Noûs

NOÛS 00:0 (2013) 1–13

The Review Paradox: On The Diachronic Costs of Not Closing Rational Belief Under Conjunction

HANNES LEITGEB*
LMU Munich

Abstract

We argue that giving up on the closure of rational belief under conjunction comes with a substantial price. Either rational belief is closed under conjunction, or else the epistemology of belief has a serious diachronic deficit over and above the synchronic failures of conjunctive closure. The argument for this, which can be viewed as a sequel to the preface paradox, is called the 'review paradox'; it is presented in four distinct, but closely related versions.

1. Introduction

Is rational (all-or-nothing) belief—the set of propositions believed by a perfectly rational agent—bound to be closed under conjunction?

There are quite a few philosophers who think the answer is *no* (such as, famously, Henry Kyburg). They do so in spite of a great tradition in doxastic/epistemic logic according to which the closure of belief under conjunction counts as a fundamental rationality postulate (cf. Hintikka 1962, Levi 1967). In the eyes of philosophers such as Kyburg, the logical tradition suffers from "conjunctivitis" (Kyburg 1970).

Some of the opponents of conjunctive closure are impressed by the Lockean thesis (cf. Foley 1993) which says that it is rational to believe a proposition X if and only if the subjective probability of X is greater than some threshold r, where the threshold in question may be vague and depending on the context. Formally: there is an r, such that for all X,

$$Bel(X) \text{ iff } P(X) > r.$$

And then they point to the obvious existence of cases in which $P(A) > r$, $P(B) > r$, whilst $P(A \cap B) \not> r$, so that, by the Lockean thesis, it must hold that $Bel(A)$, $Bel(B)$, but *not* $Bel(A \cap B)$—the propositions A and B are to be believed, though their conjunction $A \cap B$ is not.

Others might argue against the closure of belief under conjunction on grounds of paradoxes such as the preface paradox (cf. Makinson 1965): it does not seem

*We are very grateful to David Makinson, David Etlin, and an anonymous referee for invaluable comments and suggestions. This work has been supported by the Alexander von Humboldt Foundation.

irrational for an author to claim in the preface of her book that she will have made some mistakes in the subsequent chapters, and at the same time to believe each of the statements $T_1, \ldots, T_n$ that are being made in these chapters. Closure under conjunction would thus leave the author with a belief in the contradictory statement $\neg(T_1 \cap \ldots \cap T_n) \cap (T_1 \cap \ldots \cap T_n)$, which would certainly not be rational. Hence, closure must be wrong.

In the following, we are going to argue that giving up on closing rational belief under conjunction comes with a substantial price. The upshot will be: either rational belief is closed under conjunction, or else the epistemology of belief has a serious diachronic deficit over and above the synchronic failures of conjunctive closure. The argument for this, which can be viewed as a sequel to the preface paradox, will be called the 'review paradox'; it will be presented in four distinct, but closely related versions.

2. The Paradox

Let us presuppose that we intend to describe an agent's doxastic states both qualitatively, in terms of categorical belief ascriptions, and quantitatively, by means of ascribing numerical degrees of belief; abandoning either of the two kinds of ascriptions is not an option. Therefore, when the agent receives some piece of evidence X, we should be able to express what is going on doxastically in qualitative and in quantitative terms simultaneously: if she learns (or updates on) X, then something is the case that will be expressed qualitatively, and at the same time something is the case that will be be expressed quantitatively.

We consider some (inferentially) perfectly rational agent. Let t be an arbitrary point of time, let Bel_t be the set of propositions believed by the agent at t, and let P_t be the same agent's degree of belief function at t; analogously for points of time t' after t and the corresponding $Bel_{t'}$, $P_{t'}$. We will assume, without further justification, that the degree of belief function of a perfectly rational agent must always be a subjective probability measure.

Each version of our argument will proceed from three premises. In the first version, P1 is a bridge principle that tells us something about how the agent's degrees of belief and her beliefs relate to each other; P2 expresses a qualitative feature of update by evidence; and P3 states how the agent updates in quantitative terms:

P1 *If the degrees of belief that the agent assigns to two propositions are identical, then either the agent believes both of them or neither of them.* That is:

For all X, Y: if $P_t(X) = P_t(Y)$ then

$$Bel_t(X) \text{ iff } Bel_t(Y).$$

P2 *If the agent already believes X, then updating on the piece of evidence X does not change her system of (all-or-nothing) beliefs at all.* That is:

For all X: if the evidence that the agent obtains between t and $t' > t$ is the proposition X, but it holds already that $Bel_t(X)$, then for all Y:

$$Bel_{t'}(Y) \text{ iff } Bel_t(Y).$$

P3 *When the agent learns, this is captured probabilistically by conditionalization.* That is:

For all X (with $P_t(X) > 0$): if the evidence that the agent obtains between t and $t' > t$ is the proposition X, then for all Y:

$$P_{t'}(Y) = P_t(Y \mid X).$$

P1 expresses that if two propositions X and Y are assigned the same degrees of belief by a perfectly rational agent, then the agent must treat X and Y equally in terms of belief. For instance, every supporter of the Lockean thesis must accept this: for if $P(X)$ is identical to $P(Y)$, then either both of them will exceed the threshold r in the Lockean thesis or neither of them will do. More generally, every theory of belief and degrees of belief according to which belief in X supervenes, or functionally depends, on the probability of X will deliver P1 as a consequence. But P1 holds on yet more general grounds: supervenience would mean that there could not a difference in the belief status of a proposition without a difference in the probability of that proposition, which is a matter of comparing different degree of belief functions, and different belief sets, with each other. But P1, which only concerns one degree of belief function and belief set at the time, is strictly weaker than this: if $P(X) = P(Y)$, then P1 leaves open whether X and Y are believed or not, it only demands the two propositions to have the same belief status, that is, to believe both of them or neither of them. Indeed, "*I believe X but I do not believe Y, even though the two of them are equally likely for me*." sounds odd independently of the fate of the Lockean thesis or of some other principle of supervenience that might relate P and Bel. Or once again in other terms, from the point of view of the central epistemological goal of truth approximation (and disregarding other more pragmatic goals that an agent might have): if the probabilities of X and Y are the same, then this means that the agent's estimates of the truth values of X and Y are the same; since rational belief aims at the truth (see Wedgwood 2002), how could a perfectly rational agent not assign the the same belief status to the two propositions?

P2 should be quite convincing as well: it states that if a perfectly rational agent already believes X, and then she updates on X as a piece of evidence, her set of believed propositions will remain the same. *I: "I believe X to be the case." You*: "*X is the case." I: "Oh my goodness, now I need to change some of my beliefs*." does sound odd again. Accordingly, in the purely qualitative theory of belief revision (cf. Gärdenfors 1988), if X is a member of the agent's present (and consistent) set K of believed propositions, then the revision $K * X$ of K by evidence X is demanded to be K again. Since X had already been believed, receiving it as a piece of evidence should not change anything as far as all-or-nothing beliefs are concerned; the agent simply ought to retain her current belief set. The same is assumed by the less idealized theory of so-called belief base revision (see Hansson 1999) in which, unlike standard belief revision theory, the closure of belief sets under conjunction is not presupposed.

P3 is the standard Bayesian postulate on probabilistic update. There are some justifications for it in the Bayesian literature, but we will not go into them here.

We take P1–P3 for granted now. Here is then, in a nutshell and stated at first only informally, the argument that we put forward: Assume that a perfectly rational agent believes A and B but does not believe $A \cap B$. Let the agent's initial degree of belief in A lie strictly between 0 and 1.[1] Suppose the agent then receives A as a piece of evidence: when the agent updates on A, by P3, her subjective probability in B will become identical to her probability in $A \cap B$. By P1, the agent must thus have the same doxastic attitude towards B and $A \cap B$ after the update. But by P2 her doxastic attitude towards each of B and $A \cap B$ must be the same after updating on A as it had been before. Initially, by assumption, the agent believed B but did not believe $A \cap B$. Contradiction. Hence, given P1–P3, a failure of closure of belief under the conjunction of A and B leads to the absurd conclusion that the agent cannot update on A: something that should be perfectly unproblematic. Therefore, either closure of belief under conjunction must hold, or one of P1–P3 needs to be given up.[2]

One may illustrate what is going on here in terms of a sequel to the preface paradox: Assume with the paradox that the author believes each of $T_1, \ldots, T_n$ without believing $T_1 \cap \ldots \cap T_n$. Let m be the maximal number less than n so that the author believes $T_1 \cap \ldots \cap T_m$ without believing $T_1 \cap \ldots \cap T_{m+1}$; clearly, there must be such a number m in the preface paradox situation. Finally, suppose that someone writes a review of the author's book in which the reviewer strengthens the author's case for $T_1 \cap \ldots \cap T_m$, without saying anything at all about T_{m+1} or any other of the author's theses (maybe the reviewer is simply not interested in them): "*What I can say about this book is that $T_1 \cap \ldots \cap T_m$ definitely is the case*." Assume that the author is rationally absorbing this report—updating on the proposition $T_1 \cap \ldots \cap T_m$ if stated in qualitative terms, and, if stated in quantitative terms, updating on $T_1 \cap \ldots \cap T_m$ by conditionalization: then given the additional assumption that P1–P3 are the case, one encounters a contradiction, as follows from the considerations above with A being $T_1 \cap \ldots \cap T_m$, and B being T_{m+1}. It seems that the author cannot rationally take in a perfectly positive review of her book. Call this the *review paradox*.

Before we make the underlying reasoning formally precise, we introduce a second version of our paradox in which some of the concepts in P1 and P3 will be relaxed a bit. Learning evidence with certainty, as covered by P3, is rarely the case in the real world, whereas learning evidence with some probability α just a little short of 1 is much more plausible. Our new P3* will take care of this. Accordingly, P1* will extend P1 to cases in which the degrees of belief of two propositions are sufficiently close to each other without being strictly identical, where 'sufficiently close' will be treated as a vague term. By these changes we will be able to avoid replies to the paradox above of the form: sure, the agent cannot rationally update by conditionalization in the story from before, but conditionalization is artificial anyway.

This second version of our argument will proceed from three premises again, amongst which the second premise P2* will simply coincide with P2 from above (which is why we will not state P2* again):

P1* *For almost all numbers r', if the degrees of belief that the agent assigns to two propositions are sufficiently similar to r', then either the agent believes both of them or neither of them*. That is:

For almost all $0 \leq r' \leq 1$, for all X, Y: if both $P_t(X)$ and $P_t(Y)$ are sufficiently close to r', then

$$Bel_t(X) \text{ iff } Bel_t(Y).$$

P3* *When the agent learns, this is captured probabilistically by Jeffrey conditionalization* (see Jeffrey 2004, section 3.2). That is:

For all X (with $P_t(X) > 0$): if between t and $t' > t$ the evidence that the agent obtains leads her to impose the probabilistic constraint

$$P_{t'}(X) = \alpha,$$

then for all Y:

$$P_{t'}(Y) = \alpha \cdot P_t(Y \mid X) + (1 - \alpha) \cdot P_t(Y \mid \neg X).$$

P1* is a strengthening of P1 that allows for cases in which two propositions X and Y are assigned only sufficiently similar degrees of belief by a perfectly rational agent, and yet the agent must still treat X and Y equally in terms of belief. Once again, every supporter of the Lockean thesis must accept this: as long as r' is not equal to the threshold r itself, it holds that if both $P(X)$ and $P(Y)$ are sufficiently close to r', then either both of them will exceed r (when $r' > r$) or neither of them will do (when $r' < r$). Therefore, P1* holds, where in this case 'almost all' means: *all except for one* (that is, r). In order to be able to derive P1* from the Lockean thesis, it would not be possible to omit this qualification in terms of 'almost all' from it, for if $P(X)$ is very close to r but less than r, whereas $P(Y)$ is very close to r but greater than r, then X is not to be believed according to the Lockean thesis whereas Y is. However, just as it was the case for P1, also P1* may be expected to hold on far more general grounds than just the Lockean thesis.

The terms 'almost all' and 'sufficiently close' in P1* are meant to be vague, but that should not bother us, much as the potential vagueness of the threshold in the Lockean thesis is generally not perceived to be a problem. In fact, there is an even stronger correspondence to the literature on vagueness: in the terminology of that literature, P1* says more or less (ignoring possible complications through the 'almost all' quantifier) that belief is *tolerant* with respect to degrees of belief (see

Shapiro 2006, p. 8); but we leave this to one side now. For the argument below, amongst other possibilities, the following manner of making P1* crisp would do: for all points of time t, for all numbers $0 \leq r' \leq 1$ except for one, there is some ("small") number $1 - \alpha$, such that for all propositions X, Y: if both $|P_t(X) - r'| < 1 - \alpha$ and $|P_t(Y) - r'| < 1 - \alpha$, then it holds that $Bel_t(X)$ if and only if $Bel_t(Y)$.[3]

P3* is one of the usual diachronic Bayesian postulates. In the extreme case in which $\alpha = 1$, Jeffrey conditionalization simply turns into standard conditionalization on the evidence. In this sense, the original postulate P3 is actually but a special case of P3*.

Analogously to the previous argument, assuming that a perfectly rational agent's beliefs in A and B are not closed under conjunction will entail an absurd conclusion again: the agent cannot update on A in the way that the probability of A becomes close to 1. In the review paradox situation, the author cannot update on a friendly review of the form: "*What I can say about this book is that I can very much confirm* $T_1 \cap \ldots \cap T_m$."

We will now spell out the argument in full formal detail, where we will deal with both variants of the argument at the same time. Given either P1–P3 or P1*–P3*, suppose some perfectly rational agent's beliefs at time t_0 are such that

$$Bel_{t_0}(A),\ Bel_{t_0}(B),\ \text{but } \textit{not}\ Bel_{t_0}(A \cap B).$$

We also presuppose that $0 < P_{t_0}(A) < 1$.

Assume that the agent receives evidence A between t_0 and t_1: in qualitative terms, this means that the evidence that the agent obtains between t_0 and $t_1 > t_0$ is the proposition A; in quantitative terms, it means that the evidence that the agent obtains between t_0 and t_1 leads her to impose the probabilistic constraint

$$P_{t_1}(A) = \alpha$$

for some α that is either identical to 1, in the first version of the argument, or at least close to 1, in the second version. We presuppose the qualitative and the quantitative way of describing the agent's evidence to be applicable simultaneously.

Leaving the exact value of 'α' open for the moment, consider next the following thought experiment: think of α gradually tending towards 1. Then, with increasing α, it must be that $P_{t_1}(B)$ and $P_{t_1}(A \cap B)$ will get ever closer to $P_{t_0}(B \mid A)$. For, by P3*, learning proceeds in terms of Jeffrey conditionalization, and hence

$$P_{t_1}(B) = \alpha \cdot P_{t_0}(B \mid A) + (1 - \alpha) \cdot P_{t_0}(B \mid \neg A),$$

which tends towards $P_{t_0}(B \mid A)$ when α tends towards 1. And by the definition of conditional probability, the same holds for:

$$P_{t_1}(A \cap B) = \alpha \cdot P_{t_0}(A \cap B \mid A) + (1 - \alpha) \cdot P_{t_0}(A \cap B \mid \neg A) = \alpha \cdot P_{t_0}(B \mid A).$$

Therefore, when α tends towards 1, both $P_{t_1}(B)$ and $P_{t_1}(A \cap B)$ tend towards the same number $P_{t_0}(B \mid A)$. In the extreme case $\alpha = 1$ (as covered by P3), it simply holds that $P_{t_1}(B) = P_{t_1}(A \cap B) = P_{t_0}(B \mid A)$. Either way, there must be an α so close to 1 that the degrees of belief that the agent assigns to A and B at t_1 are sufficiently similar to the number $r' = P_{t_0}(B \mid A)$. In the second version of the argument, we are simply going to suppose that this number r' is amongst the "almost all numbers" over which P1* quantifies; so this is really another modest constraint on what A, B, and P_{t_0} are like. For instance, if 'almost all' means 'all except for r', then the additional assumption will be that A, B, and P_{t_0} are so that $P_{t_0}(B \mid A) \neq r$, and we add this assumption to the presumed failure of the closure of *Bel* under the conjunction of A and B; the additional constraint is modest then in the sense that $P_{t_0}(B \mid A)$ can still be "almost" any number: any number with just one exception.

Now assume that the agent's evidence imposes on her the probabilistic constraint from above for such an α. From $P_{t_1}(B)$ and $P_{t_1}(A \cap B)$ being identical or at least sufficiently similar to $r' = P_{t_0}(B \mid A)$, it follows with $t = t_1$ and P1/P1* that

$$\text{(i.i)}\ Bel_{t_1}(B) \text{ iff } Bel_{t_1}(A \cap B),$$

and the agent updating on A entails with $t = t_0$, $t' = t_1$, $Bel_{t_0}(A)$ (by assumption), and P2/P2* that both

$$\text{(ii.i)}\ Bel_{t_1}(B) \text{ iff } Bel_{t_0}(B)$$

and

$$\text{(ii.ii)}\ Bel_{t_1}(A \cap B) \text{ iff } Bel_{t_0}(A \cap B)$$

must be the case.

By assumption again, it holds that $Bel_{t_0}(B)$, which implies with (ii.i) and (i.i) that

$$Bel_{t_1}(A \cap B),$$

but then again, by assumption, $Bel_{t_0}(A \cap B)$ does not hold, which entails with (ii.ii) that

$$\text{not } Bel_{t_1}(A \cap B).$$

So we end up with a contradiction.

With P1–P3 or P1*–P3* and some failure of closing belief under conjunction being in place (as well as a minor additional assumption in the second version of the argument as mentioned before), it cannot happen that our perfectly rational agent adapts to evidence as described: she cannot update, in qualitative terms, on the proposition A, and at the same time, as far as the probabilistic side is concerned, update by conditionalizing on A or by Jeffrey conditionalizing on A with an α sufficiently close to 1.

Before we turn to the conclusions that one ought to draw from this, we will briefly discuss two further variants of the paradox in which premises P1/P1* are modified.[4]

3. A Variation

Let us replace P1 from the last section by this principle:

Q1 *If the degree of belief that the agent assigns to a proposition is identical to* 1, *then the agent believes the proposition.* That is:

For all X: if $P_t(X) = 1$, then $Bel_t(X)$.

Q1 is entailed by many theories of belief or acceptance. Indeed, "*I assign maximal degree of belief to X, but I do not believe X*" sounds strange again.

Accordingly, replace P1* from the last section by

Q1* *If the degree of belief that the agent assigns to a proposition is sufficiently close to* 1, *then the agent believes the proposition.* That is:

For all X: if $P_t(X)$ is sufficiently close to 1, then $Bel_t(X)$.

Here, 'sufficiently close' is a vague term again. For the argument below, for instance, it would be sufficient to make Q1* crisp by: for all points of time t, there is some number α (that is "close" to 1), such that for all propositions X: if $P_t(X) > \alpha$, then it holds that $Bel_t(X)$. This will then amount to an instance of the right-to-left direction of the Lockean thesis.

Q1 follows from our original P1 if given the additional assumption that there exists at least one proposition of probability 1 (e.g., a tautology) that is believed by the agent. Q1* follows from P1* given the same assumption together with the premise that 1 is amongst the "almost all" numbers r' over which P1 quantifies.

Other than Q1/Q1*, we will only presuppose P2(=P2*) and P3/P3* as used before; so Q2 = P2, Q3 = P3, Q2* = P2 = P2*, Q3*=P3*.

Now we reason as follows: Assuming Q1/Q1*, Q2(=Q2*), and Q3/Q3*, suppose there exist two propositions A, B of positive probability that are probabilistically independent of each other if measured relative to a perfectly rational agent's degree of belief function at t_0. That is:

$$P_{t_0}(A \cap B) = P_{t_0}(A) \cdot P_{t_0}(B),$$

or equivalently

$$P_{t_0}(B) = P_{t_0}(B \mid A) = P_{t_0}(B \mid \neg A) \text{ and } P_{t_0}(A) = P_{t_0}(A \mid B) = P_{t_0}(A \mid \neg B).$$

And let us suppose again that the agent believes each of A and B at t_0 but does not believe their conjunction:

$$Bel_{t_0}(A),\ Bel_{t_0}(B),\ \text{not } Bel_{t_0}(A \cap B).$$

Finally, assume that the agent's stream of evidence makes her update first on A (between t_0 and t_1), and then on B (between t_1 and t_2), taking each of their probabilities either to 1—in the Q1–Q3 version—or very close to 1—in the Q1*–Q3* version—where we exploit the independence of A and B and apply Q3/Q3* first for a suitable α and then for a suitable α'. It follows from the properties of conditionalization (Q3) and Jeffrey conditionalization (Q3*) that the independence of A and B will not be affected by this sequence of updates.

Formally: whatever the value of α is like, updating first on A leaves the probability of B the same, by B being independent of A relative to P_{t_0}:

$$P_{t_1}(B) = \alpha \cdot P_{t_0}(B \mid A) + (1-\alpha) \cdot P_{t_0}(B \mid \neg A) = \alpha \cdot P_{t_0}(B) + (1-\alpha) \cdot P_{t_0}(B) = P_{t_0}(B).$$

At the same time, the probability of A becomes α, of course:

$$P_{t_1}(A) = \alpha.$$

Furthermore, A is still independent of B relative to P_{t_1} by the definition of conditional probability and B being independent of A at P_{t_0}, as follows from

$$\begin{aligned} P_{t_1}(A \cap B) &= \alpha \cdot P_{t_0}(A \cap B \mid A) + (1-\alpha) \cdot P_{t_0}(A \cap B \mid \neg A) \\ &= \alpha \cdot P_{t_0}(A \cap B \mid A) = \alpha \cdot P_{t_0}(B) \end{aligned}$$

and $\alpha \cdot P_{t_0}(B)$ being identical to $P_{t_1}(A) \cdot P_{t_1}(B)$, by what had been shown above. For analogous reasons as before, updating P_{t_1} on B now leaves the probability of A the same while the probability of B becomes α':

$$P_{t_2}(A) = P_{t_1}(A) = \alpha \text{ and } P_{t_2}(B) = \alpha'.$$

In the case of the argument from Q1–Q3, of course, both α and α' are 1, and then the two updates are nothing but instances of conditionalization on A and B, respectively.

In any case, by Q2/Q2* and assuming α and α' to be identical to, or at least sufficiently close to, 1, the agent must continue to believe each of A, B, while still not believing their conjunction $A \cap B$. But if α and α' are 1 or sufficiently close to 1, then also the probability of $A \cap B$ must be 1 or sufficiently close to 1; for it follows from the axioms of probability that $P(A \cap B) \geq P(A) + P(B) - 1$. Thus, $A \cap B$ must in fact be believed by the agent in view of Q1/Q1* from above. Contradiction. Therefore, given either Q1–Q3 or Q1*–Q3* and a failure of closing belief under the conjunction of two probabilistically independent propositions, the agent could not update on these propositions one after the other, which is again absurd.

In the review paradox situation, this would correspond to the reviewer stating (in the Q1–Q3 case) "*I can say that* $T_1 \cap \ldots \cap T_m$ *is definitely the case. The same holds for* T_{m+1}" or (in the Q1*–Q3* case) "*I can very much confirm* $T_1 \cap \ldots \cap T_m$. *I can also very much confirm* T_{m+1}", where the author's claims $T_1 \cap \ldots \cap T_m$ and T_{m+1} happen to be independent of each other from the viewpoint of the author's degree of belief function. Another pair of reviews that our poor perfectly rational author is not able to enjoy.

4. Conclusions

What we showed in section 2 on the basis of P1–P3 and P1*–P3* was: if $Bel_{t_0}(A)$, $Bel_{t_0}(B)$, and not $Bel_{t_0}(A \cap B)$ (and $0 < P_{t_0}(A) < 1$), then *our perfectly rational agent can never simultaneously update her beliefs by A and also update her degree of belief function by assigning the maximal or at least some sufficiently high probability to A*. (In the "sufficiently high" case, this was subject to a weak additional constraint on $P_{t_0}(B \mid A)$ that we will simply suppress in what follows). Similarly, in the last section, we showed that if one relies on Q1–Q3 or Q1*–Q3*, *a perfectly rational agent can never update in the respective manner first on A, and then on B, where the two propositions are probabilistically independent*. Note that these problems affect all situations of this very general type; they do not just affect special lottery-type contexts.

Obviously, this is absurd: Why couldn't a perfectly rational agent update on evidence in these ways? How else should, e.g., the author in the review paradox react to the positive reviews of his book as described in sections 2 and 3? Either the relevant premises cannot all be true, or

$$Bel_{t_0}(A),\ Bel_{t_0}(B),\ \text{and not}\ Bel_{t_0}(A \cap B)$$

cannot hold if the agent in question is perfectly rational.

Which one should be given up? As always, different philosophers might give different diagnoses: A radical Bayesian, such as Richard Jeffrey, might take the whole misery to be yet another indication that the concept of all-or-nothing belief itself ought to be abandoned; they might say that not even dropping the closure of belief under conjunction can save the epistemologist of belief, and the whole qualitative talk of 'learning (or updating on) a proposition' needs to be given up accordingly. We will not argue against this way out of the review paradox here, but following it would certainly be against the rules of the game of this paper, because it had been presupposed from the start that abandoning either of the two kinds of belief (and corresponding learning) ascriptions was not an option.

Or *belief* and *learning* of (or *update* on) a proposition are to be kept as concepts, but one of the premises from before or conjunctive closure is jettisoned. But which one(s)?

Perhaps P1/P1*/Q1/Q1* should get rejected, which would mean that belief and degrees of belief would not line up as nicely as, e.g., the defenders of the Lockean thesis might have thought. It would not be good enough to know then that a perfectly rational agent believes two propositions to the same, or pretty much the same, degree, in order to infer that she would not believe one of these propositions without believing the other; nor would it be sufficient to know that such an agent assigns the maximal or at least a super-high degree of belief to a proposition in order to conclude that the agent believes that proposition to be true. In the case of the argument from section 2, in spite of the fact that $P_{t_1}(B)$ is identical or very close to $P_{t_1}(A \cap B)$ after updating on A, it would not be ruled out anymore that B is believed by the agent while $A \cap B$ is not; accordingly, *mutatis mutandis*, for the arguments from the last section.

Or P2(=P2*=Q2=Q2*) is being attacked, in which case one should be prepared to accept changes of belief that are grounded in evidence (the propositional contents of) which had been believed from the start. This would not just go against standard presumptions on qualitative belief revision, effectively it would mean that the system of all-or-nothing beliefs would not be able to register the occurrence of certain pieces of evidence—because there would not be any changes of belief about them—even when these pieces of evidence might trigger changes of belief in some other propositions if we still grant the combined consequences of probabilistic update (given P3/P3*/Q3/Q3*) and the assumption of P1/P1*/Q1/Q1*. For the same reason, the epistemology of belief would not be able to distinguish between cases in which some believed proposition comes along as evidence and nothing ought to be done about this by the agent, and the same believed proposition comes along as evidence and some of the agent's beliefs ought to be revised. For instance, if the evidence has the form that is described by P3* with an α that hardly exceeds the agent's present degree of belief in X, then, presumably, the agent's system of beliefs should not be affected. But if α is really close to 1, then the agent's belief system might be affected, even though in both cases X would have been believed by the agent even before the probabilistic update. As far as the argument from section 2 is concerned, with P2/P2* being dropped, one would no longer be able to conclude that the agent's belief in B and her disbelief in $A \cap B$ are being preserved when the agent receives the believed proposition A as input; analogously for the arguments from the last section.

One way of putting some pressure on P2 might be to question its validity as far as it applies to *doxastic* (or *modal*) belief contents: for instance, at first one might believe both X and also that *it might be the case that not* X, but after receiving X as a piece of the evidence one might end up believing X without believing that it might be the case that not X. Or first one believes X and also that *there is a chance that not* X, while when the evidence comes along, one believes X but no longer that *there is a chance that not* X. If so, then in either of these cases receiving X as evidence would in fact trigger some change of belief, and hence P2 would be false.[5] However, even if this were the case, it would not be clear at all whether this would lead us out of paradox: for the only instances of P2 that were required in order to get the paradoxes going were about belief contents of the form $T_1 \cap \ldots \cap T_m$ or T_{m+1} which might well be non-doxastic (and non-modal) propositions about, say, the Dead Sea or celestial bodies or natural numbers. Accordingly, if P2 were restricted just to propositions of that sort, would not the same paradoxical reasoning go through as before? Furthermore, in the Jeffrey conditionalization versions of the paradoxes, the evidence did not actually have to push the probabilities of the propositions in question to a degree of 1: hence believing that it might be the case that not X as well as believing that there is a chance that not X might both be rational before and after receiving the evidence, which means that in these cases there are not any obvious changes of beliefs with respect to doxastic (or modal) belief contents either.

In any case, giving up on P2(=P2*=Q2=Q2*) would certainly be bad news for those who subscribe to the traditional laws of the rational dynamics of all-or-nothing belief, if they also aim to play by the rules of this paper and hence do not reject simultaneous and interlocking descriptions of belief dynamics in qualitative *and* quantitative terms.

Or P3/P3*/Q3/Q3* is denied, which would go against the Bayesian mainstream.

Or: one returns to the principle of closure of belief under conjunction, which, just as dropping either of the previous premises, would have the virtue of saving a perfectly rational agent's beliefs from dynamic incoherence as exemplified by the considerations from above—that is, given the previous premises: from the embarrassment of challenging her belief in B or her lack of belief in $A \cap B$ when the evidence strengthens her degree of belief in a proposition A which she already believes to be true. By closure, our perfectly rational agent would simply never find herself in a position at time t_0 in which she believes A and B without also believing $A \cap B$.[6]

Amongst these options, restoring closure of belief under conjunction and/or abandoning P2(=P2*=Q2=Q2*) seem to be most promising emergency exits, and not just because failure of conjunctive closure and P2(=P2*=Q2=Q2*) have been the only assumptions that remained invariant throughout all four versions of the paradox. Hence, retaining both the concepts of belief and degree of belief in our epistemology, and taking the other premises for granted, the short story is: *if rational belief has the synchronic property of not being closed under conjunction, then also the rational dynamics of all-or-nothing belief must be quite different from what it is usually taken to be*. Even when a qualitative theory of belief abandons the requirement of closure under conjunction, maybe in order to be closer to a probabilistic theory of belief, differences between the two still emerge when we pass to belief change. Either the traditional epistemology of rational belief preserves closure under conjunction, or it has a more serious problem than it is normally thought to have. Either one takes one step back to the tradition or one moves even further away from it, with not much space left in between. One man's modus ponens about this will be another man's modus tollens.

Notes

[1] The existence of such a proposition A should be unproblematic: for instance, I rationally believe that I will be in my office tomorrow, even though I would not accept a bet on this proposition by which I would win one dollar if I were to be in my office tomorrow, while I would lose a billion dollars if not. By the standard Bayesian interpretation of degrees of belief in terms of betting quotients, this shows that it is rationally possible for me to believe a proposition without assigning to that proposition the maximal degree of belief of 1. Note also that the extreme version of the Lockean thesis—$Bel(X)$ iff $P(X) = 1$—would in fact guarantee the closure of rational belief under conjunction from the start; there would be nothing left to argue for in the present paper.

[2] In their effort to criticize the Lockean thesis, (Lin and Kelly, 2012, pp.958–961) also present a puzzle in which an agent's probability measure is updated by a proposition that is already believed. But they consider a particular example measure that proves it possible to run into a problem, where we are interested in an argument with general premises and an absurd conclusion that shows that one always runs into a problem given certain assumptions; they apply the Lockean thesis, which we do not; their preservation principle of "hypothetico-deductive monotonicity", which they show to be invalidated in their example, is a bridge principle for probability and belief that differs from our purely qualitative preservation principle P2 which is just the rather trivial 'if A is in K (and K is consistent), then $K * A = K$' (in belief revision terms); unlike them we do not presuppose that belief is functionally

determined by a probability measure; and closure under conjunction is not their concern, while it is the central topic in our case. In contrast to the additional versions of our paradox that will be stated further down below, Lin and Kelly restrict themselves to update by conditionalization, and they do not deal with the potential vagueness of thresholds in bridge principles for belief and probability. This said, their case is very similar to ours in addressing static postulates on belief and probability (such as the Lockean thesis) from a dynamic point of view.

[3] One might worry that P1* would be susceptible to a sorites-type of reasoning that would lead to absurdity: Start with a belief in a proposition X_1 that has probability x. Then find another proposition X_2 whose probability is $x - \epsilon$, where ϵ is sufficiently small as to make no difference to whether something counts as a belief or not (by the lights of P1*). Then find another proposition X_3 the probability of which is $x - 2\epsilon$. And so on. One might believe that eventually one would find a proposition X_k whose probability is small enough not to count as believed. If so, somewhere along the way there would have to be a pair of "adjacent" propositions, differing in probability by only ϵ, with the first believed but the second disbelieved, contra P1*.

Fortunately this is not actually the case: First of all, the tolerance principle that is enshrined in P1* only holds for almost all numbers, not for all of them, which is why there would be no guarantee for this sequence of reasoning steps to go through for each of x, $x - \epsilon$, $x - 2\epsilon$, and so on. Secondly, and more importantly, there is no guarantee either that at each of the steps a modification of the probability in question by one and the same amount of ϵ would count as "sufficiently small". P1* only demands for almost all x the existence of some such ϵ, but not necessarily the same such ϵ for different x. For instance, consider the Lockean thesis with a threshold of 0.9: subtracting an ϵ of 0.05 from an initial probability x of 0.96 would work precisely one time without changing belief into disbelief, but then for the resulting second probability $x - \epsilon$, that is, 0.91, subtracting by 0.05 would not be licensed anymore by P1*, only subtracting e.g. by an ϵ' of 0.005 would be. No sorites problem emerges from this. (We are grateful to an anonymous referee for bringing this to our attention.)

[4] This variation of the paradox was suggested to us by David Makinson in personal communication.

[5] We thank an anonymous referee for raising this concern.

[6] One can prove all of P1–P3, P1*–P3*, Q1–Q3, Q1*–Q3* to be consistent with closure of belief under conjunction. For all of these principles can be shown to follow from the joint theory of belief and degrees of belief in (Leitgeb, unpublished). And that theory is known to have a great variety of models, including also a great variety of models in which some proposition is being believed in spite of its probability being less than 1.

References

Foley, Richard. (1993) *Working Without a Net*, Oxford: Oxford University Press.

Gärdenfors, Peter. (1988) *Knowledge in Flux*. Cambridge, Mass.: The MIT Press.

Hansson, Sven Ove. (1999) *A Textbook of Belief Dynamics. Theory Change and Database Updating*, Dordrecht: Kluwer.

Hintikka, Jaakko. (1962) *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, Cornell: Cornell University Press.

Jeffrey, Richard. (2004) *Subjective Probability: The Real Thing*, Cambridge: Cambridge University Press.

Kyburg, Henry E. Jr. (1970) "Conjunctivitis", in M. Swain (ed.), *Induction, Acceptance, and Rational Belief*, Dordrecht: D. Reidel, 55–82.

Leitgeb, Hannes. "Reducing Belief Simpliciter to Degrees of Belief" and "The Stability Theory of Belief", unpublished manuscripts.

Levi, Isaac. (1967) *Gambling with the Truth. An Essay on Induction and the Aims of Science*, Cambridge, Mass.: The MIT Press.

Lin, Hanti and Kevin T. Kelly. (2012) "Propositional Reasoning that Tracks Probabilistic Reasoning", *Journal of Philosophical Logic* 41, 957–981.

Makinson, David. (1965) "The Paradox of the Preface", *Analysis* 25/6, 205–207.

Shapiro, Stewart. (2006) *Vagueness in Context*, Oxford: Clarendon Press.

Wedgwood, Ralph. (2002) "The Aim of Belief", *Philosophical Perspectives* 16, 267–297.