4-1-2023

People Prefer More Information About Uncertainty, But Perform Worse When Given This Information: An Explanation of the Paradoxical Phenomenon

Jieqiong Zhao
Arizona State University, jieqiong.zhao@asu.edu

Olga Kosheleva
The University of Texas at El Paso, olgak@utep.edu

Vladik Kreinovich
The University of Texas at El Paso, vladik@utep.edu

Comments:
Technical Report: UTEP-CS-23-13

People Prefer More Information About Uncertainty, But Perform Worse When Given This Information: An Explanation for the Paradoxical Phenomenon

Jieqiong Zhao, Olga Kosheleva, and Vladik Kreinovich

Abstract In a recent experiment, decision makers were asked whether they would prefer having more information about the corresponding situation. They confirmed this preference, and such information was provided to them. However, strangely, the decisions of those who received this information were worse than the decisions of the control group – that did not get this information. In this paper, we provide an explanation for this paradoxical situation.

1 Formulation of the Problem

When making a decision, it is desirable to have as much information as possible. To make a decision, it is desirable to have as much information about the decision making situation as possible: additional information can help make a better decision.

This desirability has been confirmed by many polls; see, e.g., [17].

This includes the need for more information about uncertainty. In complex situations, to make decisions, we usually use computers. In general, computers process numbers. Thus, the information about a situation usually consists of several numbers, e.g., the values of the corresponding physical quantities. These values usually

Jieqiong Zhao
School of Computing and Augmented Intelligence, Arizona State University
699 S Mill Ave, Tempe, AZ 85281, USA, e-mail: jieqiong.zhao@asu.edu

Olga Kosheleva
Department of Teacher Education, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: olgak@utep.edu

Vladik Kreinovich
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: vladik@utep.edu

come from measurements – it could be direct measurement or so-called indirect measurement, i.e., processing of measurement results.

Measurements are never absolutely accurate, there is always a non-zero difference $\Delta x \stackrel{\text{def}}{=} \widetilde{x} - x$ between the measurement result $\widetilde{x}$ and the actual (unknown) value x of the corresponding quantity; see, e.g., [14]. This difference is known as the *measurement error*.

In view of the desirability to get as much information as possible, it is desirable to provide the users not only the measurement results, but also with the information about the measurement error. The desirability of this information was indeed confirmed by the polls; see, e.g., [17].

What do we know about the measurement errors: probabilistic and interval uncertainty. Ideally, we should know what are the possible values of the measurement error, and what is the frequency (probability) with which each of these values will appear in the actual measurements. In other words, ideally, we should know the probability distribution of the measurement errors.

To determine this probability distribution, we need to *calibrate* the measuring instrument, i.e., to compare the values $\widetilde{x}_k$ measured by this instrument and the values x'_k measured (for the same actual quantity) by a much more accurate measuring instrument (known as *standard*). The standard measuring instrument (SMI) is so much more accurate that, in comparison with the measurement errors of the tested measuring instrument (MI), we can safely ignore the measurement errors of the SMI and assume that SMI's measurement results x'_k are equal to the corresponding actual values x_k. Thus, the differences $\widetilde{x}_k - x'_k$ can be safely assumed to equal to the values of the measurement errors $\widetilde{x}_k - x_k$. So, from the resulting sample of these differences, we can determine the desired probability distribution.

This procedure is rather time-consuming and expensive – since the use of the complex standard measuring instruments is not cheap. As a result, in practice, such a detailed calibration is often not performed. Instead, the manufacturer of the measuring instruments provides a guaranteed upper bound Δ on the absolute value of the measurement error: $|\Delta x| \le \Delta$. Under this information, after we perform the measurement and get the measurement result $\widetilde{x}$, the only information that we have about the actual (unknown) value of the measured quantity is that this value lies somewhere in the interval $[\underline{x}, \overline{x}] \stackrel{\text{def}}{=} [\widetilde{x} - \Delta, \widetilde{x} + \Delta]$. This uncertainty is known as *interval uncertainty*; see, e.g., [4, 6, 8, 9].

Probability distribution of the measurement error is often Gaussian. In many practical cases, the measurement error is caused by the large number of small independent factors. It is known that the probability distribution of the joint effect of many independent small factors is close to Gaussian; the corresponding mathematical result is known as the Central Limit Theorem; see, e.g., [16]. Thus, we can safely assume that the measurement error is distributed according to the normal (Gaussian) distribution.

This assumption is in good accordance with the empirical data, according to which in the majority of the cases, the distribution is indeed Gaussian [12, 13].

The Gaussian distribution is uniquely determined by two parameters: mean μ and standard deviation σ. By testing the measuring instrument – i.e., by comparing its results with the results of a much more accurate ("standard") measuring instrument, we can determine the bias – as the arithmetic average of the observed differences between the measurements by two instruments:

$$\mu \approx \frac{1}{n} \cdot \sum_{k=1}^{n} \left(\widetilde{x}_k - x'_k\right).$$

Once we know μ, we can recalibrate the scale – i.e., subtract this mean value μ from each measurement result, i.e., replace each value $\widetilde{x}$ with $\widetilde{x}' \stackrel{\text{def}}{=} \widetilde{x} - \mu$. After that, the mean will be 0.

Also, we can determine the standard deviation as the mean squared value of the (re-calibrated) difference

$$\sigma \approx \sqrt{\frac{1}{n} \cdot \sum_{k=1}^{n} \left(\widetilde{x}'_k - x'_k\right)^2}.$$

So, it makes sense to assume that the measurement error is normally distributed with mean 0 and the known standard deviation σ.

Confidence intervals. It is known that with high probability, all the values of the normally distributed random variable are located within the interval

$$[\mu - k_o \cdot \sigma, \mu + k_0 \cdot \sigma]$$

for an appropriate k_0: for $k_0 = 2$, this is true with probability 95%, for $k_0 = 3$, this is true with probability 99.9%, and for $k_0 = 6$, this is true with probability $1 - 10^{-8}$. These intervals are particular case of *confidence intervals*. These intervals is what is often supplied to the users as a partial information about the probability distributions.

What happened when decision maker received confidence intervals: paradoxical situation. Since the decisions makers expressed the desire to receive information about uncertainty, they were supplied, in addition to measurement results, with the corresponding confidence intervals. And here, a strange thing happened. One would expect that this additional information would help the decision makers make better decisions – or at least did not degrade the quality of their decisions. However, in reality, the decisions of those users who got this additional information were worse than the decisions make by users from the control group – that did not receive this information. How can we explain this paradoxical phenomenon?

2 Our Explanation

Preliminaries. To understand the above paradoxical phenomenon, let us recall how rational people make decisions. Let us start with the case of full information, when we know all the probabilities.

How rational people make decisions under full information: reminder. According to decision theory (see, e.g., [1, 2, 5, 7, 10, 11, 15]), preferences of each rational person can be described by assigning, to each alternative x, a numerical value $u(x)$ called its *utility*. To assign these numerical values, one need to select two alternatives: a very good alternative A_+ which is better than anything that the decision maker will actually encounter, and a very bad alternative A_- which is worse than anything that the decision maker will actually encounter. Once these two alternatives are selected, we can form, for each value p from the interval $[0,1]$, a lottery $L(p)$ in which we get A_+ with probability p and A_- with the remaining probability $1-p$.

For any actual alternative x, when p is close to 0, the lottery $L(p)$ is close to A_- and is, thus, worse than x: $A_- < x$. When p is close to 1, the lottery $L(p)$ is close to A_+ and is, thus, better than x: $x < A_+$. When we move from 0 to 1, there is a threshold value at which the relation $L(p) < x$ is replaced with $x < L(p)$. This threshold value $u(x)$ is what is called the utility of x. By definition of utility, each alternative x is equivalent, to the decision maker, to a lottery $L(u(x))$ in which we get A_+ with probability $u(x)$ and we get A_- with the remaining probability $1-u(x)$.

In general, the larger the probability p of getting the very good alternative A_+, the better the lottery. Thus, if we need to select between several lotteries of this type, we should select the lottery with the largest values of the probability p. Since each alternative x is equal to the lottery $L(p)$ with probability $p = u(x)$, this means that we should always select the alternative with the largest value of utility.

What is the utility of an action in which we get n possible outcomes, for each of which we know its probability p_i and its utility u_i? By definition of utility, each outcome is equivalent to a lottery $L(u_i)$ in which we get A_+ with probability u_i and A_- wioth probability $1-u_i$. Thus, to the user, the action is equivalent to a 2-stage lottery in which we first select each i with probability p_i and then, depending on what i we selected on the first stage, select A_+ with probability u_i and A_- with probability $1-u_i$. As a result of this two-stage lottery, we get either A_+ or A_-, and the probability u of getting A_+ is determined by the formula of full probability:

$$u = p_1 \cdot u_1 + \ldots + p_n \cdot u_n.$$

Thus, by definition of utility, this value u is the utility of the action under consideration. The right-hand side of the formula for u is the expected value of the utility. So, we can conclude that the utility of an action with random consequences is equal to the expected value of the utility of different consequences.

Comment. The numerical value of the utility depends on the selection of the alternatives A_+ and A_-. If we select a different pair A'_+ and A'_-, then we will have different

numerical values $u'(x)$. It turns out that for every two pairs of alternatives, there exist real numbers $a > 0$ and b for which, for each x, we have $u'(x) = a \cdot u(x) + b$. In other words, utility is defined modulo a linear transformation.

This is similar to the fact that, e.g., the numerical value of the moment of time also depends on what starting point we use to measure time and what measuring unit we use, and all the scales are related to each other by an appropriate linear transformation $t' = a \cdot t + b$ for some $a > 0$ and b.

How rational people make decisions under interval uncertainty. As we have mentioned, in many practical situations, we only know the values of the quantities (that describe the state of the world) with interval uncertainty. In this case, we can describe the consequences – and their utility – also under interval uncertainty. In other words, for each possible decision x, instead of the exact value $u(x)$ of the corresponding utility, we only know the interval $[\underline{u}(x), \overline{u}(x)]$ of possible utility values.

How can we make a decision in this case? To make a decision, we need to assign, to each interval $[\underline{u}, \overline{u}]$, an equivalent numerical value $u(\underline{u}, \overline{u})$. As we have mentioned, utility is defined modular a linear transformation. There is no fixed selection of the alternatives A_+ and A_-, so it makes sense to require that the function $u(\underline{u}, \overline{u})$ remains the same for all the scales, i.e., that if $u = u(\underline{u}, \overline{u})$, then for all $a > 0$ and b, we should have $u' = u(\underline{u}', \overline{u}')$, where $u' = a \cdot u + b$, $\underline{u}' = a \cdot \underline{u} + b$, and $\overline{u}' = a \cdot \overline{u} + b$.

Let us denote $\alpha \stackrel{\text{def}}{=} u(0,1)$. If we know that the utility is between 0 and 1, then the situation is clearly better (or at least as good) than when utility is 0, and worse (or at least as good) then when utility is 1. Thus, we must have $0 \le \alpha \le 1$. Every interval $[\underline{u}, \overline{u}]$ can be obtained from the interval $[0,1]$ by a linear transformation $u \mapsto a \cdot t + b$ for $a = \overline{u} - \underline{u}$ and $b = \underline{u}$. Thus, due to invariance, from $\alpha = u(0,1)$, we can conclude that

$$u(\underline{u}, \overline{u}) = a \cdot \alpha + b = \alpha \cdot (\overline{u} - \underline{u}) + \underline{u} = \alpha \cdot \overline{u} + (1 - \alpha) \cdot \underline{u}.$$

This formula was first proposed by the Nobelist Leo Hurwicz and is thus known as *Hurwicz optimism-pessimism criterion*; see, e.g., [3, 5, 7]. The name comes from the fact that for $\alpha = 1$, this means only taking into account the best-case value $\overline{u}$ – the case of extreme optimism, while for $\alpha = 0$, this means only taking into account the worst-case value $\underline{u}$ – the case of extreme pessimism. The value α is different for different decision makers, depending on their level of optimism and pessimism.

Finally, an explanation. Now we are ready to produce the desired explanation. Let us consider the simplest possible setting, when the decision maker is directly provided with the information about his/her utility $u(x)$ of each possible decision x, and the difference $\Delta u(x) \stackrel{\text{def}}{=} \widetilde{u}(x) - u(x)$ between the estimated utility $\widetilde{u}(x)$ and the actual utility value $u(x)$ (corresponding to the actual –unknown – values of the corresponding quantities) is distributed according to normal distribution with 0 mean and standard deviation σ.

In the *ideal* case, when the decision maker knows this distribution, his/her equivalent utility of each possible decision x is equal to the expected value of the random utility value $u(x) = \widetilde{u}(x) - \Delta u(x)$, i.e., to the value $\widetilde{u}(x)$.

This was the ideal case. In the above experiment, we never report the whole distribution to the decision maker. Instead, we report either a single value $\widetilde{u}(x)$ or the confidence interval $[\widetilde{u}(x) - k_0 \cdot \sigma, \widetilde{u}(x) + k_0 \cdot \sigma]$.

In the first case, when we supply no information about uncertainty, the decision maker uses the provided value $\widetilde{u}(x)$ in his/her decisions. It so happens that this value is exactly what we would get if we knew the exact distributions. In other words, in this case, the decision maker makes an optimal decision.

On the other hand, if we provide the decision maker with the confidence interval, the decision maker – using Hurwicz criterion – will assign, to each possible decision x, an equivalent value

$$\alpha \cdot (\widetilde{u}(x) + k_0 \cdot \sigma) + (1 - \alpha) \cdot (\widetilde{u}(x) - k_0 \cdot \sigma) = \widetilde{u}(x) + (2\alpha - 1) \cdot k_0 \cdot \sigma.$$

Thus, for almost all possible values α from the interval $[0, 1]$ – with the only exception of the value $\alpha = 0.5$ – this value will be different from the optimal value (corresponding to the full information). So, the decision based on such values will be not as good as the optimal decision – and this is exactly what we observe in the above-described seemingly paradoxical experiment.

Comment. Note that the worsening of the decision happens when we provide the decision make with *partial* information about uncertainty. If we provide the decision maker with full information, the decision will, of course, be optimal.

Acknowledgments

This work was supported in part by the National Science Foundation grants 1623190 (A Model of Change for Preparing a New Generation for Professional Practice in Computer Science), HRD-1834620 and HRD-2034030 (CAHSI Includes), EAR-2225395, and by the AT&T Fellowship in Information Technology.

It was also supported by the program of the development of the Scientific-Educational Mathematical Center of Volga Federal District No. 075-02-2020-1478, and by a grant from the Hungarian National Research, Development and Innovation Office (NRDI).

References

1. P. C. Fishburn, *Utility Theory for Decision Making*, John Wiley & Sons Inc., New York, 1969.
2. P. C. Fishburn, *Nonlinear Preference and Utility Theory*, The John Hopkins Press, Baltimore, Maryland, 1988.
3. L. Hurwicz, *Optimality Criteria for Decision Making Under Ignorance*, Cowles Commission Discussion Paper, Statistics, No. 370, 1951.
4. L. Jaulin, M. Kiefer, O. Didrit, and E. Walter, *Applied Interval Analysis, with Examples in Parameter and State Estimation, Robust Control, and Robotics*, Springer, London, 2001.

5. V. Kreinovich, "Decision making under interval uncertainty (and beyond)", In: P. Guo and W. Pedrycz (eds.), *Human-Centric Decision-Making Models for Social Sciences*, Springer Verlag, 2014, pp. 163–193.
6. B. J. Kubica, *Interval Methods for Solving Nonlinear Constraint Satisfaction, Optimization, and Similar Problems: from Inequalities Systems to Game Solutions*, Springer, Cham, Switzerland, 2019.
7. R. D. Luce and R. Raiffa, *Games and Decisions: Introduction and Critical Survey*, Dover, New York, 1989.
8. G. Mayer, *Interval Analysis and Automatic Result Verification*, de Gruyter, Berlin, 2017.
9. R. E. Moore, R. B. Kearfott, and M. J. Cloud, *Introduction to Interval Analysis*, SIAM, Philadelphia, 2009.
10. H. T. Nguyen, O. Kosheleva, and V. Kreinovich, "Decision making beyond Arrow's 'impossibility theorem', with the analysis of effects of collusion and mutual attraction", *International Journal of Intelligent Systems*, 2009, Vol. 24, No. 1, pp. 27–47.
11. H. T. Nguyen, V. Kreinovich, B. Wu, and G. Xiang, *Computing Statistics under Interval and Fuzzy Uncertainty*, Springer Verlag, Berlin, Heidelberg, 2012.
12. P. V. Novitskii and I. A. Zograph, *Estimating the Measurement Errors*, Energoatomizdat, Leningrad, 1991 (in Russian).
13. A. I. Orlov, "How often are the observations normal?", *Industrial Laboratory*, 1991, Vol. 57, No. 7, pp. 770–772.
14. S. G. Rabinovich, *Measurement Errors and Uncertainty: Theory and Practice*, Springer Verlag, New York, 2005.
15. H. Raiffa, *Decision Analysis*, McGraw-Hill, Columbus, Ohio, 1997.
16. D. J. Sheskin, *Handbook of Parametric and Nonparametric Statistical Procedures*, Chapman and Hall/CRC, Boca Raton, Florida, 2011.
17. J. Zhao, Y. Wang, M. Mancenido, E. K. Chiou, and R. Maciejewski, "Evaluating the impact of uncertainty visualization on model reliance", *IEEE Transactions on Visualization and Computer Graphics*, to appear.