


ARTICLE

Reasoning Without the Conjunction Closure

Alicja Kowalewska 

Carnegie Mellon University, Pittsburgh, PA, USA

Email: akowalew@andrew.cmu.edu

(Received 26 May 2020; revised 7 September 2020; accepted 3 November 2020)

Abstract

Some theories of rational belief assume that beliefs should be closed under conjunction. I motivate the rejection of the conjunction closure, and point out that the consequences of this rejection are not as severe as it is usually thought. An often raised objection is that without the conjunction closure people are unable to reason. I outline an approach in which we can – in usual cases – reason using conjunctions without accepting the closure in its whole generality. This solution is based on the notion of *confidence levels*, which can be defined using probabilities. Moreover, on this approach, reasoning has a scalable computational complexity adaptable to cognitive abilities of both rationally bounded and perfectly rational agents. I perform a simulation to assess its error rate, and compare it to reasoning with conjunction closure.

Keywords: Rational belief; conjunction closure; approximate reasoning; Lockean Thesis; Lottery Paradox; Preface Paradox

1. Introduction

One common assumption about binary beliefs is the *Conjunction Closure (CC)*. According to (CC) a rational agent is justified in believing conjunctions of her beliefs:

$$Bel(P_1) \wedge Bel(P_2) \wedge \dots \wedge Bel(P_n) \Rightarrow Bel(P_1 \wedge P_2 \wedge \dots \wedge P_n). \quad (CC)$$

Intuitively, it seems that believing a conjunction is equivalent to believing all of its conjuncts. If you believe *A* and you believe *B*, it would be natural to assume that you also believe *A* and *B*. It is also often suggested that without (CC) we wouldn't be able to conduct even basic reasoning. Imagine a dialogue between friends:

“Would you like to go for a walk with me?”

“If I have a free afternoon and the weather is good, we can go.”

“Do you have a free afternoon tomorrow?”

“Yes, I believe so.”

“And do you think that the weather will be good tomorrow?”

“I believe it will be quite good.”

“Ok, so let's meet tomorrow!”

“Wait, why tomorrow? I haven't said that tomorrow suits me.”

“Well, you've said that if you have a free afternoon and the weather is good ...”

“Right!”

“And that you believe that tomorrow you will have a free afternoon and the weather will be good ...”

“No, I haven’t said that! I believe that I will have a free afternoon and I believe that the weather will be good, but not that I will have a free afternoon **and** the weather will be good.”

The results of denying (CC) seem quite odd. That’s why usually philosophers accept (CC) and incorporate it within their theories of rational belief (see Pollock 1986; Evinne 1999; Adler 2004; Smith 2016; Leitgeb 2017).

However, there are also some problems related to embracing (CC) such as the well-known issue of risk accumulation, and so, there is also a group of people who reject (CC) (see Foley 1992; Christensen 2004; Hawthorne 2009).

The most popular objection to the views which reject (CC) concerns deductive reasoning. Supposedly, without (CC) it is impossible (or, on some approaches, overly difficult) to conduct such reasoning.

The objective of the paper is to investigate the rationale for rejecting (CC), and to defend the view that rejects (CC) by providing a framework which allows us to easily reason with conjunctions without a real acceptance of (CC).

First, in section 2, I introduce a widely known version of the difficulty about conjunction, I describe some of the available solutions to the problem that manage to keep (CC), and I point out their weaknesses. Then, in section 3 I present independent reasons for the rejection of (CC). In section 4, I respond to the most common worries about this view, and I consider some difficulties related to current (CC)-free theories. Finally, in section 5 I outline an approach that is able to handle the worries about the feasibility of reasoning without (CC).

2. The difficulty

In contemporary epistemology we can distinguish at least two notions of belief: we may talk about *binary* (i.e. *all-or-nothing*) *belief*, and about *credences* (i.e. *graded belief*, *degrees of belief*, or *subjective probability*). The relationship between these notions is currently a widely debated issue.

A popular bridge principle linking the two concepts, introduced by Foley (1992), is the *Lockean Thesis* (LT). It can be formulated in terms of *necessary* and *sufficient* requirements for rational binary belief. (*LT-nec*) states that believing a proposition *P* is justified only if the agent’s subjective probability assigned to *P* is above a certain threshold *t*:

$$Bel(P) \Rightarrow Pr(P) > t. \quad (LT-nec)$$

(*LT-suf*) states that if an agent’s subjective probability assigned to a proposition is above the threshold, the agent is justified to believe it:

$$Pr(P) > t \Rightarrow Bel(P). \quad (LT-suf)$$

A weak version of LT, (*LT-nec*) with 0.5 threshold, is especially popular among epistemologists:¹

$$Bel(P) \Rightarrow Pr(P) > 0.5. \quad (LT-nec-0.5)$$

¹There are some theories of rational belief that don’t embrace this assumption, but this counts as their controversial aspect. For instance Normic Support Theory (Smith 2016) or a view on which to believe a proposition is for the proposition to be part of the overall theory that one gives the highest credence to (compared with competitors).

It seems so plausible, because what it says is simply that we're not permitted to believe a proposition whenever we think that its negation is more or equally probable. Denying (*LT-nec-0.5*) would require some principled explanation. After all, why would anyone believe P , if $\neg P$ seems more (or equally) likely? It still may be the case that, in spite of a higher probability, an agent lacks sufficient justification for believing $\neg P$ and can opt to suspend judgment about P , but it is rather clear that in such a situation she shouldn't believe P .²

There is one more widely accepted requirement for rational binary belief – the *No Contradiction Rule (NCR)*:

$$\neg(\text{Bel}(P) \wedge \text{Bel}(\neg P)). \quad (\text{NCR})$$

It states that believing both a proposition and its negation can never be rational. So far (*NCR*) hasn't raised too many objections.³

Even though they seem quite plausible taken separately, *Conjunction Closure*, *Lockean Thesis*, and *No Contradiction Rule* jointly lead to paradoxes such as the *Lottery* and the *Preface*.

The *Lottery Paradox*, first introduced by Kyburg (1961), goes as follows: because of high objective chances and (*LT-suf*) one believes about any lottery ticket that it will lose. By (*CC*) one believes that all tickets will lose. At the same time, one knows that a lottery is fair and some ticket will win. This set of beliefs violates (*NCR*).

A similar paradox arises when an author rationally believes in every claim she has written in the book. By (*CC*) she believes all the claims are true. But she also knows that everyone is fallible, and a mistake is highly probable. Therefore, by (*LT-suf*), she believes that some claim must be false. Her beliefs violate (*NCR*). Authors often apologize for mistakes in the preface, so the problem is usually called the *Preface Paradox* (see Makinson 1965).

These two well-known paradoxes illustrate a general schema that leads to problematic results. It seems there are cases when we are justified in believing a great number of certain propositions:

$$\text{Bel}(P_1) \wedge \text{Bel}(P_2) \wedge \dots \wedge \text{Bel}(P_n), \quad (\text{propositions})$$

and at the same time we can rationally believe that not all of them are true:

$$\text{Bel}(\neg \wedge P_i). \quad (\text{neg-conjunction})$$

However, (*propositions*) and (*neg-conjunction*) are incompatible with (*CC*) and (*NCR*). At least one of these four claims must be given up. But which one, and why?

There has been numerous attempts to provide a solution to the problem with conjunction, and full discussion of this topic goes beyond the scope of the paper. Some of the most prominent options include:

- **Rejecting the Lockean Thesis:** without (*LT-suf*), in the case of the *Lottery Paradox*, (*propositions*) is no longer justified. In the case of the *Preface Paradox* (*neg-conjunction*) becomes unsupported. This way both problems seem to be resolved. However, as Worsnip (2016) points out, successful application of this strategy

²Note we're talking here about *subjective* not *objective* probabilities, and we take them to be equivalent to credences.

³For the purpose of this paper I consciously put aside the issue of dialetheism.

also requires denying the weakened version of LT – (*LT-nec-0.5*),⁴ and thus allows to justifiably believe P instead of $\neg P$ while one has higher credence in $\neg P$ than P .

- **Context-based solutions:** these solutions suggest that while we can believe in each conjunct when thinking about them separately, we should stop believing conjuncts when we consider them together. This way (*propositions*) is unjustified and the problems seem to evaporate. One objection to this approach is that it may lead us to change our single belief without a proper reason, i.e. when we consider the rest of our beliefs, even though they might not be connected in any way (see Staffel 2016; Schurz 2019).

3. Rejecting the closure

A simple way out of these problems would be to just reject (*CC*). Let's see why this approach can be more plausible than many think.

The source of the difficulties with conjunction lies in not requiring certainty for belief. Intuitively, we agree that one can rationally believe a proposition even if the ascribed subjective probability is below 1. But when we consider the conjunction of these propositions, the risk aggregates, and the probability decreases. The conjunction of many rational beliefs will usually have an extremely low probability. For instance, the probability of a conjunction composed of 50 probabilistically independent conjuncts with probability 0.95 is only 0.077. (With the number of conjuncts approaching infinity, the probability converges to 0.) However, if we accept (*CC*) – and, at the same time, don't employ any other solution to this problem – we are forced to admit that believing an improbable conjunction is justified, which seems to conflict with our intuitions. It also imposes that we need to reject (*LT-nec-0.5*).

In fact, the risk aggregation shouldn't really surprise us. The lower we set the probability threshold for beliefs, the more mistakes we allow. When we agree on small uncertainty, and this permission is used very often, we should simply expect an error. Even a small mistake repeated many times can give miserable results.

A famous example of that is a failure of the Patriot system which was protecting the U.S. Army in Dhahran, Saudi Arabia, in 1991.⁵ The time was stored in a 24-bit register and was incremented every 0.1 s. Due to binary representation of 0.1 each incrementation added an error of around 0.000000095 s. After 100 h the error added up to 0.34 s; given the speed of R-17, that caused 0.5 km location inaccuracy in Patriot. The system didn't detect a Scud rocket which resulted in 28 killed and 100 wounded.

This mistake shows that when values contain even a very small error, we shouldn't expect all mathematical equations that use them, such as $0.100000095 + 0.100000095 = 0.2$, to be correct. Similarly, if we allow imperfections in rational belief, we cannot expect that classical logic will hold universally in our belief systems. This is simply the cost of not requiring certainty.

4. Reply to common worries

4.1. Incoherent set of beliefs

The primary argument against rejecting (*CC*) is that we no longer can prohibit seemingly incoherent belief sets. For instance, one's set of beliefs could include

⁴This solution to the *Preface Paradox* judges (*propositions*), i.e. the belief that the book is error-free, as justified. But the probability of this claim is low, it seems to be below 0.5. Therefore, this approach is incompatible with (*LT-nec*) with 0.5 threshold.

⁵<https://hownot2code.com/2016/11/09/r-17-vs-patriot-a-rounding-issue-bugs-in-a-missile-defense-system/>

$\{A, B, \neg(A \wedge B)\}$,⁶ while there can't be any possible world where A , B and $\neg(A \wedge B)$ hold simultaneously.

This result sounds quite suspicious, but is it really so unacceptable? Well, it depends on the assumptions. This seems to be a problem if you insist that the full conjunction of your beliefs is plausible, and you are justified in believing it. On the contrary, if you assume that some of your beliefs are wrong (due to risk accumulation), then it doesn't make sense to always expect a possible world in which all of your beliefs would be true. It should be rather clear that the argumentation for rejecting (CC) presented so far is also the argumentation for the latter assumption.

4.2. Possibility of reasoning

Another worry one may have is that giving up the use of classical logic is the price we cannot afford. A typical objection to the idea of rejecting (CC) is that without it reasoning would be impossible. An agent could believe A , B , and that $A \wedge B$ implies C :

$$Bel(A), Bel(B), Bel(A \wedge B \Rightarrow C) \quad (1)$$

and, at the same time, rationally refuse to believe C :

$$\neg Bel(C). \quad (2)$$

Foley (1992) argues that we can bear the consequence of not being able to reason with beliefs since there are other concepts, such as acceptance, that can be used to reason. Nonetheless, this account raises a worry that the role of beliefs would become very limited. Moreover, it isn't clear that the other concepts used for reasoning wouldn't fall prey to analogous difficulties. It seems that this argumentative strategy only redirects the problems with beliefs to different concepts.

However, it is important to note that rejecting (CC) doesn't mean that we should *never* believe any conjunctions of our beliefs. Getting rid of (CC) we only admit that it's not *always* the case that we should believe conjunctions of our multiple beliefs. In fact, it might be true that in everyday life we *usually* do justifiedly believe in conjunctions of some of our beliefs.

If the conjunction is well justified and sufficiently probable then, of course, an agent should believe it. By rejecting (CC), we just don't enforce a belief in a complex, improbable conjunction based on uncertain beliefs in all its conjuncts. For instance, the probability of 10 independent propositions, each with probability 0.99, is around 0.9, and it might be still plausible to believe it. But the probability of the conjunction of 50 independent propositions, each with probability of 0.9, is around 0.005, so being permitted to believe this conjunction doesn't seem desirable. With this in mind, it is clear that even without (CC), some reasoning is still possible. Furthermore, if we take into account that often the conjunctions that people use to reason are not very complex, the problem ceases to disturb us, because usually the probabilities of conjunctions will be sufficiently high for justification.

4.3. Calculating probabilities

Now, some may still argue that even if *theoretically possible*, reasoning without (CC) is much more difficult, and it seems *unfeasible* for people. Before believing any conjunction, each time we would be expected to calculate its numerical probability, preferably

⁶Note that on most approaches that reject (CC) this set of beliefs is rarely rational. For instance on a view that endorses *LT-nec* with a sufficiently high threshold, such as 0.8, it can never be justified.

with infinite accuracy. For example, Christensen (2004) develops a view in which binary beliefs are heavily dependent on the corresponding degrees of belief, but despite this his approach doesn't necessarily require a full rejection of the all-or-nothing concept. However, one can still argue that some objections to the reductionist views apply here *mutatis mutandis* – to use binary beliefs we would have to calculate the corresponding probabilities, which seems to be an overly demanding task.

A natural response is to remember that the philosophical discussion about beliefs is concerned with a normative model, which is supposed to serve as an ideal. Therefore, we shouldn't be particularly worried that humans aren't able to reach perfect rationality. At this point we could engage in a lengthy discussion about the concept of rationality, and whether *ought* should imply *can*, but let's put those issues aside, because the opponents may present us with, in my opinion, a more difficult objection anyway.

They can insist that a theory which requires people to calculate probabilities, simply doesn't match how we really gain beliefs. The problem isn't just that we aren't able to do it perfectly, but rather that normally we don't do it at all. Thus, it doesn't make sense to think that people even try to imitate this kind of perfect reasoning. In other words, the worry is that this theory only applies to idealized agents, and cannot be easily modified to become a guide that people use on a daily basis, even imperfectly, trying to meet the requirements of rationality.

One way out would be to develop an additional account of *limited rationality* better suited to describe human reasoning. It appears that the notion of binary belief is mostly useful in a practical context. We often need it for personal communication, making decisions, or simplifying our reasoning. At the same time, in everyday life, extreme accuracy doesn't seem to be a primary concern, so we might be willing to allow higher rates of mistakes to be able to make reasoning much easier and faster. (Even if we care about accuracy, we'll have to agree on some imprecision anyway to simply make the reasoning feasible.)

(CC) is clearly an example of a rule which really simplifies a reasoning process. It is also possible that normally the probability of a simple conjunction doesn't drop much lower than the probability of its conjuncts. In such cases, if we don't require very high accuracy, results given by (CC) can be acceptable. In fact, this might be the reason why we use it so often, and why it seems so intuitive.

This leads us to a dual approach: we would require calculations of probability to achieve *perfect* rationality, and we would allow the use of (CC) in cases when one only aims at *limited* rationality. But this strategy also has its drawbacks. It seems quite extreme – there is nothing between *perfect* and *limited* rationality. While people might not be able to reason perfectly, we would like to do it as well as we can. Alas, we recognize that some results given by (CC) (namely that we should believe highly improbable conjunctions) aren't reasonable. So it appears that the application of this rule doesn't take full advantage of human capabilities, and therefore it's doubtful that it can serve as a guide to rationality.

It seems that another theory more capable of filling the gap between *perfect* and *feasible* is needed. In the next section I will briefly describe and discuss Hawthorne's (2009) idea which allows us to describe perfect rationality but doesn't require the use of numerical probabilities.

4.4. Comparing confidence

Hawthorne develops a (CC)-free logic of belief based on relative probabilities.⁷ The underlying idea is that instead of assigning specific values of probabilities, we can

⁷For further discussion about relative probabilities see Krantz *et al.* (1971), Fine (1973) and Halpern (2003).

only compare our confidence in different propositions. Then, if we already believe in a proposition which we are less confident about, we should also believe in a proposition under consideration. He offers a full formalization of this system, and argues that it can be also interpreted by numerical probability measures with a belief threshold. Whether we are supposed to accept a conjunction depends on its plausibility as compared with our other beliefs.

At first, this approach seems really promising, because one no longer needs to calculate the probability of each conjunction, which makes the reasoning easier. But is it really so easy? Are people able to simply compare their confidence in two complex conjunctions? If the difference in the confidence between those conjunctions is significant, this might be quite easy, but otherwise doing it in one step doesn't seem feasible. Intuitively, we are often unable to immediately compare our credences in more complex conjunctions. For instance, to evaluate $A: A_1 \wedge A_2 \wedge \dots \wedge A_{10}$ in relation to $B: B_1 \wedge B_2 \wedge \dots \wedge B_{10}$, most probably we would first estimate separately our confidence in A and our confidence in B , to be able to later easily compare them. But this suggests that we may still need to *calculate* the probabilities of conjunctions despite using the comparative framework. Therefore, real benefits of this strategy are not so clear.

Imagine the following situation: you got a big jar of cents as a payment, and you need to decide whether there is enough money. A strategy similar to *counting numerical probability* would suggest to simply count all the cents. A strategy analogous to *the confidence comparison* would advise you to compare a jar to another which was already accepted. If you were lucky and the previously accepted jar was definitely smaller, you could easily accept the new one. But if the sizes of accepted and unaccepted jars were quite comparable to this one, you still wouldn't know what to do. Assuming that you really need to get enough money, you would have to count the cents as well.

An alternative approach, which I'll describe in a moment, would suggest to simply weigh the jar, count the average weight of a coin, and approximate the amount of money. Depending on the accuracy of the chosen weighting machine, the results would be more or less accurate. However, if the weighing machine had infinite accuracy (and assuming all the coins had identical weight), there wouldn't be any error.

5. Levels of confidence

5.1. Introducing confidence levels

Let's get into details of this alternative. This approach draws from the strategy with numerical probabilities, but instead of presenting a confidence as a *real number*, it requires to assign a proposition a corresponding *level of confidence*. So, the main difference is that we no longer treat a credence as a continuous variable, but rather take the distinguished levels as a conceptually discrete measure of an agent's confidence.⁸ Depending on the situation, abilities or goals, one can choose the number of levels that should be distinguished. Confidence levels may be described in many different ways; for instance an ordered scale with 7 levels (see Figure 1), can be represented as:

- integer values from 1 to 7,
- letters from A to G,
- or descriptively: *certain, almost certain, quite confident, no clue, quite confident that not, almost certain that not, certain that not.*

⁸If for some reason, one has a strong intuition that degrees of belief are indeed continuous, one can still benefit from this approach using it as an approximation method.




Fig. 1. A scale with 7 confidence levels.

Each level corresponds to some probability – these probabilities can be chosen in a specific way to form an even scale, for instance: 1, 0.9, 0.7, 0.5, 0.3, 0.1, 0. In what follows I will preserve values 0 and 1 as separate categories which do not get rounded to other values and to which no other values are rounded. This is because we want to be able to account for logical truths and falsehoods. Other intervals will be taken to be of equal size.

On the one hand, we can describe the relationship between binary beliefs and confidence levels using the equivalents of *(LT-nec)* and *(LT-suf)* for this approach. To rationally believe a proposition *P*, an agent must assign *P* a level of confidence *L(P)* higher than the threshold level *l_t* that we chose for binary belief:

$$Bel(P) \Rightarrow L(P) > l_t. \tag{LT-nec-lev}$$

Similarly, if one assigns *P* a level of confidence *L(P)* higher than the threshold level *l_t*, then one should believe *P*.⁹

$$L(P) > l_t \Rightarrow Bel(P). \tag{LT-suf-lev}$$

On the other hand, we can define how credences should be translated to confidence levels. If one has already assigned a more precise probability value to a proposition than the chosen levels scale can capture, one needs to round this value to the nearest non-extreme probability connected with a level. Formally, we will write:

$$L_p(Credence) = lp. \tag{Lev-prob}$$

For instance, in 7-levels scale, if a credence in a proposition is 0.97 or 0.85, it should be rounded to probability 0.9, which represents level *almost certain*: $L_p(0.97) = 0.9$, $L_p(0.85) = 0.9$.¹⁰

5.2. Simple rules instead of calculations

Now, some may wonder how this approach resolves the difficulty of determining probabilities of conjunctions (necessary for accessing which conjunctions should be

⁹If you think there are some additional requirements for a belief, and therefore you reject *(LT-suf)*, you may reject *(LT-suf-lev)* as well. In what follows, whenever I claim that one *should believe* a conjunction, you may replace that with *can believe* or *should believe, conditional on the additional requirement(s)*.

¹⁰A reader may insist that in some contexts the agent’s uncertainty is expressed in terms of a set of probability measures. Such a move is debatable: there are issues with scoring rules and decision theory for imprecise probabilities. At least until such tools for imprecise probabilities are in better shape, I would suggest reasoning with precise values obtained by taking averages. Perhaps, one could use weighted averages, for instance, if one wants to model risk-aversion or higher-order uncertainty. Then, one can follow the steps described above. Thanks to the anonymous reviewer for pointing out these issues.

believed). Aren't precise calculations still needed? We could use the probabilities assigned to the confidence levels, count the probability of a conjunction, round it to the nearest level (strictly speaking, its probability) and check whether it exceeds the threshold. But this would still require a similar amount of calculations.

It turns out that all those calculations are usually not needed. To make reasoning easier and faster it's enough to remember only a few rules. Consider the 7-levels scale described above with a threshold level for (*LT-nec-lev*) set to *no clue* (i.e. 0.5). To make sure that our confidence in a conjunction C (consisting of probabilistically independent propositions)¹¹ exceeds the threshold for binary belief, we only need to check that one of the following conditions holds:

1. C consists of max. 1 conjunct that has level *quite confident*, max. 1 conjunct that has level *almost certain*, and the rest of them have level *certain*,
2. C consists of max. 4 conjuncts that have level *almost certain*, and the rest of them have level *certain*.

Given an n -levels scale, for each belief threshold l_t there is a set of analogous conditions. A condition represents a combination of conjuncts, i.e. how many conjuncts are assigned to each level. These combinations must ensure that the confidence in the whole conjunction exceeds the threshold. To reduce the number of conditions, the combination is also required to be maximal – no more uncertain conjuncts could be added without the confidence in the conjunction falling below the threshold. For instance, ∞ premises at $lp_1 = 1$ and 4 premises at $lp_2 = 0.9$ constitute a maximal combination which results in a conjunction above the threshold 0.5. 5 premises at $lp_1 = 1$ and 3 premises at $lp_2 = 0.9$ also result in a conjunction above 0.5, however this combination is not maximal.

Generating and remembering only maximal combinations helps to reduce the memory complexity, but at the same time is enough to be able to check if any combination in question will also result in a conjunction above the threshold. One simply needs to make sure that a combination in question is smaller than a maximal one.

More formally, given an n -levels scale with confidence levels l_1, l_2, \dots, l_n and corresponding probabilities lp_1, lp_2, \dots, lp_n , each condition describes the numbers x_1, x_2, \dots, x_n of conjuncts which are assigned confidence levels 1 to n . To account for the fact that the confidence in the conjunction C exceeds the threshold l_t , i.e:

$$L(C) > l_t, \quad (3)$$

$$lp_C > lp_t, \quad (4)$$

the values x_1, x_2, \dots, x_n must meet the following constraint :

$$Lp(lp_1^{x_1} \times lp_2^{x_2} \times \dots \times lp_n^{x_n}) > lp_t. \quad (5)$$

¹¹In the case of a conjunction made of probabilistically dependent propositions A_1, A_2, \dots, A_n , one needs to assess the confidence level assigned to each conjunct using a similar procedure to the chain rule for conditional probabilities: $L'(A_1) = L(A_1)$, $L'(A_2) = L(A_2|A_1)$, $L'(A_3) = L(A_3|A_2 \wedge A_1)$, and so on. These conditional levels L' , in what follows, should simply be treated in our calculations the way the levels of confidence in given conjuncts were treated before. Due to the fact that $L'(A_i)$ is an approximate value, and determining it can include rounding, the final result may depend on the ordering of conjuncts. However, the more confidence levels one uses, the smaller this dependence and possible differences. Therefore, this simply seems to be yet another source of error resulting from approximations.

Note that the conjuncts are assumed to be probabilistically independent (footnote 10 describes required changes for probabilistically dependent conjuncts), therefore simple exponentiation and multiplication of conjuncts' lp is enough to determine lp_C . Because we also assume that each level corresponds to an interval of equal size, and to a probability that lays in the middle of the interval, we can rewrite (5):

$$lp_1^{x_1} \times lp_2^{x_2} \times \dots \times lp_n^{x_n} > \frac{1}{2}(lp_t + lp_{t-1}). \tag{6}$$

As was already mentioned, the combination of conjuncts also needs to be maximal. So, if one more conjunct that has level l_2 was added,¹² (6) shouldn't hold:

$$lp_1^{x_1} \times lp_2^{x_2+1} \times \dots \times lp_n^{x_n} \leq \frac{1}{2}(lp_t + lp_{t-1}) \tag{7}$$

$$lp_1^{x_1} \times lp_2^{x_2} \times \dots \times lp_n^{x_n} \leq \frac{lp_t + lp_{t-1}}{2 \times lp_2}. \tag{8}$$

For instance, in the case of the already described 7-levels scale (where $lp_1 = 1, lp_2 = 0.9, lp_3 = 0.7, lp_4 = 0.5, lp_5 = 0.3, lp_6 = 0.1, lp_7 = 0$) with the threshold level set to *no clue*, i.e. $lp_t = 0.5$, we obtain two constraints:

$$1^{x_1} 0.9^{x_2} 0.7^{x_3} 0.5^{x_4} 0.3^{x_5} 0.1^{x_6} 0^{x_7} > 0.6 \tag{9}$$

$$1^{x_1} 0.9^{x_2} 0.7^{x_3} 0.5^{x_4} 0.3^{x_5} 0.1^{x_6} 0^{x_7} \leq 0.6 \tag{10}$$

Now, it is fairly easy to generate all possible solutions that meet (9) and (10):

	x_1	x_2	x_3	x_4	x_5	x_6	x_7
1.	∞	1	1	0	0	0	0
2.	∞	4	0	0	0	0	0

The solutions presented in the table correspond to the conditions 1. and 2. mentioned above. Because $1^\infty 0.9^1 0.7^1 > 0.6$, we know that the conjunction which exceeds the 0.5 threshold, can consist of ∞ conjuncts at l_1 (*certain*), 1 conjunct at l_2 (*almost certain*), and 1 conjunct at l_3 (*quite confident*). Analogously $1^\infty 0.9^4 > 0.6$, so the conjunction can also consist of ∞ conjuncts at l_1 (*certain*), and 4 conjuncts at l_2 (*almost certain*).

A similar procedure can be applied for different thresholds and scales. For example, for $l_t = 0.5$ and 11-levels scale there are 6 analogous conditions, and for 15-levels scale 18 conditions.

The low number of conditions and the fact that each condition describes a certain pattern, make it plausible that people are able to remember these patterns, especially the most common ones, and then easily recognize them.

¹² l_2 is the highest uncertain level.

By checking these conditions we may easily verify whether our confidence in the conjunction exceeds the threshold, and therefore whether it should be believed. But if you want to determine not only the binary attitude towards a conjunction, but also its confidence level, you can use a similar method. All you have to do is find the highest threshold l_t for which at least one condition still holds.¹³ This way you make sure that the confidence in the conjunction exceeds level l_t , and therefore you should ascribe the conjunction level l_{t+1} .

5.3. Performance of the level-based approach

One can ask how much we lose by using only a few levels of confidence instead of the infinite spectrum of real numbers, i.e. how much the sets of believed conjunctions will differ on these two approaches. This can be studied through simulations. To give the reader the gist of the procedure, I briefly explain key steps.

- Pick parameters: a probability corresponding to a necessary threshold for binary belief t , say 0.5, and a maximal number of conjuncts for simulated conjunctions, say 4.
- Sample, with uniform distribution, 1 000 000 numbers above the threshold – here, from [0.5, 1]. These, intuitively, are the precise probabilities of some (probabilistically independent) believed propositions.
- Randomly pick 10 000 sequences of length 2–4 from the sample. These correspond to conjunctions of believed propositions.
- Calculate the real probability of the conjunction by multiplying probabilities of conjuncts. For example, if a conjunction consists of two conjuncts whose probabilities are p_1 and p_2 the real probability of a conjunction equals p_1p_2 .
- Calculate the estimate confidence level of the conjunction. Round the conjuncts to their nearest level probabilities $Lp(p_1)$, $Lp(p_2)$. Multiply and round again: $Lp(Lp(p_1)Lp(p_2))$.¹⁴
- Check whether the conjunction evaluated according to these two methods exceeds the threshold t . Test the outcome according to the following table:

$p_1p_2 > 0.5$	$Lp(Lp(p_1)Lp(p_2)) > Lp(0.5)$	Result
Yes	Yes	True positive (tp)
Yes	No	False negative (fn)
No	Yes	False positive (fp)
No	No	True negative (tn)

- Collect and graph the results. You may also play around with parameters.

Figure 2 illustrates the different performance metrics for beliefs in conjunctions relative to the number of confidence levels. Tested conjunctions consisted from 2 to 4

¹³To do it even more efficiently, one can take advantage of a binary search algorithm.

¹⁴As already mentioned, in real life people might remember simple rules which determine whether the conjunction exceeds the threshold (e.g. whether $Lp(Lp(p_1)Lp(p_2)) > Lp(0.5)$), without performing this calculation directly each time.

conjuncts, and the threshold level necessary for belief was set to 0.5.¹⁵ Accuracy reports the ratio of joint correct beliefs (in conjunctions) and correct lack thereof. Precision indicates the ratio of true beliefs among your beliefs. Recall tells you how often you believe the things you should. F-score is the harmonic mean of precision and recall.

accuracy	precision	recall	f-score
$\frac{tp + tn}{all}$	$\frac{tp}{tp + fp}$	$\frac{tp}{tp + fn}$	$\frac{2 \times precision \times recall}{precision + recall}$

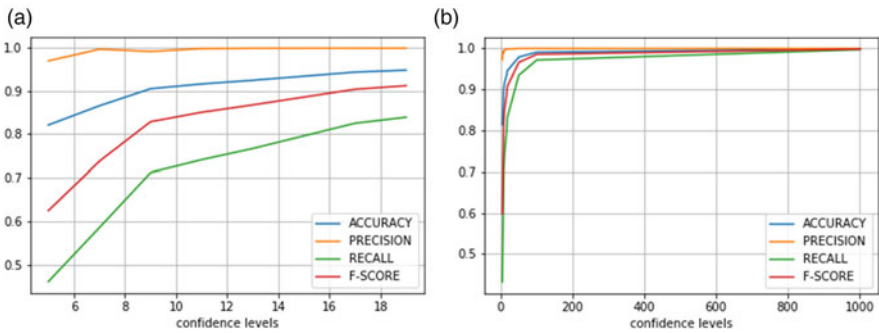


Fig. 2. Performance metrics for level-based beliefs in conjunction. (a) Low number of levels; (b) High number of levels.

With the growing number of levels, accuracy becomes quite high. For 9 levels it is around 90%, for 51 levels it is 98%, for 1001 levels it is 99.9%, and as the number of levels tends to infinity, it tends to 100%. Therefore, this approach allows us to account for both *perfect* and *typical human* reasoning. A version of this account with 5 confidence levels is similar to a slightly enhanced theory of binary beliefs (certainly false, considered false, undecided, considered true, certainly true), and the version with in the limit to the theory of graded beliefs.

Compare this to the performance of using conjunction closure. In this context, it would tell us to believe in all the conjunctions under consideration (because we formed them from propositions above the threshold). There are no negatives, recall is simply 1, precision will equal accuracy and they depend on the number of conjuncts. If we look at conjunctions built from 2–4 conjuncts, both will be around 32% (see Table 1). When we restrict our attention to conjunctions of only two conjunct they will be around 60%.¹⁶

One could also ask how the level-based approach handles single beliefs. Sometimes it happens that the probability assigned to a proposition is above the (*LT-nec*) threshold but because we round it down to the closest level it doesn't satisfy (*LT-nec-lev*). This produces some false negatives. For instance, with 7-levels scale and the threshold set to 0.5 (*no clue*), beliefs with probabilities assigned from the [0.5, 0.6] interval get rounded down to the level *no clue*, and therefore don't exceed the threshold level.

¹⁵For other parameters the results will differ; higher thresholds require more accuracy levels to achieve a similar accuracy score, while simulation of longer conjunctions decreases the accuracy. The simulation code is available here: <https://github.com/alako/Conjunction-simulation>.

¹⁶This is somewhat sensitive to the choice of threshold as well. With higher thresholds, accuracy goes down. With threshold at 0.95 it reaches ca. 22% for 2–4 conjuncts, and ca. 50% for two conjuncts.




Fig. 3. Performance metrics for level-based single beliefs. (a) Low number of levels; (b) High number of levels.

Table 1. The comparison of metrics for conjunctions with 2–4 conjuncts.

Method	Accuracy	Precision	Recall	f-score
(CC)	0.32	0.32	1	0.48
7 levels	0.87	0.98	0.59	0.74
15 levels	0.93	0.997	0.77	0.87

However, the more levels we distinguish, intervals assigned to each level get smaller. Figure 3 shows the metrics for a single belief. As there are no false positives and true negatives accuracy equals recall.

6. Results and summary

Let’s get back to the initial issues – *Lottery* and *Preface Paradoxes*. How does the current approach handle them? In the case of *Lottery Paradox*, each ticket has a high probability of losing. No matter whether we accept (*LT-suf*), the belief in the conjunction, that is the claim that all the tickets will lose, won’t be justified, because its probability is too low to satisfy (*LT-nec-0.5*). The same result will be obtained whether we use the approach with precise numerical probabilities or a few levels of confidence. Thus, avoiding belief in the conjunction, we avoid the paradox. Yet, at the same time, we preserve the reliability of many applications of conjunction closure.

The approach handles *Preface Paradox* in an analogous manner. Even if the author believes in all the claims separately, she is not committed to the belief that her book is error-free, because the probability of this claim doesn’t exceed the 0.5 threshold. Again, this is achieved without abandoning the usual reliability of conjunction closure.

Now, let’s come back to the dialogue presented at the beginning of the paper. Does it still seem so odd in the light of our considerations?

The initial dialogue didn’t include information about the confidence in the believed propositions. However, as I argued, beliefs violate conjunction closure only if the risk accumulates and there is enough uncertainty about the conjunction. In cases with two conjuncts this happens rarely – only when confidences in the conjuncts are quite low. Therefore, a more realistic conversation could go as follows:

- “Would you like to go for a walk with me?”
- “If I have a free afternoon and the weather is good, we can go.”
- “Do you have a free afternoon tomorrow?”

“I think so, but I’m not sure. I need to visit a doctor this week, I have to check if she is available on other days.”

“And do you think that the weather will be good tomorrow?”

“I believe it will be quite good. Though yesterday I also expected nice weather and it was raining cats and dogs.”

“Ok, so I understand you aren’t convinced about tomorrow. We can meet another day as well.”

“All right then. Let’s be in touch!”

We started with looking at *Lottery* and *Preface Paradoxes*, which need to be addressed if one wants to hold (CC). I argued that rejecting (CC) is a sensible approach after all. I defended this strategy by offering a view that is able to explain the usual reliability of conjunction closure without its real acceptance. One advantage of the approach based on confidence levels is that it allows to model both normative reasoning of an idealized agent and simplified human reasoning in the same manner. Furthermore, even this simplified reasoning process (with small number of levels) achieves much higher accuracy than the standard approach with (CC).¹⁷

References

- Adler J.E. (2004). ‘Belief’s Own Ethics.’ *Erkenntnis* 61(1), 123–42.
- Christensen D. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
- Evnine S.J. (1999). ‘Believing Conjunctions.’ *Synthese* 118(2), 201–27.
- Fine T.L. (1973). *Theories of Probability*. London: Academic Press.
- Foley R. (1992). ‘The Epistemology of Belief and the Epistemology of Degrees of Belief.’ *American Philosophical Quarterly* 29(2), 111–24.
- Halpern J.Y. (2003). *Reasoning About Uncertainty*. Cambridge, MA: MIT Press.
- Hawthorne J. (2009). ‘The Lockean Thesis and the Logic of Belief.’ In F. Huber and C. Schmidt-Petri (eds), *Degrees of Belief*, pp. 49–74. Dordrecht: Springer.
- Krantz D., Luce D., Suppes P. and Tversky A. (1971). *Foundations of Measurement, Vol. I: Additive and Polynomial Representations*. New York, NY: Academic Press.
- Kyburg H.E. (1961). ‘Probability and the Logic of Rational Belief.’ *Journal of Symbolic Logic* 35(1), 127.
- Leitgeb H. (2017). *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford: Oxford University Press.
- Makinson D.C. (1965). ‘The Paradox of the Preface.’ *Analysis* 25(6), 205–7.
- Pollock J.L. (1986). ‘The Paradox of the Preface.’ *Philosophy of Science* 53(2), 246–58.
- Schurz G. (2019). ‘Impossibility Results for Rational Belief.’ *Noûs* 53(1), 134–59.
- Smith M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Staffel J. (2016). ‘Beliefs, Buses and Lotteries: Why Rational Belief can be Stably High Credence.’ *Philosophical Studies* 173(7), 1721–34.
- Worsnip A. (2016). ‘Belief, Credence, and the Preface Paradox.’ *Australasian Journal of Philosophy* 94(3), 549–62.

Alicja Kowalewska pursues her PhD in Philosophy at Carnegie Mellon University. She is also a member of LoPSE research group at the University of Gdansk. Her interests focus on formal epistemology and decision theory. She is curious about the applications of philosophical considerations, especially in legal and IT contexts. Previously she studied computer science at Gdansk University of Technology and worked in the field of machine learning (NLP).

¹⁷For many valuable discussions and comments on earlier drafts of this paper, I am very grateful to Rafal Urbaniak. I would also like to thank Marcello Di Bello, Mattias Skipper, members of the LoPSE research group at the University of Gdansk, and anonymous referees for helpful suggestions.