On the Quantification of Crowd Wisdom, Optimal Crowd Size, and the Mean-Median
Dilemma

Jan Lorenz¹

¹ Department of Psychology and Methods, Jacobs University Bremen, Germany

Author Note

This research was supported in part by grants from the German Research Council:
DFG 265108307 and DFG 396901899.

Correspondence concerning this article should be addressed to Jan Lorenz, Jacobs
University, Campus Ring 1, 28759 Bremen, Germany. E-mail: post@janlo.de

Abstract

Crowd wisdom is a fascinating metaphor in the realm of collective intelligence. However, even for the simple case of estimation tasks of one continuous value, the quantification of the phenomenon lacks some conceptual clarity. Two interrelated questions of quantification are at stake. First, how we can quantify the degree of crowd wisdom in a sample and how it can be distinguished from the individual wisdom of its members? Two measures will be presented and discussed. One can also be used to quantify optimal crowd sizes. Even purely statistical, it turns out that smaller crowds are more advisable when intermediate systematic errors in estimating crowds are frequent. In such cases, larger crowds are more likely to be outperformed by a single estimator. Second, how can we best aggregate the collective decision from a sample of estimates, with the mean or the median? Arguments are not only statistical but also related to the question if democratic decision-making can have an epistemic quality. A practical result of this study is that we should usually aggregate democratic decisions by the median, but have a backup with the mean when the decision space has two natural bounds and societies polarize.

Keywords: wisdom-of-crowd indicator, fraction of outperformed estimates, epistemic democracy, collective decision, accuracy

Word count: 7565

On the Quantification of Crowd Wisdom, Optimal Crowd Size, and the Mean-Median Dilemma

Significance Statement

There is no crowd wisdom when individual estimates do not bracket the true value, but there is also no crowd wisdom when everyone knows it. We present tools to measure crowd wisdom in that sense. For collective decisions where we expect crowds to be systematically biased by more than half a standard deviation it turns out that it is often advisable to limit crowds to an optimal size which minimizes the probability that a randomly selected individual performs better than the crowd. When crowd wisdom can be expected we should use it to realize the epistemic potential of democracy by aggregating our individual estimates to make close to correct collective decisions. Standard reasoning would suggest that we should aggregate not by averaging but by taking the middlemost (or median) value of our estimates to avoid incentives for malicious misreporting. However, this rationale switches when the number to decide is between two natural bounds and estimates are polarized. We need the average to compromise in such situations. Unfortunately, this can make polarization more persistent. These results build theoretical foundations for democratic institutions of direct collective decisions in numbers.

Introduction

We speak of the wisdom-of-crowds effect (as popularized by Surowiecki, 2004) when the collective decision of a crowd is better than the decision of the individual. Nowadays, the term “crowd wisdom” is used as an explanation for the functioning of a variety of institutions like crowd sourcing, crowd funding, or even democracy, and the wisdom-of-crowd literature spans several fields of application. Cognitive and social psychology study the performance of groups in judgment and decision making (see Gigone & Hastie, 1997; Davis-Stober, Budescu, Dana, & Broomell, 2014; Laan, Madirolas, & Polavieja, 2017). Philosophy and political science develop epistemic theories of democracy

to understand if and how democratic procedures such as deliberation and voting help democratic decisions approximate a procedure-independent standard of correctness (see Cohen, 1986; Goodin & Spiekermann, 2018; Landemore & Elster, 2012). This article provides a conceptual framework to quantify crowd wisdom from a statistical perspective and extends it to questions of epistemic democracy.

In experiments and field studies, there is a large variety of problems in which crowds and individuals decide. Many empirical studies focus on the most simple case of binary or multiple discrete choice, where a group has to find the correct decision among a finite set of options (see Galesic, Barkoczi, & Katsikopoulos, 2018; Couzin et al., 2011; Frey & Rijt, 2020; Kao & Couzin, 2014; Prelec, Seung, & McCoy, 2017). Also, many theories of epistemic democracy start from binary choice building on Condorcet’s jury theorem (see List & Goodin, 2001). However, these also extend to more general discrete and, in particular, continuous choice spaces which we will treat here (Pivato, 2017).

Our focus will also be on simple one-shot situations, where a crowd has to aggregate a collective decision through sampling of estimates where we do not assume anything about individual competence, confidence, or past performance and learning of estimators. Therefore, the results here are not about selecting, weighting, or training individuals to increase crowd wisdom, nor are they about group structure (see, e.g., Golub & Jackson, 2010) or communication protocols. In the words of group judgment research, we will only look at estimates of *statisticized groups* and not at group processes.

Statisticized groups are also relevant for the question of epistemic democracy. Pivato (2011) shows that many voting procedures can be reconceptualized as statistical estimators for the “truth” in a setting where voters only have noisy signals about the true state of the world. However, democracy is not only about aggregating the most correct collective decisions but also about collective decisions under conflicting preferences of estimators. So, any sample of estimators may be prone to estimators not acting with epistemic motivation

but self-interested and strategically. This aspect was already famously introduced by Galton (1907a) calling the median a “*democratic*” aggregation rule for continuous problems while the mean would give “*voting power to ‘cranks’ in proportion to their crankiness.*”

The focus of this paper is about such wisdom-of-crowd problems of continuous values, that means, estimation tasks for continuous variables which provide more nuance than problem of discrete decision. What matters in continuous decisions is not if the decision is exactly right or wrong, but how close to correct it is to the (yet unknown) truth. The results in this paper are based on this feature of continuous spaces.¹

The seminal example is the weight-judging competition at the West of England Fat Stock and Poultry Exhibition 1906 in Plymouth reported from one of the founding father of statistics – Francis Galton (1907c). Competitors had to guess the weight of the meat of an ox after it has been slaughtered and dressed. Figure 1A shows the histogram of estimates (as Wallis (2014) extracted them from Galton’s notebook) together with the truth, the median and the mean of estimates. A similar estimation game happened at the lottery at the festival “ViertelFest” August 22-24, 2008 in Bremen where people could buy lots which could bring them prices. Each buyer got another chance to win by estimating the total number of lots sold at the festival. Further on, before the lottery the organizers asked an expert, the organizer of Bremen’s biggest lottery about his estimate for reference. Figure 1B shows the histogram of estimates, the truth, the median, the arithmetic mean analog to Galton’s data but also the geometric mean and the expert’s estimate. Appendix A provides more context and analysis of these particular estimation games. Lorenz (2021) documents the data and all production code in the Rmarkdown file of this paper.

¹ Of course, also discrete decision spaces may be equipped with certain gradual measures of goodness, e.g. modeled through utility or loss functions in statistical decision theory (Berger, 1989). The difference between discrete and continuous may be not exactly sharp from a practical perspective. Nevertheless, the typical cases of a binary and a continuous choice sets are fundamentally different, and continuous decisions allow other interesting theoretical insights about the quantification of crowd wisdom.

Both examples in Figure 1 are impressive instances of the wisdom of crowd phenomenon, but also show remarkable differences in the shape of the distribution of estimates. It is not clear how to assess if one of these crowds is wiser than the other. Further on, the Viertelfest sample shows that mean and median can produce drastically different collective estimates. In the following, we pursue these two different questions of quantification: (1) How to quantify the degree of crowd wisdom in samples which distinguishes crowd wisdom from average individual wisdom (Gigone & Hastie, 1997)? (2) How to best aggregate a collective estimate of such samples of estimates? To that end, we will first look how classical statistical decision theory approaches the first question. Then, we use the two-dimensional definition of accuracy of measurement methods to conceptualize crowd wisdom. We will look at the wisdom-of-crowd indicator and the fraction of outperformed estimates as two different measures to quantify the degree of crowd wisdom. The latter can be used to conceptualize and quantify an optimal crowd size. The second part of the paper is about the choice of the aggregation function as the median or the mean. We will point out that Galton’s call for the median as the “democratic” function should be refined for situations when estimators tend to be bipolarized and, in particular, when the decision space has natural upper and lower bounds as, for example, probabilities or percentages have. The paper ends with a discussion of the practical application of the results.

All crowds are wise in statistical decisions theory

In terms of statistical decision theory (Berger, 1989), we aggregate the *collective decision* $\text{colD}(x)$ from a *sample* of estimates $x_1, \dots, x_n \in \mathbb{R}$ by an aggregation function which is a statistical estimator of the *true value* θ . The classical aggregation function is the arithmetic mean $\text{colD}(x) = \bar{x}$ which coincides with the idea of estimating the expected value $E(X)$ of the underlying random variable X . All these definitions and notations and the following ones are summarized in Table 1.

The idea of statistical decision theory is to find the optimal decision under uncertainty. A crucial ingredient is the quantification of the gain or loss of the decision-maker. For estimating a continuous value the cost function is a function of the difference of the decision $\text{colD}(x)$ and the true value $\text{colD}(x) - \theta$. The absolute value and the square function are the most used candidates (see, e.g., Laan et al., 2017; Becker, Porter, & Centola, 2019; Jayles et al., 2017). The absolute value because it is the most natural candidate and the square function because it has nice theoretical properties.²

With the square cost function and the arithmetic mean as the aggregation function, the phenomenon of crowd wisdom can be conceptualized by the relation of the *collective error* $\text{colErr} = (\bar{x} - \theta)^2$, the *mean squared error* $\text{MSE}(x, \theta) = \frac{1}{n} \sum_{i=1}^n (x_i - \theta)^2$, and the *variance* $\text{Var}(x) = \frac{1}{n} \sum_{i=1}^n (\bar{x} - x_i)^2$. The collective error is also called statistical bias³ or population bias (Vul & Pashler, 2008). In terms of measurement theory, collective error is caused by systematic error while variance is caused by random error.

The mathematical relation of the mean squared error, the variance and the statistical bias is the *bias-variance decomposition of squared error* which states that the mean of squared errors is equal to the squared bias plus the variance of estimates, as used in statistics (O’Sullivan, 1986), machine learning (Geman, Bienenstock, & Doursat, 1992), or ensemble learning (Brown, Wyatt, Harris, & Yao, 2005). More suggestive of the wisdom of

² However, cost functions may also relate the error of the decision to its practical consequences. Larger absolute errors maybe proportionally more or less costly, overestimating maybe more costly than underestimating or the other way round, or thresholds may play a role, e.g., when the decision must lie in a certain range. For example, the experimental setting of Lorenz, Rauhut, and Kittel (2015) uses a payoff function which has a sharp peak at the correct answer and flattens to zero in discrete steps with increasing error. This payoff function can be interpreted as an inverse of a cost function where cost flattens to a constant for arbitrarily large error.

³ Note, that the term “bias” has a purely statistical meaning here and should not be confused with systemic or cognitive biases which relate to institutional or mental processes which might cause statistical bias.

crowds, Page (2007) calls it the *diversity prediction theorem*

$$\text{colErr}(x, \theta) = \text{MSE}(x, \theta) - \text{Var}(x)$$

which states that the collective error is the average individual error (MSE) minus diversity (variance). This implies that higher diversity decreases the collective error, or “Diversity is good for collective intelligence!” This message is a bit suggestive as it blurs over the fact that increases in variance (diversity) usually also imply an increase of the average individual error. For example, let us consider a certain sample of estimates with a certain collective error and a certain mean of squared error. Increasing diversity in this sample could be realized by adding a random number from a standard normal distribution. Obviously, the collective error remains unchanged (neglecting random fluctuation), but diversity (variance) of the sample increases (by the variance sum law). The increase of diversity thus did not decrease the collective error, because at the same time as diversity increases, also the mean squared error increases. Less suggestive, the message to extract from the diversity prediction theorem is: In case of large individual errors, diversity is needed for a small collective error.

Analogously, Davis-Stober et al. (2014) approach crowd wisdom by comparing the average of the estimates of several estimators with the choosing strategy (terminology of Laan et al., 2017) which is to select one estimator at random and use its estimate as the collective decision. Using the square cost function, the expected cost using the choosing strategy coincides with the mean squared error $\text{MSE}(x)$ and consequently crowds are almost always wiser than their individuals.⁴ Laan et al. (2017) point out that the average of estimates of a crowd are generally better than the choosing strategy whenever the cost function is convex (which the square function is). This result is a mathematical conclusion

⁴ Davis-Stober et al. (2014) also consider weighted arithmetic means and weighted selection in the choosing strategy, and different correlations of the estimates of different individuals with the true value where the true value (called the criterion) is also considered a random variable. These aspects are not considered here.

of Jensen's inequality and holds for any sample of random numbers. Therefore Laan et al. (2017) critically ask: *"Should we say that the collection of random numbers possesses collective intelligence?"*

Table 1: Notations and definitions.

High trueness, low precision: The crowd as a measurement device

In the following, we want to answer the question “*When is a crowd wise?*” (Davis-Stober et al., 2014) not with “always” as classical statistical decision theory suggests. After Galton (1907c) found the wisdom-of-crowd phenomenon in the data he collected in Plymouth, he concluded: “*This result is, I think, more creditable to the trustworthiness of a democratic judgment than might have been expected.*” When we want to distinguish crowd wisdom from individual wisdom, we need to quantify the aspect how unexpected the accuracy of the aggregated estimate is compared to individual estimates.

To that end, we distinguish two extreme cases to assess when we consider crowds as not wise. (1) When the true value is not even sandwiched between the minimal and maximal estimate, then there is no wisdom in the crowd but a systematic error in the sample. (2) When all individuals estimate correctly, then the crowd is not wiser than each individual. Thus, there is no crowd wisdom which goes beyond individual wisdom.

When we see a crowd of estimators as a measurement device we can quantify crowd-wisdom by using the two-dimensional concept of accuracy for measurement devices. The International Organization for Standardization (ISO) defines in ISO 5725⁵ *accuracy* of a measurement method by trueness and precision: “*Trueness* is the closeness of agreement between the arithmetic mean of a large number of measurements and the true or accepted reference value. *Precision* instead is the closeness of agreement between different measurements.” Trueness resembles an inverse of the collective error and precision an inverse of the variance.

According to the two extreme cases and ISO’s terminology, the wisdom-of-crowd phenomenon appears when estimates from a crowd show high trueness (low collective

⁵ ISO 5725 (six parts) “Accuracy (trueness and precision) of measurement methods and results”

error), but low precision (high variance). Lower crowd wisdom appears obviously with low trueness and low precision, but also with high trueness and high precision. The latter situation is captured by case (2). The lowest amount of crowd wisdom appears with low trueness but high precision as captured in case (1). While individual wisdom increases with trueness (reducing bias) and precision (reducing uncertainty), crowd wisdom is maximal with perfect trueness and low precision (high diversity). Looking on individuals, the worst situation appears with low trueness and low precision, leading to bias and uncertainty. From a crowd wisdom perspective, low trueness and high precision is the worst situation, because it creates a false consensus effect for an external observer, or a *collective tunnel vision* (Rauhut, Lorenz, Schweitzer, & Helbing, 2011).

What could be a one-dimensional measure of crowd wisdom in the sense of the two extreme cases when crowds are not wise? We could compute how many standard deviations⁶ the truth deviates from the collective decision (also called standard score, $\frac{|\bar{x}-\theta|}{SD(x)}$). This would allow to compare crowds of samples operating on different magnitudes, e.g. the crowds estimating the weights of mice and elephants.

However, many real world settings of collective estimation produce samples for which the arithmetic mean is not the best aggregation function and the variance not the most appropriate measure of dispersion. For samples with substantial skewness the mean would shift towards the tail and deviates from where most estimates lie. Further on, the standard deviation would not compare under- and overestimation well. Samples with fat tails where very few very extreme outliers exist would be assessed very different from samples without such outliers. The two samples in Figure 1 are good examples. Galton's sample is close to normal while the Viertelfest sample is very skew with the mean performing much worse than the median.

⁶ The standard deviation (SD) is the square root of the variance.

In the following, we look at the *wisdom-of-crowd indicator* and the *fraction of outperformed estimates*. Two measures acknowledging this non-trivial conception of crowd wisdom. Both measures are theoretically bounded by zero and one and allow the comparison of crowd wisdom in very different samples.

Wisdom-of-crowd indicator

The *wisdom-of-crowd indicator* (Lorenz, Rauhut, Schweitzer, & Helbing, 2011) $\text{WoC}(x, \theta)$ for a sample of estimates x and the true value θ is defined in Table 1. The measure relies on the median as aggregation function. The idea is to bracket the truth in the ordered sample with an interval centered around the median estimate. The indicator equals one, when this interval is minimally small. The indicator equals zero, if the truth is not even bracketed by the minimal and the maximal estimates, $\theta < \hat{x}_1$ or $\theta > \hat{x}_n$.

The wisdom-of-crowd indicator can also be described as the fraction of ordered estimates outside of the smallest set of estimates which is centered around the median and truth-bracketing.

The wisdom-of-crowd indicators of the samples in Figure 1 are $\text{WoC}(x^G, 1198) = 0.896$ for Galton’s sample and $\text{WoC}(x^V, 10788) = 0.906$ for the ViertelFest sample. Hence, the ViertelFest sample shows a slightly higher degree of crowd wisdom: 90.6% (instead 89.6%) of all estimates lie outside of the smallest centered interval bracketing the truth.

Fraction of outperformed estimates

Based on the results of statistical decision theory Davis-Stober et al. (2014) concluded: “*Given our results, we conclude that, in general, extraordinary evidence is needed to justify choosing an expert’s judgment over the aggregate of a crowd.*” However, this result is based on the expected error of a randomly selected person as the expert. So, they also acknowledged “*In addition, one could consider alternative generalizations, such as comparing the crowd performance with the best performing individual.*” The fraction of

outperformed estimates $\text{FoE}(x, \theta, \text{colD})$ takes a similar approach.

The idea is to count all estimates which are further away from the truth than the collective decision $\text{colD}(x)$ and divide by the total number of estimates. This reflects the probability that the collective decision is better than a decision made by choosing a random estimate. In the following, we will call those individuals whose estimates are closer to the truth than a collective estimate *experts*. This is a post hoc definition of an expert and does not rely on an assessment of deep competence in judgment but just declares good de facto estimators as de facto experts.⁷ Crowd wisdom in a sample is thus maximal when there are no experts. The formal definition of the fraction of outperformed estimates is shown in Table 1. Notice that the formal definition compares distances to the truth of individual and collective estimates, but the result would be identical for squared distances as well as all other monotone functions of distance.

The fraction of outperformed agents modifies the cost functions based on the distance through a competitive perspective: It becomes important to be closer to the truth than others and cost is not directly related to distance to the truth anymore.

For the median as aggregation function, the fraction of outperformed estimates of the samples in Figure 1 are $\text{FoE}(x^G, 1198, \text{median}) = 0.882$ for Galton’s sample and $\text{FoE}(x^V, 10788, \text{median}) = 0.92$ for the Viertelfest sample. Thus, this measures confirms the assessment of the wisdom-of-crowd indicator that the ViertelFest sample shows a slightly higher degree of crowd wisdom when the median is used. The median outperforms 92% (instead of 88.2%) of all estimates.

The fraction of outperformed estimates also allows to assess the degree of crowd wisdom according to other measures than the median. The arithmetic mean is a better

⁷ A de facto expert according to the definition is not a “superforecaster” in the sense of Tetlock and Gardner (2015).

aggregation function in Galton’s sample with $\text{FoE}(x^G, 1198, \text{mean}) = 0.996$ and the geometric mean in the Viertelfest sample with $\text{FoE}(x^V, 10788, \text{geomean}) = 0.986$. The median seems to be the most appropriate common aggregation function for both samples. This fits to the theoretical fact that the median coincides with the geometric mean of a lognormal distribution, while it coincides with the arithmetic mean for a normal distribution. Appendix A provides some more details about the two example samples.

Optimal crowd sizes which maximize the fraction of outperformed estimates

The competitive conception of crowd wisdom which underlies the fraction of outperformed estimates can be used approach another central question of crowd wisdom, the optimal crowd size (see Karotkin & Paroush, 2003). From the mechanism designer’s perspective, one aspect of this question appears obviously when estimates are costly, e.g., because experts must be paid. Another aspect, which we focus here, is that a large crowd would for sure realize the collective error in the collective decision of the population, while a smaller crowd would leave some chances that the collective decision is closer to the truth by chance. The concept of the fraction of outperformed estimates can help to quantify this idea.

Besides computing the probability that a randomly selected estimate from a sample is better than the aggregated estimate from the crowd, we may also want to know the probability that the aggregated estimate of a crowd of three, five, or more randomly selected estimates is better than the median of the full sample. Further on, we may ask if we can expect that experts (in the sense of post hoc assessment as discussed before) exist for every estimation task?

To that end, we consider individuals with their particular estimates to be random draws from a population. A collective estimate of a crowd of k individuals is the aggregation of k independent draws from the same population. An estimate is thus a

random variable X with a certain probability density function f_X . The collective decision of k estimates from X is called $\text{colD}([X]_k)$. It is another random variable aggregated from k independent and identically distributed random variables $X_1, \dots, X_k$. Consequently, the probability density function $f_{\text{colD}([X]_k)}$ can thus be derived from f_X and the aggregation function for the collective estimate. The probability that a single estimate is closer to the truth than the mean of k estimates is thus $\Pr(|X - \theta| < |[\bar{X}]_k - \theta|)$. This probability also gives us a measure for the expected number of experts in a sample of k individuals by multiplying the probability by k . Further on, we can use this a framework to analyze the impact of bias and crowd size on the probability to select an expert by chance.

The following analysis we assume that the correct value is zeron and estimates come from a normal distribution with standard deviation one and a mean of b which stands for “bias”. Appendix B shows how the probabilities can be computed. Lorenz (2021) documents the code. Figure 2 shows the impact of an increasing group size in this situation. Figure 2A show that the expected number of experts increases without saturation in the case of no bias ($b = 0$). However, the probability to draw an expert converges to zero with $k \rightarrow \infty$ as 2B shows. For a group of 1,000 estimators the probability to find an expert is 0.0202, and consequently the expected number of experts among these is 20.

This picture changes with bias as shown in 2C for a bias of $b = 0.5$. In the limit of $k \rightarrow \infty$ the probability to draw an expert now saturates at around 0.34. (Consequently, the number of experts would increase linearly with k .) Interestingly, the probability to draw an expert does not decline monotonically. There is an interior minimum at $k = 19$ for which the probability is lowest (the value is 0.337). Thus, for a bias of 0.5 it is in the competitive sense of the fraction of outperformed estimates optimal to ask 19 individuals and take their average estimate. The explanation for this intermediate optimum under bias is the following. For a lower crowd size the group is not large enough to optimally outperform a randomly drawn individual as in the unbiased situation. For large groups the dispersion of

the collective decision vanishes, such that we end up with a collective estimate very close to the bias with high certainty. The optimal crowd size is achieved when the variance decreasing effect is balanced optimally against the effect that larger groups will deliver a wrong answer for sure.

In a next step, we computed this optimal crowd size for different biases. Figure 3 shows the impact of the bias on the optimal crowd size and the probability to draw an expert. This theoretical analysis delivers the following insights on the optimal crowd size. In a large crowd with a large bias, half of the people are post hoc experts because they outperform the mean estimate⁸. Precisely, with increasing bias the probability to draw an expert from a large crowd increases and saturates at 0.5. For a bias of 1.5 standard deviations it is already 0.50.

In a large crowd with small bias, experts are rare, and the optimal size of a group is very large. For $b < 0.26$ the optimal crowd size is larger than 100, for $b < 0.14$ it is larger than 500 and for $b < 0.10$ larger than 1,000. In a crowd with no bias, the share of experts goes to zero, although their number increases slowly without saturation. When the bias is intermediate between 0.5 and 1.5 standard deviations, the optimal crowd size is between 19 and 2. In this range of bias there is a notable difference of the probability to draw an expert which is better than a large crowd ($k \rightarrow \infty$) and a crowd of optimal size. Thus, in this range of bias there is a risk of drawing too many estimates and receive worse results with higher likelihood.

We can also compute optimal crowd sizes for the two empirical samples when we assume that the crowd of size k is sampled from the empirical sample. To that end, we must compute the distribution of the median estimate of such crowds of size k and compute the probability that the median is better than one random estimate.⁹ For

⁸ Note, that this result generalizes only for symmetric distributions.

⁹ The computation of the distribution of the median is documented here

Galton's sample the optimal crowd size is 579 realizing the lowest possible probability to draw an expert of 0.097. For the Viertelfest sample it is 2,869 with a minimal probability to draw an expert of 0.075.

Mean vs. median under natural bounds and polarization

The former analyses all took the aggregation function (the mean, median, or geometric mean) as given. However, the choice is ours and it can have drastic consequences on the wisdom of crowd. Now, we treat this second quantification question: How to best aggregate a collective estimate of such samples of estimates? We focus on mean versus median, the most common measures (Hora, Fransen, Hawkins, & Susel, 2013; Stone, 1961).

The empirical samples in Figure 1 seem to underpin Galton (1907a), who argued that the median shall always be used because of its “democratic” aggregation. The mean shall not be used, because it would give *“voting power to ‘cranks’ in proportion to their crankiness.”* This argument has two aspects. In the terminology of modern statistics, it recalls that the median is a more robust measure of central tendency which is not effected by the extremeness of outliers. From the perspective of epistemic democracy, it states that the median cannot be manipulated by few malicious estimators who strategically over- or understate estimates to steer the mean to a certain direction. Arguably, the democracy aspect does not play a crucial role in estimation games, because there are no incentives for manipulation. Nevertheless, the aspect is important because in many practical estimation problems actors may have mixed interests not solely focused on approximating the truth as close as possible. For example, politicians with a short-term interest to push public investments may have an interest that the tax estimation is better over- than underestimated; economic advisers believing in self-fulfilling prophecies may prefer overestimation of economic growth; the oil industry would prefer lower than correct

<https://math.stackexchange.com/questions/3212165/sample-k-of-n-numbers-with-replacement-what-is-the-probability-for-a-cert> . Lorenz (2021) provides the code.

assessments of the impact of climate gas emissions on the increase of world temperature, while climate activists would prefer the opposite.¹⁰ Estimation problems with mixed preferences are probably the more common and more relevant ones. Mixed preferences likely do both, they trigger cognitive biases and create incentives to report values maliciously wrong.

In the two empirical examples, the median is the preferable aggregation function against malicious estimates, but is this always the case? In the following, we focus on the interplay of dispersion and the existence of “natural bounds” of the decision space.¹¹ Galton’s sample and the Viertelfest sample both have a natural bound of zero, but arguably, Galton’s sample has no effective bounds because no estimates are close to the lower bound.

Besides those problems with one or no effective natural bounds, there is a third class of continuous problems: Those with two natural bounds, a lower and an upper one. Examples can be found in the studies of Granovskiy, Gold, Sumpter, and Goldstone (2015) and Lorenz et al. (2015) where percentages should be estimated. Also other natural bounds limiting estimates may exist, e.g. the seats in a plane or the height of a room. An example for a decision in the realm of epistemic democracy is the decision on the magnitude of income or wealth tax that maximizes productivity and/or state revenue. Here we can expect that estimators have mixed epistemic and egoistic preferences.

By taking into account that practical problems may be transformed through scaling (multiplying by a constant) and shifting (adding a constants), we distinguish three

¹⁰ An example of wisdom-of-crowd methods in climate change is the pooling of expert views on sea level rise provided by Bamber and Aspinall (2013).

¹¹ A “natural bound” is a number which is externally determined, e.g. a minimum of zero for counts or weights as in the two examples. It is not just the minimal number in a sample or some value which nobody regards as reasonable because this is cannot be quantified precisely.

fundamental classes of continuous problems. Those with two, one, or no natural bounds with the canonical decision spaces $[0, 1]$, $[0, +\infty[$, and $[-\infty, +\infty[$. For each of these decision spaces we can specify a typical distribution: The Beta distribution, the lognormal distribution and the normal distribution. We take these three distributions for a mix of pragmatic reasons: They are common in the literature and have relatively tractable analytical forms.¹² When bounds are not effective because few values lie close to them, all the three distributions can look indistinguishable in samples of small or medium size. Further on, the lognormal and the Beta distribution may be indistinguishable when most values lie close to zero.

Figure 4 shows for each of the three distributions four examples. All examples are parameterized to have an arithmetic mean of 0.6, but different standard deviation from 0.15 to 0.45.¹³ Each of the twelve subplots shows a sample of 10,000 random draws and the sample's mean and median. In these examples we do not assume a specific true value θ but discuss how mean and median deviate and what true values could be more likely.

The Figure shows that small standard deviations indeed render the three distributions quite similar. Also, the median lies close to the mean of 0.6. This is similar to the empirical example of Galton (1907c). For larger standard deviations, the lognormal distribution appears more and more right-skew and the median departs from the mean by moving closer to zero, as in the ViertelFest dataset. The mean stays constant by design.

¹² More consistent would be the logit-normal distribution instead of the Beta distribution, but this has no analytical forms for its moments. The Gamma distribution instead of the lognormal distribution would be more related to the Beta distribution, but not to the normal distribution.

¹³ For the normal distribution the mean μ and the standard deviation σ are the standard parameters. The parameters of the lognormal distribution, $\mu_{\log} = \log(\mu^2/\sqrt{\sigma^2 + \mu^2})$ and $\sigma_{\log} = \sqrt{\log(\sigma^2/\mu^2 + 1)}$ represent mean and standard deviation of the logarithm of the underlying random variable. The Beta distribution has the probability density function $x^{\alpha-1}(1-x)^{\beta-1}/B(\alpha, \beta)$ where the denominator is the Beta function. The parameters are $\alpha = \mu(\mu(1-\mu)/\sigma^2 - 1)$ and $\beta = (1-\mu)(\mu(1-\mu)/\sigma^2 - 1)$.

For the Beta distribution with an intermediate standard deviations of 0.25 the shape remains single-peaked but becomes more left-skew. Left-skewness appears in our example because the mean is larger than 0.5. For a mean below 0.5 the distribution would become right-skew. A value above 0.5 is chosen here to make the difference to the lognormal distribution more clear. For even higher standard deviations, the distribution becomes U-shaped and bimodal with modes at zero and one, where the mode at one is always higher than the mode at zero. For means less than 0.5 this would reverse. A U-shaped distribution describes a bipolarized society. In our example, bipolarization triggers a shift of the median towards one. In the extreme case of full bipolarization 60% of the population estimates one and 40% estimates zero. The median aggregates the majority decision to the value one, while the mean aggregates to 0.6. Arguably, the mean’s “compromise” may be closer to correct in many real-world cases.

The mean-median dilemma of epistemic democracy

Galton (1907a) argued, that the median is most robust against malicious “cranks”. When arbitrarily low and arbitrarily high estimates are permitted than a single malicious estimator who has a good sense of the mean of all other estimates can steer the final mean to any value.¹⁴ A dramatic example of malicious manipulation in the real world is the Libor scandal (Wikipedia contributors, 2021). The manipulation strategy does not work for a single estimator with the median. A look on the three problem classes from Figure 4 shows that this holds mainly for problems with no effective bounds for the estimates. When there is a natural bound which only permits positive estimates than manipulation of the mean by one malicious estimate would only work for larger target outcomes while manipulation

¹⁴ The formula for a malicious estimator who wants the collective mean to be y is to set the malicious estimate to $ny - (n - 1)x$ where n is the number of estimators and x is the mean of the estimates of all $n - 1$ other estimates.

towards zero is limited.¹⁵ Further on, the manipulation strategy works only with extremely wrong estimates which should appear as obviously not based on truth-seeking preferences.

The picture of the robustness of the median somehow reverses in bipolarized situations when estimates have to be between two natural bounds. In such situations, the median can be more prone to manipulation. Let us assume a probability shall be estimated which is 0.5, but the population is divided into fifty estimators who think it is zero, fifty estimators who think it is one and only one estimator who correctly thinks it is 0.5. In this situation, the mean as well as the median correctly aggregates the probability to the true value 0.5, but the central estimator is able to manipulate the median to any value including both extremes. Further on, only one of the hundred extreme estimators needs to switch sides to produce an extreme and most wrong collective decision. Instead, the mean is quite robust and can only be manipulated to values between 0.49 and 0.51 by these manipulations.

Figure 5 demonstrates that the robustness of the mean over the median also appears for less than fully polarized distributions beyond a certain critical polarization threshold. The figure picks up the case of the Beta distribution and assume that most estimates come from this distribution while a small fraction of estimates is zero to manipulate the collective decision to the closest possible value. In this example we further assume that the true value is 0.5 and coincides with the mean of the Beta distribution where the majority of estimates come from. On the horizontal axis we increase the polarization of the distribution by increasing the standard deviation (SD) of the Beta distribution. For reference: With $SD = 0$ the distribution is fully condensed on the true value while $SD(x) = 0.5$ is the most polarized situation with half of it condensed at each extreme side. $SD(x) = 0.1$ implies a distribution $B(12, 12)$, $SD(x) = 1/\sqrt{12} \approx 0.289$ implies the uniform distribution $B(1, 1)$, and $SD(x) = 0.4$ implies approximately $B(0.28, 0.28)$. The Figure

¹⁵ When the geometric mean is used instead, manipulation arbitrarily close to zero is possible.

shows that the median is more robust against the manipulators for low standard deviations while the mean becomes dramatically better under polarization. The critical polarization appears for the uniform distribution in this particular example. In the case of a polarized crowd the mean of estimates might not represent any single estimate in the crowd, which can be seen as a disadvantage of aggregation by the mean (O’Hagan et al., 2006). However, in a political context it could appear as a reasonable compromise.

It should be noted, that the phenomenon also appears when the truth and mean of the Beta distribution are not in the center of the decision space. It also appears when the truth deviates slightly from the Beta distribution’s mean. In this case, a tiny group of manipulators may shift the collective decision accidentally closer to the truth, but still a slightly larger groups can manipulate the median as dramatically as without systematic error. We skip a more detailed analysis here.

The results above create a *mean-median dilemma* of epistemic democracy. Let us assume a tax rate should be set by decision makers to maximize societal benefits by raising the maximal revenue (e.g., following theories of optimal taxation Mirrlees, 1971). When estimators are themselves affected by the tax in different magnitudes, this creates a typical situation of mixed truth-seeking and egoistic preferences. A mechanism designer’s problem is to choose an appropriate aggregation rule.

Let us assume the aggregation rule is the mean and estimators estimate initially non-extreme values centered around the true value of, e.g., 30%. Estimators personally affected by the tax might start to submit lower and lower estimates to drag the tax rate down. As a reaction estimators not personally affected by the tax rate may submit higher values to ensure the tax revenue and avoid that other taxes have to be raised. Thus, the mean erodes truth-seeking preferences and creates more extreme estimates through this vicious cycle. The reason is as Galton (1907a) said: Cranks are given voting power in proportion to their crankiness. This calls for an implementation of the median instead.

However, this may create another problem. If the group of estimators polarizes because of other processes, e.g. social influence (see Flache et al., 2017; Lorenz et al., 2021), then the decisions aggregated by the median will be much more extreme and less likely to lie close to the optimal intermediate value. In this case, a reversion to the mean as decision rule would come closer to optimal.

A solution for the mean-median dilemma could lie in using the median in non-polarized situations to avoid incentives for polarization but to switch to using the mean when polarization appears. The operationalization and a game-theoretic analysis is beyond the scope of this study.

Conclusion and Discussion of Practical Application

Many studies of crowd wisdom focus on binary choice where the concept of accuracy is as simple as the probability for the correct decision. For continuous estimation problems, the quantification of crowd wisdom has intertwined aspects in the choice of the aggregation function and measurement of the degree crowd wisdom which distinguishes it from the average individual wisdom. The standard tools of statistical decision theory render every crowd wiser than its individuals and does not provide a gradual quantification of the phenomenon which enables comparison of different types of samples. The ISO standard two-dimensional definition of accuracy provides a useful concept in which crowd wisdom can be characterized with high trueness and low precision in a sample of estimates. The wisdom-of-crowd indicator and the fraction of outperformed estimates are two measures which capture the phenomenon of crowd wisdom in a gradual way. The later can be used to quantify optimal crowd sizes based on the expectation of systematic biases in the estimation tasks. When intermediate systemic biases are frequent, larger crowds can make it more likely that these crowds are outperformed by individual estimators.

How can the results about optimal crowd size be applied by a mechanism designer

setting up crowds of independent estimators? First of all, the real-world cost function should be of a competitive nature as the fraction of outperformed estimates suggests. That means, benefits are drawn from having final estimates closer to the truth than most others, while the cost of failing in that are not increasing with distance. Second, a crucial parameter for the computation of the optimal crowd size is the bias, a parameter which is only known when the true value is known and when it is known there is no need for a wise crowd anymore. However, in the case of repeated estimation tasks of similar type a post hoc analysis could show how large biases typically are. When biases lie frequently in the range of 0.5 to 1.5 standard deviations, a limited crowd size between 17 and 2 is advisable. The exact numbers apply for normally distributed estimates. Other distributions would require a new computation. When biases are usually smaller it is reasonable to gather as many estimates as possible. When biases are often larger, then it is not advisable to rely on crowds of independent estimators.

The idea of the wisdom of crowds is increasingly used for the epistemic interpretation that democratic decision making can “track the truth” (Goodin & Spiekermann, 2018; Landemore & Elster, 2012; List & Goodin, 2001). A practical procedure in an epistemic democracy could be the direct decision on continuous quantities like tax rates, allowances, minimal wages, or basic incomes through an aggregation of estimates of the electorate. This would be the application of wisdom of crowd aggregation as analyzed here for policy making. In such political situations, preferences of voters cannot be assumed to be purely truth-seeking. Aspects of robustness against manipulation become relevant for the selection of the appropriate aggregation function, and a mechanism designer may wonder which aggregation function would be best, the mean or the median.

The results presented here, suggest that for non-polarized issues with no effective bounds, the median is better because it shows the highest barriers for manipulation. However, this changes for issues where the decision space has two natural bounds and the

electorate is polarized. In these situations, the median tends to produce extreme decisions by tight majorities. A societal compromise is better aggregated with the arithmetic mean than with the median in this case. This would then arguably be a better decision from an epistemic perspective. To function well, such democratic procedures might need options to switch from median aggregation to aggregation by the mean when polarization appears. One way would be to implement a polarization threshold where the aggregation automatically switch from the median to the mean. The example from Figure 5 would suggest for collective decision bounded by zero and one, that such a threshold could be when the standard deviation is $\frac{1}{\sqrt{12}} \approx 0.289$ which is the standard deviation of a uniform distribution. The aggregation rule would thus first compute the standard deviation and in a second step the collective decision by the median (when the standard deviation of individual decisions is lower) or the mean (when it is above the threshold). However, this threshold comes from an example where the correct value is 0.5. Further on, such discontinuous breaks in a decision rule should be analyzed game-theoretically in more detail before implementation.

An important additional insight is that there is a mean-median dilemma: Aggregation by the mean creates incentives for strategic radicalization of estimators in a nonpolarized electorate. So, the median is preferable in the first place but when polarization appears and the aggregation switches to aggregation by the mean, polarization may easily become persistent. The mean would then produce compromise but does not depolarize the electorate.

Appendix A: Details about the two estimation games

This appendix provides a bit more context and analysis of the particular estimation games shown in Figure 1.

The data of the seminal estimation game that Galton (1907c) described was recently

dug out. Wallis (2014) extracted the full sample of estimates from Galton’s notebook out of the archive and made it publicly available.¹⁶ We call Galton’s sample x^G . The median estimate is $\tilde{x}^G = 1,208$. Figure 1A shows the histogram of estimates together with the truth, the median and the mean of estimates. Interestingly, the mean of estimates $\bar{x}^G = 1,196.7$ is even closer to the truth than the median.¹⁷

The other estimation game used here happened at the lottery “Haste mal ’nen Euro? – Tombola”¹⁸ at the festival “ViertelFest” August 22-24, 2008 in Bremen. This was an ordinary tombola, where people could buy lots which could bring them prizes. As the topic of the festival was “Swarm Intelligence”, I asked the organizers to include an estimation game similar to the one of the West of England Fat Stock and Poultry Exhibition. Each buyer of a lot got another chance to win. They could estimate the number of lots which would have been sold in total after the three days of the festival. The estimates closest to the true value got tickets for the circus. Further on, before the lottery the organizers asked an expert, the organizer of Bremen’s biggest lottery (Bürgerpark-Tombola) about his estimate for reference. The data is accessible online.¹⁹

We call the ViertelFest sample x^V . In total, 1,226 estimates of the number of lots were handed in and 10,788 lots had been sold²⁰. The expert estimated 19,100 lots. The 5- and

¹⁶ He provided data at

https://www2.warwick.ac.uk/fac/soc/economics/staff/academic/wallis/publications/galton_data.xlsx , downloaded Mar 16, 2020. Wallis also found some small errors in Galton’s computation.

¹⁷ Hooker (1907) already made this point. He estimated the mean as 1,196 based on the percentiles reported by Galton (1907c). In a reply, Galton (1907b) reported it to be 1,197.

¹⁸ Colloquial for “Would you spend one Euro? – Tombola”.

¹⁹ Blog post in German: <http://janlo.de/wp/2010/06/22/die-weisheit-der-bremer/>. Direct link to the data: https://docs.google.com/spreadsheets/d/1HiYhUrYrsbeybJ10mwsae_hQCawZlUQFOOZzcugXzgA/edit#gid=0

²⁰ Thus, 11.4% of the lots were used to participate in the estimation game.

95-percentiles spans values from 1,213 to 99,774, and thus almost three orders of magnitude. The median of estimates is $\tilde{x}^V = 9,843$ and comes much closer to the true value than the expert. The arithmetic mean $\bar{x}^V = 53,164$ is even worse.

Figure 1B shows the truth, the expert’s estimate, the median, the arithmetic, and the geometric mean and the histogram of estimates from the ViertelFest with the x-axis clipped such that all estimates larger than 150,000 are not visible. There are 29 estimates larger than 150,000 (2.4%) ranging from 157,853 to 29,530,000. These large estimates (especially the largest one) make the arithmetic mean such a bad aggregation function in this case. Nevertheless, even when these 29 values were removed the mean would be 18,756 and still way too high.

The reason, why the median’s performance is so different from the mean’s is that the distribution is heavily right-skew with a fat right tail. This is typical for distributions when numbers are bounded by zero but span several orders of magnitude. In this situation, humans face a problem of “logarithmic” nature. That means, they cope first with finding the right magnitude.²¹ In such situations, the distribution is better assumed to be lognormally instead of normally distributed. Consequently, the geometric mean²² is often a better measure of central tendency than the arithmetic mean. The geometric mean of the ViertelFest estimates is $\text{geomean}(x^V) = 10,510$ and thus only 278 less than the correct answer.²³ Also Galton’s sample is naturally bounded by zero as the data from the ViertelFest, but here answers do not span several orders of magnitudes and no answers are close to this lower bound. This makes the data close to normally distributed.²⁴

²¹ Logarithmic thinking also seems to be the natural human intuition for mapping numbers to a line, while the linear mapping is a cultural invention needing formal education (Dehaene, Izard, Spelke, & Pica, 2008).

²² The geometric mean is the exponential of the arithmetic mean of the logarithmized sample data.

²³ The six questions in Lorenz et al. (2011) confirm the superiority of the geometric over the arithmetic mean.

²⁴ A closer look shows that the distribution is slightly left-skew as already Galton noticed. Interestingly,

These examples underpin that an ideal procedure of knowledge aggregation should matches the nature of the knowledge distribution (Laan et al., 2017). Indeed, for Galton’s sample the normal distribution is a better fit than the log-normal distribution according to all goodness-of-fit measures.²⁵ For the ViertelFest data instead the log-normal fit is better than the normal fit by all goodness-of-fit measures.²⁶

Appendix B: Computing the optimal crowd size when estimates are normally distributed

Here we derive how to compute $\Pr(|X - \theta| < |[\bar{X}]_k - \theta|)$ for a random variable X with probability density function f_X , where $[\bar{X}]_k$ is the collective decision of k estimates independently sampled from X . Consequently, the its probability density function $f_{[\bar{X}]_k}$ can be derived from f_X .

Here, we consider X to be normally distributed with mean μ and standard deviation σ . For the sake of statistical simplicity we also assume that we draw the estimators with replacement. That means that there is also no dependence between the individual and the aggregated estimates.

Let us further assume without loss of generality that $\sigma = 1$ and $\theta = 0$. Then, $b = \mu$ is

this is the opposite of what one would expect from a problem with just one natural bound of zero. Nash (2014) explains this using an augmented quincunx model of sequential and probabilistic cue categorization arguing that left-skewness appears because most people knew what an average weight would be and the fact that the ox was comparably heavy.

²⁵ The function `gofstat` in the R-package `fitdistr` delivers the measures: Kolmogorov-Smirnov statistic (KS), the Cramer-von Mises statistic (CvM), the Anderson-Darling statistic (AD), Akaike’s Information Criterion (AIC), and the Bayesian Information Criterion (BIC). The values for the normal fit (lognormal fit in parentheses) are: KS 0.0814 (0.0947), CvM 1.78 (2.34), AD 10.5 (13.6), AIC 9,002 (9,035), BIC 9,012 (9,044).

²⁶ Normal fit (lognormal fit in parantheses): KS 0.475 (0.0725), CvM 92.6 (1.39), AD Inf (6.92), AIC 36,955 (26,849), BIC 36,965 (26,859).

the (unsquared) bias in the population and $f_X(x) = \phi(x - b)$, with $\phi(x) = \frac{e^{-x^2/2}}{\sqrt{2\pi}}$ being the probability density function of the standard normal distribution. For means of independent normally distributed random variables, it holds that $[\bar{X}]_k$ is also normally distributed with standard deviation $\frac{1}{\sqrt{k}}$. Hence, it holds $f_{[\bar{X}]_k}(x) = \frac{1}{\sqrt{k}}\phi(\frac{x-b}{\sqrt{k}})$. Now, the probability of receiving a better estimate from one randomly selected estimator than the arithmetic mean of k estimators is

$$\Pr(|X| < |[\bar{X}]_k|) = \Pr(|X| - |[\bar{X}]_k| < 0) = \int_{-\infty}^0 f_{|X| - |[\bar{X}]_k|}(x) dx = \int_{-\infty}^0 f_{|X|} * f_{-|[\bar{X}]_k|}(x) dx \quad (1)$$

where ‘ $*$ ’ is the convolution of the two functions. The two functions in the convolution are

$$f_{|X|} = \begin{cases} \phi(x - b) + \phi(-x + b) & \text{if } x - b \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$f_{-|[\bar{X}]_k|} = \begin{cases} 0 & \text{if } x - b \geq 0, \\ \frac{1}{\sqrt{k}} \left(\phi(-\frac{x-b}{\sqrt{k}}) + \phi(\frac{x-b}{\sqrt{k}}) \right) & \text{otherwise.} \end{cases}$$

Figure 6 shows an example of these function for a bias $b = 0.50$. The gray area represents the probability that $|X| < |[\bar{X}]_k|$. In this paper, we compute the values of the integral over the convolution of Equation (1) by numerical convolution and integration using a discretization $dx = 0.01$.

References

- Bamber, J. L., & Aspinall, W. P. (2013). An expert judgement assessment of future sea level rise from the ice sheets. *Nature Climate Change*, advance online.
<https://doi.org/10.1038/nclimate1778>
- Becker, J., Porter, E., & Centola, D. (2019). The wisdom of partisan crowds. *Proceedings of the National Academy of Sciences*, 201817195.
<https://doi.org/10.1073/pnas.1817195116>

- Berger, J. O. (1989). Statistical decision theory, 217–224.
https://doi.org/10.1007/978-1-349-20181-5_26
- Brown, G., Wyatt, J., Harris, R., & Yao, X. (2005). Diversity creation methods: A survey and categorisation. *Information Fusion*, 6(1), 5–20.
<https://doi.org/doi.org/10.1016/j.inffus.2004.04.004>
- Cohen, J. (1986). An epistemic conception of democracy. *Ethics*, 97(1), 26–38.
<https://doi.org/10.1086/292815>
- Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., . . . Leonard, N. E. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062), 1578–1580.
<https://doi.org/10.1126/science.1210280>
- Davis-Stober, C. P., Budescu, D. V., Dana, J., & Broomell, S. B. (2014). When is a crowd wise?. *Decision*, 1(2), 79–101. <https://doi.org/10.1037/dec0000004>
- Dehaene, S., Izard, V., Spelke, E., & Pica, P. (2008). Log or Linear? Distinct Intuitions of the Number Scale in Western and Amazonian Indigene Cultures. In *Science* (Vol. 320, pp. 1217–1220). <https://doi.org/10.1126/science.1156540>
- Flache, A., Mäs, M., Feliciani, T., Chattoe-Brown, E., Deffuant, G., Huet, S., & Lorenz, J. (2017). Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*, 20(4), 2. <https://doi.org/10.18564/jasss.3521>
- Frey, V., & Rijt, A. van de. (2020). Social influence undermines the wisdom of the crowd in sequential decision making. *Management Science*.
<https://doi.org/10.1287/mnsc.2020.3713>
- Galesic, M., Barkoczi, D., & Katsikopoulos, K. (2018). Smaller crowds outperform larger

- crowds and individuals in realistic task conditions. *Decision*, 5(1), 1–15.
<https://doi.org/10.1037/dec0000059>
- Galton, F. (1907a). One vote, one value. *Nature*, 75(1948), 414.
<https://doi.org/10.1038/075414a0>
- Galton, F. (1907b). The Ballot-Box. *Nature*, 75(1952), 509–510.
- Galton, F. (1907c). Vox populi. *Nature*, 75(1949), 450–451.
<https://doi.org/10.1038/075450a0>
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1), 1–58. <https://doi.org/10.1162/neco.1992.4.1.1>
- Gigone, D., & Hastie, R. (1997). Proper analysis of the accuracy of group judgments. *Psychological Bulletin*, 121(1), 149. <https://doi.org/10.1037/0033-2909.121.1.149>
- Golub, B., & Jackson, M. O. (2010). Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1), 112–149.
<https://doi.org/10.1257/mic.2.1.112>
- Goodin, R. E., & Spiekermann, K. (2018). *An epistemic theory of democracy*. Oxford University Press.
- Granovskiy, B., Gold, J. M., Sumpter, D. J. T., & Goldstone, R. L. (2015). Integration of social information by human groups. *Topics in Cognitive Science*, 7, 469–493.
<https://doi.org/10.1111/tops.12150>
- Hooker, R. H. (1907). Mean or Median. *Nature*, 75(1951), 487.
- Hora, S. C., Fransen, B. R., Hawkins, N., & Susel, I. (2013). Median aggregation of distribution functions. *Decision Analysis*, 10(4), 279–291.

<https://doi.org/10.1287/deca.2013.0282>

Jayles, B., Kim, H.-r., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., ... Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47), 12620–12625. <https://doi.org/10.1073/pnas.1703695114>

Kao, A. B., & Couzin, I. D. (2014). Decision accuracy in complex environments is often maximized by small group sizes. *Proceedings of the Royal Society B: Biological Sciences*, 281(1784). <https://doi.org/10.1098/rspb.2013.3305>

Karotkin, D., & Paroush, J. (2003). Optimum committee size: Quality-versus-quantity dilemma. *Social Choice and Welfare*, 20(3), 429–441.

Laan, A., Madirolas, G., & Polavieja, G. G. de. (2017). Rescuing collective wisdom when the average group opinion is wrong. *Frontiers in Robotics and AI*, 4, 56. <https://doi.org/10.3389/frobt.2017.00056>

Landemore, Hélène, & Elster, J. (Eds.). (2012). *Collective wisdom: Principles and mechanisms*. Cambridge University Press.

List, C., & Goodin, R. E. (2001). Epistemic democracy: Generalizing the condorcet jury theorem. *Journal of Political Philosophy*, 9(3), 277–306. <https://doi.org/10.1111/1467-9760.00128>

Lorenz, J. (2021). On the quantification of crowd wisdom. OSF. Retrieved from <http://osf.io/zwxg2>

Lorenz, J., Neumann, M., & Schröder, T. (2021). Individual attitude change and societal dynamics: Computational experiments with psychological theories. *Psychological Review*. Lorenz, Jan: Department of Psychology; Methods, Jacobs University

- Bremen gGmbH, Campus Ring 1, Bremen, Germany, 28759, post@janlo.de:
American Psychological Association. <https://doi.org/10.1037/rev0000291>
- Lorenz, J., Rauhut, H., & Kittel, B. (2015). Majoritarian democracy undermines truth-finding in deliberative committees. *Research & Politics*, (2), 1–10.
<https://doi.org/10.1177/2053168015582287>
- Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22), 9020–9025. <https://doi.org/10.1073/pnas.1008636108>
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(114), 175–208. Retrieved from
<http://ideas.repec.org/a/bla/restud/v38y1971i114p175-208.html>
- Nash, U. W. (2014). The curious anomaly of skewed judgment distributions and systematic error in the wisdom of crowds. *PLoS ONE*, 9(11), 1–17.
<https://doi.org/10.1371/journal.pone.0112386>
- O’Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., ... Rakow, T. (2006). *Uncertain judgements: Eliciting experts’ probabilities*. John Wiley & Sons.
- O’Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. *Statistical Science*, 502–518.
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton University Press.
- Pivato, M. (2011). Voting rules as statistical estimators. *Social Choice and Welfare*, 1–50.
<https://doi.org/10.1007/s00355-011-0619-1>

- Pivato, M. (2017). Epistemic democracy with correlated voters. *Journal of Mathematical Economics*, 72, 51–69. <https://doi.org/10.1016/j.jmateco.2017.06.001>
- Prelec, D., Seung, H. S., & McCoy, J. (2017). A solution to the single-question crowd wisdom problem. *Nature*, 541(7638), 532–535. Retrieved from <https://doi.org/10.1038/nature21054>
- Rauhut, H., Lorenz, J., Schweitzer, F., & Helbing, D. (2011). Reply to Farrell: Improved individual estimation success can imply collective tunnel vision. *Proceedings of the National Academy of Sciences*, 108(36), E626. <https://doi.org/10.1073/pnas.1111007108>
- Stone, M. (1961). The opinion pool. *Annals of Mathematical Statistics*, 4, 1339–1342. <https://doi.org/10.1214/aoms/1177704873>
- Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Doubleday Books.
- Tetlock, P., & Gardner, D. (2015). *Superforecasting – the art and science of prediction*. Random House Books.
- Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, 19(7), 645–647. <https://doi.org/10.1111/j.1467-9280.2008.02136.x>
- Wallis, K. F. (2014). Revisiting Francis Galton’s forecasting competition. *Statist. Sci.*, 29(3), 420–424. <https://doi.org/10.1214/14-STS468>
- Wikipedia contributors. (2021). Libor scandal — Wikipedia, the free encyclopedia. Retrieved from

https://en.wikipedia.org/w/index.php?title=Libor_scandal&oldid=1019350224

Notation	Explanation
θ	true value
$x_1, \dots, x_n$ (x as vector)	sample of estimates from n estimators
$\hat{x}_1, \dots, \hat{x}_n$ ($\hat{x}$ as vector)	the ordered sample of $\{x_1, \dots, x_n\}$
$\text{colD}(x)$	general aggregation function $\mathbb{R}^n \rightarrow \mathbb{R}$ for the collective decision
$\bar{x}$, $\text{mean}(x)$	$\frac{1}{n} \sum_{i=1}^n x_i$ arithmetic mean
$\text{geomean}(x)$	$\sqrt[n]{\prod_{i=1}^n x_i}$ geometric mean
$\tilde{x}$, $\text{median}(x)$	$\hat{x}_{\frac{n+1}{s}}$ if n odd, $(\hat{x}_{\frac{n}{s}} + \hat{x}_{\frac{n}{s}+1})/2$ if n even
$\text{collErr}(x, \theta)$	$(\bar{x} - \theta)^2$ collective error/population bias
$\text{Var}(x)$	$\frac{1}{n} \sum_{i=1}^n (\bar{x} - x_i)^2$ variance/diversity
$\text{MSE}(x, \theta)$	$\frac{1}{n} \sum_{i=1}^n (x_i - \theta)^2$ mean squared error/ average individual error
$\text{WoC}(x, \theta)$	$\frac{\max\{i \mid \hat{x}_i \leq \theta \leq \hat{x}_{n-i+1}\}}{\lceil n/2 \rceil}$ (zero if nominator is empty set), wisdom-of-crowd indicator
$\text{FoE}(x, \theta, \text{colD})$	$\#\{i \mid \text{colD}(x, \theta) - \theta < x_i - \theta \}/n$ fraction of outperformed estimates
X, f_X	the estimate of a randomly selected estimator as a random variable and its probability density function
$[X]_n$	vector of n independent and identical random variables of the type of X
$\text{colD}([X]_n), [\bar{X}]_n, [\tilde{X}]_n$	the random variables of the collective decision of $[X]_n$

226 estimates, x-axis clipped, Data: Jan Lorenz <http://janlo.de/wp/2010/06/22/die-weisheit-der-bremer/>

Figure 1. (A) Histogram of estimates for the weight-judging competition at the West of England Fat Stock and Poultry Exhibition in Plymouth 1907. (B) Histogram of estimates for the estimation game “How many lots will be sold at the end of the festival” at the ViertelFest Aug 22-24, 2008 in Bremen.

Figure 2. (A) The number of experts in a sample of size k and (B) the probability of drawing an expert against a sample of k when there is no bias ($b = 0$). (C) The same probability when the bias is $b = 0.5$. Assumptions are the estimates come from a standard normal distribution shifted by b , the collective decision is the arithmetic mean of the sample, and the truth is $\theta = 0$.

Figure 3. (A) The impact of the bias b on the optimal crowd size for which the probability to draw an expert is lowest. (B) The probability to draw an expert under the optimal crowd size. The blue line the probability for large groups ($k = \infty$), the black line the probability for the optimal k . The deviation of the blue and the black lines represents the risk of having worse collective estimates by collecting too many estimates. This region ($0.5 < b < 1.5$) is focused in the inset in the second panel.

Figure 4. Samples from Beta, lognormal, and normal distributions with mean 0.6 and standard deviations 0.15, 0.25, 0.35, and 0.45. The smaller blue pin is the sample's mean, the larger red pin its median. The y-axis is the same over all plots. The two extreme bins in the Beta plot for standard deviation 0.45 exceed the y-axis' maximum.

Figure 5. Robustness of the distance of the collective decision to the true value by the mean and the median depending on polarization (measured by standard deviation). Assumptions: Estimates of the majority of the population are Beta-distributed with a mean coinciding with the true value $\theta = 0.5$. A certain small fraction of the population estimates zero to manipulate the collective decision to be as low as possible.

Figure 6. Example for the probability density functions of X (standard normal distribution with bias $b = 0.5$), $[\bar{X}]_k$ (average of $k = 11$ versions of X , labelled X_{11}), $|X|$, $-|[\bar{X}]_k|$, and $|X| - |[\bar{X}]_k|$.