David C. Makinson

Logical questions behind the lottery and preface paradoxes: lossy rules for uncertain inference

Article (Accepted version) (Refereed)

Original citation:
Makinson, David C. (2012) *Logical questions behind the lottery and preface paradoxes: lossy rules for uncertain inference.* Synthese, 186 (2). pp. 511-529. ISSN 0039-7857
DOI: 10.1007/s11229-011-9997-2

Logical Questions behind the Lottery and Preface Paradoxes: Lossy Rules for Uncertain Inference

David Makinson

Abstract

We reflect on lessons that the lottery and preface paradoxes provide for the logic of uncertain inference. One of these lessons is the unreliability of the rule of conjunction of conclusions in such contexts, whether the inferences are probabilistic or qualitative; this leads us to an examination of consequence relations without that rule, the study of other rules that may nevertheless be satisfied in its absence, and a partial rehabilitation of conjunction as a 'lossy' rule. A second lesson is the possibility of rational inconsistent belief; this leads us to formulate criteria for deciding when an inconsistent set of beliefs may reasonably be retained.

Keywords

Lottery paradox, preface paradox, uncertain inference, conjunction, rationality, inconsistency, lossy rules.

1. The Lottery and its Moral

Kyburg first aired the lottery paradox in a meeting of the Association of Symbolic Logic in 1959, then the following year at the International Congress for the History and Philosophy of Science. Its first appearance in print was the year after that, buried in the middle of his 1961 book *Probability and the Logic of Rational Belief*. There was no index to mark its page, and little to indicate that it would eventually become more widely known than much of the other material in the volume. We give a slightly more concise formulation from the last pages of his 1970 *Probability and Inductive Logic*.

> Consider a fair, 1,000,000-ticket lottery, with one prize. On almost any view of probability, the probability that a given ticket (say ticket n° 1) will win is 0.000001, and the probability that it will not in is 0.999999. Surely if a sheer probability is ever sufficient to warrant the acceptance of a hypothesis, this is a case. It is hard to think of grounds on which to base a distinction between this case and the cases of thoroughly acceptable statistical hypotheses. The same argument, however, goes through for ticket n° 7, ticket n° 156, etc. In fact for any *i* between 1 and 1,000,000 inclusive, the argument goes though and we should rationally be entitled – indeed obligated – to accept the statement 'ticket *i* will not win'. A commonly accepted principle of acceptability is that if *S* and *T* are acceptable statements, then their conjunction is also acceptable. But this means that in the lottery case, since each statement of the form 'ticket *i* will not win' is acceptable, so is the conjunction of 1,000,000 of these statements, which is equivalent to the statement that *no* ticket will win the lottery, which contradicts the statement which we initially took to be acceptable that one ticket would win the lottery.

Kyburg was quite clear about the lesson that should be drawn: we must give up or at least qualify the principle of conjunction for rational belief – the rule that whenever a reasonably accepted body of statements contains two statements then it also contains, or is at least committed to accepting, their conjunction. Those brought up on a diet of deductive logic tend accept that principle without a raising an eyebrow; what the lottery paradox shows is that if we admit that our beliefs are in general no more than highly probable, then it must be abandoned, or at least moderated. To insist on the principle in all contexts is, as Kyburg punned in a paper of 1970, to suffer from conjunctivitis.

Of course, this reaction is not the only one conceivable. One could take the view that the extremely low probability of a ticket winning should not lead us to *believe* that it will lose. More generally, it is sometimes said, we should never accept any belief unless it is utterly certain; in the case of the lottery we should reserve all judgements until after the result has been drawn. Such a view has occasionally been taken in the literature – for pointers see the overview of Wheeler 2007. But as early as 1963 in his paper 'Probability and randomness' based on the original oral presentation of the paradox, Kyburg met this response, observing that its demands are so high that it would leave us with almost no beliefs at all about the world around us.

2. The Preface and another Lesson

Meanwhile in 1964 on the other side of the Atlantic a graduate student more interested in deduction than in probability, and quite unaware of Kyburg's work, was studying a textbook on the foundations of mathematics and noticed something rather odd. In the preface, the author went through the usual expressions of gratitude and shouldering of responsibility, but acknowledged that errors must surely remain. In other words, the body of the book was making a large number of assertions $x_1,\ldots,x_n$, each of which considered individually was believed to be true. But the preface stated confidently that at least one of the beliefs will be in error. Thus the book as a whole was advancing an inconsistent set of beliefs.

This puzzle was written up in a very short note, published in 1965 in the journal *Analysis* under the title 'The paradox of the preface'. The moral drawn, however, was not quite the same as emphasized by Kyburg. It was that the author of the book:

> …is being rational even though inconsistent. More than this: he is being rational even though he believes each of a certain collection of statements, which *he knows* are logically incompatible…this appears to present a living and everyday example of a situation which philosophers have commonly dismissed as absurd; that it is sometimes rational to hold incompatible beliefs.

We thus have two distinct challenges to received opinion. One puts in question the long-accepted rule of conjoining beliefs; the other queries the hallowed principle that it can never be rational to hold an inconsistent set of beliefs.

Distinct, but not in conflict. Indeed, neither author would have had quarrel with the essential point made by the other. When discussing his paradox in *Probability and the Logic of Rational Belief*, Kyburg remarks that although our beliefs about the lottery taken together do not *contain* any contradiction, nevertheless a contradiction may perfectly well be *deduced* from them without impugning their rationality – in other words, they form inconsistent (though rational) set. Conversely, when discussing the preface paradox in his *Analysis* paper, Makinson tentatively suggests that we make a distinction between the rationality of accepting *each* of the individual beliefs and the rationality of the *set* of all them: it can be irrational to accept their conjunction even when accepting the conjuncts. Thus, the morals drawn are complementary rather than conflicting.

Considering both of the paradoxes together provides us with a perspective rather broader than is available when they are studied alone, and each throws light on the other. The failure of consistency as a requirement on rational belief is not just a product of the vagueness that is inherent in everyday qualitative reasoning as used in the preface; it also arises in the quantitative context of the lottery paradox. Similarly, the failure of the rule of conjunction is not an artefact of probability functions with their highly specific features as defined by the Kolmogorov postulates, nor indeed of numbers in any shape or form, since it also arises in the purely qualitative context of the preface.

Historically, both paradoxes have roots older than those mentioned here. The first appendix recalls partial anticipations by Venn and Ramsey respectively.

3. Challenges for the Logician

From the standpoint of a logician there is a way in which Kyburg's lesson may be extended. As well as considering *belief sets*, we can look at *consequence* (alias inference) *relations*. In deductive logic, we are accustomed to treating the relation of classical consequence as satisfying the rule of conjunction in the conclusion, also known as the 'AND' rule, or $\wedge+$. Writing $\vdash$ for that relation, lower-case roman letters for propositions and upper-case ones for sets of them, this is the rule that whenever $A \vdash x$ and $A \vdash y$ then $A \vdash x \wedge y$.

Kyburg's moral may then be expressed as follows: while the rule $\wedge+$ is sound for purely deductive consequence relations, it is not so for non-deductive ones, i.e. for those associated with uncertain reasoning. Indeed, in a 1997 paper 'The rule of adjunction and reasonable inference', he began to look at the matter in this way, though not quite in the systematic fashion in which a logician might do.

This apparently small reformulation in terms of consequence relations rather than belief sets opens broad perspectives. New questions arise, in particular:

- If the rule of conjunction is not sound for uncertain reasoning, what kinds of rules *are* sound? In other words, what is the logic of uncertain inference? Is it different in the qualitative and quantitative cases, or the same?

- Again, given the unsoundness of $\wedge+$ for uncertain reasoning, should it simply be abandoned altogether in such contexts, or can it rescued in some way for *partial use*?

Even on the level of belief sets, an important issue arises:

- If it can sometimes be rational to hold beliefs that are jointly inconsistent, how frequent is this phenomenon? *When* is joint inconsistency tolerable, and when should it be taken as a signal to revise?

In response to the first question, considerable formal work has been carried out the last few years, with interesting results and some intriguing still-open problems; in section 4 we outline the state of play. On the second question, several rescue attempts have been made in the literature; in section 5 we review them and in section 6 give our own proposal for the restrained use of conjunction as what we call a 'lossy' rule, with calculable error bounds. To the writer's knowledge, the third question has not received much attention; in section 7 we suggest two tentative criteria for tolerating inconsistency. Proofs and some of the more technical material are placed in appendices.

4. Consequence Relations Failing $\wedge$+

4.1. Two systems for uncertain reasoning

It was only in the 1980s that serious work began on the behaviour of consequence relations for uncertain reasoning. To begin with, the focus was entirely on qualitative approaches and went under the name of 'nonmonotonic logic' – an unfortunate tag, both because of its negative character and because failure of monotony is a feature of probabilistic consequence relations as much as of the qualitative ones. The central system developed was that of preferential consequence, due to Kraus, Lehmann and Magidor (1990), henceforth KLM, building on initial ideas of Shoham in his doctoral thesis and subsequent book Shoham (1988). It can be given both a syntactic presentation and a semantic one in terms of 'preferential models', with equivalence between the two established by the KLM representation theorem. For our purposes, we need only consider the syntactic side. For a less laconic introduction see e.g. Makinson (2005).

In contrast to the situation for classical logic, one cannot give formal rules that determine a specific, favoured relation of uncertain inference. That is simply not a formal matter. But we can give closure conditions (Horn rules) that appear to be reasonable constraints. That is what KLM did. From the syntactic side, a preferential consequence relation may be defined as any relation |~ (in words: snake) between formulae (of classical propositional logic) satisfying the following Horn rules (where $\vdash$ is classical consequence, not to be confused with the snake, and $\approx$ is classical equivalence):

a |~ a	(reflexivity)
whenever a |~ x and $x \vdash y$, then a |~ y	(RW: right weakening)
whenever a |~ x and $a \approx b$, then b |~ x	(LCE: left classical equivalence)
whenever a |~ $x \wedge y$, then $a \wedge x$ |~ y	(VCM: very cautious monotony)
whenever a |~ x and b |~ x, then $a \vee b$ |~ x	($\vee$+: disjunction in the premises)
whenever a |~ x and a |~ y, then a |~ $x \wedge y$	($\wedge$+: conjunction in conclusion).

This system is known as P (for 'preferential', in allusion to the semantics). From the point of view of the discussion above, its immediate difficulty is that it includes the rule $\wedge$+, impugned for uncertain reasoning whether quantitative or qualitative. From the more specific perspective of probabilistic inference, its inclusion of the rule $\vee$+ is also unsound.

To see this, define probabilistic consequence in the natural way as threshold conditional probability. That is, consider relations $|\!\sim_{pt}$ modulo any probability function p (satisfying the usual Kolmogorov postulates) on the set of classical propositional formulae into the real interval [0,1] and any threshold $t > 0$ in that interval, defined by putting $a \mid\!\sim_{pt} x$ iff either $p(x|a) = p(a \wedge x)/p(a) \geq t$ or (limiting case) $p(a) = 0$. Equivalently, iff $p(a \wedge x) \geq t \cdot p(a)$. Take a Horn rule, like any of those mentioned above, to be probabilistically sound iff it preserves probabilistic consequence. That is, iff for an arbitrary instantiation of the rule, every relation $|\!\sim_{pt}$ satisfying the premises satisfies the conclusion of the rule. Then it is easy to show that the first four of the above rules are indeed probabilistically sound, but the last two, $\vee$+ and $\wedge$+, are not.

What, then, is the logic of such probabilistic inference? In particular, what Horn rules (alias closure conditions) does it satisfy? A conjecture was made by Hawthorne in 1996. Although $\vee$+ and $\wedge$+ are not probabilistically sound, they have weakened versions that are:

Whenever $a \mid\sim x$, $b \mid\sim x$ and $a \vdash \neg b$, then $a \vee b \mid\sim x$ (XOR: exclusive $\vee$+)
Whenever $a \mid\sim x$ and $a \wedge \neg y \mid\sim y$, then $a \mid\sim x \wedge y$ (WAND: weak $\wedge$+).

XOR weakens $\vee$+ by adding the condition that a and b are mutually inconsistent ($a \vdash \neg b$), while WAND weakens $\wedge$+ by strengthening the condition $a \mid\sim y$ to $a \wedge \neg y \mid\sim y$. The system consisting of the first four of the above rules, plus these weakened versions of $\vee$+ and $\wedge$+, is sound with respect to probabilistic consequence. But is it complete? To answer, we summarize results concerning this system, called O, obtained by Hawthorne and Makinson (2007) and Paris and Simmonds (2009).

4.2. How complete is Hawthorne's system?

In part, the answer to the question depends on what, in this context, is meant by completeness. The system is certainly not complete for non-Horn rules – witness for example the alternative-conclusion ('almost Horn') rule of negation rationality, acronym NR. This says that whenever $a \vee b \mid\sim x$ and $a \vdash \neg b$, then either $a \mid\sim x$ or $b \mid\sim x$. It is easy to verify that this quasi-converse of XOR is also sound for probabilistic consequence as defined above but does not follow from the Hawthorne postulates (see the second appendix for two short proofs of the latter, one direct and quantitative, the other indirect and qualitative).

Nor is the system complete for Horn rules with infinitely many premises. An elegant counterexample can be constructed with countably many of them; it does its job essentially because of Archimedes' theorem: for any real $\varepsilon > 0$ there is a positive integer n with $n \cdot \varepsilon > 1$ (see Hawthorne and Makinson (2007) for details).

So do we have completeness over the restricted but familiar family of finite-premise Horn rules? One hopeful sign is that it can be shown that whenever such a rule is probabilistically sound then it is derivable from the KLM postulates for preferential consequence relations (though not, of course, conversely, given the presence of $\vee$+ and $\wedge$+ in the latter), which certainly limits the field of potential counterexamples. Another good sign is that some well-known Horn rules that are derivable in the KLM system but not that of Hawthorne, notably cumulative transitivity (alias cut) and cautious monotony, are not probabilistically sound. Encouraged by these facts, and discouraged by the repeated failure of their candidate counterexamples, Hawthorne and Makinson (2007) conjectured a positive answer to the completeness question.

That conjecture was shattered by Paris and Simmonds (2009), who found examples of finite-premise Horn rules that are probabilistically sound but not derivable in Hawthorne's system. Indeed, they found an infinite (and very intricate) family of such rules, which taken together with those in O form a complete postulate set. We mention one of their simplest counterexamples. It may be called (2,3)XOR – in contrast with (plain) XOR, which can be seen as (1,2)XOR – and it says: whenever a_1, a_2, a_3 are pairwise inconsistent formulae such that $a_i \vee a_j$ |~ x for all distinct $i,j \leq 3$, then $a_1 \vee a_2 \vee a_3$ |~ x.

Once this rule is articulated, it is straightforward to verify its probabilistic soundness, making use of XOR together with the already-known probabilistic soundness of the non-Horn rule of negation rationality, mentioned above. However, showing the independence of the rule (2,3)XOR from system O is a more delicate affair. Paris and Simmonds (2009) obtain the result as a corollary of sophisticated constructions using non-standard integers and deep facts about the real numbers. In the third appendix we provide a direct and elementary proof.

4.3. Qualitative semantics without $\wedge$+

These results tell us a good deal about the rules that are sound for *probabilistic* uncertain inference. But what about *qualitative* reasoning: what rules are appropriate there? We have suggested that the system should be weaker than that for KLM preferential consequence, since the paradox of the preface teaches us that $\wedge$+ should fail. But how much weaker should it be?

Intuitively, it should include all the one-premise rules in Hawthorne's system O, presumably also the two-premise ones (exclusive $\vee$+ and weak $\wedge$+), perhaps too the non-Horn rule of negation rationality. But there are other Horn rules that not derivable in O but are obtainable in the KLM system P – notably unrestricted $\vee$+, cumulative transitivity, and cautious monotony. Are *they* qualitatively acceptable?

It is difficult even to begin answering this question rigorously, because to do so we need a qualitative semantics for uncertain inference that, unlike the preferential semantics for the KLM postulates, fails the rule $\wedge$+. Can it be constructed? One candidate semantics, based on the notion of 'large' subsets of a given set, is sketched in the fourth appendix. However, it is rather unsatisfying because it seems to do little more than translate the rules of O to the semantic level without making it much clearer which among them might be more acceptable or desirable.

In this connection, it may be noticed that we do not need the full force of the Kolmogorov postulates to ensure that all of the rules of system O plus the almost-Horn rule of negation rationality are probabilistically sound: a considerably weakened version suffices for the job (see the fifth appendix). Moreover, this weakened set makes no use of the arithmetic operation of addition. It does still employ division to define the notion of conditional probability. However, if we treat conditional probability as a primitive two-place function rather than as being defined from one-place probability, then it becomes possible to eliminate all numerical operations (addition, division, multiplication) from the postulates, leaving us with an axiom system of an entirely non-numerical nature but still strong enough to support system O and negation rationality. This is done in section 5 of Makinson (2011).

5. Three Unsuccessful Rescue Expeditions for $\wedge$+

Even when we put aside the deductive context and focus on uncertain inference, the failure of $\wedge$+ is hard to live with. Efforts have been made at rehabilitation, partial or total, pragmatic or formal. In this section we recall three rescue attempts in the literature and comment on their shortcomings. In the following section we describe a fourth one, which we suggest is more successful within its limited ambitions.

5.1. Quick and dirty

The first rescue effort is pragmatic rather than formal in character. It is implicit in much of the earlier literature on nonmonotonic reasoning (including work of the present writer) and emphasizes the limited capacities of human processing in real-life situations.

There are of course problems of uncertain reasoning that do not permit the assignment of meaningful probabilities to statements. Even when such assignments are possible, practical constraints may prevent us from carrying out any numerical calculations beyond those that can be done in the head or on the fly.

In such conditions we must manage our beliefs with minimal book-keeping. We need a 'quick and dirty' inference system. We should no longer worry about how high our degree of belief might really be, so long as it is over a given threshold. Once a conclusion is accepted, it is accorded full honours and treated as if it had our unreserved belief or unit probability, unless and until it is thrown out. One of these honours is the right to be conjoined.

Despite its initial attractions, which have given it some appeal among computer scientists working on inference for artificial intelligence, this is a rather coarse procedure in several respects. In the first place, while there are indeed many situations in which we are limited in the way described, there are also plenty of others in which we can nuance our levels of acceptance to at least some degree, and some where it makes sense to use probabilities. In such contexts it is rather foolhardy not to make full use of our resources. Secondly, it is not entirely clear whether the 'quick and dirty' rationale is intended to rehabilitate the rule of conjunction fully, or only partially. In the former case, it goes too far, since it leaves us open to the lottery paradox again. In the latter case, it leaves us unclear when we should refrain from applying it. Finally, it appears to take us beyond $\wedge+$. If uncertain conclusions are really to be treated as if they had unreserved belief, then why not accept also monotony and thus all the Horn rules valid for classical consequence? Taking these features into consideration, a finer analysis seems to be needed.

5.2. Redefining soundness using limits

Whereas the first rescue attempt was rough and pragmatic, the second is the very opposite – fully formal, subtle and infinitely discriminating in a way that appeals to mathematicians. Moreover, it drives a clear wedge between $\wedge+$ and monotony.

The leading idea is to move to an epsilon/delta account of soundness, i.e. one in terms of limits. Consider any Horn rule for consequence (to fix ideas, think of $\wedge+$ and monotony as contrasting examples), and define it to be *uniformly eventually sound* iff the conditional probability of the conclusion of the rule may be brought as close to one as desired by bringing the probability of the premises sufficiently close to one. To be precise: iff for every threshold $t \in (0,1)$, there is an $s \in (0,1)$ such that for every probability function p: if $1 > p(x_i|a_i) \geq s$ for each premise $a_i \mid\!\sim x_i$ of the rule, then $p(y|b) \geq t$, where $b \mid\!\sim y$ is the conclusion of the rule. This kind of definition of probabilistic soundness was proposed as long ago as 1966 and 1975 by Adams, and again with different analytical tools by Pearl (1988, 1989). They showed that it rehabilitates $\wedge+$ but does not support monotony, and is characterized by the KLM postulate set P for qualitative uncertain inference.

The construction is mathematically elegant. But with its $\forall t \exists s \forall p$ form it uses rather more machinery than may seem appropriate for the task at hand. More important, by rehabilitating the rule of conjunction without reserve or qualification it goes too far, since it leaves us open to the lottery paradox again.

5.3. Probable small loss

The third rescue attempt has yet a different character. Proposed by Johnson and Parikh (2008), it does not even attempt to drive a wedge between conjunction and monotony, but to salvage monotony in a statistical manner and thereby also provide a partial recovery of the rule of conjunction.

Monotony may be expressed in various ways. In our context, where we are looking at consequence relations taking individual propositions as premises (rather than sets of them), it is most conveniently enunciated as the Horn rule of strengthening the premise: whenever a |~ x then $a \wedge b$ |~ x. As is notorious, uncertain inference fails this rule – hence the negative name 'nonmonotonic logics' for systems such as the KLM 'preferential' P. Moreover, as is well known, the failure is central in the sense that adding monotony to P gives us full classical consequence. Indeed, even adding it to Hawthorne's system O, gives us $\wedge$+ (in just two steps, using monotony and weak $\wedge$+), thus yielding in turn $\vee$+ and thereby the whole of P, and hence again classical consequence. So any rescue of monotony in the context of O is *ipso facto* a rehabilitation of the conjunction rule and more.

The Johnson/Parikh idea is most easily expressed in the field-of-sets mode, reading a,b,x,y as subsets of some fixed finite universe u rather than as propositions in a finite language. Suppose we are reasoning over a finite universe u, and a,x are subsets of u such that the probability $p(x|a)$ is high (e.g. over 0.95). Suppose for simplicity that our probability function p is based on a uniform distribution. Then for almost all of the possible values of b (say, 99% of them), $p(x|a \cap b)$ is almost as high (say, at least 0.94). Roughly speaking, there is usually little loss occasioned by the additional condition. In unpublished work, Johnson and Parikh have generalised from the particular values of $p(x|a)$ and $p(x|a \cap b)$ used as illustrations above to give bounds in the general case.

Curiously, this idea – randomly selecting a value for an elementary letter occurring only in the conclusion of the Horn rule – cannot be applied *directly* to conjunction, where all the letters in the conclusion of the rule already appear in the premises. But as mentioned, any partial justification of monotony will automatically be a partial justification of $\wedge$+, since the latter is derivable from the former in system O. And being only partial, it does not expose us again to the preface paradox.

However, in the present writer's view, the good news for monotony is more than counterbalanced by some very bad news: there is also a sense in which even a single application of that rule can be absolutely disastrous. No matter how high the conditional probability of x given a, so long as it is less than unity the conditional probability of x given $a \wedge b$ can be *zero* for suitable choices of b. Thus the rule of monotony leaves us open to truly catastrophic one-step inferences, whether by an innocently injudicious choice of b or an insidious one in the course of adversarial debate. For this reason, even a partial salvage of $\wedge$+ that proceeds by refloating monotony goes too far.

6. A Fourth Rescue: Small Maximal Loss

We suggest a fourth rescue effort for the rule of conjunction in the conclusion. Like the epsilon/delta approach, it clearly separates $\wedge+$ and monotony, but in a quite distinct manner. Like the statistical approach, it pays close attention to diminishing numerical values, but again in a different way.

We have already noted that a single application of monotony can give rise to a total loss of probability: no matter how close $p(x|a)$ is to unity, the value of $p(x|a\wedge b)$ can be zero for suitable choice of b (e.g. as $\neg x$). But the rule of conjunction is not open to that kind of disaster. Provided the conditional probability of each of x and y, given a, is above 0.5, that of $x\wedge y$ given a cannot be zero. More generally, if the former are 'sufficiently high', the latter must be 'reasonably high'. Thus there seems to be a sense in which $\wedge+$ is less disastrous for uncertain consequence relations than is monotony. Can we make this precise?

Consider any n-premise Horn rule, authorizing passage from $\{a_i \mid\!\sim x_i : 1 \leq i \leq n\}$ to $b \mid\!\sim y$, and let p be a probability function. We use a concept that goes back to Boole (1854) in the context of unconditional probabilities, and to Adams (1966) for conditional ones. Say that the rule is *sum-sound under p* iff the improbability of y given b is less than or equal to the sum of the improbabilities of the x_i given their respective a_i. In other words, iff $1-p(y|b) \leq \Sigma\{1-p(x_i|a_i): 1 \leq i \leq n\}$. In the limiting case that $n = 0$, where the right hand side becomes $\Sigma\varnothing$, this is understood as requiring that $1-p(y|b) \leq 0$, i.e. that $p(y|b) = 1$.

When a rule is sum-sound under *every* probability function p, we say simply that it is *sum-sound*. This property contrasts with plain (probabilistic) soundness as we defined it earlier, which in effect requires that $1-p(y|b) \leq max\{1-p(x_i|a_i): i \leq n\}$ for every probability function p. Sum replaces max.

For zero and one-premise rules, the two concepts coincide. Moreover, when more than one premise is in play, plain soundness implies sum-soundness, since $max\{1-p(x_i|a_i): i \leq n\} \leq \Sigma\{1-p(x_i|a_i): i \leq n\}$. Thus all of the postulates of the system O are sum-sound. In contrast, we already know that monotony is not sound in the plain sense, and thus being a one-premise Horn rule, is not sum-sound.

On the other hand, it is not difficult to show that $\wedge+$ is sum-sound. That is, for every probability function p we have $1-p(x\wedge y|a) \leq (1-p(x|a))+(1-p(y|a))$ or, as one might prefer to write it, $imp(x\wedge y|a) \leq imp(x|a)+imp(y|a)$. Indeed, this follows from a stronger full equality which one might call *sharp sum-soundness*: $imp(x\wedge y|a) = imp(x|a) + imp(y|a) - imp(x\vee y|a)$. This is established by the following short calculation: LHS $= p(\neg(x\wedge y)|a) = p(\neg x\vee\neg y|a) = p(\neg x|a) + p(\neg y|a) - p(\neg x\wedge\neg y|a) = 1-p(x|a) + 1-p(y|a) - p(\neg(x\vee y)|a) =$ RHS as desired.

Sum-soundness may be written in variously arithmetically equivalent ways, for example, $p(x\wedge y|a) \geq p(x|a)+p(y|a)-1$, and likewise for the sharp version we have $p(x\wedge y|a) = p(x|a)+p(y|a)-p(x\vee y|a)$. The verifications of equivalence are routine.

In this precise sense, the loss of probability occasioned by application of the conjunction rule is limited. The improbability in the conclusion is no more than the sum of the improbabilities in the two premises. *A fortiori*, it is no more than twice the maximal improbability in the premises.

It follows immediately, as noted by Adams (1966), that in an extended session of reasoning with final conclusion $b \mid\sim y$, if we apply the conjunction rule n times while none of the other steps in the session lose conditional probability, then $p(y|b) \geq 1-n\varepsilon$, where ε is the greatest value of $1-p(x|a)$ for any premise $a \mid\sim x$. A more refined analysis of the derivation tree for $b \mid\sim y$ permits a tighter estimate, defined recursively: see the sixth appendix.

This can be seen as a partial justification for applying the rule $\wedge$+ in probabilistic contexts. We can employ the rule $\wedge$+ so long as we do so in moderation, and beginning from premises that are well above the standard of probability required by our threshold. The rule of conjunction in the conclusion is thus 'lossy', just as procedures for compressing the data for a digital image are. The amount of loss is variable, but we have a relatively small maximal loss on each application, and we can tolerate it if the initial images/premises are very sharp and we do not apply the compression/inference too often.

The sum-soundness of $\wedge$+ prompts a search for bounds of other Horn rules available in the KLM system P but also failing for probabilistic inference. In particular, unrestricted disjunction in the premises ($\vee$+), cumulative transitivity (CT) and cautious monotony (CM) are all in that situation. How lossy are they?

We know (Hawthorne and Makinson 2007) that each of these rules is derivable from the system O supplemented by $\wedge$+. Now this alone does not quite imply that any lower bound for $p(x\wedge y|a)$ also serves as a lower bound for each of those three rules, but an inspection of the derivations reveals that they have the following three features: (1) $\wedge$+ is applied only once, (2) weak $\wedge$+ is not applied at all, while (3) exclusive $\vee$+ is applied only with one of its premises a classical consequence. Under these conditions, it can be shown that any lower bound for $\wedge$+ also serves as a lower bound for the derived rules. Details are given in the seventh appendix, which also provides tighter lower bounds for all three rules.

These observations suggest further questions. Certain almost-Horn rules are often considered as supplements to the KLM system, in particular (in ascending order of strength): negation rationality, disjunctive rationality and rational monotony. We have already noticed that the first of these is probabilistically sound, and so has zero loss; but the others are not so. Open problem: are there any interesting bounds to their losses?

7. Inconsistent Sets of Beliefs

As we have seen, the lesson drawn by Makinson from the preface paradox, also implicit in Kyburg's discussion of the lottery, is that it can sometimes be rational to maintain an inconsistent set of beliefs. But, it might be asked, doesn't this licence logical irresponsibility, authorizing agents to persist in their beliefs regardless of their internal conflicts? If not, what criteria might be used to sort out defensible cases of such obstinacy from indefensible ones?

Before tackling these questions, we relax for a moment in English literature, where there are many famous dismissals of consistency. For example, Ralph Waldo Emerson is known for declaring that "a foolish consistency is the hobgoblin of little minds". Oscar Wilde described it as "the last refuge of the unimaginative", and Aldous Huxley remarked that "The only completely consistent people are the dead". Rather strong language! Are these lights of the literary world advocating logical irresponsibility, or are they merely being flippant? In general, neither one nor the other. In such dicta, consistency is usually understood in a particular way, as meaning *consistency of beliefs over time*. In other words, to be 'consistent' in one's beliefs is understood as never revising them, even when they run into difficulty. The quotations may thus be read as incitations to change one's beliefs when one discovers a fact incompatible with them, rather than ignore the new item and remain fixed in the old frame.

One obvious feature of both lottery and preface paradoxes is the large number of 'internal' beliefs involved (the claims about individual tickets, statements in the body of the book). This is evidently relevant to the feeling that the entire collection of beliefs, internal plus external (the rules of the lottery, the preface of the book) may rationally be accepted despite their joint inconsistency. If we gradually reduce the number of tickets or size of the book, the feeling diminishes. Kyburg began with a million tickets, writing in all the zeros to emphasize the magnitude; with just a dozen or even a hundred, the example would be rather less convincing. Size matters.

Clearly this is a 'slippery slope', reminiscent of many others in epistemology and methodology dating from the bald man conundrum of Greek antiquity up to the choice of significance levels for evaluating a hypothesis in the light of experimental data as laid down in modern statistics. In general, there are no criteria for privileging any specific cut-off point. If in practice we need to make binary decisions that depend on such a point, we must accept an element of convention or arbitrariness in setting it. But although there are no grounds for decreeing a particular point of cleavage, there can be reasons for making it rather high or rather low. It is partly a matter of comparing costs and benefits – in the present context, those of accepting an inconsistent belief ensemble compared with those of not doing so.

For the preface, there are essentially three alternatives. One is to drop the belief, expressed in the foreword, that there are surely residual errors somewhere inside. But this is to fly in the face of experience, and the cost is undue confidence in the total text. A second is to throw assertions out from the body of the book – not just one, but sufficiently many to make us really sure that there is no longer serious danger of any error. The cost is that we will end up with a very slim book indeed. The third is to revise the entire book over and over again, re-checking sources and grounds for every single assertion made. The third alternative is surely the one to follow – but only for so long as the costs of doing so remain bearable. These include energy, time and opportunities lost for doing other things, along with diminishing marginal returns in reduction of error.

In the case of the lottery, there is essentially only one alternative to accepting all beliefs in the inconsistent set. It was already mentioned briefly in the first section: suspend *all* of the many equi-probable beliefs about individual tickets until the lottery has actually been drawn and the prize-winner announced. It may at first seem that mere suspension of belief has little or no cost, but this can be an illusion. It can rob us of grounds, imperfect but relevant, for action and may also encourage the growth of unrealistic hopes. For example, I may be worried by my friend's extravagant spending on lottery tickets but remain silent, crossing my fingers in the feeble fancy that one of his tickets may win before he bankrupts himself.

But surely, it may be retorted, the potential cost of inconsistency is enormous, far outweighing any of the above. We may find ourselves inferring, in accord with the principles of classical logic, any proposition whatsoever.

The reply is ‘yes’, ‘no’, and ‘in between’. On the one hand, if we use the full resources of classical logic and sufficiently many of our individual beliefs as premises, then we certainly do run that risk. At the other extreme, if we refrain entirely from using the conjunction rule, following Kyburg’s proposal in a radical manner, then we face no risk at all – at the cost of abandoning a lot of potentially useful conclusions. However, if we use conjunction as a ‘lossy rule’, i.e. one with small maximal loss of probability per application, then the risk is both moderate and calculable while more conclusions become available for use.

Our suggestion, therefore, is that at least two criteria are in play together. A known inconsistency in a belief set may be tolerated when the number of items involved is large (strictly speaking, when the least size of any inconsistent subset is large), and the costs of eliminating the inconsistency by contraction or revision are high. In such a case, the rule of conjunction should be applied sparingly, given its lossy nature.

Appendices

1. Prehistory of the Lottery and Preface Paradoxes

It is interesting to compare the lottery and preface paradoxes, as formulated respectively by Kyburg (1961) and Makinson (1965), with a discussion by John Venn (1876) and a very brief remark of F.P. Ramsey (1929).

In the second edition of his book *The Logic of Chance*, chapter XIV on fallacies in probabilistic reasoning, sections 34-36, Venn discusses the “undue neglect of small probabilities”. He considers a situation where we have a large number of mutually exclusive but together exhaustive possibilities, each of very small probability, but with one a little more probable than any of the others. To facilitate the comparison with Kyburg, we can think of many people, each holding a single lottery ticket, except for one individual who has two of them. We thus have a large number of propositions of the form ‘individual x has a winning ticket’. Should we believe any of them?

Some might urge, says Venn, that “we ought forthwith to accept the one which, as compared with the others, is the most plausible, whatever its absolute worth may be”. But in his view, “This seems distinctly an error ... The only rational position surely is that of admitting that the truth is somewhere amongst the various alternatives, but confessing plainly that we have no such preference for one over another as to permit our saying anything else than that we disbelieve each one of them.” (chapter XIV section 36). In other words, we should disbelieve each, but accept that their disjunction holds. In terms of their negations and perhaps over-interpreting Venn’s ‘disbelieve’ as ‘believe the negation’: believe each, but disbelieve their conjunction.

Ramsey’s remark is contained in a posthumously published two-page manuscript of 1929 entitled “Knowledge”. There he comments that we “cannot without self-contradiction say p and q and r and … and one of p, q, r,… is false … But we can be nearly certain that one is false and nearly certain of each…”.

It is tempting to see Venn as anticipating the lottery paradox and Ramsey as providing the germ of the paradox of the preface. However, they are rather less than explicit about the general principles at stake. Neither of the two authors actually articulates the rule of conjunction as a closure condition on belief sets or as an inference rule, nor formulates in general terms the principle of consistency for uncertain but rational belief. Nevertheless, the passages do implicitly serve to put those principles in serious question, and can thus be seen as partial anticipations of their later counterparts.

2. Two Ways of Showing that Negation Rationality is not Derivable in O

There are two easy ways of showing that the 'almost-Horn' rule of negation rationality NR (whenever $a \vee b$ |~ x, where a and b are classically inconsistent with each other, then either a |~ x or b |~ x) is not derivable in O. One is direct and probabilistic, the other indirect and qualitative.

For the probabilistic argument, consider the intersection $|\!\sim_{pqt} = |\!\sim_{pt} \cap |\!\sim_{qt}$ of any two probabilistic consequence relations (as defined in section 4.1 above) with the same threshold t but different probability functions p, q. Since the postulates of O are sound for each of the relations $|\!\sim_{pt}$, $|\!\sim_{qt}$ considered separately (noted in section 4.1) and these postulates are also Horn rules, it follows that they are sound for the meet relation $|\!\sim_{pqt}$. But for suitable choices of probability functions p, q and mutually inconsistent a, b we can have $a \vee b \,|\!\sim_{pqt} x$ but neither $a \,|\!\sim_{pqt} x$ nor $b \,|\!\sim_{pqt} x$. For example, take the language with just two elementary letters e, e', and define functions p, q on the state-descriptions $\pm e \wedge \pm e'$ by putting $p(e \wedge e') = 0.4$, $p(e \wedge \neg e') = 0.1$, $p(\neg e \wedge e') = 0.4$, $p(\neg e \wedge \neg e') = 0.1$ while $q(e \wedge e') = 0.1$, $q(e \wedge \neg e') = 0.4$, $q(\neg e \wedge e') = 0.1$, $q(\neg e \wedge \neg e') = 0.4$. Choosing $a = e$, $b = \neg e$ (so they are inconsistent), $x = e \leftrightarrow \neg e'$, and the threshold $t = 0.5$, we get the desired configuration.

The qualitative argument, used in Hawthorne and Makinson (2007), employs terms that are not defined in the present paper but will be familiar to some readers. It goes as follows: we already know that the postulates of O are sound in every preferential model, and we also know that there are preferential models failing negation rationality, so we may conclude.

3. Direct Proof of the Incompleteness of Hawthorne's System O

We sketch a direct proof that (2,3)XOR is independent of O. For simplicity, we use a field of sets rather than a propositional language, but this can easily be translated into the other framework. In a field of sets, classical consequence becomes inclusion and inconsistency between formulae becomes disjointedness. In this context, we write $\cap,\cup$ in place of $\wedge,\vee$, and the rule (2,3)XOR says: whenever a_1, a_2, a_3 are pairwise disjoint sets such that $a_i \cup a_j$ |~ x for all distinct $i,j \leq 3$, then $a_1 \cup a_2 \cup a_3$ |~ x.

Let $u = \{1,\ldots,6\}$, and consider the field of all its subsets. Define a relation |~ over the field as follows: a |~ x iff either (1) $a \subseteq x$ or (2) there is a two-element subset s of the odd numbers in u such that $x \supseteq s \subseteq a \subseteq (s \cup s' \cup x)$, where $s' = \{k+1 \colon k \in s\}$. To visualize this, draw a diagram consisting of two concentric circles, with the odds 1,3,5 in the inner ring and the evens 2,4,6 in the outer one, each even k+1 shadowing its odd k. The sets s are the two-element sets included in the inner ring, while the sets s' are their projections in the outer one. It is straightforward, though a bit tedious, to verify that all the postulates of O are sound for the relation |~ so defined. But the rule (2,3)XOR fails. Choose $a_1 = \{1,2\}$, $a_2 = \{3,4\}$, $a_3 = \{5,6\}$ (so they are disjoint), and put $x = \{1,3,5\}$ (thus with a singleton overlap with each a_i). In the diagram x is thus the inner ring, and the a_i are pie slices. With these values it is straightforward to check that $a_i \cup a_j$ |~ x for all distinct $i,j \leq 3$, but $a_1 \cup a_2 \cup a_3$ |/~ x.

Two features of this proof are intriguing: (1) although it concerns a system and a rule designed for probabilistic consequence, it is itself purely qualitative; (2) it uses a 64-element field of sets (all subsets of a six-element set) – it does not seem possible to do it with less.

4. A 'Large Subsets' Semantics for Qualitative Uncertain Inference Failing $\wedge$+

One suggestion for such a semantics, in the spirit of Schlechta (2004), would be to define models as tuples $M = (U,f,v)$ where U is an arbitrary non-empty set, f is a function taking subsets A of U to non-empty collections of subsets of A, and v is a function taking formulae of propositional logic to subsets of U that behaves classically on the classical propositional connectives. Intuitively, U is understood as a set of 'states' (more pompously 'possible worlds'); for $A \subseteq U$, $f(A)$ is read as the collection of all 'large' subsets of A; and v is seen as a satisfaction function, i.e. $v(a)$ is the set of all states in U in which formula a is deemed to be true.

Each such model M determines a consequence relation $|\!\sim_M$ defined by putting $a \,|\!\sim_M x$ iff $v(a\wedge x) \in f(v(a))$; intuitively, iff the set of states satisfying $a\wedge x$ is a large subset of the set of states satisfying a. It is not difficult to formulate conditions on the function f that suffice to validate the postulates of O plus, if desired, the almost-Horn condition of negation rationality and thus also the Paris/Simmonds rule (2,3)XOR.

While formally neat, such a semantics is not very satisfying from an intuitive or philosophical point of view. In effect, the conditions on the model function f turn out to do little more than translate the rules of O to the semantic level, giving us some further insight into the situation, but not a great deal. Moreover, while the conditions corresponding to certain of the rules (notably right weakening, very cautious monotony and weak $\wedge$+) appear eminently reasonable, others (for XOR, negation rationality) remain just as questionable as the rules themselves.

5. A Sub-Kolmogorov Basis for the Logic O plus NR

In order to establish the probabilistic soundness of all rules in O plus NR, we evidently have to make appeal to the Kolmogorov postulates for probability. But we do not have to use their full force. In particular, we can weaken the sum postulate: $p(a\vee b) = p(a)+p(b)$ whenever $a \vdash \neg b$. The following four conditions on $p: L\to[0,1]$ suffice to do the job:

1. $p(\bot) = 0$
2. $p(a) \leq p(b)$ whenever $a \vdash b$
3. $p(a\vee b) = p(a)$ whenever $a \vdash \neg b$ and $p(b) = 0$
4. $p(a\wedge x)/p(a) \leq p((a\vee b)\wedge x)/p(a\vee b) \leq p(b\wedge x)/p(b)$ whenever $p(a) \neq 0 \neq p(b)$, $a \vdash \neg b$ and $p(a\wedge x)/p(a) \leq p(b\wedge x)/p(b)$.

The first two conditions are straightforward. The third is an explicit restriction of the sum postulate to the case that $p(b) = 0$, in which case $p(a) = p(a)+p(b)$. It eliminates reference to the addition operation, and could be called *zero increase*. The fourth condition is a consequence of the Kolmogorov postulates taken together, and might be termed *proportional interpolation*. Written more succinctly using the usual notation for conditional probability, it says: $p(x|a) \leq p(x|a\vee b) \leq p(x|b)$ whenever $p(x|a) \leq p(x|b)$ and $p(a) \neq 0 \neq p(b)$, $a \vdash \neg b$.

These four conditions are *strictly* weaker than the Kolmogorov postulates, as is easily seen by considering the function p: $L \rightarrow [0,1]$ that puts $p(a) = 0$ whenever $a \approx \bot$ and $p(a) = r$ otherwise, where r is any fixed element of [0,1]. For on the one hand, that function does not satisfy the Kolmogorov postulates (so long as the language has at least one elementary letter), while on the other hand it does satisfy conditions (1) through (4).

6. Improved Bound for Iterated Application of $\wedge$+

Consider any derivation tree for b |~ y from premises a_i |~ x_i using rules from O plus $\wedge$+. Let p be any probability function. We index each node α of the tree with a real number $\varepsilon(\alpha) \in [0,1]$, calling $\varepsilon(\alpha)$ the *error estimate* for α, by the following recursion:

- If α is a leaf node labelled by a premise a_i |~ x_i, then $\varepsilon(\alpha) = 1 - p(x_i|a_i)$.
- If α is a leaf node obtained by an application of the zero-premise rule of reflexivity, then $\varepsilon(\alpha) = 0$.
- If α is an interior node obtained by an application of a one-premise rule to a node β, then $\varepsilon(\alpha) = \varepsilon(\beta)$.
- If α is an interior node obtained by an application of a two-premise rule of O to nodes β and γ, then $\varepsilon(\alpha) = \max(\varepsilon(\beta),\varepsilon(\gamma))$.
- If α is an interior node obtained by an application of the two-premise rule $\wedge$+ to nodes β, γ, then $\varepsilon(\alpha) = \varepsilon(\beta)+\varepsilon(\gamma)$.

Then it is easy to show that for every node α labelled by an expression a |~ x, $\varepsilon(\alpha) \geq 1-p(x|a)$. In other words, $\varepsilon(\alpha)$ serves as an upper bound on the improbability of x given a, equivalently $1-\varepsilon(\alpha)$ is a lower bound on the probability of x given a. In particular, for the root node ρ of the derivation, labelled by b |~ y, we have $\varepsilon(\rho) \geq 1-p(y|b)$.

The proof is a straightforward induction, piggy-backing on the recursive definition of the error estimate function. For the basis, case 1 is given explicitly in the definition, and case 2 holds because $p(a|a) = 1$ and so $1-p(a|a) = 0$. For the induction step, cases 3 and 4 hold because the rules of O are not lossy, and case 5 holds because $\varepsilon(\beta)+\varepsilon(\gamma) \geq 1-p(x|a) + 1-p(y|a)$ by the induction hypothesis and we already know (fifth appendix) that $1-p(x|a) + 1-p(y|a) \geq 1-p(x \wedge y|a)$.

Incidentally, the success of this induction illustrates the usefulness of sum-soundness, even though it is weaker than the full equality $imp(x \wedge y|a) = imp(x|a) + imp(y|a) - imp(x \vee y|a)$ verified in the text. A simple induction of this kind does not seem to go through easily for the full equality.

It would be interesting to know whether this recursive bound on iterated applications of $\wedge$+ can be improved, and whether there is an explicit one that can rival it.

7. Lower Bounds for CM, CT, $\vee$+

We first show that any lower bound for $\wedge$+ (from a |~ x and a |~ y to a |~ $x \wedge y$) that is expressed as a function of the premise probabilities $p(x|a)$, $p(y|a)$, also serves as a lower bound for each of CM, CT, $\vee$+, expressed as a function of their premise probabilities. Then we show that for each of the last three rules, the bound $\wedge$+ may be tightened somewhat. The looser bound remains of interest after the tighter one is established, as its proof illustrates a technique that also applies to any other rules whose derivation from O-plus-$\wedge$+ satisfies the three conditions mentioned at the end of section 6 of the main text.

The proof proceeds by analysing derivations of the three rules in O-plus-$\wedge$+, writing $a \mid\!\sim_{\varepsilon} x$ as shorthand for $p(x|a) \geq 1-\varepsilon$. Note that as the rules of O are probabilistically sound we may apply them taking the index of the conclusion to be the maximum of the indices of the premises. For the one-premise rules, this means that the index stays unchanged, and for the zero-premise rule of reflexivity the index is 0.

Suppose that $a \mid\!\sim_{f(\varepsilon1,\varepsilon2)} x\wedge y$ whenever $a \mid\!\sim_{\varepsilon1} x$ and $a \mid\!\sim_{\varepsilon2} y$. We want to show that $\varepsilon_3 = f(\varepsilon_1,\varepsilon_2)$ also serves as a lower bound for CM, CT, $\vee$+.

The simplest of the three derivations is for CM (from $a \mid\!\sim x$ and $a \mid\!\sim y$ to $a\wedge x \mid\!\sim y$). Given $a \mid\!\sim_{\varepsilon1} x$ and $a \mid\!\sim_{\varepsilon2} y$ we need to get $a\wedge x \mid\!\sim_{\varepsilon3} y$. From the two assumptions we have by $\wedge$+ that $a \mid\!\sim_{\varepsilon3} x\wedge y$ and so by VCM $a\wedge x \mid\!\sim_{\varepsilon3} y$.

The derivation of CT (from $a \mid\!\sim x$ and $a\wedge x \mid\!\sim y$ to $a \mid\!\sim y$) is almost as simple. Given $a \mid\!\sim_{\varepsilon1} x$ and $a\wedge x \mid\!\sim_{\varepsilon2} y$ we need to get $a \mid\!\sim_{\varepsilon3} y$. Note that by reflexivity $a\wedge\neg x \mid\!\sim_0 a\wedge\neg x$ so by right weakening $a\wedge\neg x \mid\!\sim_0 \neg x\vee y$. On the other hand, $a\wedge x \mid\!\sim_{\varepsilon2} y$ gives $a\wedge x \mid\!\sim_{\varepsilon2} \neg x\vee y$ by right weakening. Combining these by XOR and LCE and noting that $max(\varepsilon_2,0) = \varepsilon_2$, we have that $a \mid\!\sim_{\varepsilon2} \neg x\vee y$. Applying $\wedge$+ to this and $a \mid\!\sim_{\varepsilon1} x$ gives $a \mid\!\sim_{\varepsilon3} x\wedge(\neg x\vee y)$ and so by right weakening $a \mid\!\sim_{\varepsilon3} y$ as desired.

The derivation of $\vee$+ (from $a \mid\!\sim x$ and $b \mid\!\sim x$ to $a\vee b \mid\!\sim x$) is the most complex of the three. Suppose $a \mid\!\sim_{\varepsilon1} x$ and $b \mid\!\sim_{\varepsilon2} x$. We need to show that $a\vee b \mid\!\sim_{\varepsilon3} x$. Applying LCE to the second supposition, $(a\vee b)\wedge b \mid\!\sim_{\varepsilon2} x$, so $(a\vee b)\wedge b \mid\!\sim_{\varepsilon2} x\vee a$ (right weakening). Also $(a\vee b)\wedge\neg b \mid\!\sim_0 x\vee a$ (reflexivity and RW); thus $a\vee b \mid\!\sim_{\varepsilon2} x\vee a$ (XOR and LCE). Similarly, applying LCE to the first supposition, $(a\vee b)\wedge a \mid\!\sim_{\varepsilon1} x$, so $(a\vee b)\wedge a \mid\!\sim_{\varepsilon1} x\vee\neg a$ (RW). Also $(a\vee b)\wedge\neg a \mid\!\sim_0 x\vee\neg a$ (reflexivity and RW); thus $a\vee b \mid\!\sim_{\varepsilon1} x\vee\neg a$ (XOR and LCE). Putting these together by $\wedge$+ we have $a\vee b \mid\!\sim_{\varepsilon3} (x\vee a)\wedge(x\vee\neg a)$ so by RW $a\vee b \mid\!\sim_{\varepsilon3} x$ as desired.

Note a subtle limitation of these proofs. They apply to the ($\wedge$+)-bound $p(x\wedge y|a) \geq 1-[imp(x|a) + imp(y|a)]$ because this is expressed as a function of the premise-probabilities $p(x|a)$, $p(y|a)$, but they do not apply, at least in a direct way, to the exact ($\wedge$+)-value given by the equality $p(x\wedge y|a) = 1-[imp(x|a) + imp(y|a) - imp(x\vee y|a]$, because the last term is not expressed as a function of the two premise improbabilities. Once again, the inequality for sum-soundness is easier to work with than the exact equality.

All three of these bounds may be tightened – CM in one way and CT, $\vee$+ in another.

For CM: we need only observe that by multiplying out we have $p(y|a\wedge x) = p(a\wedge x\wedge y)/p(a\wedge x) = p(x\wedge y|a)/p(x|a)$ and we can replace the numerator by any known bound for it, e.g. by $p(x\wedge y|a) \geq 1-[imp(x|a)+imp(y|a)]$ or the equality $p(x\wedge y|a) = 1-[imp(x|a)+imp(y|a)-imp(x\vee y|a]$, in the latter case giving an exact bound for CM.

The first of these two improved bounds is equivalent to one found by Bourne & Parsons (1998). Using linear algebra they showed that $p(y|a\wedge x) \geq 1-[imp(y|a) / p(x|a)]$. To verify the equivalence, recall that $1-[imp(x|a)+imp(y|a)] = p(x|a)+p(y|a)-1$, so that our bound on CM may be written as $p(y|a\wedge x) \geq [p(x|a)+p(y|a)-1] / p(x|a) = 1-[(1-p(y|a)) / p(x|a)] = 1-[imp(y|a) / p(x|a)]$ as given by Bourne & Parsons.

For CT: it suffices to note that whenever $p(a\wedge x) > 0$ we have $p(y|a) \geq p(x\wedge y|a) = p(x|a)\cdot p(y|a\wedge x)$, that is, the conclusion probability is at least as high as the product of the premise probabilities. This was observed by Adams 1998 (section 6.6, page 128). As Jim Hawthorne has remarked (personal communication), this bound is always at least as high as one minus the sum of the premise improbabilities. To see this, note that $r\cdot s \geq r+s-1 = 1-[(1-r)+(1-s)]$ for any reals $r,s \in [0,1]$; for supposing $r\cdot s < r+s-1$ one gets $1-s < r-r\cdot s = r(1-s)$ which implies $r > 1$. A simple instantiation in $r\cdot s \geq r+s-1$ gives us $p(x|a)\cdot p(y|a\wedge x) \geq p(x|a)+p(y|a\wedge x)-1 = 1-[imp(x|a)+imp(y|a\wedge x)]$ as desired.

For unrestricted $\vee$+: a careful analysis of the derivation of unrestricted $\vee$+ from O-plus-CT in Hawthorne and Makinson (2007, Appendix, Fact 4.1.4) reveals that for it we likewise have the product bound $p(x|a\vee b) \geq p(x|a)\cdot p(x|b)$. This may also be obtained from its well-known limiting case that $p(x|a\vee b) \geq p(x|a)$ whenever $p(x|b) = 1$, as follows. Observe that $p(x|a\vee b) = p(x\wedge(x\vee b)|a\vee b) = p(x\vee b|a\vee b)\cdot p(x|(a\vee b)\wedge(x\vee b))$. Now applying the limiting case, we have $p(x\vee b|a\vee b) \geq p(x\vee b|a) \geq p(x|a)$ and likewise $p(x|(a\vee b)\wedge(x\vee b)) = p(x|(b\vee(x\wedge a)) \geq p(x|b)$. Thus $p(x|a\vee b) \geq p(x|a)\cdot p(x|b)$ and we are done. This bound for $\vee$+ was also given with a very intricate, indirect, proof in Adams (1996, Table I and appendix).

In contrast, neither $\wedge$+ nor CM respects the product bound. For a simple counterexample to both, let p be any probability function with each $p(\pm x\wedge\pm y) > 0$ and put $a = \neg(x\wedge y)$. Then $p(x|a)$, $p(y|a) > 0$ while $p(x\wedge y|a) = p(y|a\wedge x) = 0$.

Question: Can we find a bound for $p(x|a\vee b)$ as an arithmetic function of $p(x|a)$ and $p(x|b)$, that is tighter than the product one? Very partial answer: in the special case that $x \vdash a\wedge b$ we have $p(x|a\vee b) \geq \min[p(x|a)/2, p(x|b)/2]$, which gives a better result when both $p(x|a)$, $p(x|b) \leq 0.5$.

Acknowledgements

Thanks to an anonymous referee who pressed for a lower bound on iterated $\wedge$+ better than one first provided, Jim Hawthorne for collaborating on the product lower bounds for CT and $\vee$+ in the last appendix, Colin Howson for general discussions on probability, and Horacio Arló Costa for drawing attention to the remark of Ramsey quoted in the first appendix.

References

Adams, E. W. (1966). Probability and the logic of conditionals. (In J.K.K. Hintikka & P. Suppes eds, *Aspects of Inductive Logic*, pp. 265-316. Amsterdam: North-Holland).

Adams, E. W. (1975). *The Logic of Conditionals*. (Dordrecht: Reidel).

Adams, E. W. (1996). Four probability-preserving properties of inferences. *Journal of Philosophical Logic*, 25, 1-24.

Adams, E. W. (1998). *A Primer of Probability Logic*. (Stanford: CSLI).

Boole, G. (1854). *Laws of Thought* (reprinted New York: Dover 2003).

Bourne, R. & Simon Parsons. Propagating probabilities in system P. *Proc. 11th Internat. FLAIRS Conf*. AAAI Org, 1998: 440-445.

Hawthorne, J. (1996). On the logic of nonmonotonic conditionals and conditional probabilities. *Journal of Philosophical Logic*, 25, 185-218.

Hawthorne, J. & Makinson, D. (2007). The quantitative/qualitative watershed for rules of uncertain inference. *Studia Logica*, 86, 249-299.

Johnson, M. & Parikh, R. (2008). Probabilistic conditionals are almost monotonic. *The Review of Symbolic Logic*, 1, 73-80.

Kraus, S., Lehmann, D. & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44, 167-207.

Kyburg, H. E. (1961). *Probability and the Logic of Rational Belief*. (Middletown Conn: Wesleyan University Press).

Kyburg, H. E. (1963). Probability and randomness. *Theoria*, 29, 27-55. Reprinted in Kyburg, H. E. (1983) *Epistemology and Inference*. (Minneapolis: Minnesota Press).

Kyburg, H. E. (1970). *Probability and Inductive Logic*. (London: Macmillan).

Kyburg, H. E. (1970a). Conjunctivitis. (In M. Swain ed. *Induction, Acceptance and Rational Belief* pp 55-82. Dordrecht: Reidel).

Kyburg, H. E. (1997). The rule of adjunction and reasonable inference. *The Journal of Philosophy*, 94, 109-125.

Makinson, D. (1965). The paradox of the preface. *Analysis*, 25, 205-207.

Makinson, D. (2005). *Bridges from Classical to Nonmonotonic Logic*. (London: College Publications).

Makinson, D. (2011). Conditional probability in the light of qualitative belief change. *Journal of Philosophical Logic*, 40, 121-153. A preliminary version appeared in Hosni & Montagna (Eds), *Probability, Uncertainty and Rationality*. Springer: CRM Series Vol 10, Edizioni della Scuola Normale Superiore, Pisa, 2010.

Paris, J. & Simmonds, R. (2009). O is not enough. *Review of Symbolic Logic*, 2, 298-309.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. (Los Altos: Morgan Kaufmann).

Pearl, J. (1989). Probabilistic semantics for nonmonotonic reasoning: a survey. In Brachman et al eds *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, pp. 505-516. (San Mateo: Morgan Kaufmann).

Ramsey, F.P. (1929). Knowledge. Manuscript published posthumously in Braithwaite, R.B. ed, *The Foundations of Mathematics and other Logical Essays*. (London: Routlege and Kegan Paul 1931, pp. 258-9. Most recently republished in Mellor, D.H. ed, *F.P. Ramsey: Philosophical Papers*, pp. 110-111 (Cambridge: Cambridge University Press 1990).

Schlechta, K. (2004). *Coherent Systems*. (Amsterdam: Elsevier).

Shoham, Y. (1988). *Reasoning About Change*. (Cambridge USA: MIT Press).

Venn, J. (1876). *The Logic of Chance*, second edition. (London: Macmillan). The passage from chapter XIV section 36 cited in the first appendix above did not occur in the first edition of 1866; but it is preserved in the third edition of 1886 (reissued in 2006 by Dover Publications, New York), where it is moved to section 33 of the same chapter.

Wheeler, G. (2007). A review of the lottery paradox. In W. L. Harper & G. Wheeler eds, *Probability and Inference: Essays in Honour of Henry E. Kyburg, Jr*. pp 1-31 (London: College Publications).

Dept. of Philosophy, Logic & Scientific Method
London School of Economics
Houghton Street, London WC2A 2AE, United Kingdom
david.makinson@gmail.com

Note: This text includes some improvements made after publication: correction of a couple of minor typos in the published version, and addition of some further material on lower bounds in section 6 and appendix 7.