

Prediction Diversity and Selective Attention in the Wisdom of Crowds

Davi A. Nobre
José F. Fontanari

Instituto de Física de São Carlos, Universidade de São Paulo
Caixa Postal 369
São Carlos, SP 13560-970, Brazil

The wisdom of crowds is the idea that the combination of independent estimates of the magnitude of some quantity yields a remarkably accurate prediction, which is always more accurate than the average individual estimate. In addition, it is largely believed that the accuracy of the crowd can be improved by increasing the diversity of the estimates. Here we report the results of three experiments to probe the current understanding of the wisdom of crowds, namely, the estimates of the number of candies in a jar, the length of a paper strip and the number of pages of a book. We find that the collective estimate is better than the majority of the individual estimates in all three experiments. In disagreement with the prediction diversity theorem, we find no significant correlation between the prediction diversity and the collective error. The poor accuracy of the crowd on some experiments leads us to conjecture that its alleged accuracy is most likely an artifact of selective attention.

Keywords: wisdom of crowds; forecast combinations; diversity prediction theorem; judgment distributions

1. Introduction

Although the notion of the wisdom of crowds is more than a century old, being brought to light by Galton's 1907 seminal study of the estimation of the weight of an ox at the West of England Fat Stock and Poultry Exhibition in Plymouth [1] (see [2] for a historical account), it is still a subject of fascination for the general public [3] and for the scientific community [4] as well. This fascination stems from the often remarkably accurate estimate given by the mean of independent individual estimates of the magnitude of some quantity. For instance, in the ox-weighing contest, the crowd overestimated the weight of the ox by less than 1% of the correct weight [1].

Explanations for the accuracy of the collective estimate based on purely statistical arguments are not satisfactory, since they assume that the individual estimates are unbiased; that is, that the errors

spread in equal proportion around the correct value of the unknown quantity [5], whereas experimental evidence, as well as common sense, points to the existence of systematic errors in the individual estimates [6, 7].

While the accuracy of the collective estimate remained a sort of mystery, attention has been directed to the observation that the collective estimate is always better than the average individual estimate. In fact, it seems that for some researchers this is the defining characteristic of the wisdom of crowds [8]. This observation can be explained by a simple statistical argument, the so-called diversity prediction theorem [4], which asserts that the error of the collective estimate is less than or equal to the average error of the individual estimates. Here, average means simply the arithmetic mean of the individual estimates. This implies that, on the average, the collective estimate is better than the estimate of a randomly selected individual in the group. This finding has considerable practical importance as it guarantees that, in the event one does not know who the experts are, it is advantageous to combine the forecasts of all members of the group. Nonetheless, the diversity prediction theorem offers no clue at all to the accuracy of the collective estimate.

However, as hinted by its name, the diversity prediction theorem seems also to have a say in the role of the diversity of the individual estimates, the so-called prediction diversity. In fact, since the theorem asserts that the (quadratic) collective error equals the average (quadratic) individual error minus the diversity of the estimates, one is tempted to think that the increase of the prediction diversity would improve the collective estimate [4, 8]. Unfortunately, this result, which reflects somewhat the zeitgeist of the twenty-first century, does not follow from the diversity prediction theorem, since the average individual error and the diversity of the estimates are not independent quantities; that is, an increase of the prediction diversity may change the average individual error with unpredictable effects on the collective error. Nevertheless, unveiling the influence of the prediction diversity on the accuracy of the collective estimate is clearly a crucial issue for the understanding of the wisdom of crowds.

Accordingly, in this paper we report the results of three experiments to probe the wisdom of crowds phenomenon, namely, the classic estimate of the number of candies in a jar, the estimate of the length of a paper strip and the estimate of the number of pages of a book. The number of estimates in each experiment was over 100. To measure the correlation between the prediction diversity and the collective error, we produced a large number of virtual experiments by selecting N estimates at random and without replacement from the original ensemble of estimates of the real experiments. We find no significant correlation between the prediction diversity and the collective

error, thus implying that diversity has no predictive value for the accuracy of the collective estimate. Most interestingly, for the candies-in-a-jar and the pages-of-a-book experiments, where the collective estimates grossly missed the correct value, there is an optimal group size that maximizes the chances of high-accuracy collective predictions, similarly to the findings on distributed cooperative problem-solving systems [9]. However, for the paper-strip experiment, where the collective estimate was already very accurate, the chances of high-accuracy collective predictions (i.e., predictions that miss the correct value by less than 5%) increase monotonically with the group size N .

Regarding the accuracy of the collective estimate, our experiments support the view of systematic errors in the collective forecast [6, 7], which depend on the skills of the subjects on the proposed tasks. For example, in the paper-strip experiment the crowd underestimated the length by only 1.8%, whereas in the pages-of-a-book experiment the underestimate was of 28.4%. Our conclusion is that the high accuracy of the wisdom of crowds, which is responsible for its popularity among the general public, is an illusion resulting from selective attention that gives prominence to the successful outcomes only.

The rest of this paper is organized as follows. In Section 2 we present a brief review of the diversity prediction theorem [4] and introduce the basic quantities used to characterize the wisdom of crowds experiments. In Section 3 we describe and analyze the results of our three experiments, emphasizing the influence of the prediction diversity on the accuracy of the collective estimate. Finally, in Section 4 we summarize our results and present our concluding remarks.

2. Page's Diversity Prediction Theorem

The diversity prediction theorem [4] is considered a main attainment by those who celebrate the power of diversity to improve the performance of groups [8] (see, however, [10]), since it shows that the (quadratic) collective error can be related in a very simple manner to the average (quadratic) individual error and to a measure of the diversity of the estimates. More pointedly, let g_i be the estimate of some unknown quantity, such as the weight of the ox in Galton's experiment, by individual $i = 1, \dots, N$. We will consider the g_i as random variables that, as far as the diversity prediction theorem is concerned, need not be independent. In addition, let G be the true value of the unknown quantity the N individuals are trying to estimate and let the collective estimate be defined as the arithmetic mean of the individual estimates, that is, $\langle g \rangle = \sum_{i=1}^N g_i / N$. (We note that Galton used the median of the individual estimates as the collective estimate in the

ox-weighting experiment [1], though the arithmetic mean proved to be a better estimator in that case [2].) Thus, the quadratic collective error is defined as

$$\gamma = (\langle g \rangle - G)^2. \quad (1)$$

Next we define the average quadratic individual error

$$\epsilon = \frac{1}{N} \sum_{i=1}^N (g_i - G)^2 \quad (2)$$

and the diversity of the estimates

$$\delta = \frac{1}{N} \sum_{i=1}^N (g_i - \langle g \rangle)^2, \quad (3)$$

so that the identity

$$\gamma = \epsilon - \delta \quad (4)$$

follows straightforwardly. This identity is Page's diversity prediction theorem, which asserts that the (quadratic) collective error equals the average (quadratic) individual error minus the prediction diversity. This result is sometimes interpreted as the proof that increasing the prediction diversity δ results in the decrease of the collective error γ [8]. Of course, since δ and ϵ cannot be varied independently of each other, this interpretation is not correct.

In fact, the discussion on the value of Page's diversity prediction theorem is reminiscent of the arguments about the relevance of the celebrated Price equation [11] for evolutionary biology. We note that Price's equation, which has a straightforward derivation from the definition of fitness, is considered by many researchers as a mere mathematical tautology [12].

Here we take a pragmatic stance and carry out experiments to check whether the increase of the diversity of the estimates is likely to result in a decrease of the collective error, without entering into the merit of the diversity prediction theorem. It is interesting to note that the diversity of estimates δ is known in the statistical literature as the precision of the estimates, that is, the closeness of repeated estimates (of the same quantity) to one another [13]. The experiments and the analyses of their results are the subjects of the next section.

3. Experiments

We have carried out three experiments in which a number of science, technology, engineering and mathematics (STEM) students of the

University of São Paulo guessed independently the number of candies in a jar, the length of a paper strip and the number of pages of a book.

3.1 Estimating the Number of Candies in a Jar

This is a classic wisdom of crowds experiment [3] in which the individuals have to guess independently the number of candies in a jar (see [14] for a variant with non-independent guesses), resulting, typically, in a group estimate superior to the vast majority of the individual guesses [8]. In particular, 105 students guessed the number of candies in a transparent jar that held $G = 636$ candies. The group estimate $\langle g \rangle = 531$ was better than 70% of the individual estimates. In our setup, there was no penalty for the guess farthest from the correct answer and the reward for the best guess (630, in our case) was the jar of candies.

Figure 1 shows the histogram of the normalized guesses $g_i / \langle g \rangle$, which we attempted to fit using a two-piece normal distribution of mean 1 (see, e.g., [15]). Notice that the most probable estimate is about half of the correct number of candies in the jar. The asymmetry of the distribution of estimates is quite noticeable and is confirmed by the positive value of the moment coefficient of skewness $\tilde{\mu}_3 = 0.73$.

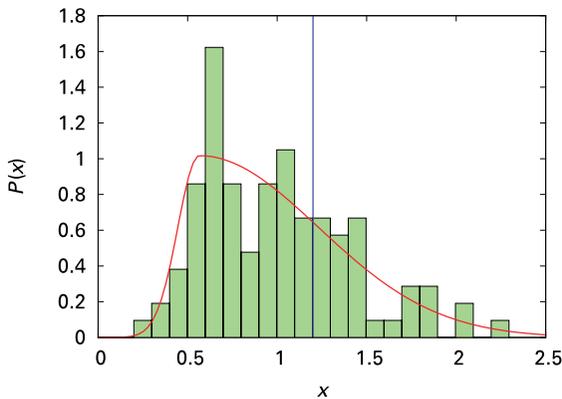


Figure 1. Histogram of the normalized estimates $x = g_i / \langle g \rangle$ for the candies-in-a-jar experiment. The vertical line indicates the ratio between the correct number of candies $G = 636$ and the estimate of the group $\langle g \rangle = 531$, that is, $G / \langle g \rangle \approx 1.2$. The solid curve is the best-fitting ($R^2 = 0.77$) two-piece normal distribution $Ae^{-(x-\mu)^2/2\sigma_1^2}$ if $x < \mu$ and $Ae^{-(x-\mu)^2/2\sigma_2^2}$ otherwise, where $A = [\sqrt{2\pi}(\sigma_1 + \sigma_2) / 2]^{-1}$ with $\sigma_1 = 0.12$, $\sigma_2 = 0.66$ and $\mu = 1 - (\sigma_2 - \sigma_1)\sqrt{2/\pi} \approx 0.565$, so that the mean of the fitting distribution equals 1.

This means that the odds of gross overestimation of the outcome are much greater than of underestimation due to the right-tailed nature of the distribution. It is interesting that Galton's ox-weighing experiment results in a left-skewed, left-tailed distribution of estimates [2]. In fact, in order to tackle the asymmetry of the distribution, Galton suggested the use of two normal distributions to fit the lower and the upper halves of the histogram of estimates, hence our choice of the two-piece normal distribution as the fitting distribution in Figure 1.

Since for a single experiment, we cannot draw any conclusion about the influence of the diversity of the estimates δ on the collective error γ , here we use the estimates of the candies-in-a-jar experiment to generate 10^4 virtual experiments. In each experiment, N estimates are drawn at random without replacement from the 105 original estimates. For each experiment, γ and δ are evaluated so we can draw the scatter plots shown in Figure 2. In particular, this figure shows the relative collective error $\gamma^{1/2}/G$ and the relative diversity $\delta^{1/2}/\langle g \rangle$ for each one of the 10^4 experiments for groups of size $N = 10, 20, 40$ and 60 . We note that the mean $\langle g \rangle$ depends on the particular experiment considered, so it assumes a different value for each point in the scatter plots. In terms of these dimensionless quantities, equation (4) is rewritten as

$$\frac{\gamma}{G^2} = \frac{\epsilon}{G^2} - \frac{\delta}{\langle g \rangle^2} \frac{\langle g \rangle^2}{G^2}, \quad (5)$$

which preserves the main points of the diversity prediction theorem, namely, the correct claim that $\gamma^{1/2}/G \geq \epsilon^{1/2}/G$ and the incorrect inference that the increase of the relative diversity $\delta^{1/2}/\langle g \rangle$ implies a decrease of the relative collective error $\gamma^{1/2}/G$. We hasten to note that although, in principle, these two quantities might be negatively correlated, this result does not follow from equation (5) since the increase of δ may result in a decrease of ϵ and $\langle g \rangle$.

The Pearson correlation coefficients between $\delta^{1/2}/\langle g \rangle$ and $\gamma^{1/2}/G$ are $r = -0.005$ for $N = 10$, $r = 0.04$ for $N = 20$, $r = 0.06$ for $N = 40$ and $r = 0.07$ for $N = 60$, so it is safe to state that the diversity of the estimates conveys very little, if any, information about the accuracy of the group estimate. We note that the center of mass of the data shown in the panels of Figure 2 is at $\delta^{1/2}/\langle g \rangle = 0.416$ and $\gamma^{1/2}/G = 0.165$ regardless of the value of N . Of course, these values coincide with those calculated using the original sample of 105 estimates.

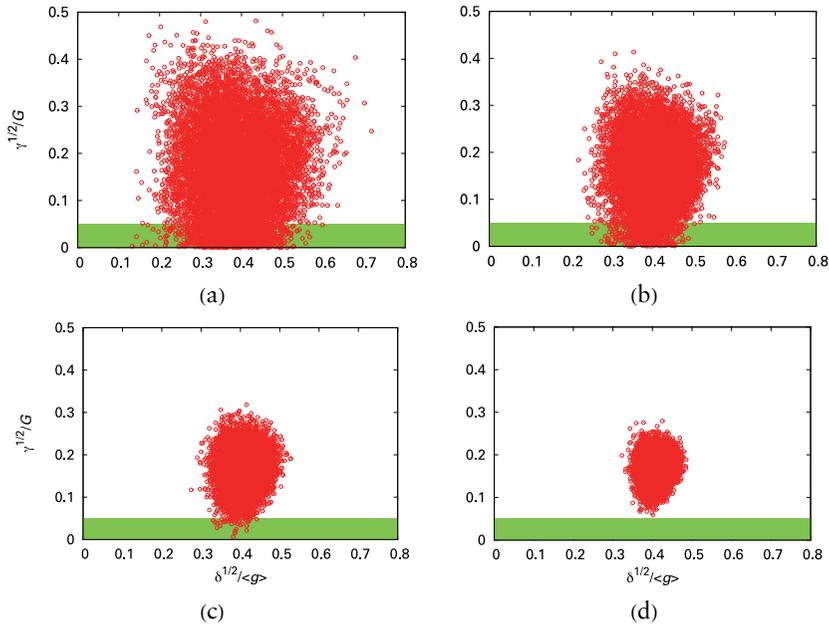


Figure 2. Scatter plots of the candies-in-a-jar experiment for groups of size (a) $N = 10$, (b) $N = 20$, (c) $N = 40$, (d) $N = 60$. The x axis is the relative standard deviation of the estimates $\delta^{1/2} / \langle g \rangle$, which measures the diversity of the estimates, whereas the y axis is the relative collective error $\gamma^{1/2} / G$, which measures the accuracy of the group estimate. The horizontal band at the bottom of the panels indicates the regions where the error of the group estimate is less than 5%.

An interesting result revealed in Figure 2 is that the probability that the error of the group estimate is less than 5%, which is given by the fraction of data points that fall within the horizontal band in the scatter plot, depends on the group size N . To quantify this effect, we present this probability in Figure 3 as a function of the group size. These startling results reveal the existence of an optimal group size ($N = 5$) that maximizes the odds of producing a high-accuracy estimate of the number of candies in the jar. In other words, picking $N = 5$ estimates at random out of the 105 original ones yields a 14% chance of producing a group estimate that misses the correct value by less than 5%. In addition, selection of a single estimate at random is more likely to result in such high-accuracy estimates than the aggregation of $N \geq 20$ estimates. The chance of such a precise estimate vanishes like $e^{-\beta N^2}$ with $\beta > 0$ for increasing N . This is so because the collective error of the original sample of 105 estimates, namely, 16.5%, does not tally with our definition of a high-accuracy estimate.

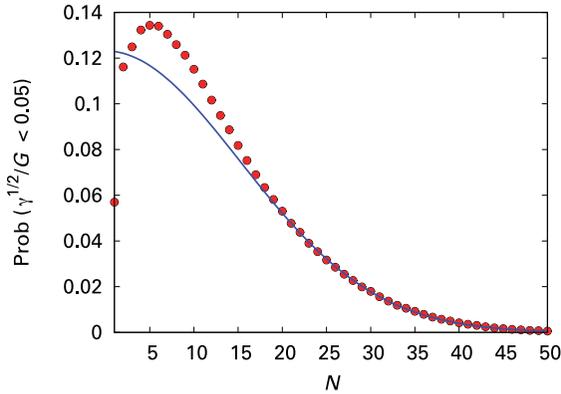


Figure 3. Probability that the collective error is less than 5% for the candies-in-a-jar experiment as a function of the group size N . The solid curve is the fitting function $\alpha e^{-\beta N^2}$ for the large N regime. The fitting parameters are $\alpha = 0.12$ and $\beta = 0.0021$.

3.2 Estimating the Length of a Paper Strip

Whereas the competence of the students to estimate the number of candies in a jar is very problematical, as our experiment was their first experience on that task, we expect them to be better skilled to size up the length of a paper strip. Accordingly, we asked 139 students to guess the length of a paper strip that measured $G = 22.4$ cm. The group estimate $\langle g \rangle = 22.0$ cm was better than 85% of the individual estimates and corresponds to an error of only 1.8%. The best guess was 22.5 cm. There was no reward or penalty for the subjects in this experiment.

The histogram of the normalized guesses $x = g_i / \langle g \rangle$ for the paper-strip experiment is shown in Figure 4 together with a best-fitting Gaussian distribution. The symmetry of the estimates around the mean and their small variance attest to the skill of the students to size up ordinary lengths. Although the noticeable asymmetry of the histogram of estimates for the candies-in-a-jar experiment leads us to follow Galton's suggestion and use a two-piece normal distribution as the fitting distribution (see Figure 1), a normal distribution seems much more suitable to fit the almost symmetric histogram of the paper-strip experiment.

The scatter plots of Figure 5 that show the properly normalized diversity of the estimates and the relative collective error for distinct group sizes reveal no meaningful correlation between these quantities. More pointedly, the Pearson correlation coefficients between $\delta^{1/2} / \langle g \rangle$ and $\gamma^{1/2} / G$ are $r = 0.28$ for $N = 10$, $r = 0.11$ for $N = 20$, $r = -0.049$ for $N = 40$ and $r = -0.11$ for $N = 60$. We note that the

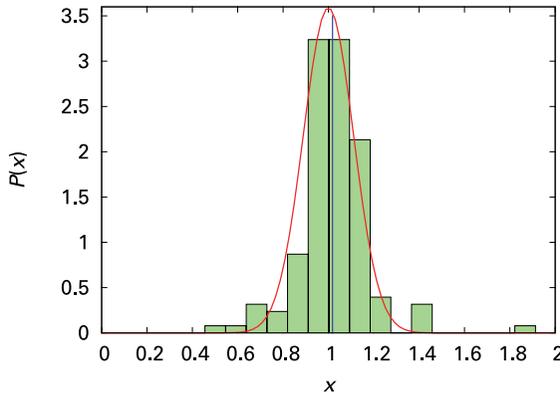


Figure 4. Histogram of the normalized estimates $x = g_i / \langle g \rangle$ for the paper-strip experiment. The vertical line indicates the ratio between the length of the paper strip $G = 22.4$ cm and the estimate of the group $\langle g \rangle = 22.0$ cm, that is, $G / \langle g \rangle \approx 1.018$. The solid curve is the best-fitting ($R^2 = 0.94$) Gaussian distribution $e^{-(x-1)^2/2\sigma^2} / \sqrt{2\pi\sigma^2}$ with $\sigma^2 = 0.012$ and mean 1.

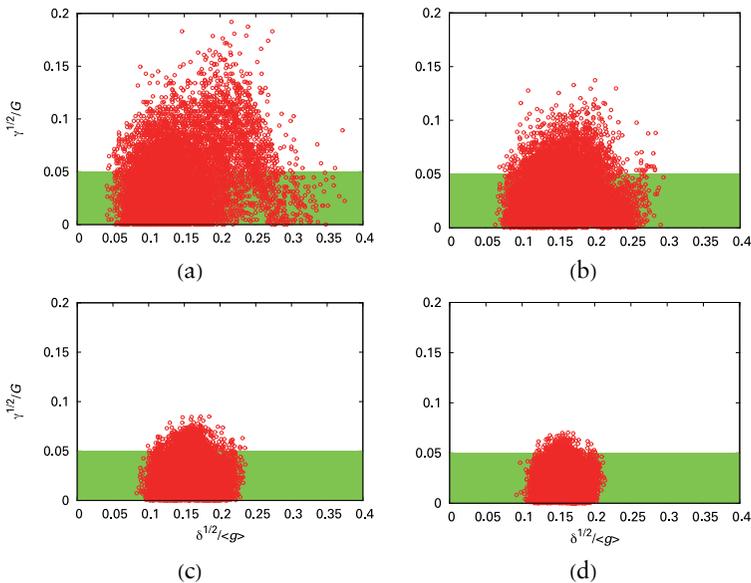


Figure 5. Scatter plots of the paper-strip experiment for groups of size (a) $N = 10$, (b) $N = 20$, (c) $N = 40$, (d) $N = 60$. The x axis is the relative standard deviation of the estimates $\delta^{1/2} / \langle g \rangle$, which measures the diversity of the estimates, whereas the y axis is the relative collective error $\gamma^{1/2} / G$, which measures the accuracy of the group estimate. The horizontal band at the bottom of the panels indicates the regions where the error of the group estimate is less than 5%.

claim that high prediction diversity leads or, more precisely, is associated with lower collective errors should be supported by a large negative correlation between $\delta^{1/2} / \langle g \rangle$ and $\gamma^{1/2} / G$. Although this correlation becomes more negative as N increases, it is too low to offer useful information on the wisdom of crowds puzzle. The center of mass of the data shown in the panels of Figure 5 is at $\delta^{1/2} / \langle g \rangle = 0.16$ and $\gamma^{1/2} / G = 0.018$ regardless of the value of N . The fraction of group estimates whose error is less than 5% shown in Figure 6 contrasts starkly with the results for the candies-in-a-jar experiment (see Figure 3), as now the group accuracy increases with increasing N . In addition, there is a 30% chance that a randomly selected estimate yields a prediction within the 5% accuracy range, which certifies the competence of the students to gauge lengths.

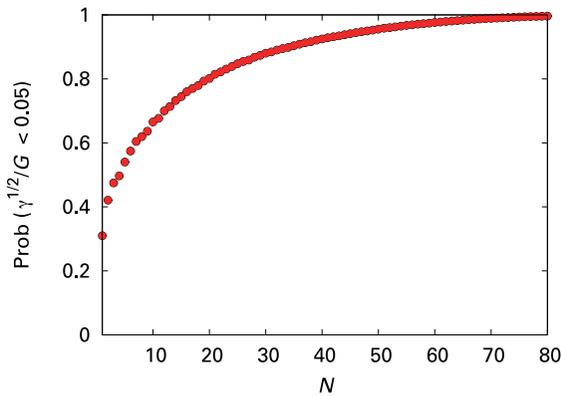


Figure 6. Probability that the collective error is less than 5% for the paper-strip experiment as a function of the group size N .

The stark difference between Figures 3 and 6 is somewhat reminiscent of Condorcet's jury theorem (see, e.g., [16]) in the sense that adding more voters (i.e., increasing the number of estimates N) may either improve or degrade the collective performance, depending on some circumstances. In Condorcet's theorem, it is the probability that an individual votes for the correct decision that determines whether a single voter or a jury will maximize the probability of making the correct decision. In our case, this role is played by the percent error of the collective estimate: if it is too large, our results indicate that discarding individual estimates at random will increase the probability of producing a high-accuracy collective estimate.

3.3 Estimating the Number of Pages of a Book

The same 139 students who were asked to estimate the length of the paper strip also estimated the number of pages of a book of $G = 784$ pages. As in the previous experiment, there was no reward or penalty for the subjects. We recall that those students were very accurate with their estimates of the length of a paper strip (see Figure 4). Surprisingly, their collective estimate of the number of pages of the book was kind of disastrous: $\langle g \rangle = 561$, which was superior to only 63% of the individual estimates and corresponds to an error of 28.4%. The best guess was 800 pages, which corresponds to an error of only 2%.

Figure 7 shows the histogram of the normalized guesses $g_i / \langle g \rangle$, which, following Galton’s original hunch, we attempted to fit using a two-piece normal distribution of mean 1. The histogram exhibits a considerable asymmetry that is measured by the moment coefficient of skewness $\tilde{\mu}_3 = 1.22$. The analysis of the prediction power of random assembled combinations of N estimates yields results qualitatively similar to those of the candies-in-a-jar experiment. In particular, the scaled collective error is $\gamma^{1/2} / G = 0.28$ and the scaled prediction diversity is $\delta^{1/2} / \langle g \rangle = 0.36$. The pages-of-a-book experiment is an excellent illustration of the largely unsung but fundamental fact that the collective estimate may not produce accurate predictions.

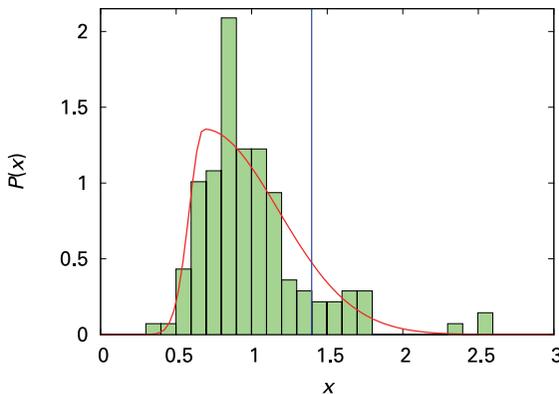


Figure 7. Histogram of the normalized estimates $x = g_i / \langle g \rangle$ for the pages-of-a-book experiment. The vertical line indicates the ratio between the correct number of pages $G = 784$ and the estimate of the group $\langle g \rangle = 561$, that is, $G / \langle g \rangle \approx 1.4$. The solid curve is the best-fitting ($R^2 = 0.84$) two-piece normal distribution $Ae^{-(x-\mu)^2/2\sigma_1^2}$ if $x < \mu$ and $Ae^{-(x-\mu)^2/2\sigma_2^2}$ otherwise, where $A = [\sqrt{2\pi} (\sigma_1 + \sigma_2) / 2]^{-1}$ with $\sigma_1 = 0.09$, $\sigma_2 = 0.49$ and $\mu = 1 - (\sigma_2 - \sigma_1)\sqrt{2/\pi} \approx 0.682$, so that the mean of the fitting distribution equals 1.

4. Discussion

The belief that cooperation can aid a group of agents to solve problems more efficiently than if those agents worked in isolation is a commonplace [17, 18, 19], although the factors that make cooperation effective still need much straightening out [20]. In fact, this is the main issue addressed by the research on distributed cooperative or parallel problem-solving systems [21], since cooperation may well lead the group astray, resulting in the so-called madness of crowds as neatly expressed by MacKay almost two centuries ago: “Men, it has been well said, think in herds; it will be seen that they go mad in herds, while they only recover their senses slowly, and one by one.” [22].

However, the notion that a collection of independently deciding individuals is likely to predict better than individuals or even experts within the group—a phenomenon dubbed the wisdom of crowds [3]—is much more controversial. The first report of this phenomenon in the literature was probably Galton’s account of the surprisingly accurate estimate of the weight of an ox given by the median of the sample of the individual guesses [1]. Although much of the evidence of the wisdom of crowds is anecdotal (see, e.g., [3]), there are a few efforts aiming at explaining this phenomenon either using a purely statistical framework [4] or using psychological arguments on the nature of the individual estimates [6, 7].

The main difficulty to approaching the wisdom of crowds phenomenon in a non-contentious manner is that it seems to have distinct meanings for different researchers. For instance, some researchers view it as the idea that a crowd can solve problems better than most individuals in it, including experts [8]. We note, however, that this is not what Page’s diversity prediction theorem asserts. In fact, equation (4) asserts that $\gamma \geq \epsilon$, where γ is the quadratic collective error and ϵ is the average quadratic individual error, which equals the expected error of the estimate of a randomly selected individual in the group. Hence, the theorem offers no guarantee that the collective estimate will be better than the estimates of most individuals in the group. Nevertheless, for the pages-of-a-book experiment we find that the collective estimate is better than 63% of the individual estimates, whereas this figure increases to 85% for the paper-strip experiment. We note that the specification “including experts” in the above definition is misleading because one may be led to believe incorrectly that the collective estimate is better than the experts’. For example, even in the paper-strip experiment, for which the collective estimate was highly accurate, it was outperformed by 15% of the individual estimates.

At this stage we should note that our disagreement over Page’s diversity prediction theorem $\gamma = \epsilon - \delta$ (see equation (4)), or, more correctly, over the interpretation of that theorem, since its proof is

straightforward, is that one cannot infer how a change in the diversity of the estimates δ will influence the collective error γ , because the mean individual error ϵ will also be affected by that change. Nevertheless, the effect of δ on γ is an important issue that can be addressed empirically. Since it is not feasible to carry out many independent experiments to calculate the correlation between these quantities, here we produced those experiments artificially by selecting N estimates at random and without replacement from the original set of estimates of our experiments. Our results (see Figures 2 and 5) indicate that there is no significant correlation between γ and δ ; that is, diversity has no predictive value at all for the accuracy of the collective estimate.

We think the reason the phenomenon of the wisdom of crowds caught on has little to do with the fact that on the average the collective estimate improves upon randomly chosen individual estimates, which can actually be quite poor, as illustrated by our pages-of-a-book experiment. For many researchers (see, e.g., [7]), the real riddle is the surprisingly good accuracy of the collective estimate that, in Galton's seminal experiment, missed the correct weight by only 0.8% [1]. A possible explanation involves the combination of forecasts [5], which, on the condition that the individual forecasts are unbiased, guarantees that the accuracy of the combined estimation increases as the number of independent estimates increases. The trouble with this approach is the implausibility of the assumption that the individual estimates are unbiased, that is, that their means coincide with the correct value of the quantity being estimated. (It is interesting that the assumption of unbiased individual estimates means that one could harvest the benefits of the wisdom of crowds by asking a single individual to make several estimates at different times [23].) On the contrary, there seems to be a systematic error in the wisdom of crowds, so that the collective estimate depends on some typical, group-dependent belief about the value of the unknown quantity [6, 7]. Since this typical value may be quite apart from the correct value, there is no guarantee of the accuracy of the collective estimate, which is precisely the conclusion we draw from the experiments reported here. Hence, the high accuracy of the collective estimate, which gives the wisdom of crowds its popular appeal, is most likely an artifact of selective attention or cherry picking that gives prominence to the successful outcomes only.

Acknowledgments

The research of J. F. Fontanari was supported in part by Grant No. 2017/23288-0, Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP) and by Grant No. 305058/2017-7, Conselho

Nacional de Desenvolvimento Científico e Tecnológico (CNPq). D. A. Nobre was supported by grant 2017/08475-9, Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP).

References

- [1] F. Galton, “Vox Populi,” *Nature*, 75(1949), 1907 pp. 450–451. doi:10.1038/075450a0.
- [2] K. F. Wallis, “Revisiting Francis Galton’s Forecasting Competition,” *Statistical Science*, 29(3), 2014 pp. 420–424. doi:10.1214/14-STS468.
- [3] J. Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*, New York: Doubleday, 2004.
- [4] S. E. Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton, NJ: Princeton University Press, 2007.
- [5] J. M. Bates and C. W. J. Granger, “The Combination of Forecasts,” *Operational Research Quarterly*, 20(4), 1969 pp. 451–468. doi:10.2307/3008764.
- [6] U. W. Nash, “The Curious Anomaly of Skewed Judgment Distributions and Systematic Error in the Wisdom of Crowds,” *PLoS ONE*, 9(11), 2014 e112386. doi:10.1371/journal.pone.0112386.
- [7] U. W. Nash, “Sequential Sampling, Magnitude Estimation, and the Wisdom of Crowds,” *Journal of Mathematical Psychology*, 77, 2017 pp. 165–179. doi:10.1016/j.jmp.2017.01.001.
- [8] M. J. Mauboussin. “Explaining the Wisdom of Crowds.” *Mauboussin on Strategy*. (Oct 6, 2020) trendfollowing.com/whitepaper/ExplainingWisdom.pdf.
- [9] J. F. Fontanari, “Imitative Learning as a Connector of Collective Brains,” *PLoS ONE*, 9(10), 2014 e110517. doi:10.1371/journal.pone.0110517.
- [10] J. F. Fontanari, “When More of the Same Is Better,” *Europhysics Letters*, 113(2), 2016 28009. doi:10.1209/0295-5075/113/28009.
- [11] G. R. Price, “Selection and Covariance,” *Nature*, 227(5257), 1970 pp. 520–521. doi:10.1038/227520a0.
- [12] S. A. Frank, “Natural Selection. IV. The Price Equation,” *Journal of Evolutionary Biology*, 25(6), 2012 pp. 1002–1019. doi:10.1111/j.1420-9101.2012.02498.x.
- [13] P.-N. Tan, M. Steinbach and V. Kumar, *Introduction to Data Mining*, Essex: Pearson Education Limited, 2014.

- [14] A. J. King, L. Cheng, S. D. Starke and J. P. Myatt, “Is the True ‘Wisdom of the Crowd’ to Copy Successful Individuals?,” *Biology Letters*, 8(2), 2011 pp. 197–200. doi:10.1098/rsbl.2011.0795.
- [15] K. F. Wallis, “The Two-Piece Normal, Binormal, or Double Gaussian Distribution: Its Origin and Rediscoveries,” *Statistical Science*, 29(1), 2014 pp. 106–112. doi:10.1214/13-STS417.
- [16] D. Austen-Smith and J. S. Banks, “Information Aggregation, Rationality, and the Condorcet Jury Theorem,” *American Political Science Review*, 90(1), 1996 pp. 34–45. doi:10.2307/2082796.
- [17] B. A. Huberman, “The Performance of Cooperative Processes,” *Physica D: Nonlinear Phenomena*, 42(1–3), 1990 pp. 38–47. doi:10.1016/0167-2789(90)90065-W.
- [18] S. H. Clearwater, B. A. Huberman and T. Hogg, “Cooperative Solution of Constraint Satisfaction Problems,” *Science*, 254(5035), 1991 pp. 1181–1183. doi:10.1126/science.254.5035.1181.
- [19] S. M. Reia, A. C. Amado and J. F. Fontanari, “Agent-Based Models of Collective Intelligence,” *Physics of Life Reviews*, 31, 2019 pp. 320–331. doi:10.1016/j.plrev.2018.10.004.
- [20] S. M. Reia, L. F. Aquino and J. F. Fontanari, “The Surprising Little Effectiveness of Cooperative Algorithms in Parallel Problem Solving,” *The European Physical Journal B*, 93(7), 2020 140. doi:10.1140/epjb/e2020-10199-9.
- [21] D. Lazer and A. Friedman, “The Network Structure of Exploration and Exploitation,” *Administrative Science Quarterly*, 52(4), 2007 pp. 667–694. doi:10.2189/asqu.52.4.667.
- [22] C. MacKay, *Extraordinary Popular Delusions and the Madness of Crowds*, London: Richard Bentley, 1841.
- [23] E. Vul and H. Pashler, “Measuring the Crowd Within: Probabilistic Representations within Individuals,” *Psychological Science*, 19(7), 2008 pp. 645–647. doi:10.1111/j.1467-9280.2008.02136.x.