

The Opinion Tetrahedron as a Tool for Coalescing Group Beliefs

J. Michael Dunn

School of Informatics, Computing, and
Engineering

and Department of Philosophy

Indiana University-Bloomington

LORI VI, 2017

Hokkaido University

Sapporo, Japan

Prelude

"What a lot of books!" she screamed. "And have you really read them all, Monsieur Bonnard?"

"Alas! I have," I replied, "and that is just the reason that I do not know anything; for there is not a single one of those books which does not contradict some other book; so that by the time one has read them all one does not know what to think about anything. That is just my condition, Madame."

- Anatole France, *The Crime of Sylvestre Bonnard*, 1917 (translated by Lafacdio Hearn) . I owe this quote to Jon Doyle.

This talk is based on two earlier papers:

- “Contradictory Information: Too Much of a Good Thing,” *J. of Philosophical Logic* (2010), vol. 39, pp. 425-452.
- “Contradictory Information: Better than Nothing? The Paradox of the Two Firefighters,” co-authored with Nicholas Kiefer, forthcoming in *Graham Priest on Dialetheism and Paraconsistency*, eds. C. Baskent, T. Ferguson, H. Omori, Outstanding Contributions to Logic Series, Springer.

In this last paper my co-author Nick Kiefer (Statistics and Economics, Cornell) and I used two methods (a paraconsistent method (the Opinion Tetrahedron) and a probability method (Linear Opinion Pooling) as ways of coalescing the contradictory opinions of two firefighters. The title of the talk today emphasizes the first.

The main aim was to show that sometimes even a contradiction can provide useful information.

This is contrary to Luciano Floridi (2011, p. 109) who says contradictions contain zero information, and that "inconsistent information is obviously of no use to a decision maker," and also to Karl Popper (1934, 1959): "The importance of the requirement of consistency will be appreciated if one realizes that a self-contradictory system is uninformative. It is so because any conclusion we please can be derived from it."

Bar-Hillel and Carnap (1953, p. 229)) famously wrote (Floridi calls it the Bar-Hillel Carnap Paradox) "It might perhaps, at first, seem strange that a self-contradictory sentence, hence one which no ideal receiver would accept, is regarded as carrying with it the most inclusive information. ... A self-contradictory sentence asserts too much; it is too informative to be true."

Explosion!

This is not something “just postulated.” There are deep seated reasons in classical logic and in classical information theory for this.

- According to classical logic a contradiction implies every sentence whatsoever.
- According to classical information theory (Shannon), the information of a sentence is the inverse $1/n$ (to base 2) of its probability n . Since the probability of a contradiction is 0, its inverse is infinite, or more properly, undefined.


The Paradox of the Two Firefighters

Suppose you are awakened in your hotel room by a fire alarm. You open the door. You see three possible ways out: left, right, straight ahead.


- Scenario 1. You see two firefighters. One says the only safe route is to your left. The other says to your right. Contradictory information!
- Scenario 2. You find no one to give directions. Incomplete information!

Question: Which scenario would you prefer?

Scenario 1?




Or Scenario 2?




Let me repeat the question: Which scenario would you prefer? Show your hands please!

Scenario 1?



Or Scenario 2?



Obvious answer: A rational agent would prefer to be in Scenario 1. Contradictory information in this case is better than no information at all.

A More Homey Example

The essence of the example of the two firefighters can be duplicated over and over again. An instance very familiar to me goes like this: My wife Sally and I are leaving the house. I reach in my pocket and cannot find my car keys. I tell Sally I think they are in a jacket pocket in the closet. She tells me they are on the piano.

Again, this is all useful information, and I will use it in my search. But it is contradictory.

Two Solutions to the Firefighter's Paradox


It's not a paradox unless it has at least two solutions. :)

1. A Paraconsistent Solution (Opinion Tetrahedron)
2. A Probability Solution (Linear Opinion Pooling)

Paraconsistent Solution The Opinion Tetrahedron

First we present Audun Jøsang's Opinion Triangle from his 1997 "Subjective Logic."

Opinion Triangle





Note that it has just one kind of “uncertainty.” If this reminds you of the Kleene-Łukasiewicz lattice $3N$ with values T, N, F, you are not mistaken. It can be embedded into the Opinion Triangle T = complete Belief, N = complete Uncertainty, F = complete Disbelief.

But those of you who know about relevance logic, or about paraconsistent logics more generally, know that there are two kinds of uncertainty, the kind that comes about from ignorance, or the kind that come about from conflict – too little information or too much information.

Two Versions of the 4-valued De Morgan Lattice DM3

Asenjo,
Sugihara,
Priest




By combining them we get

BEST of BOTH!

This plays the same fundamental role among Monteiro's De Morgan lattices (distributive lattices with an order inverting mapping $\sim$ of period two) that the 2-element Boolean algebra plays among Boolean algebras. Białynicki-Birula and Rasiowa's studied De Morgan lattices under the name quasi-Boolean algebras.

DM4


Belnap-Dunn 4-valued Logic




Approximation (knowledge, information) Order

Two Kinds of Uncertainty:

Too little information, too much information




Add a line and visualize it as an “Opinion Tetrahedron”!




- Establish coordinates by dropping altitudes from each vertex to the center of the opposite side and by convention assign each the length 1.0 (measuring from 0 at the base to 1 at the vertex). They intersect at 0.25.
- A point (b, d, u, c) in the Opinion Tetrahedron is to be understood as follows: b = degree of belief, d = degree of disbelief, u = degree of uncertainty (ignorance), c = degree of contradiction. $0 \leq b, d, u, c \leq 1$.

4 Values as Elements of
Lattice DM4

$$T = \{t\}$$



4 Values as Apexes of
Opinion Tetrahedron



A sentence is given a value (b,d,u,c)

Before the firefighters, you might evaluate each of R, S, and L as $(1,0,0,0)$, optimistically assuming that there is no reason why any hallway would not lead to an exit stairway. Or you might more cautiously evaluate each as $(1/3,0,2/3,0)$, assuming that at least one of the hallways must lead to an exit stairway.

After the firefighters give their "pitches," if you are an optimist you disregard the conflict and focus on the fact that the two firefighters agree that straight is not an exit. So both R and L get the value $(1/2,1/2,0,0)$, but S gets the value $(0,1,0,0)$. If you are a pessimist you focus on the conflict and think that they must be incompetent and/or have flawed evidence, and you give both R and L say the value $(1/3,0,0,2/3)$.and S something like the value $(\varepsilon,1-\varepsilon,0,0)$ (where ε , varies with your degree of pessimism). But in any event, the degree of belief in S shifts downward substantially after you listen to the two firefighters.

Probability Solution

Linear Opinion Pooling

The Bayesian decision maker (you) will consider the messages from the firefighters as expressing their own beliefs about the possible escape paths, and will look for a way to combine this information with your own beliefs and come to a decision about the route. To set this up, let us cast the statements in terms of reports of probability distributions. Here, the distributions reported by the firefighters are rather trivial - the probabilities are zeros and ones - but the setting is useful. First, what is the space on which the probabilities are defined? There are 3 hallways, L, S, and R. Each can be an escape route or not, denoted by 1 or 0 respectively.

Thus there are 2^3 possibilities, $(L, S, R)=(0, 0, 0)$ to $(1, 1, 1)$; arrange these in lexicographic order and index the probabilities as $(p_1, \dots, p_8)$:

(L,	S,	R)	Probability
(0,	0,	0)	p_1
(0,	0,	1)	p_2
(0,	1,	0)	p_3
(0,	1,	1)	p_4
(1,	0,	0)	p_5
(1,	0,	1)	p_6
(1,	1,	0)	p_7
(1,	1,	1)	p_8

The cases in which hallway L works are

$(1,0,0), (1,0,1), (1,1,0), (1,1,1),$

so the probability that hallway L works is $p_L = p_5 + p_6 + p_7 + p_8$. For the decision at hand, the decision maker is only interested in the probabilities p_L, p_S, p_R , three probabilities but not itself a probability distribution. Suppose you have no information at all about the relative likelihood of the hallways (e.g. an exit sign!); then it makes sense to assign probability $1/8$ to each of the 8 possible outcomes, resulting in aggregate probabilities of $1/2$ associated with each hallway

The first firefighter is reporting $P^1=(p_L, p_S, p_R) = (1, 0, 0)$ and the second $P^2=(0, 0, 1)$. You wish to combine this information with your own initial beliefs $P^0=(1/2, 1/2, 1/2)$ to obtain posterior, updated beliefs P^* . A natural and attractive way to proceed is with a weighted average, the "linear opinion pool,"

$$P^* = w_0P^0 + w_1P^1 + w_2P^2,$$

where w_0, w_1, w_2 are the weights you assign to yourself (w_0) and the two firefighters (w_1, w_2).

Each weight w_i is to be multiplied component-wise across the triple P^i , and then the results are added component-wise to obtain P^* . The weights must add to 1, and it is plausible for you, being ignorant and fair minded, to assign the weights equally as $1/3$ each.

We do a calculation to illustrate:

$$\begin{aligned} P^* &= 1/3(1/2, 1/2, 1/2) + 1/3(1, 0, 0) + 1/3(0, 0, 1) = \\ &(1/6, 1/6, 1/6) + (1/3, 0, 0) + (0, 0, 1/3) = (3/6, 1/6, 3/6) = \\ &(1/2, 1/6, 1/2). \end{aligned}$$

The weights, like the probabilities themselves, are subjective. They reflect your assessments of the relative reliability of the 3 information sources, and can depend on impressions (is one of them delirious?) but not on the probability assessments. More on this below. You may wish to give more weight to the firefighters' assessments, thinking that they are better informed than yourself, but it is unlikely that you will give your own assessment zero weight, if only to be certain that all three hallways are included in the posterior support.

The opinion pooling literature works with the entire probability distributions (here, over the 8 possible states) rather than the probabilities of events (combinations of states) as we just did above.

Thus the prior distribution $P_f^0 = (1/8, \dots, 1/8)$, an 8 component vector, and the posterior distribution is $P_f^* =$
 $w_0(1/8, 1/8, 1/8, 1/8, 1/8, 1/8, 1/8, 1/8) + w_1(0, \mathbf{1}, 0, 0, 0, 0, 0, 0) + w_2(0, 0, 0, 0, \mathbf{1}, 0, 0, 0) =$
 $1/3(1/8, 1/8, 1/8, 1/8, 1/8, 1/8, 1/8, 1/8) + 1/3(0, 1, 0, 0, 0, 0, 0, 0) + 1/3(0, 0, 0, 0, 1, 0, 0, 0)$
 $= (1/24, 9/24, 1/24, 1/24, 9/24, 1/24, 1/24, 1/24).$

The entropy of a distribution P is $H(P) = -\sum_{i=1}^K p_i \log_2 p_i$. Note that $\log_2(1/8) = -3$, hence the entropy of the prior distribution $H(P_f^0) = 3.00$. We calculate the posterior distribution (noting that $1/24$ occurs as a component 6 times in P_f^* and that $9/24$ occurs 2 times):

$$H(P_f^*) =$$

$$\begin{aligned} & - \{[6 \times (1/24 \times \log_2(1/24))] + [2 \times (9/24 \times \log_2(9/24))]\} = \\ & - \{(6/24 \times \log_2(1/24)) + [(18/24 \times \log_2(9/24))]\} = \\ & - \{[1/4 \times \log_2(1/24)] + [3/4 \times \log_2(9/24)]\} = \\ & - \{(-4.58496 \dots / 4) + [.75 \times (-1.41504 \dots)]\} = \\ & - \{-1.145 + -1.058\} = \end{aligned}$$

$$2.207519 \dots$$

The two contradictory firefighters thus provide a clear reduction in uncertainty, from 3 (the number of hallways (L, S, R)) to approximately 2.21.

You likely noticed that our example of the two firefighters has two firefighters, and that the inconsistency between L and R is divided between them. This is why linear opinion pooling is so appropriate. And it is easy to see how the example might be extended in various ways to more than two firefighters, and how linear opinion pooling might be applied to such examples. Linear opinion pooling is based on the idea of combining various individual views, each of which is consistent but the combination might well be inconsistent.

But can linear opinion pooling be applied if there is just one firefighter? What if the firefighter is alone and you are not even there. How can just a single firefighter give contradictory information in a way that might be useful (to himself)?

Dietrich and List (2016) give a nice example related to this: "Finally, in a purely intra-personal case, an agent may seek to reconcile different 'selves' by aggregating their conflicting opinions on the safety of mountaineering, in order to decide whether to undertake a mountain hike and which equipment to buy."

They speak of two 'selves' in quotes for what is popularly called "being of two minds." The single firefighter says (talking to himself -- remember you are not even there) says something like this. "I have been up and down each of the three hallways. I am somewhat disoriented because of the fire, but I am sure that I remember that there is only one way out. Oh, and I also remember that it is the left hall. No, I also remember that it is the right hall. I am of two minds." But I will toss a between left and right.

So linear pooling can be used even when pooling the views of a single individual.

And we shall next see that the Opinion Tetrahedron can be used to coalesce the views of multiple individuals, as well as those of a single individual.

A model for quantifying uncertainty in the sense of ignorance

- It may be apocryphal, but I have heard that in the early days of the U.S. Weather Service, the chance of “rain” was determined by taking a vote of the forecasters. Suppose 100 forecasters are polled, 67 say yes, 33 say no. Then the probability of rain was reported as 0.67.
- But suppose that 58 say yes, 31 say no, and 11 hesitate to offer an opinion. We could interpret this as the degree of belief is 0.58, the degree of disbelief is 0.31, and the degree of uncertainty is 0.11.
- Both of these are particular kinds of normalized (to sum to 1.0) weightings according to number of sources.

A model for quantifying uncertainty of two different kinds: ignorance and conflict

Let us return to the forecasters where there was some degree of uncertainty. Remember that 58 said yes to rain, 31 say no, and 11 hesitated to offer an opinion. But suppose this time the forecasters could not just abstain – they could also vote both yes and no. Perhaps 7 of them had been in meetings and had no chance to study the prospects of rain, while 4 had worked very hard and had produced persuasive evidence on both sides of the question. We could interpret this as the degree of belief is 0.58, the degree of disbelief is 0.31, the degree of “uncertainty” in the sense of ignorance is 0.07, and the degree of “uncertainty” in the sense of conflict is 0.04.

It is common to include in a survey besides Yes / Agree and No /Disagree, a neutral option (Don't know / Undecided / No Opinion, whatever), or equivalently an odd number of graded responses, say 1-3.

There has been much discussion and even research about the wisdom of giving the respondent a neutral option, since it allows a easy way out and does not force the respondent to take sides (which they might have to do eventually, as in an election, or buying a product).

I can only imagine the discussion that would be provoked by adding yet one more kind of “neutral option.” But nonetheless it seems abstractly an interesting option.

Two kinds of conflict – within a single individual and between separate individuals.

It would be possible to quantify a 4-valued poll

Agree

Disagree

Ignorant

Conflicted

even more finely. Suppose we offer a scale of 1-10:

Sample Poll

Circle one number in each row. Note: the numbers circled must add up to 10.

President Trump is doing the greatest job ever.

Agree	0	1	2	3	4	5	6	7	8	9	10
Disagree	0	1	2	3	4	5	6	7	8	9	10
Ignorant	0	1	2	3	4	5	6	7	8	9	10
Conflicted	0	1	2	3	4	5	6	7	8	9	10

A good discussion in the polling literature on having two different kinds of “neutral points” is S. M. Nowlis, B. E. Kahn, and R. Dhar (2002), “Indifference versus Ambivalence: The Effect of a Neutral Point on Consumer Attitude and Preference Measurement.”

“Indifference” is not the same as “Ignorance,” nor is “Ambivalence” quite the same thing as conflict. Nowlis *et al* were talking about “opinion” polls where “opinions” have often to do with preferences, not beliefs *per se*.

WWW: Better Than Nothing?

Suppose you want to find an answer to a certain yes/no question on the WWW. Which of the following scenarios do you prefer?

- A. You google and get no (relevant) response.
- B. You google and get multiple conflicting responses.


I , and I expect you, would prefer to be in scenario B. In this circumstance you can at least try to sort the situation out.

You can count the relative number of opinions on either side.

And then you can weight them using factors such as the following:

- 1) You can somehow assess the credentials (authority/motives) of the sources on either side.
- 2) You can look at the arguments, if any, provided by the sources.
- 3) You can try to find cited “facts,” try to reproduce cited experiments. Etc.

With respect to 1) (count the opinions on each side), let's get back to our two firefighters. Now suppose that a third firefighter shows up and points right saying that this is the only safe way to go. Now I have some reason to run right.



Perhaps such investigations will sort things out so that you can use linear pinion pooling.


But in a worse case, where you do not have the mind, or time, to do that, at least you might assign a value within the Opinion Tetrahedron, if only based on your gut reaction.

It is worth emphasizing that the utility of contradictions is due not just to their content but also to their pragmatic context.

Some of the tools described on the previous slide might be taken as “logical fallacies.” E.g., counting the number of sources can be interpreted as an “argument from repetition” -- Fifty Million Frenchmen Can’t Be Wrong. Anatole France responded: “If fifty million people say a foolish thing, it is still a foolish thing.” But the aim here is not to prove P by the number of sources that say it, but rather to take a “vote” to determine the subjective likelihood that P . An improvement would be to check for duplications (one source merely repeating another source), and to have “trust” weightings based on reliability (expertise, honesty, lack of bias, etc.).

As an example of the use of duplication to manipulate political views see “Twitter bombs”: <http://truthy.indiana.edu/press>.

Postlude




Jorges Luis Borges (1941) *La biblioteca de Babel* “The Library of Babel”

The impious maintain that nonsense is normal in the Library and that the reasonable (and even humble and pure coherence) is an almost miraculous exception. They speak (I know) of the "feverish Library whose chance volumes are constantly in danger of changing into others and **affirm, negate and confuse everything** like a delirious divinity."

I repeat: it suffices that a book be possible for it to exist. Only the impossible is excluded. For example: **no book can be a ladder, although no doubt there are books which discuss and negate and demonstrate this possibility** and others whose structure corresponds to that of a ladder.

Have we discovered that Luis Borges invented the Internet? 😊



Thank you very much

どうもありがとうございます

domo arigatou gozaimasu

Some references

- Yehoshua Bar-Hillel and Rudolf Carnap (1953), "Semantic Information," *The British Journal for the Philosophy of Science*, 4, 147-157.
- Nuel D. Belnap. (1977): A Useful Four-Valued Logic, in: J.M. Dunn, G. Epstein (eds.), *Modern Uses of Multiple-Valued Logic*. Reidel, Dordrecht, 8-37.
- Franz Dietrich and Christian List (2016), "Probabilistic Opinion Pooling," in *Oxford Handbook of Probability and Philosophy*, eds. Christopher Hitchcock and Alan Hájek, Oxford University Press. Print version 2016, online version 2017.
- J. M. Dunn (1976): "Intuitive Semantics for First-Degree Entailments and Coupled Trees," *Philosophical Studies*, 29, pp. 149-168.
- Luciano Floridi (2011), *The Philosophy of Information*, Oxford University Press.
- Audun Jøsang (1997), "Artificial Reasoning with Subjective Logic," *Proceedings of the Second Australian Workshop on Commonsense Reasoning*, Perth.
- Audun Jøsang (1999), "An Algebra for Assessing Trust in Certification Chains," *Proceedings of the NDSS'99 Network and Distributed Systems Security Symposium*, The Internet Society, San Diego.

- Kalman J. Kaplan (1972), "On the Ambivalence-indifference Problem in Attitude theory and Measurement: A suggested modification of the semantic differential technique. *Psychological Bulletin*, 77, 361-372.
- Nicholas M. Kiefer (2011), "Default Estimation, Correlated Defaults and Expert Information," *Journal of Applied Econometrics*, 26, 173-192.
- Stephen M. Nowlis, Barbara E. Kahn, and Ravi Dhar (2002), "Indifference versus Ambivalence: The Effect of a Neutral Point on Consumer Attitude and Preference Measurement," *The Journal of Consumer Research*, 29,, pp. 319-334.
- Karl R. Popper (1934), *Logik Der Forschung: Zur Erkenntnistheorie Der Modernen Naturwissenschaft*, Vienna: J. Springer. English translation (1959) *The Logic of Scientific Discovery*, London: Hutchinson.
- Mervyn Stone (1961), "The Opinion Pool," *Annals of Mathematical Statistics*, 32,1339-1342.