro tanto* valuable because it facilitates independence between the experts, which in turn raises the probability that the expert consensus (if such a consensus exists) is correct.

explore how this trade-off is best made. This may be the right response in some cases. But I want to suggest that a more attractive response is sometimes available, one that allows groups of epistemically humble individuals to harness the “wisdom of the crowd” to the full extent.

The first thing to observe is that, in setting out the problem above, I assumed that each individual member voted for whichever answer they considered more likely to be correct at the time of casting their vote. This is what allowed us to establish a connection between, on the one hand, whether the individual members were epistemically humble and, on the other hand, whether their votes were independent. It is natural to think, then, that if we can find a way to sever the connection between what the individual members believe and how they vote, we might be able to avoid the problem altogether. The question is whether there is a principled way of doing so. What are the members to base their votes on if not their beliefs?

We can find the beginning of an answer by looking back at how the Humbleteers arrived at their votes. Above, I said that when the ordinary members of the Humbleteers learn what their Team Captain voted, they each come to trust their Team Captain’s answer more than whatever answer they themselves might have initially been inclined towards. But even after the Humbleteers have updated their beliefs in light of their Team Captain’s answer, they might still be aware of (or be able to recall) which answer they themselves judged more likely to begin with. And if they base their votes on their initial judgments rather than their beliefs at the time of casting their vote, they will obtain the same level of independence as the Intransigents, and hence, be in just as good a position to harness the “wisdom of the crowd.”

The idea that we can distinguish between what an agent believes, all things considered, and what the agent judges to be the case, independently of other people’s influence, has recently come up in several places, especially in the epistemological literature. For example, Barnett (2019) appeals to a notion of “disagreement-insulated inclination” to shed light on the role of disagreement within philosophy; Skipper and Steglich-Petersen (2021) appeal to an agent’s “first-order judgment” to address a problem about conciliatory responses to in-group disagreement; and Worsnip (forthcoming) appeals to an agent’s “personal take” on an issue to address a problem that can arise in communities where some members are conciliatory in their response to a disagreement while others are steadfast. Despite some differences in how these notions are characterized and motivated, they are all meant to capture a version of the idea that in addition to asking what people believe, all things considered, we can meaningfully ask what they judge to be the case, bracketing certain kinds of evidence, such as evidence of disagreement or testimonial evidence in general.⁷

The thought I want to pursue, then, is this: when we find ourselves in situations like that of the Humbleteers, we should base our votes not on what we *believe* to be the case, all things considered, but on what we *judge* to be the case, independently

⁷ See Barnett (ms) for further discussion of how to understand this distinction, and for additional examples of how the distinction might be put to use.

of other people's influence. Doing so, the hope is, will allow us to harness the "wisdom of the crowd" while remaining epistemically humble in our response to testimonial evidence.

Although I believe this thought is on the right track, it stands in need of refinement. First of all, it cannot be right to say that the Humbleteers should bracket *all* testimonial evidence when casting their vote. Suppose that each member of the Humbleteers has access to their own external advisor, who is highly reliable, and who is not going to partake in the vote. If the members bracket the testimonial evidence coming from such an external advisor, they miss out on an opportunity to greatly improve the reliability of their votes at no risk of compromising the independence between them. That is clearly not a good idea.

On the other hand, it also cannot be right to say that the Humbleteers should *only* bracket testimonial evidence from people who partake in the vote. Suppose that the members have access to the *same* external advisor (rather than different ones). In this case, including the testimonial evidence from the external advisor will completely undermine the independence between the votes, thereby leaving the Humbleteers in no position to harness the "wisdom of the crowd." That is clearly not a good idea either.

A more promising suggestion is to say that the Humbleteers should bracket a piece of testimony if, and only if, including it compromises the independence between their votes. This comes closer to the mark, but it still cannot be quite right. Suppose that we divide the Humbleteers into pairs and assign each pair its own external advisor. Then, if the members bracket the testimonial evidence coming from these external advisors, they miss out on an opportunity to greatly improve the reliability of their votes while only compromising the independence between their votes to a comparatively small degree. This will at least sometimes have a negative overall effect on the group's reliability, as illustrated by the following example:

Example 2.1. Consider a group of 100 voters, each of whom is 55% reliable. The voters are divided into pairs, and each pair is assigned an external advisor who is 90% reliable. Given this, if the voters bracket the testimony from the external advisors, the group's reliability will be equivalent to that of a hundred independent voters, each of whom is 55% reliable: $\sum_{i=51}^{100} \frac{100!}{i!(100-i)!} .55^i (1 - .55)^{100-i} \approx 82\%$.

By contrast, if the voters include the testimony coming from the external advisors, the group's reliability will be equivalent to that of a group of fifty independent voters, each of whom is 90% reliable, which is much higher: $\sum_{i=26}^{50} \frac{50!}{i!(50-i)!} .9^i (1 - .9)^{50-i} \approx 99.9\%$.

This leaves us with the following proposal: the Humbleteers should bracket a piece of testimonial evidence if, and only if, the negative effect of including it, which results from the compromised independence between the votes, outweighs the positive effect of including it, which results from the improved reliability of the votes. The same holds whenever a group of people need to form a collective judgment on some factual matter, at least if they want to harness the "wisdom of the crowd."




Fig. 1 Comparison of the strategies of including vs. bracketing the Team Captain's testimony, holding fixed the reliability of the ordinary team members (at 55%) and varying the reliability of the Team Captain

We can think of this proposal as telling us to balance two opposing effects against each other. On the one hand, including a piece of testimonial evidence may increase the reliability of one's vote, which is beneficial to the group's reliability (other things being equal). On the one hand, bracketing that same piece of testimony facilitates voting independence and hence puts the group in a better position to harness the "wisdom of the crowd." And the relative strength of these opposing effects determines which voting strategy performs best in a given context.

Here is an example of how this balancing act might play out in the case of the Humbleteers:

Example 2.2. Suppose that the ordinary members of the Humbleteers are 55% reliable and let x be the reliability of the Team Captain ($55\% < x \leq 100\%$). Then, for sufficiently low values of x ($x \lesssim 86\%$), the strategy of bracketing the Team Captain's testimony results in a higher group reliability than the strategy of including it. Conversely, for sufficiently high values of x ($x \gtrsim 86\%$), the strategy of including the Team Captain's testimony results in a higher group reliability than the strategy of bracketing it. And, for $x \approx 86\%$, the two strategies result in the same group reliability (see Fig. 1 for an illustration).

The thing to observe in this example is that when the Team Captain is relatively unreliable in comparison to the ordinary team members, the strategy of bracketing the Team Captain's testimony performs better than the strategy of including it,

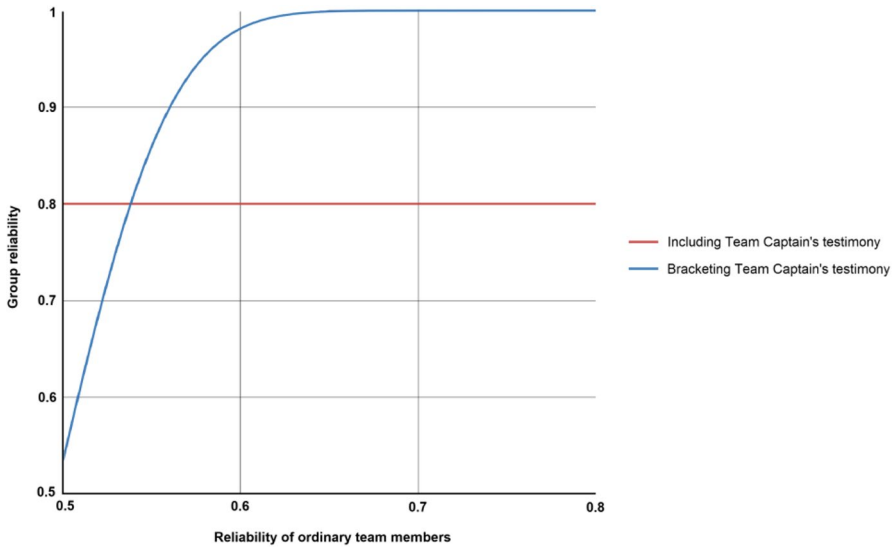


Fig. 2 Comparison of the strategies of including vs. bracketing the Team Captain’s testimony, holding fixed the reliability of the Team Captain (at 80%) and varying the reliability of the ordinary team members

since the “wisdom of the crowd” effect is strong enough to outweigh the opposing effect stemming from the compromised reliability of the individual votes. But as we increase the Team Captain’s reliability, the balance eventually tips in favor of including the Team Captain’s testimony, since the “wisdom of the crowd” effect no longer outweighs the opposing effect stemming from the compromised reliability of the individual votes.

We can illustrate the same mechanism with a different example in which we vary the reliability of the team members while keeping the Team Captain’s reliability fixed:

Example 2.3. Suppose that the Team Captain of the Humbleteers is 80% reliable and let x be the reliability of the ordinary team members ($50\% \leq x < 80\%$). Then, for sufficiently low values of x ($x \lesssim 54\%$), the strategy of including the Team Captain’s testimony results in a higher group reliability than the strategy of bracketing it. Conversely, for sufficiently high values of x ($x \gtrsim 54\%$), bracketing the Team Captain’s testimony results in a higher group reliability than including it. And, for $x \approx 54\%$, the two strategies result in the same group reliability (see Fig. 2 for an illustration).

Again, the thing to observe here is that when the team members are relatively unreliable in comparison to the Team Captain, the strategy of including the Team Captain’s testimony performs better than the strategy of bracketing it, since the “wisdom of the crowd” effect is not strong enough to outweigh the opposing effect stemming from the compromised reliability of the individual votes. But as we increase

the team members' reliability, the balance eventually tips in favor of bracketing the Team Captain's testimony, since the "wisdom of the crowd" effect becomes strong enough to outweigh the opposing effect stemming from the compromised reliability of the individual votes.

Needless to say, these examples are highly idealized. In the real world, it may be much more difficult to determine whether bracketing a piece of testimony will have a net positive effect on a group's reliability. This is so not least because social dependence relations—how people in large groups or networks communicate and interact—can often be difficult to gauge. Nonetheless, I take the preceding examples to show that the central proposal of the paper does say something informative about when and when not to bracket a piece of testimony before casting one's vote. This may amount to little more than a proof of concept. But if the proof is sound, it gives us grounds for optimism that we can, at least sometimes, get the best of both worlds: harness the "wisdom of the crowd" while treating testimonial evidence with due epistemic humility.

Funding Open Access funding provided by Inland Norway University Of Applied Sciences

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Barnett, Z. (2019). Philosophy without belief. *Mind*, 128, 109–138.
- Barnett, Z. (ms). Five roles for inclination. Unpublished manuscript.
- Christensen, D. (2007). Epistemology of disagreement: The good news. *The Philosophical Review*, 116, 187–217.
- Daoust, M-K. (forthcoming). Optimizing individual and collective reliability: a puzzle. *Social Epistemology*. Online first.
- Dellsén, F. (2020). The epistemic value of expert autonomy. *Philosophy and Phenomenological Research*, 100, 344–361.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41, 478–502.
- Elgin, C. (2010). Persistent disagreement. In R. Feldman & T. A. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.
- Enoch, D. (2014). A defense of moral deference. *The Journal of Philosophy*, 111, 229–258.
- Feyerabend, P. (1965). Problems of empiricism. In Colodny (ed.), *Beyond the Edge of Certainty. Essays in Contemporary Science and Philosophy*, pp. 154–260. Prentice-Hall.
- Goodin, R. (2006). The epistemic benefit of multiple biased observers. *Episteme*, 3, 166–174.
- Hazlett, A. (2016). The social value of non-deferential belief. *Australasian Journal of Philosophy*, 94, 131–151.
- Hong, L., & Page, S. (2001). Problem solving by heterogenous agents. *Journal of Economic Theory*, 97, 123–163.
- Hong, L., & Page, S. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101, 16385–16389.

- Howell, R. (2014). Google morals, virtue, and the asymmetry of deference. *Noûs*, 48, 389–415.
- Hull, D. (1988). *Science as a Process*. University of Chicago Press.
- Kelly, T. (2005). The epistemic significance of disagreement. *Oxford Studies in Epistemology* 1, Gendler & Hawthorne (eds.), pp. 167–96.
- Kelly, T. (2010). Peer disagreement and higher-order evidence. *Social Epistemology: Essential Readings*, Goldman & Whitcomb (eds.), Oxford: Oxford University Press, 183–217.
- Mayo-Wilson, C., Zollman, K., & Danks, D. (2011). The independence thesis: when individual and social epistemology diverge. *Philosophy of Science* 78, 653–77.
- Owen, G., Grofman, B., & Feld, S. (1989). Proving a distribution-free generalization of the Condorcet jury theorem. *Mathematical Social Sciences*, 17, 1–16.
- Skarsaune, K.-O. (2016). Moral deference and authentic interaction. *Journal of Philosophy*, 113, 346–357.
- Skipper, M., & A. Steglich-Petersen (2021). When conciliation frustrates the epistemic priorities of groups. F. Broncano-Berrocal and J.A. Carter (eds.), *The Epistemology of Group Disagreement*, Routledge.
- Worsnip, A. (forthcoming). Compromising with the uncompromising: political disagreement under non-compliance. *Forthcoming in Journal of Political Philosophy*.
- Zollman, K. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 2, 17–35.