

The Lockean Thesis and the Logic of Belief

James Hawthorne
University of Oklahoma
hawthorne@ou.edu

Draft -- submitted to Franz Huber and Christoph Schmidt-Petri (eds.), *Degrees of Belief*, Springer.

1. Introduction

In a penetrating investigation of the relationship between *belief* and quantitative *degrees of confidence* (or *degrees of belief*) Richard Foley (1992) suggests the following thesis:

... it is epistemically rational for us to believe a proposition just in case it is epistemically rational for us to have a sufficiently high degree of confidence in it, sufficiently high to make our attitude towards it one of belief.

Foley goes on to suggest that *rational belief* may be just *rational degree of confidence* above some threshold level that the agent deems sufficient for belief. He finds hints of this view in Locke's discussion of probability and degrees of assent, so he calls it the *Lockean Thesis*.¹

The Lockean Thesis has important implications for the logic of belief. Most prominently, it implies that even a logically ideal agent whose degrees of confidence satisfy the axioms of probability theory may quite rationally believe each of a large body of propositions that are jointly inconsistent. For example, an agent may legitimately believe that on each given occasion her well-maintained car will start, but nevertheless believe that she will eventually encounter a

¹ Foley cites Locke's *An Essay Concerning Human Understanding*, (1975), Book IV, Chapters xv and xvi. Foley discusses the thesis further in his (1993).

dead battery.² Some epistemologists have strongly resisted the idea that such beliefs can be jointly rationally coherent. They maintain that *rationality*, as a normative standard, demands consistency among all of the agent's beliefs – that upon finding such an inconsistency, the rational agent must modify her beliefs. The advocates of consistent belief allow that a real agent may legitimately fall short of such an ideal, but the legitimacy of this short-fall is only sanctioned, they maintain, by the mitigating circumstance of her limited cognitive abilities – in particular by an her lack of the kind of logical omniscience one would need in order to compute enough of the logical consequences of believed propositions to uncover the inconsistencies. But if the Lockean Thesis is right, the logic of belief itself permits a certain degree of inconsistency across the range of an agent's beliefs, even for idealized, logically omniscient agents.³

² This is an instance of the preface paradox. In (Hawthorne and Bovens, 1999) we explored implications of the Lockean Thesis for the preface and lottery paradoxes in some detail. I'll more fully articulate the logic underlying some of the main ideas of the earlier paper, but my treatment here will be self-contained.

³ Closely related is the issue of whether *knowledge* should be subject to logical closure – i.e. whether a rational agent is committed to knowing those propositions he recognizes to be logically entailed by the other propositions he claims to know. (See John Hawthorne's (2004) for an insightful treatment of this matter.) This issue is, however, somewhat distinct from the issue of whether an agent may legitimately maintain inconsistent collections of beliefs. For, knowledge requires more than rational belief – e.g. it requires truth. So one might well maintain that everything an agent claims to *know* should be jointly consistent (if not, then closure must be rejected!), and yet hold that an agent may *legitimately believe* each of some jointly inconsistent

So we are faced with two competing paradigms concerning the nature of rational belief. I don't intend to directly engage this controversy here.⁴ My purpose, rather, is to spell out a qualitative logic of belief that I think provides a more compelling model of coherent belief than the quantitative logic based on probabilistic degrees of confidence. I'll show that this qualitative logic fits the Lockean Thesis extremely well. More specifically, this logic will draw on two qualitative doxastic primitives: the relation of *comparative confidence* (i.e. the agent is *at least as confident that A as that B*) and a predicate for *belief* (i.e. the agent *believes that A*). It turns out that this qualitative model of belief and confidence shares many of the benefits associated with the probabilistic model of degrees of confidence. For, given any such comparative confidence relation and associated belief predicate, there will be a probability function and associated threshold that models them in such a way that belief satisfies the Lockean Thesis.

2. The Ideal Agents and the Qualitative Lockean Thesis

The Lockean Thesis is clearly not intended as a description of how real human agents form beliefs. For one thing, real agents don't often assign numerical degrees of confidence to propositions. And even when they do, their probabilistic confidence levels may fail to consistently link up with belief in the way the Lockean Thesis recommends. Indeed, real agents

collection of propositions that he doesn't claim to know.

⁴ I recommend David Christensen's (2004) excellent treatment of these issues. Whereas Christensen draws on the logic of numerical probabilities in developing his view, I'll show how to get much the same logic of belief from a more natural (but related) logical base. So the present paper might be read as offering a friendly amendment to Christensen's account.

are naturally pretty bad at probabilistic reasoning – they often fail miserably at even simple deductive reasoning. So clearly the Lockean Thesis is intended as an idealized model of belief, a kind of normative model, somewhat on a par with the (competing) normative model according to which an agent is supposed to maintain logically consistent beliefs. But even as a normative model, the Lockean Thesis may seem rather problematic, because the model of probabilistic coherence it depends on seems like quite a stretch for real agents to even approximate. For, although we seldom measure our doxastic attitudes in probabilistic degrees, the Lockean Thesis seems to insist that rationality requires us to attempt to do so – to assign propositions weights consistent with the axioms of probability theory. Such a norm seems much too demanding as a guide to rationality?

To see the point more clearly, think about the alternative *logical consistency norm*. It's proponents describe an *ideally rational agent* as maintaining a logically consistent bundle of beliefs. Here the *ideal agent* is a component of the normative model that real agents are supposed to attempt to emulate, to the best of their abilities, to the extent that it is practical to do so. They are supposed to follow the normative ideal by being on guard against inconsistencies that may arise among their beliefs, revising beliefs as needed to better approximate the ideal. If instead we take probabilistic coherence as a normative model, how is the analogous account supposed to go? Perhaps something like this: Real agents should try to emulate the *ideal agent* of the model (to the best of their abilities) by attempting to assign probabilistically coherent numerical weights to propositions; they should then *believe* just those propositions that fall above some numerical threshold for belief appropriate to the context, and should revise probabilistic weights and beliefs as needed to better approximate the ideal.

The problem is that this kind of account of how probabilistic coherence should function as a

normative guide seems pretty far-fetched as a guide for real human agents. It would have them try to emulate the normative standard by actually constructing numerical probability measures of their belief strengths as a matter of course. Real agents seldom do anything like this. Perhaps there are good reasons why they should try.⁵ But a more natural model of confidence and belief might carry more authority as a normative ideal.

As an alternative to the quantitative model of coherent belief, I will spell out a more compelling *qualitative logic of belief and confidence*. I'll then show that probabilistic measures of confidence lie just below the surface of this qualitative logic. Thus, we may accrue many of the benefits of the probabilistic model without the constant commitment to the arduous task of assigning numerical weights to propositions.

A very natural qualitative version of the Lockean Thesis will better fit the qualitative doxastic logic I'll be investigating. Here it is:

⁵ There are, of course, arguments to the contrary. Dutch book arguments attempt to show that if an agent's levels of confidence cannot be numerically modeled in accord with the usual probabilistic axioms, she will be open to accepting bets that are sure to result in net losses. And the friends of rational choice theory argue that an agent's preferences can be rationally coherent only if his levels of confidence may be represented by a probability function. The import of such arguments is somewhat controversial (though I find the de pragmatized versions in Joyce (1999) and Christensen (2004) pretty compelling). In any case, the present paper will offer a separate (but somewhat related) de pragmatized way to the existence of an underlying probabilistic representation. So let's put the usual arguments for probabilistic coherence aside.

Qualitative Lockean Thesis: An agent is epistemically warranted in believing a proposition *just in case* she is epistemically warranted in having a sufficiently high *grade of confidence* in it, sufficiently high to make her attitude towards it one of belief.

This qualitative version of the thesis draws on the natural fact that we believe some claims more strongly than others – that our confidence in claims comes in relative strengths or *grades*, even when it is not measured in numerical *degrees*. For instance, an agent may (warrantedly) *believe* that F without being *certain* that F. Certainty is a higher grade of confidence than mere belief. Also, an agent may *believe* both F and G, but be *more confident* that F than that G. Belief and confidence may be graded in this way without being measured on a numerical scale.

I will describe a logic for ‘ α believes that B’ and for ‘ α is at least as confident that B as that C’ (i.e. ‘ α believes B at least as strongly as C’) that ties the belief predicate and the confidence relation together by way of this Qualitative Lockean Thesis. In particular, I will show how two specific rules of this logic tie *belief* to *confidence* in a way that is intimately connected to the *preface* and *lottery* paradoxes. It will turn out that any *confidence relation* and associated *belief predicate* that satisfies the rules of this logic can be modeled by a probability function together with a numerical threshold level for belief – where the threshold level depends quite explicitly on how the qualitative logic treats cases that have the logical structure of *preface* and *lottery* paradoxes.⁶ In effect what I’ll show is that probability supplies a kind of formal representation that models the qualitative logic of belief and confidence. The *qualitative semantic rules* for the logic of belief and confidence turn out to be sound and complete with respect to this probabilistic

⁶ Henry Kyburg first discussed the lottery paradox in his (Kyburg, 1961). Also see Kyburg’s (1970). The preface paradox originates with David Makinson (1965).

model theory.

How good might this qualitative account of belief and confidence be at replacing the onerous requirements of probabilistic coherence? Let's step through the account of what the normative ideal recommends for real agents one more time, applying it to the qualitative model. The idea goes like this: Real agents should try to emulate the *ideal agents* of the model (to the best of their abilities) by being on guard against incoherent comparative confidence rankings (e.g. against being simultaneously more confident that A than that B *and* more confident that B than that A), and against related incoherent beliefs (e.g. against believing both A and not-A); and they should revise their beliefs and comparative confidence rankings as needed to better approximate this ideal. The plausibility of this kind of account will, of course, largely depend on how reasonable the proposed coherence constraints on *confidence* and *belief* turn out to be. To the extent that this account succeeds, it inherits whatever benefits derive from the usual probabilistic model of doxastic coherence, while avoiding much of the baggage that attends the numerical precision of the probabilistic model. Thus, it should provide a much more compelling normative model of qualitative *confidence* and *belief*.

3. The Logic of Comparative Confidence

Let's formally represent an agent α 's *comparative confidence relation* among propositions (at a given moment) by a binary relation ' \geq_α ' between statements.⁷ Intuitively ' $A \geq_\alpha B$ ' may be read

⁷ The syntax of the logic I'll be describing employs sentences which, for a given assignment of meanings, become statements that express propositions, as is usual in a formal logic. So from this point on I'll speak in terms of sentences and statements. On this usage, to say that an agent

in any one of several ways: ‘ α is at least as confident that A as that B’, or ‘ α believes A at least as strongly as she believes B’, or ‘A is at least as plausible for α as is B’. For the sake of definiteness I will generally employ the first of these readings, but you may choose your favorite comparative doxastic notion of this sort. The following formal treatment of \geq_α should fit any such reading equally well. Furthermore, I invite you to read ‘ $A \geq_\alpha B$ ’ as saying “ α is *warranted* in being at least as confident that A as that B” (or “ α is *justified* in believing A as strongly as B), if you take that to be the better way of understanding of the important doxastic notion whose logic needs to be articulated.

One comment about the qualifying term ‘warranted’ (or ‘justified’) in the context of the discussion of *confidence* and *belief*. I am about to specify logical rules for ‘ \geq_α ’ (and later for ‘believes that’) – e.g., one such rule will specify that ‘*is at least as confident as*’ should be transitive: if $A \geq_\alpha B$ and $B \geq_\alpha C$, then $A \geq_\alpha C$. Read without the qualifying term ‘warranted’, this rule says, “if α is at least as confident that A as that B, and α is at least as confident that B as that C, then α is at least as confident that A as that C.” Read this way, α is clearly supposed to be a logically ideal agent. In that case you may, if you wish, presume that the ideal agent is warranted in all of her comparative confidence assessments and beliefs. Then, to the extent that the logic is compelling, real agents are supposed to attempt to live up to this logical ideal as best they can. Alternatively, if you want to think of α as a realistic agent, the qualifier ‘warranted’ may be employed throughout, and takes on the extra duty of indicating a logical norm for real agents. For example, the transitivity rule is then read this way: “if α is warranted in being at least as confident that A as that B, and α is warranted in being at least as confident that B as that C, then

believes statement S just means that she *believes the proposition expressed by statement S*.

α is warranted in being at least as confident that A as that C.” In any case, for simplicity I’ll usually suppress ‘warranted’ in the following discussion. But feel free to read it in throughout, if you find the norms on these doxastic notions to be more plausible when expressed that way.

In this section I will specify the logic of the *confidence* relation. Closely associated with it is the *certainty* predicate ‘ $\text{Cert}_\alpha[A]$ ’ (read “ α is certain that A”). Certainty is easily definable from comparative confidence. To be certain that A is to be at least as confident that A as that a simple tautology of form ‘ $(A \vee \neg A)$ ’ holds – i.e., by definition, ‘ $\text{Cert}_\alpha[A]$ ’ will just mean ‘ $A \geq_\alpha (A \vee \neg A)$ ’. For now we stick strictly to *confidence* and *certainty*. We will pick up *belief* in a later section.

3.1 The Rudimentary Confidence Relations

To see that we can spell out the logic of belief and confidence in a completely rigorous way, let’s define confidence relations as semantic relations between object language sentences of a language L for predicate logic with identity. A weaker language would do – e.g. a language for propositional logic. But then you might wonder whether for some reason the following approach wouldn’t work for a stronger language. So, for the sake of definiteness, I’ll directly employ the stronger language. Indeed, the logic of belief and confidence presented here should work just fine for any object language together with it’s associated logic – e.g. for your favorite modal logic. Furthermore, I appeal to a formal language only because it helps provide a well understood formal model of the main idea. The object language could just as well be a natural language, provided that the notion of deductive logical entailment is well defined there.

So, a confidence relation \geq_α is a semantic relation between sentences of a language. The following semantic rules (or axioms) seem to fit the intuitive reading of this notion quite well.

Definition 1: Rudimentary Confidence Relations: Given a language L for predicate logic with identity, the *rudimentary confidence relations* on L are just those relations \geq_α that satisfy

the following rules (where ' $\models A$ ' say that A is a logical truth of L):

First, define ' $\text{Cert}_\alpha[A]$ ' (read ' α is (warranted in being) certain that A') as $A \geq_\alpha (A \vee \neg A)$;

For all sentences A, B, C, D, of L,

1. it's never the case that $\neg(A \vee \neg A) \geq_\alpha (A \vee \neg A)$ (*nontriviality*);
2. $B \geq_\alpha \neg(A \vee \neg A)$ (*minimality*);
3. $A \geq_\alpha A$ (*reflexivity*);
4. if $A \geq_\alpha B$ and $B \geq_\alpha C$, then $A \geq_\alpha C$ (*transitivity*);
- 5.1 if $\text{Cert}_\alpha[C \equiv D]$ and $A \geq_\alpha C$, then $A \geq_\alpha D$ (*right equivalence*);
- 5.2 if $\text{Cert}_\alpha[C \equiv D]$ and $C \geq_\alpha B$, then $D \geq_\alpha B$ (*left equivalence*);
- 6.1 if for some E, $\text{Cert}_\alpha[\neg(A \cdot E)]$, $\text{Cert}_\alpha[\neg(B \cdot E)]$, and
 $(A \vee E) \geq_\alpha (B \vee E)$, then $A \geq_\alpha B$ (*subtractivity*);
- 6.2 if $A \geq_\alpha B$, then for all G such that
 $\text{Cert}_\alpha[\neg(A \cdot G)]$ and $\text{Cert}_\alpha[\neg(B \cdot G)]$, $(A \vee G) \geq_\alpha (B \vee G)$ (*additivity*);
7. if $\models A$, then $\text{Cert}_\alpha[A]$ (*tautological certainty*).

Also, define ' $A \approx_\alpha B$ ' (read " α is *equally confident* in A and B") as ' $A \geq_\alpha B$ and $B \geq_\alpha A$ ';

define ' $A >_\alpha B$ ' (read " α is *more confident* in A than in B"), as ' $A \geq_\alpha B$ and not $B \geq_\alpha A$ ';

and define $A \sim_\alpha B$ (read " α 's comparative confidence that A as compared to B is indeterminate"), as ' $\text{not } A \geq_\alpha B \text{ and not } B \geq_\alpha A$ '.

These rules are a weakened version of the axioms for *qualitative probability* (sometimes called *comparative probability*).⁸ From these axioms together with some definitions one can prove a

⁸ For a standard treatment of the *qualitative probability* relations see (Savage, 1972). The axioms

number of intuitively plausible things about comparative confidence. For example, the following relationships follow immediately from the definitions together with *transitivity* and *reflexivity* (but draw on none of the other rules): (i) for any two statements, either $A >_{\alpha} B$ or $B >_{\alpha} A$ or $A \approx_{\alpha} B$ or $A \sim_{\alpha} B$; (ii) $A \geq_{\alpha} B$ just in case either $A >_{\alpha} B$ or $A \approx_{\alpha} B$; (iii) ' $>_{\alpha}$ ' is transitive and asymmetric, and ' \approx_{α} ' is an equivalence relation (i.e. transitive, symmetric, and reflexive); (iv) whenever two statements are considered equally plausible by α (i.e. whenever $A \approx_{\alpha} B$) they share precisely the same confidence relations (\geq_{α} , $>_{\alpha}$, \approx_{α} , and \sim_{α}) to all other statements. The following claims are also easily derived⁹: (v) if $\text{Cert}_{\alpha}[A]$, then, for all B , $A \geq_{\alpha} B \geq_{\alpha} \neg A$ and $(B \cdot A) \approx_{\alpha} B \approx_{\alpha} (B \vee \neg A)$; (vi) if $\text{Cert}_{\alpha}[B \supset A]$, then $A \geq_{\alpha} B$; (vii) if $A \geq_{\alpha} B$ then $\neg B \geq_{\alpha} \neg A$.

Let's look briefly at each rule for the *rudimentary confidence relations* to see how plausible it is as a constraint on comparative confidence – i.e., to see how well it fits our intuitions about comparative confidence. Rules 1 and 2 are obvious constraints on the notion of comparative

given here are weaker in that they only require confidence relations to be partial preorders (i.e. reflexive and transitive), whereas such relations are usually specified to be total preorders (i.e. complete and transitive). Also, the present axioms have been adapted to apply to sentences of a language, whereas Savage's version applies to *sets of states* or *sets of possible worlds*. Although that approach is formally somewhat simpler, it tends to hide important philosophical issues, such as the issue of the *logical omniscience* of the agent. Notice that our approach only draws on the notion of *logical truth* in rule 7. The other rules are quite independent of this notion. This will permit us to more easily contemplate how the rules may apply to logically more realistic agents.

⁹ The derivations of these draw on rule 7 only to get certainty for some very simple tautologies – e.g. $\models A \equiv ((A \cdot B) \vee (A \cdot \neg B))$, and $\models \neg((A \cdot B) \cdot (A \cdot \neg B))$.

confidence. Rule 2, *minimality*, just says that every statement B should garner at least as much confidence as a simple contradiction of form $\neg(A \vee \neg A)$. The agent should have no confidence at all in such simple contradictions – they lay at the bottom of the confidence ordering. Given the definition of ‘ $>_{\alpha}$ ’, rule (1), *nontriviality*, taken together with *minimality* is equivalent to ‘ $(A \vee \neg A) >_{\alpha} \neg(A \vee \neg A)$ ’, which says that the agent is (warranted in being) *more confident* in any simple tautology of form $(A \vee \neg A)$ than in the simple contradiction gotten by taking its negation. If this rule failed, the agent’s “confidence ordering” would indeed be trivial. Indeed, given the remaining rules, the agent would be equally confident in every statement.¹⁰

Rule 3, *reflexivity*, merely requires that the agent find each statement to be at least as plausible as itself. This should be uncontroversial.

Rule 4, *transitivity*, is more interesting, but should not really be controversial. It says that whenever α is (warranted in being) at least as confident that A as that B, and is (warranted in being) at least as confident that B as that C, then α is (warranted in being) at least as confident that A as that C. This rule seems unassailable as a principle of comparative confidence. Ideal agents follow it, and it seems perfectly reasonable to expect real agents to try to conform to it.

All of the rules up to this point should be uncontroversial. Indeed, of all the rules for the *rudimentary confidence relations*, I only expect there to be any significant concern over *tautological certainty* (rule 7), which seems to require a kind of logical omniscience. We’ll get to that. But none of the rules described thus far require anything we wouldn’t naturally expect of real rational agents.

¹⁰ Because, if $\neg(A \vee \neg A) \geq_{\alpha} (A \vee \neg A)$, from certainty in some very simple tautologies it follows that for each B and C, $B \geq_{\alpha} \neg(A \vee \neg A) \geq_{\alpha} (A \vee \neg A) \geq_{\alpha} C$; thus $B \geq_{\alpha} C$, and similarly $C \geq_{\alpha} B$.

The usual axioms for qualitative probability are stronger than the rules presented here. In place of *reflexivity*, the usual axiom include an axiom of *complete comparison*, which says that for any pair of sentences A and B, the agent is either at least as confident that A as that B, or she is at least as confident that B as that A:

3*. $A \geq_{\alpha} B$ or $B \geq_{\alpha} A$ (*completeness*, a.k.a. *totality* or *comparability*).

Completeness says that the agent can make a determinate confidence comparison between any two statements. This rule is rather controversial, so I've not made it a necessary constraint on the *rudimentary confidence relations*. However, I will argue in a bit that the rudimentary confidence relations should always be *extendable* to confidence relations that satisfy *completeness*. More about this later.

Notice that *completeness* would supersede *reflexivity*, since *completeness* implies ' $A \geq_{\alpha} A$ or $A \geq_{\alpha} A$ ' – i.e. $A \geq_{\alpha} A$. When any binary relation is both *reflexive* and *transitive* it is called a *preorder* (alternatively, a *quasi-order*). Adding *completeness* to *transitivity* yields a *total preorder*.¹¹ Where *completeness* is present, the relationship of *confidence ambiguity*, ' $A \sim_{\alpha} B$ ',

¹¹ Terminology about order relations can be confusing because usage is not uniform. By a '(weak) preorder' I mean a reflexive and transitive relation. The term 'weak partial order' is often used this way too, but is also often used to mean a reflexive, transitive, and antisymmetric relation. (Antisymmetry says, 'if $A \geq B$ and $B \geq A$, then $A = B$ ', where '=' is the *identity relation*, not just the equivalence we've denoted by ' \approx '.) Applied to statements, antisymmetry would be too strong. It would mean that whenever α is equally confidence that A as that B (i.e. whenever $A \approx_{\alpha} B$), A and B must be the same statement, or at least be logically equivalent statements. The term '(weak) total preorder' means a preorder that also satisfies *completeness*. The term 'weak

will be vacuous; there is never ambiguity in confidence comparisons between two statements. I'll discuss *completeness* more a bit later. For now, suffice it to say that the *rudimentary confidence relations* are not required to satisfy it.

Rules 5, *substitutivity of equivalences* (*left* and *right*), make good sense. The two parts together say that whenever an agent is *certain* that statements X and Y are materially equivalent (i.e. *certain* that they agree in truth value), then all of her comparative confidence judgments involving Y should agree with those involving X.

The two rules 6 taken together say that incompatible disjuncts (added or subtracted) should make no difference to the comparative confidence in statements. To see the idea behind rule 6.1, *subtractivity*, consider a case where the agent is certain that some statement E is incompatible with each of two statements A and B. (There will always be such an E – e.g. the simple contradiction $(C \cdot \neg C)$.) If the agent is at least as confident that $(A \vee E)$ as that $(B \vee E)$, then intuitively, she should be at least as confident that A as that B. The removal of the “disjunctively tacked on” incompatible claim E should have no effect on the agent’s comparative confidence with respect to A and B.

Furthermore, whenever the agent considers a statement G to be incompatible with the truth of A and with the truth of B, her relative confidence in the disjunctions, $(A \vee G)$ as compared to $(B \vee G)$, should agree with her relative confidence in A as compared to B. Only the agent’s confidence in A as compared to B should matter. Disjunctively tacking on the incompatible claim G should have no influence on her assessment. This is just what rule 6.2, *additivity*, says. Both

order’ is often used this way too, but is also often used to mean that the relation is antisymmetric as well. (The *weak/strict* distinction picks out the difference between \geq and $>$.)

subtractivity and *additivity* are substantive rules. But both are completely reasonable constraints on an agent's comparative confidence assessments.

I'll soon suggest two additional rules that I think a more complete account of the notion of comparative confidence should satisfy. But all of the rules for the rudimentary relations appear to be sound, reasonable constraints on comparative confidence. I think that rule 7, *tautological certainty*, is the only constraint stated so far that should raise any controversy. It says that if a sentence is logically true, the ideal agent will be certain that it's true – i.e. as confident in it as in a simple tautology of form ' $(A \vee \neg A)$ '. It thereby recommends that when a real agent attempts to assess her comparative confidence in some given pair of statements, she should (to the extent that it's practical for her to do so) seek to discover whether they are logically true, and should become certain of those she discovers to be so. The ideal agent always succeeds, and the real agent is supposed to attempt it, to the extent that it's practical to do so. Put this way the rule sounds pretty innocuous. Rules of this kind are common in epistemic and doxastic logics. How problematic is it as component of a normative guide?

First let's be clear that failure to be sufficiently logically talented does not, on this account, warrant calling the real agent *irrational* – it only implies that she is *less than ideally rational*. But some will argue that this ideal is too far beyond our real abilities to count as an appropriate doxastic norm. Let's pause to think a bit more about this, and about the kind of norm we are trying to explicate.

Notice that we might easily replace rule 7 by a weaker version. We might, for example, characterize a broader class of confidence relations by reading ' $\models A$ ' as ' A is a logical truth of the sentential logic of L '. In that case the agent need not be certain of even the simplest predicate logic tautologies. However, even the computation of propositional logic tautologies is in general

NP hard, and so in many cases outstrips the practical abilities of real agents. Perhaps a better alternative would be to only require certainty for some easily computable class of logical truths – e.g., read ‘ $\models A$ ’ as “A is a logical truth computable via a truth tree consisting of no more than 16 branches”; or perhaps read it as “the number of computation steps needed to determine the logical truth of A is bounded by a (specified) polynomial of the number of terms in A”. Some such weaker rule, which together with the other rules characterizes a broader class of rudimentary confidence relations, might well provide a more realistic normative constrain on the comparative confidence assessments of real agents.

The full development of a less demanding doxastic logic that better fits the abilities of real agents would certainly be welcomed. But even if/when we have such a logic available, the more demanding ideal we are exploring here will continue to have an important normative role to play. To see the point imagine that such a *real-agent-friendly* logic of confidence and belief has been worked out, and consider some collection Γ of confidence relationships (between statements) or beliefs that this logic endorses as *rationally coherent for real agents*.¹² Wouldn’t we still want to know whether the *realistic coherence* of Γ arises *only* because of the limited logical abilities of the agents we are modeling? Wouldn’t we want to know whether a somewhat more demanding *real-agent logic*, suited to somewhat more logically adept agents, would pronounce a different verdict on the coherence of Γ , perhaps assessing this collection as *rationally incoherently* for the

¹² Suppose, for example, that this logic endorses as *rationally coherent*, beliefs like those that take the form of the *preface-paradox* – where an agent believes each of a number of claims, S_1 through S_n (e.g. where S_i says that page i of her book is free of error) and she also believes $\neg(S_1 \cdot \dots \cdot S_n)$ (e.g. that not all pages of her n page book are error free).

more adept? That is, wouldn't we still want to know whether the assessment of Γ as *coherent* results only from the limited deductive abilities of real agents, or whether such confidence relations and beliefs would continue to count as *jointly coherent*, regardless of limitations? Only a normative ideal that doesn't model deductive-logical limitations can answer *these* questions.¹³

There will always be some cognitive differences among real people. Some will be more logically adept than others, and the more adept reasoners should count as *better reasoners* for it. And it seems unlikely that there is a plausible way to draw a firm line to indicate where "good enough reasoning" ends. That is, it seems doubtful that we can develop a *logic of real reasoning* that would warrant the following kind of claim: "Reasoning that reaches the logical depth articulated by *this logic* is as good as we can plausibly want a real reasoner to be, and any actual agent who recognizes more logical truths than that will just *not count as any better* at maintaining *belief coherence*." The point is that no matter how successful a *real-agent logic* is at describing plausible norms, if the norm falls short of *tautological certainty*, there may always be some agents who exceed the norm to some extent, and they should count as *better* for it. Thus, although the ideal of *tautological certainty* may be an unattainable standard for a real agent, it nevertheless provides a kind of least upper bound on classes of rationally coherent comparative confidence relations.

It turns out that any relation that satisfies rules 1-7 behaves a lot like comparisons of probabilistic degrees of confidence. That is, each of these relations is *probabilistically* sound in the following sense:

¹³ Indeed, later we will see that the logic we are investigating here, ideal as it is, affirms the rational coherence of *preface-like* and *lottery-like* beliefs, even for logically ideal agents.

Given any probability function P_γ (that satisfies the usual probability axioms),¹⁴ the relation \geq_γ defined as “ $A \geq_\gamma B$ just when $P_\gamma[A] \geq P_\gamma[B]$ ” satisfies rules 1-7.

Thus, if an agent were to have a *probabilistic confidence function* that provides a numerical measure of her degree of confidence in various statements, this function would automatically give rise to a *rudimentary confidence relation* for her. However, some confidence relations that satisfy 1-7 cannot be represented by any probability function – i.e. rules 1-7 are not *probabilistically complete*. Two additional rules will place enough of a restriction on the rudimentary confidence relations to close this gap.

3.2 The Completed Confidence Relations

Rudimentary confidence relations allow for the possibility that an agent cannot determine a definite confidence comparison between some pairs of statements. When this happens, the confidence relation is *incomplete* – i.e. for some A and B, neither $A >_\alpha B$, nor $A \approx_\alpha B$, nor $B >_\alpha A$. Real agents may well be unable to assess their comparative confidence in some pairs of statements. Nevertheless, there is a perfectly legitimate role for *completeness*¹⁵ to play as a normative guide. I’ll argue that a reasonable additional constraint on comparative confidence is this: an agent’s comparative confidence relation should in principle be *consistently extendable* to a relation that compares all statements. For, if no such consistent extension is even *possible*, then the agent’s *definite confidence relationships* must be *implicitly incoherent*.

¹⁴ Here are the usual axioms for probabilities on sentences of a formal language L. For all R and S: (i) $P[S] \geq 0$; (ii) if $\models S$, then $P[S] = 1$; (iii) if $\models \neg(R \cdot S)$, then $P[R \vee S] = P[R] + P[S]$.

¹⁵ This notion of *completeness* should not be confused with the notion of *probabilistic completeness* for a confidence relation described at the end of the previous subsection.

To see this, suppose that $A \sim_\alpha B$, and suppose that no extension of her definite confidence relations ($>_\alpha$ and \approx_α) to any definite relationship between A and B would yield a confidence relation consistent with rules 1-7. That means that her definite confidence relationships imply *on their own* (from rules 1-7) that $A \sim_\alpha B$ *must hold* – because no definite confidence relationship between A and B is coherently possible, given her other definite confidence relationships. The agent maintains coherence only by refusing to commit to a definite relationship between A and B. Thus, in such a case, the agent’s inability to assess a determinate confidence relationship between A and B is not merely a matter of it “being a hard case”. Rather, her refusal to make an assessment is forced upon her. It is her only way to stave off explicit incoherence among her other determinate comparative confidence assessments. This seems a really poor reason for an agent to maintain indeterminateness. Rather, we should recommend that when a real agent discovers such implicit incoherence, she should revise her confidence relation to eliminate it. Her revised confidence relation might well leave the relationship between A and B indeterminate – but this should no longer be due to the *incoherence of the possibility* of placing a definite confidence relationship between them.

Thus, insofar as the rules for the rudimentary confidence relations seem reasonable as a normative standard, it also makes good sense to add the normative condition that a coherent rudimentary confidence relation should be extendable to a *complete* rudimentary confidence relation, a relation that satisfies rule 3*. (I’ll show how to handle this formally in a moment.) There will often be lots of possible ways to extend a given vague or indeterminate confidence relation, many different ways to fill in the gaps. I am *not* claiming that the agent should be willing to embrace some particular such extension of her confidence relation, but only that some such extension should be consistent with the confidence orderings she does have.

Let's now restrict the class of rudimentary confidence relations to those that satisfy an additional two-part rule that draws on *completeability* together with one additional condition. The most efficient way to introduce this additional rule is to first state it as part of a definition.

Definition 2: Properly Extendable Rudimentary Confidence Relations: Let us say that a *rudimentary confidence relation* \geq_α on language L is *properly extendable* just in case there is a *rudimentary confidence relation* \geq_β on some language L^+ an extension of L that agrees with the determinate part of \geq_α (i.e., whenever $A \approx_\alpha B$, $A \approx_\beta B$; and whenever $A >_\alpha B$, $A >_\beta B$) on the shared language L, and also satisfies the following rule for all sentences of L^+ :

(X) (i) (*completeness*): either $A \geq_\beta B$ or $B \geq_\beta A$; and

(ii) (*separating equiplausible partitions*): If $A >_\beta B$, then, for some integer n, there are n sentences S_1, \dots, S_n that β takes to be mutually incompatible (i.e., $\text{Cert}_\beta[\neg(S_i \cdot S_j)]$ for $i \neq j$), and jointly exhaustive (i.e., $\text{Cert}_\beta[S_1 \vee \dots \vee S_n]$) and in all of which β is equally confident (i.e. $S_i \approx_\beta S_j$ for each i, j), such that for each of them, S_k , $A >_\beta (B \vee S_k)$.

(Any set of sentences $\{S_1, \dots, S_n\}$ such that $\text{Cert}_\beta[\neg(S_i \cdot S_j)]$ and $\text{Cert}_\beta[S_1 \vee \dots \vee S_n]$ is called an n-ary *equiplausible partition* for β .)

The 'X' here stands for 'eXtendable'. The idea is that when a confidence relation is rationally coherent, there should in principle be some *complete* extension that includes partitions of the "space of possibilities", where the parts of the partition S_k are, in β 's estimation, equally plausible, but where there are so many alternatives that β can have very little confidence in any one of them. Indeed, for any statement A in which β has more confidence than another statement B, there is some large enough such partition that her confidence in each partition statement must be so trifling that she remains more confident in A than she is in the disjunction of any one of

them with B. (This implies that the partition is fine-grained enough that at least one disjunction $B \vee S_k$ must *separate* A from B in that $A >_{\beta} (B \vee S_k) >_{\beta} B$.)

More concretely, consider some particular pair of statements A and B, where β is more confident that A than that B, and where \geq_{β} is a complete rudimentary confidence relation. Suppose there is a fair lottery consisting of a very large number of tickets, n, and let ‘ S_i ’ say that ticket i will win. Further suppose that with regard to this lottery, β is certain of its fairness (i.e. $S_i \approx_{\beta} S_j$ for every pair of tickets i and j), she is certain that no two tickets can win, and she is certain that at least one will win. Then rule X will be satisfied for the statements A and B provided that the lottery consists of so many tickets (i.e. n is so large) that β remains more confident in A than in the disjunction of B any one claim S_i asserting that a specific ticket will win. To satisfy rule X we need only suppose that for each pair of sentences A and B such that $A >_{\beta} B$, there is such a lottery, or that there is some similar partition into extremely implausible possibilities (e.g. let each S_i describe one of the $n = 2^m$ possible sequences of *heads* and *tails* in an extremely long sequence of tosses of a fair coin).

That explains rule X. But what if there are no such lotteries, nor any similar large equiplausible partitions for an agent to draw on in order to satisfy rule X? I have yet to explain the notion of being *properly extendable*, and that notion is designed to deal with this problem. According to the definition, the agent β who possesses a “properly extended” confidence relation has a rich enough collection of equiplausible partitions at hand to satisfy rule X for all sentences A and B. But in general an agent α may not be so fortunate. For example, α may not think that there are any such lotteries, or any such events that can play the role of the needed partitions of her “confidence space”. Nevertheless, α ’s comparative confidence relation will have much the same structure as β ’s, provided that α ’s confidence relation could be gotten by starting with β ’s,

and then throwing away all of those partitions that aren't available to α (e.g. because the relevant statements about them are not expressed in α 's language). In that case, although α herself doesn't satisfy rule X, her confidence relation is *properly extendable* to a relation that does.¹⁶ Indeed, when \succeq_α is *properly extendable*, there will usually be many possible ways (many possible β s) that extend α 's confidence relation so as to satisfy rule X.

¹⁶ To put it another way, α may think that there are no fair lotteries (or similar chance events) anywhere on earth. Thus, rule X does not apply to her directly. But suppose that α 's language could in principle be extended so that it contains additional statements that describe some new *possible chance events* (they needn't be real or actual) not previously contemplated by α , and not previously expressible by α 's language. (Perhaps in order to describe these events α would have to be in some new referential relationship she is not presently in. E.g. suppose there is some newly discovered, just named star, Zeta-prime, and suppose someone suggests that a culture on one of its planets runs lotteries of the appropriate kind, the "Zeta-prime lotteries".) Now, for α 's confidence relation to be *properly extendable*, it only need be *logically possible* that some (perhaps extremely foolish) agent β , who agrees with α as far as α 's language goes, satisfies rule X by employing the newly expressible statements. Notice that we do not require α herself to be willing to extend her own confidence relation so as to satisfy rule X. E.g., when α 's language is extended to describe these new (possible) lotteries, α herself might extend her own confidence relation to express *certainty* that the suggested lotteries don't really exist (or she may think they exist, but take them to be biased). How α would extend her own confidence relation is not in any way at issue. All that matters for our purposes is that her confidence relation *could in principle* coherently (with rules 1-7) be extended to satisfy rule X for some logically possible agent β .

Now we are ready to supplement rules 1-7 with this additional rule. Here is how to do that:

Definition 3: Confidence Relations: Given a language L for predicate logic, the (fully refined) *confidence relations* \geq_α on L are just the *rudimentary confidence relations* (those that satisfy rules 1-7) that are also *properly extendable*.

To recap, rule X is satisfied by a relation \geq_α provided it can be extended to a *complete* relation \geq_β on a language that describes, for example, enough fair single-winner lotteries that whenever $A >_\beta B$, there is some lottery with so many tickets that disjoining with B any claim S_i that says “ticket i will win” leaves $A >_\beta (B \vee S_i)$.

Every probabilistic *degree of confidence function* behaves like a (fully refined) *confidence relation* – i.e. the rules for confidence relations are probabilistically sound in the following sense:

Theorem 1: Probabilistic Soundness of the *confidence relations*: Let P_α be any probability function (that satisfies the usual axioms). Define a relation \geq_α as follows: $A \geq_\alpha B$ just when $P_\alpha[A] \geq P_\alpha[B]$. Then \geq_α satisfies rules 1-7 and is *properly extendable* to a relation that also satisfies rule X.

Conversely, every confidence relation can be modeled or *represented* by a probabilistic *degree of confidence function*:

Theorem 2: Probabilistic Representation of the *confidence relations*: For each relation \geq_α that satisfies rules 1-7 and is *properly extendable*, there is a probability function P_α that models \geq_α as follows:

- (1) if $P_\alpha[A] > P_\alpha[B]$, then $A >_\alpha B$ or $A \sim_\alpha B$;
- (2) if $P_\alpha[A] = P_\alpha[B]$, then $A \approx_\alpha B$ or $A \sim_\alpha B$.

Furthermore, if \geq_α itself satisfies rule X (rather than merely being *properly extendable* to a rule X satisfier), then P_α is unique and $P_\alpha[A] \geq P_\alpha[B]$ if and only if $A \geq_\alpha B$.

Notice that, taken together, (1) and (2) imply the following:

(3) if $A >_{\alpha} B$, then $P_{\alpha}[A] > P_{\alpha}[B]$;

(4) if $A \approx_{\alpha} B$, then $P_{\alpha}[A] = P_{\alpha}[B]$.

And this further implies that

(5) if $A \sim_{\alpha} B$, then for each C and D such that $C >_{\alpha} A >_{\alpha} D$ and $C >_{\alpha} B >_{\alpha} D$,

both $P_{\alpha}[C] > P_{\alpha}[A] > P_{\alpha}[D]$ and $P_{\alpha}[C] > P_{\alpha}[B] > P_{\alpha}[D]$.

That is, whenever $A \sim_{\alpha} B$, the representing probabilities must either be equal, or lie *relatively close together* – i.e. both lie below the smallest representing probability for any statement C in which α is determinately more confident (i.e. such that both $C >_{\alpha} A$ and $C >_{\alpha} B$) and both lie above the largest representing probability for any statement D in which α is determinately less confident (i.e. such that both $A >_{\alpha} D$ and $B >_{\alpha} D$).

Thus, probabilistic *degree of confidence functions* simply provide a way of modeling qualitative confidence relations on a convenient numerical scale. The probabilistic model will not usually be unique. There may be lots of ways to model a given confidence relation probabilistically. However, in the presence of equiplausible partitions, the amount of wiggle room decreases, and disappears altogether for those confidence relations that themselves satisfy the conditions of rule X.

The probabilistic model of a refined confidence relation will not usually be unique. There will usually be lots of ways to model a given confidence relation probabilistically – because there will usually be lots of ways to extend a given confidence relation to a complete-equiplausibly-partitioned relation. So in general each comparative confidence relation is represented by a set of representing probability functions.

A common objection to “probabilism” (the view that belief-strengths should be probabilistically coherent) is that the probabilistic model is overly precise, even as a model of *ideally* rational agents. Proponents of probabilism often respond by suggesting that, to the extent that vagueness in belief strength is reasonable, it may be represented by sets of degree-of-belief functions that cover the reasonable range of numerical imprecision. Critics reply that this move is (at best) highly questionable – it gets the cart before the horse. Probabilism first represents agents as having overly precise belief strengths, and then tries to back off of this defect by taking the agent to actually be a whole chorus of overly precise agents.

“Qualitative probabilism” not only side-steps this apparent difficulty -- it entirely resolves (or dissolves) this issue. In the first place, qualitative probabilism doesn’t suppose that the agent has numerical degrees of belief – it doesn’t even suppose that the agent can determine definite confidence-comparisons between all pairs of statements. Secondly, the Representation Theorem(s) show how confidence relations give rise to sets of degree-of-belief functions that reflect whatever imprecision is already in the ideal agent’s confidence relation. We may *model* the agent with any one of these (overly precise) functions, keeping in mind that the *quantitative* function is only one of a number of equally good numerical representations. So, for qualitative probabilism, the appeal to sets of representing probability functions is a natural consequence of the incompleteness (or indeterminateness) of the ideal agent’s relative confidence relation – rather than a desperate move to add indefiniteness back into a model that was overly precise from the start.

4. The Integration of Confidence and Belief

Now let’s introduce the notion of *belief* and tie it to the *confidence* relation in accord with the

Qualitative Lockean Thesis. I'll represent belief as a semantic predicate, $\text{Bel}_\alpha[S]$, that intuitively says, “ α believes that S ”, or “ α is warranted in believing that S ”.

Clearly whenever α is certain of a claim, she should also believe it. Thus, the following rule:

(8) If $\text{Cert}_\alpha[A]$ then $\text{Bel}_\alpha[A]$ (*certainty-implies-belief*).

Now, the obvious way to tie *belief* to *confidence* in accord with the Lockean Thesis is to introduce the following rule:

(9) If $A \geq_\alpha B$ and $\text{Bel}_\alpha[B]$, then $\text{Bel}_\alpha[A]$ (*basic confidence-belief relation*).

This rule guarantees that there is a confidence relation threshold for belief. For, given any statement that α believes, whenever α is at least as confident in another statement R as she is in that believed statement, she must believe (or be warranted in believing) R as well. And given any statement that α fails to believe, whenever α is at least as confident in it as she is in another statement S , she must fail to believe S as well.¹⁷

Taking into account the probabilistic modelability of the *confidence* relations, guaranteed by Theorem 2, rule (9) also implies that the quantitative version of the Lockean Thesis is satisfied. That is, for each confidence relation \geq_α and associated belief predicate Bel_α (satisfying rules (8) and (9)), there is at least one probability function P_α and at least one threshold level q such that one of the following conditions is satisfied:

¹⁷ However, this does not imply that there must be a “threshold statement”. There may well be an infinite sequence of statements with descending confidence levels for α , $R_1 >_\alpha R_2 >_\alpha \dots >_\alpha R_n > \dots$, all of which α believes. And there could also be another infinite sequence of statements with ascending confidence levels for α , $S_1 <_\alpha S_2 <_\alpha \dots <_\alpha S_n, \dots$, all of which α fails to believe. (I.e., for countable sets of sentences there need be no greatest lower bound or least upper bound *sentence*.)

- (i) for any sentence S , $\text{Bel}_\alpha[S]$ just in case $P_\alpha[S] \geq q$, or
- (ii) for any sentence S , $\text{Bel}_\alpha[S]$ just in case $P_\alpha[S] > q$.¹⁸

Furthermore, if the confidence relation \geq_α itself satisfies rule X, then P_α and q must be unique.

4.1 The *Preface* and the n-Bounded Belief Logics

We are not yet done with articulating the logic of belief. The present rules don't require the belief threshold to be at any specific level. They don't even imply that the probabilistic threshold q that models *belief* for a given *confidence relation* has to be above $1/2$; q may even be 0, and every statement may be believed. So to characterize the logic of belief above some reasonable level of confidence, we'll need additional rules. I'll first describe these rules formally, and then I'll explain them more intuitively in terms of how they capture features of the *preface paradox*.

The following rule seems reasonable:

- (1/2): if $\text{Cert}_\alpha[A \vee B]$, then not $\text{Bel}_\alpha[\neg A]$ or not $\text{Bel}_\alpha[\neg B]$.

That is, if the agent is certain of a disjunction of two statements, then she may believe the negation of *at most one* of the disjuncts.

There are several things worth mentioning about this rule. First, in the case where B just is A , the rule says that if the agent is certain of A , then she cannot believe $\neg A$. Second, taking B to be $\neg A$, the rule implies that the agent cannot believe both a statement and its negation. Furthermore, the (1/2) rule is probabilistically sound: for any probability function P_α and any specific threshold value $q > 1/2$, the corresponding confidence relation \geq_α and belief predicate Bel_α , defined as 'A

¹⁸ The sequence of probabilities associated with the sequence of statements in the previous note, $P[R_1] > P[R_2] > \dots >_\alpha P[R_n] > \dots$, is bounded below (by 0 at least), so has a greatest lower bound, call it q .

$\geq_\alpha B$ iff $P_\alpha[A] \geq P_\alpha[B]$, and $\text{Bel}_\alpha[C]$ iff $P_\alpha[C] \geq q > 1/2$ ', satisfies all of the previous rules, including the (1/2) rule.

Now consider a rule that is somewhat stronger than the (1/2) rule:

(2/3): if $\text{Cert}_\alpha[A \vee B \vee C]$, then not $\text{Bel}_\alpha[\neg A]$ or not $\text{Bel}_\alpha[\neg B]$ or not $\text{Bel}_\alpha[\neg C]$.

According to this rule, if the agent is certain of a disjunction of *three* statements, then she may believe the negations of *at most two* of the disjuncts. But this rule doesn't bar her from believing the negations of each of a larger number of claims for which the disjunction is certain. Notice that the (1/2) rule is a special case of the (2/3) rule – the case where C just is B. So the (2/3) rule implies the (1/2) rule. Also, in the case where C is $\neg(A \vee B)$, $\text{Cert}_\alpha[A \vee B \vee C]$ must hold because $(A \vee B \vee C)$ is a tautology. In that case the (2/3) rule says that the agent must fail to believe one of the claims $\neg A$ or $\neg B$ or $(A \vee B)$ (i.e. $\neg\neg(A \vee B)$). Furthermore, the (2/3) rule is probabilistically sound: for any probability function P_α and any specific threshold value $q > 2/3$, the corresponding confidence relation \geq_α and belief predicate Bel_α , defined as ' $A \geq_\alpha B$ iff $P_\alpha[A] \geq P_\alpha[B]$, and $\text{Bel}_\alpha[C]$ iff $P_\alpha[C] \geq q > 2/3$ ', satisfies all the previous rules together with the (2/3) rule.

More generally, consider the following $((n-1)/n)$ rule for any fixed $n \geq 2$:

$((n-1)/n)$: if $\text{Cert}_\alpha[A_1 \vee \dots \vee A_n]$, then not $\text{Bel}_\alpha[\neg A_1]$ or ... or not $\text{Bel}_\alpha[\neg A_n]$.

According to this rule, if the agent is certain of a disjunction of n statements, then she may believe the negations of *at most* $n-1$ of the disjuncts. But this rule doesn't bar her from believing each of a larger number of statements for which the disjunction is certain. Notice that for any $m < n$, the $((m-1)/m)$ rule is a special case of the $((n-1)/n)$ rule. So the $((n-1)/n)$ rule implies all $((m-1)/m)$ rules for $n > m \geq 2$. Furthermore, the $((n-1)/n)$ rule is probabilistically sound in that, given any probability function P_α and any specific threshold value $q > (n-1)/n$, the corresponding

confidence relation \geq_α and belief predicate Bel_α , defined as ‘ $A \geq_\alpha B$ iff $P_\alpha[A] \geq P_\alpha[B]$, and $\text{Bel}_\alpha[C]$ iff $P_\alpha[C] \geq q > (n-1)/n$ ’, satisfies all of the previous rules together with the $((n-1)/n)$ rule.

Let’s say that any confidence relation \geq_α together with the associated belief predicate Bel_α that satisfy rules 1-9 and the $((n-1)/n)$ rule (for a given value of n) *satisfies an n-bounded belief logic*. Clearly the n -bounded belief logics form a nested hierarchy; each confidence-belief pair that satisfies an n -bounded logic satisfies all m -bounded logics for all $m < n$. Whenever an agent whose confidence-belief pair satisfies an n -bounded logic is certain of a disjunction of n statements, she may believe the negations of *at most* $n-1$ of the disjuncts. However it remains possible for such agents to believe the negations of each of a larger number of statements and yet be certain of their disjunction. The *preface* “paradox” illustrates this aspect of an n -bounded logic quite well.

Suppose that an agent writes a book consisting of $k-1$ pages. When the book is completed, she has checked each page, and believes it to be error free. Let E_i say there is an error on page i . Then we may represent the agent’s doxastic state about her book as follows: $\text{Bel}_\alpha[\neg E_1]$, ..., $\text{Bel}_\alpha[\neg E_{k-1}]$. On the other hand, given the length of the book and the difficulty of the subject, the agent also believes that there is an error on at least one page: $\text{Bel}_\alpha[E_1 \vee \dots \vee E_{k-1}]$. (And she may well say in the preface of her book that there is bound to be an error somewhere – thus the name of this paradox.)

One might think that such a collection of beliefs is incoherent on its face – that real agents maintain such collections of beliefs only because real agents fall short of logical omniscience. Indeed, if an agent is warranted in believing the conjunction of any two beliefs, and if she holds preface-like beliefs, as just described, then she must also be warranted in believing a pretty simple logical contradiction. For, she is warranted in believing $(\neg E_1 \cdot \neg E_2)$, and then warranted in

believing $(\neg E_1 \cdot \neg E_2 \cdot \neg E_3)$, and then ..., and then warranted in believing $(\neg E_1 \cdot \neg E_2 \cdot \dots \cdot \neg E_{k-1})$, and then warranted in believing $(\neg E_1 \cdot \neg E_2 \cdot \dots \cdot \neg E_{n-1} \cdot (E_1 \vee E_2 \vee \dots \vee E_{k-1}))$. So, if the correct *logic of belief* warrants belief in the conjunction of beliefs, and if believing simple contradictions is a doxastic failure of the agent, then preface-like beliefs can only be due to an agent's logical fallibility, her inability to see that her beliefs imply that she should believe a contradiction, which should in turn force her to give up at least one of those beliefs.

The confidence-belief logic I've been articulating puts the preface paradox in a different light. If belief behaves like certainty, like probability 1, then the conjunction rule for beliefs should indeed hold. However, the confidence-belief logic we've been investigating permits belief to behave like probability above a threshold $q < 1$. It allows that the agent may well believe two statements without believing their conjunction, just as happens with probabilities, where it may well be that $P_\alpha[A] \geq q$ and $P_\alpha[B] \geq q$ while $P_\alpha[A \cdot B] < q$. Similarly, according to the confidence-belief logic, the agent is not required to believe the conjunction of individual beliefs. So the kind of doxastic state associated with the preface "paradox" is permissible. However, there are still important constraints on such collections of beliefs. The $((n-1)/n)$ rule is one such constraint.

It will turn out that the confidence-belief logic is probabilistically modelable – that for each confidence-belief pair, there is a probability function and a threshold level that models it. Given that fact, it should be no surprise that the confidence-belief logic behaves like probability-above-a-threshold with regard to conjunctions of beliefs. For, whenever a confidence-belief pair is modelable by a probability function P_α at a threshold level q , if $q > (n-1)/n$, then the $((n-1)/n)$ rule must hold.¹⁹

¹⁹ To see this, let $[\geq_\alpha, \text{Bel}_\alpha]$ be a belief-confidence pair that is probabilistically modelable at some

To see what this means for the logic of belief, suppose, for example, that the agent's confidence-belief pair behaves like probability with a belief bound q just a bit over $9/10$. Then she must satisfy the (9/10) rule:

(9/10): if $\text{Cert}_\alpha[A_1 \vee \dots \vee A_{10}]$, then not ($\text{Bel}_\alpha[\neg A_1]$ and ... and $\text{Bel}_\alpha[\neg A_{10}]$)

She is certain of tautologies, so in the *preface paradox* case we've been discussing for any k page book, $\text{Cert}_\alpha[E_1 \vee \dots \vee E_k \vee \neg(E_1 \vee \dots \vee E_k)]$ always holds. Now, according to the (9/10) rule, if the number of pages in her book is $k \leq 9$, she cannot be in the doxastic state associated with the preface paradox – i.e. she *cannot* ($\text{Bel}_\alpha[\neg E_1]$ and $\text{Bel}_\alpha[\neg E_2]$ and ... and $\text{Bel}_\alpha[\neg E_9]$ and $\text{Bel}_\alpha[E_1 \vee \dots \vee E_9]$).²⁰ However, provided that her book contains $k > 9$ pages, the 10-bounded logic of confidence-belief pairs (associated with the (9/10) rule) *permits* her to be in a doxastic state like that of the preface paradox – she *may* $\text{Bel}_\alpha[\neg E_1]$ and $\text{Bel}_\alpha[\neg E_2]$ and ... and $\text{Bel}_\alpha[\neg E_9]$ and

threshold level $q > (n-1)/n$. So, $\text{Bel}_\alpha[A]$ holds *just when* $P_\alpha[A] \geq q$; and $A \geq_\alpha B$ *just when* $P_\alpha[A] \geq P_\alpha[B]$. Suppose that $\text{Cert}_\alpha[A_1 \vee \dots \vee A_n]$. We show that not $\text{Bel}_\alpha[\neg A_i]$ for at least one of the A_i . $\text{Cert}_\alpha[A_1 \vee \dots \vee A_n]$ implies that $P_\alpha[A_1 \vee \dots \vee A_n] = 1$ (since $\text{Cert}_\alpha[A]$ holds *just when* $A \geq_\alpha (A \vee \neg A)$). Thus, $1 = P_\alpha[A_1 \vee \dots \vee A_n] \leq P_\alpha[A_1] + \dots + P_\alpha[A_n] = (1 - P_\alpha[\neg A_1]) + \dots + (1 - P_\alpha[\neg A_n]) = n - (P_\alpha[\neg A_1] + \dots + P_\alpha[\neg A_n])$. So $P_\alpha[\neg A_1] + \dots + P_\alpha[\neg A_n] \leq (n - 1)$. Now, if for each A_i , $P_\alpha[\neg A_i] \geq q > (n-1)/n$, then we would have $n-1 = n \cdot ((n-1)/n) < n \cdot q \leq P_\alpha[\neg A_1] + \dots + P_\alpha[\neg A_n] \leq (n-1)$ – contradiction! Thus, $P_\alpha[\neg A_i] < q$ for at least one of the A_i – i.e. not $\text{Bel}_\alpha[\neg A_i]$ for at least one of the A_i . This establishes the *probabilistic soundness* of rule $((n-1)/n)$ for all thresholds $q > n-1/n$.

²⁰ Here E_1 through E_9 are A_1 through A_9 , respectively, of the (9/10) rule. And in this example A_{10} of the (9/10) rule corresponds to $\neg(E_1 \vee \dots \vee E_9)$.

$\text{Bel}_\alpha[E_1 \vee \dots \vee E_9]$. But, of course, the logic doesn't *require* her to be in such a state.

More generally, any given n -bounded logic *permits* its confidence-belief pairs to satisfy each of $\{\text{Cert}_\alpha[A_1 \vee \dots \vee A_k], \text{Bel}_\alpha[\neg A_1], \dots, \text{Bel}_\alpha[\neg A_k]\}$ (for some agents α and statements A_i) when $k > n$, but absolutely forbids this (for all agents α and statements A_i) when $k \leq n$. This behavior is characteristic of any confidence-belief logic that arises from a probability function with a belief bound just above $(n-1)/n$.

4.2 The *Lottery* and the $(n+1)^*$ -Bounded Belief Logics

The rules described thus far characterize lower bounds, $(n-1)/n$, on the confidence threshold required for belief. A further hierarchy of rules characterizes upper bounds on the belief modeling confidence thresholds. I'll first describe these rules formally, and then explain them more intuitively in terms of how they capture features of a version of the *lottery paradox*.

Consider the following rule:

(2/3)*: if $\text{Cert}_\alpha[\neg(A \cdot B)]$ and $\text{Cert}_\alpha[\neg(A \cdot C)]$ and $\text{Cert}_\alpha[\neg(B \cdot C)]$, then $\text{Bel}_\alpha[\neg A]$ or $\text{Bel}_\alpha[\neg B]$ or $\text{Bel}_\alpha[\neg C]$.

This rule says that if the agent is certain that *no two of the three* statements A , B , and C is true, then she should also believe the negation of at least one of them. This rule is probabilistically sound in that, given any probability function P_α and any specific threshold value $q \leq 2/3$, the corresponding confidence relation \geq_α and belief predicate Bel_α , defined as ' $A \geq_\alpha B$ iff $P_\alpha[A] \geq P_\alpha[B]$, and $\text{Bel}_\alpha[C]$ iff $P_\alpha[C] \geq q$, where $q \leq 2/3$ ', satisfies all of the previous rules together with the (2/3) rule. If the agent's threshold for belief is no higher than $2/3$, then she has to believe that at least one of a mutually exclusive triple is false.²¹

²¹ For, suppose P_α models the $[\geq_\alpha, \text{Bel}_\alpha]$ pair with a threshold for belief $q \leq 2/3$. For mutually

More generally, consider the following $(n/(n+1))^*$ rule for any fixed $n \geq 2$:

$(n/(n+1))^*$: for any $n+1$ sentences S_1, \dots, S_{n+1} , if for pairs $(i \neq j)$, $\text{Cert}_\alpha[\neg(S_i \cdot S_j)]$, then

$\text{Bel}_\alpha[\neg S_1]$ or $\text{Bel}_\alpha[\neg S_2]$ or ... or $\text{Bel}_\alpha[\neg S_{n+1}]$.

This rule is probabilistically sound in that, given any probability function P_α and any specific threshold value $q \leq n/(n+1)$, the corresponding confidence relation \geq_α and belief predicate Bel_α , defined as ‘ $A \geq_\alpha B$ iff $P_\alpha[A] \geq P_\alpha[B]$, and $\text{Bel}_\alpha[C]$ iff $P_\alpha[C] \geq q$, where $q \leq n/(n+1)$ ’, satisfies all of the previous rules together with the $(n/(n+1))^*$ rule.²²

According to this rule, if the agent is certain that at most one of the $n+1$ statements is true, then she must believe the negation of *at least* one of them. But notice that when the $(n/(n+1))^*$ rule holds for a confidence-belief pair, the agent is permitted to withhold belief for the negations of fewer than $n+1$ mutually exclusive statements. The rule only comes into play when collections of $n+1$ or more mutually exclusive statements are concerned – and then it requires belief for the negation of at least one of them.

Let’s say that any confidence relation \geq_α together with associated belief predicate Bel_α that satisfies rules 1-7 and the $(n/(n+1))^*$ rule (for a given value of n) *satisfies an $(n+1)^*$ -bounded*

exclusive A , B , and C we have $1 \geq P_\alpha[A \vee B \vee C] = P_\alpha[A] + P_\alpha[B] + P_\alpha[C] = (1 - P_\alpha[\neg A]) + (1 - P_\alpha[\neg B]) + (1 - P_\alpha[\neg C]) = 3 - (P_\alpha[\neg A] + P_\alpha[\neg B] + P_\alpha[\neg C])$, which entails that $P_\alpha[\neg A] + P_\alpha[\neg B] + P_\alpha[\neg C] \geq 2$. So at least one of $P_\alpha[\neg A]$, $P_\alpha[\neg B]$, or $P_\alpha[\neg C]$ must be at least as great as $2/3 \geq q$ (since if each of these three probabilities is less than $2/3$, their sum must be less than 2); so at least one of $\neg A$, $\neg B$, and $\neg C$ must be believed.

²² A probabilistic argument similar to that in the previous note shows the soundness of this rule.

belief logic. A version of the *lottery* “paradox” illustrates the central features of an $(n+1)^*$ -bounded logic quite well.

Lotteries come in a variety of forms. Some are designed to guarantee at least one winner. Some are designed to permit at most one winner. And, of course, some lotteries have both features. Lotteries are usually designed to give each ticket the same chance of winning. But for the purposes of illustrating the $(n/(n+1))^*$ rule we need not suppose this. Indeed, for our purposes we need only consider lotteries that are *exclusive* – where no two tickets can win. I’ll call such lotteries ‘*exclusive lotteries*’. (These lotteries may also guarantee at least one winner – but for our purposes we need not assume that they do).

Let ‘ W_i ’ stand for the statement that ticket i will win, and suppose an agent α is certain that no two tickets can win this particular lottery – i.e. for each pair, $i \neq j$, $\text{Cert}_\alpha[\neg(W_i \cdot W_j)]$. Let’s say that α is in an *m-ticket optimistic state* just in case:

for some collection of at least m tickets (which may be arbitrarily labeled as ‘ticket 1’, ..., ‘ticket m ’), α deems it *genuinely possible* that W_1 (i.e. not $\text{Bel}_\alpha[\neg W_1]$), and ..., and α deems it *genuinely possible* that W_m (i.e. not $\text{Bel}_\alpha[\neg W_m]$).

Consider an agent α whose confidence relation and belief predicate is modeled by a probability function with an explicit threshold value q for belief. Suppose $q = .99$. It is easy to see how α might come to be in an *m-ticket optimistic state* if the *exclusive lottery* has relatively few tickets. For instance, in a lottery with three tickets, she might believe that ticket A has a .40 chance of winning, that ticket B has a .30 chance of winning, and that ticket C has a .20 chance of winning, leaving a .10 chance that none of the tickets will win. Then, for any given ticket i , α does not believe that ticket i will *not* win, since, for each i , her degree of confidence in $\neg W_i$ is smaller

than $q = .99$. Hence, she is in a 3-ticket optimistic state with respect to the 3-ticket lottery. However, for larger and larger lotteries exclusivity will force her to assign lower and lower degrees of confidence to at least some of the W_i . For a sufficiently large lottery, a lottery of 100 or more tickets, her degree of confidence in $\neg W_i$ must come to exceed $q = .99$ for at least one ticket i . Thus, she must believe $\neg W_i$ for at least one ticket i . (If, in addition, she is equally confident regarding the possibility of each ticket winning – i.e. if $W_i \approx_\alpha W_j$ for each i and j – then she must believe of each ticket that it will *not* win.)

The point is that when the quantitative Lockean Thesis holds for α at a threshold level q (or higher) for belief, then the following rule is sound for any value of $m \geq 1/(1-q)$:

$(n/(n+1))^*$: if for each $i \neq j$, $\text{Cert}_\alpha[\neg(W_i \cdot W_j)]$, then $\text{Bel}_\alpha[\neg W_1]$ or ... or $\text{Bel}_\alpha[\neg W_m]$.

This is just the $(n/(n+1))^*$ rule stated another way. (Using a bit of algebra to calculate q in terms of m , the above condition holds just when $q \leq (m-1)/m$. Then, substituting $n+1$ for m , the above rule is just the $(n/(n+1))^*$ rule for $n+1$ statements.) However, for each value of $m < 1/(1-q)$, m -ticket optimistic states remain rationally coherent for α . For then the belief threshold q is above $(m-1)/m$, and the agent may well remain optimistic about the possibility of each of the m tickets winning – i.e. it may well be that not $\text{Bel}_\alpha[\neg W_1]$ and ... and not $\text{Bel}_\alpha[\neg W_m]$.

Notice that for each given value of n , the $(n/(n+1))^*$ rule is perfectly compatible with the $((n-1)/n)$ rule described in the previous section. However the starred and unstarred rules don't fit together at all well when the starred rule takes a fractional value equal to or smaller than the unstarred rule. To see why, consider a confidence-belief relation that has, for a given n , both the rules $(n/(n+1))^*$ and $(n/(n+1))$. These two rules together would say this:

for any $n+1$ sentences S_1, \dots, S_{n+1} ,

if $\text{Cert}_\alpha[S_1 \vee \dots \vee S_{n+1}]$ and for pair (each $i \neq j$), $\text{Cert}_\alpha[\neg(S_i \cdot S_j)]$, then $\text{Bel}_\alpha[\neg S_i]$ for at

least one S_i and not $\text{Bel}_\alpha[\neg S_j]$ for at least one S_j .

Such a rule would *rule out the possibility* of a partition (i.e. a mutually exclusive and exhaustive set of statements S_i) for which the agent is equally confident in each (i.e. where for each i and j , $S_i \approx_\alpha S_j$). That is, for such an agent, in cases where exactly one ticket is certain to win, *no* $n+1$ *ticket lottery could possibly be fair* (could possibly permit equal confidence in each ticket's winning). The *logic alone* would rule that out! Similarly, the starred rule cannot have a *lower* fractional value than the unstarred rule, for the same reason. Thus, the tightest bounds on belief thresholds that properly fits these n -bounded rules corresponds to those confidence-belief logics that have both an $((n-1)/n)$ rule and an $(n/(n+1))^*$ rule.²³

5. The Logic of Belief

Let's now pull together the rules studied in the previous sections to form one grand logic of confidence and belief. Here is how to do that formally:

Definition 4: the Rudimentary n -Level Confidence-Belief Pairs: Given a language L for predicate logic, the *rudimentary n -level confidence-belief pairs* on L are just the pairs $[\geq_\alpha, \text{Bel}_\alpha]$ consisting of a *rudimentary confidence relations* and a *belief predicate* that together satisfy the following rules:

(8) if $\text{Cert}_\alpha[A]$ then $\text{Bel}_\alpha[A]$;

(9) If $A \geq_\alpha B$ and $\text{Bel}_\alpha[B]$, then $\text{Bel}_\alpha[A]$;

²³ See (Hawthorne, 1999, 2007) and (Hawthorne and Makinson, 2007) for a related treatment of the logics of classes of nonmonotonic conditionals that behave like conditional probabilities above a threshold. Rules very similar to the $((n-1)/n)$ and $(n/(n+1))^*$ rules apply there.

Level n rules (for fixed $n \geq 2$):

(10) for any n sentences S_1, \dots, S_n ,

if $\text{Cert}_\alpha[S_1 \vee S_2 \vee \dots \vee S_n]$, then not ($\text{Bel}_\alpha[\neg S_1]$ and $\text{Bel}_\alpha[\neg S_2]$ and ... and $\text{Bel}_\alpha[\neg S_n]$);

(11) for any n+1 sentences S_1, \dots, S_{n+1} ,

if for each $i \neq j$, $\text{Cert}_\alpha[\neg(S_i \cdot S_j)]$, then $\text{Bel}_\alpha[\neg S_1]$ or $\text{Bel}_\alpha[\neg S_2]$ or ... or $\text{Bel}_\alpha[\neg S_{n+1}]$.

The rules for the rudimentary n-level confidence-belief pairs are probabilistically sound in the sense that for any probability function P_α and any specific threshold level $q > 1/2$ such that $n/(n+1) \geq q > (n-1)/n$, the corresponding relation \geq_α and belief predicate Bel_α (defined as $A \geq_\alpha B$ iff $P_\alpha[A] \geq P_\alpha[B]$, and $\text{Bel}_\alpha[C]$ iff $P_\alpha[C] \geq q$) must satisfy rules 1-7 and 8-11. However, as with the rudimentary confidence relations, some confidence-belief pairs are not constrained enough by these rules to be modelable by a probability function. But that is easy to fix using precisely the same kind of rule that worked for selecting the (refined) confidence relations from the rudimentary ones.

Definition 5: the Properly Extendable Rudimentary n-Level Confidence-Belief Pairs:

Let us say that a rudimentary n-level confidence-belief pair $[\geq_\alpha, \text{Bel}_\alpha]$ on a language L is *properly extendable* just in case there is a rudimentary confidence-belief pair $[\geq_\beta, \text{Bel}_\beta]$ on some language L^+ an extension of L that agrees with the determinate part of \geq_α and Bel_α (i.e. whenever $A \approx_\alpha B$, $A \approx_\beta B$; whenever $A >_\alpha B$, $A >_\beta B$; and whenever $\text{Bel}_\alpha[C]$, $\text{Bel}_\beta[C]$) on the shared language L, and also satisfies the following rule for all sentences of L^+ :

(XX) (i) (*completeness*): Either $A \geq_\beta B$ or $B \geq_\beta A$; and

(ii) (*separating equiplausible partitions*): If $A >_\beta B$, then, for some (large enough) n, there are n sentences S_1, \dots, S_n that β takes to be mutually incompatible

(i.e., $\text{Cert}_\beta[\neg(S_i \cdot S_j)]$ for $i \neq j$), and jointly exhaustive (i.e., $\text{Cert}_\beta[S_1 \vee \dots \vee S_n]$, where β is equally confident in each (i.e. $S_i \approx_\beta S_j$ for each i, j), such that for each of them, $A >_\beta (B \vee S_k)$.

This is exactly like Definition 2 for *properly extendable rudimentary confidence relations*, but adds to it that the extended belief predicate agrees with the belief predicate on the sentences in the shared language L .

Now we may specify the (refined) *n-level confidence-belief pairs* in the obvious way.

Definition 6: the n-Level Confidence-Belief Pairs: Given a language L for predicate logic, the (refined) *n-level confidence-belief pairs* on L are just the *rudimentary n-level confidence-belief pairs* $[\geq_\alpha, \text{Bel}_\alpha]$ on L that are *properly extendable*.

The logic of the *n-level confidence belief pairs* is sound and complete with respect to probability functions and corresponding belief thresholds.

Theorem 3: Probabilistic Soundness for n-level confidence-belief pairs: Let P_α be any probability function (that satisfies the usual axioms). Define the relation \geq_α as follows: $A \geq_\alpha B$ just when $P_\alpha[A] \geq P_\alpha[B]$. And for any $q \geq 1/2$, define Bel_α in terms of threshold level q in any one of the following ways:

- (i) $q = (n-1)/n$ for fixed $n \geq 2$, and for all S , $\text{Bel}_\alpha[S]$ just when $P_\alpha[S] > q$, or
- (ii) $n/(n+1) > q > (n-1)/n$ for fixed $n \geq 2$, and for all S , $\text{Bel}_\alpha[S]$ just when $P_\alpha[S] \geq q$, or
- (iii) $n/(n+1) > q > (n-1)/n$ for fixed $n \geq 2$, and for all S , $\text{Bel}_\alpha[S]$ just when $P_\alpha[S] > q$, or
- (iv) $q = (n-1)/n$ for fixed $n \geq 3$, and for all S , $\text{Bel}_\alpha[S]$ just when $P_\alpha[S] \geq q$.

Then the pair $[\geq_\alpha, \text{Bel}_\alpha]$ satisfies rules 1-9 and the level n versions of rules (10) and (11).

Theorem 4: Probabilistic Completeness for n-level confidence-belief pairs: For each n-level confidence-belief pair $[\geq_\alpha, \text{Bel}_\alpha]$ (i.e. each pair satisfying Definition 6), there is a probability function P_α and a threshold q that models \geq_α and Bel_α as follows: for all sentences A and B , (1) if $P_\alpha[A] > P_\alpha[B]$, then $A >_\alpha B$ or $A \sim_\alpha B$; (2) if $P_\alpha[A] = P_\alpha[B]$, then $A \approx_\alpha B$ or $A \sim_\alpha B$; and one of the following clauses holds:

- (i) $q = (n-1)/n$ for fixed $n \geq 2$, and $P_\alpha[S] > q$ just when $\text{Bel}_\alpha[S]$, or
- (ii) $n/(n+1) > q > (n-1)/n$ for fixed $n \geq 2$, and $P_\alpha[S] \geq q$ just when $\text{Bel}_\alpha[S]$, or
- (iii) $n/(n+1) > q > (n-1)/n$ for fixed $n \geq 2$, and $P_\alpha[S] > q$ just when $\text{Bel}_\alpha[S]$, or
- (iv) $q = (n-1)/n$ for fixed $n \geq 3$, and $P_\alpha[S] \geq q$ just when $\text{Bel}_\alpha[S]$.

Furthermore, if \geq_α itself satisfies rule X, then P_α and q are unique.

Theorem 4 shows us precisely how the Qualitative Lockean Thesis is satisfied. It tells us that each confidence relation and belief predicate that satisfies the n-level rules (10) and (11) (for a specific value of n) may be modeled by a probability function and a suitable threshold level q in the range between $n/(n+1)$ and $(n-1)/n$ (as specified by one of (i)-(iv)). Furthermore, at the end of section 3 we saw that any given probabilistic model may be *overly precise*, specifying definite relative confidence relationships that go beyond those the agent is willing to accept. This point continues to hold for probabilistic models with thresholds of confidence and belief. Thus, an agent's confidence-belief pair may be better represented (or modeled) by a set of probability-function–threshold-level pairs that capture the agent's incomplete (indefinite) assessment of comparative confidence relationships among some statements.

6. Concluding Remarks

A *qualitative logic of confidence and belief* fits well with the Lockean Thesis. This logic is based in the logic of the *at least as confident as* relation (i.e., in the logic of qualitative probability) extended to accommodate a qualitative *belief* threshold. It turns out that this logic may be effectively modeled by quantitative probability functions together with numerical thresholds for *belief*. Thus, for this logic, the Qualitative Lockean Thesis is recapitulated in an underlying quantitative model that satisfies the Quantitative Lockean Thesis.

The version of qualitative probabilism associated with the Qualitative Lockean Thesis needn't suppose that the agent has anything like precise numerical degrees of belief. Indeed, it doesn't even suppose that the agent can determine definite confidence-comparisons between all pairs of statements. Rather, the Representation Theorems show how a qualitative confidence relation and corresponding belief predicate may give rise to a *set* of degree-of-belief functions and associated numerical thresholds, where the *set* reflects whatever imprecision is already in the ideal agent's qualitative confidence relation and qualitative belief predicate. We may *model* the agent with any one of these (overly precise) functions and numerical thresholds, keeping in mind that the *quantitative* function is only one of a number of equally good numerical representations. So, for qualitative probabilism, the appeal to sets of representing probability functions is a natural consequence of the incompleteness (or indeterminateness) of the ideal agent's relative confidence relation – rather than merely a device for adding indefiniteness back into a quantitative model that was overly precise from the start.

I'll now conclude with a few words about how this logic of confidence and belief may be further extended and developed.

The logic presented here only provides a static model of confidence and belief. A dynamic model would add an account of confidence/belief updating – an account of the logic of an agent's

transitions from one confidence/belief model to another based on the impact of evidence. The deep connection with probability makes it relatively easy to see how standard accounts of probabilistic belief dynamics – e.g., Bayesian updating, and Jeffrey Updating – may be adapted to qualitative confidence and belief. Since an agent’s qualitative confidence-relation/belief-predicate pair can be modeled as a set of probability functions with numerical belief thresholds, schemes for updating quantitative degrees of confidence suggest approaches to updating qualitative confidence and belief as well.

One way in which real belief may be more subtle than the model of belief captured by the Lockean Thesis as explored thus far is that real belief seems to have a contextual element. The level of confidence an agent must have in order for a statement to qualify as *believed* may depend on various features of the context, such as the subject matter and the associated doxastic standards relevant to a given topic, situation, or conversation. The logic investigated here is easily extended to handle at least some facets of this context-sensitivity of *belief*. To see how, consider the following modification of the Lockean Thesis:

Contextual Qualitative Lockean Thesis: An agent is epistemically warranted in believing a statement in a context ψ *just in case* she is epistemically warranted in having a sufficiently high *grade of confidence* in the statement – sufficiently high to make her attitude towards it one of *belief* in context ψ .

The idea is that rather than represent the doxastic state of an agent α by a single confidence/belief pair, we may represent it as a confidence relation together with a list of belief predicates, $[\geq_{\alpha}, Bel_{\alpha}^{\phi}, Bel_{\alpha}^{\psi}, \dots, Bel_{\alpha}^{\chi}]$, where each belief predicate Bel_{α}^{ψ} is associated with a specific kind of context ψ , where each pair $[\geq_{\alpha}, Bel_{\alpha}^{\psi}]$ constitutes an n-level Confidence/Belief pair (as specified in Definition 6) appropriate to the context. Then we simply specify that α believes S in context ψ

($\text{Bel}_\alpha[S]$ in context ψ) just when $\text{Bel}_\alpha^\psi[S]$ holds.²⁴

Variations on this approach may be employed to represent additional subtleties. For example perhaps *only certain kinds of statements* (e.g. those about a specific subject matter) are “doxastically relevant or appropriate” for a given context. We may model this by restricting the contextually sensitive predicate Bel_α^ψ to only those statements considered relevant or appropriate in the context. Thus, although $Q \geq_\alpha R$ and $\text{Bel}_\alpha^\psi[R]$ holds, α may fail to believe Q in context ψ because this context itself excludes Q from consideration. (E.g., relative to the context we form a new confidence relation \geq_α^ψ by dropping context-irrelevant statements like Q from the full confidence relation \geq_α . We may then characterize the belief predicate Bel_α^ψ appropriate to the context so as to satisfy Definition 6 for the confidence/belief pair $[\geq_\alpha^\psi, \text{Bel}_\alpha^\psi]$.) Thus, the qualitative logic of confidence and belief that attends the Qualitative Lockean Thesis should be sufficiently flexible to represent a range of additional features of confidence and belief.²⁵

REFERENCES

²⁴ Alternatively, we might represent this contextuality in terms of a single two-place belief relation $\text{Bel}[S, \psi]$ between statements and contexts together with a single three-place confidence relation between pairs of statements and contexts.

²⁵ Thanks to Luc Bovens and Franz Huber and Martin Montminy for their many helpful comments on drafts of this paper. This paper is in part a response to Greg Wheeler’s (2005) critique of (Hawthorne and Bovens, 1999). I thank Greg for pressing me on these issues, both in his paper and through private communications. I owe a special debt to Luc Bovens for the insights he contributed to our earlier paper, which provided the impetus for this one.

- Christensen, David 2004: *Putting Logic in its Place*. Oxford: Oxford University Press.
- Foley, Richard 1992: "The Epistemology of Belief and the Epistemology of Degrees of Belief", *American Philosophical Quarterly*, 29 (2), 111-121.
- 1993: *Working Without a Net*. Oxford: Oxford University Press.
- Hawthorne, James 1996: "On the Logic of Nonmonotonic Conditionals and Conditional Probabilities", *Journal of Philosophical Logic*, 25 (2), 185-218.
- Hawthorne, James 2007: "Nonmonotonic Conditionals that Behave Like Conditional Probabilities Above a Threshold", *Journal of Applied Logic*, 5, (4), 625-637.
- Hawthorne, James and Bovens, Luc 1999: "The Preface, the Lottery, and the Logic of Belief", *Mind*, 108, 430, 241-264.
- Hawthorne, James and Makinson, David 2007: "The Quantitative/Qualitative Watershed for Rules of Uncertain Inference", *Studia Logica*, 86 (2), 247-297.
- Hawthorne, John 2004: *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Joyce, James 1999: *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Kyburg, Henry E., Jr. 1961: *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- 1970: "Conjunctivitis", in Marshall Swain (ed.), *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel.
- Locke, John 1975: *An Essay Concerning Human Understanding*, P.H. Nidditch, ed. Oxford: Clarendon Press.
- Makinson, David C. 1965: "The Paradox of the Preface", *Analysis*, 25, 205-207.
- Wheeler, Gregory 2005: "On the Structure of Rational Acceptance: Comments on Hawthorne

and Bovens”, *Synthese*, 144 (2), 287-304.

Savage, Leonard J. 1972: *The Foundations of Statistical Inference*. New York: Dover.