Reducing Belief Simpliciter to Degrees of Belief

Hannes Leitgeb

March 2010

Abstract

We prove that given reasonable assumptions, it is possible to give an explicit definition of belief simpliciter in terms of subjective probability, such that 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. Belief simpliciter is not to be eliminated in favour of degrees of belief, rather, by reducing it to assignments of consistently high degrees of belief, both quantitative and qualitative belief turn out to be governed by one unified theory. Turning to possible applications and extensions of the theory, we suggest that this will allow us to see: how the Bayesian approach in general philosophy of science can be reconciled with the deductive or semantic conception of scientific theories and theory change; how primitive conditional probability functions (Popper functions) arise from conditionalizing absolute probability measures on maximally strong believed propositions with respect to different cautiousness thresholds; how the assertability of conditionals can become an all-or-nothing affair in the face of non-trivial subjective conditional probabilities; how knowledge entails a high degree of belief but not necessarly certainty; and how high conditional chances may become the truthmakers of counterfactuals.

1 Introduction

[THIS IS A PRELIMINARY AND INCOMPLETE DRAFT OF JUST THE TECHNICAL DETAILS...]

Belief is said to come in a quantitative version—degrees of belief—and in a qualitative one—belief simpliciter. More particularly, *rational* belief is said to have such a quantitative and a qualitative side, and indeed we will only be interested in notions of belief here which satisfy some strong logical requirements. Quantitative belief is given in terms of numerical degrees that are usually assumed to obey the laws of probability, and we will follow this tradition. Belief simpliciter, which only recognizes belief, disbelief, and suspension of judgement, is closed under deductive inference as long as every proposition that an agent is committed to believe is counted as being believed in an idealised sense; this is how epistemic logic conceives of belief, and we will subscribe to this view in the following. Despite of these logical differences between the two notions of belief, it would be quite surprising if it did not turn out that quantitative and qualitative belief were but aspects of one and the same underlying substratum; after all, they are both concepts of *belief*. However, this still allows for a variety of possibilities: they could be mutually irreducible conceptually, with only some more or less tight bridge laws relating them; or one could be reducible to the other, without either of them being eliminable from scientific or philosophical thought; or either of them could be eliminable. So which of these options should we believe to be true?

The concept of quantitative belief is being applied successfully by scientists, such as cognitive psychologists, economists, and computer scientists, but also by philosophers, in particular, in epistemology and decision theory; eliminating it would be detrimental both to science and philosophy. On the other hand, it has been suggested (famously, by Richard Jeffrey) that the concept of belief simpliciter can, and should, be eliminated in favour of keeping only quantitative belief. But this is not advisable either: (i) Epistemic logic, huge chunks of cognitive science, and almost all of traditional epistemology rely on the concept of belief in the qualitative sense; by abandoning it one would simply have to sacrifice too much. (ii) Beliefs held by some agent are the mental counterparts of the scientific theories and hypotheses that are held by a scientist or a scientific community; they can be true or false just as those theories and hypotheses can be (taking for granted a realist view of scientific theories). But not many would recommend banning the concept *holding a scientific theory/hypothesis* from science or philosophy of science. (iii) The concept of belief simpliciter, which is a classificatory concept, occupies a more elementary scale of measurement than the numerical concept of quantitative belief does, which is precisely one of the reasons why it is so useful. That is also why giving up on any of the standard properties of rational belief, such as closure under conjunction (the Conjunction property)—if X and Y are believed, then $X \wedge Y$ is believed—as some have suggested in response to lottery-type paradoxes (see Kyburg...), would not be a good idea: for without these properties belief simpliciter would not be so much less complex than quantitative belief anymore (however, see Hawthorne & Makinson...). But then one could have restricted oneself to quantitative belief from the start, and in turn one would lack the simplifying power of the qualitative belief concept. (iv) Beliefs involve dispositions to act under certain conditions. For instance, if I believe that my original edition of Carnap's *Logical Syntax* is on the bookshelf in my office, then given the desire to look something up in it, and with the right background conditions being satisfied, such as not being too tired, not being distracted by anything else, and so on, I am disposed to go to my office and pick it up. The same belief also involves lots of other dispositions, and what holds all of these dispositions together is precisely that belief. If one looks at the very same situation in terms of degrees of belief, then with everything else in place, it will be a matter of what my degree of belief in the proposition that Carnap's *Logical Syntax* is in my office is like whether I will actually go there or not, and similarly for all other relevant dispositions. Somehow the continuous scale of degrees of belief must be cut down to a binary decision: acting in a particular way or not. And the qualitative concept of belief is exactly the one that plays that role, for it is meant to express precisely the condition other than desire and background conditions that needs to be satisfied in order for to me to act in the required way, that is, for instance, to walk to the office and to pick up Carnap's monograph from the bookshelf. Decision theory, which is a probabilistic theory again, goes some way of achieving this without using a qualitative concept of belief, but it does not quite give a complete account. Take assertions as a class of actions. One of the linguistic norms that govern assertability is: If all of $A_1, \ldots, A_n$ are assertable for an agent, then so is $A_1 \wedge \ldots \wedge A_n$. One may of course attack this norm on different grounds, but the norm still seems to be in force both in everday conversation and in scientific reasoning. Here is plausible way of explaining why we obey that norm by means of the concept of qualitative belief: Given the right desires and background conditions, a descriptive sentence gets asserted by an agent if and only if the agent believes the sentence to be true. And the assertability of a sentence A is just that very necessary epistemic condition for assertion—belief in the truth of A—to be satisfied. (Williamson... states an analogous condition in terms of knowledge rather than belief; but it is again a qualitative concept that is used, not a quantitative one.) But if an agent believes all of $A_1, \ldots, A_n$, then the agent believes, or is at least epistemically committed to believe, also $A_1 \wedge \ldots \wedge A_n$. That explains why if $A_1, \ldots, A_n$ are assertable for an agent, so is $A_1 \wedge \ldots \wedge A_n$. And it is not clear how standard decision theory just by itself, without any additional resources at hands, such as a probabilistic explication of belief, would be able to give a similar explanation. The assertability of indicative conditionals $A \to B_i$ makes for a similar case. Here, one of the linguistic norms is: If all of $A \to B_1, \ldots, A \to B_n$ are assertable for an agent, then so is $A \to (B_1 \wedge \ldots \wedge B_n)$. This may be explained by invoking the Ramsey test for conditionals (see...) as follows: Given the right desires and background conditions, $A \to B_i$ gets asserted by an agent if and only if the agent accepts $A \to B_i$, which in turn is the case if and only if the agent believes B_i to be true conditional on the supposition of A. Again, the assertability of a sentence, $A \to B_i$, is just that respective necessary epistemic condition—belief in B_i on the supposition of A—to be satisfied. But, if an agent believes all of $B_1, \ldots, B_n$ conditional on A, then the agent believes, or is epistemically committed to believe, also $B_1 \wedge \ldots \wedge B_n$ on the supposition of A. Therefore, if $A \to B_1, \ldots, A \to B_n$ are assertable for an agent, so is $A \to (B_1 \wedge \ldots \wedge B_n)$. Ernest Adams' otherwise marvellous probabilistic theory of indicative conditionals (...), which ties the acceptance of any such conditional to its corresponding conditional subjective probability and hence to the quantitative counterpart of conditional belief, does not by itself manage to explain such patterns of assertability. While from Adams' theory one is able to derive that the uncertainty (1 minus the corresponding conditional probability) of $A \to (B_1 \wedge \ldots \wedge B_n)$ is less than or equal the sum of the uncertainties of $A \to B_1, \ldots, A \to B_n$, and thus if all of the conditional probabilities that come attached to $A \to B_1, \ldots, A \to B_n$ tend to 1 then so does the conditional probability that is attached to $A \to (B_1 \wedge \ldots \wedge B_n)$, it also follows that for an increasing number n of premises, ever greater lower boundaries $1 - \delta$ of the conditional probabilities for $A \to B_1, \ldots, A \to B_n$ are needed in order to guarantee that the conditional probability for $A \to (B_1 \wedge \ldots \wedge B_n)$ is bounded from below by a given $1 - \epsilon$. No *uniform* boundary emerges that one might use in order to determine for a conditional—whether premise or conclusion, whatever the number of premises, or whether in the context of an inference at all—its assertability simpliciter. But since there is only assertion simpliciter, at some point a condition must be invoked that discriminates between what is a case of asserting and what is not. Once again the concept of (conditional) qualitative belief gives us exactly what we need.

The upshot of this is: Neither the concept of quantitative belief nor the concept of qualitative belief ought to be eliminated from science or philosophy. But this leaves open, in principle, the possibility of *reducing* one to the other without eliminating either of them—using traditional terminology: one concept might simply turn out to be logically prior to the other. Now, reducing degrees of belief to belief simpliciter seems unlikely (no pun intended!), simply because the formal structure of quantitative belief is so much richer than the one of qualitative belief. But for the same reason, at least *prima facie*, one would think that the converse ought to be feasible: by abstracting in some way from degrees of belief, it ought to be possible to explicate belief simpliciter in terms of them. Belief simpliciter would thus be qualitative only at first glance; its deeper logical structure would turn out to be quantitative after all. One obvious suggestion of how to explicate belief simpliciter on the basis of degrees of belief is to maintain that having the belief that X is just having assigned to X a degree of belief strictly above some threshold level less than 1 (this is called the Lockean thesis by Richard Foley... more about which below). If that threshold is also greater than or equal to $\frac{1}{2}$, then belief would simply amount to high subjective probability. But since the probability of $X \wedge Y$ might well be below the threshold even when the probabilities of X and Y are not, one would thus have to sacrifice logical properties such as the Conjunction property, which one should not, as mentioned above. While the Lockean thesis seems materially fine, for qualitative belief *does* seem to be close to high subjective probability, it does not get the logical properties of qualitative belief right. Or one identifies the belief that X with having a degree of belief of 1 in X: call this the 'probability

1 proposal'. While this does much better on the logical side, it is not perfect on that side either. Truth for propositions is certainly closed under taking conjunctions of arbitrary cardinality, however, being assigned probability 1 is not so except for those cases in which probability assignments simply coincide with truth value assignments; but in the presence of uncertainy, subjective probability measures do not. If qualitative belief inherits this general conjunction property from truth—maybe because truth is what qualitative beliefs aim at, whether directly or indirectly—then an explication of qualitative belief in terms of probability 1 is simply not good enough. More importantly, apart from such logical considerations, the proposal is materially wrong. As Roorda (...) pointed out, our pre-theoretic notions of belief-in-degrees and belief simpliciter have the following epistemic and pragmatic properties: (i) One can believe X and Y without assigning the same degree of belief to them. But then at least one of X and Y must have a probability other than 1. For instance, I believe that my desk will still be there when I enter my office tomorrow, and I also believe that every natural number has a successor, but should I therefore be forced to assign the same degree of belief to them? (ii) One can believe X without being disposed to accept every bet whatsoever on X, although the latter ought be that case by the standard Bayesian understanding of probabilities if one assigns probability 1 to X, at least as long as the stakes of the bet are not too extravagant. For example, I do believe that I will be in my office tomorrow. But I would refrain from accepting a bet on this if I were offered 1 Pound if I were right, and if I were to lose lose 1000 Pound if not. (Alternatively, one could abandon the standard interpretation of subjective probabilities in terms of betting quotients, but breaking with such a successful tradition comes with a price of its own. However, later we will see that our theory will allow for a reconciling offer in that direction, too.) Roorda's presents a third argument against the probability 1 proposal based on considerations on fallibilism, but with it we are going to deal later. This shows that Ramsey's term 'partial belief' for subjective probability is in fact misleading (or at least ambiguous, about which more later): for *full* belief, that is, belief simpliciter, does not coincide with having a degree of belief of 1, and hence a degree of belief of less than 1 should not be regarded as *partial* belief. All of these points also apply to a much more nuanced version of the probability 1 proposal which was developed by Bas van Fraasen, Horacio Arlo-Costa, and Rohit Parikh, according to which within the quantitative structure of primitive conditional probability measures (Popper functions) one can always find so-called belief cores, which are propositions with particularly nice and plausible logical properties; by taking supersets of those one can define elegantly notions of qualitative belief in different variants and strengths. But the same problems as mentioned before emerge, since all belief cores can be shown to have absolute probability 1. Additionally, the axioms of Popper functions are certainly more controversial than those of the standard absolute or unconditional probability measures, and since two distinct belief cores differ only in terms of some set

of absolute probability 0, one wonders whether in many practically relevant situations in which only probability measures on finite spaces are needed and where often there are no non-empty zero sets at all—or otherwise the corresponding worlds with zero probabilistic weight would simply have been dropped from the start—the analysis is too far removed from the much more mundane reality of real-world reasoning and epistemological thought experiments. On the other hand, we will see that the *logical* properties of belief cores are enormously attractive: we will return to this later, when we will show that it is actually possible to restore most of them in the new setting that we are going to propose.

Summing up: Reducing qualitative belief to quantitative belief does not seem to work either. In the words of Jonathan Roorda (...), "The depressing conclusion ... is that no explication of belief is possible within the confines of the probability model". Roorda himself then goes on to suggest an explication that is based on *sets* of subjective probability measures rather than just one probability measure as standard Bayesianism has it. In contrast, we will bite the bullet and stick to just one probability measure below.

Given all of these problems, the only remaining option seems to be: neither of quantitative or qualitative belief can be reduced to the other; while there are certainly bridge principles of some kind that relate the two, it is impossible to understand qualitative belief just in terms of quantitative belief or the other way round. A view like this has been proposed and worked out in detail, for example, by Isaac Levi (...) and recently by James Hawthorne (...). And apart from extreme Bayesians who believe that one can do without the concept of qualitative belief, it is probably fair to say that something like this is the dominating view in epistemology these days.

In what follows, we are going to argue *against* this view: we aim to show that it is in fact possible to reduce belief simpliciter to probabilistic degrees of belief by means of an explicit definition, without stripping qualitative belief of any of its constitutive properties, without revising the intended interpretation of subjective probabilities in any way, without running into any of the difficulties that we found to affect the standard proposals for quantitative explications of belief, and without thereby intending to eliminate the concept of belief simpliciter in favour of quantitative belief. Both notions of belief will be preserved; it is just that having the qualitative belief that A will turn out to be definable in terms of assignments of consistently high degrees of belief, where what this means exactly will be clarified below. We will also point out which consequences this has for various problems in philosophy of science, epistemology, and the philosophy of language. And for the convinced Bayesian, who despises qualitative belief, the message will be: within your subjective probability measure you find qualitative belief anyway; so you might just as well use it.

Before we turn to the details of our theory, we will first sketch the underlying idea of the explication.

2 The Basic Idea

Our starting point is again what Richard Foley (..., pp. 140f) calls the *Lockean thesis*, that is:

> to say that you believe a proposition is just to say that you are sufficiently confident of its truth for your attitude to be one of belief

and consequently

> it is rational for you to believe a proposition just in case it is rational for you to have a sufficiently high degree of confidence in it, sufficiently high to make your attitude toward it one of belief.

He takes this to be derivative from Locke's views on the matter, as exemplified by

> most of the Propositions we think, reason, discourse, nay act upon, are such, as we cannot have undoubted Knowledge of their Truth: yet some of them border so near upon Certainty, that we make no doubt at all about them; but *assent* to them firmly, and act, according to that Assent, as resolutely, as if they were infallibly demonstrated, and that our Knowledge of them was perfect and certain (Locke..., p. 655, Book IV, Chapter XV; his emphasis)

and

> *the Mind if it will proceed rationally, ought to examine all the grounds of Probability*, and see how they make more or less, *for or against* any probable Proposition, before it assents to or dissents from it, and upon a due ballancing the whole, reject, or receive it, with a more or less firm assent, proportionably to the preponderancy of the greater grounds of Probability on the one side or the other. (Locke..., p. 656, Book IV, Chapter XV; his emphasis)

We take this account of belief simpliciter in terms of high degrees of belief to be right in spirit. However, as we know from lottery paradox situations, it is not yet good enough: there are logical principles for belief (such as the Conjunction principle) which we regard as just as essential to the belief in X as assigning a sufficiently high subjective probability to X, and it is precisely these logical principles that which are invalidated if the Lockean thesis is turned into a definition of belief. Instead, we take the Lockean thesis to characterise a more preliminary notion of belief, or what one might call *prima facie* belief:

Definition 1 *Let P be a subjective probability measure. Let X be a proposition in the domain of P: X is believed prima facie as being given by P if and only if $P(X) > r$.*

Of course, more needs to be said about the threshold value r here, but let us postpone this discussion.

In analogy with the case of *prima facie* obligations in ethics, a proposition is believed *prima facie* in view of the fact that it has an epistemic feature that speaks in favour of it being a belief proper—that is, to have a sufficiently high subjective probability—and as long as no other of its epistemic properties tells against it being such, it will in fact be properly believed.

Accordingly, as far as belief itself is concerned, we suggest to drop just the right-to-left direction of the Lockean thesis, so that high subjective probability is still a necessary condition for belief but it is not anymore demanded to be a sufficient one. Thus, ultimately, all beliefs simpliciter will be among the *prima facie* candidates for beliefs. The left-to-right direction is going to ensure that beliefs remain reasonably cautious—how cautious will depend on the "cautiousness parameter" r—and that they inherit all the dispositional consequences of having sufficiently high degrees of belief. On the other hand, the right-to-left direction was the one that got us into lottery-paradox-like trouble. Instead of it, we will regard all the standard logical principles for belief as being constitutive of belief from the start. Unlike the definition of *prima facie* belief which expresses a condition to be satisfied by single beliefs, these logical principles do not apply to beliefs taken by themselves but rather to systems of beliefs taken as wholes. Therefore, when putting together the left-to-right direction of the Lockean thesis with these logical postulates, we need to formulate the result as a constraint on an agent's belief *system* or *class*. Furthermore, we will not just do this for absolute or unconditional belief—the belief that X is the case—but also for conditional belief, that is, belief under a supposition, as in: the belief that X is the case under the supposition that Y is the case. Indeed, generalizing the left-to-right direction of the original Lockean thesis to cases of conditional belief will pave the way to our ultimate understanding of belief. And arguably belief simpliciter under a supposition is just as important for our epistemic lives as belief simpliciter taken absolutely or unconditionally. This will give us then something of the following form:

- If P is an agent's degree-of-belief function at a time t, and if Bel is the class of believed propositions by the agent at t (and both relate to the same underlying class of propositions), then they have the following properties:

  (1) Probabilistic constraint:

  - $*$ P is a probability measure.

    $\vdots$

    (Additional constraints on P.)

  (2) Logical constraints:

- $*$ For all propositions Y, Z: if $Y \in Bel$ and Y logically entails Z, then $Z \in Bel$.
- $*$ For all propositions Y, Z: if $Y \in Bel$ and $Z \in Bel$, then $Bel(Y \cap Z)$.
- $*$ No logical contradiction is a member of Bel.

  $\vdots$

  (Other standard logical principles for Bel and their extensions to conditional belief.)

(3) Mixed constraints:

- $*$ For all propositions $X \in Bel$, $P(X) > r$.
- $*$ (An extension of this to conditional belief.)

  $\vdots$

  (Additional mixed constraints on P and Bel.)

While the conjunction of (1), (2), and (3) might well do as a meaning postulate on '*Bel*' and '*P*', obviously this is not an explicit definition of '*Bel*' on the basis of '*P*' anymore. Is there any hope of turning it into an explicit definition of belief again?

Immediately, David Lewis' (...) classic method of defining theoretical terms, which builds on work by Ramsey and Carnap, comes to mind: given P, define '*Bel*' to be *the* class, such that the conditions on Bel and P above are the case. But of course this invites all the standard worries about such definitions by definite description: First of all, for given P, there might simply not be any such class Bel at all. Fortunately, we will be able to prove that this worry does *not* get confirmed. Secondly, at least for many P, there might be more than just one class Bel that satisfies the constraints above. Worse, for some P, there might even be two such classes that contain mutually inconsistent propositions. We will prove later that this is not so, in fact, for every given P and for every two distinct classes Bel which satisfy the conditions above (relative to that P) it is always the case that one of the two contains the other as a subset. Even with that in place, one would still have to decide which class Bel in the resulting chain of belief classes ought to count as the "actual" belief class as being given by P in order to satisfy the uniqueness part of our intended definition by definite description. But then again, what if there were a largest such class Bel? That class would have all the intended properties, and it would contain every proposition that is a member of any class Bel as above. It would therefore maximize the extent by which *prima facie* beliefs in the sense defined before are realized in terms of actual beliefs. In other words: it would approximate as closely as possible the right-to-left direction of the Lockean thesis that we were forced to drop in view of the logical principles of belief. The class would thus have every right to be counted as *the* class of beliefs at a time t of an agent whose subjective probability measure at that time is P, and no restriction of bounded variables to “natural” classes as in Lewis’ original proposal would be necessary at all. *If such a largest belief class exists*, of course—but as we will prove later, indeed it does.

What we will have found then is that the following is a materially adequate and explicit definition of an agent’s beliefs in terms of the agent’s subjective probability measure:

- If P is an agent’s subjective probability measure at a time t that satisfies the additional constraints..., then a proposition (in the domain of P) is believed as being given by P if and only if it is a member of the largest class Bel of propositions that satisfies the following properties:

  (1) Belief constraints:

    * For all propositions Y, Z: if $Y \in Bel$ and Y logically entails Z, then $Z \in Bel$.
    * For all propositions Y, Z: if $Y \in Bel$ and $Z \in Bel$, then $Bel(Y \cap Z)$.
    * No logical contradiction is a member of Bel.

    $\vdots$

    (Other standard logical principles for Bel and their extensions to conditional belief.)

  (2) Mixed constraints:

    * For all propositions $X \in Bel$, $P(X) > r$.
    * (An extension of this to conditional belief.)

    $\vdots$

    (Additional mixed constraints on P and Bel.)

So we will have managed to define belief simpliciter just in terms of ‘P’ and logical and set-theoretical vocabulary. In fact, it will turn out to be possible to characterize the defining conditions of belief just in terms of a simple and independently appealing quantitative condition on P and elementary set-theoretic operations and relations.

Belief simpliciter will therefore have been reduced to degrees of belief. In the following two sections, we are going to execute this strategy in all formal details. The remaining sections will be devoted to applications and extensions of the theory.

3 The Reduction of Belief I: Absolute Beliefs

The goal of this section and the subsequent one is to enumerate a couple of postulates on quantitative and qualitative beliefs and their interaction; and we will assume that the fictional epistemic agent *ag* that we will deal with has belief states of both kinds available which obey these postulates. The terms 'P' and 'Bel' that will occur in these postulates should be thought of as primitive 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 on subjective probability and belief in the form of postulates or axioms, it will turn out that they will be strong enough to constrain qualitative belief in a way such that the concept of qualitative belief ends up being definable explicitly just on the basis of 'P', that is, in terms of quantitative belief (and a cautiousness parameter) only. When we state the theorems from which this follows, 'P' and 'Bel' will become variables, so that we will able to say: For all P, Bel, it holds that P and Bel satisfy so-and-so if and only if.... Accordingly, in the definition of belief simpliciter itself, 'P' will be a variable again, and 'Bel' will be a variable the extension of which is defined on the basis of 'P' (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.

3.1 Probabilistic Postulates

Consider an epistemic agent ag which we keep fixed throughout the article. Let W be a (non-empty) set of logically possible worlds. Say, 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 demanded to coincide with some power set algebra, instead $\mathfrak{A}$ might simply not count certain subsets of W as propositions at all.

We will extend the standard logical terminology that is normally defined for formulas or sentences to 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, when we refer to the negation of a proposition X we do refer to 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.

Let P be ag's degree-of-belief function (quantitative 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 $X_1, X_2 \ldots$ 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 $P(X) > 0$.

As far as our familiar treatment of conditional probabilities in terms of the ratio formula for absolute or unconditional probabilities is concerned, we should stress that the elegant theory of primitive conditional probability measures (Popper functions) would allow $P(Y|X)$ to be defined and non-trivial even when $P(X) = 0$ (that is, as we will sometimes say, when X is a *zero set* as being given by P). But the theory is still not accepted widely, and we want to avoid the impression that the theory in this paper relies on Popper functions in any sense. We shall nevertheless have occasion to return to Popper functions later in some parts of the paper.

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 in fact not uncontroversial even within the Bayesian camp itself, although in purely mathematical contexts, such as measure theory, σ-additivity is usually beyond doubt (but see Schurz & Leitgeb...); we shall simply take it for granted now. For many practical purposes, $\mathfrak{A}$ may simply be taken to finite, and then σ-additivity reduces to finite additivity again which is indeed 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 simpliciter instead under the mere assumption of finite additivity, that is, assuming just P1 but not P2. 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).

3.2 Belief Postulates

Let us turn now from quantitative belief to qualitative belief: Each belief simpliciter—or more briefly: each *belief*—that *ag* holds at *t* is assumed to have a set in $\mathfrak{A}$ as its propositional content. As a first approximation, 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 are entailed by the modal axiom K and by applications of necessitation to tautologies), *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 $Bel(Y \cap Z)$.

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 certainly involves a good deal of abstraction. On the other hand, if $\mathfrak{A}$ is finite, then the last principle simply reduces to the case of finite conjunctions again. In any case, 4. has the following obvious consequence: There is a *least set* (a strongest proposition) Y, such that $Bel(Y)$; that Y is 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 enables us to represent the sum of *ag*'s beliefs in terms of such a unique proposition or a unique set of possible worlds. In the semantics of doxastic or epistemic logic, our 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 yet another independent reason for assuming 4.: In light of lottery paradox or preface paradox situations, with which we will deal later, it is thought quite commonly that if the set of beliefs simpliciter is presupposed to be closed under conjunction, then this prohibits any probabilistic analysis of belief simpliciter from the start. We will show that beliefs simpliciter can in fact be reduced to quantitative belief *even though 4. expresses the strongest form of closure under conjunction whatsoever that a set of beliefs can satisfy*. 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 (as one can show would be the case if we dropped countable additivity as being one of our assumptions). In a nutshell: 4. prohibits our agent from having anything like an ω-inconsistent set of beliefs.

Finally, we add

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

as our agent *ag* does not believe a contradiction. Once again, this will be granted in order to mimick the same assumption that in epistemic logic is sometimes made: one justification for it is the thought that if a rational agent is shown to believe a contradiction, then he 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.

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 in the following 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. Accordingly, we now call all and only the members Y of Bel_W to be believed absolutely or unconditionally, and Bel_W the *absolute* or *unconditional belief set*.

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. In particular, this will yield an explication of absolute or unconditional belief in terms of subjective probabilities, which is the main focus of this section. In the next section we will add some postulates which will impose constraints even on beliefs conditional on propositions in $\mathfrak{A}$ that contradict B_W, and ultimately we be able to state a corresponding explication of conditional belief in general. 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 we will in the section after this one, 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 one. 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.*

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:

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

On the other hand, we assume the Consistency postulate to hold only for beliefs conditional on W at this point (in the next section this will be generalised). So just as in the case of 5. above, we only demand:

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

By now the axioms should look quite uncontroversial, if given our logical approach to belief. Assuming B1 is unproblematic at least under a suppositional reading of conditional belief: under the (matter of fact) supposition of X, with X being consistent with what the agent believes, the ideally rational agent ag holds X true at time t. Of course, B3 is redundant really in light of B4, but we shall keep it as well for the sake of continuity with the standard treatment of belief. As before, B4 now entails for every $X \in \mathfrak{A}$ for which $\neg Bel(\neg X|W)$ 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' 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. 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. 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, one 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 in this section by means of '$\neg Bel(\neg X|W)$', we see that we may just as well replace this qualification by '$X \cap B_W \neq \varnothing$', and this is what we are going to do in the following.

So far there are no postulates on how belief sets conditional on different propositions relate to each other logically. At this point we demand one such condition to be satisfied which corresponds to the standard AGM (...) postulates K*3 and K*4 on belief revision if B_W takes over the role of AGM's syntactic belief set K, and if 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 $Z \supseteq Y \cap B_W$.

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 is entailed by the AGM postulates on revision by propositions which are consistent with what the agent believes at the time, and it can be justified in terms of 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, 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.

Equivalently:

B6 (Expansion)
For all $Y \in \mathfrak{A}$, such that for all $Z \in \mathfrak{A}$, if $Bel(Z|W)$ then $Y \cap Z \neq \emptyset$:
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:

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

The superset claim that is implicit in the 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 our last formulation of B6 (trivially) that for all $Y \cap B_W \neq \emptyset$, B_Y is non-empty, simply because $B_Y = Y \cap B_W$ in that case.

AGM's K*3 and K*4 have not remained unchallenged, of course. One typical worry is that revising by some new evidence or suppositional information Y may 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 defence 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 defence is: a postulate such as B6 might be true of belief simpliciter, and without it qualitative belief would not have the simplifying power that is essential to it. But there might nothing like it that would hold of quantitative belief, 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. We will return to this issue later where we will see in what sense our theory allows us to reconcile B6 above with the worry about them that we were addressing in this paragraph.

This ends our list of postulates on qualitative belief.

3.3 Mixed Postulates and the Explication of Absolute Belief

Finally, we turn to our promised necessary probabilistic condition for having a belief—the left-to-right direction of the Lockean thesis—and indeed for having a belief conditional on any proposition consistent with all the agent *ag* believes at *t*; this will make *ag*'s 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' 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. The principle says:

$BP1^r$ (Likeliness) For all $Y \in \mathfrak{A}$ such that $Y \cap B_W \neq \emptyset$ 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 generalisation of the left-to-right direction of the Lockean thesis to the case of beliefs conditional on propositions Y which are consistent with all absolute beliefs. The antecedent clause '$P(Y) > 0$' in $BP1^r$ is there 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 (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 an subjective probability of more than r is a necessary condition for a proposition to be believed absolutely, although it will become clear below that this is far from being 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. $BP1^r$ 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) a threshold value that might depend on Y (or on Y and Z). But without any further qualification, a principle such as the latter 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—this r really derives from considerations on the concept of belief itself rather than from considerations on the contents of belief. (Remark: It would be possible to weaken '>' to '≥' in $BP1^r$; not much will depend on it, except that whenever we are going to use $BP1^r$ with $r \geq \frac{1}{2}$ below, one would rather have to choose some $r' > \frac{1}{2}$ instead and then demand that '$\ldots P(Z|Y) \geq r'$' is the case).

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 were invalidated, then there would be Z and Y, such that our agent *ag* believes Z conditional on Y, but where $P(Z|Y) \leq \frac{1}{2}$: 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. While the former is difficult to accept—and the more difficult the lower the value of $P(Z|Y)$—the latter might be acceptable if one presupposes a voluntaristic conception of belief such as van Fraassen's (...). But it would still be questionable then why the agent would choose to believe Z, rather than $\neg Z$, but not choose to assign to Z a higher degree of belief than to $\neg Z$ (assuming this voluntary conception of belief would apply to degrees of belief, too). Richard Foley (...) has argued that the Preface Paradox would show that a principle such as $BP1^{\frac{1}{2}}$ would in fact be too strong: a probability of $\frac{1}{2}$ could not even amount to a necessary condition on belief. We will return to this when we discuss the Lottery Paradox and Preface Paradox in section **??**. Instead of defending $BP1^{\frac{1}{2}}$ or any other particular instance of $BP1^r$ at this point, we will simply move on now, taking for granting one such $BP1^r$ has been chosen. We will argue later that choosing $r = \frac{1}{2}$ is in fact the right choice for the *least possible* threshold value that would give us an account of 'believing that', even though taking any greater threshold value less than 1 would still be acceptable. However, for weaker forms of subjective commitment, such as 'supecting that' or 'hypothesizing that', r ought to be chosen to be less than $\frac{1}{2}$.

For the moment this exhausts our list of postulates (with two more to come later). Let us pause for now and focus instead on jointly necessary and sufficient conditions for our postulates up to this point to be satisfied, which will lead us to our first representation theorem by which pairs $\langle P, Bel \rangle$ that jointly satisfy our postulates get characterized transparently. In order to do so, we will need the following additional probabilistic concept which will turn out to be crucial for the whole theory:

Definition 2 *(P-Stability^r) Let P be a probability measure on a set algebra $\mathfrak{A}$ over W. 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. So a P-stable^r proposition has 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, as we shall see later in section 3.4, 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. With X being P-stable^r, and Y being such that $Y \cap X \neq \varnothing$ and $P(Y) > 0$, it holds 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 with $P(X \cap Y) > r \cdot [P(X \cap Y) + P(\neg X \cap Y)]$, which yields $P(X \cap Y) > \frac{r}{1-r} \cdot P(\neg X \cap Y)$. $X \cap Y$ is some proposition in $\mathfrak{A}$ that is a subset of X, and by assumption it needs to be non-empty. $\neg X \cap Y$ is just some proposition in $\mathfrak{A}$ which is a subset of $\neg X$. If $P(X \cap Y)$ were 0, then the inequality above could not be satisfied irrespective of what $\neg X \cap Y$ would be like; and if $P(X \cap Y)$ is greater than 0, then *a fortiori* $X \cap Y \neq \varnothing$ and also $P(Y) > 0$ are the case. So really X is P-stable^r if and only if for all $Y, Z \in \mathfrak{A}$, such that Y is a subset of X with $P(Y) > 0$ and where Z is a subset of $\neg X$, it holds that $P(Y) > \frac{r}{1-r} \cdot P(Z)$. In words: The probability of any subset of X that has positive probability at all is greater than the probability of any subset of $\neg X$ if the latter is multiplied by $\frac{r}{1-r}$. In the special case in which $r = \frac{1}{2}$, this factor is just 1, and hence 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 $\neg X$. So P-stability^r is also a *separation property*, which divides the class of subpropositions of a proposition from the class of subpropositions of its negation in terms of probability.

Here is a property of P-stable^r propositions X that we will need on various occasions: if $P(X) < 1$, then there is no non-empty $Y \subseteq X$ with $Y \in \mathfrak{A}$ and $P(Y) = 0$. For assume otherwise: 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 hold that $P(X|Y \cup \neg X) = \frac{P(X \cap Y)}{P(Y \cup \neg X)} > r$, which contradicts $P(X \cap Y) \leq P(Y) = 0$. For the same reason, non-empty propositions of probability 0 cannot be P-stable^r, or in other words: non-empty P-stable^r propositions X have positive probability.

Using this new concept, we can show 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 $\mathfrak{A}$ as explained above, let $P : \mathfrak{A} \to [0, 1]$, and let $0 \leq r < 1$. 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 a (uniquely determined) $X \in \mathfrak{A}$, 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 $Z \in \mathfrak{A}$:*

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

(and hence, $B_W = X$).

Proof. From left to right: P1 is satisfied by assumption. Now we let $X = B_W$, where B_W exists and has the intended 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 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 that there is an $X' \in \mathfrak{A}$, such $X' \neq X$, 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 $X' = B_W$, and hence with $X = B_W$ from above that $X' = X$, which is a contradiction.

From right to left: 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 characterisation 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 B1–B4 are satisfied by the assumed characterisation 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 $P(Z \cap Y) \geq P(Y \cap X)$. From X being P-stable^r and $P(Y) > 0$ we have $P(X|Y) > r$.

Taking this 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 equivalence being satisfied.

This simple theorem will prove to be fundamental for all subsequent arguments in this paper. We start by exploiting it first 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: take X to be W itself. $P(W) > 0$ and P-stability^r follow then 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 $Y \cap W = Y \neq \emptyset$: 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 absolutely thereby by our agent is 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 what is more important, we actually find that a much more general pattern is emerging: Let P be given again as before. Now choose *any* non-empty P-stable^r proposition X, and define conditional belief in all cases in which $Y \cap X \neq \emptyset$ 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 probability of less than 1: Even if $P(X) < 1$, 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. It does *not* follow that if a proposition has a probability in the interval $[P(X), 1]$, then this just by itself implies that the proposition is also believed absolutely, since it is not entailed that the proposition is then also a superset of the P-stable^r proposition X that had been chosen initially.

Since P-stable^r propositions play such a distinguished role in this, the questions arise: Do P-stable^r sets other W exist at all for many P? More generally: Do non-trivial exist for many P, that is, such with a probability strictly between 0 and 1? Subsection 3.4 below will show that the answers are affirmative. And how difficult is it to determine whether a proposition is a non-empty P-stable^r set?

About the last question: At least in the case where W is finite, it turns out not to be difficult at all: Let $\mathfrak{A}$ be the power set algebra on W, and let P be defined on $\mathfrak{A}$. By definition, X is P-stable^r if and only if for all $Y \in \mathfrak{A}$ with $Y \cap X \neq \varnothing$ and $P(Y) > 0$, $P(X|Y) = \frac{P(X \cap Y)}{P(Y)} > r$. We have seen already that all sets with probability 1 are P-stable^r. So let us focus just on how to generate all non-empty P-stable^r sets X that have a probability of less than 1. As we observed before, such sets do not contain any subsets of probability 0, which in the present context means that if $w \in X$, $P(\{w\}) > 0$.

For any given such non-empty X with $P(X) < 1$, as we have shown before, 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:

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

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 P-stable^r subsets of W. We only give a sketch for the case $r = \frac{1}{2}$: All sets of probability 1 are P-stable^r, so we disregard them. All other non-empty P-stable^r sets do not have singleton subsets of probability 0, so let us also disregard all worlds whose singletons are zero sets. Assume that after dropping all worlds with zero probabilistic mass, there are exactly n members of W left, and $P(\{w_1\}), P(\{w_2\}), \ldots, P(\{w_n\})$

is already in (not necessarily strictly) decreasing order. If $P(\{w_1\}) > P(\{w_2\}) + \ldots + P(\{w_n\})$ then $\{w_1\}$ is P-stable$^{\frac{1}{2}}$, and one moves on to the list $P(\{w_2\}), \ldots, P(\{w_n\})$. If $P(\{w_1\}) \leq P(\{w_2\}) + \ldots + P(\{w_n\})$ then consider $P(\{w_1\}), P(\{w_2\})$: If both of them are greater than $P(\{w_3\}) + \ldots + P(\{w_n\})$ then $\{w_1, w_2\}$ is P-stable$^{\frac{1}{2}}$, and one moves on to the list $P(\{w_3\}), \ldots, P(\{w_n\})$. If either of them is less than or equal to $P(\{w_3\}) + \ldots + P(\{w_n\})$ then consider $P(\{w_1\}), P(\{w_2\}), P(\{w_3\})$: And so forth, until the final P-stable$^{\frac{1}{2}}$ set W has been generated. This recursive procedure yields precisely all non-empty P-stable$^{\frac{1}{2}}$ sets of probability less than 1 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 terminate in finite time.)

What Theorem 3 gives us therefore is not just a construction procedure but even, in the finite case, an *efficient* construction procedure for a class *Bel* from any given probability measure P, so that the two together satisfy all of our postulates. P2 still has not played a role so far. But Theorem 3 does more: it also shows that whatever our agent *ag*'s actual probability measure P and his actual class *Bel* of conditionally believed pairs of propositions are like, as long as they satisfy our postulates from above, 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 were 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 identifying the very P-stable^r proposition X that would give us the agent's actual belief class *Bel*, we could define explicitly the set of all pairs $\langle Z, Y\rangle$ in that class *Bel* for which $Y \cap X \neq \emptyset$ holds by means of that proposition X and thus, ultimately, by the given measure 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.

So are we in the 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, we need the following additional theorem first:

Theorem 4 *Let* $P : \mathfrak{A} \to [0, 1]$ *such that* P1 *is satisfied. Let* $r \geq \frac{1}{2}$. *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 *in* $\mathfrak{A}$ *with probability less than 1, that is, there is* *no countably infinite sequence*

$$X_0 \supsetneq X_1 \supsetneq X_2 \supsetneq \ldots$$

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

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

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' being taken the other way round 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 showed before P-stable^r propositions with probability less than 1 do not have non-empty subsets with probability 0. 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 $(X \cap \neg X') \cup \neg X$ being positive. The same must hold, *mutatis mutandis*, for $P(X'|(X' \cap \neg X) \cup \neg X')$. So we have

  $$P(X|(X \cap \neg X') \cup \neg X) > r \geq \frac{1}{2}$$

  and

  $$P(X'|(X' \cap \neg X) \cup \neg X') > r \geq \frac{1}{2},$$

  where $r \geq \frac{1}{2}$ by assumption.

  Next we show that

  $$P(X \cap \neg X') > P(\neg X).$$

  For suppose otherwise, that is $P(X \cap \neg X') \leq P(\neg X)$: 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:

$\frac{1}{2} > P(\neg X|(X \cap \neg X') \cup \neg X) = \frac{P(\neg X)}{P((X \cap \neg X') \cup \neg X)} \geq \frac{P(X \cap \neg X')}{P((X \cap \neg X') \cup \neg X)} = P(X \cap \neg X'|(X \cap \neg X') \cup \neg X)$, by our initial supposition; 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

$$P(X' \cap \neg X) > P(\neg X').$$

Finally, from this (and P1) we can derive: $P(X \cap \neg X') > P(\neg X) \geq P(X' \cap \neg X) > P(\neg X') \geq P(X \cap \neg X')$, which is a contradiction.

- Ad IV: Assume for contradiction that there is a sequence $X_0 \supsetneq X_1 \supsetneq X_2 \supsetneq \ldots$ 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 $i \geq 0$, 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 $P(B) = P(\bigcup_{i=o}^{\infty} A_i) = \sum_{i=o}^{\infty} P(A_i)$ 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 $\frac{P(A_i)}{P(A_i)+P(\neg X_0)} = \frac{P(X_0 \cap (A_i \cup \neg X_0))}{P(A_i \cup \neg X_0)} = P(X_0|A_i \cup \neg X_0)$ 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 two 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. Secondly, in light 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. Furthermore, 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 the latter are all P-stable^r. If we only look at non-empty P-stable^r propositions with a probability of less than 1, we find therefore that they constitute a sphere system that satisfies the Limit Assumption (by well-orderedness) for every proposition in $\mathfrak{A}$, in the sense of Lewis (...). Note that P2 (Countable Additivity) was needed in IV. in order to derive the well-foundedness of the chain of P-stable^r propositions of probability less than 1.

For given P (and given $\mathfrak{A}$ and W), such that P satisfies P1–2, and for given $r \in [0,1)$, let us denote the class of all non-empty P-stable^r propositions X with $P(X) < 1$ by: $\mathcal{X}_P^r$. We know from Theorem 4 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 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 $\emptyset$). In case the union of all X_α^r is W, each world $w \in W$ 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.

Furthermore, by P1–2, Theorem 4, and the fact that no non-empty P-stable^r 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 $P(X)$, that is: 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). Accordingly, since every ordinal number has a unique successor, there is a bijective mapping between the set of intervals of the form $(P(X_\alpha^r), P(X_{\alpha+1}^r))$ for $\alpha < \beta_P^r$ and the set β_P^r. See Figure 1.

From this we can determine a boundary for the ordinal type of β_P^r:

Observation 5 *Let P be a countably additive probability measure on a σ-algebra $\mathfrak{A}$ over W. Let $\frac{1}{2} \leq r < 1$.*

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 $\beta_P^r \geq \omega + 1$: 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 $0 \leq n < \omega$, let $Y_n = X_{n+1}^r \setminus X_n^r$, and let $Z_n = \bigcup_{m \geq n} Y_m^r$. We know that for all n it holds that $Z_n \in \mathfrak{A}$, by Theorem 4 and the definition of 'X_α^r' it is the case that $Z_n \subseteq X_\omega^r$, by assumption we have $P(X_\omega^r) < 1$, and furthermore $P(Z_n) < 1$ and the sequence (Z_n) is strictly monotonically decreasing. So there is a sequence $X_\omega^r \supsetneq Z_0 \supsetneq Z_1 \supsetneq \ldots$ of sets in $\mathfrak{A}$ with probability less 1, with X_ω^r being P-stable^r, in contradiction with IV of Theorem 4. ■

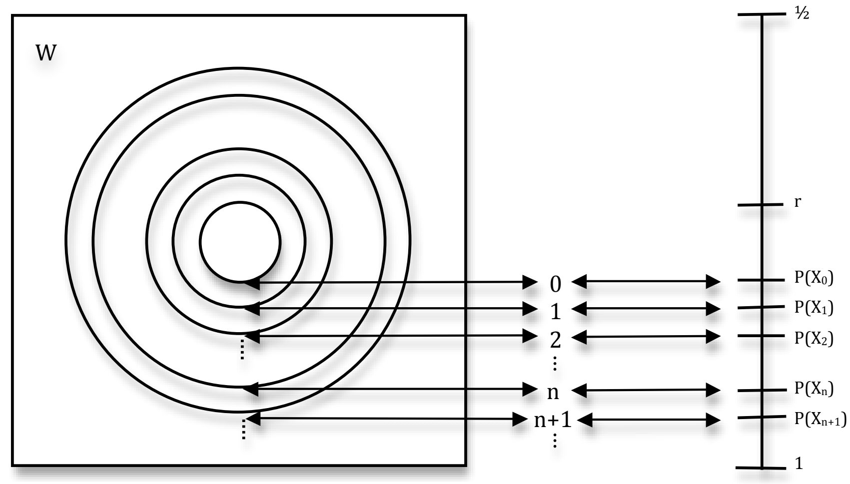

Figure 1: P-stable sets for $r \geq \frac{1}{2}$

We also find that, given P is countably additive, if there are countably infinitely many non-empty P-stable^r propositions X with probability less than 1, 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 $Y \cap \bigcup_{\alpha<\omega} X^r_\alpha \neq \emptyset$ for $Y \in \mathfrak{A}$ with $P(Y) > 0$, then there must be an X^r_α with $\alpha < \omega$, such that $Y \cap X^r_\alpha \neq \emptyset$. Because X^r_α is P-stable^r, it follows that $P(X^r_\alpha|Y) > r$. But by P1, $P(\bigcup_{\alpha<\omega} X^r_\alpha|Y) \geq P(X^r_\alpha|Y)$, 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 $\omega + 1$, which was ruled out by Observation 5. So $P(\bigcup_{\alpha<\omega} X^r_\alpha) = 1$.

Since, as we saw before, no non-empty P-stable^r propositions X with probability less than 1 contains a non-empty zero set as a subset, that union could not do so either. So in the case in which β^r_P is infinite, that union of all non-empty P-stable^r propositions with probability less than 1 would then have to be the least P-stable^r proposition with probability 1.

Now back to our remaining open questions. Let us start with: what should we choose as r?

For the proof of III. in Theorem 4 it was crucial that $r \geq \frac{1}{2}$. 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 $X \cap X' = \emptyset$. This

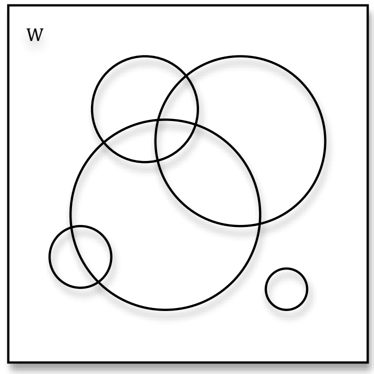

Figure 2: P-stable sets for $r < \frac{1}{2}$

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. It seems advisable then, for the sake of a better theory, to demand that $r \geq \frac{1}{2}$, for this will allow us to derive as a *law* that a situation such as that cannot occur. Of course, this is far from being a knock-down argument against $r < \frac{1}{2}$, but it certainly puts a bit of methodological pressure on it. 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 (..., p. 350) puts it, the assumed role of full belief is "to form a single, unequivocally endorsed picture of what things are like": If $r \geq \frac{1}{2}$, 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 quite reasonable for him to have assigned to that proposition a probability that is greater than the probability of its negation—we do have a plausible case against choosing r in that way. (But we will see later that $r < \frac{1}{2}$ is an attractive choice if '*Bel*' is taken to express not belief but some weaker epistemic attitude.)

Apart from presupposing $r \geq \frac{1}{2}$, 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 6 *Let P be a probability measure on an algebra $\mathfrak{A}$ over W. Let $X \in \mathfrak{A}$, and assume that $\frac{1}{2} \leq r < s < 1$. 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 $Y \cap X \neq \emptyset$, $P(X|Y) > s$. But then it also holds for all $Y \in \mathfrak{A}$ with $Y \cap X \neq \emptyset$ that $P(X|Y) > r$, since $r < s$ by assumption, so X is P-stable^r as well. ■

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 $\frac{1}{2} \leq r < 1$, 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. 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 $r \geq \frac{1}{2}$*, without making any particular choice of the value of 'r' mandatory.

With that 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 subjective probability measure P (and a threshold value r)? We need one more postulate before we answer this.

Degrees of belief conditional on a proposition of probability 0 are brought in line with beliefs conditional on a contradiction in the following manner:

BP2 (Zero Supposition) For all $Y \in \mathfrak{A}$: If $Y \cap B_W \neq \emptyset$ and $P(Y) = 0$, then $B_Y = \emptyset$.

Since P is an absolute probability measure that does not allow for conditionalization on a proposition of probability 0 at all, it makes sense to restrict belief simpliciter accordingly in the way that supposing any such proposition of probability 0 amounts to believing a contradiction. For intuitively there is no reason to think that supposing a proposition qualitatively ought to less zero-intolerant—using Jonathan Bennett's corresponding term (...) which he applies to indicative conditional whose antecedent has subjective probability 0—than the quantitative supposition of a proposition. 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 Popper functions 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 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, 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 is indeed usually taken for granted or otherwise W would be redefined by dropping all worlds whose singleton sets have zero probability.

Here is an important consequence of BP2: Let $Y \in \mathfrak{A}$ be such that $P(Y) = 1$. Y must then have non-empty intersection with B_W, in light of P1 and $P(B_W) > 0$. Therefore, by B6, $B_Y = Y \cap B_W \subseteq B_W$. Assume that Y is a proper subset of B_W: then both $Y \cap B_W$ and $\neg Y \cap B_W$ are non-empty. Since $P(Y) = 1$, it follows that $P(\neg Y) = 0$ and hence with BP2: $B_{\neg Y} = \emptyset$. But since $\neg Y$ has non-empty intersection with B_W, BP6 entails that $B_{\neg Y} = \neg Y \cap B_W$. Therefore, $\neg Y \cap B_W = \emptyset$, which contradicts $\neg Y \cap B_W$ being non-empty. So we find that by BP2 (and the rest of our postulates), every $Y \in \mathfrak{A}$ for which $P(Y) = 1$ holds is such that $B_Y = B_W$. This also entails that, since $B_Y \subseteq Y$ for all such Y by the definition of 'B_Y', if B_W has probability 1 itself, then B_W must be the least proposition in $\mathfrak{A}$ with probability 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 we given just *ag*'s subjective probability measure P (and a threshold value r). As explained already in section 3, apart from satisfying our postulates, the class *Bel* ought to be so that the resulting class of absolute beliefs is maximised, as this approximates *prima facie* belief, and hence, the right-to-left direction of the original Lockean thesis, to the greatest possible extent. This corresponds to the following postulate:

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 with respect to

the class of absolute beliefs, that is, pairs of the form $\langle Z, W\rangle$, that it determines.

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

The logical character of BP3 is obviously different from the one of our previous postulates, but then again adding postulates that maximize or minimize classes subject to axiomatic constraints is of course not unheard of; for example, famously, Hilbert (...) uses this strategy in his axiomatization of geometry.

The term 'the largest' in BP3 is well-defined given the postulates P1–P2, B1–B6, BP1^r, BP2 Theorem 3, or in view of Theorem 4, and by what we pointed out before: Because of Theorem 3, B_W must be a non-empty P-stable^r proposition in $\mathfrak{A}$ in order to satisfy P1, B1–B6, and BP1^r. If there is at least one non-empty P-stable^r proposition with probability less than 1, then we know that amongst all the non-empty P-stable^r propositions that are candidates for the maximally strong believed proposition B_W according to Theorem 3 (which relied on P2), there must be a least one by Theorem 4: this least P-stable^r proposition X_{least}, which then has a probability of less than 1, and which does not have any non-empty zero sets as subsets and hence satisfies BP2, must therefore determine the largest class of absolute beliefs once II. in Theorem 3 is turned into a (partial) definition of conditional belief again, since its class of supersets is the largest one possible. On the other hand, if there are no non-empty P-stable^r propositions with probability less than 1, then by P1, B1–B6, and BP1^r again, $P(B_W)$ must be a non-empty P-stable^r proposition with probability 1, and from our considerations on BP2 above we know that B_W must really be the least set of probability 1 in $\mathfrak{A}$.

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 $Y \cap B_W \neq \emptyset$. 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 restriction of generality, $B'_W \subsetneq B''_W$. But then, first of all, the class of all pairs $\langle Z, Y\rangle$ for which $Y \cap B'_W \neq \emptyset$ is distinct from the class of all pairs $\langle Z, Y\rangle$ for which $Y \cap B''_W \neq \emptyset$, 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, when 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 light 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 7 *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 $Y \cap X_{least} \neq \emptyset$, 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 $Z \in \mathfrak{A}$:*

$$Bel(Z \,|\, W) \text{ if and only if } Z \supseteq X_{least}.$$

Where the previous postulate was reminiscent of Hilbert's axiomatisation of geometry, with respect to its open parameter r the last corollary is closer in spirit to something like Zermelo's (...) quasi-categoricity result for second-order set theory: according to Zermelo's 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.

VI. of Corollary 7 can now be turned into an explicit definition of all relevant conditional beliefs just on the basis of P (and logical and set-theoretic notions). Since in the next section we will extend this result to arbitrary conditional beliefs, whether or not they are beliefs conditional on proposition that are consistent with what the agent believes, we refrain from stating the resulting definition here. However, we do exploit Corollary 7 by deriving from it a particularly important special case: the concept of absolute belief can be defined explicitly in terms of P alone.

In order to do so, we will take one final step. We restrict the probability measures P that we are interested in such that the existence claim in VI. is always satisfied. While our explicit definition of belief will then just hold conditional on that additional restriction, since the restriction is not overly demanding in our belief context (though it would be in other contexts, say, in measure theory, where one needs measures for integration), we will still end up with a definition that assigns the right reference to '*Bel*' for a wide range of subjective probability measures.

This is thus the restriction on P that we 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 $P(1) = 0$: $X \subseteq Y$. 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 which is a proper subset of X.

Given this Least Certain Set Restriction, there is always a least non-empty P-stable^r proposition in $\mathfrak{A}$: 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 $\mathfrak{A}$.

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 that 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: for all $X \in \mathfrak{A}$, $P(X) = 0$ if and only if $X = \emptyset$. Here the empty happens to be the least set of probability 1. Regularity (Strict Coherence) does not enjoy general support, even though Carnap, Shimony, Stalnaker and others 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 favours Regularity. (But see Levi... for contrary arguments.)

- All countably additive probability measures on a countably infinite σ-algebra: The conjunction of all sets of probability 1 is then a countably infinite conjunction, so it is a member of the given σ-algebra, and it is again the least set of probability 1.

These examples demonstrate that a great variety of probability measures satisfy P1, P2, and our additional constraint, and many—if not most—of the typical philosophical toy examples of subjective probability measures are covered by these examples.

We end up with the following materially adequate explicit definition of absolute belief for countably additive probability measures that satisfy this additional constraint of the Least Certain Set Restriction:

Definition 8 *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 $\frac{1}{2} \leq r < 1$:

Y is believed (to a cautiousness degree of r) as being given by P if and only if Y is a superset of X_{least}.

By 'materially adequate' we mean here: By Corollary 7, since 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, BP2 being acceptable for the sake of simplicity, our definition of belief is true if given a probability measure that satisfies the Least Certain Set Restriction, if the definition is taken as a descriptive sentence. What is more, since all of P1, B1–B6, BP1^r, BP3 are not just true but even conceptually necessary or analytic of belief, the definition is so as well (conditional on the presupposition of P2 and BP2).

Note that from the theory above we know that the definiens 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 the belief predicate in any way.

If we finally define for any given $P : \mathfrak{A} \to [0, 1]$, $Y \in \mathfrak{A}$ is believed *a priori* as being given by P if and only if $P(Y) = 1$, then we end up with three notions of belief of increasing strength for all P that satisfy P1, P2, and the Least Certain Set Restriction: *prima facie* belief, belief (to a cautiousness degree of r), and belief *a priori*. For any two of these concepts, under very special circumstances, that is, for very special P, they can in fact determine precisely the same beliefs (later we will deal with an example). But under "normal" circumstances, for realistic P, they will differ extensionally, and belief in the sense of Definition 8 is the concept that we offer as an explication of our pre-theoretic notion of qualitative belief.

3.4 Examples

Finally, here are some examples. In all of them, $\mathfrak{A}$ is simply the full power set algebra of W. If W contains exactly two worlds, then the situation is trivial insofar as for given $\frac{1}{2} \leq r < 1$,

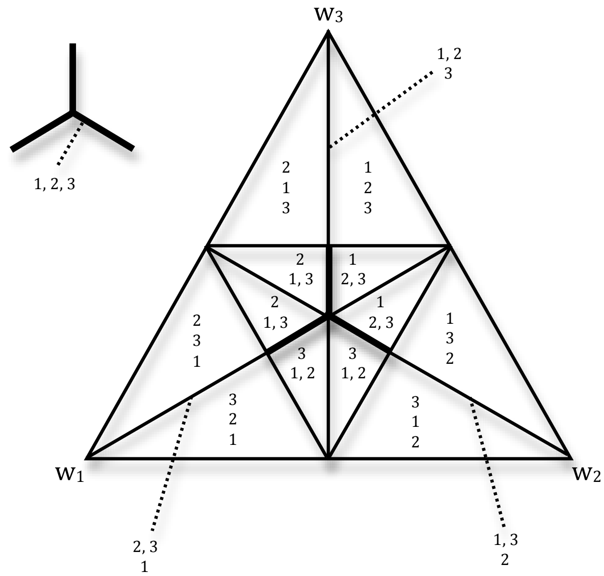

Figure 3: Rankings determined by P

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 of 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 a 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 Figure 3.

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 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 before, 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. 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. 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 seen, these form a sphere 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 8. 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 Figure 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.

If $r > \frac{1}{2}$, then a diagram similar to Figure 3 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 Figure 3 why 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, what is true in general: 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 generalise the distinguished line segments in our diagram to the higher-dimensional case.

For analogous reason, 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, that is, geometrical measure, 0. This is because, for any such P: If there were a unique world whose singleton had least probability, then W without that world would be P-stable$^{\frac{1}{2}}$; 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. We conclude: *Almost all probability measures over a finite algebra have a least P-stable$^{\frac{1}{2}}$ set with a probability less than 1.*

Here is another example with 7 worlds and concrete numbers: Let $W = \{w_1, \ldots, w_7\}$ and $P(\{w_1\}) = 0.54$, $P(\{w_2\}) = 0.342$, $P(\{w_3\}) = 0.058$, $P(\{w_4\}) = 0.03994$, $P(\{w_5\}) =$

0.018, $P(\{w_6\}) = 0.002$, $P(\{w_7\}) = 0.00006$. Then the resulting sphere system of non-empty P-stable$^{\frac{1}{2}}$ sets is: $\{w_1\}, \{w_1, w_2\}, \{w_1, \ldots, w_4\}, \{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}, \{w_1, \ldots, w_7\}$. However, if we switch e.g. to $r = \frac{3}{4}$, then the corresponding sphere system of non-empty P-stable$^{\frac{3}{4}}$ sets is: $\{w_1, w_2\}, \{w_1, \ldots, w_4\}, \{w_1, \ldots, w_5\}, \{w_1, \ldots, w_6\}, \{w_1, \ldots, w_7\}$. In line with Observation 6, the latter sphere system is a subclass of the former one. With a cautiousness degree of $r = \frac{1}{2}$, the proposition $\{w_1\}$ is the strongest one that is believed as being given by P, while relative to a cautiousness degree of $r = \frac{3}{4}$, the proposition $\{w_1, w_2\}$ is the strongest one that is believed as being given by the same probability measure, as entailed by Definition 8.

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

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

It is also easy to see that every finite sphere system can be realized in this way in terms of P-stable$^{\frac{1}{2}}$ propositions of probability less than 1, and hence every AGM-style belief revision operator on a logically finite language. So there are really *lots* of different types of sphere systems of P-stable$^{\frac{1}{2}}$ propositions.

Once we have covered conditional belief in full in the next section, we will return to some of these examples and analyse them in terms of conditional belief accordingly. Moreover, eventually, we will give some of these examples an intended interpretation by assuming that the possible worlds in question satisfy particular statements.

4 The Reduction of Belief II: Conditional Beliefs

Now we finally generalise 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.

P1–P2 remain unchanged, of course. Our generalisations of B1–B5 simply result from dropping the antecedent '$\neg Bel(\neg X|W)$' condition that 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 for belief revision if 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 $Z \supseteq Y \cap B_X$.

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 it can be justified in terms of standard possible worlds accounts of similarity orderings (as for David Lewis' conditional logic) or 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 thus 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 generalised 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 generalise BP2 in the same way, and additionally we strengthen it by assuming also the converse of the resulting generalisation:

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

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: If $B_Y = \varnothing$ and $Y \cap B_W \neq \varnothing$, then $P(Y) = 0$—is that the resulting principle would have been empty: if $Y \cap B_W \neq \varnothing$, then by BP6 the proposition B_Y would have to be non-empty, in contradiction with $B_Y = \varnothing$, so the antecedent of that direction would always have to be false.

We have seen in the last section that BP2, and hence BP2*, entails (given the other postulates): all $Y \in \mathfrak{A}$ for which $P(Y) = 1$ holds are such that $B_Y = B_W$, and B_W is the least proposition in $\mathfrak{A}$ of probability 1. The additional strengthening has it that the propositions the supposition of which leads to inconsistency qualitatively are precisely those for which conditionalization is undefined quantitatively. As mentioned before, if we had started with primitive conditional probability measures, which do allow for conditionalization on zero sets, then BP2* should not be taken for granted, but in the context of absolute probability measures BP2* is natural to postulate in order to treat qualitative and quantitative supposition similarly.

We are now ready to prove the main result of our theory on conditional beliefs in general. The "soundness" direction of the following representation theorem incorporates the corrsponding direction of Grove's (...) 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 in ... into our present context and apply it, since his construction of spheres involves taking unions of propositions that might not be members of our σ-algebra $\mathfrak{A}$ anymore. That is why the proof of that part of the theorem differs from Grove's proof quite significantly.

Here is the theorem:

Theorem 9 *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:*

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 (uniquely determined) 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 $Y \cap X \neq \emptyset$ holds (which exists), then for all $Z \in \mathfrak{A}$:*

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

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

Proof. The right-to-left 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 $Y \cap X \neq \emptyset$. The existence of that least member follows from Theorem 4, from the fact that every non-empty P-stable^r propositions with probability less than 1 is a subset of the least set in $\mathfrak{A}$ with probability 1, 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. The proof of B6* is straight forward (and analogous to Groves Theorem in...), as is the proof of BP2*.

So we can concentrate on the left-to-right direction: P1–P2 are satisfied by assumption. Now we define $\mathcal{X}$ by transfinite recursion as the class of all sets X_α of the following kind: For all ordinals $\alpha < \beta^r_P + 1$ (the successor ordinal of the ordinal that was defined in the last section), 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 transfinite induction. For assume that all X_γ are in $\mathfrak{A}$ for $\gamma < \alpha < \beta^r_P + 1$: by the results of the last section, β^r_P 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 $B_{W \setminus \bigcup_{\gamma<\alpha} X_\gamma} \in \mathfrak{A}$; hence, $X_\alpha \in \mathfrak{A}$.

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

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

(d) For all $\alpha < \beta^r_P + 1$: For all $Y \in \mathfrak{A}$ with $Y \cap X_\alpha \neq \emptyset$, it holds that $B_Y \subseteq X_\alpha$. This is because: If $Y \cap X_\alpha \neq \emptyset$, then by (c) there is a $\gamma \leq \alpha$, such that $Y \cap B_{W \setminus \bigcup_{\delta<\gamma} X_\delta} \neq \emptyset$, 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 $Y \cap \bigcup_{\delta<\gamma} X_\delta = \emptyset$, and hence $Y \subseteq W \setminus \bigcup_{\delta<\gamma} X_\delta$. By B6*, $B_{[W \setminus \bigcup_{\delta<\gamma} X_\delta] \cap Y} = Y \cap B_{W \setminus \bigcup_{\delta<\gamma} X_\delta}$, which is equivalent to $B_Y = Y \cap B_{W \setminus \bigcup_{\delta<\gamma} X_\delta}$ by what we have shown before. Finally, because $Y \cap B_{W \setminus \bigcup_{\delta<\gamma} X_\delta} \subseteq B_{W \setminus \bigcup_{\delta<\gamma} X_\delta} \subseteq X_\alpha$ by (c) again, it follows that $B_Y \subseteq X_\alpha$.

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

Proof: First of all, either there P-stable^r propositions in $\mathfrak{A}$ with probability less than 1 or not: If so, then as shown in the last section their (countable) union is the least proposition $X \in \mathfrak{A}$ with probability 1; if not, then as observed before, BP2* entails with the other postulates that B_W is the least $X \in \mathfrak{A}$ of probability 1. In either case, there exists the least proposition $X \in \mathfrak{A}$ with probability 1.

Secondly, we turn to the proof of: $X \in \mathcal{X}$, and X is the only member of $\mathcal{X}$ with probability 1. Assume for contradiction that all sets X_α with $\alpha < \beta^r_P + 1$ have probability less than 1. Since they are all P-stable^r by (e), it follows from (b) that there is a well-ordered chain of (not necessarily strictly) increasing P-stable^r sets of probability less than 1, where the length of that chain is $\beta^r_P + 1$. That chain could not be a chain of *strictly* increasing P-stable^r sets of probability less than 1, by Observation 5 and by the definition of β^r_P which is the ordinal type of *all* P-stable^r sets of probability less than 1 whatsoever. So there must be $\alpha < \alpha' < \beta^r_P + 1$, such that $X_\alpha = X_{\alpha+1}$. Hence, by (b) again: $X_\alpha = X_\alpha \cup B_{W \setminus X_\alpha}$. Because $P(X_\alpha) < 1$, it holds that $P(W \setminus X_\alpha) > 0$ by P1, so by the right-to-left direction of BP2* it follows that $B_{W \setminus X_\alpha} \neq \emptyset$. Since $B_{W \setminus X_\alpha} \subseteq W \setminus X_\alpha$ by the definition of '$B_{W \setminus X_\alpha}$' and B1*–B4*, a contradiction follows. Hence, we have that there must be at least one set X_α with $\alpha < \beta^r_P + 1$ that has probability 1. Since $\beta^r_P + 1$ is an ordinal, there must be a least $\alpha < \beta^r_P + 1$, such that $P(X_\alpha) = 1$. By Observation 5, either β^r_P is finite or equal to ω. We will deal with these cases separately:

In the former case, there is a $\gamma < \beta^r_P + 1$, such that $\alpha = \gamma + 1$, and, by (b) again: $X_\alpha = X_\gamma \cup B_{W \setminus X_\gamma}$. If there is a set $Y \in \mathfrak{A}$, such that $P(Y) = 1$ and Y is not a superset of X_α, then $X_\alpha \cap \neg Y$ is non-empty, where $X_\alpha \cap \neg Y$ is a zero set since $\neg Y$ is. Because X_γ is P-stable^r with a probability of less than 1, it cannot contain any non-empty zero set, as shown in the previous section. So $X_\gamma \cap \neg Y$ is empty, and therefore $B_{W \setminus X_\gamma} \cap \neg Y$ must be non-empty. This implies by B6*: $B_{[W \setminus X_\gamma] \cap \neg Y} = \neg Y \cap B_{W \setminus X_\gamma}$. But $[W \setminus X_\gamma] \cap \neg Y$ is a set of probability 0 since $\neg Y$ is, which means by BP2* that $B_{[W \setminus X_\gamma] \cap \neg Y}$ is empty, which is a contradiction. Therefore, all $Y \in \mathfrak{A}$ with $P(Y) = 1$ are supersets of X_α, and so X_α is the least set in $\mathfrak{A}$ of probability 1. Furthermore, if $\alpha < \beta^r_P$, then $X_{\alpha+1} \in \mathcal{X}$, and by (b) again: $X_{\alpha+1} = X_\alpha \cup B_{W \setminus X_\alpha}$. But $W \setminus X_\alpha$ has probability 0 then, hence by BP2* it must hold that $B_{W \setminus X_\alpha}$ is empty, and so $X_{\alpha+1} = X_\alpha$. Thus, X_α, the least set in $\mathfrak{A}$ of probability 1, remains to be the only set in $\mathcal{X}$ with probability 1.

In the other case, where $\beta^r_P = \omega$, if $\alpha < \omega$, then by the same reasoning as before, X_α, the least set in $\mathfrak{A}$ of probability 1, remains to be the only set in $\mathcal{X}$ with probability 1. Finally, if $\alpha = \omega$, then all sets X_γ with $\gamma < \omega$ must be P-stable^r sets with probability less than 1. If these sets are not pairwise distinct, they must be equal from some ordinal less than ω by (b), hence there is such an X_γ, such that $X_\alpha = X_\gamma \cup B_{W \setminus X_\gamma}$, which entails just as before that X_α is the least set in $\mathfrak{A}$ of probability 1 and the only set in $\mathcal{X}$ that has probability 1. On the other hand, if the sets X_γ with $\gamma < \omega$ are pairwise distinct, then by Observation 5, their union $\bigcup_{\gamma<\omega} X_\gamma$ must be equal to the union of all P-stable^r sets with probability less than 1. And as shown immediately after Observation 5, that union is the least set in $\mathfrak{A}$ of probability 1. By definition, $X_\alpha = X_\omega = \bigcup_{\gamma<\omega}[X_\gamma] \cup B_{W \setminus \bigcup_{\gamma<\omega} X_\gamma}$, and since $W \setminus \bigcup_{\gamma<\omega} X_\gamma$ is then a zero set, $B_{W \setminus \bigcup_{\gamma<\omega} X_\gamma}$ is empty as follows from BP2*, and therefore X_α is again identical to the least set in $\mathfrak{A}$ with probability 1, and it is the only set in $\mathcal{X}$ with probability 1 since $\alpha = \omega = \beta^r_P$ is the last ordinal less than $\beta^r_P + 1$ in the present case. This concludes (f): the least set X with probability 1 is a member of $\mathfrak{A}$ and indeed of $\mathcal{X}$, and X is the only member of $\mathcal{X}$ with probability 1.

Now 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 $\alpha < \beta^r_P + 1$ be least, such that $Y \cap X_\alpha \neq \emptyset$: because of (b), X_α is then with respect to the subset relation the least member of $\mathcal{X}$ for which this holds. We 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), $X_\alpha = \bigcup_{\gamma<\alpha} B_{W \setminus \bigcup_{\delta<\gamma} X_\delta} \cup B_{W \setminus \bigcup_{\gamma<\alpha} X_\gamma}$. If Y had non-empty intersection with any set of the form $B_{W \setminus \bigcup_{\delta<\gamma} X_\delta}$ for $\gamma < \alpha$, then $Y \cap X_\gamma \neq \emptyset$, by (c) again, in contradiction with the way in which α was defined before. Therefore, $Y \cap X_\alpha = Y \cap B_{W \setminus \bigcup_{\gamma<\alpha} X_\gamma} \neq \emptyset$. The latter implies with B6* that $B_{[W \setminus \bigcup_{\gamma<\alpha} X_\gamma] \cap Y} = Y \cap B_{W \setminus \bigcup_{\gamma<\alpha} X_\gamma}$. As in the proof of (d), $Y \cap \bigcup_{\gamma<\alpha} X_\gamma$ is empty, and thus $[W \setminus \bigcup_{\gamma<\alpha} X_\gamma] \cap Y = Y$. So we have $B_Y = Y \cap B_{W \setminus \bigcup_{\gamma<\alpha} X_\gamma} = Y \cap X_\alpha$ and we are done.

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

Uniqueness follows from: if there are two such classes $\mathcal{X}, \mathcal{X}'$ with the stated properties, then they must differ with respect to at least one P-stable^r sets with probability less than 1. Without restriction 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 just as before 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 $B_{W \setminus X_\gamma} = X_\alpha \setminus X_\gamma$ could not be the same as being given by $\mathcal{X}$ and $\mathcal{X}'$, which would again be a contradiction. ■

Theorem 9 generalises Theorem 3 of the last section to conditional beliefs in general—Theorem 3 simply dealt with the special case of a sphere system of just one P-stable^r set.

It remains to generalise 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*′: $Bel \supseteq Bel'$.

Using this, we can derive:

Corollary 10 *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 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 $Y \cap X \neq \emptyset$ holds (which exists), then for all $Z \in \mathfrak{A}$:*

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

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

This follows immediately from Theorem 9, except that we have to show: adding 'BP3*' to I. of Theorem 9 is equivalent to determining $\mathcal{X}$ as in IV. of Corollary 10.

But that is a consequence of the following independent observation:

Observation 11 *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) of II. of Theorem 9 is satisfied. Let Bel, Bel' be defined in terms of $\mathcal{X}, \mathcal{X}'$, respectively, as stated in II. of Theorem 9. 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 $Bel'(Z|Y)$. ■

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 9 is satisfied must lead to the maximal class *Bel* of pairs of propositions in $\mathfrak{A}$, if *Bel* is given as in in II. of Theorem 9. But that is exactly what we did in IV. of Corollary 10. Note that unlike the case of absolute belief, where where we were only interested in the least P-stable^r proposition B_W, the additional Least Certain Set Restriction on P is even *entailed* by our postulates on subjective probability and belief. So when we finally turn now IV. of Corollary 10 into an explicit definition of belief on the basis of P, but this time of *conditional belief in general*, then doing so "just" for probability measures for which there exist least propositions of probability 1 is not an actual constraint (given our postulates are plausible). After all, only such probability measures can be combined with any class *Bel* at all, such that all of our postulates are satisfied jointly by them.

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

Definition 12 *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 $\mathcal{X}$ 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 $Y, Z \in \mathfrak{A}$ and $\frac{1}{2} \leq r < 1$:

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 at the end of the previous section: the definition is a true, and even conceptually true, sentence, if taken as a descriptive statement and if given our postulates.

In analogy with the case of absolute beliefs, we could now define notions of *prima facie* conditional belief and conditional belief *a priori* again, and again we would end up with three notions of belief of increasing strength: *prima facie* conditional belief, conditional belief (to a cautiousness degree of r), and conditional belief *a priori*. Of course, conditional belief in the sense of Definition 12 is the concept that we propose as an explication of our pre-theoretic notion of conditional belief simpliciter.

[APPLICATIONS AND EXTENSIONS LEFT OUT.]