Annals of Pure and Applied Logic 164 (2013) 1338–1389

Contents lists available at SciVerse ScienceDirect

Annals of Pure and Applied Logic

www.elsevier.com/locate/apal

Reducing belief simpliciter to degrees of belief

Hannes Leitgeb

Ludwig-Maximilians-Universität München, München, Germany

A R T I C L E   I N F O

Article history:
Received 11 February 2012
Received in revised form 7 October 2012
Accepted 10 October 2012
Available online 30 July 2013

MSC:
03A10
03A05
68T27
60A05

Keywords:
Subjective probability
Doxastic logic
Belief revision
Formal epistemology

A B S T R A C T

Is it possible to give an explicit definition of belief (simpliciter) in terms of subjective probability, such that believed propositions are guaranteed to have a sufficiently high probability, and yet it is neither the case that belief is stripped of any of its usual logical properties, nor is it the case that believed propositions are bound to have probability 1? We prove the answer is 'yes', and that given some plausible logical postulates on belief that involve a contextual "cautiousness" threshold, there is but one way of determining the extension of the concept of belief that does the job. The qualitative concept of belief is not to be eliminated from scientific or philosophical discourse, rather, by reducing qualitative belief to assignments of resiliently high degrees of belief and a "cautiousness" threshold, qualitative and quantitative belief turn out to be governed by one unified theory that offers the prospects of a huge range of applications. Within that theory, logic and probability theory are not opposed to each other but go hand in hand.

1. Introduction: Qualitative vs. quantitative belief

It is well known that rational belief can be ascribed to agents by means of different concepts that occupy different scales of measurement. Indeed, sometimes even within one and the same area in philosophy or science, concepts of belief on different scales are used next to each other.[1]

Epistemology is a paradigm case example: In the traditional analysis of knowledge, *belief* has always been one of the central components; for whatever knowledge may be, mainstream epistemology certainly has it that if a proposition is known by a person, the same proposition must also be believed by that person. Clearly this is about belief on a qualitative or classificatory scale. At the same time, Bayesianism in formal epistemology regards a person's *degrees of belief* as crucial for the explication of concepts such as the confirmation of a hypothesis in the light of evidence; for the evidence is said to confirm the hypothesis if and only if the degree of belief that is assigned to the hypothesis is raised by assuming the evidence. This is again about belief, in some sense, but now expressed on a quantitative or numerical scale. So epistemologists

E-mail address: Hannes.Leitgeb@lrz.uni-muenchen.de.

[1] For more on scales of measurement in general, see the classic Krantz et al. [23] and its subsequent volumes on measurement theory.

0168-0072/$ – see front matter
http://dx.doi.org/10.1016/j.apal.2013.06.015

are divided in whether they apply a qualitative or a quantitative concept of belief in their theories, and some of them even apply qualitative and quantitative concepts of belief within one and the same theory (Isaac Levi is a classical example; cf. [27,28]).[2]

When we say *qualitative belief* here—also called 'belief simpliciter', 'flat-out-belief', 'all-or-nothing belief', 'binary belief', or simply 'belief'[3]—we mean belief in the sense that an agent either believes that A is the case, or he believes that $\neg A$ is the case, or he is agnostic about the whole alternative, that is, he neither believes A nor $\neg A$ and in this sense suspends judgment on A. One the other hand, by *quantitative belief* we mean the assignment of numerical degrees of belief to propositions, so that any such degree measures the strength of an agent's belief in a proposition. Typically, A is believed to degree 1 means that the agent is certain of A being true,[4] assigning a degree of belief of 0 to A corresponds to the agent being certain that $\neg A$ is true (and thus A to be false), and any degree of belief that is numerically in between these values represents the agent's strength of belief in A lying somewhere in between these states of certainty.

The compresence of a qualitative and a quantitative account of belief is to be observed not just in epistemology but also in many other areas: philosophy of science, philosophical logic, cognitive psychology, artificial intelligence, amongst others. And there are various different attitudes towards the corresponding bifurcation of scales of measurement for belief concepts: For instance, so-called radical Bayesians such as Richard Jeffrey (see, e.g., [21]) have proposed to simply *eliminate* the qualitative concept of belief from philosophical and scientific discourse; then fields such as epistemology would be left again with a concept of belief on just one level of measurement, that is, the quantitative one. But let us suppose we do not want to go down any such eliminative road.[5] So both the qualitative and the quantitative belief concept are to remain within our conceptual repertoire. Then the obvious question is: *How do qualitative and quantitative belief relate to each other?* Possible answers range from one of the two being reducible to the other—for example, qualitative belief being reducible to quantitative belief (see proposals (1) and (2) further below), or maybe *vice versa* (as argued by Harman [14])—to one of the two being reducible to the other *plus* something else (as in (3) below) to the denial of reducibility in either way or maybe even to the denial of the existence of any general and informative relationships between the two at all (see (4) below), with lots of further possible answers in between these options.

But the question from above, if taken just by itself, is still quite vague and too general in view of the many different *ways* in which qualitative and quantitative belief might relate to each other. For instance: One might be interested to know how the belief *states* to which these concepts refer relate to each other *ontologically*. Here is one possible answer to that kind of question: Assume that a cognitive agent instantiates just one type of belief state to which both the qualitative and the quantitative concept of belief refer. Then this would be similar to the case of temperature where there is but one empirical phenomenon that may be described simultaneously in terms of qualitative concepts, such as ... *is warm*, and quantitative concepts, such as *the temperature of ... is x degrees centigrade*. Accordingly, in such a case, one would expect the qualitative description to be a coarse-grained version of the quantitative description: which qualitative description applies should depend only on what quantitative description applies, maybe supplemented by

[2] Actually, there are levels of measurement other than 'qualitative' and 'quantitative' on which belief can be studied. In particular, as we will see later, if one invokes a qualitative concept of *conditional* belief—believing that B is the case *given that A is the case*—then this conditional notion will in fact come with an implicit *ordering* of possible worlds and hence it may just as well be said to occupy an *ordinal* scale of measurement that lies strictly between the qualitative and the quantitative scale. Spohn [45] suggests to go even beyond the ordinal scale by demanding that also comparisons of *differences* between ranks of (dis-)belief should make epistemic sense, but where still not all of the usual arithmetical operations on numbers make sense for such ranks; and so forth.

[3] Sometimes, in the relevant literature on this topic, the more technical term 'acceptance' is used instead of 'belief', which has the advantage of making clear from the start that not necessarily all properties of the commonsensical notion of belief need to be preserved in a precise logical reconstruction thereof. But we will stick to the more traditional term 'belief' in this paper.

[4] Speaking probabilistically, one should actually qualify this in terms of: certain *up to a set of probability 0*.

[5] We will not give any detailed argument here for why we think that one *should* not go down that road. Let it suffice to say that this eliminative procedure would necessitate such a drastic revision of all those fields in which the qualitative concept of belief is being employed that its costs would be enormous.

some parameters of context. In our example, whether a physical body is warm (relative to a person at a time) only depends on whether its temperature is above a certain threshold temperature, where that threshold is given contextually by how sensitive one is about temperature, by the temperature of the object that a person had touched immediately before, and so on. Something like that might be expected to hold then in the case of belief, too, if both the qualitative and the quantitative concept of belief refer to one and the same mental state. Since the quantitative concept of belief is likely to allow for more fine-grained descriptions of that very mental state, quantitative belief may be expected to gain some sort of conceptual priority over qualitative belief then—just as numerical temperature ascriptions are conceptually prior to qualitative ones (in science, though not in everyday life)—and speaking in terms of qualitative belief might become "merely" an instrument by which one may talk about quantitative beliefs in a suitably simplified manner.

But there are further ontological possibilities: A cognitive agent might actually instantiate both qualitative and quantitative mental states of belief, without either of the two being ontologically reducible to the other one. For example, imagine one part of the brain to determine one's qualitative beliefs and a different part of the brain to determine one's quantitative beliefs. Then the qualitative concept of belief would refer to the qualitative belief states, the quantitative concept of belief would refer to the quantitative belief states, and any answer to the question of how the two kinds of states relate to each other would depend on how the two corresponding parts of the brain interact causally. In such circumstances, qualitative belief would not just end up being an epiphenomenon of quantitative belief but it would rather be a mental state *sui generis*. However, even in such a case, the two kinds of belief states might actually be "coordinated" so that the qualitative belief states would still be equivalent to coarse-grainings of the quantitative belief states in the threshold sense explained before. E.g., the two parts of the brain might be set up in the way that both sides are constantly aiming at achieving a joint state that makes it look *as if* one's qualitative beliefs were determined completely by one's quantitative beliefs if supplied by some contextual parameters again. This might be to the benefit of the agent for otherwise his two belief systems might recommend pairwise incompatible courses of action or the like. And there are many further ontological options.

In any case: If one understands our question from above from this ontological point of view, then it is inevitable that in order to answer it successfully one will need to carry out some empirical investigations into belief along the lines of: what are the agent's belief-generating systems like? We will not take a stand here on any such ontological issues.[6]

Instead we will approach our original question from above—how do qualitative and quantitative belief relate to each other?—from a purely *logical* and hence also, to some extent, from a *normative* point of view: *In view of the fact that we have a reasonably clear picture of what the logics of qualitative and quantitative belief are like, what conclusions can we draw from this on how qualitative and quantitative belief ought to relate to each other, assuming that they satisfy their respective logics? How do they relate to each other in the case of an agent who is a perfect reasoner?* At least at first glance answering these questions should not commit us to any of the ontological possibilities from before. And since the notions of belief that we are after in this paper are notions of *rational* belief, asking for relationships between qualitative and quantitative belief that follow from their logical properties—where 'logical' is understood in a broad sense that will cover also axioms of belief revision and probability—is certainly more than just a side issue.

So what are these logics of belief on the qualitative and on the quantitative side? In the following, we will take as given what are usually considered the standard logical laws or rationality postulates for qualitative and quantitative belief; we will simply assume them without further scrutiny or much discussion.

[6] This said, without any argument, we do regard this last ontological option—two distinct kinds of mental states, but rational qualitative belief states still being coarse-grainings of rational quantitative belief states—to be the most likely one, as far as human agents are concerned.

For belief simpliciter, this will mean that the axioms and rules of doxastic-epistemic logic (Hintikka [19] is the classic reference, but see also postulates CR1–CR2 in Hempel [17]) are considered *sacrosanct*; amongst these we find axiom schemata and schematic rules such as:

$$\text{K:}\quad Bel(A \to B) \to \big(Bel(A) \to Bel(B)\big)$$

$$\text{D:}\quad \neg Bel(\bot)$$

$$\text{Necessitation:}\quad \frac{A}{Bel(A)}$$

which we have expressed here, just as usual, in the style of modal logic, where *Bel* is the sentential modal operator for belief. So the closure of the set of believed propositions under logical consequence is taken as a given; in particular, if A is a logical truth, then A is believed; if A is believed, and A logically implies B, then B is believed; and if both A and B are believed, then also $A \wedge B$ is believed. All of these claims follow from the axioms and rules of doxastic logic.[7] And, obviously, none of these are meant as descriptions of actual human agents, for we might well fail to believe, e.g., some logical truth A that is simply too complicated for us to "see" as being a logical truth, or where A is simply not interesting enough for us to be considered at all, or the like. Instead, rationality postulates such as the ones above may be taken to apply only to *perfectly ideal* agents, or they may be understood as expressing a doxastic *commitment* (see [29]) that a real-world agent carries around without necessarily being able to live up to it all the time. The details are tricky,[8] but they are not our topic here; instead we will simply presuppose the standard accounts of the logic of belief without further qualification, as it is common practice in philosophical logic. Axioms and rules of any normal systems of modal logic, such as the one above, can be justified on the basis of a possible worlds semantics for the modality in question, here, a sentential belief operator. But for our purposes it will not be necessary to deal with the justification of such logical systems in any further detail.

In one respect, we will not exploit the full power of doxastic logic in the rest of this paper: for we are going to ignore all cases of *introspective* belief. E.g., although $Bel(Bel(\top))$ is of course provable in the system of doxastic logic that was sketched above, this type of logical principle or rationality postulate will not play any role later. Nor will we add typical postulates of introspection such as the modal axiom scheme 4, that is, $Bel(A) \to Bel(Bel(A))$ or the like.

But in a different respect we will presuppose more than merely the postulates of doxastic logic: Just as there is both absolute (unconditional) probability and conditional probability on the quantitative side, there is also both absolute (unconditional) belief and conditional belief on the qualitative side: Even when an agent does not (absolutely or unconditionally) believe B, it might be the case that *on the supposition that A is the case* the agent does in fact believe that B is the case. This is very much like applying the first step in the natural deduction rule of conditional proof: assuming A, in combination with whatever else has been assumed or derived before, B is derivable. The only difference between modern logical accounts of conditional belief and this part of the classical rule of conditional proof is that according to the former it is even possible to suppose a statement A that is *inconsistent* with what one believes absolutely or unconditionally—inconsistent with whatever was given outside of the suppositional context—without necessarily a logical contradiction following from this. Instead, the act of supposing A functions in the way that at first enough of one's other beliefs are withdrawn in order to "make room" for A in one's belief system, before A is then added as an assumption, that is, as it were, as a merely hypothetical belief. It is

[7] Not everyone would agree with what we count as "the" logic of belief here: e.g., famously, Henry Kyburg (cf. [24]) denies the closure of belief under conjunction in response to his famous Lottery Paradox to which we will return later.

[8] For it is simply not an easy task to formulate in precise terms the constraints that logic imposes on any *real-world agent's* beliefs: Harman [14] has raised some doubts on whether it is possible at all to derive epistemological constraints on inference from formal laws of logic; recently, Field [9] addressed Harman's worries and presented a tentative way of how they can be resolved.

well known that there is not necessarily a unique way of achieving this goal of revising one stock of beliefs by a statement A in view of different possibilities of how to withdraw propositions from a belief set, but at least all ways of achieving this have been argued to plausibly satisfy one and the same set of rationality postulates. These are usually spelled out in terms of a revision operator $*$ that takes one's (deductively closed) set K of presently believed statements and the input formula A as arguments and which maps them to a new (deductively closed) set $K * A$ that is the result of revising K by A. In the theory of belief revision (cf. [1]—the so-called AGM account—and also [12]), the input formula A is often not so much considered as *assumed* but rather as a piece of evidence that is to be *learned*, but the axioms of belief revision allow for both interpretations, and we will stick to the suppositional one throughout the article.[9] Here is the corresponding classical AGM axiomatization of belief revision or of belief under a supposition:

Closure: $K * A = Cn(K * A)$ (where Cn is the deductive closure operator).
Success: $A \in K * A$.
Inclusion: $K * A \subseteq K + A$ (where $K + A$ is the result of adding A to K and then closing deductively).
Preservation: If $\neg A \notin K$, then $K * A = K + A$.
Consistency: If A is consistent, so is $K * A$.
Equivalence: If $(A \leftrightarrow B) \in Cn(\varnothing)$, then $K * A = K * B$.
Superexpansion: $K * (A \wedge B) \subseteq (K * A) + B$.
Subexpansion: If $\neg B \notin K * A$, then $(K * A) + B \subseteq K * (A \wedge B)$.

We will turn to these axioms later in more detail, though spelled out in a different syntactic format.[10] One justification for this set of postulates is given by Grove's [13] representation theorem: Every operator $*$ that satisfies all of these postulates can be represented in terms of what is called a sphere system (cf. [30]) or, equivalently, a total pre-order of possible worlds according to which possible worlds get ranked in terms of their plausibility; K will then coincide with the set of formulas that are true in all those possible worlds that are the most plausible ones overall; and $K * A$ is the set of formulas that are true in all those possible worlds which are the most plausible ones amongst those that satisfy A.[11] And *vice versa* every sphere system or total pre-order of possible worlds will determine an operator $*$ in this way that satisfies all of the AGM postulates from above.

On the qualitative side, it is really the laws and rules of doxastic logic *and* the axioms of AGM belief revision that we will take for granted. Indeed, for our purposes, it will be sufficient to presuppose just the latter, since, as long as one disregards introspective belief, they do contain the former as a proper part.

Now for the logic of quantitative belief: here the standard view is the so-called Bayesian one according to which degrees of belief obey the axioms of the probability calculus together with the usual definition

[9] Some of the axioms of belief revision are in fact more easily defendable if read suppositionally: E.g., the Success postulate $A \in K * A$ is problematic if A is taken as being learned—sometimes one's "evidence" is flawed and should be rejected rather than taken on board—but the same postulate is plausible if A is taken as being supposed—once A has been assumed, A becomes something that is believed hypothetically in that suppositional context and will not be rejected as long as one remains in that suppositional context.

[10] This does lead to minor differences: in particular, we are going to start from whatever given non-empty set W of possible worlds, where it will not be presupposed that W is the set of all logically possible worlds for some given propositional language. In contrast, AGM belief revision actually demands that in any ranked model every *logically* possible world for the language in question occurs somewhere in the ranking; we will not make this assumption later when we will introduce our own versions of the belief revision postulates for propositions, that is, sets of worlds (rather than sentences or formulas). As a consequence, we will not commit ourselves to a counterpart of the Consistency postulate above.

[11] The question whether for every non-empty set of worlds in any such a model there is always a non-empty subset of maximally plausible worlds needs special care; but we will not need to discuss this here.

of conditional probability.[12] So it is generally taken for granted (where for the moment we will ascribe probabilities to sentences or formulas):

Domain/Co-domain:

$$P : \mathcal{L} \to [0,1]$$

(where $\mathcal{L}$ is a language that is closed under Boolean connectives).
Logical Truth: If A is logically true, then $P(A) = 1$.
Finite Additivity[13]: If A and B are logically inconsistent with each other, then

$$P(A \vee B) = P(A) + P(B)$$

Ratio Formula:

$$P(B|A) = \frac{P(B \wedge A)}{P(A)} \quad \text{if } P(A) > 0$$

Hence, the degree of belief that is to be assigned to a logically true statement is 1; if A logically implies B then the degree of belief for A is less than or equal to the degree of belief for B; and the degrees of belief of A and of its negation $\neg A$ add up to 1. All of these claims can be derived from the axioms above. There are various kinds of justifications for these axioms, the most famous one being the so-called Dutch book argument by which one proves (given some background assumptions) that one is guaranteed to lose money in some system of bets that appear to be acceptable in view of one's degrees of belief if and only if one's distribution of degrees of belief over statements does not conform to the axioms of probability. Just as before, these axioms are not meant to describe how human agents actually distribute their degrees of belief—the logical structure of some A might be so complicated that we fail to understand that A is logically true and hence we might not assign 1 to it, or we are not interested enough to assign a degree of belief to A at all, and so on—but rather they express what degrees of belief should be like. And once again, as it is standard procedure in probability theory, we ignore these complications and simply presuppose that the axioms of probability apply to an agent's degrees of belief without further qualification, the agent thereby being perfectly rational or at least being committed to probabilistic degrees of belief even when his actual degrees of belief do not quite meet this standard.

With these characterizations of the logics of qualitative and quantitative belief in hand, our original question can now be rendered more precise: *What "logical", normatively reasonable, bridge principles relating qualitative and quantitative belief can we expect to hold over and above the logics for qualitative and quantitative belief taken separately? And what conclusions on the relationship between qualitative and quantitative belief can we draw from the conjunction of such bridge principles with these logics of belief?*

In the relevant literature, there is actually a variety of proposals for such bridge principles for qualitative and quantitative belief, where by 'bridge principle' we mean any postulate that includes simultaneously expressions for qualitative and for quantitative belief.[14] But almost all of the *classical* proposals do belong to one the following categories or are very close to them:

(1) The *Probability 1 Proposal*:

$$\mathit{Bel}(A) \text{ iff } P(A) = 1$$

[12] Once again, not everyone will agree with this: for instance, a proponent of the Dempster–Shafer theory (see e.g. [51]) will regard *its* axioms to be right ones for degrees of belief.
[13] We omit *Countably Infinite* Additivity here, for the sake of simplicity.
[14] Hilpinen [18] gives a very nice summary of the traditional theories; and Swain [47] collects many important primary sources including early versions of the theories by Levi, Kyburg, and Jeffrey. Christensen [6] and Huber and Schmidt-Petri [20] do the same, respectively, for the more recent theories.

According to this principle, an agent believes A if and only if he is certain of A. This equivalence could be understood as a definition of its left-hand side in terms of its right-hand side, but it does not have to be interpreted in such a way.[15]

In whatever manner it is being interpreted, if the Probability 1 Proposal is granted, then from it—in conjunction with the axioms of probability—one can derive the basic logical axioms for belief. For the set of propositions with probability 1 is closed under logical consequence (from finite sets of premises). This said, in spite of this logical attraction, the proposal still seems *materially* wrong, at least if taken as a principle that is supposed to be generally valid. One simply does not get the extension of the concept of belief right in this way: As Roorda [36] points out, one can believe A without being disposed to accept every bet whatsoever on A, although the latter should be that case if one assigns probability 1 to A by the standard Bayesian understanding of probabilities in terms of betting quotients, at least as long as the stakes of the bet are not too extravagant. For example, I honestly believe that I will be in my office tomorrow. But I would refrain from accepting a bet on this if I were offered 1 Euro if I were right, and if I were to lose 1 000 000 Euro if not. Alternatively, one could abandon the standard interpretation of subjective probabilities in terms of betting quotients, but breaking in this way with the mainstream tradition would come with a huge price of its own. So the Probability 1 Proposal simply seems to be too restrictive: it should be possible to believe A even in a case when one is not assigning the maximal possible degree of belief to A. This does not exclude the *possibility* of cases where A is believed and A is also assigned probability 1—it is just that this should not be so by necessity.

(2) The *Lockean Thesis*:

The obvious suggestion of how to avoid the problem that was affecting the Probability 1 Proposal is to weaken its right-hand side, that is, to maintain that having the belief that A is equivalent to merely assigning to A a degree of belief strictly above some threshold level r *less than 1*:

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

With P being fixed, r measures the agent's cautiousness with respect to his qualitative beliefs: the greater r, the more cautious the agent concerning his beliefs, for more is then demanded of believed propositions in terms of their probability. *Vice versa*, the lower r, the braver the agent as regards his beliefs, for the less is demanded of believed propositions in terms of their probability.

This proposal was termed the "Lockean Thesis" by Richard Foley [10] who traces it back to John Locke's *Essay Concerning Human Understanding*.[16] If the corresponding threshold r is also greater

[15] Roorda [36] calls this "the received view", although it is difficult to find proponents of it in the more recent literature on this topic. If anything, only its left-to-right direction is being adopted, and even this comes usually with certain qualifications: E.g., Levi [28] accepts the left-to-right direction for his so-called credal probability measures. But that is in a context in which a credal state may involve more than just one credal probability measure, and where propositions of credal probability 1 are not meant to be incorrigible (as they may cease to have probability 1 in the future in view of possible future revisions of qualitative belief). Van Fraassen [48], Arló-Costa [2], and Arló-Costa and Parikh [3] also take the left-to-right direction of the Probability 1 Proposal for granted. However, they do not just presuppose standard probability measures but really primitive conditional probability measures (Popper functions): probability measures that allow for the conditionalization on sets of probability 0; see Makinson [35] for an excellent recent overview. As they show, one can then always find so-called belief cores which are propositions with particularly nice logical properties; by taking supersets of those one can define elegantly notions of qualitative belief in different variants and strengths. Since all such belief cores have absolute probability 1, they end up with the left-to-right direction of the proposal above. Finally, in a spirit similar to that of the Probability 1 Proposal, Williamson [50] suggests to determine so-called epistemic probabilities by conditionalization on the conjunction of everything one knows, which has the consequence that if A is known then A has epistemic probability 1. But of course knowledge is not belief and epistemic probabilities are not meant to be degrees of belief either in Williamson's account.

[16] The most famous defender of the Lockean Thesis is Kyburg; see, e.g., Kyburg [24], the last chapter of Kyburg [25] which contains a comparison between Kyburg's account and some rival ones, including Levi's, and the much more recent Kyburg and Teng [26]. We should add that for Kyburg the probability of a proposition is actually an *interval* of real numbers, but not much hangs on this as far as our present purposes are concerned. More recently, Foley [10] has argued in favor of the Lockean Thesis, and Hawthorne and Bovens [15] and Hawthorne and Makinson [16] have studied the logic of absolute belief and the logic of conditional belief, respectively, that result from taking versions of the Lockean Thesis for granted. Also Sturgeon [46] defends the Lockean Thesis but combines it with the possibility of starting from a set of probability measures rather than from a single one.

than or equal to $\frac{1}{2}$, then belief would simply be equivalent to high subjective probability, which sounds right, at least at first glance. The equivalence could be understood to hold by definition, that is, by the very meaning of the expression '(rational) belief' being given by its right-hand side, or it could be supported on different grounds. Either way it leads to a logical worry: if the threshold can be chosen independently of what one's probability measure is like, the probability of $A \wedge B$ might well drop below the threshold even when the probabilities of A and B do not. This is illustrated by the famous Lottery Paradox (Kyburg [24]; see also Wheeler [49], and Douven and Williamson [8])[17]: Consider a fair lottery with, say, 100 tickets, which is certain to take place; set the threshold value r to 0.98. Then for each ticket i the statement that i will not win will end up being believed by the agent, by the uniformity of the probability measure taken together with the Lockean Thesis for that threshold r. From the closure of belief under conjunction it then follows that the agent would have to believe: ticket 1 will not win and ticket 2 will not win and ... and ticket 100 will not win. But that statement has probability 0 and hence is not to be believed, by the Lockean Thesis again. So we have a contradiction.

Kyburg's own reaction to his paradox was to sacrifice one of the standard logical properties of qualitative belief, that is, the closure of belief under conjunction, while keeping the Lockean Thesis intact. In view of the rules of the game of our own paper, this is not an option, for we take the logics of belief that were stated above as a given. And they did contain the principle of closure of belief under conjunction. But it is still clear that, *prima facie*, the Lockean Thesis has a lot to speak in its favor. In a nutshell: While the Probability 1 Proposal was logically fine but materially wrong, the Lockean Thesis seems to be materially fine—for qualitative belief *does* seem to be close to something like high subjective probability—but then again, as things stand, it does not seem to get the logical closure properties of qualitative belief right.[18]

(3) *Decision-Theoretic Accounts*:

The Probability 1 Proposal and the Lockean Thesis are pretty much the simplest possible bridge principles for qualitative and quantitative belief that one could think of. If both of them are problematic, at least without adding further qualifications, then one natural way out would be to look for a more complex set of joint principles for the two kinds of belief. One way of realizing this ambition is in terms of a decision-theoretic account of qualitative belief.[19]

The underlying idea is this: Consider *believing* A as some of kind of action (the action $bel\, A$). Which actions should you take? As decision theory has it, only those that maximize expected utility or the expected utility of which is greater than that of some relevant alternative actions or the like. So given an agent's subjective probability measure P—say, now defined over the algebra of subsets of a set W of possible worlds—and given also some utility measure u that assigns some sort of utilities (maybe "theoretical" or "doxastic", rather than "practical", ones) to the outcomes of actions in worlds, it should be the case that

$$Bel(A) \text{ iff } \sum_{w \in W} P\big(\{w\}\big) \cdot u(bel\, A, w) \text{ has property such-and-such}$$

[17] The Preface Paradox (cf. [34]) is another well-known example that illustrates the problem of joining the Lockean Thesis with the logical closure of belief.

[18] Incidentally, one can show that the Lockean Thesis *can* get the logical closure of belief right *if only the choice of the threshold value of 'r' is assumed to depend on the probability measure P in question*: It is possible to prove that the conjunction of the Lockean Thesis with the usual logical closure conditions on belief is perfectly satisfiable, and that it is satisfied if and only if belief is determined from probability in precisely the same manner in which it will be determined according to the theory in this paper—by means of a P-stable^r set. (In fact, already the right-to-left direction of the Lockean Thesis together with the logical closure of belief would be sufficient for this.) We will have to leave this alternative axiomatic way of getting at our theory to a different paper. But we will at least sketch the underlying formal point later; see footnote 26. In the theory below, we will only take the left-to-right direction of the Lockean Thesis for granted; however, we will postulate that left-to-right direction even for conditional belief, whereas if we presupposed the full Lockean Thesis it would be sufficient to deal with absolute belief in order to develop the theory.

[19] Such accounts can be found, e.g., and in different forms and of different complexity, in Hempel [17], Levi [27], Kaplan [22], Maher [33], Frankish [11].

where 'such-and-such' needs to be replaced appropriately, so that the expected utility of believing A is salient in some sense if compared to the expected utilities of alternative actions in some relevant class; the class could be, e.g., $\{bel\, A, bel\, \neg A\}$ or some other class.

The most straightforward interpretation of such a decision-theoretic account of belief would again be in terms of a definition of its left-hand side on the basis of its right-hand side, though once again there is some leeway here. In any case, what beliefs a perfectly rational agent will end up with according to this proposal will not just depend on the agent's degrees of belief but also, crucially, on what his utility measure is like (and also the set of relevant alternative acts). Depending on the properties of u, such a decision-theoretic account might well collapse into one of the previous proposals—for instance, Hempel's [17] classical decision-theoretic account of acceptance turns out to be equivalent to the Lockean Thesis from above—or it might differ completely from the previous suggestions. Indeed, some A might end up being believed while A's probability is less than $\frac{1}{2}$, which certainly seems to stretch the meaning of 'believed' beyond its normal usage.[20] Furthermore, there is nothing in the decision-theoretic format just by itself that would guarantee that any of the standard logical properties of (rational) belief follow from it: e.g., while believing A might maximize expected utility, and while believing B might do so, too, it might be the case that believing $A \wedge B$ does not. If the logic of belief is taken to be sacrosanct (as in our case), then one way of accommodating that thought in a decision-theoretic context is to compute not so much the expected utilities of single belief acts but rather of the act of choosing a unique *logically strongest believed proposition* (as e.g. in [27]) or the act of choosing a full belief *system* or a theory that is required to be closed logically *by fiat* (as this is the case, e.g., in [33]). Although this takes care of logical closure, it generates some new concerns: In the former case, it is not yet clear why the expected utility of a proposition B should be high or salient just because B is *entailed* by a proposition A that is the logically strongest believed proposition and whose expected utility is indeed maximal or at least salient in some sense. And in the other case, whether one ought to believe a proposition or not should not be *that* complicated: whether one ought to believe the single proposition B should not depend on something as elaborate and abstract as the expected utilities of choices between various *full systems* of belief, even in the case of highly idealized agents. At least there should be some provably equivalent but more effective way of determining one's rational beliefs.

In any case, if worries such as these cannot be answered appropriately, this might lead to the more general worry that it might not have been the right move—normatively, let alone descriptively—to conceive of beliefs as resulting from decisions about acts of believing in the first place. Unfortunately, we will not be able to review all the different decision-theoretic proposals in the literature here on such grounds.

(4) The *Nihilistic Proposal*:

Finally, if no proposal seems to work, one might draw the same conclusion as Roorda [36], who maintains "The depressing conclusion … is that no explication of belief is possible within the confines of the probability model".[21] In fact, maybe there are simply no interesting and generally valid bridge principles for qualitative and quantitative belief at all, let alone a general definition of the qualitative belief in A in terms of the subjective probability of A.[22]

And clearly the worry about this (lack of a) proposal is: Could it really be that *belief* is so different in nature from *degree of belief* that no sufficiently simple and general bridge principle could be formulated that would relate the two types of belief in an informative and yet transparent manner? In the

[20] But see Section 6.2.4 of Maher [33] for the contrary view as regards the acceptance of scientific theories.

[21] Ross and Schroeder [37] defend the same claim, and similar views are implicit in work done by other philosophers, such as Wolfgang Spohn and Timothy Williamson.

[22] Roorda himself then goes on to suggest an explication of belief that is relative to a *set* of subjective probability measures, which he calls the "extended probability model", rather than just one probability measure as standard Bayesianism has it. Sturgeon [46] makes a similar move. In contrast, we are going to bite the bullet and stick to just one probability measure in the theory below.

following, we will answer this question to the negative by developing a joint theory of qualitative and quantitative belief that is based solely on simple and intrinsically plausible principles from which we will be able to prove that qualitative belief is actually definable explicitly in terms of quantitative belief (and a cautiousness threshold). However, this will be achieved in a way that does not coincide with any of the approaches that were mentioned so far.

Obviously, we do not want to imply that these four classes of approaches, (1)–(4), are exhaustive. Most notably, Lin and Kelly [31,32] have recently developed a beautiful theory of qualitative and quantitative belief that does not belong to any of the four categories. On the one hand, they study systematically different possible bridge principles for belief and probability in terms of their general mathematical (especially, geometrical) properties; and on the other hand, they present and defend their own bridge principles and reduction methods. One of the differences between their proposal and ours is that they do not take the full AGM account of belief revision for granted. A detailed comparison between the two theories has to be left for a different paper.[23]

Our own proposal in this paper will derive from the following bridge law (in conjunction with the logics of belief from above), which postulates that having a sufficiently high degree of belief is a *necessary* condition for belief,

$$\text{If } Bel(A) \text{ then } P(A) > r$$

where 'r' denotes again a threshold value that is determined contextually in some way. So this is just the left-to-right direction of the Lockean Thesis. It is entailed both by the Lockean Thesis and by the Probability 1 Proposal, and it is consistent at least in principle with a decision-theoretic account of belief. Unlike the Probability 1 Proposal, the principle above does not force believed propositions to be certain from the viewpoint of the agent; in contrast with the Lockean Thesis it will not lead to contradictions in a lottery-type situation; it is but a local constraint on single beliefs and hence differs from global decision policies about whole belief systems; and it is simple and transparent, in opposition to what is claimed by the Nihilistic Proposal.

With this pretty weak bridge principle being in place, and presupposing the logics of beliefs as stated above, we will show that there is actually not a lot of room to maneuver as regards the relation between qualitative and quantitative belief. We will prove that once a subjective probability measure is given, then—up to extensional equivalence—there is more or less just one way of how to determine qualitative belief so that the logics of belief and our final bridge principle from before are validated. The vague phrase 'more or less' will be made precise later in terms of some auxiliary bridge principles that will again apply to qualitative and quantitative belief jointly.

In Section 2 we will first carry out this programme for absolute belief and for an important restricted case of conditional belief; in Section 3 we will then do the same for conditional belief in general. From the representation theorems that we are going to prove in the two sections it will follow that there is in fact even an explicit definition of belief in terms of degrees of belief that is materially adequate as long as all of our assumptions are satisfied. In other words: *belief can be explicated on the basis of degrees of belief (and a threshold)*. By this explicating definition, believed propositions must have a high enough probability without that probability necessarily being 1, and the class of believed propositions will obey all the usual logical closure conditions that one finds in the literature on doxastic logic and belief revision. The resulting theory of belief is thereby, at the same time, qualitative and quantitative, and logical and probabilistic. The theory comes with the promise of allowing for lots of interesting applications in all those areas of philosophy or science in which states of belief are of principal interest, whether expressed in qualitative or in quantitative

[23] At least a brief note on this can be found in Lin and Kelly [31], p. 10, where they comment on our theory in light of their own background theory.

terms. However, we will not be able to review these applications in this paper; that task will have to be left for a different occasion. But we will mention some of them very briefly in the final Section 4, where we will also highlight some potential problems for the theory, and where we will indicate how the theory can be extended into various directions.

Here is a little example that may serve as an illustration of what is yet to come:

Example 1. Let $W = \{w_1, \ldots, w_8\}$ be a set of eight possible worlds. Let P be the probability measure on the power set algebra on W that is given by: $P(\{w_1\}) = 0.54$, $P(\{w_2\}) = 0.342$, $P(\{w_3\}) = 0.058$, $P(\{w_4\}) = 0.03994$, $P(\{w_5\}) = 0.018$, $P(\{w_6\}) = 0.002$, $P(\{w_7\}) = 0.00006$, $P(\{w_8\}) = 0$.

Then if the cautiousness threshold r is $\frac{1}{2}$, our theory is going to entail: $X \subseteq W$ is to be believed (unconditionally) by the agent if and only if $\{w_1\} \subseteq X$. In other words: The logically strongest believed proposition with respect to P and $r = \frac{1}{2}$ is $\{w_1\}$. Furthermore, e.g., on the supposition of $\{w_1, w_2\}$, the agent ought to be believe $\{w_1\}$, and on the supposition of $\{w_3, w_4, w_6\}$, the agent ought to be believe $\{w_3, w_4\}$.

However, if one switches to, e.g., $r = \frac{3}{4}$, then $X \subseteq W$ is to believed (unconditionally) by the agent if and only if $\{w_1, \ldots, w_5\} \subseteq X$. That is: The logically strongest believed proposition with respect to P and $r = \frac{3}{4}$ is $\{w_1, \ldots, w_5\}$. It is no longer so that on the supposition of $\{w_1, w_2\}$ the agent ought to be believe $\{w_1\}$, but, e.g., on the supposition of $\{w_3, w_4, w_6\}$, the agent still ought to be believe $\{w_3, w_4\}$. And conditional on $\{w_5, w_6, w_7, w_8\}$ the agent should believe $\{w_5\}$.

We will be able to derive all of these claims from the theory that will be developed in Sections 2 and 3. And we will take up this little toy example throughout the paper.

Let me conclude this introduction with a couple of remarks on some closely related theories in the existing literature on this topic.

Amongst the philosophical theories, the theory that comes closest to the theory that is developed below is Skyrms [38,39], who investigates the notion of objective chance and its applications in terms of a more fundamental concept: the *resiliency* of the (subjective) probability of a statement. In the simplest case—Skyrms' theory is actually much more general—the degree of resiliency of a (non-probabilistic) statement A is the infimum of the set of conditional probabilities of the form $P(A|B)$ where 'B' varies over all statements that are consistent with A in a given language. As we will see in Subsection 2.4, the concept that will be fundamental for our own theory is a categorical notion of stability with respect to a probability measure P and a threshold r: a proposition X will be defined P-stable^r if $P(X|Y) > r$ for all Y which are consistent with X and which have positive probability. Clearly, the two concepts are closely related, even when the underlying aims of the two theories differ: Skyrms' to explicate objective chance, mine to explicate belief. Indeed, it will follow from the postulates of the theory to be stated below that the logically strongest proposition that is believed by a perfectly rational agent must be P-stable^r. While the results below are new, in particular the two representation theorems, some of Skyrms' results do overlap with ours; e.g., note (2) on pp. 712f of [38], in which Skyrms is dealing with *acceptance*, and his theorem on p. 713 go some way towards Theorem 3 below.

Example 2 *(The example above and P-stability^r).* For P as before and $r = \frac{1}{2}$, it will turn out that the non-empty P-stable$^{\frac{1}{2}}$ sets will be:

$$\{w_1\},\ \{w_1, w_2\},\ \{w_1, \ldots, w_4\},\ \{w_1, \ldots, w_5\},\ \{w_1, \ldots, w_6\},\ \{w_1, \ldots, w_7\},\ \{w_1, \ldots, w_8\}$$

only the last two of which have probability 1. For instance, $\{w_1\}$ is P-stable$^{\frac{1}{2}}$ since for every Y with $Y \cap \{w_1\} \neq \varnothing$ and $P(Y) > 0$, it holds that $P(\{w_1\} \mid Y) > \frac{1}{2}$.

With $r = \frac{3}{4}$ the non-empty P-stable$^{\frac{3}{4}}$ sets will be:

$$\{w_1, \ldots, w_5\},\ \{w_1, \ldots, w_6\},\ \{w_1, \ldots, w_7\},\ \{w_1, \ldots, w_8\}$$

For instance, $\{w_1, \ldots, w_5\}$ is P-stable$^{\frac{3}{4}}$ since for every Y with $Y \cap \{w_1, \ldots, w_5\} \neq \varnothing$ and $P(Y) > 0$, it holds that $P(\{w_1, \ldots, w_5\}|Y) > \frac{3}{4}$.

The logically strongest believed propositions with respect to P and the two thresholds $r = \frac{1}{2}$ and $r = \frac{3}{4}$ were $\{w_1\}$ and $\{w_1, \ldots, w_5\}$, respectively, and indeed these are both P-stable^r sets for the thresholds in question.

It should not come as a big surprise either that *belief* will be tied to stable probabilities, even independently of any considerations on logical closure conditions or the like: for belief as a disposition to act over a period of time requires stability. If it were all too easy to get rid of the belief in X by conditionalizing on potential new evidence—maybe evidence that would even be consistent with everything that one believes—then one would hardly ever be in a position to carry out an action successfully on the basis of that belief. This is because just by acquiring such evidence rationally in the course of that action, one would be at permanent risk of losing the very mental state on which that action was partially based, that is, the belief in X. At the same, of course, it should not become impossible to revise one's beliefs by updating on new evidence either, but fortunately this will not be the case as long as one's logically strongest believed proposition has a probability of less than 1. In the example above: for $r = \frac{1}{2}$, the agent believes all propositions that are entailed by (supersets of) $\{w_1\}$ which has a probability less than 1; if a new piece if evidence E, say, $\{w_2, w_3\}$ comes along, then on the quantitative belief side P will be conditionalized on $\{w_2, w_3\}$, and on the qualitative belief side the new logically strongest believed proposition after updating on E will be $\{w_2\}$—since the agent holds the conditional belief in $\{w_2\}$ given $\{w_2, w_3\}$ before the update—as will follow from the theory that we are going to develop below. Hence, by updating on E, the belief in $\{w_1\}$ will be retracted even though $\{w_1\}$ is P-stable^r.

In the relevant literature in computer science, Snow's [42–44] account of *atomic bound systems* and Benferhat et al.'s [5] similar account of *big-stepped probabilities* actually deal with a special case of the theory below. Both of them consider special probability measures P on a finite set of atoms for which the following *atomic bound condition* holds: there is a strict total order $<$ over the atoms so that the probability of an atom a is greater than the sum of probabilities of all atoms b that lie above a in that strict total order. For those probability measures induce in an obvious sense nonmonotonic inference relations that satisfy well-known attractive logical closure conditions (that are known to be equivalent to those of AGM belief revision operators and which thus correspond to our postulates below on conditional belief). Indeed, for all sets X, Y of atoms, if X is non-empty, the uniquely determined $<$-least member of X is a member of Y if and only if $P(Y|X) > \frac{1}{2}$. Probability measures that satisfy the atomic bound condition can be generated by distributing probabilities over a given strict total order in exponential steps: e.g., if n is the number of atoms, the $<$-greatest atom receives probability $\frac{1}{2^n-1}$, the next one below gets probability $\frac{2}{2^n-1}$, the next one probability $\frac{4}{2^n-1}$, and so on, to the $<$-least atom overall, which receives the largest probability, $\frac{2^{n-1}}{2^n-1}$.

Accordingly, our second representation theorem for conditional belief—Theorem 13 in Section 3—will show that, given a probability measure P, our postulates on conditional belief will always determine uniquely a total pre-order or a "sphere system" of worlds which will be given by a sequence of P-stable^r sets: for one can show that every sequence of P-stable^r sets is a well-ordering with respect to the subset relation, and so every world can be assigned an ordinal rank according to its first "appearance" in a P-stable^r set in that well-ordering. Furthermore, in the finite case, the probability of a singleton set of a world w will be greater than the sum of the probabilities of the singleton sets of all worlds w' that lie above w in that ordinal ranking, exactly as in the atomic bound condition above. So Snow's and Benferhat et al.'s theory is a special case of ours in the sense that if P is one of these "big-stepped" probabilities, then their theory and ours coincide (for, in our case, a threshold of $r = \frac{1}{2}$); and for such measures, it will always hold that Y is believed on the condition of X if and only if the conditional probability of Y given X exceed $\frac{1}{2}$.

The essential difference between our approach and Snow's and Benferhat et al.'s is that we will be able to develop the theory below *for all probability measures P whatsoever*, and not just for special ones. This will be possible because the theory will not determine total orders from probability measures and thresholds but total *pre-orders* or *rankings* of worlds that allow for ties between worlds, that is, where different worlds whose singletons have different probabilities may nevertheless have the same rank. On probability measure that are not "big-stepped" in the sense from before it will still always be the case that *if* Y is believed on the condition of X *then* the conditional probability of Y given X exceeds $\frac{1}{2}$, but the converse will not necessarily be satisfied anymore.

For instance:

Example 3 *(The example above and total pre-orders of worlds).* For P as before and $r = \frac{1}{2}$, it will turn out that the ranking in question is given by:

$$w_1 < w_2 < w_3, w_4 < w_5 < w_6 < w_7 < w_8$$

(where later in the paper we will actually omit a world such as w_8 from any such ranking because of the probability of its singleton set being zero). So here one can observe a tie between the worlds w_3 and w_4. The rank of a world is determined by means of the least P-stable^r set in which it occurs: e.g., the rank of w_1 is 0 because, as we mentioned before, $\{w_1\}$ is the least non-empty P-stable$^{\frac{1}{2}}$ set; accordingly, the rank of w_2 is 1, the rank of w_3 is equal to the rank of w_4, that is, 2, and so forth.

And for $r = \frac{3}{4}$ we will have:

$$w_1, w_2, w_3, w_4, w_5 < w_6 < w_7 < w_8$$

Note that in the case $r = \frac{1}{2}$, $P(\{w_3\})$ is greater than $P(\{w_5\}) + P(\{w_6\}) + P(\{w_7\}) + P(\{w_8\})$, and also $P(\{w_4\})$ is greater than $P(\{w_5\}) + P(\{w_6\}) + P(\{w_7\}) + P(\{w_8\})$, but neither $P(\{w_3\}) > P(\{w_4\}) + P(\{w_5\}) + P(\{w_6\}) + P(\{w_7\}) + P(\{w_8\})$ nor $P(\{w_4\}) > P(\{w_3\}) + P(\{w_5\}) + P(\{w_6\}) + P(\{w_7\}) + P(\{w_8\})$ does hold. That is exactly why none of w_3 and w_4 happens to be ranked below the other. Subsection 2.6 will make precise why and how P-stable^r sets determine total pre-orders of worlds, and Subsection 2.5 will present the simple algorithm by which P-stable^r sets can be computed efficiently in terms of inequalities between probabilities. None of our results on P-stability^r nor the representation theorems below are contained in Snow's or Benferhat et al.'s work.

There are also related theories and notions of *robustness* or stability in other areas, such as in statistics (which partially inspired Skyrms' theory) or in game theory; but to the best of our knowledge, all of them differ from what will follow below. E.g., in game theory, the concept of *strong belief* (see [4]) is nothing but certainty of a hypothesis under all histories consistent with the hypothesis (this is spelled out in the context of primitive conditional probability measures that allows for conditionalization on zero sets, which makes for an important formal difference); so an agent strongly believes X if he assign probability 1 to X being the case at the beginning of a game, and continues to do so no matter how the game develops as long as X is not falsified by the evidence, where the evidence may even have probability 0. Our probabilistic explication of belief will be similar to that, but it will not presuppose certainty—just high probability.

2. The reduction of belief I: Absolute belief and restricted conditional belief

The goal of this section and the subsequent one is to enumerate, and study, a couple of postulates on quantitative and qualitative beliefs: first taken separately—what we called the 'logics of belief' in the introductory section—and then concerning their interaction—what we called 'bridge principles' before. We will assume that we are dealing with a fictional epistemic agent ag, who, at some point of time t, has belief states of both the qualitative (Bel) and the quantitative (P) type available so that these states obey the given postulates. Other than that, the specifics of t and ag will be irrelevant, which is why we are going to suppress reference to t and ag in the postulates themselves and only refer to them in the informal explanation and justification of the postulates. The terms 'P' and 'Bel' that will occur in these postulates should be thought of as primitive or undefined at first, with each postulate expressing a constraint either on the reference of 'P' or on the reference of 'Bel' or on the references of 'P' and 'Bel' simultaneously. Even though initially we will present these constraints in the form of postulates or axioms, it will turn out that they will be strong enough to constrain qualitative belief so that the concept of qualitative belief ends up being definable explicitly just on the basis of 'P' (and a threshold numeral and mathematical terms). When we state the theorems from which this will follow, 'P' and 'Bel' will become variables, so that we will able to say: For all P, Bel, it holds that P and Bel satisfy such-and-such if and only if such-and-such is the case. Accordingly, in our ultimate definition of belief simpliciter itself, 'P' will be a variable again, and 'Bel' will be a variable the reference of which is defined on the basis of 'P' (and a threshold numeral and mathematical vocabulary). We will keep using the same symbols 'P' and 'Bel' for all of these purposes, but their methodological status should always become clear from the context.

As far as the postulates themselves are concerned, we have to choose a formal background framework that will allow us to move back and forth smoothly between qualitative and quantitative postulates on belief. For that purpose, we make the following decisions: (1) On both sides, we opt for *propositions* or *sets of possible worlds* as the contents of belief rather than sentences or other syntactic items; and (2) we will assimilate the syntactic format of qualitative belief ascriptions to that of quantitative belief ascriptions. Because of (1), when we will say '$P(X) = y$' or '$Bel(X)$', it is really a proposition X that is assigned probability y or that is believed simpliciter. The logical postulates that are going to govern what will be expressed by '$Bel(X)$' will thus be *propositional* counterparts of the logical axioms and rules for doxastic logic that we sketched in the previous section; accordingly, the axioms of probability will be reformulated below for propositions, that is, for sets (as it is standard in mathematical probability theory), and not for formulas (as it was the case in the previous section). In view of (2), we are going to ascribe conditional beliefs to our agent ag by means of statements of the form '$Bel(Y|X)$' that express that the agent believes Y given X, which will mimic in a straightforward manner the usual way of ascribing conditional probabilities by means of expressions of the form '$P(Y|X) = y$'. The logic of conditional beliefs on the qualitative side will still turn out to be that of AGM belief revision, but the usual AGM postulates will be reformulated in terms of the binary 'Bel' symbol instead of the revision operator symbol '$*$' that was used in the last section.

2.1. Probabilistic postulates

We start with the postulates on quantitative belief. Consider our epistemic agent ag, which we will keep fixed throughout the article. Let W be a (non-empty) set of possible worlds; in the intended interpretation, W is the set of worlds which might in principle become doxastically possible from the viewpoint of ag at some point of time, that is, which ag may come to regard as a possible way the actual world might be. Now fix a point of time t and assume that at t our agent ag is capable in principle of entertaining all and only propositions (sets of worlds) in a class $\mathfrak{A}$ of subsets of W, where $\mathfrak{A}$ is formally a σ-algebra over W, that is: W and $\varnothing$ are members of $\mathfrak{A}$; if $X \in \mathfrak{A}$ then the relative complement of X with respect to W, $W \setminus X$, is also a member of $\mathfrak{A}$; for $X, Y \in \mathfrak{A}$, $X \cup Y \in \mathfrak{A}$; and finally if all of $X_1, X_2, \ldots, X_n, \ldots$ are members of $\mathfrak{A}$, then $\bigcup_{n\in\mathbb{N}} X_n \in \mathfrak{A}$. It follows that $\mathfrak{A}$ is closed under countable intersections, too. $\mathfrak{A}$ is not required to coincide with the power set algebra on W, that is, $\mathfrak{A}$ might not count certain subsets of W as propositions at all.

We will extend the standard logical terminology that is normally defined only for formulas or sentences to the propositions in $\mathfrak{A}$: so when we speak of a proposition as a logical truth we actually have in mind the unique proposition W, when we say that a proposition is consistent we mean that it is non-empty, and one proposition X logically entailing another one Y just coincides with X being a subset of Y, that is, every world that satisfies X (is a member of X) also satisfies Y (is a member of Y). When we refer to the negation of a proposition X we mean its complement relative to W (and we will denote it by '$\neg X$'), the conjunction of two propositions is of course their intersection, and so on. We shall speak of conjunctions and disjunctions of propositions even in cases of *infinite* intersections or unions of propositions.

Finally, let P be ag's degree-of-belief function at time t. Following the Bayesian take on quantitative belief, we postulate:

P1 (Probability)

P is a probability measure on $\mathfrak{A}$, that is, P has the following properties:

$P: \mathfrak{A} \to [0,1]$; $P(W) = 1$; P is finitely additive: if X_1, X_2 are pairwise disjoint members of $\mathfrak{A}$, then $P(X_1 \cup X_2) = P(X_1) + P(X_2)$.

Conditional probabilities are introduced by: $P(Y|X) = \frac{P(Y\cap X)}{P(X)}$ whenever $P(X) > 0$.

In particular, the conditional probability $P(Y|W)$ is nothing but the absolute probability of Y again.

As far as this familiar treatment of conditional probabilities in terms of the ratio formula for absolute or unconditional probabilities is concerned, we should stress that, e.g., the elegant theory of primitive conditional probability measures (or Popper functions, cf. [35]) would allow $P(Y|X)$ to be defined and non-trivial even when $P(X) = P(X|W) = 0$ (that is, as we will sometimes say, when X is a *zero set* as being given by P). But that theory is still not accepted widely, and we want to avoid the impression that our own theory in this paper relies on Popper functions in any sense. Additionally, in many practically relevant situations in which only probability measures on finite spaces are needed, and where often there are no non-empty sets of probability zero at all—the corresponding worlds with zero probabilistic weight having been dropped from the start—the analysis of conditional probabilities in terms of Popper functions seems to be a bit remote from the much more mundane reality of real-world reasoning and epistemological thought experiments. This said, as will be mentioned then, some of the postulates further below might actually be dropped or simplified if we started from Popper functions on the probabilistic side.

To P1 we add:

P2 (Countable Additivity)

P is countably additive (σ-additive): if $X_1, X_2, \ldots, X_n, \ldots$ are pairwise disjoint members of $\mathfrak{A}$, then $P(\bigcup_{n\in\mathbb{N}} X_n) = \sum_{n=1}^{\infty} P(X_n)$.

Countable Additivity or σ-additivity is not uncontroversial even within the Bayesian camp itself, although in purely mathematical contexts, such as measure theory, σ-additivity is beyond doubt; we shall simply take it for granted now. For many practical purposes, $\mathfrak{A}$ may be taken to be finite anyway, and then σ-additivity reduces to finite additivity again which is uncontroversial for all Bayesians whatsoever. In our context, Countable Additivity serves just one purpose: it simplifies the theory. However, in future versions of the theory one might want to study belief instead under the mere assumption of finite additivity, that is, assuming just P1 but not P2.[24]

The reason why we formulate P1 and P2 separately is that some, but not all of our results later will depend on P2, so it is simply handy to have a label at hand for 'probability theory except for countably infinite additivity' (that is: 'P1').

Example 4 *(The example from Section 1 reconsidered).* Let $W = \{w_1, \ldots, w_8\}$ be a set of eight possible worlds. One might think of these eight possibilities to coincide with all Boolean combinations of three propositions A, B, C, so that, e.g., $\{w_1\}$ might be $A \wedge B \wedge \neg C$, whereas, say, $\{w_8\}$ might be $A \wedge B \wedge C$, and the like (see Fig. 1).

[24] Extending the theory in that direction is feasible: Dropping P2 may be seen to correspond, roughly, to what happens to David Lewis' "spheres semantics" of counterfactuals when the so-called Limit Assumption is dropped (to which Lewis himself does not subscribe, while others do; see [30]).

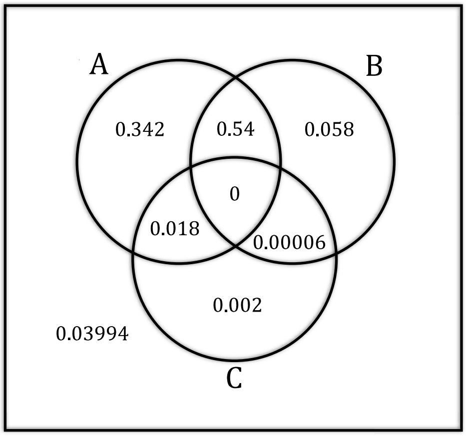

Fig. 1. An example measure.

Let P be the following probability measure on the power set algebra $\mathfrak{A}$ on W: $P(\{w_1\}) = 0.54$, $P(\{w_2\}) = 0.342$, $P(\{w_3\}) = 0.058$, $P(\{w_4\}) = 0.03994$, $P(\{w_5\}) = 0.018$, $P(\{w_6\}) = 0.002$, $P(\{w_7\}) = 0.00006$, $P(\{w_8\}) = 0$. By P1, P is determined uniquely in this way, since the probability of any set (say, $\{w_1, w_3, w_8\}$) is thus given by finite additivity (e.g., $P(\{w_1, w_3, w_8\}) = P(\{w_1\}) + P(\{w_3\}) + P(\{w_8\}) = 0.54 + 0.058 + 0 = 0.598$). P2 does not play a role here since W is finite.

Fig. 1 depicts what this probability space is like.

2.2. Belief postulates

Let us turn now to qualitative belief: Each belief simpliciter that ag holds at t is assumed to have a set in $\mathfrak{A}$ as its propositional content. In other words: quantitative and qualitative beliefs take their contents from the same space. Assume that by 'Bel' we are going to denote the class of propositions that our ideally rational agent believes to be true at time t. Instead of writing '$Y \in Bel$', we will rather say: $Bel(Y)$; and we call Bel our agent ag's *belief set* at time t. In line with elementary principles of doxastic or epistemic logic (which we sketched in the introductory section), Bel is assumed to satisfy the following postulates:

1. $Bel(W)$.
2. For all $Y, Z \in \mathfrak{A}$: if $Bel(Y)$ and $Y \subseteq Z$, then $Bel(Z)$.
3. For all $Y, Z \in \mathfrak{A}$: if $Bel(Y)$ and $Bel(Z)$, then $Bel(Y \cap Z)$.

So the agent's belief set is closed under logic. Actually, we are going to strengthen the principle on finite conjunctions of believed propositions to the case of the conjunction of *all* believed propositions *whatsoever*:

4. For $\mathcal{Y} = \{Y \in \mathfrak{A} \mid Bel(Y)\}$, $\bigcap \mathcal{Y}$ is a member of $\mathfrak{A}$, and $Bel(\bigcap \mathcal{Y})$.

This principle is a joint constraint on Bel and on the σ-algebra $\mathfrak{A}$. It does involve a good deal of idealization even in the case of a perfectly rational agent such as ag: for 4. is much like assuming that some system of arithmetic is closed under the infinitary ω-rule which may yield some not recursively axiomatizable set of theorems. On the other hand, if $\mathfrak{A}$ is finite, then 4. simply reduces to the case of finite conjunctions again. In any case, 4. has the following obvious and pleasant consequence: There is a *least set* (a strongest proposition) Y, such that $Bel(Y)$; Y is of course just the conjunction of all propositions believed by ag at t. We will denote this very proposition by: B_W. The main reason why we presuppose 4. is that it will enable us in this way to represent the sum of ag's beliefs as a unique proposition or a unique set of possible worlds. In the semantics of doxastic or epistemic logic, this set B_W would correspond to the set of accessible worlds from the viewpoint of the agent's current mindset. Accordingly, using the terminology that is quite common in areas such as belief revision or nonmonotonic reasoning, one might think of the members of B_W as being precisely the most plausible candidates for what the actual world might be like, if seen from the viewpoint of ag at time t.

Our postulate 4. imposes also another constraint on $\mathfrak{A}$: While it is not generally the case that the algebra $\mathfrak{A}$ contains arbitrary conjunctions of members of $\mathfrak{A}$, 4. together with our other postulates does imply that $\mathfrak{A}$ is closed under taking arbitrary *countable* conjunctions of *believed* propositions: for if all the members of any countable class of propositions are believed by ag at t, then their conjunction is a member of $\mathfrak{A}$ by $\mathfrak{A}$ being a σ-algebra, and the conjunction is a member of Bel by its being a superset of B_W and by 2. above.

There is another, independent, reason for assuming 4.: Because of lottery-type considerations, it is quite commonly thought that if the set of beliefs simpliciter is presupposed to be closed under finite conjunctions, then this prohibits any such probabilistic analysis of belief simpliciter from the start in which beliefs would be determined from "high" probabilities. However, as we will show, beliefs simpliciter can in fact be reduced to quantitative belief in such a manner *even though 4. expresses a much stronger form of closure under conjunction than just finite conjunction*. So we will not be accused of playing tricks by building up some kind of non-standard model for qualitative belief in which certain types of conjunction rules are applicable to certain sets of believed propositions but where other types of conjunction rules may not be applied. In a nutshell: 4. (together with 5. below) prohibits our agent from having anything like an ω-inconsistent set of beliefs.

Finally, we add

5. (Consistency)
   $\neg Bel(\varnothing)$.

as our sufficiently ideal agent ag does not believe a contradiction. One justification for this is the thought that even a realistically rational agent, when he is shown to believe a contradiction, will aim to change his mind; if ag's actual beliefs are considered to coincide with the (in principle) outcome of such a rationalization process, then 5. should be fine again. In any case, 5. should be acceptable for all perfectly rational agents (as we assume ag to be).

So much for belief if taken *unconditionally*. But we will require more than just qualitative belief in that sense—indeed, this will turn out to be the key move: Let us assume that ag also holds *conditional* beliefs, that is, beliefs conditional on certain propositions in $\mathfrak{A}$. We will interpret such conditional beliefs in suppositional terms: they are beliefs that the agent has *under the supposition of certain propositions*, where the only type of supposition that we will be concerned with will be supposition *as a matter of fact*, that is, suppositions which are usually expressed in the indicative, rather than the subjunctive mood: *Suppose that X is the case. Then I believe that Y is the case.* If X is any such "assumed" proposition, we take Bel_X to be the class of propositions that our ideally rational agent believes to be true at time t conditional on X; instead of writing '$Y \in Bel_X$', we will say somewhat more transparently: $Bel(Y|X)$. Accordingly, we call Bel_X our agent ag's *belief set conditional on X* at t, and we call any such class of propositions for whatever $X \in \mathfrak{A}$ a *conditional belief set* at t of our agent ag. In this extended context, Bel itself should now be regarded as a class of *ordered pairs* of members of $\mathfrak{A}$, rather than as a set of members of $\mathfrak{A}$ as before; instead of '$\langle Y, X\rangle \in Bel$' we may simply say again: $Bel(Y|X)$. And we may identify ag's belief set at t from before with one of ag's conditional belief sets at t: the class of propositions that ag believes to be true at t *conditional on the tautological proposition W*, that is, with the class Bel_W (just as $P(Y|W)$ coincided with the absolute or unconditional degree of belief in Y). Accordingly, we now refer to all and only the members Y of Bel_W as believed absolutely or unconditionally, and to Bel_W as the *absolute* or *unconditional belief set*. In the terms of belief revision theory (the principles of which we sketched in the introductory section): Bel_W corresponds to their set K; and $Bel(Y|X)$ corresponds to Y being a member of the belief set that results from revising Bel_W (or K) by X.

In the present section we will be interested only in conditional beliefs in Y given X where X is consistent with everything that the agent believes absolutely (or conditionally on W) at that time; equivalently: where X is consistent with B_W; equivalently: where $\neg X$ is not believed by the agent. (We will see later that all of these conditions are pairwise equivalent.) Ultimately, this will lead us to an explication of *absolute* or *unconditional* belief in terms of subjective probabilities, which will be the main focus of this section. In Section 3 we will then strengthen the postulates on conditional beliefs so that they will impose constraints even on beliefs conditional on propositions that contradict B_W, and ultimately we be able to state a corresponding explication of *conditional belief in general* in terms of probabilities. Even in the cases in which we will consider a belief suppositional on a proposition that is inconsistent with the agent's current absolute beliefs, as will be the case in Section 3, we will still regard the supposition in question to be a matter-of-fact supposition, in the sense that in natural language it would be expressed in the indicative rather than the subjunctive mood. As in: *I believe that John is not in the building. But suppose that he is in the building: then I believe he is in his office*. However, for now, we will focus solely on suppositions that are consistent with what the agent believes, a case that is much easier to handle than supposition in general. Accordingly, the postulates in the present section will be weaker than the postulates in Section 3, and it will be important to observe that even these weaker postulates will allow us to derive a substantial representation theorem for belief.

For every $X \in \mathfrak{A}$ that is consistent with what the agent believes, Bel_X is a set of the very same kind as the original unconditional or absolute belief set of propositions from above. And for every such $X \in \mathfrak{A}$, Bel_X will therefore be assumed to satisfy postulates of the very same type as suggested before for absolute beliefs. In order to make clear that we are dealing only with suppositions in this section that are consistent with what the agent believes unconditionally, we will add conditions of the form '$\neg Bel(\neg X|W)$' antecedently when we state these postulates:

B1 (Reflexivity)
If $\neg Bel(\neg X|W)$, then $Bel(X|X)$.

B2 (One Premise Logical Closure)
If $\neg Bel(\neg X|W)$, then for all $Y, Z \in \mathfrak{A}$: if $Bel(Y|X)$ and $Y \subseteq Z$, then $Bel(Z|X)$.

B3 (Finite Conjunction)
If $\neg Bel(\neg X|W)$, then for all $Y, Z \in \mathfrak{A}$: if $Bel(Y|X)$ and $Bel(Z|X)$, then $Bel(Y \cap Z|X)$.

B4 (General Conjunction)
If $\neg Bel(\neg X|W)$, then for $\mathcal{Y} = \{Y \in \mathfrak{A} \mid Bel(Y|X)\}$, $\bigcap \mathcal{Y}$ is a member of $\mathfrak{A}$, and $Bel(\bigcap \mathcal{Y}|X)$.

In spite of this generalization to certain kinds of conditional beliefs, we will still assume the Consistency postulate to hold only for absolute beliefs or beliefs conditional on W at this point (only in the next section this will be generalized, too). So just as in the case of 5. above, we only demand:

B5 (Consistency)
$\neg Bel(\varnothing|W)$.

Assuming B1 above is unproblematic at least under a suppositional reading of conditional belief: under the supposition of X, our ideally rational agent *ag* must hold X true at time t *in the context of the supposition of* X. This is much like in conditional proofs in which the statement that was first assumed may then also be concluded. B3 above is really redundant in view of B4, but we shall keep for the sake of continuity. As before, B4 now entails for every $X \in \mathfrak{A}$ for which $\neg Bel(\neg X|W)$ is the case that there is a *least set* (a strongest proposition) Y, such that $Bel(Y|X)$, which by B1 must be a subset of X. For any such given X, we will denote this very proposition by: B_X. For $X = W$, this is consistent with the notation 'B_W' as introduced before.

Clearly, we have then for all X with $\neg Bel(\neg X|W)$ and for $Y \in \mathfrak{A}$:

$$Bel(Y|X) \text{ if and only if } Y \supseteq B_X$$

from left to right by the definition of 'B_X', and from right to left by B2 and the definition of B_X again. So determining B_X suffices in order to pin down completely our agent's beliefs. Furthermore, it also follows that

$$Y \supseteq B_X \text{ if and only if } Bel(Y|B_X)$$

since if the left-hand side holds, then the right-hand side follows from B1 and B2, and if the right-hand side is the case then the left-hand side must be true by the definition of 'B_X' and the previous equivalence. So we find that actually for all $Y \in \mathfrak{A}$,

$$Bel(Y|X) \text{ if and only if } Bel(Y|B_X)$$

hence what is believed by *ag* conditional on X may always be determined just by means of considering all and only the members of $\mathfrak{A}$ which *ag* believes conditional on the subset B_X of X. Accordingly, $Bel(Z|W)$ iff $Z \supseteq B_W$ iff $Bel(Z|B_W)$. We will use these equivalences at several points, and when we do so we will not state this explicitly anymore.

By B5, W itself is such that $\neg Bel(\neg W|W)$ (since $\neg W = \varnothing$), hence all of B1–B4 apply to $X = W$ unconditionally, and consequently B_W must be non-empty. Using this and the first of the three equivalences above, we can thus derive

$$\neg Bel(\neg X|W) \text{ if and only if } X \cap B_W \neq \varnothing$$

For this reason, instead of qualifying the postulates by means of '$\neg Bel(\neg X|W)$', in the following we will do so by means of the simpler '$X \cap B_W \neq \varnothing$'.

So far there are no postulates on how belief sets conditional on *different propositions* relate to each other logically. At this point we demand just one such condition to be satisfied which corresponds to the standard AGM postulates K*3 and K*4 (Inclusion and Preservation) on belief revision taken together, where B_W is the propositional counterpart of AGM's syntactic belief set K, and where the revised belief set in the sense of AGM gets described in terms of conditional belief:

B6 (Expansion)
For all $Y \in \mathfrak{A}$ such that $Y \cap B_W \neq \varnothing$:
For all $Z \in \mathfrak{A}$, $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap B_W$.

(Note that given B5 and the equivalence of '$Y \cap B_W \neq \varnothing$' with '$\neg Bel(\neg Y|W)$', B6 actually entails B1–B4 from above.)

In words: if the proposition Y is consistent with B_W, then *ag* believes Z conditional on Y if and only if Z is entailed by the conjunction of Y with B_W. This is really just a postulate on "revision by expansion" in terms of propositional information that is consistent with the sum of what the agent believes; nothing is said at all about revision in terms of information that would contradict some of the agent's beliefs, which will be the topic of the next section.

As mentioned before, a principle like B6 corresponds to the AGM postulates on belief revision by propositions which are consistent with what the agent believes, and it can be justified in terms of total plausibility rankings of possible worlds: say that conditional beliefs express that the most plausible of their antecedent-worlds are among their consequent-worlds; then if some of the most plausible worlds overall are Y-worlds,

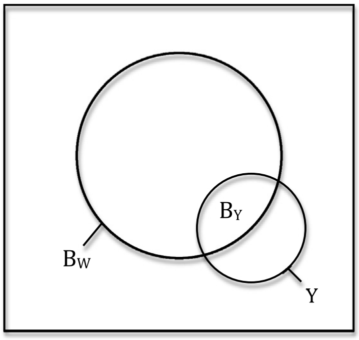

Fig. 2. The expansion operation.

these worlds must be precisely the most plausible Y-worlds, and therefore in that case the most plausible Y-worlds are Z-worlds if and only if all the most plausible worlds overall that are Y-worlds are Z-worlds.

From our previous considerations on $Bel(Z|W)$ being equivalent to $Z \supseteq B_W$, it is clear that this is an equivalent way of stating B6:

B6 (Expansion)
For all $Y \in \mathfrak{A}$, such that for all $Z \in \mathfrak{A}$, if $Bel(Z|W)$ then $Y \cap Z \neq \varnothing$:
For all $Z \in \mathfrak{A}$, $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap B_W$.

Supplying conditional belief with our intended suppositional interpretation again: If Y is consistent with everything *ag* believes absolutely, then supposing Y as a matter of fact amounts to nothing else than adding Y to one's stock of absolute beliefs, so that what the agent believes conditional on Y is precisely what the agent would believe absolutely if the strongest proposition that he believes were the intersection of Y and B_W. That is, we may reformulate B6 one more time in the form (see Fig. 2):

B6 (Expansion)
For all $Y \in \mathfrak{A}$ such that $Y \cap B_W \neq \varnothing$: $B_Y = Y \cap B_W$.

The superset claim that is implicit in this equality statement follows from the postulates above because $Bel(B_Y|Y)$ holds by the definition of 'B_Y' and then the original formulation of B6 above can be applied. The corresponding subset claim follows from the definition of B_Y again since $Bel(Y \cap B_W|Y)$ follows from the original version of B6. Similarly, the original version of B6 above can be derived from our last version of that principle and the other postulates that we assumed. It follows from the last formulation of B6 (trivially) that for all $Y \cap B_W \neq \varnothing$, B_Y is non-empty, simply because $B_Y = Y \cap B_W$.

Although AGM's K*3 and K*4 determine expansion to be the obvious qualitative counterpart of probabilistic conditionalization—supposing Y means restricting the space of doxastically possible worlds to the doxastically possible worlds in Y—they have not remained unchallenged, of course. One typical worry is that revising by some new piece of evidence or suppositional information Y might be thought to lead to more beliefs than what one would get deductively by adding Y to one's current beliefs, in view of possible *inductively* strong inferences that the presence of Y might warrant. One line of defense of AGM here is: if the agent's current beliefs are themselves already the result of the inductive expansion of what the agent is certain about, so that the agent's beliefs are really what he *expects* to be the case, then revising his beliefs by consistent information might reduce to merely adding it to his beliefs and closing off deductively. Another line of defense is: a postulate such as B6 might be true of belief simpliciter simply because without it qualitative belief would not have the simplifying power that is essential to it. Inductive reasoning in terms of *quantitative* belief is yet another matter, and the mentioned criticism of the conjunction of K*3 and K*4 might simply result from mixing up considerations on qualitative and quantitative belief.

Example 5 *(The example from Section 1 reconsidered).* By choosing $B_W \subseteq W$ in some arbitrary way, it becomes possible to determine *Bel* so that all of our postulates B1–B6 are satisfied.

For instance, let $B_W = \{w_1\}$, and turn B6 from above into a definition of *Bel* for all $Y \cap \{w_1\} \neq \varnothing$ and for all $Z \subseteq W$: $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap B_W$. It follows then that all of our belief postulates hold true, and e.g. $Bel(\{w_1\}|W)$, $Bel(\{w_1\}|\{w_1, w_2\})$ are the case.

But we might just as well choose $B_W = \{w_1, w_2\}$ and define for all $Y \cap \{w_1, w_2\} \neq \varnothing$ and for all $Z \subseteq W$: $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap B_W$. Then again all of our belief postulates are satisfied, and e.g. $Bel(\{w_1, w_2\}|W)$ (but not $Bel(\{w_1\}|W)$), not $Bel(\{w_1\}|\{w_1, w_2\})$ are the case.

And of course one might choose B_W in yet some different manner; any subset of W will do really, if *Bel* is then defined as before, as long as that subset is non-empty: for $B_W = \varnothing$ was excluded by B5.

2.3. The main bridge postulate

Finally, we turn to our promised necessary condition for having a belief—the left-to-right direction of the Lockean Thesis. Indeed we will formulate this condition more generally for all beliefs conditional on any proposition that is consistent with all the agent *ag* believes at *t*. This will make *ag*'s conditional degrees of beliefs at *t* and (some of) his conditional beliefs simpliciter at *t* compatible in a sense. The resulting bridge principle between qualitative and quantitative belief will involve a numerical constant 'r' again the value of which we will leave indeterminate at this point—just assume that r is some real number in the half-open interval $[0, 1)$. Note that the principle is not yet meant to give us anything like a definition of '*Bel*' (nor of any terms defined by means of '*Bel*', such as 'B_W') on the basis of 'P'. It only expresses a joint constraint on the references of '*Bel*' and 'P', that is, on our agent's *ag*'s actual conditional beliefs and his actual subjective probabilities at the given time t.

The principle says (where '*BP*' signals that this is for *B*elief and *P*robability simultaneously):

BP1^r (Likeliness)
For all $Y \in \mathfrak{A}$ such that $Y \cap B_W \neq \varnothing$ and $P(Y) > 0$:
For all $Z \in \mathfrak{A}$, if $Bel(Z|Y)$, then $P(Z|Y) > r$.

BP1^r is just the obvious generalization of the left-to-right direction of the Lockean Thesis to the case of beliefs conditional on propositions Y that are consistent with all of the agent's absolute beliefs. The antecedent clause '$P(Y) > 0$' in BP1^r is present to make sure that the conditional probability $P(Z|Y)$ is well-defined. By using W as the value of 'Y' and B_W as the value of 'Z' in BP1^r, and then applying the definition of B_W (a set which exists by B1–B4) and P1, it follows that $P(B_W|W) = P(B_W) > r$. Therefore, from the definition of 'B_W' and P1 again, having a subjective probability of more than r is a necessary condition for a proposition to be believed absolutely, although it will become clear later that this is not necessarily a sufficient condition.

r is a non-negative real number less than 1 which functions as a threshold value and which at this stage of our investigation can be chosen freely; so BP1^r can be viewed as a scheme the instances of which are given by determining the value of 'r'. BP1^r really says: conditional beliefs (with the relevant Ys) entail having corresponding conditional probabilities of more than r. One might wonder why there should be *one* such threshold r *for all* propositions Y and Z as stated in BP1^r at all, rather than having *for all* Y (or for all Y and Z) *some* threshold value that might depend on Y (or on Y and Z). But without any further qualification, this revised principle would be almost empty, because as long as for Y and Z it is the case that $P(Z|Y) > 0$, there will always be an r such that $P(Z|Y) > r$. In contrast, BP1^r postulates a conditional probabilistic boundary from below that is *uniform* for all conditional beliefs whatsoever—such an r should really derive from contextual considerations on the concept of belief itself rather than from contextual considerations on the contents of belief.[25]

For illustration, in BP1^r, think of r as being equal to $\frac{1}{2}$: If degrees of beliefs and beliefs simpliciter ought to be compatible in some sense at all, then the resulting BP1$^{\frac{1}{2}}$ is pretty much the weakest possible expression of any such compatibility that one could think of: if ag believes Z (conditional on one of Y's referred to above), then ag assigns an subjective probability to Z (conditional on Y) that exceeds the subjective probability that he assigns to the negation of Z (conditional on Y). If BP1$^{\frac{1}{2}}$ were invalidated, then there would be Z and Y, such that our agent ag believes Z conditional on Y, but where $P(Z|Y) \leqslant \frac{1}{2}$: if $P(Z|Y) < \frac{1}{2}$, then ag would be in a position in which he regarded $\neg Z$ as more likely than Z, conditional on Y, even though he believes Z, but not $\neg Z$, conditional on Y. On the other hand, if $P(Z|Y) = \frac{1}{2}$, then ag would be in a position in which he regarded $\neg Z$ as equally likely as Z, conditional on Y, even though he believes Z, but not $\neg Z$, conditional on Y. At least the former should be very difficult to accept—and the more difficult the lower the value of $P(Z|Y)$.

Instead of defending BP1$^{\frac{1}{2}}$ or any other particular instance of BP1^r at this point, we will simply move on, taking for granting one such BP1^r has been chosen. Within our theory, choosing $r = \frac{1}{2}$ will in fact be the right choice for the *least possible* threshold value that would give us an account of 'believing that'; but taking any greater threshold value less than 1 will be permissible, too. However, for weaker forms of subjective commitment than belief, such as 'suspecting that' or 'hypothesizing that', r might well be chosen to be less than $\frac{1}{2}$, and some of our formal results below will not depend on r being greater than or equal to $\frac{1}{2}$.

For the moment this exhausts our list of postulates (with two more auxiliary bridge principles to come later).

2.4. *P-stability and the first representation theorem*

For now let us pause as regards the introduction of postulates and let us instead focus on finding jointly necessary and sufficient conditions for all of our postulates up to this point to be satisfied. This will lead us to our first representation theorem: the theorem will thus characterize in transparent terms those pairs $\langle P, Bel\rangle$ whose coordinate entries jointly satisfy all of the postulates so far.

In order to formulate the theorem, we will need the following probabilistic concept which will turn out to be crucial for the whole theory:

Definition 1 *(P-stability^r).* Let P be a probability measure on a set algebra $\mathfrak{A}$ over W, let $0 \leqslant r < 1$. For all $X \in \mathfrak{A}$:

X is P-stable^r if and only if for all $Y \in \mathfrak{A}$ with $Y \cap X \neq \varnothing$ and $P(Y) > 0$: $P(X|Y) > r$.

If we think of $P(X|Y)$ as the degree of X under the supposition of Y, then a P-stable^r proposition X has the property that whatever proposition Y one supposes, as long as Y is consistent with X and probabilities conditional on Y are well-defined, it will be the case that the degree of X under the supposition of Y exceeds r. For such a non-empty P-stable^r set X, one of the Ys that we could choose above is of course the full set W of possible worlds: such a P-stable^r X must therefore have an absolute probability that is greater than r. What P-stability^r adds to this is that this will *remain* to be so as long as one supposes propositions that are consistent with X and on which conditionalization is defined at all. So a P-stable^r proposition has

[25] Remark: It would be possible to weaken '>' to '$\geqslant$' in BP1^r; in what follows, not much will depend on this, except that whenever we are going to use BP1^r with $r \geqslant \frac{1}{2}$ below, one would rather have to choose some $r' > \frac{1}{2}$ instead and then demand that '$\ldots P(Z|Y) \geqslant r'$' is the case.

a special *stability property*: it is characterized by its stably high probabilities under all suppositions of a particularly salient type. Trivially, the empty set is P-stable^r. W is P-stable^r, too, and more generally all propositions X in $\mathfrak{A}$ with probability $P(X) = 1$ are P-stable^r. More importantly, and perhaps surprisingly, as we shall see later in Section 2.10, there are in fact lots of probability measures for which there are lots of *non-trivial* P-stable^r propositions which have a probability *strictly between 0 and 1*.

A different way of thinking of P-stability^r is the following one. Let X be P-stable^r: For all Y being such that $Y \cap X \neq \varnothing$ and $P(Y) > 0$, it holds then that $P(X|Y) = \frac{P(X\cap Y)}{P(Y)} > r$, which is equivalent to: $P(X \cap Y) > r \cdot P(Y)$. But by P1 this is again equivalent to $P(X \cap Y) > r \cdot [P(X \cap Y) + P(\neg X \cap Y)]$, which yields $P(X \cap Y) > \frac{r}{1-r} \cdot P(\neg X \cap Y)$. By letting the value of 'Y' vary over all members of $\mathfrak{A}$ that have non-empty intersection with X as well as non-zero probability, we can take the value of '$X \cap Y$' to vary over precisely the non-empty subsets of X that are members of $\mathfrak{A}$, and the value of '$\neg X \cap Y$' to vary over precisely the subsets of $\neg X$ that are members of $\mathfrak{A}$ and which have positive probability (if there are any), when we give the following equivalent characterization of P-stability^r:

> X is P-stable^r if and only if for all $Y, Z \in \mathfrak{A}$, such that $Y \in \mathfrak{A}$, $Y \neq \varnothing$, $Y \subseteq X$, and where $Z \in \mathfrak{A}$, $Z \subseteq \neg X$, $P(Z) > 0$, it holds that:
>
> $$P(Y) > \frac{r}{1-r} \cdot P(Z)$$

In the special case in which $r = \frac{1}{2}$, the factor $\frac{r}{1-r}$ is just 1, and hence X is P-stable$^{\frac{1}{2}}$ if and only if the probability of any non-empty subset of X is greater than the probability of any non-zero subset of $\neg X$. So P-stability^r is also a *separation property* that divides the class of subpropositions of a proposition from the class of subpropositions of its negation in terms of their probabilities.

Here is a property of non-empty P-stable^r propositions X that we will need on various occasions, which is why it is worth stating it explicitly:

Observation 2. For all $X \in \mathfrak{A}$ with X non-empty and P-stable^r:

> If $P(X) < 1$, then there is no non-empty $Y \subseteq X$ with $Y \in \mathfrak{A}$ and $P(Y) = 0$.

For assume this is not so: then $Y \cup \neg X$ has non-empty intersection with X since Y has, and at the same time $P(Y \cup \neg X) > 0$ because $P(\neg X) > 0$. By X being P-stable^r, it would therefore have to be the case that $P(X|Y \cup \neg X) = \frac{P(X\cap Y)}{P(Y\cup\neg X)} > r$, in contradiction with $P(X \cap Y) \leqslant P(Y) = 0$. This also implies that non-empty propositions of probability 0 cannot be P-stable^r, or in other words: all non-empty P-stable^r propositions X have positive probability.

Using the new concept of P-stability^r, we can formulate the following first and rather simple representation theorem on belief (there will be another more intricate one in the next section which will extend the present one to conditional belief in general):

Theorem 3. *Let Bel be a class of ordered pairs of members of a σ-algebra $\mathfrak{A}$, let $P : \mathfrak{A} \to [0, 1]$, and let $0 \leqslant r < 1$. Then the following two statements are equivalent*:

I. *P and Bel satisfy* P1, B1–B6, *and* BP1^r.
II. *P satisfies* P1, *and there is an $X \in \mathfrak{A}$, such that X is a non-empty P-stable^r proposition, and*:
   – *For all $Y \in \mathfrak{A}$ such that $Y \cap X \neq \varnothing$, for all $Z \in \mathfrak{A}$*:

$$Bel(Z|Y) \textit{ if and only if } Z \supseteq Y \cap X$$

   (*and hence, $B_W = X$*).

Furthermore, if I is the case, then X in II is actually *uniquely determined*.

Proof. From I to II: P1 is satisfied by assumption. Now we let $X = B_W$, where B_W exists and has by definition the property of being the strongest believed proposition by B1–B4: First of all, as derived before by means of B5, B_W is non-empty; and B_W is P-stable^r: For let $Y \in \mathfrak{A}$ with $Y \cap B_W \neq \varnothing$, $P(Y) > 0$: since $B_W \supseteq Y \cap B_W$, it thus follows from B6 that $Bel(B_W|Y)$, which by BP1^r and $P(Y) > 0$ entails that $P(B_W|Y) > r$, which was to be shown. Secondly, let $Y \in \mathfrak{A}$ be such that $Y \cap B_W \neq \varnothing$, let $Z \in \mathfrak{A}$: then it holds that $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap B_W$ by B6, as intended. Finally, uniqueness: Assume $X' \in \mathfrak{A}$, X' is non-empty, P-stable^r, and for all $Y \in \mathfrak{A}$ with $Y \cap X' \neq \varnothing$, for all $Z \in \mathfrak{A}$, it holds that $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap X'$. But from the latter it follows that $Bel(B_W|W)$ if and only if $B_W \supseteq W \cap X' = X'$, and hence, with $Bel(B_W|W)$ from the definition of B_W, we may conclude $B_W \supseteq X'$. On the other hand, by choosing X' as the value of 'Z' and W as the value of 'Y', we have $Bel(X'|W)$ if and only if $X' \supseteq W \cap X'$, and thus $Bel(X'|W)$; but by the definition of B_W again this entails: $X' \supseteq B_W$. Therefore, $X' = B_W = X$.

From II to I: Suppose P satisfies P1, and there is an X, such that X and Bel have the required properties. Then, first of all, all the instances of B1–B5 for beliefs conditional on W are satisfied: for it holds that $W \cap X = X \neq \varnothing$ because X is non-empty by assumption, so $Bel(Z|W)$ if and only if $Z \supseteq W \cap X = X$, by assumption, therefore B5 is the case, and the instances of B1–B4 for beliefs conditional on W follow from the characterization of beliefs conditional on W in terms of supersets of X. Indeed, it follows: $B_W = X$. So, for arbitrary $Y \in \mathfrak{A}$, $\neg Bel(\neg Y|W)$ is really equivalent to $Y \cap X \neq \varnothing$, as we did already show after our introduction of B1–B5, and hence all instances of B1–B4 are satisfied by the assumed characterization of beliefs conditional on any Y with $Y \cap X \neq \varnothing$ in terms of supersets of $Y \cap X$. B6 holds trivially, by assumption and because of $B_W = X$. About BP1^r: Let $Y \cap X \neq \varnothing$ and $P(Y) > 0$. If $Bel(Z|Y)$, then by assumption $Z \supseteq Y \cap X$, hence $Z \cap Y \supseteq Y \cap X$, and by P1 it follows that $P(Z \cap Y) \geqslant P(Y \cap X)$. From X being P-stable^r and $Y \cap X \neq \varnothing$ and $P(Y) > 0$ we also have $P(X|Y) > r$. Taking these two facts together, and by the definition of conditional probability in P1, this implies $P(Z|Y) > r$, which we needed to show. □

Note that P2 (Countable Additivity) did not play any role in this; but of course P2 may be added to both sides of the proven equivalence with the resulting extended equivalence being satisfied.

This simple theorem will prove to be fundamental for all subsequent arguments in this paper. As we will see, what it says is: let P and r be fixed; then it does not really matter whether one deals with P, r, and some belief set Bel, such that all of our postulates (except maybe for P2) are satisfied by them, or whether one deals with P, r, and one "cooks" up Bel by choosing some non-empty P-stable^r set in $\mathfrak{A}$ to be the logically strongest believed proposition B_W.

We start by exploiting the theorem in a rather trivial fashion: Let us concentrate on its right-hand side, that is, condition II of Theorem 3. Disregarding for the moment any considerations on qualitative belief, let us just assume that we are given a probability P over a set algebra $\mathfrak{A}$ on W. We know already that one can in fact always find a non-empty set X, such that X is a P-stable^r proposition: just take any proposition with probability 1. In the simplest case: for now, take X to be W itself. $P(W) > 0$ and P-stability^r are then seen to be the case immediately. Now consider the very last equivalence clause of II and turn it into a (conditional) *definition* of $Bel(.|Y)$ for all the cases in which $Y \cap W = Y \neq \varnothing$: that is, for all $Z \in \mathfrak{A}$, define $Bel(Z|Y)$ to hold if and only if $Z \supseteq Y \cap W = Y$. In particular, $Bel(Z|W)$ holds then if and only if $Z \supseteq W$ which obviously is the case if and only if $Z = W$. $B_W = W$ follows, all the conditions in II of Theorem 3 are satisfied, and thus by Theorem 3 all of our postulates from above must be true as well. What this shows is that given a probability measure, it is always possible to define belief simpliciter in a way such that all of our postulates turn out to be the case. What would be believed thereby by our agent would be maximally cautious: having such beliefs, *ag* would believe absolutely just W, and therefore trivially every absolute belief would have probability 1. Accordingly, he would believe conditionally on the respective Ys from above just what is logically entailed by them, that is, all supersets of Y.

As we pointed out in the introduction, this is *not* in general a satisfying explication of belief. But we actually find a much more general pattern to be emerging from Theorem 3: Let P and r be given again as before. Now choose *any* non-empty P-stable^r proposition X, and define conditional belief in all cases in which $Y \cap X \neq \varnothing$ by: $Bel(Z|Y)$ if and only if $Z \supseteq Y \cap X$. Then $B_W = X$ follows again, and all of our postulates hold by Theorem 3—including B3 (Finite Conjunction) and B4 (General Conjunction)—even though it might well be that $P(X) < 1$ and *hence even though there might be beliefs whose propositional contents have a subjective probability of less than 1 as being given by P*. Such beliefs are not maximally cautious anymore—exactly as it is the case for most of the beliefs of any real-world human agent *ag*. Of course this does not mean that according to the current construction *all* believed propositions would have to be assigned a probability of less than 1: When $P(X) < 1$, although there will always be believed propositions that have a probability of precisely 1—for instance, W—it only follows that there exist believed propositions that have a probability of less than 1; X itself is an example. And every believed proposition must then have a probability that lies somewhere in the closed interval $[P(X), 1]$, so that $P(X)$ becomes a lower threshold value; furthermore, since X is P-stable^r, $P(X)$ itself is strictly bounded from below by r.[26]; for instance, r is not necessarily given by the result of applying P to some distinguished proposition or the like—it could be chosen before any considerations on P or Bel commence.

Note that the following is *not* entailed by Theorem 3: *Every* believed proposition is to be P-stable^r. In fact, given our postulates, P-stability^r is only required for *precisely one* believed proposition: the logically strongest one that is believed at all, or equivalently the conjunction of all believed propositions. Accordingly, in view of Theorem 3: what P-stability^r demands is that if some proposition is consistent with *everything* that the agent believes, then conditionalizing on such a proposition should not drag down the probability of any believed proposition below the threshold r.[27]

Since P-stable^r propositions play such a distinguished role in this, the question arises: How difficult is it to determine whether a proposition is a non-empty P-stable^r set? To this question we will turn in the next subsection.

2.5. Computing P-stable^r sets

At least in the case where W is finite, it turns out not to be difficult at all to determine all and only P-stable^r sets: For let W be finite. Without loss of generality, let $\mathfrak{A}$ be the power set algebra on W, and let P be a probability measure on $\mathfrak{A}$. We have seen already that all sets with probability 1 are P-stable^r and that the empty set is trivially P-stable^r. So let us focus just on how to generate all non-empty P-stable^r sets X that are non-trivial, that is, which have a probability of less than 1. As we observed before (Observation 2),

---

[26] In fact, one can show more: Let $r \geqslant \frac{1}{2}$; in a situation in which X is non-empty and P-stable^r, with $P(X) < 1$, and where X is the strongest proposition B_W believed, it holds: a proposition Y has a probability in the interval $[P(X), 1]$ if and only if Y is believed. The right-to-direction is obvious, since if Y is believed, then $Y \supseteq B_W = X$, and the rest follows by the monotonicity property of probability (as entailed by P1). And from left to right: Assume $P(Y) \geqslant P(X)$ but Y is not believed; then $Y \not\supseteq B_W = X$, that is, $\neg Y \cap X$ is non-empty. Thus, $[\neg Y \cap X] \cup \neg X$ has non-empty intersection with X and its probability is greater than 0, because $1 > P(X) = 1 - P(\neg X)$ and so $P(\neg X) > 0$ (by P1). But from X being P-stable^r it would then follow that $P(X|[\neg Y \cap X] \cup \neg X) > r \geqslant \frac{1}{2}$, that is, by P1 again, $P(\neg Y \cap X) > \frac{P([\neg Y \cap X] \cup \neg X)}{2} = \frac{P(\neg Y \cap X)}{2} + \frac{P(\neg X)}{2}$, and hence $P(\neg Y \cap X) > P(\neg X)$. However, by $P(Y) \geqslant P(X)$ and P1, $P(\neg X) \geqslant P(\neg Y)$. So we would get $P(\neg Y \cap X) > P(\neg Y)$, which contradicts P1. So Y is in fact believed.

Summing up: In such circumstances, one will always get an instance of the *full* Lockean Thesis satisfied. This said, the threshold value in such an instance of the Lockean Thesis would not necessarily be r but rather $P(X)$, with a $\geqslant$-version of the inequality in the Lockean Thesis. Therefore the choice of the threshold would follow to be sensitive to the properties of the given probability measure P—one could not choose the threshold independently of the measure P anymore. We will not be able to go into this in any further detail; a separate paper will be devoted to a corresponding alternative derivation of the present theory from a version of the Lockean Thesis in which the threshold value will *depend* on the probability measure in question. In the current paper, we will keep presupposing just the left-to-right direction of the Lockean Thesis, in which the threshold value of 'r' can be chosen *independently* of what P is like.

[27] If the right-hand side of our definition of P-stability were strengthened in the way that in order for X to be P-stable^r its probability would have to remain above r conditional on every Y consistent *with at least some superset of* X, then all P-stable^r sets would have to be of probability 1: for certainly W is a superset of X. But that would seem to be overly cautious, for precisely the reason for which the 'Probability 1 Proposal' in the introduction was found wanting.

such sets do not contain any non-empty subsets of probability 0, which in the present context means that if $w \in X$, $P(\{w\}) > 0$.

For any such X, as we have also proven in Subsection 2.4, it follows that X is P-stable^r if and only if for all $Y, Z \in \mathfrak{A}$, such that Y is a subset of X (and hence, in the present case, $P(Y) > 0$) and where Z is a subset of $\neg X$, it holds that $P(Y) > \frac{r}{1-r} \cdot P(Z)$. Therefore, in order to check for P-stability^r in the current context, it suffices to consider just sets Y and Z which have the required properties and for which $P(Y)$ is minimal and $P(Z)$ is maximal. In other words, we have for all non-empty X with $P(X) < 1$:

$$X \text{ is } P\text{-stable}^r \text{ if and only if for all } w \text{ in } X \text{ it holds that } P(\{w\}) > \frac{r}{1-r} \cdot P(W \setminus X).$$

In particular, for $r = \frac{1}{2}$, this is (where X was assumed non-empty and $P(X)$ was assumed to be less than 1):

$$X \text{ is } P\text{-stable}^{\frac{1}{2}} \text{ if and only if for all } w \text{ in } X \text{ it holds that } P(\{w\}) > P(W \setminus X).$$

Thus it turns out to be very simply to decide whether a set X is P-stable^r and even more so if it is P-stable$^{\frac{1}{2}}$.

From this it is easy to see that in the present finite context there is also an efficient procedure that computes all non-empty non-trivial P-stable^r subsets X of W. We only give a sketch for the case $r = \frac{1}{2}$: Since such sets X do not have singleton subsets of probability 0, let us also disregard all worlds whose singletons are zero sets. Assume that after dropping all worlds of zero probabilistic mass, there are exactly n members of W left, and $P(\{w_1\}), P(\{w_2\}), \ldots, P(\{w_n\})$ are already in (not necessarily strictly) decreasing order. If $P(\{w_1\}) > P(\{w_2\}) + \cdots + P(\{w_n\})$ then $\{w_1\}$ is the first P-stable$^{\frac{1}{2}}$ set determined, and one moves on in the list $P(\{w_2\}), \ldots, P(\{w_n\})$, keeping $\{w_1\}$ "on stock" (the next possible candidate for a P-stable$^{\frac{1}{2}}$ set would be $\{w_1, w_2\}$), and hence now one considers $P(\{w_1\}), P(\{w_2\})$. However, if $P(\{w_1\}) \leqslant P(\{w_2\}) + \cdots + P(\{w_n\})$ then consider $P(\{w_1\}), P(\{w_2\})$ from the start: If both of them are greater than $P(\{w_3\}) + \cdots + P(\{w_n\})$ then $\{w_1, w_2\}$ is the next P-stable$^{\frac{1}{2}}$ set, and one moves on to the list $P(\{w_3\}), \ldots, P(\{w_n\})$. If either of them is less than or equal to $P(\{w_3\}) + \cdots + P(\{w_n\})$ then consider $P(\{w_1\})$, $P(\{w_2\})$, $P(\{w_3\})$: And so forth, until the set $\{w_1, w_2, \ldots, w_n\}$ has been reached which then coincides with the least subset of W of probability 1, that is, the smallest set that is but a trivial instance of P-stability$^{\frac{1}{2}}$. This recursive procedure yields precisely all non-empty non-trivial P-stable$^{\frac{1}{2}}$ sets, and it does so in polynomial time complexity. (The same procedure can be applied in cases in which W is countably infinite and $\mathfrak{A}$ is the full power set algebra on W, but then, of course, the procedure will not necessarily terminate in finite time.)

What Theorem 3 gives us, therefore, is not just a representation result, but even, in the case of a given finite probability space for P, an *efficient construction procedure* for all classes *Bel*, so that *Bel* and the given P together satisfy all of our postulates. P2 still has not played a role so far.

Example 6 *(The example from Section 1 reconsidered).* By means of the algorithm from above it is easy to compute all non-empty and non-trivial P-stable$^{\frac{1}{2}}$ sets: $\{w_1\}, \{w_1, w_2\}, \{w_1, \ldots, w_4\}, \{w_1, \ldots, w_5\}$, $\{w_1, \ldots, w_6\}$.

$\{w_1, \ldots, w_7\}$ is also P-stable$^{\frac{1}{2}}$, but it is trivial in the sense of having probability 1. And we leave out w_8 from the start, since $P(\{w_8\}) = 0$.

Accordingly, for instance, $P(\{w_1\})$ is greater than the sum of all other probabilities of singletons, $P(\{w_2\})$ is greater than $P(\{w_3\}) + \cdots + P(\{w_7\})$, both $P(\{w_3\})$ and $P(\{w_4\})$ are greater than $P(\{w_5\}) + \cdots + P(\{w_7\})$, and so on. But it is neither the case that $P(\{w_3\})$ is greater than $P(\{w_4\}) + \cdots + P(\{w_7\})$, nor is it the case that $P(\{w_4\})$ is greater than $P(\{w_3\}) + P(\{w_5\}) + \cdots + P(\{w_7\})$, which is why neither $\{w_1, w_2, w_3\}$ nor $\{w_1, w_2, w_4\}$ are P-stable$^{\frac{1}{2}}$.

On the other hand, for $r = \frac{3}{4}$ the corresponding non-empty and non-trivial P-stable$^{\frac{3}{4}}$ sets are: $\{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}$.

With BP1^r in place, as stated by Theorem 3, it is no longer possible to determine our agent's beliefs by choosing $B_W \subseteq W$ *arbitrarily*: for B_W must be a non-empty P-stable^r set. While this is the case for our previous choices $B_W = \{w_1\}$ (for $r = \frac{1}{2}$, but not for $r = \frac{3}{4}$), and $B_W = \{w_1, \ldots, w_5\}$ (for both $r = \frac{1}{2}$ and $r = \frac{3}{4}$), it would not be possible to choose, e.g., $B_W = \{w_1, w_3\}$ or $B_W = \{w_3, w_7, w_8\}$ for the same purpose, whatever the value of 'r', since these sets can be shown not to be P-stable^r for any r.

2.6. Some properties of P-stable^r sets

But Theorem 3 does more: it also shows that whatever our agent ag's *actual* probability measure P at t and his *actual* class *Bel* of conditionally-believed pairs of propositions at t are like, as long as they satisfy our postulates from above (with a given threshold r), then it must be possible to partially reconstruct *Bel* by means of some P-stable^r proposition X as explained before, where: X is then simply identical to B_W; and by 'partially' we mean that it would only be possible to reconstruct beliefs that are conditional on propositions Y which are consistent with $X = B_W$. For this is just the left-to-right direction of the theorem. Hence, if we had any additional means of determining from P the very P-stable^r proposition X that (partially) determines ag's actual belief set *Bel*, we could define explicitly the set of all pairs $\langle Z, Y\rangle$ in that set *Bel* for which $Y \cap X \neq \varnothing$ holds by means of that proposition X and hence, ultimately, by means of P. Amongst those conditional beliefs, in particular, we would find all of ag's absolute beliefs, and therefore the set of absolutely believed propositions could be defined explicitly in terms of P (and r).

So are we in a position to identify the P-stable^r proposition X that gives us ag's actual beliefs, simply by being handed only ag's subjective probability measure? That is the first open question that we will deal with in the remainder of this section. The other open question is: What should r be like in our postulate BP1^r above?

In order to address these two questions appropriately, we first need another theorem which summarizes two important properties of (non-empty and non-trivial) P-stable^r sets[28]:

Theorem 4. *Let $P : \mathfrak{A} \to [0,1]$ such that* P1 *is satisfied. Let $\frac{1}{2} \leqslant r < 1$. Then the following is the case*:

III. *For all $X, X' \in \mathfrak{A}$: If X and X' are P-stable^r and at least one of $P(X)$ and $P(X')$ is less than* 1, *then either $X \subseteq X'$ or $X' \subseteq X$ (or both).*

IV. *If P also satisfies* P2, *then there is no infinitely descending chain of sets in $\mathfrak{A}$ that are all subsets of some P-stable^r set X_0 of $\mathfrak{A}$ with probability less than* 1, *that is, there is no countably infinite sequence*

$$X_0 \supsetneq X_1 \supsetneq X_2 \supsetneq \cdots$$

of sets in $\mathfrak{A}$ (and hence no infinite sequence of such sets in general), such that X_0 is P-stable^r, $P(X_0) < 1$, and each X_n is a proper superset of X_{n+1} (hence $P(X_n) < 1$ for all $n \geqslant 0$).

A fortiori, *given* P2, *there is no infinitely descending chain of P-stable^r sets in $\mathfrak{A}$ with probability less than* 1.

[28] These properties correspond to some of the properties of so-called belief cores in Arló-Costa and Parikh [3] (see also [48] and [2]), which are special sets of absolute probability 1 in a context where probabilities are determined by a primitive conditional probability measure or a Popper function. In fact, this is not a mere coincidence: Once our theory has been generalized in Section 3 to the case of arbitrary suppositions, one can show that by defining $P^*(Y|X) = P(Y|B_X)$, a Popper function P^* is defined from our P and *Bel* (and given r); and by this definition our P-stable^r sets are transformed into belief cores as being given by P^*. One can also show that every Popper function on a finite space can be represented in this way in terms of an absolute probability measure and *Bel* (for any given threshold r).

Proof.

- Ad III: First of all, let X and X' be P-stable^r, and $P(X) = 1$, $P(X') < 1$: as observed before, there is then no non-empty subset Y of X', such that $P(Y) = 0$. But if $X' \cap \neg X$ were non-empty, then there would have to be such a subset of X'. Therefore, $X' \cap \neg X$ is empty, and thus $X' \subseteq X$. The case for X and X' taken the other way around is analogous.
  So we can concentrate on the remaining logically possible case. Assume for contradiction that there are P-stable^r members X, X' of $\mathfrak{A}$, such that $P(X), P(X') < 1$, and neither $X \subseteq X'$ nor $X' \subseteq X$. Therefore, both $X \cap \neg X'$ and $X' \cap \neg X$ are non-empty, and they must have positive probability since as we know P-stable^r propositions with probability less than 1 do not have non-empty zero sets as subsets. We observe that $P(X|(X\cap\neg X')\cup\neg X)$ is greater than r by X being P-stable^r, $(X\cap\neg X')\cup\neg X \supseteq (X\cap\neg X')$ having non-empty intersection with X, and the probability of $(X\cap\neg X')\cup\neg X \supseteq (X\cap\neg X')$ being positive. The same must hold, *mutatis mutandis*, for $P(X'|(X' \cap \neg X) \cup \neg X')$. So we have
  $$P\big(X\big|\big(X \cap \neg X'\big) \cup \neg X\big) > r \geqslant \frac{1}{2}$$
  and
  $$P\big(X'\big|\big(X' \cap \neg X\big) \cup \neg X'\big) > r \geqslant \frac{1}{2}$$
  where $r \geqslant \frac{1}{2}$ by assumption.
  Next we show that
  $$P\big(X \cap \neg X'\big) > P(\neg X)$$
  For suppose otherwise, that is (i) $P(X \cap \neg X') \leqslant P(\neg X)$: Since by P1 (and $P((X \cap \neg X') \cup \neg X) > 0$), it must be the case that $P(X \cap \neg X'|(X \cap \neg X') \cup \neg X) + P(\neg X|(X \cap \neg X') \cup \neg X) = 1$, and since we know from before that the second summand must be strictly less than $\frac{1}{2}$, the first summand has to strictly exceed $\frac{1}{2}$. On the other hand, it also follows that: $\frac{1}{2} > P(\neg X|(X \cap \neg X') \cup \neg X) = \frac{P(\neg X)}{P((X\cap\neg X')\cup\neg X)} \geqslant \frac{P(X\cap\neg X')}{P((X\cap\neg X')\cup\neg X)} = P(X \cap \neg X'|(X \cap \neg X') \cup \neg X)$, the inequality following from our supposition (i); but this contradicts our conclusion from before that $P(X \cap \neg X'|(X \cap \neg X') \cup \neg X)$ exceeds $\frac{1}{2}$.
  Analogously, it follows also that
  $$P\big(X' \cap \neg X\big) > P\big(\neg X'\big)$$
  Finally, from this (and P1) we can derive: $P(X\cap\neg X') > P(\neg X) \geqslant P(X'\cap\neg X) > P(\neg X') \geqslant P(X\cap\neg X')$, which is a contradiction.
- Ad IV: Assume for contradiction that there is a sequence $X_0 \supsetneq X_1 \supsetneq X_2 \supsetneq \cdots$ of sets in $\mathfrak{A}$ with probability less 1, with X_0 being P-stable^r as described. None of these sets can be empty, or otherwise the subset relationships holding between them could not be proper. Now let $A_i = X_i \setminus X_{i+1}$ for all $i \geqslant 0$, and let $B = \bigcup_{i=0}^{\infty} A_i$. Note that every A_i is non-empty and indeed has positive probability, since as observed before P-stable^r sets with probability less than 1 do not contain subsets with probability 0. Furthermore, for $i \neq j$, $A_i \cap A_j = \varnothing$. Since $\mathfrak{A}$ is a σ-algebra, B is in fact a member of $\mathfrak{A}$. By P2, the sequence $(P(A_i))$ must converge to 0 for $i \to \infty$, for otherwise $P(B) = P(\bigcup_{i=0}^{\infty} A_i) = \sum_{i=0}^{\infty} P(A_i)$ would not be a real number. Because, by assumption, X_0 has a probability of less than 1, $P(\neg X_0)$ is a real number greater that 0. It follows that the sequence of real numbers of the form $\frac{P(A_i)}{P(A_i)+P(\neg X_0)} = \frac{P(X_0\cap(A_i\cup\neg X_0))}{P(A_i\cup\neg X_0)} = P(X_0|A_i\cup\neg X_0)$ also converges to 0 for $i \to \infty$, where for every i, $(A_i\cup\neg X_0)\cap X_0 \neq \varnothing$ and $P(A_i \cup \neg X_0) > 0$. But this contradicts X_0 being P-stable^r. □

We may draw three conclusions from this. First of all, in view of IV, P-stable^r sets of probability less than 1 have a certain kind of groundedness property: they do not allow for infinitely descending sequences of subsets. For the infinite case, this result finally required P2 (Countable Additivity) as a presupposition. (IV above together with Observation 2 also implies that every P-stable^r set is *finite*.[29])

Secondly, in view of III and IV taken together, the whole class of P-stable^r propositions X in $\mathfrak{A}$ with $P(X) < 1$ is well-ordered with respect to the subset relation. In particular, if there is a non-empty P-stable^r proposition with probability less than 1 at all, there must also be a *least* non-empty P-stable^r proposition with probability less than 1. A different way of expressing this fact is: If we only look at non-empty P-stable^r propositions with a probability of less than 1, we find that they constitute a so-called *sphere system* that satisfies the Limit Assumption (by well-orderedness) in the sense of Lewis [30].

Finally, by III (and P1), we immediately have the following, which we put on record for further use:

Observation 5. If $\frac{1}{2} \leqslant r < 1$, then: All P-stable^r propositions X in $\mathfrak{A}$ with $P(X) < 1$ are subsets of all propositions in $\mathfrak{A}$ of probability 1. (And we know already that the latter are all P-stable^r.)

For given P (and given $\mathfrak{A}$ and W), such that P satisfies P1–P2, and for given $r \in [0,1)$, let us denote the class of all non-empty P-stable^r propositions X with $P(X) < 1$ (that is, which are non-trivial) by: $\mathcal{X}_P^r$. What Theorem 4 says is that $\langle \mathcal{X}_P^r, \subseteq \rangle$ is then a well-order. So by standard set-theoretic arguments, there is a bijective and order-preserving mapping from $\mathcal{X}_P^r$ into a uniquely determined ordinal β_P^r, where β_P^r is a well-order of ordinals with respect to the subset relation (which is of course also the order relation for ordinals); β_P^r measures the length of the well-ordering $\langle \mathcal{X}_P^r, \subseteq \rangle$. Hence, $\mathcal{X}_P^r$ is identical to a strictly increasing sequence of the form $(X_\alpha^r)_{\alpha < \beta_P^r}$. X_0^r is then the least non-empty P-stable^r proposition in $\mathfrak{A}$ with probability less than 1, if there is one at all. If there are none, then β_P^r is simply equal to 0 (that is, the ordinal $\varnothing$). Each world in the union of all X_α^r can be assigned a uniquely determined ordinal rank: the least ordinal α, such that $w \in X_\alpha^r$. So we find that the non-empty P-stable^r propositions X with probability less than 1, if they exist, determine ordinal rankings of those possible worlds that are members of at least one of them.

Additionally, in the case of a finite probability space, one could also assign all worlds w that are not a member of any non-trivial P-stable^r set, but whose singletons $\{w\}$ still have positive probability, the ordinal rank β_P^r. In that case, only worlds whose singletons $\{w\}$ are zero sets would not be assigned any ordinal rank at all.

Furthermore, by P1–P2, Theorem 4, and the fact that no non-empty P-stable^r set of probability less than 1 has a non-empty subset of probability zero, each such X in $\mathcal{X}_P^r$ determines a number $P(X) \in (r, 1]$ and no non-empty P-stable^r proposition of probability less than 1 other than X could determine the same number $P(X)$. By P1, the greater the set X with respect to the subset relation, the greater its probability $P(X)$. So we have: for $\alpha < \alpha' < \beta_P^r$ it holds that $r < P(X_\alpha^r) < P(X_{\alpha'}^r)$. It follows that there is also a bijective and order-preserving mapping from the set of probabilities of the members of $\mathcal{X}_P^r$ to the set of ordinals below β_P^r (that is, to the set β_P^r). The situation is summarized by Fig. 3.

From this we can determine a boundary for β_P^r in case $r \geqslant \frac{1}{2}$:

Observation 6. Let P be a countably additive probability measure on a σ-algebra $\mathfrak{A}$ over W. Let $\frac{1}{2} \leqslant r < 1$. The ordinal β_P^r (see above) is either finite or equal to ω.

(Hence, the class $\mathcal{X}_P^r$ of all non-empty P-stable^r propositions X with probability less than 1 is countable.)

Proof. Assume for contradiction that $\beta_P^r \geqslant \omega + 1$: then there certainly exist non-empty P-stable^r propositions X with probability less than 1. Now, for X_α^r as defined above, and for all $0 \leqslant n < \omega$, let $Y_n = X_{n+1}^r \setminus X_n^r$, and let $Z_n = \bigcup_{m \geqslant n} Y_m^r$. We know that for all n it holds that $Z_n \in \mathfrak{A}$; by Theorem 4 and the definition

[29] We are grateful to Martin Krombholz for calling our attention to this.

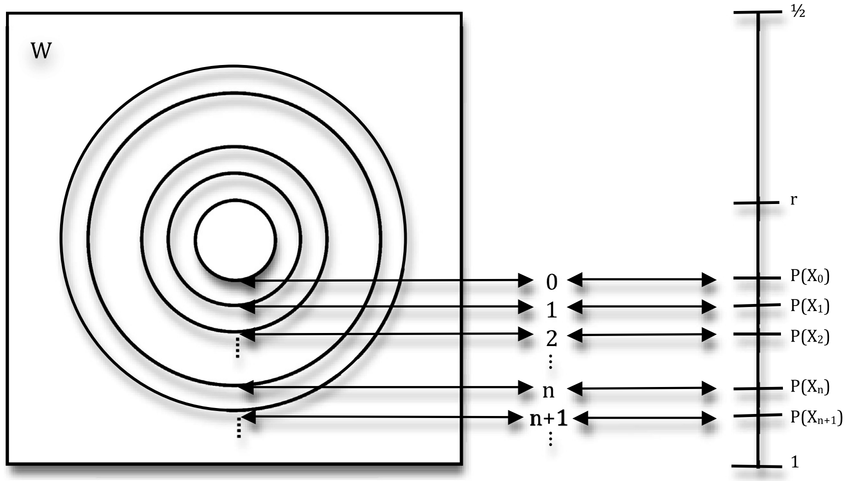

Fig. 3. P-stable sets for $r \geqslant \frac{1}{2}$.

of 'X^r_α' it is the case that for all n, $Z_n \subseteq X^r_\omega$, by assumption we have $P(X^r_\omega) < 1$, and furthermore for all n, $P(Z_n) < 1$, and the sequence (Z_n) is strictly monotonically decreasing. So there is a sequence $X^r_\omega \supsetneq Z_0 \supsetneq Z_1 \supsetneq \ldots$ of sets in $\mathfrak{A}$ with probability less 1, with X^r_ω being P-stable^r, in contradiction with IV of Theorem 4. $\square$

We also find that, given that P is countably additive, if there are countably infinitely many non-empty P-stable^r propositions X with probability less than 1—so that's the case in which $\beta^r_P = \omega$—then the union of all non-empty P-stable^r propositions X with probability less than 1 is itself P-stable^r, non-empty, and it must have probability 1. For: The countable union $\bigcup_{\alpha<\omega} X^r_\alpha$ is a member of our σ-algebra $\mathfrak{A}$. If $Y \cap \bigcup_{\alpha<\omega} X^r_\alpha \neq \varnothing$ for $Y \in \mathfrak{A}$ with $P(Y) > 0$, then there must be an X^r_α with $\alpha < \omega$, such that $Y \cap X^r_\alpha \neq \varnothing$. Because X^r_α is P-stable^r, it follows that $P(X^r_\alpha|Y) > r$. But by P1, $P(\bigcup_{\alpha<\omega} X^r_\alpha|Y) \geqslant P(X^r_\alpha|Y)$, hence $P(\bigcup_{\alpha<\omega} X^r_\alpha|Y) > r$. So $\bigcup_{\alpha<\omega} X^r_\alpha$ is P-stable^r (and non-empty, of course). If $P(\bigcup_{\alpha<\omega} X^r_\alpha)$ were less than 1, then β^r_P would have to be at least of the order type $\omega + 1$, as $\bigcup_{\alpha<\omega} X^r_\alpha$ is certainly a proper superset of any single set X^r_α by what we have seen before; but $\beta^r_P \geqslant \omega + 1$ was ruled out by Observation 6. So $P(\bigcup_{\alpha<\omega} X^r_\alpha) = 1$.

Since, as was noted in terms of Observation 2, no non-empty P-stable^r proposition X with probability less than 1 contains a non-empty zero set as a subset, no union of such sets X could do so either. So in the case in which β^r_P is infinite, the union of all non-empty P-stable^r propositions with probability less than 1 would then not just have probability 1, it would also not have any non-empty zero subset. But that means that union must then be the least set in $\mathfrak{A}$ with probability 1, which thus must exist in that case.

Summing up: With the X^r_α sets for $\alpha < \beta^r_P$ being all and only the non-empty P-stable^r sets of probability less than 1, if $\beta^r_P = \omega$, then it holds that $P(\bigcup_{\alpha<\omega} X^r_\alpha) = 1$ and $\bigcup_{\alpha<\omega} X^r_\alpha$ is the least member in $\mathfrak{A}$ that has this property.

Example 7 *(The example from Section 1 reconsidered).* Clearly, the non-empty and non-trivial P-stable$^{\frac{1}{2}}$ sets are well-ordered with respect to $\subseteq$: $\{w_1\} \subseteq \{w_1, w_2\} \subseteq \{w_1, \ldots, w_4\} \subseteq \{w_1, \ldots, w_5\} \subseteq \{w_1, \ldots, w_6\}$.

Accordingly, for the non-empty and non-trivial P-stable$^{\frac{3}{4}}$ sets: $\{w_1, \ldots, w_5\} \subseteq \{w_1, \ldots, w_6\}$.

Fig. 4 depicts the ordinal ranks (natural numbers) of all worlds of positive probabilistic mass for the case $r = \frac{1}{2}$.

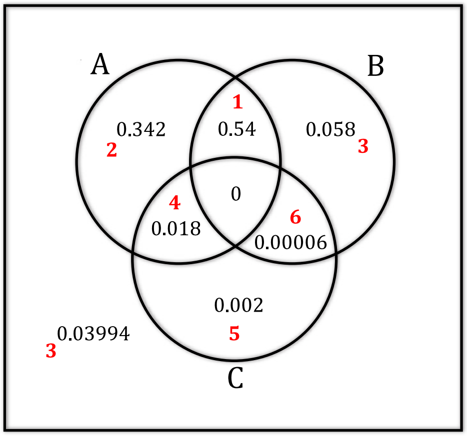

Fig. 4. Ordinal ranks of worlds for the example measure.

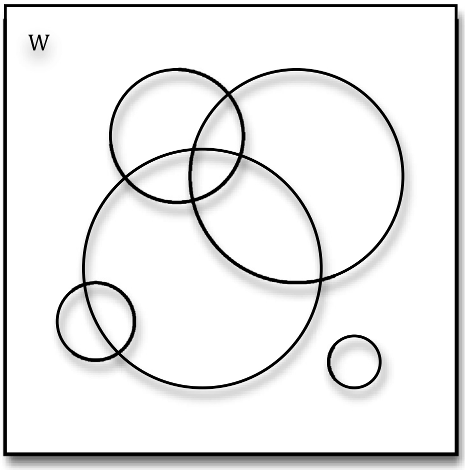

Fig. 5. P-stable sets for $r < \frac{1}{2}$.

2.7. *Choosing the value of 'r'*

Now back to our remaining open questions. Let us start with: what should one choose as the threshold value of 'r' in BP1^r?

For the proof of III in Theorem 4 it was crucial that $r \geqslant \frac{1}{2}$. Indeed, one can show by means of examples that if $r < \frac{1}{2}$ then III can be invalidated: it is possible then that there are P-stable^r members X, X' of $\mathfrak{A}$, such that neither $X \subseteq X'$ nor $X' \subseteq X$. In fact, it is even possible that there are non-empty P-stable^r members X, X' of $\mathfrak{A}$, such that $X \cap X' = \varnothing$. This means: if our agent *ag*'s probability measure P is held fixed for the moment, and if $r < \frac{1}{2}$, then depending on what P is like, our postulates P1–P2, B1–B6, and BP1^r might allow for two classes *Bel* such that all of these postulates are satisfied for each of them (by Theorem 3) and yet some absolute beliefs according to the one class *Bel* contradict some absolute beliefs according to the other class *Bel*, although both are based on one and the same subjective probability measure P. Fig. 5 visualizes a situation like that, where the circles represent different sets B_W (non-empty P-stable^r sets) for different sets *Bel*.

It seems advisable then to demand that $r \geqslant \frac{1}{2}$ in order to be able to derive as a *law* that a situation like that cannot occur. For if P is fixed, then one might think that our postulates should suffice to rule out systems of qualitative belief that contradict each other. As van Fraassen [48], p. 350, puts it, the assumed role of full belief is "to form a single, unequivocally endorsed picture of what things are like": If $r \geqslant \frac{1}{2}$, then while Theorem 4 does not yet pin down such a "single, unequivocally endorsed picture of what things are like", at least the linearity condition III guarantees the following: given P, if X and X' are possible choices of strongest possible believed propositions B_W such that P1–P2, B1–B6, and BP1^r are satisfied, that is, by Theorem 3, if X and X' are both non-empty P-stable^r members of $\mathfrak{A}$, then either everything that *ag* believes absolutely according to $B_W = X$ would also be believed if it were the case that $B_W = X'$ or *vice versa*. Combining this with what we said about $r < \frac{1}{2}$ initially when we introduced BP1^r above—that is, that if an agent believes a proposition it is hardly acceptable for the same agent to assign to that proposition a probability that is less than or equal to the probability of its negation—we do have a plausible case against choosing r in that way.

But that does not mean that choosing $r < \frac{1}{2}$ would not be an attractive choice if '*Bel*' were taken to express not belief but some weaker epistemic attitude. Let us assume that '$Bel(Y|X)$' is not meant to express belief but rather something like: *Supposing X, proposition Y is an interesting or salient thesis that is to be investigated further*. We might then be interested in determining *systems* of such interesting hypotheses (given the supposition X) which "cohere" with each other logically in the way believed propositions do; so B1–B5 would then be plausible again. But BP1^r with $r \geqslant \frac{1}{2}$ would be too much to ask of propositions that are merely considered interesting or salient: however, demanding of them that their probabilities are at least above some minimal threshold $r < \frac{1}{2}$ would still be reasonable, for if a proposition is way too unlikely, it is probably not worth further scrutiny either. In this way, all of our postulates so far would be plausible, but the threshold value of 'r' would *not* be assumed to be greater than or equal to $\frac{1}{2}$ in such a case. We leave the detailed treatment of this alternative application of the present theory to a different occasion.

In the present context, we conclude $r \geqslant \frac{1}{2}$. Apart from presupposing $r \geqslant \frac{1}{2}$, is it possible to exclude other possible values of 'r'? Before we answer this question, the following elementary observation informs us about some of the consequences that the answer will have:

Observation 7. Let P be a probability measure on an algebra $\mathfrak{A}$ over W. Let $X \in \mathfrak{A}$, and assume that $\frac{1}{2} \leqslant r < s < 1$. Then it holds:

If X is P-stable^s, X is P-stable^r.

Proof. If X is P-stable^s, then for all $Y \in \mathfrak{A}$ with $Y \cap X \neq \varnothing$, $P(X|Y) > s$. But then it also holds for all $Y \in \mathfrak{A}$ with $Y \cap X \neq \varnothing$ that $P(X|Y) > r$, since $r < s$ by assumption, so X is P-stable^r as well. □

Example 8 *(The example from Section 1 reconsidered).* In line with Observation 7, all of the P-stable$^{\frac{3}{4}}$ sets are also P-stable$^{\frac{1}{2}}$ sets. This is clear for the empty set and for the two sets of probability 1, that is, $\{w_1,\dots,w_7\}$ and W, but it also holds for the non-empty and non-trivial cases of P-stability$^{\frac{3}{4}}$: both of $\{w_1,\dots,w_5\}$, $\{w_1,\dots,w_6\}$ are to be found amongst the P-stable$^{\frac{1}{2}}$ sets $\{w_1\}$, $\{w_1,w_2\}$, $\{w_1,\dots,w_4\}$, $\{w_1,\dots,w_5\}$, $\{w_1,\dots,w_6\}$.

Hence, the smaller the threshold value r, the more inclusive is the class of P-stable^r sets that it determines. What this tells us, in conjunction with our previous results, is that if we choose r *minimally* such that $\frac{1}{2} \leqslant r < 1$, that is, if we choose $r = \frac{1}{2}$, then we do not exclude any of the logically possible options for B_W.

Should our agent *ag* exclude some of them? By determining the value of 'r', one lays down how brave a belief can be maximally, or how cautious a belief needs to be minimally, in order not to cease to count as a belief from the start. Choosing $r = \frac{1}{2}$ is the bravest possible option. At the same time, beliefs in this sense would not necessarily seem *too* brave: after all, with P being given, *Bel* would still be constrained by BP1$^{\frac{1}{2}}$. In particular, if Y is believed in this sense, then the subjective probability of Y would have to be greater than $\frac{1}{2}$. And of course *Bel* would have to satisfy all of the standard logical properties of belief simpliciter, as expressed by B1–B6. Indeed, for many purposes this might well be the right choice. But then again, maybe, for other purposes a more cautious notion of belief is asked for, which would correspond to choosing a value for 'r' that is greater than $\frac{1}{2}$. In many cases, the value of 'r' might be determined by the epistemic and pragmatic context in which our agent ag is about to reason and act, and different contexts might ask for different values of 'r'. In yet other cases, the value of 'r' might only be determined vaguely; and so on. And all of these options would still be covered by what we call pre-theoretically 'belief'. We suggest therefore to explicate belief *conditional on any given threshold value* $r \geqslant \frac{1}{2}$, without making any particular choice of the value of 'r' mandatory.

But we should add that the closer r get to 1, the more extreme the distribution of probabilities over singletons of worlds must be in order for there to be a P-stable^r set of probability less than 1, as follows immediately from our characterization of such P-stable^r sets in Section 2.5, that is: for all w in X it holds that $P(\{w\}) > \frac{r}{1-r} \cdot P(W \setminus X)$. For instance, for $r = 0.9$, if X is P-stable^r and has a probability of less than 1, then for all w in X it holds that $P(\{w\}) > 9 \cdot P(W \setminus X)$. For instance, in the example of Section 1, only $\{w_1, \ldots, w_6\}$ satisfies this condition, and that it satisfies the condition relies on the fact that the probability of $\{w_7\}$ is much smaller than the probability of each of $\{w_1\}, \ldots, \{w_6\}$.

2.8. Two final bridge principles

With one of our two open questions settled (or rather dismissed), we are in the position to address the other one: Can we always identify the P-stable^r proposition X that yields our agent's ag's actual beliefs, if we are given only ag's actual subjective probability measure P (and a threshold value r)? We need two more auxiliary bridge postulates before we answer this question.

First of all, we must correct an asymmetry between quantitative and qualitative belief conditional on propositions of probability zero. By the standard definition of conditional probabilities in terms of the ratio formula (as given by P1), the conditional probability of X given Y is undefined if $P(Y) = 0$, whatever X is like. However, qualitative conditional beliefs given any such Y are well-defined, at least as long as we are dealing with a Y that is consistent with everything the agent believes absolutely (at t). For instance, B1 entails then that $Bel(Y|Y)$ holds. In order to bring in line conditional beliefs with degrees of belief conditional on a zero set, we are going to postulate that supposing any such proposition of probability zero amounts to believing a contradiction on the qualitative side. In a nutshell: Where $P(.|Y)$ is undefined, $Bel(.|Y)$ will get trivialized.

This said, rather than *restricting qualitative belief* in such a way, it would actually be more attractive to *liberate quantitative probability* such that the (non-trivial) conditionalization on zero sets becomes possible: that is, as mentioned before, one might want to use a Popper function P from the start. But then again the current theory has the advantage of relying just on the much more common absolute probability measures, and since the theory is not particularly affected by using BP2 below as an additional assumption, we shall stick to conditional belief being constrained as expressed by BP2. So BP2 is acceptable really just for the sake of simplicity. At least, if P is regular (strictly coherent), that is, every non-empty proposition in $\mathfrak{A}$ has positive probability, then BP2 is of course superfluous, and for many practically relevant scenarios, Regularity may indeed be taken for granted or otherwise W could be redefined by dropping all worlds whose singleton sets have zero probability.

So we assume:

BP2 (Zero Supposition)

For all $Y \in \mathfrak{A}$: If $P(Y) = 0$, then $Y \cap B_W = \varnothing$.

Equivalently: if $P(Y) = 1$ then $B_W \subseteq Y$.

Note that we might just as well have included our usual constraint '$Y \cap B_W \neq \varnothing$' antecedently, which would simply have led to an equivalent, but clumsy, reformulation: 'If $Y \cap B_W \neq \varnothing$ and $P(Y) = 0$, then

$Y \cap B_W = \varnothing$' is logically equivalent to BP2 above, because it is logically equivalent to 'If $P(Y) = 0$ then, if $Y \cap B_W \neq \varnothing$ then $Y \cap B_W = \varnothing$'.

In the case where $P(B_W) < 1$, we could actually derive the corresponding instance of BP2 from our previous postulates: For if $P(Y) = 0$ and $Y \cap B_W \neq \varnothing$, then B_W would have a non-empty subset $Y \cap B_W$ of probability $P(Y \cap B_W) \leqslant P(Y)$ (by P1), that is, B_W would have a non-empty zero subset; but that could not be the case by Observation 2.

However, in the case where $P(B_W) = 1$, BP2 does have bite. In particular, if B_W has probability 1 itself, then by the equivalent reformulation of BP2 above, B_W must be the least proposition in $\mathfrak{A}$ with probability 1. Hence, in the case where $P(B_W) = 1$, BP2 forces our probability space to be such that there exists a least set with probability 1, which is a non-trivial constraint that is not satisfied generally.[30] Furthermore, even though we know that every proposition of probability 1 is P-stable^r, we have just seen that with BP2 only one such proposition of probability 1 can be of the form B_W: the *least* proposition of probability 1. And so, if $P(B_W) = 1$, B_W cannot have non-empty subsets of probability 0 either, just as it was seen to be the case before when $P(B_W) < 1$.

Now we are in the position to answer our remaining question from above affirmatively, by identifying the P-stable^r proposition X that yields *ag*'s actual beliefs if given just *ag*'s subjective probability measure P (and a threshold value r). This will be achieved by means of our final auxiliary bridge postulate.

Let us start by reconsidering our agent's actual probability measure P. *Bel* is of course *ag*'s actual belief set; but for the moment let us disregard *Bel* and consider instead *all possible sets Bel′* which, together with P, satisfy all of our postulates so far. (*Bel* is just one amongst these sets *Bel′*.) Then it is easy to see that the larger such a set *Bel′* is with respect to the absolute beliefs that it determines, the more instances of the *right-to-left* direction of the Lockean Thesis for absolute beliefs (as stated in Section 1) are satisfied by it: For let *Bel′*, *Bel″* be two candidate sets that both satisfy, if taken together, respectively, with our given P, all postulates so far: By BP1^r, we have for both them the left-to-right direction of the Lockean Thesis for absolute beliefs, that is,

For all $Z \in \mathfrak{A}$, if $Bel'(Z|W)$, then $P(Z|W) > r$.
For all $Z \in \mathfrak{A}$, if $Bel''(Z|W)$, then $P(Z|W) > r$.

Now assume for all $Z \in \mathfrak{A}$, if $Bel'(Z|W)$ then $Bel''(Z|W)$ (or equivalently $B'_W \supseteq B''_W$): Then every instance of the *right-to-left* direction of the Lockean Thesis for absolute beliefs that holds for *Bel′* is also satisfied by *Bel″*, because, when

$$\text{If } P(Z|W) > r, \text{ then } Bel'(Z|W)$$

holds for $Z \in \mathfrak{A}$, it follows then that also

$$\text{If } P(Z|W) > r, \text{ then } Bel''(Z|W)$$

holds. Moreover, if the set of absolute beliefs as determined by *Bel″* is actually *strictly* larger than the set of absolute beliefs as determined by *Bel′*, then there are also properly *more* instances of the right-to-left direction of the Lockean Thesis for absolute beliefs which are satisfied by *Bel″* than by *Bel′*: For then there is a Z, such that $Bel''(Z|W)$ but not $Bel'(Z|W)$; since $Bel''(Z|W)$, BP1^r entails that $P(Z|W) > r$; so

$$\text{If } P(Z|W) > r, \text{ then } Bel''(Z|W)$$

[30] The Lebesgue measure on the unit interval would be a counterexample.

is the case, while

$$\text{If } P(Z|W) > r, \text{ then } Bel'(Z|W)$$

is not.

So indeed we have: The larger a set Bel' is with respect to the absolute beliefs that is determines, the more instances of the right-to-left direction of the Lockean Thesis for absolute beliefs are satisfied by it. Equivalently (by B1–B4): *The smaller the corresponding logically strongest believed proposition B'_W is, the more instances of the right-to-left direction of the Lockean Thesis for absolute beliefs are satisfied by Bel'.*

Is there always a greatest such set Bel' and accordingly a least such set B'_W? Given our postulates, and with $\frac{1}{2} \leqslant r < 1$ (as was presupposed in order for 'Bel' to express *belief*), the answer turns out to be 'yes': From Theorem 3 we know that such sets B'_W must be non-empty and P-stable^r (or otherwise one of our postulates would be invalidated). Furthermore, from Theorem 4 we know that (i) if there are non-empty P-stable^r sets of probability less than 1 at all, then they are well-ordered with respect to $\subseteq$; and (ii) if there are no such non-empty P-stable^r sets of probability less than 1, then the non-empty P-stable^r sets are precisely the propositions of probability 1, hence $P(B'_W)$ must be 1, in which case we know from BP2, as shown above, that B'_W is then the least proposition of probability 1 and thus also the least P-stable^r set. In either case, *the least possible proposition of the form B'_W coincides with the least non-empty P-stable^r set*; and the latter is well-defined, as we have just shown. For the record:

Observation 8. From our postulates (and $\frac{1}{2} \leqslant r < 1$) it follows that there is a least non-empty P-stable^r set in $\mathfrak{A}$. If there are no non-empty P-stable^r sets of probability less than 1, then the least non-empty P-stable^r set is the least set in $\mathfrak{A}$ which has probability 1, which must exist then by our postulates.

Consequently, given our agent ag's actual probability measure P, amongst all sets Bel' that together with P satisfy all of our postulates so far, *there is always one that is greatest with respect to the absolute beliefs that it determines, namely, the belief set Bel' that is generated by the least P-stable^r set in $\mathfrak{A}$ which exists given our postulates on ag's actual P and Bel so far*. And this very belief set Bel' exemplifies as many of the right-to-left direction instances of the Lockean Thesis as possible, in the context where all of our postulates have been assumed to be satisfied.

As we explained in Section 1, the Lockean Thesis does seem materially right, though logically problematic. Presupposing our postulates so far made clear what logical closure properties on belief we find non-negotiable, and we also assumed the left-to-right direction of the Lockean Thesis; but with that in place, it would in fact be attractive to reclaim as many instances of the right-to-left direction of the Lockean Thesis as possible, in order to approximate as closely as possible the Lockean Thesis in its full original strength. And given our postulates, we have just shown that there is actually a unique way of achieving this, by making our agent ag's belief set as large as possible with respect to the absolute beliefs that it determines. So if we assume that our agent ag's *actual belief set Bel* is precisely that largest belief set, then this amounts to the assumption that given P, r, and all of our postulates up to this point, Bel is such that it approximates the Lockean Thesis as closely as possible, which does sound reasonable.

Let us formulate this in terms of our last postulate (which is a bridge principle again):

BP3 (Maximality)

Among all classes Bel' of ordered pairs of members of $\mathfrak{A}$, such that P and Bel' jointly satisfy P1–P2, B1–B6, BP1^r, and BP2, the class Bel is the largest with respect to the class of absolute beliefs, that is, the class of pairs of the form $\langle Z, W\rangle$, that it determines.

In other words, for all such Bel': $Bel \cap \{\langle Z, W\rangle \mid Z \in \mathfrak{A}\} \supseteq Bel' \cap \{\langle Z, W\rangle \mid Z \in \mathfrak{A}\}$.

The logical character of BP3 is obviously different from that of our previous postulates, but then again adding postulates that maximize or minimize classes subject to axiomatic constraints is of course not unheard of either. For example, famously, David Hilbert used this strategy in his axiomatization of geometry.

The term 'the largest' in BP3 is well-defined given our postulates P1–P2, B1–B6, BP1^r, BP2, by what we pointed out before.

Since we did not just deal with absolute belief in this section but also with belief conditional on any proposition that is consistent with everything the agent believes absolutely, one might wonder why we did not demand *Bel* in BP3 to be largest even with respect to the class of pairs $\langle Z, Y\rangle$ for which $Y \cap B_W \neq \varnothing$. However, let $B'_W \neq B''_W$ derive from two distinct candidates Bel', Bel'', such that both satisfy all of our postulates apart from BP3: by Theorem 3, without loss of generality, $B'_W \subsetneq B''_W$. But then, first of all, the class of all pairs $\langle Z, Y\rangle$ for which $Y \cap B'_W \neq \varnothing$ is distinct from the class of all pairs $\langle Z, Y\rangle$ for which $Y \cap B''_W \neq \varnothing$, so it would not be clear with respect to which of two classes our intended belief class *Bel* ought to be the largest. Furthermore, there are propositions $Z \in \mathfrak{A}$ (as, e.g., $B''_W \setminus B'_W$), such that Z has non-empty intersection with B''_W but not with B'_W; while BP6 would tell us whether $Bel''(.|Z)$, it would not give us any information whatsoever on $Bel'(.|Z)$. For these reasons, it will only be in the next section, where we will deal with conditional beliefs in general, that we will be in the position to strengthen Maximality so that it extends to *all* pairs $\langle Z, Y\rangle$ for $Z, Y \in \mathfrak{A}$ whatsoever. The resulting class *Bel* will again be defined uniquely, and the set of absolute beliefs that it determines will correspond to what is required by BP3 and the rest of the postulates of the present section.

With BP3 on board, and in view of our previous results, we may conclude from our postulates that in each and every case our agent's set B_W is nothing but the least non-empty P-stable^r set in $\mathfrak{A}$. In other words, our postulates (including BP3) entail the explicit definability of *ag*'s absolute beliefs, and indeed the definability of all of his beliefs conditional on any Y that is consistent with B_W, by means of the following corollary to our results mentioned before:

Corollary 9. *Let Bel be a class of ordered pairs of members of a σ-algebra $\mathfrak{A}$, let $P : \mathfrak{A} \to [0,1]$. Then the following two statements are equivalent*:

V. *P and Bel satisfy* P1–P2, B1–B6, BP1^r, BP2, BP3.

VI. *P satisfies* P1–P2, *there exists a* (*uniquely determined*) *least non-empty P-stable^r proposition X_{least} in $\mathfrak{A}$, and*:
   - *For all $Y \in \mathfrak{A}$ such that $Y \cap X_{least} \neq \varnothing$, for all $Z \in \mathfrak{A}$*:

$$Bel(Z|Y) \text{ if and only if } Z \supseteq Y \cap X_{least}$$

   - *In particular*: $B_W = X_{least}$, *and for all $Z \in \mathfrak{A}$*:

$$Bel(Z|W) \text{ if and only if } Z \supseteq X_{least}$$ [31]

VI of Corollary 9 can now be turned into an explicit definition of all relevant conditional beliefs just on the basis of P (and r and mathematical notions). Since in the next section we will extend this result to arbitrary conditional beliefs, whether or not they are beliefs conditional on propositions that are consistent with what the agent believes, we refrain from stating the resulting definition here. However, we do exploit

[31] Where the postulate BP3 was reminiscent of Hilbert's axiomatization of geometry, with respect to its open parameter r the present corollary is closer in spirit to something like Ernst Zermelo's famous quasi-categoricity result for second-order set theory: according to Zermelo's classical theorem, the cumulative hierarchy of sets is pinned down uniquely conditional on the specification of an ordinal number of a certain kind. The real number r in BP1^r above takes over the function of such an ordinal number in Zermelo's theorem, for only conditional on it the class *Bel* is specified uniquely.

Corollary 9 by deriving from it a particularly important special case: the explicit definition of the concept of *absolute* belief.

2.9. Defining absolute belief in terms of degrees of belief

In order to do so, we will have to take one final step. When we state our probabilistic definition of belief, we will restrict the probability measures P that we are interested in so that the initial existence claim in VI above is always satisfied. While our explicit definition of belief will then just hold conditional on that additional restriction, since the restriction is something that one would have to buy anyway if one grants all of our postulates, this should not cause major concern. Furthermore, we will see below that the constraint will not be overly demanding in all those contexts in epistemology and in related disciplines in which belief is a topic of interest (although it would be much too restrictive in other areas, such a measure theory, where one needs various measure spaces for integration that would be ruled out by the constraint).

This is thus the restriction on P that we will use. Call it the 'Least Certain Set Restriction': There is a member $X \in \mathfrak{A}$, such that $P(X) = 1$, and for every $Y \in \mathfrak{A}$, with $P(Y) = 1$: $X \subseteq Y$. That is: There is a least set of probability 1 in $\mathfrak{A}$. Equivalently, by P1, there is a member $X \in \mathfrak{A}$, such that $P(X) = 0$, and for every $Y \in \mathfrak{A}$, with $P(Y) = 0$: $Y \subseteq X$. Or in other words: there is a greatest set of probability 0 in $\mathfrak{A}$ (which is just the complement of the least set of probability 1). It is easy to see that the least proposition X of probability 1 cannot have a non-empty subset $Y \in \mathfrak{A}$, such that $P(Y) = 0$: for otherwise, $X \wedge \neg Y$, which is a member of $\mathfrak{A}$ again, would be a set of probability 1 that is a proper subset of X.

Given this Least Certain Set Restriction, there is always a least non-empty P-stable^r proposition in $\mathfrak{A}$ (for $\frac{1}{2} \leqslant r < 1$), for the same reasons as stated before: Either there is a non-empty P-stable^r proposition of probability less than 1, and then there is a least non-empty P-stable^r proposition anyway by Theorem 4. Or all and only non-empty P-stable^r propositions have probability 1: but then by the Least Certain Set Restriction there is a least set with probability 1, and that set is thus the least non-empty P-stable^r proposition in $\mathfrak{A}$. For the record again:

Observation 10. From the Least Certain Set Restriction (and $\frac{1}{2} \leqslant r < 1$) it follows that there is a least non-empty P-stable^r set in $\mathfrak{A}$. If there are no non-empty P-stable^r sets of probability less than 1, then the least non-empty P-stable^r set is the least set in $\mathfrak{A}$ which has probability 1, which must exist then by the Least Certain Set Restriction.

Observation 10 is but the counterpart of Observation 8 in a context in which our postulates have not been presupposed to hold.

Standard examples of countably additive probability measures for which there are least sets of probability 1 are:

- All probability measures on finite algebras $\mathfrak{A}$, and hence also all probability measures on algebras $\mathfrak{A}$ that are based on a finite set W of worlds.
- All countably additive probability measures on the power set algebra of a set W where W is countably infinite: In that case the conjunction of all sets of probability 1 is a member of the algebra of propositions again, and of course it is then the least set of probability 1.
- All countably additive probability measures (on a σ-algebra) that are regular (strictly coherent); that is, where: for all $X \in \mathfrak{A}$, $P(X) = 0$ if and only if $X = \varnothing$. Here W itself happens to be the least set of probability 1. Regularity does not enjoy general support, even though authors such as Carnap, Shimony, Stalnaker and others have argued for it as a plausible constraint on subjective probability measures, some of them in view of a special variant of the Dutch book argument that favors Regularity.

These examples demonstrate that a great variety of probability measures satisfy P1, P2, and this one additional constraint, and most of the typical philosophical toy examples of subjective probability measures are covered by this range of options.

We end up with the following explicit definition of absolute belief relative to any countably additive probability measure that satisfies the Least Certain Set Restriction, and relative to an appropriate threshold:

Definition 11. Let $P : \mathfrak{A} \to [0,1]$ be a countably additive probability measure on a σ-algebra $\mathfrak{A}$, such that there exists a least set of probability 1 in $\mathfrak{A}$. Let X_{least} be the least non-empty P-stable^r proposition in $\mathfrak{A}$ (which exists).

Then we say for all $Y \in \mathfrak{A}$ and $\frac{1}{2} \leqslant r < 1$:

$Bel^r_P(Y)$ ("Y is believed, to a cautiousness degree of r, as being given by P") if and only if $Y \supseteq X_{least}$.

From our previous considerations—summarized by Corollary 9—we know that this definition is *materially adequate*: Taking for granted that all of P1, B1–B6, BP3 are plausibly true, BP1^r is true conditional on the choice of r as a cautiousness threshold, and with P2 and BP2 being acceptable for the sake of simplicity, what we mean by 'materially adequate' here is: our definition of belief is a true statement, and it is provably so. In other words: If all of our postulates are true, which we believe to be the case, then Definition 11 can be proven to get the extension of '*Bel*' right as far as an agent's absolute or unconditional beliefs are concerned. What is more, if all of P1, B1–B6, BP1^r, BP3 are not just true but even necessary or analytic of belief, the definition is so as well (conditional on the additional presupposition of P2 and BP2).[32]

This is a reduction of absolute belief to degrees of belief, in the traditional sense of reduction by definition.

Example 9 *(The example from Section 1 reconsidered).* By Definition 11, we do have indeed: $X \subseteq W$ is to believed by the agent (to a cautiousness degree of $\frac{1}{2}$) as being given by P if and only if $\{w_1\} \subseteq X$, since $\{w_1\}$ is the least P-stable$^{\frac{1}{2}}$ set.

Accordingly, $X \subseteq W$ is to believed by the agent (to a cautiousness degree of $\frac{3}{4}$) as being given by P if $\{w_1, \dots, w_5\} \subseteq X$, because $\{w_1, \dots, w_5\}$ is the least P-stable$^{\frac{3}{4}}$ set.

2.10. Some further examples

Finally, we will illustrate the theory so far in terms of some further examples. In all of them, $\mathfrak{A}$ will simply be the full power set algebra of W.

If W contains exactly two worlds, then the situation is trivial insofar as for given $\frac{1}{2} \leqslant r < 1$, the singleton $\{w\} \subseteq W$ is the least non-empty P-stable^r proposition if $P(\{w\}) > r$, and W itself is such otherwise.

So let us turn to the first non-trivial case, that is, where W is a set $\{w_1, w_2, w_3\}$ of three elements. For simplicity, let $r = \frac{1}{2}$. Let us view all probability measures on that set W as being represented by points in a triangle, such that $P(\{w_1\})$, $P(\{w_2\})$, $P(\{w_3\})$ become the scalar factors of the convex combination of three given vectors that we associate with the worlds w_1, w_2, w_3. Then, depending on where P is represented in that triangle, P determines different classes of P-stable^r sets. See Fig. 6.

The diagram should be read as follows: The vertices of the outer equilateral triangle represent the probability measures that assign 1 to the singleton set of the respective world and 0 to all other singleton sets. Each non-vertex on any of the edges of the outer equilateral triangle represents a probability measure that assigns 0 to exactly one of the three worlds. Each edge of the inner equilateral triangle separates the representatives of probability measures of the following kinds: probability measures that assign to the

32 From what was shown before it follows immediately that the definitions of Definition 11 could actually be replaced by 'Y is a superset of *some* non-empty P-stable^r proposition in $\mathfrak{A}$' without thereby changing the extension of '*Bel*' in any way.

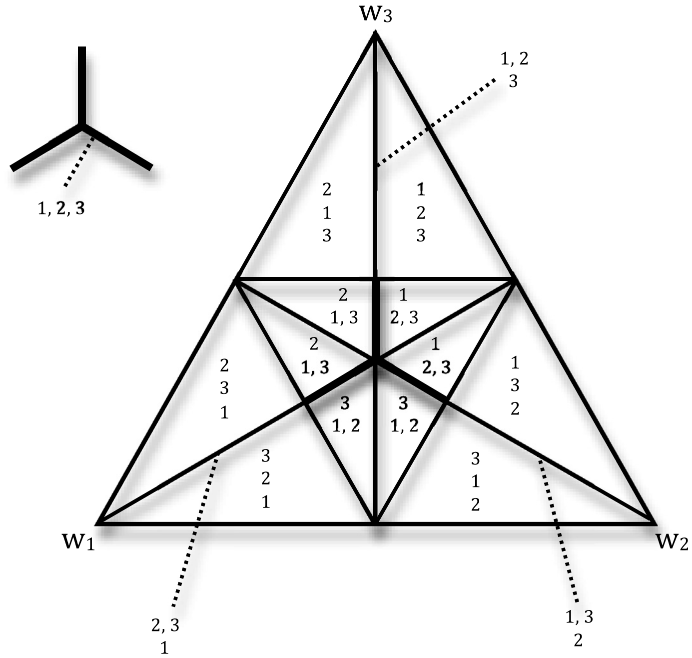

Fig. 6. Rankings determined by P.

singleton set of some world a probability that is *greater* than the sum of probabilities that it assigns to the singleton sets of the two other worlds; and probability measures that assign to the singleton set of some world a probability that is *less* than the sum of probabilities that it assigns to the singleton sets of the two other worlds. For instance, to the left-below of the left edge of the inner equilateral triangle we find such probability measures represented which assign to $\{w_1\}$ a greater probability than to the sum of what they assign to $\{w_2\}$ and $\{w_3\}$. Each straight line segment that connects a vertex with the mid-point of the opposite edge of the outer equilateral triangle separates the representatives of probability measures of the following kinds: probability measures that assign to the singleton set of one world a *greater* probability than to the singleton set of another world; and the probability measures that do so *the other way round*. Accordingly, the straight line segment that connects w_3 and the mid-point of the edge from w_1 to w_2 separates the probability measures that assign more probability to $\{w_1\}$ than to $\{w_2\}$ from those which assign more probability to $\{w_2\}$ than to $\{w_1\}$. The center point of both equilateral triangles represents the probability that is uniform over $W = \{w_1, w_2, w_3\}$.

Given all of that, and using the construction procedure for P-stable$^{\frac{1}{2}}$ sets that we have sketched in Subsection 2.5, it is easy to read off for each point, and hence for the probability measure that this point represents, all the non-empty P-stable$^{\frac{1}{2}}$ sets that are determined by it. To this we will turn now.

First of all, the points on the outer equilateral triangle are special: The probability measure represented by the vertex for w_i has $\{w_i\}$ as its least non-empty P-stable$^{\frac{1}{2}}$ set, all supersets of that set are non-empty and P-stable$^{\frac{1}{2}}$, too, and all of them have probability 1. Secondly, the probability measures represented by (the inner part of) the edge between the vertices that belong to two worlds w_i and w_j have either $\{w_i\}$, or $\{w_j\}$, or $\{w_i, w_j\}$ as their least non-empty P-stable$^{\frac{1}{2}}$ set, depending on whether the representing point is closer to the vertex of w_i than to the vertex of w_j, or *vice versa*, or equidistant of both of them; all supersets of each of them, respectively, are non-empty and P-stable$^{\frac{1}{2}}$ again, and all of them have probability 1.

But the really interesting part of the diagram concerns the interior of the outer equilateral triangle: Since relative to the probability measures that are represented as such only W has probability 1 (and hence is P-stable$^{\frac{1}{2}}$), we can concentrate solely on non-empty P-stable$^{\frac{1}{2}}$ sets with probability less than 1. As we have learned from Theorem 4, these form a sphere system or a ranked system of sets. In the diagram, we denote these sphere systems by enumerating in different lines the numeral indices of worlds of equal rank in the sphere system, starting with the worlds of rank 0 which we take to correspond to the entries in the bottom line of each numerical inscription. For example: Consider the interior of the two smallest rectangular triangles that are adjacent to w_1. Probability measures which are presented by points in the upper one yield a sphere system of three non-empty P-stable$^{\frac{1}{2}}$ sets: $\{w_1\}, \{w_1, w_3\}, \{w_1, w_2, w_3\}$. So w_1 has rank 0, w_3 has rank 1, and w_2 has rank 2. Accordingly, probability measures represented by points in the lower one of the two triangles determine a sphere system of the three non-empty P-stable$^{\frac{1}{2}}$ sets $\{w_1\}, \{w_1, w_2\}, \{w_1, w_2, w_3\}$. In either of these two cases, the probability measures in question would yield an absolute belief in every proposition that includes w_1 as a member, by Definition 11. And this makes quite good sense intuitively: for all of these measures are pretty close to the vertex for w_1. While the measures in the upper rectangular triangle prefer w_2 over w_3 in the ranking that they determine, this is just the other way around for the measures in the lower rectangular triangle. And again this makes perfect sense: after all, from the viewpoint of the “dominating” vertex for w_1, the measures in the upper rectangular triangle inhabit the w_3 half of the total equilateral triangle, while the measures in the lower rectangular triangle inhabit the w_2 half of the total equilateral triangle.

The further one moves geometrically towards the center point of the two equilateral triangles, the more coarse-grained the orderings become that are given by the sphere systems of the probability measures thus represented, and the smaller the class of absolutely believed propositions gets. Probability measures which are presented by points on any of the designated straight line segments within the interior of the outer equilateral triangle require special attention: Probability measures whose points lie on the boldface part in the diagram are treated separately in the little graphics left to the triangle; they all lead to the three worlds ranked equally. For three of the straight line segments we have denoted the sphere systems that they determine explicitly. The points on the three edges of the inner equilateral triangle—or rather the six halfs of those (without their midpoints which fall into the boldfaced lines)—yield sphere systems which coincide with those of the areas to which they are adjacent on the inside, which is why we did not say anything about them explicitly in Fig. 3. Finally, for the three straight line segments in the interior of the inner equilateral triangle we did not say anything about “their” sphere systems either because they simply inherit them from the rectangular triangle areas that they separate.

In particular, we should highlight that the ranking of worlds that is determined by the *uniform* probability measure that is represented by the center of the equilateral triangles is just w_1, w_2, w_3 ranked equally. That is: for such a measure, the theory predicts that what the agent ought to be believe is W and *there is no logically stronger proposition* which ought to be believed. Of course, this uniform probability case corresponds precisely to the case of the *lottery paradox*: the theory does not lead any contradiction here but what it does is to recommend to be cautious in such circumstances and not to believe of any single ticket that it will lose. The same observation applies if the number of worlds is greater than 3.

If $r > \frac{1}{2}$, then a diagram similar to Fig. 6 can be drawn, with all of the interior straight line segments being pushed towards the three vertices to an extent that is proportional to the magnitude of r.

One might wonder about why in Fig. 6 sphere systems with one world of rank 0 and two worlds of rank 1 are determined only by points or probability measures in one-dimensional line segments rather than in two-dimensional areas. In one sense, this is really just a consequence of dealing with precisely three worlds. If W had four members, then sphere systems with one world of rank 0, two worlds of rank 1, and hence one world of rank 2 would be represented in terms of proper areas again. However, it is true in general that sphere systems with precisely two worlds of maximal rank can only be represented by points or probability measures of areas of dimension $n-1$, if W has n members. For then the probabilities of these two worlds of maximal rank must be the same, which means the points of the represented probability measures must lie on one of the distinguished hyperplanes that generalize the distinguished line segments in our diagram to the higher-dimensional case.

For analogous reasons, the following is true: The set of points in the diagram which represent probability measures for which a set of probability 1 is the least P-stable$^{\frac{1}{2}}$ set has Lebesgue measure (geometrical measure) 0. In words: *Almost all probability measures have non-empty and non-trivial P-stable$^{\frac{1}{2}}$ sets!* This is because, for any such P for which all non-empty P-stable$^{\frac{1}{2}}$ sets are trivial (have probability 1): If there were a unique world whose singleton had least probability amongst all singletons, then W without that world would be P-stable$^{\frac{1}{2}}$ and non-trivial; so there must be at least two worlds whose singleton sets have the same probability, and the rest follows in the same way as before. This is not just so in the case of three possible worlds, such as in our diagram, but the same must hold for all finite probability space. Since this fact is important in showing that our whole theory is not just applicable non-trivially to an extremely narrow class of probability spaces, let us formulate this more explicitly:

> *For all finite algebras $\mathfrak{A}$, almost all probability measures over $\mathfrak{A}$ have a least P-stable$^{\frac{1}{2}}$ set with a probability less than 1.*

This said, if r is greater than $\frac{1}{2}$, then the Lebesgue measure of geometrical representatives of probability measures P that do not allow for a P-stable^r set of probability less than 1 is a positive number, and it gets ever closer to 1 if r get ever closer to 1. That is: the more cautious one aims to be in terms of one's threshold r, the harder it gets to find a probability measure that allows for non-trivial P-stable^r sets. (All of this holds for finite probability spaces.)

Finally, here is a simple *infinite* example: Let $W = \{w_1, w_2, w_3, \ldots\}$ be countably infinite, let $\mathfrak{A}$ be again the power set algebra on W, and let P be the unique regular countably additive probability measure that is given by: $P(\{w_1\}) = \frac{1}{2} + \frac{1}{4}$, $P(\{w_2\}) = \frac{1}{8} + \frac{1}{16}$, $P(\{w_3\}) = \frac{1}{32} + \frac{1}{64}$, and so on. Then the resulting non-empty P-stable$^{\frac{1}{2}}$ sets are:

$$\{w_1\},\ \{w_1, w_2\},\ \{w_1, w_2, w_3\},\ \ldots,\ \{w_1, w_2, \ldots, w_n\},\ \ldots \text{ and } W$$

and belief is determined accordingly.

Once we have covered conditional belief in full in the next section, we will return to some of these examples and analyze them in terms of conditional belief accordingly.

3. The reduction of belief II: Conditional belief in general

As promised, we finally generalize the postulates of the previous section to the case of beliefs that are conditional on propositions which may even be inconsistent with what our agent *ag* believes absolutely.

3.1. The generalized postulates

P1–P2 remain unchanged, of course. Our generalizations of B1–B5 simply result from dropping the antecedent '$\neg Bel(\neg X|W)$' condition which they contained:

B1* (Reflexivity)
$Bel(X|X)$.

B2* (One Premise Logical Closure)
For all $Y, Z \in \mathfrak{A}$: if $Bel(Y|X)$ and $Y \subseteq Z$, then $Bel(Z|X)$.

B3* (Finite Conjunction)
For all $Y, Z \in \mathfrak{A}$: if $Bel(Y|X)$ and $Bel(Z|X)$, then $Bel(Y \cap Z|X)$.

B4* (General Conjunction)
For $\mathcal{Y} = \{Y \in \mathfrak{A} \mid Bel(Y|X)\}$, $\bigcap \mathcal{Y}$ is a member of $\mathfrak{A}$, and $Bel(\bigcap \mathcal{Y}|X)$.

The Consistency postulate stays the same:

B5* (Consistency)
$\neg Bel(\varnothing|W)$.

The same arguments as before apply: B4* now entails that for every $X \in \mathfrak{A}$ there is a *least set* Y, such that $Bel(Y|X)$, which by B1* must be a subset of X. We denote this proposition again by: B_X. This is consistent with the corresponding notations that we used in the last section. Once again, we have

$$Bel(Y|X) \text{ if and only if } Y \supseteq B_X \text{ if and only if } Bel(Y|B_X)$$

The following postulate extends our previous Expansion postulate B6 to all cases of conditional belief whatsoever. It corresponds to the standard AGM postulates K*7 and K*8 (Superexpansion and Subexpansion) for belief revision taken together and translated again into the current context:

B6* (Revision)
For all $X, Y \in \mathfrak{A}$ such that $Y \cap B_X \neq \varnothing$:
For all $Z \in \mathfrak{A}$, $Bel(Z|X \cap Y)$ if and only if $Z \supseteq Y \cap B_X$.

So *any* $X \in \mathfrak{A}$ can now take over the role of W in our original B6 postulate on expansion. Equivalently:

B6* (Revision)
For all $X, Y \in \mathfrak{A}$, such that for all $Z \in \mathfrak{A}$, if $Bel(Z|X)$ then $Y \cap Z \neq \varnothing$:
For all $Z \in \mathfrak{A}$, $Bel(Z|X \cap Y)$ if and only if $Z \supseteq Y \cap B_X$.

That is: if the proposition Y is consistent with B_X—equivalently: Y is consistent with everything ag believes conditional on X—then ag believes Z conditional on the conjunction of Y and X if and only if Z is logically entailed by the conjunction of Y with B_X. Just as the original B6 postulate this can be justified in terms of standard possible worlds accounts of similarity orderings (as for David Lewis' conditional logic) or total plausibility rankings (as in belief revision and nonmonotonic reasoning): say, what a conditional belief expresses is again that the most plausible antecedent-worlds are consequent-worlds; then if some of the most plausible X-worlds are Y-worlds, these worlds must be precisely the most plausible $X \cap Y$-worlds, hence the most plausible $X \cap Y$-worlds are Z-worlds if and only if all the most plausible worlds X-worlds that are Y-worlds are Z-worlds.

Analogously to the last section, this is yet another equivalent statement of B6*:

B6* (Revision)
For all $X, Y \in \mathfrak{A}$ such that $Y \cap B_X \neq \varnothing$:
$B_{X \cap Y} = Y \cap B_X$.

The generalized version BP1^{r*} of our previous BP1^r postulate arises simply by dropping the '$Y \cap B_W \neq \varnothing$' restriction again. So we have:

BP1^{r*} (Likeliness)
For all $Y \in \mathfrak{A}$ with $P(Y) > 0$:
For all $Z \in \mathfrak{A}$, if $Bel(Z|Y)$, then $P(Z|Y) > r$.

Finally, we generalize BP2, and additionally we strengthen it by assuming also the converse of the resulting generalization:

BP2* (Zero Supposition)
For all $Y \in \mathfrak{A}$: $P(Y) = 0$ if and only if $B_Y = \varnothing$.

We have also replaced the original '$Y \cap B_W = \varnothing$' in BP2 by '$B_Y = \varnothing$': in the cases that were relevant in the last section, that is, where $Y \cap B_W \neq \varnothing$, the two were equivalent by B6; but now we need a more generally applicable formulation, since we can now no longer presuppose that $Y \cap B_W \neq \varnothing$.[33]

The reason why the original BP2 principle did not include the corresponding right-to-left direction of BP1^{r*} with the qualification '$Y \cap B_W \neq \varnothing$—that is, why we did not postulate also: If $Y \cap B_W \neq \varnothing$, then if $B_Y = \varnothing$ (that is, by B6, $Y \cap B_W = \varnothing$), then $P(Y) = 0$—is that the resulting principle would have been vacuously true. However, now the right-to-left direction is non-trivial.

As we have seen in our discussion of BP2 in Subsection 2.8, the postulate BP2, and hence the stronger BP2*, entail (given the other postulates): if $P(B_W) = 1$, all $Y \in \mathfrak{A}$ for which $P(Y) = 1$ holds are such that $B_Y = B_W$, and then B_W is the least proposition in $\mathfrak{A}$ of probability 1. The additional strength of BP2* has it that the propositions whose qualitative supposition lead to inconsistent beliefs are now precisely those for which conditionalization is undefined quantitatively. This completes our effort of bringing in line the quantitative and the qualitative supposition of zero sets, as explained in Subsection 2.8.[34]

And in combination with B6* it leads to further conclusions:

Observation 12. Our postulates entail that for all $X \in \mathfrak{A}$: B_X does not contain a non-empty zero subset.

Proof. For assume otherwise, that is: there is a $Y \subseteq B_X$ $(\subseteq X)$, $Y \neq \varnothing$, and $P(Y) = 0$. By B6*, $B_Y = B_{X \cap Y} = Y \cap B_X = Y$. So B_Y would have to be non-empty, too. But by BP2*, since $P(Y) = 0$, it must be that $B_Y = \varnothing$, which is a contradiction. □

3.2. The second representation theorem

We are now ready to prove the main result of our theory on conditional beliefs in general. The "soundness" direction (right-to-left) of the following representation theorem incorporates the corresponding direction of Grove's [13] representation theorem for belief revision operators in terms of sphere systems.

However, since all the propositions or sets of worlds that we are about to consider are required to be members of our given algebra $\mathfrak{A}$, it is not possible to simply translate the more difficult "completeness" part of Grove's representation theorem into our present context and then to apply it in our own left-to-right direction, since Grove's construction of spheres involves taking unions of propositions that would not be guaranteed to be members of our given σ-algebra $\mathfrak{A}$. Grove simply did not work in a probabilistic context in which propositions are required to be members of some given σ-algebra. That is also why our proof of that part of the theorem differs quite significantly from Grove's proof.

Here is thus our second, more general, representation theorem for conditional belief in general:

Theorem 13. *Let Bel be a class of ordered pairs of members of a σ-algebra $\mathfrak{A}$, and let $P : \mathfrak{A} \to [0, 1]$. Then the following two statements are equivalent*:

[33] Note that we did *not* assume that $B_Y = \varnothing$ implies $Y = \varnothing$; if we had, then BP2* would force P to be regular, that is, to assign zero probability to $\varnothing$ only.
[34] As mentioned back then, if we had started with a primitive conditional probability measure, which would in fact allow for conditionalization on zero sets, then BP2* should *not* be taken for granted, and it would not be needed either. But in the context of absolute probability measures, BP2* is natural to postulate in order to treat qualitative and quantitative supposition similarly.

I. *P and Bel satisfy* P1–P2, B1*–B6*, BP1^{r*}, BP2*.

II. *P satisfies* P1–P2, *$\mathfrak{A}$ contains a least set of probability* 1, *and there is a class $\mathcal{X}$ of non-empty P-stable^r propositions in $\mathfrak{A}$, such that* (i) *$\mathcal{X}$ includes the least set of probability* 1 *in $\mathfrak{A}$,* (ii) *all other members of $\mathcal{X}$ have probability less than* 1, *and*:

  – *For all $Y \in \mathfrak{A}$ with $P(Y) > 0$: if, with respect to the subset relation, X is the least member of $\mathcal{X}$ for which $Y \cap X \neq \varnothing$ holds (which exists), then for all $Z \in \mathfrak{A}$:*

$$Bel(Z|Y) \text{ if and only if } Z \supseteq Y \cap X$$

  Additionally, for all $Y \in \mathfrak{A}$ with $P(Y) = 0$, for all $Z \in \mathfrak{A}$: $Bel(Z|Y)$.

Furthermore, if I is the case, then $\mathcal{X}$ in II is actually uniquely determined.

Proof. The II-to-I direction is like the one in Theorem 3, except that one shows first that the equivalence for Bel entails for all $Y \in \mathfrak{A}$ with $P(Y) > 0$ that $B_Y = Y \cap X$, where X is the least member of $\mathcal{X}$ for which $Y \cap X \neq \varnothing$. The existence of that least member follows from Theorem 4, from the fact that every non-empty P-stable^r proposition with probability less than 1 is a subset of the least set in $\mathfrak{A}$ with probability 1 (by Observation 5), and from the fact that the least set of probability 1 in $\mathfrak{A}$ must have non-empty intersection with every proposition of positive probability (by P1). The proof of B6* is straightforward (and analogous to [13]), as is the proof of BP2* (by the 'Additionally, ...' assumption in II above).

So we can concentrate on the I-to-II direction: P1–P2 are satisfied by assumption. Now we define $\mathcal{X}$ by recursion as the class of all sets X_α of the following kind: For all ordinals $\alpha < \beta_P^r + 1$ (the successor of the ordinal that was defined in Subsection 2.6), let

$$X_\alpha = \bigcup_{\gamma<\alpha} [X_\gamma] \cup B_{W \setminus \bigcup_{\gamma<\alpha} X_\gamma}$$

(So, in particular, $X_0 = B_W$.)

At first we make a couple of observations about this class $\mathcal{X}$:

(a) Every member of $\mathcal{X}$ is also a member of $\mathfrak{A}$. By induction. For assume that all X_γ are in $\mathfrak{A}$ for $\gamma < \alpha < \beta_P^r + 1$: by Observation 6, β_P^r is countable and so are its predecessors, and therefore by $\mathfrak{A}$ being a σ-algebra, $\bigcup_{\gamma<\alpha} X_\gamma \in \mathfrak{A}$; thus $W \setminus \bigcup_{\gamma<\alpha} X_\gamma \in \mathfrak{A}$, and therefore $B_{W \setminus \cup_{\gamma<\alpha} X_\gamma} \in \mathfrak{A}$; hence, $X_\alpha \in \mathfrak{A}$.

(b) For all $\gamma < \alpha < \beta_P^r + 1$: $X_\gamma \subseteq X_\alpha$. This follows directly from the definition of the members of $\mathcal{X}$. From this it also follows that for all $\alpha + 1 < \beta_P^r + 1$: $X_{\alpha+1} = X_\alpha \cup B_{W \setminus X_\alpha}$.

(c) For all $\alpha < \beta_P^r + 1$: $X_\alpha = \bigcup_{\gamma<\alpha} B_{W \setminus \cup_{\delta<\gamma} X_\delta} \cup B_{W \setminus \cup_{\gamma<\alpha} X_\gamma}$. By induction. Assume that for all $\gamma < \alpha$: $X_\gamma = \bigcup_{\delta<\gamma} B_{W \setminus \cup_{\eta<\delta} X_\eta} \cup B_{W \setminus \cup_{\delta<\gamma} X_\delta}$. Substituting this for the first occurrence of 'X_γ' in the original definition of X_α, we conclude: $X_\alpha = \bigcup_{\gamma<\alpha} [\bigcup_{\delta<\gamma} B_{W \setminus \cup_{\eta<\delta} X_\eta} \cup B_{W \setminus \cup_{\delta<\gamma} X_\delta}] \cup B_{W \setminus \cup_{\gamma<\alpha} X_\gamma}$. But this can be simplified, by eliminating double counting of ordinals, to: $X_\alpha = \bigcup_{\gamma<\alpha} [B_{W \setminus \cup_{\delta<\gamma} X_\delta}] \cup B_{W \setminus \cup_{\gamma<\alpha} X_\gamma}$, which was to be shown.

(d) For all $\alpha < \beta_P^r + 1$: For all $Y \in \mathfrak{A}$ with $Y \cap X_\alpha \neq \varnothing$, it holds that $B_Y \subseteq X_\alpha$. This is because: If $Y \cap X_\alpha \neq \varnothing$, then by (c) there is a $\gamma \leqslant \alpha$, such that $Y \cap B_{W \setminus \cup_{\delta<\gamma} X_\delta} \neq \varnothing$, and by the well-orderedness of the ordinals, there must be a least such ordinal γ. Note that for that least ordinal γ it holds that $Y \cap \bigcup_{\delta<\gamma} X_\delta = \varnothing$, by (c) again, and hence $Y \subseteq W \setminus \bigcup_{\delta<\gamma} X_\delta$. By B6*, $B_{[W \setminus \cup_{\delta<\gamma} X_\delta] \cap Y} = Y \cap B_{W \setminus \cup_{\delta<\gamma} X_\delta}$, which is equivalent to $B_Y = Y \cap B_{W \setminus \cup_{\delta<\gamma} X_\delta}$ by what we have shown before. Finally, because $Y \cap B_{W \setminus \cup_{\delta<\gamma} X_\delta} \subseteq B_{W \setminus \cup_{\delta<\gamma} X_\delta} \subseteq X_\alpha$ by (c) again, it follows that $B_Y \subseteq X_\alpha$.

(e) For all $\alpha < \beta_P^r + 1$: X_α is P-stable^r. This can be derived from applying Definition 1: For all $Y \in \mathfrak{A}$, if $Y \cap X_\alpha \neq \varnothing$ and $P(Y) > 0$, then by (d), $B_Y \subseteq X_\alpha$, and hence by the standard properties of 'B_Y': $Bel(X_\alpha|Y)$. But this implies by BP1^{r*} that $P(X_\alpha|Y) > r$. That is: X_α is P-stable^r.

(f) There exists a least proposition $X \in \mathfrak{A}$ with probability 1, $X \in \mathcal{X}$, and X is the only member of $\mathcal{X}$ with probability 1.

First of all, assume for contradiction that all sets X_α with $\alpha < \beta_P^r + 1$ have probability less than 1. Since they are all P-stable^r by (e), it follows from (b) that there is a well-ordered sequence of (not necessarily strictly) increasing P-stable^r sets of probability less than 1, where the length of that whole sequence is $\beta_P^r + 1$. That sequence could not be one of *strictly* increasing P-stable^r sets of probability less than 1 throughout the sequence, for by the definition of β_P^r in Subsection 2.6, β_P^r itself was already the ordinal type of *all* P-stable^r sets of probability less than 1 whatsoever. So there must be $\alpha < \alpha' < \beta_P^r + 1$, such that $X_\alpha = X_{\alpha+1}$. Hence, by (b) again: $X_\alpha = X_\alpha \cup B_{W \setminus X_\alpha}$, and therefore $B_{W \setminus X_\alpha} \subseteq X_\alpha$. Since additionally $B_{W \setminus X_\alpha} \subseteq W \setminus X_\alpha$ by the definition of '$B_{W \setminus X_\alpha}$' and B1*–B4*, it follows that $B_{W \setminus X_\alpha} = \varnothing$. But because $P(X_\alpha) < 1$ by assumption, it also holds that $P(W \setminus X_\alpha) > 0$ by P1, so by the right-to-left direction of BP2* it follows that $B_{W \setminus X_\alpha} \neq \varnothing$, which is a contradiction. Hence, we have that there must be at least one set X_α with $\alpha < \beta_P^r + 1$ that has probability 1.

Secondly, consider any such set X_α with $\alpha < \beta_P^r + 1$ and $P(X_\alpha) = 1$: By (c), any such set X_α is a union of sets of the form 'B_X'. By Observation 12 and P2, any such set X_α does not have any non-empty zero subsets. But that means, by P1, any such set X_α must be identical to the unique least set $X \in \mathfrak{A}$ with probability 1, which therefore must exist. So we have that there exists a least proposition $X \in \mathfrak{A}$ with probability 1, $X \in \mathcal{X}$, and X is the only member of $\mathcal{X}$ with probability 1.

Now we can conclude the proof of II. We derive first the main equivalence claim for *Bel*. Let $Y \in \mathfrak{A}$ with $P(Y) > 0$: By P1 and (f), there is a member of $\mathcal{X}$ with which Y has non-empty intersection. Let $\alpha < \beta_P^r + 1$ be least, such that $Y \cap X_\alpha \neq \varnothing$: because of (b), X_α is then, with respect to the subset relation, the least member of $\mathcal{X}$ for which this holds. We will now show that $B_Y = Y \cap X_\alpha$, from which the relevant part of II follows by means of the definition of 'B_Y' and B1*–B4*. From (d) we know already that $B_Y \subseteq X_\alpha$ and hence with B1*–B4*, $B_Y \subseteq Y \cap X_\alpha$. Now consider $Y \cap X_\alpha$ again, which by assumption is non-empty: By (c), $X_\alpha = \bigcup_{\gamma<\alpha} B_{W \setminus \cup_{\delta<\gamma} X_\delta} \cup B_{W \setminus \cup_{\gamma<\alpha} X_\gamma}$. If Y had non-empty intersection with any set of the form $B_{W \setminus \cup_{\delta<\gamma} X_\delta}$ for $\gamma < \alpha$, then $Y \cap X_\gamma \neq \varnothing$, by (c) again, in contradiction with the way in which α was defined before. Therefore, $Y \cap X_\alpha = Y \cap B_{W \setminus \cup_{\gamma<\alpha} X_\gamma} \neq \varnothing$. The latter implies with B6* that $B_{[W \setminus \cup_{\gamma<\alpha} X_\gamma] \cap Y} = Y \cap B_{W \setminus \cup_{\gamma<\alpha} X_\gamma}$. By the defining property of α again, $Y \cap \bigcup_{\gamma<\alpha} X_\gamma$ is empty, and thus $[W \setminus \bigcup_{\gamma<\alpha} X_\gamma] \cap Y = Y$. So we have $B_Y = Y \cap B_{W \setminus \cup_{\gamma<\alpha} X_\gamma} = Y \cap X_\alpha$, and we are done.

Finally, consider $Y \in \mathfrak{A}$ with $P(Y) = 0$: By BP2*, $B_Y = \varnothing$, from which the last line of II follows by means of the definition of 'B_Y' and B1*–B4* again.

Uniqueness follows from: if there are two distinct such classes $\mathcal{X}$, $\mathcal{X}'$ with the stated properties, then they must differ with respect to at least one P-stable^r sets of probability less than 1. Without loss of generality, let X_α be the first member of $\mathcal{X}$ that is not also a member of $\mathcal{X}'$: since X_α is P-stable^r and has probability less than 1, it follows from Observation 6 that α is finite. If $\alpha = 0$, then B_W could not be the same as being given by $\mathcal{X}$ and $\mathcal{X}'$, which would be a contradiction. If α is a successor ordinal $\gamma + 1$, then, by (b), $B_{W \setminus X_\gamma} = X_\alpha \setminus X_\gamma$ as well as $B_{W \setminus X'_\gamma} = X'_\alpha \setminus X'_\gamma$. But $X_\gamma = X'_\gamma$, and hence $B_{W \setminus X_\gamma}$ could not be the same as being given by $\mathcal{X}$ and $\mathcal{X}'$, which would again be a contradiction. □

The right-to-left direction of the theorem gives us a recipe of how to build models for the conjunction of all our postulates from any σ-additive probability measure that satisfies the Least Certain Set Restriction: Just pick some non-empty P-stable^r sets (if there are any), take also the least set of probability 1, and put them together into a set $\mathcal{X}$. Then by defining *Bel* by means of what we say in II we end up with such a model. And the left-to-right direction of the theorem shows that *every* model of all of our postulates taken together can be built in such manner.

Example 10 *(The example from Section 1 reconsidered).* We know already that the non-empty and non-trivial P-stable$^{\frac{1}{2}}$ sets for our example P are: $\{w_1\}, \{w_1, w_2\}, \{w_1, \ldots, w_4\}, \{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}$.

By means of Theorem 13, it is now easy to define Bel so that all of our postulates on probability and conditional belief in general are satisfied (including all of our bridge axioms for them). For instance, pick $\{w_1, w_2\}, \{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}$, add the least set $\{w_1, \ldots, w_7\}$ of probability 1, and the resulting set

$$\mathcal{X} = \big\{\{w_1, w_2\}, \{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}, \{w_1, \ldots, w_7\}\big\}$$

will do, amongst many others. Note that for all $Z \subseteq W$, $Bel(Z|\{w_8\})$, if we follow the 'Additionally, ...' clause in II of Theorem 13, since $P(\{w_8\}) = 0$. In other words: $B_{\{w_8\}} = \varnothing$.

3.3. The final bridge principle and the explication of conditional belief in general

Theorem 13 generalizes Theorem 3 from Section 2 to conditional beliefs in general—accordingly, Theorem 3 simply dealt with the special case of a sphere system of just *one* P-stable^r set.

It remains to generalize BP3 in the now obvious way:

BP3* (Maximality)
Among all classes Bel' of ordered pairs of members of $\mathfrak{A}$, such that P and Bel' jointly satisfy P1–P2, B1*–B6*, BP1^{r*}, BP2*, the class Bel is the largest one.
In other words, for all such Bel': $Bel \supseteq Bel'$.

The rationale for BP3* is once again: to make as many right-to-left instance of the Lockean Thesis come out as true as possible (given that all other postulates up to this point are satisfied). However, in the present context, these instances have the more general form

$$\text{if } P(Z|Y) > r, \text{ then } Bel'(Z|Y)$$

with Y, Z in $\mathfrak{A}$, since we are now dealing with beliefs conditional on any proposition whatsoever.

Using BP3* as an additional constraint, we find:

Corollary 14. *Let Bel be a class of ordered pairs of members of a σ-algebra $\mathfrak{A}$, let $P:\mathfrak{A} \to [0,1]$. Then the following two statements are equivalent*:

III. *P and Bel satisfy* P1–P2, B1*–B6*, BP1^{r*}, BP2*, BP3*.
IV. *P satisfies* P1–P2, $\mathfrak{A}$ *contains a least set of probability* 1, *and if $\mathcal{X}$ is such that* (*and indeed is uniquely determined by*): (i) *$\mathcal{X}$ includes the least set of probability* 1 *in* $\mathfrak{A}$, (ii) *and all the other members of $\mathcal{X}$ are precisely all of the non-empty P-stable^r propositions in $\mathfrak{A}$ which have probability less than* 1, *then*:
– *For all $Y \in \mathfrak{A}$ with $P(Y) > 0$: if, with respect to the subset relation, X is the least member of $\mathcal{X}$ for which $Y \cap X \neq \varnothing$ holds* (*which exists*), *then for all $Z \in \mathfrak{A}$*:

$$Bel(Z|Y) \textit{ if and only if } Z \supseteq Y \cap X$$

Additionally, for all $Y \in \mathfrak{A}$ with $P(Y) = 0$, for all $Z \in \mathfrak{A}$: $Bel(Z|Y)$.

The crucial point here is: The class $\mathcal{X}$ in II includes as members *all* non-empty P-stable^r sets of probability less than 1 *whatsoever*.

Corollary 14 follows immediately from Theorem 13, except that we have to show: adding 'BP3*' to I of Theorem 13 is equivalent to including in $\mathcal{X}$ (as stated in IV) all non-empty P-stable^r sets of probability less than 1.

But that equivalence is a consequence of the following independent observation:

Observation 15. Let P be a countably additive probability measure on a σ-algebra $\mathfrak{A}$ over W. Assume that $\mathfrak{A}$ contains a least set of probability 1. Let $\mathcal{X}, \mathcal{X}'$ be classes of non-empty P-stable^r propositions for which (i) and (ii) in II of Theorem 13 are satisfied. Let Bel, Bel' be defined in terms of $\mathcal{X}, \mathcal{X}'$, respectively, as stated in II of Theorem 13. Then it holds:

If $\mathcal{X} \subseteq \mathcal{X}'$, then for all $Y, Z \in \mathfrak{A}$: If $Bel(Z|Y)$ then $Bel'(Z|Y)$.

Proof. Let $\mathcal{X} \subseteq \mathcal{X}'$. For Y with $P(Y) = 0$ there is nothing to show. So let Y be such that $P(Y) > 0$: If $Bel(Z|Y)$, then by definition $Z \supseteq Y \cap X$ with X being the least member of $\mathcal{X}$ for which $Y \cap X \neq \varnothing$ holds. But since X is also a member of $\mathcal{X}'$, the least member X' of $\mathcal{X}'$ for which $Y \cap X' \neq \varnothing$ holds must then be a subset of X; hence, $Z \supseteq Y \cap X'$ and therefore $Bel'(Z|Y)$. $\square$

From this it follows that choosing $\mathcal{X}$ to be the *greatest* class of all non-empty P-stable^r propositions in $\mathfrak{A}$ such that (i) and (ii) of II of Theorem 13 is satisfied must yield the *greatest* possible class Bel of pairs of propositions in $\mathfrak{A}$, if Bel is given as in II of Theorem 13. But that is exactly what we did in IV of Corollary 14. Accordingly, the total pre-order or ranking of worlds that corresponds to this greatest class $\mathcal{X}$ is the finest-grained one that is possible given P, r, and with all of our postulates up to, but not including, BP3* being satisfied. BP3* then says that this very $\mathcal{X}$ is precisely the one that determines the agent *ag*'s *actual* conditional belief set *Bel*.

As in the case of absolute belief, where what we needed was the existence of a least P-stable^r proposition, this additional Least Certain Set Restriction on P which we then introduced in Subsection 2.9, before we actually defined absolute belief probabilistically, is *entailed again by our postulates* on subjective probability and belief. So when we now finally turn IV of Corollary 14 into an explicit definition of belief again on the basis of P—but this time of *conditional belief in general*—then doing so "just" for probability measures for which there exists a least proposition of probability 1 is not an actual constraint, given that our postulates are true.

This is thus the intended, materially adequate, explicit definition of conditional belief:

Definition 16. Let $P : \mathfrak{A} \to [0,1]$ be a countably additive probability measure on a σ-algebra $\mathfrak{A}$, such that there exists a least set of probability 1 in $\mathfrak{A}$, and let $\frac{1}{2} \leqslant r < 1$.

Let $\mathcal{X}$ be uniquely determined by: (i) $\mathcal{X}$ includes the least set of probability 1 in $\mathfrak{A}$, (ii) and all the other members of $\mathcal{X}$ are precisely all the non-empty P-stable^r propositions in $\mathfrak{A}$ which have probability less than 1. Then we say for all $Y, Z \in \mathfrak{A}$:

$Bel^r_P(Z|Y)$ ("Z is believed conditional on Y, to a cautiousness degree of r, as being given by P") if and only if either (i) $P(Y) > 0$ and Z is a superset of the intersection of Y with the least non-empty P-stable^r proposition X_{least} in $\mathfrak{A}$ that has a non-empty intersection with Y (which exists), or (ii) $P(Y) = 0$.

By 'materially adequate' we mean the same as in Subsection 2.9: the definition, if taken as a descriptive statement (and given our postulates), is true. We have finally succeeded in reducing even full conditional belief to degrees of belief.

Example 11 *(The example from Section 1 reconsidered).* In the case in which the cautiousness threshold r is $\frac{1}{2}$, the sphere system that is going to define *Bel* is the set of all non-empty and non-trivial P-stable$^{\frac{1}{2}}$ sets together with the least set of probability 1. That is:

$$\mathcal{X} = \big\{\{w_1\}, \{w_1, w_2\}, \{w_1, \ldots, w_4\}, \{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}, \{w_1, \ldots, w_7\}\big\}$$

and Definition 16 then yields, e.g.: on the supposition of $\{w_1, w_2\}$, the agent believes $\{w_1\}$, and on the supposition of $\{w_3, w_4, w_6\}$, the agent believes $\{w_3, w_4\}$.

Analogously, in the case $r = \frac{3}{4}$:

$$\mathcal{X} = \big\{\{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}, \{w_1, \ldots, w_7\}\big\}$$

which implies by Definition 16: it is no longer the case that on the supposition of $\{w_1, w_2\}$ the agent believes $\{w_1\}$; but, e.g., on the supposition of $\{w_3, w_4, w_6\}$, the agent still believes $\{w_3, w_4\}$. And conditional on $\{w_5, w_6, w_7, w_8\}$ the agent does believe $\{w_5\}$.

More generally, it is easy to see that every finite sphere system can be realized in terms of the class of all P-stable$^{\frac{1}{2}}$ propositions of probability less than 1 for some measure P, and hence every AGM-style belief revision operator on a logically finite language is realizable in terms of a probability measure to which one applies Definition 16. So there are really *lots* of different types of sphere systems of P-stable$^{\frac{1}{2}}$ propositions (and accordingly for all other threshold values between $\frac{1}{2}$ and 1).

3.4. Some further examples, once again

Let us reconsider $W = \{w_1, w_2, w_3\}$, with $r = \frac{1}{2}$. Fig. 6 then already depicts the rankings of worlds, or the corresponding sphere systems, for the different possible probability measures on the power set algebra of W. For example: Consider again the interior of the two smallest rectangular triangles that are adjacent to w_1. Probability measures which are presented by points in the upper one determine a sphere system of three non-empty P-stable$^{\frac{1}{2}}$ sets: $\{w_1\}, \{w_1, w_3\}, \{w_1, w_2, w_3\}$. So w_1 has rank 0, w_3 has rank 1, and w_2 has rank 2. Accordingly, Bel as given by Definition 16 for $r = \frac{1}{2}$ and for any such P will have the property: $Bel(\{w_1\}|W)$, $Bel(\{w_2\}|\{w_2, w_3\})$.

In our little infinite example, the ranking of worlds simply coincided with their natural number index: so, e.g., $Bel(\{w_2\}|\{w_2, w_4, w_6, w_8, \ldots\})$ follows if Bel is determined from $r = \frac{1}{2}$ and our given P by means of Definition 16.

4. Prospects and potential problems

The theory in this paper can, and ought to be, supplemented in various ways, which we will have to leave to future work.

4.1. Prospects

First of all, as mentioned briefly before (see footnotes 18 and 26), it is possible to derive the very same theory on an alternative axiomatic basis: If the *full* Lockean Thesis (or actually its full right-to-left direction would be sufficient) is combined with the usual logical closure conditions on *absolute* belief, then these assumptions turn out to be perfectly consistent with each other as long as the threshold value of 'r' in the Lockean Thesis is made to depend on the probability measure P in question. And absolute belief can be shown to be determined then again by means of a P-stable^r set. In that alternative approach, while one needs to assume the right-to-left direction of the Lockean Thesis (in contrast with the theory in this paper), one may at the same time weaken the logical closure conditions that were employed here just to the case of absolute or unconditional belief.

Secondly, there are a great number of applications of the theory in all areas in which either qualitative or quantitative belief is prominent. In fact, the example measure P that we used for the purposes of illustration throughout the paper was taken from an article by Jon Dorling (see [7]) in which that measure is used to reconstruct in Bayesian terms an episode in 19th century astronomy: in our Fig. 1, A represents a relevant fragment of Newtonian mechanics, B represents the auxiliary hypothesis that tidal friction is negligible as far as the secular acceleration of the moon is concerned, and C represents the observation of the secular acceleration of the moon. The three of them taken together are inconsistent, which is why $\{w_8\}$ did have probability 0. Indeed, P is supposed to represent the degree of belief function of some reasonable and informed astronomer of the time. Plugging in that very P and a threshold r, say $\frac{3}{4}$, into our definition of belief, it follows: $Bel^r_P(A|C)$ and $Bel^r_P(\neg B|C)$. For conditional on $\{w_5, w_6, w_7, w_8\}$ the agent should believe $\{w_5\}$, as demonstrated in Example 11. In other words: Given the observation of the secular acceleration of the moon, a rational astronomer of the time ought to still believe in Newtonian mechanics but also disbelieve the negligibility of tidal friction; and of course that is exactly what happened in astronomy.

We can only mention some of the other topics to which the theory does have applications: the acceptability of indicative conditionals; the determination of parameters for Jeffrey conditionalization; the relation between knowledge and subjective probability; the relation between 'Y is an interesting hypothesis to be studied further, given X' and subjective probability; the acceptability of counterfactuals; expected chance and belief update; probabilistic truth conditions for counterfactuals by which a counterfactual can be true at a world without the corresponding conditional chance being equal to 1; the interpretation of Popper functions; vagueness and probability; preference and judgment aggregation; inference to the most cautious explanation; a representation theorem for probability in terms of qualitative belief; the relationship between action, decision theory, and qualitative belief; a correspondence between iterated belief revision and Jeffrey conditionalization; and a correspondence between introspective belief and probabilistic reflection principles. Some of these applications presuppose that 'Bel' is not interpreted as belief, and 'P' is not interpreted as *subjective* probability; but our theorems will still apply as long as our postulates are plausibly satisfied.

Other than these applications, which are all to questions and problems in philosophy, the theory should also be applicable to some more practical questions in computer science. In particular: Given a probabilistic database; how is it possible to break down the, possibly enormous, amounts of quantitative data that are stored in such a database in qualitative terms which might be more accessible to everyday users? The theory contained in this paper is able to generate systematically answers to questions of the form 'Shall I believe X?' or 'Shall I believe Y if X (that is, Y under the supposition of X)' if they are posed to a (complete) set of probabilistic data. And all of these answers will then come with the quality assurance of logical closure. For instance, if both 'if X then Y' and 'if X then Z' are accepted, the same will be the case for 'if X then $Y \wedge Z$'; and the like.

4.2. Potential problems

Finally, in future work, we will have to deal in detail with potential problems of the theory. Most importantly: We have seen in Subsection 2.10 that the classical lottery paradox does not necessarily cause trouble for the theory—if P is a uniform probability measure on a finite set W of worlds, then W is the only P-stable^r set, so only W is to be believed then. While not particularly exciting, this certainly makes good sense once the logical closure conditions for belief is taken as a given (as it has been the case in this paper); if it makes good sense even independently of the assumption of logical closure has to be left for further study.

But there are still problems in the same *ballpark* as the lottery paradox which might lead to more serious concerns: We showed in Subsection 2.7 that if X is P-stable^r, then X is also P-stable$^{\frac{1}{2}}$. And one of our results in Subsection 2.4 was that X is P-stable$^{\frac{1}{2}}$ if and only if the probability of any subset of X that has positive probability at all is greater than the probability of (any subset of) $\neg X$. Since the logically strongest proposition B_W that is believed by a rational agent has to be P-stable^r by the conjunction of our postulates, as followed from our first representation theorem (Theorem 3 in Subsection 2.4), it also follows that the probability of any subset of B_W that has positive probability at all is greater than the probability of (any subset of) $\neg B_W$. Now assume that there is a finite partition of B_W into pairwise disjoint sets $Y_1, \ldots, Y_n$, such that $B_W = Y_1 \cup \cdots \cup Y_n$. Furthermore, assume that P is uniform over these partition sets, that is, $P(Y_1) = \cdots = P(Y_n)$; then since each such probability $P(Y_i) \leqslant \frac{1}{n}$, also $P(\neg B_W) \leqslant \frac{1}{n}$, which means $P(B_W) \geqslant 1 - \frac{1}{n}$. In other words: The greater the number n of such partition cells, which might correspond to the outcomes of a fair lottery with n tickets which are all *believed possible*, the closer the probability of $P(B_W)$ must be to 1, and therefore also the closer the probability of *any* believed proposition X must be to 1. This is also exactly what is to be expected from the type of limitative theorems that one can find in Douven and Williamson [8].[35]

For closely related reasons, in the infinite case and assuming P2, it follows from IV in Theorem 4 of Subsection 2.6 that if for every non-empty proposition $Y \subseteq B_W$ there is a *proper* subset Y' of Y that is also a member of the given algebra $\mathfrak{A}$ of propositions, then the probability of B_W must be *equal to 1*. This was to be expected again in view of a limitative theorem that is applicable in a more general context, that is, Theorem 2 in Smith [41].

What all of this shows is: There are various types of probability spaces for which the definition of absolute belief that was part of the theory in this paper closely approximates or even collapses into the traditional Probability 1 Proposal. (There are similar worries also in the case of conditional belief, of course.) And this is so even though it also followed from Subsection 2.10 that in the case of a finite probability space almost all probability measures will *not* lead to such a collapse. In a nutshell: The more fine-grained an agent's qualitative beliefs, the more cautious an agent must be about choosing the threshold value of 'r' in our left-to-right direction of the Lockean Thesis ($\mathrm{BP1}^r$ and $\mathrm{BP1}^{r*}$ from above), and hence the closer this will get him to belief as determined by the Probability 1 Proposal.

But what if a reasonable and realistic agent's belief system always needs to be sufficiently fine-grained? Then our theory would simply amount to (more or less) the Probability 1 Proposal. In order to steer clear of a conclusion like that, it will be necessary to supply the present theory with an additional degree of freedom: the choice of a (more or less coarse-grained) *partition set* of the set W of (maximally specific) worlds where this partition set will have to be given by the belief *context* much in the same way in which the threshold value of 'r' is determined contextually. In each application of the theory, at first such a partition will have to be determined; the partition cells of any such partition will then function as the (more or less coarse-grained) "pseudo-worlds" in this context; and our theory, including the definition of belief, will then be applied just to the probability measure P restricted to propositions that can be built from these partition cells, that is, to a subalgebra of $\mathfrak{A}$. Belief as determined from P, r, *and such a partition* can differ completely from what follows from the Probability 1 Proposal, even if belief as determined from the original measure P over W would have approximated the Probability 1 Proposal for reasons such as the ones stated before. This doubly contextual approach, in view of the contextual determination both of r and of the partition, certainly needs proper defense. We will not attempt such a defense here but only point out for now that a similar relativization to partition sets is well known from theories such as Levi's theory of acceptance in Levi [27] and its follow-ups, and also Skyrms' theory of objective chance in his Skyrms [40], and that such a relativization seems to be unavoidable anyway under very general conditions if one is to determine belief from probability while sticking to the logical closure of belief, as follows from results in sections 13 and 14 of Lin and Kelly [32].

There are further problems to be dealt with: What can we say about the Preface Paradox in the setting of the present theory? And is the following an issue: It is easy to see that given our theory it holds that if $Bel^r_P(Y|X)$, then also $Bel^r_{P(\cdot|X)}(Y)$. In words: If relative to P (and r) one believes Y on the qualitative supposition of X, then after supposing X quantitatively by means of conditionalizing P on X, one also ends up believing Y absolutely relative to the measure after conditionalization (and r). However, the converse

[35] We would like to thank Timothy Williamson for an illuminating discussion on this.

does *not* hold necessarily—the ranking of worlds in X that is given by $P(\cdot|X)$ (and r) can be shown always to be a refinement of the original ranking on X that was determined by P (and r) and indeed it might be a *proper* refinement. In the finite case, the converse of the claim above holds only if all singleton sets $\{w\}$ and $\{w'\}$ of worlds w and w' of the same rank as determined by *Bel* as determined by P (and r) also have the same probability $P(\{w\}) = P(\{w'\})$. This is obvious given our theory: for if w had the same rank as w' as being given by *Bel* (as determined by P and r) while at the same time $P(\{w\}) < P(\{w'\})$—which we know to be perfectly possible—then $P(\{w\}|\{w, w'\}) > \frac{1}{2}$ and $P(\{w'\}|\{w, w'\}) < \frac{1}{2}$, so with $r = \frac{1}{2}$ it must hold that w is *more preferred* than w' according to Bel' as being determined by $P(\cdot|\{w, w'\})$ (and r). A similar result also follows in a much more general setting, as proven by Lin and Kelly [31], p. 10. Is this a problem for any probabilistic reconstruction of an AGM account of conditional belief (as Lin and Kelly [31] seem to suggest)? Or is this "merely" a reflection of the fact that probabilistic conditionalization operates on a more fine-grained scale than qualitative conditional belief? We leave answers to questions like these to future work.

For now we hope to have laid the foundations of what will hopefully become a valuable contribution to the "peace project" between logic and probability theory.

Acknowledgements

We are very grateful to David Etlin, Martin Krombolz, David Makinson, Arthur Paul Pedersen, Richard Pettigrew, Gerhard Schurz, Martin Smith, Philip Welch, Timothy Williamson, and an anonymous referee for their invaluable comments on this paper, and we are grateful to the Alexander von Humboldt Foundation for supporting this work.

References

[1] C.E. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, J. Symbolic Logic 50 (1985) 510–530.
[2] Horacio Arló-Costa, Bayesian epistemology and subjective conditionals: On the status of the export–import laws, J. Philos. 98 (11) (2001) 555–598.
[3] Horacio Arló-Costa, Rohit Parikh, Conditional probability and defeasible inference, J. Philos. Logic 34 (2005) 97–119.
[4] Pierpaolo Battigalli, Strong belief and forward induction reasoning, J. Econom. Theory 106 (2002) 356–391.
[5] Salem Benferhat, Didier Dubois, Henri Prade, Possibilistic and standard probabilistic semantics of conditional knowledge, J. Logic Comput. 9 (1997) 873–895.
[6] David Christensen, Putting Logic in Its Place, Clarendon Press, Oxford, 2004.
[7] Jon Dorling, Bayesian personalism, the methodology of scientific research programmes, and Duhem's problem, Stud. History Philos. Sci. A 10 (3) (1979) 177–187.
[8] Igor Douven, Timothy Williamson, Generalizing the lottery paradox, British J. Philos. Sci. 57 (4) (2006) 755–779.
[9] Hartry Field, What is the normative role of logic? Proc. Aristot. Soc. 83 (2009) 251–268.
[10] Richard Foley, Working Without a Net, Oxford University Press, Oxford, 1993.
[11] Keith Frankish, Mind and Supermind, Cambridge University Press, Cambridge, 2004.
[12] Peter Gärdenfors, Knowledge in Flux, MIT Press, Cambridge, MA, 1988.
[13] Adam Grove, Two modellings for theory change, J. Philos. Logic 17 (1988) 157–170.
[14] Gilbert Harman, Change in View, MIT Press, Cambridge, MA, 1986.
[15] James Hawthorne, Luc Bovens, The preface, the lottery, and the logic of belief, Mind 108 (1999) 241–264.
[16] James Hawthorne, David Makinson, The quantitative/qualitative watershed for rules of uncertain inference, Studia Logica 86 (2007) 247–297.
[17] Carl G. Hempel, Deductive-nomological vs statistical explanation, in: H. Feigl, G. Maxwell (Eds.), Minnesota Studies in the Philosophy of Science III, University of Minnesota Press, Minneapolis, 1962.
[18] Risto Hilpinen, Rules of Acceptance and Inductive Logic, Acta Philos. Fenn., vol. 22, North-Holland, Amsterdam, 1968.
[19] Jaakko Hintikka, Knowledge and Belief: An Introduction to the Logic of the Two Notions, Cornell University Press, Cornell, 1962.
[20] Franz Huber, Christoph Schmidt-Petri (Eds.), Degrees of Belief (Synthese Library), Springer, 2009, p. 342.
[21] Richard Jeffrey, Subjective Probability: The Real Thing, Cambridge University Press, Cambridge, 2004.
[22] Mark Kaplan, Decision Theory as Philosophy, Cambridge University Press, Cambridge, 1996.
[23] David H. Krantz, R. Duncan Luce, Patrick Suppes, Amos Tversky, Foundations of Measurement, vol. I, Academic Press, New York, 1971.
[24] Henry E. Kyburg Jr., Probability and the Logic of Rational Belief, Wesleyan University Press, Middletown, 1961.

[25] Henry E. Kyburg Jr., Probability and Inductive Logic, MacMillan, Toronto, 1970.

[26] Henry E. Kyburg Jr., Cho Man Teng, Uncertain Inference, Cambridge University Press, Cambridge, 2001.

[27] Isaac Levi, Gambling with the Truth. An Essay on Induction and the Aims of Science, MIT Press, Cambridge, MA, 1967.

[28] Isaac Levi, The Enterprise of Knowledge. An Essay on Knowledge, Credal Probability and Chance, MIT Press, Cambridge, MA, 1980.

[29] Isaac Levi, Commitment and change of view, in: José Luis Bermúdez, Alan Millar (Eds.), Reason and Nature: Essays in the Theory of Rationality, 2002, pp. 209–232.

[30] David K. Lewis, Counterfactuals, Blackwell, Oxford, 1973.

[31] Hanti Lin, Kevin T. Kelly, Propositional Reasoning that Tracks Probabilistic Reasoning, in: 8th Formal Epistemology Workshop, USC, Los Angeles, CA, 2011.

[32] Hanti Lin, Kevin T. Kelly, A geo-logical solution to the lottery paradox, Synthese 186 (2) (2012) 531–575.

[33] Patrick Maher, Betting on Theories, Cambridge University Press, Cambridge, 1993.

[34] David Makinson, The paradox of the preface, Analysis 25 (6) (1965) 205–207.

[35] David C. Makinson, Conditional probability in the light of qualitative belief change, J. Philos. Logic 40 (2) (2011) 121–153.

[36] Jonathan Roorda, Revenge of Wolfman: A probabilistic explication of full belief, unpublished manuscript, www.princeton.edu/~bayesway/pu/Wolfman.pdf, 1995.

[37] Jacob Ross, Mark Schroeder, Belief, credence and pragmatic encroachment, Philos. Phenomenol. Res., in press, http://dx.doi.org/10.1111/j.1933-1592.2011.00552.x.

[38] Brian Skyrms, Resiliency, propensities, and causal necessity, J. Philos. 74 (11) (1977) 704–713.

[39] Brian Skyrms, Causal Necessity, Yale University Press, New Haven, 1980.

[40] Brian Skyrms, Pragmatics and Empiricism, Yale University Press, New Haven, 1984.

[41] Martin Smith, A generalised lottery paradox for infinite probability spaces, British J. Philos. Sci. 61 (4) (2010) 821–831.

[42] Paul Snow, Is intelligent belief really beyond logic? in: Proceedings of the Eleventh International Florida Artificial Intelligence Research Society Conference, American Association for Artificial Intelligence, 1998, pp. 430–434.

[43] Paul Snow, Diverse confidence levels in a probabilistic semantics for conditional logic, Artificial Intelligence 113 (1999) 269–279.

[44] Paul Snow, Belief, logic, and partial truth, Comput. Intell. 19 (2) (2003) 201–213.

[45] Wolfgang Spohn, Ordinal conditional functions. A dynamic theory of subjective states, in: William L. Harper, Brian Skyrms (Eds.), Causation in Decision, Belief Change, and Statistics, vol. II, Kluwer, Dordrecht, 1988, pp. 105–113.

[46] Scott Sturgeon, Reason and the grain of belief, Nous 42 (1) (2008) 139–165.

[47] Marshall Swain (Ed.), Induction, Acceptance and Rational Belief, Reidel, Dordrecht, 1970, pp. 55–82.

[48] Bas C. Van Fraassen, Fine-grained opinion, probability, and the logic of full belief, J. Philos. Logic 24 (1995) 349–377.

[49] Gregory Wheeler, A review of the lottery paradox, in: W. Harper, G. Wheeler (Eds.), Probability and Inference: Essays in Honor of Henry E. Kyburg Jr., 2007, pp. 1–31.

[50] Timothy Williamson, Conditionalizing on knowledge, British J. Philos. Sci. 49 (1998) 89–121.

[51] Ronald R. Yager, Liping Liu, Classic Works of the Dempster–Shafer Theory of Belief Functions, Stud. Fuzziness Soft Comput., vol. 219, Springer, Berlin, 2008.