The Preface, the Lottery, and the Logic of Belief

JAMES HAWTHORNE AND LUC BOVENS

John Locke proposed a straightforward relationship between qualitative and quantitative doxastic notions: belief corresponds to a sufficiently high degree of confidence. Richard Foley has further developed this Lockean thesis and applied it to an analysis of the preface and lottery paradoxes. Following Foley's lead, we exploit various versions of these paradoxes to chart a precise relationship between belief and probabilistic degrees of confidence. The resolutions of these paradoxes emphasize distinct but complementary features of coherent belief. These features suggest principles that tie together qualitative and quantitative doxastic notions. We show how these principles may be employed to construct a quantitative model—in terms of degrees of confidence—of an agent's qualitative doxastic state. This analysis fleshes out the Lockean thesis and provides the foundation for a logic of belief that is responsive to the logic of degrees of confidence.

I

We often express our doxastic states in qualitative terms. We speak of doubting, believing, assenting to, accepting, and being certain of various propositions. But in some contexts quantitative terms seem more fitting. We speak of the strength of our beliefs, of how likely we think they are, of our degrees of confidence in various propositions. We move easily between these qualitative and quantitative doxastic notions, but sometimes there are subtle tensions. These tensions are brought out starkly in two related paradoxes, the *lottery* and the *preface.*[1]

A standard version of the lottery paradox goes like this. You receive an entry to the Publisher's Hucksterhouse Sweepstakes saying that you are the lucky winner of ten million dollars—if you hold the winning ticket. You have no illusion that your life is about to change. You may well believe that the ticket in your hands is of no value at all, that it just is not the winning ticket. At the same time, you may have no doubt that the sweepstakes is fair, that each entry ticket has an equal chance of winning and that someone will surely collect the money. If, however, you believe that *your* ticket will not win, and you consider any two tickets to have the same

[1] The lottery paradox was first discussed by Henry Kyburg (1961). (Also see Kyburg 1970.) The preface paradox was first described by David Makinson (1965).

Mind, Vol. 108 . 430 . April 1999

chance of winning, then you (should) also believe that your neighbour's ticket will not win, and believe that your neighbour's neighbour's ticket will not win, and … . This seems inconsistent with your belief that someone is bound to win.

A standard version of the preface paradox runs as follows. You have just finished writing your *magnum opus*, a tome of several thousand pages. You have carefully edited your work page by page. As a matter of fact, in your last reading, you stuck to your resolution to move on to the next page only when quite confident that all of the claims on the previous page were accurate. Subsequently, as you draft the preface to your work you add the usual disclaimer: "I wish to express my gratitude to my colleagues for their many thoughtful comments and suggestions on earlier drafts; any remaining errors—and there are bound to be some—are entirely my own". If this assertion truly expresses your belief that some errors remain, if it is not just self-effacing modesty, then it may seem inconsistent with your confidence in the veracity of each page, considered one by one.[2]

These seeming inconsistencies may lead one to think that such *preface* beliefs and *lottery* beliefs are symptoms of human fallibility, that they violate *logical* constraints on rationally coherent belief. But accusations of irrationality or incoherence seem rather audacious here. When stepping onto a commercial airliner, we believe that it will not crash; and we have equally good reason to believe, of each commercial airline flight in the coming year, that it too will not crash. But we may also believe that some commercial airliner will crash on some trip in the coming year. If such beliefs are jointly incoherent, then incoherent belief is indeed ubiquitous.[3]

The fact that lottery and preface beliefs are common in everyday life is no argument that they are rational in any *logical* sense. Much work has been done in philosophy and psychology to unveil common patterns of human logical fallibility, and someone might argue that the *preface* and the *lottery* draw our attention to just such patterns. One diagnosis is that such fallibility stems from a person's resistance to conjoining individual

[2] The preface paradox is sometimes treated as a paradox of self-reference, but we will not be treating it this way here. This is not to say that self-referential readings are uninteresting. But there is paradox enough in the non-self-referential version. One can shortcut the self-referential reading by not including the preface as part of the *work*. By the *work* we mean the part of the author's writings in which the pages are typically numbered by Arabic numerals, not those with Roman numerals.

[3] One might try to avoid this sort of curious doxastic situation by reserving *belief* to propositions of which one is absolutely certain. Almost no one is *certain* of each flight that it will not crash and is also *certain* that at least one airliner will crash in the coming year. And anyone who is certain of all of these propositions at once may quite plausibly be charged with holding irrational or incoherent beliefs. However, if all belief is to be restricted to *certainties*, then it seems that most rational people believe very little indeed.

beliefs. If only people were willing to conjoin their beliefs, the inconsistencies might become apparent, and they should then feel compelled to abandon or revise some of their beliefs in order to secure rational coherence. Indeed, the compartmentalization of beliefs is a common pattern of irrationality. For instance, people often hold mutually inconsistent, self-serving moral convictions, and their resistance to conjoining these beliefs is a kind of self-deception that facilitates their self-serving ways. However, it is not at all clear that the *preface* and the *lottery* are examples of this kind of irrationality. To the contrary! We will argue that even ideally rational agents may coherently hold beliefs like those of the *lottery* and the *preface.* If such beliefs are jointly coherent for ideally rational agents, then they may also be coherently held by common folk.

If *preface* and *lottery* beliefs are redescribed in quantitative doxastic terms, their paradoxical features evaporate. In the *lottery* we realize that the likelihood that any given ticket will win is extremely low, yet this in no way contradicts our certainty that some ticket will win. In the *preface* we judge that the likelihood that any given page still contains an error is extremely low, yet this is perfectly consistent with our high degree of confidence that at least one error has been missed in a lengthy book.

These observations, however, do not close the book on the *lottery* and the *preface*. Rather, they invite further reflection on the relationship between qualitative and quantitative doxastic notions. In a penetrating investigation of this relationship, 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 sufficiently high degree of confidence in it, sufficiently high to make our attitude towards it one of belief. (p. 111)

Foley goes on to suggest that *rational belief* is 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*.[4] Foley employs the *lottery* and the *preface* to probe the Lockean Thesis. We will follow his lead and probe yet more deeply.

The Lockean thesis is a normative rather than a descriptive thesis. Actual human agents generally do not have precise degrees of confidence in propositions. To the extent that people do employ numerical degrees of confidence, such confidence levels may not consistently link up with belief in the way that the Lockean thesis recommends. However, as a normative claim the Lockean thesis is impeccable. We will show how the link

[4] Foley cites Locke's *An Essay Concerning Human Understanding*, especially Book IV, Chs. xv–xvi.

that it suggests between beliefs and degrees of confidence may be exploited to provide a completely satisfactory treatment of the *preface* and the *lottery* and to explicate significant features of the logic of belief.

As paradoxes, the *preface* and the *lottery* are clearly akin. But how closely related are they? Do they afford different insights into epistemic logic, or are they just different ways of making the same point? We will show that the *preface* and the *lottery* illuminate complementary facets of the relationship between qualitative and quantitative doxastic notions. By exploiting various versions of preface and lottery scenarios we will show that qualitative doxastic notions can be modeled quite precisely by quantitative notions, and *vice versa*. On the qualitative side we employ the notions of *believing that Q*, *being certain that Q*, and *deeming Q and R equally plausible*.[5] On the quantitative side, we consider the probabilistic doxastic notion of *degree of confidence.* We will chart a precise relationship between these qualitative and quantitative doxastic notions, a relationship that will provide a foundation for a very plausible account of the logic of rationally coherent belief.

II

We begin by considering some implications of the preface paradox for the relationship between *belief* and *degrees of confidence*. Let "E_i" stand for the proposition that there is an error on page *i* of a given work. For a fixed number of pages *n*, we say that an agent α is in an *n*-page *preface state* just in case:

α believes $\neg E_1$, α believes $\neg E_2$, … , α believes $\neg E_n$,
but also α believes ($E_1 \quad E_2 \quad \dots \quad E_n$).

The preface paradox dissipates in the following two limiting cases. First, suppose that an agent has written a one-page flyer. It would be odd for her to be in a preface state about it. This would amount to believing E_1 but also believing $\neg E_1$, which violates a very plausible constraint on rational belief. Second, suppose that someone is an extreme skeptic in the sense that he believes a proposition just in case he is absolutely certain of it. If, considering each page of his book individually, he is *certain* that *that* page has no errors, then it would be quite odd for him to be *certain* that at least one page of his book has an error. Indeed, in the context of *certain* belief,

[5] *Being certain that Q* is eliminable in the presence of the notion of *equal plausibility*; it may be defined as *deeming Q and* **T** *equally plausible*, for some obvious tautology, **T**. Also, it turns out that the notion of *deeming Q and R equally plausible* is eliminable in the presence of the notions of *belief* and *certainty*, although the technique for doing this is rather complex. Details are given in Appendix 2.

it is rational for an agent to believe with certainty the conjunction of any pair of propositions that he believes with certainty. Hence, this agent should be certain that $(\neg E_1 \,\&\, \neg E_2 \ldots \&\, \neg E_n \,\&\, (E_1 \vee E_2 \vee \ldots \vee E_n))$—that is, he should believe an obvious contradiction. This surely violates a constraint on rational belief.

As we move away from these limiting cases, intuitively it becomes more and more reasonable for an agent to enter a preface state. On the one hand, consider an agent who is not a skeptic of the sort who requires certainty for belief. It would be odd for her to be in a preface state about a one-page flyer. It would be a bit less odd for her to be in a preface state about a three-page pamphlet. As the number of pages increases it seems more and more reasonable for an agent to be in a preface state. For the other limiting case, consider an average-sized book. We would not expect an extreme skeptic to be in a preface state. But as we encounter agents who require less and less certitude for belief, we may find it more and more reasonable for them to be in a preface state. In the next two sections we will show how the Lockean thesis may explain all of this quite handily.

III

Suppose an agent α has explicit degrees of confidence, especially with regard to the veracity of her book. She assigns degrees of confidence to propositions on a scale from 0 to 1, and assigns 1 to those propositions of which she is certain. Suppose also that whenever α is certain that $\neg(A \,\&\, B)$, her degree of confidence in $(A \vee B)$ is the sum of her degree of confidence in A and her degree of confidence in B. If in addition α is certain of all logical truths, we call her an *ideally rational* agent. The degree-of-confidence-function of any ideally rational agent can be represented by a classical probability function.[6] Now let us also suppose that α satisfies the Lockean thesis for some specific threshold value q. That is, for α *belief* is merely a convenient way to categorize those

[6] More formally, to say that an agent's degrees of confidence in propositions can be represented by a probability function just amounts to saying that her degree-of-confidence function P is constrained by the following rules: (i) For any proposition S, $P(S) \geq 0$; (ii) For any proposition S, if S is a logical truth, then $P(S) = 1$; (iii) For any two propositions R and S, if $\neg(R \,\&\, S)$ is a logical truth, then $P(R \vee S) = P(R) + P(S)$. Such an agent is clearly an idealization. In particular, she is logically omniscient in the sense that when contemplating any logical truth her degree of confidence in it will be 1. However, it will suffice for most purposes in this paper if our ideally rational agent has just enough logical knowledge to recognize some easily decidable logical truths and equivalences of propositional logic. We will discuss the implications of our analysis for the doxastic states of real human agents in the penultimate section.

propositions for which her degree of confidence is no less than some threshold value q that she considers significantly high. If asked whether she believes proposition S, α may even explicitly report that her degree of confidence in S is no less than q, and that since she takes q to be an adequate threshold value for *belief*, she does indeed believe S.[7]

It is easy to see how a preface state can arise for an agent whose doxastic states satisfy the Lockean thesis in the way that α's do. If probability function P represents α's degrees of confidence, then for her to be in a preface state just amounts to this: for each page i, $P(\neg E_i) \geq q$, but also $P(E_1 \quad \ldots \quad E_n) \geq q$. This may occur when $P(\neg E_1 \,\&\, \neg E_2 \,\&\, \ldots \,\&\, \neg E_n)$ becomes very small, so that $1 - P(\neg E_1 \,\&\, \neg E_2 \,\&\, \ldots \,\&\, \neg E_n) = P(\neg(\neg E_1 \,\&\, \neg E_2 \,\&\, \ldots \,\&\, \neg E_n)) = P(E_1 \quad E_2 \quad \ldots \quad E_n)$ rises above the threshold value q for α's belief. There is, however, a strict constraint on the rational coherence of such preface states. If such a state arises, then the number of pages n must be no smaller than $q/(1-q)$. (See Appendix 1, Theorem 1, Case 1.) So, when $q = .99$, no rationally coherent preface state may arise unless the book contains at least 99 pages. Thus, if the Lockean thesis holds for α at threshold level q, then *rational coherence* implies the soundness of the following rule for all values of $n < q/(1-q)$:

> if α believes $\neg E_1$, α believes $\neg E_2$, … , α believes $\neg E_n$,
>
> then α does not believe $(E_1 \quad E_2 \quad \ldots \quad E_n)$.[8]

However, for each value of $n \geq q/(1-q)$ preface states may remain perfectly coherent for α.

This constraint on the rational coherence of preface states does *not* presuppose that the errors in the book are independent of one another. Independence among errors would mean that the commission of one error does not tend to induce (or impede) the commission of additional errors. For

[7] For a real agent the threshold level at which degrees of confidence correspond to belief may well be context or domain dependent. An agent may "believe" such mundane claims as that the train will be on schedule this morning without having a very high degree of confidence. The same agent may require a much higher degree of confidence before coming to *believe*, say, a scientific claim. But no special difficulty for our analysis ensues. The agent simply has two different standards for belief, but may easily compare these standards in terms of her degrees of confidence. She may readily report that although she "believes" that the train will be on time, under the standard she employs when contemplating belief in scientific claims, she does not *believe* that the train will be on time. Our analysis applies to any single belief standard, and may be applied to each of a number of standards, one by one.

[8] Similar rules are satisfied by certain kinds of nonmonotonic conditionals that represent defeasible support relations among propositions. These conditionals behave like conditional probabilities above a given threshold level, conditionals for which "if C, then B" means roughly, "if C holds, then B is very probably true". For details see Hawthorne (1996), particularly the rules for System Q (Definition 18).

contexts in which α takes errors to be independent, if a preface state arises, then n must be no smaller than $\log(1 - q)/\log(q)$. (See Appendix 1, Theorem 2.) So, in a context where α takes errors to be independent and her belief threshold is $q = .99$, no coherent preface state can arise unless the book contains at least 459 pages.

IV

In the previous section we contemplated the doxastic states of an agent who has an explicit degree-of-confidence function for propositions concerning the veracity of her book and an explicit threshold for belief. We were then able to determine constraints on her preface states in terms of the number of pages of the book. Now consider a second agent, β, who has no explicit degree-of-confidence function at all. His doxastic states simply involve believing or not believing various propositions. But suppose that β is also *ideally rational* in that his beliefs satisfy the Lockean thesis *in principle*. That is, there exists some degree-of-confidence function and some threshold value that would yield a categorization of propositions into *believed* and *not-believed* that corresponds to β's actual doxastic states. When this condition is satisfied, we (still) call the agent *ideally rational* and call the agent's beliefs *rationally coherent.*[9] There may well be more than one degree-of-confidence function and threshold level that will accurately model β's beliefs in conformity with the Lockean thesis; uniqueness is not required. But let us suppose now that neither β nor anyone else has an inkling of which degree-of-confidence functions and threshold values would properly model β's doxastic states in this way.[10]

[9] The supposition that β's beliefs satisfy the Lockean thesis in this sense is actually quite a weak one. InAppendix 2 we describe two additional notions of *rationally coherent* belief that may seem to be somewhat weaker than the probabilistic notion of *rational coherence* (i.e. weaker than compatibility with a degree-of-confidence function and a threshold value). However, it turns out that the beliefs of any agent who is *ideally rational* in the sense that her beliefs conform to either of these other notions of *rational coherence* will automatically satisfy the Lockean thesis.

[10] One way in which a real agent may fail to satisfy the Lockean thesis *in principle* is that he may employ a lower standard for belief in some contexts or for some domains than for others. Although he may not utilize an explicit degree-of-confidence function, diverse belief standards may show up as follows: he "believes" that the morning train will be on time, he does not *believe* that birds evolved from dinosaurs, but he admits to greater confidence in the truth of the second claim than in the truth of the first. Such an agent may readily admit that relative to the standard he applies to scientific claims, he does not *believe* that the train will be on time. Here again our analysis is intended to apply to any single belief standard, and may be applied to each of a number of standards, one by one. (Cf. note 7.)

Now, it happens that β is the author of a series of works—flyers, pamphlets, papers, moderately sized books, and tomes—and he has written each with integrity. For each work, he believes of each page that it contains no errors. Nevertheless, he is in a preface state with respect to some works, but not with respect to others. β's preface states will enable us to measure a rather precise upper bound on the threshold value for belief that is implicit in his doxastic states. This bound on the threshold value will apply to all possible degree-of-confidence functions that may model β's doxastic states.

To measure the upper bound on the threshold value for belief, choose a work from β's series and ask β whether he believes that there is an error in this work. If he does not believe there to be an error, then he is not in a preface state with respect to this work. From this nothing follows that permits one to estimate a threshold value for β's beliefs—β may just be so certain that each page is free of errors that no preface state arises. Now, suppose that we can find some work of k pages about which β is in a preface state. This information permits one to begin to model his preface beliefs in terms of a threshold value q, which applies to any implicit degree-of-confidence function that models β's belief states. The Lockean thesis implies that for β's beliefs to be coherent the threshold value q that suffices for belief (in any probabilistic model of β's beliefs) must be no larger than $k/(k + 1)$—i.e., $q \leq k/(k + 1)$—where k is the number of pages involved. (See Appendix 1, Theorem 1, Case 1.)

Suppose we next find that β is also in a preface state with respect to some larger work, a work of m pages where $m > k$. Then his beliefs for this book can be modeled in terms of the same threshold value q as for his k-page work. For, when $q \leq k/(k + 1)$, it follows that $q \leq m/(m + 1)$.

Suppose we now find that β is in a preface state with respect to some smaller work, a work of n pages where $n < k$. This imposes additional constraints on the threshold value for a probabilistic model of β's qualitative doxastic states. All of β's preface states with regard to the works considered thus far can be modeled with a threshold value $q \leq n/(n + 1)$. Thus, the number of pages n in the smallest work we can find for which β is in a preface state provides a least upper bound for the threshold value q for belief. If, for instance, the smallest such work one can find is a 26-page paper, then the threshold value in a probabilistic model of these beliefs must be a value of $q \leq 26/27$ 0.96.

This consistency constraint on threshold values does not presuppose that the errors are independent. For contexts in which β takes errors to be independent, it can be shown that if the smallest work for which β is in a preface state has n pages, then the least upper bound for q is the value of q such that $q + q^n = 1$. (See Appendix 1, Theorem 2.) So, if the smallest work for which β is in a preface state is a 26-page paper and β takes whatever errors may occur there to be independent, then we may infer a threshold value of $q \lesssim 0.91$.

Thus, there is no paradox in the *preface.* For agents whose beliefs satisfy the Lockean thesis a preface state may be perfectly coherent, provided only that a certain constraint is met: the number of pages of the book with respect to which the agent is in a preface state must be sufficiently large in relation to his implicit threshold value for belief.

V

Let us now turn to the *lottery.* Real lotteries come in a variety of forms. Some are designed to guarantee at least one winner. Let us call this an *exhaustive* lottery. Some are designed to permit at most one winner. Let us call this an *exclusive* lottery. And of course some lotteries have both features. Lotteries are usually designed to give each ticket the same chance of winning. Let us call this an *equiprobable* lottery. The design of a lottery may or may not be fully transparent to the agent. When a lottery's design is transparent to an agent, his doxastic state should suitably correspond: if it is exhaustive, the agent is certain that at least one ticket will win; if it is exclusive, the agent is certain that at most one ticket will win; if the lottery is equiprobable, the agent deems it equally plausible that any two tickets will win. We shall see that doxastic states with respect to lotteries can provide additional insight into the connection between beliefs and degrees of confidence.

Let us begin by considering a particularly weak kind of doxastic state. A *weak lottery context* for an agent is a context in which she is certain that the lottery is exclusive (i.e. that there will be at most one winner). An agent in a weak lottery context need not believe that there will be at least one winner, and she need not deem winning to be equally plausible for pairs of tickets.

The following convention will prove convenient in the following discussion: for any proposition *S*, to say that an agent *deems it genuinely possible that S* is just to say that upon considering *S*, her doxastic attitude is that she does not believe $\neg S$.[11] Now, let "W_i" stand for the

[11] In epistemic logic the locution "for all I know, it is possible that *S*" is commonly used for "I do not know that $\neg S$". By analogy we might have used the locution "for all I believe, it is possible that *S*" to stand for "I do not believe that $\neg S$", but this locution is less natural than its epistemic counterpart. Furthermore, in ordinary language we often make a distinction between those possibilities that we think likely enough to be worthy of some consideration and those possibilities that we consider so unlikely as hardly to merit thinking about. The former are often called *genuine possibilities* or *real possibilities*, the latter are usually referred to as *mere possibilities.* Thus, our convention fits well with the vernacular.

proposition that ticket *i* will win and suppose an agent α is in a weak lottery context with respect to some n-ticket lottery—that is, for each pair of tickets *i* and *j*, α is certain that $\neg(W_i \,\&\, W_j)$. We will say that α is in an *m-ticket optimistic state* (where $m \leq n$) just in case:

> for at least *m* tickets, α deems it *genuinely possible* that W_1, α deems it *genuinely possible* that W_2, …, and α deems it *genuinely possible* that W_m.

Just as preface states allow one to determine upper bounds on threshold values in terms of numbers of pages, *m*-ticket optimistic states in weak lottery contexts will allow one to determine lower bounds on threshold values in terms of numbers of tickets.

Consider an agent α who has explicit degrees of confidence with regard to lotteries, and has a threshold value *q* for belief, say, 0.99. It is easy to see how α might come to be in an *m*-ticket optimistic state. She may well be in a weak lottery context for a lottery with few tickets. For instance, in a lottery with three tickets, she might believe that ticket *A* has a 0.40 chance of winning, that ticket *B* has a 0.30 chance of winning, and that ticket *C* has a 0.20 chance of winning, which leaves a 0.10 chance that no ticket 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 $q = 0.99$. 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. Thus, for a sufficiently large lottery her degree of confidence in $\neg W_i$ must come to exceed *q* for at least some tickets *i*.

If α is in a weak lottery context with respect to, say, a one-million ticket lottery, what is the maximum number *m* for which she can be in an *m*-ticket optimistic state? That is, what is the maximum number of tickets such that the agent may deem winning *genuinely possible* for each? To deem winning *genuinely possible* for ticket *i*, α's degree of confidence in W_i must exceed 0.01 (assuming $q = 0.99$); so there can be at most 99 such tickets. Hence, she can be in an *m*-ticket optimistic state for at most 99 tickets in the one-million ticket lottery. Indeed, the precise size of the lottery is irrelevant to this observation. When α is in a 99-ticket optimistic state, her degree of confidence that any of the other tickets may win must be quite close to 0.

For a given belief threshold *q*, how large can *m* be and still permit an *m*-ticket optimistic state in a weak lottery context? The Lockean thesis implies that an *m*-ticket optimistic state can occur only if $m < 1/(1 - q)$. (See Appendix 1, Theorem 3.) This is a strict constraint on the rational coherence of such states. Stated another way, if the Lockean thesis holds for α at threshold level *q*, then rational coherence implies the soundness of the following rule for any value of $m \geq 1/(1 - q)$:

> if for each $i \neq j$, α is certain that $\neg(W_i \,\&\, W_j)$, and
>
> if α does not believe $\neg W_1$, α does not believe $\neg W_2$, … , α does not believe $\neg W_{m-1}$,
>
> then for each $k \geq m$, α believes $\neg W_k$.[12]

However, for each value of $m < 1/(1 - q)$, *m*-ticket optimistic states in weak lottery contexts remain rationally coherent for α.

VI

Suppose that an agent β has no explicit degree-of-confidence function, but β has qualitative doxastic states with respect to various lotteries. If β's beliefs are compatible with some degree-of-confidence function and threshold value, so that the Lockean thesis is satisfiable in principle, then his *m*-ticket optimistic states will enable us to determine lower bounds on a threshold value for a degree-of-confidence function that models his beliefs.

Imagine a collection of lotteries with respect to which β is in a weak lottery context. That is, for each lottery he is certain that no two tickets will win. Choose a lottery and ask β about his beliefs with regard to whether the various tickets may win. If β is not in any *m*-ticket optimistic state for this lottery (for any value of *m*)—that is, if he believes of each ticket that it will not win—then nothing follows from this information that permits one to model a threshold value for belief. But as soon as we find a lottery in which β deems each of *m* tickets to have a *genuine possibility* of winning (for some value of *m*), we gain some information about β's implicit threshold value, *q*, for belief.

Suppose we find β to be in an *m*-ticket optimistic state with regard to some particular lottery. Then, β's beliefs about this lottery can be modeled with a degree-of-confidence function and a threshold level *q* only if the value of $q > (m - 1)/m$. (See Appendix 1, Theorem 3.) So, the number of tickets m for which β can coherently maintain an optimistic state imposes a lower bound on his implicit threshold value *q* for belief. And the largest value of m for which β expresses an *m*-ticket optimistic state provides a greatest lower bound on *q*. For example, suppose we discover a lottery for which β is in a 10-ticket optimistic state—that is, there are 10 tickets in the lottery such that β believes of each that it might win. Furthermore, suppose that among all of the lotteries we have available to test β's beliefs,

[12] The comment in note 8 also applies to this rule.

β is not in an *m*-ticket optimistic state for any $m > 10$. Then the greatest lower bound we have found on β's threshold value *q* is $(m - 1)/m = (10-1)/10 = 0.90 < q$. The total number of tickets *n* in the lottery does not matter at all.

No one has ever suggested that paradox looms for agents in weak lottery states. When a lottery is merely exclusive (i.e. when there is no assurance of a winner), the rational coherence of an agent who deems each ticket to have a *genuine possibility* of winning does not appear to be at risk, regardless of the number of tickets involved. Nevertheless, the Lockean thesis does place strict constraints on the rational coherence of such beliefs: the number of separate tickets that a person deems to have a *genuine possibility* of winning must remain sufficiently small in relation to his implicit threshold for belief.

VII

Our analyses of the *preface* and the *lottery* show how to model qualitative doxastic states in terms of a threshold value for a degree-of-confidence function, and how to determine bounds on the threshold value from the agent's qualitative doxastic states. The *preface* provides a least upper bound and the lottery provides a greatest lower bound. Each provides its respective bound in virtue only of certain formal features. Nothing at all turns on whether the doxastic states involve beliefs about tickets or pages. All that really counts is the numbers of propositions involved in the respective doxastic states.

Nevertheless, the content of *preface* and *lottery* stories helped guide our intuitions to these more significant logical points. Can we exploit the *preface* and the *lottery* still further? If the stories are told in the right way, a preface scenario may be constructed that includes the formal features of the *lottery* and vice versa. Hence one might construct a single preface story or a single lottery story that affords both lower and upper bounds. In the case of the *preface* such a story would look quite contrived. For instance, consider the condition of exclusivity. It is possible to tell some story in which the agent is certain that there is at most one mistake in her book, but such a story would certainly seem odd. However, lotteries that combine the features of the *preface* and the *weak lottery* are quite natural. So, let us consider a lottery scenario which combines the formal features of both, and thus simultaneously affords both a least upper bound and a greatest lower bound.

Define a *strong lottery context* as a context in which an agent α is certain that there is to be at least one winner (i.e. α is certain that the lottery is exhaustive) and certain that there is to be at most one winner (i.e. α is certain that the lottery is exclusive).[13] Suppose that α is in a *strong lottery context* with respect to each lottery in an ensemble of lotteries. For some small lottery with k tickets, α is in a k-ticket *optimistic state*—that is, for each ticket, he deems it *genuinely possible* that it might win. For some large lottery with m tickets α is in a *pessimistic state*—i.e. he believes of each ticket that it will not win. Now suppose we can find a pair of lotteries, one with n tickets and the other with $n + 1$ tickets, such that α is in an n-ticket optimistic state for the n-ticket lottery and in a pessimistic state for the $n + 1$-ticket lottery. These lotteries provide tight lower and upper bounds on the threshold value q for α's beliefs: $(n - 1)/n < q \leq n/(n + 1)$. (See Appendix 1, Theorem 4.) These bounds correspond to the bounds provided by our earlier analyses of the *lottery* and the *preface*, respectively.

Although this procedure yields fairly tight intervals, it is still somewhat crude. It relies only on two qualitative doxastic notions, viz. *belief* and *certainty.* We can home in on the value of q much more precisely by supplementing the former two notions with our third qualitative doxastic notion, viz. *deeming equally plausible*. (Cf. note 5.)

Lotteries are most commonly designed to give each ticket an equal chance of winning. Let a *strong equiplausible lottery context* for an agent be a context in which the agent is certain that the lottery is *exhaustive* and *exclusive,* and deems it equally plausible that any two tickets will win.

It is common for people to buy blocks of lottery tickets. An agent may well believe that he hasn't a prayer of winning with a single ticket, but may deem winning *genuinely possible* for a large block of tickets. Consider a 100-ticket lottery. Suppose that agent β is in a pessimistic state with respect to a single ticket, and for blocks of two, three, and four tickets—he doesn't deem winning to be a *genuine possibility* for such small blocks. But suppose β is in an optimistic state with respect to a block of five or more tickets—he deems it *genuinely possible* that one of the tickets will win. In a *strong equiplausible lottery context* it follows from β's pessimistic state that, for each block of 96 or more tickets, he believes that it contains a winner; and it follows from β's optimistic state that, for each block of 95 or fewer tickets, β does not believe it to contain a winner. Hence, we can conclude that β's threshold value for belief must be greater than 0.95 and no greater than 0.96—that is, $0.95 < q \leq 0.96$. To home in more precisely on q we need only confront β with a larger lottery, say a 1000-ticket lottery, for which he is in a *strong equiplausible lottery context.* β should now be in a pessimistic state with respect to blocks of forty

[13] Here, however, the agent may not presume that tickets have an equal chance of winning.

or fewer tickets, and in an optimistic state with respect to blocks of fifty or more tickets. We probe β's doxastic states for blocks between forty and fifty tickets. If β is in a pessimistic state for a block of 46 tickets and in an optimistic state for a block of 47 tickets, then we can infer that his threshold for belief is $0.953 < q \leq 0.954$. By appealing to larger and larger lotteries, we can approximate the threshold value q to any degree desired. In general, if in a *strong equiplausible lottery context* for an n ticket lottery an agent believes she will not win with only $m - 1$ tickets, but deems it *genuinely possible* that she may win with m tickets, then the agent's threshold value for belief is some number q such that $1 - (m/n) < q \leq 1 - [(m - 1)/n]$.

VIII

According to the Lockean thesis the interconnection between qualitative and quantitative doxastic states applies quite generally. So far, we have explored this relationship strictly within the context of preface and lottery scenarios. Can we extend our observations to contexts involving beliefs about other matters? Clearly, if an agent α explicitly employs a degree-of-confidence function for her doxastic judgments, then she can obtain a coherent notion of belief simpliciter by adopting a threshold value for belief. But the converse situation is less clear. Consider an agent β who has no explicit degree-of-confidence function, but has only qualitative doxastic states. Provided that β satisfies the Lockean thesis at least in principle, can we ascertain a quantitative degree-of-confidence function and a threshold value that model his doxastic states? We will describe a general procedure for doing so.

In order to ascertain degrees of confidence for propositions in general, let us suppose we can find a lottery for which agent β is in a *strong equiplausible lottery context*—that is, β is certain that the lottery is exhaustive and exclusive and deems winning equally plausible for any two tickets in the lottery. In addition, suppose that the lottery is sufficiently large that β believes of each ticket that it will not win. We have already shown that by drawing on a sufficiently large lottery we can home in on a threshold value q by determining a lower bound q_L and an upper bound q_U that form an interval around q as narrow as one may desire. So let us assume that we have already found desirably close upper and lower bounds for β's threshold value: $q_L < q \leq q_U$. Now, for any given proposition S, to ascertain an approximation of β's degree of confidence in S we apply the following procedure.

For each proposition *S*, either β does not believe ¬*S* or else β believes ¬*S*. The procedure will attend to each of these disjuncts in turn. Suppose first that β does not believe ¬*S*—that is, β deems *S* to be *genuinely possible*. Employ any *n*-ticket lottery for which β is in a *strong equiplausible lottery context* and believes ¬W_1 (i.e. β does not deem W_1 to be *genuinely possible*). Since β is certain that one of the tickets will win and deems *S* *genuinely possible*, he also deems it *genuinely possible* that $(S \;\&\; (W_1 \quad \ldots \quad W_n))$. And, since β does not deem it *genuinely possible* that W_1, he does not deem it *genuinely possible* that $(S \;\&\; W_1)$. Hence, there must be some value *k* between 1 and *n* such that β deems it *genuinely possible* that $(S \;\&\; (W_1 \quad \ldots \quad W_k))$ but does not deem it *genuinely possible* that $(S \;\&\; (W_1 \quad \ldots \quad W_{k-1}))$.[14] This is just to say that there is some value *k* between 1 and *n* such that β does not believe $\neg(S \;\&\; (W_1 \quad \ldots \quad W_k))$ and he does believe $\neg(S \;\&\; (W_1 \quad \ldots \quad W_{k-1}))$. It can then be shown that $(1 - q_U)\,(n/k) < \mathrm{P}(S) < (1 - q_L)\,(n/(k-1))$. (See Appendix 1, Theorem 5.) And it also follows that $[1 - (1 - q_U)\,(n/k)] > \mathrm{P}(\neg S) > [1 - (1 - q_L)\,(n/(k-1))]$.

Suppose, alternatively, that β does believe ¬*S*. Then he does not believe that *S*. Now repeat the process described in the previous paragraph with "¬*S*" and "*S*" interchanged.

This procedure may be employed to home in on an agent's implicit degree of confidence in any proposition within her understanding. The narrower the bounds that have previously been determined for *q* and the larger the size *n* of the lottery employed in this procedure, the more tightly bounded will P(*S*) and P(¬*S*) become and the greater will be the precision of one's estimates of the agent's degrees of confidence.

IX

In his defense of the Lockean thesis, Foley suggests that qualitative belief talk is a way of categorizing degrees of confidence in propositions. But Foley is not sanguine about our ability to measure an agent's threshold value for belief:

> There doesn't seem to be any principled way to identify a precise threshold … . Still, we will want to be able to say something, even if vague, about the threshold above which our degrees of confidence in a proposition must rise if we are to believe that proposition. What to say is not obvious, however, since there doesn't

[14] For a large lottery a real agent may be hard put to identify a precise number of tickets, *k*, that marks a boundary between possibilities he deems *genuine* and possibilities he believes will not happen. There may be a region of vagueness. We will discuss the relevance of our analysis of the *preface* and the *lottery* to real agents shortly.

> seem to be a non-arbitrary way to identify even a vague threshold. We deal with other kinds of vagueness by stipulation. Why not do the same here? (p. 112)

It is not quite clear what problem Foley has in mind here. If an agent has explicit degrees of confidence, then *belief* might well be just a matter of her stipulating a confidence level that she finds high enough to be of special significance. If, on the other hand, an agent has no explicit degree-of-confidence function, then there seems to be no need for stipulation at all. He believes what he believes. To the extent that his beliefs are consistent with the Lockean thesis, his preface and lottery beliefs may be employed to measure explicit bounds on his implicit threshold value for belief. So, it turns out, there is no need to resort to stipulation here.

Real human agents are, of course, subject to a certain amount of vagueness with regard to what they believe. A real agent may, for example, hesitate to identify a precise boundary between lottery-ticket block-sizes for which he deems winning *genuinely possible* and block-sizes that he believes will not win. Such vagueness is to be expected in the subjective states of real people. Perhaps Foley is suggesting that stipulation can be of some use here. Ramsey's method for eliciting subjective probabilities and utilities and the von Neumann-Morgenstern method for modeling strength of preference by means of utility functions meet with similar challenges when applied to real people. Stipulation may play a legitimate role in eliminating some vagueness when we employ such models to represent a real agent's doxastic states, provided that the resulting representation does no violence to the unambiguously definite attitudes of the agent we are trying to model. And, when desirable, there are also ways of incorporating some degree of vagueness into models of doxastic states.[15]

In our discussion of the *lottery* and the *preface* we have argued that even an *ideally rational* agent may coherently hold preface-like and lottery-like beliefs. If such beliefs may coherently be held by ideally rational agents, then they may also coherently be held by common folk. Like most other formal models of ideal rationality, our models presuppose that an ideally rational agent possesses a kind of logical omniscience. Our models assume that the ideal agent is certain of logical truths and that logical truths have

[15] Consider, for example, how von Neumann-Morgenstern utilities can be adapted to accommodate vagueness. If an agent prefers *a* to *b* to *c* and is indifferent between *b* and a lottery in which she will get *a* with probability P and *c* with probability (1 – P), then any linear transformation of a function that assigns 1 to *a*, P to *b*, and 0 to *c* is a von Neumann-Morgenstern utility function that models the agent's strength of preference. However, if the agent wavers between *b* and lotteries with probabilities in the neighbourhood of P, this vagueness can be modeled, too, by means of a function which maps to intervals of the reals (e.g. by a function that represents the set of all utility functions that are consistent with the agent's definite preferences).

degree of confidence 1 (see note 6). Real agents, of course, are not *ideally* rational to this extent, and haven't a prayer of becoming so. Real agents cannot possibly come to recognize every logical truth they encounter as logically true. So, this kind of idealization may raise questions as to the relevance of such models to human doxastic states. Foley himself raises this issue (1993; see especially Ch. 5, Sec. 4): he extends his treatment of the Lockean thesis and advances his broader aim of providing a realistic characterization of rational human belief. He does this in part through an investigation of the legitimacy of certain proposed constraints on ideal rationality—for example, that ideally rational belief should be logically consistent, or probabilistically coherent, or should be so configured as not to permit the kind of bets that result in Dutch books. Foley concludes that none of these conditions are requisite for human rational belief, and in this we agree with him fully. However, he also appears to argue that any account of ideal rationality that incorporates logical omniscience can have no relevance whatever to human rationality. We find this challenge to the relevance of such idealizations to be grossly overstated.

If a real person fails to be ideally rational (according to some such account), is she then to be considered *irrational*? Foley would have us believe that the proponents of such accounts think so. But the proper epithet for one who fails to be *ideally* rational is "not *ideally* rational" or "less than *ideally* rational"; it is not "irrational". If Foley's only point were to remind us of this, we wouldn't quibble. But Foley is pushing the further contention that such models of ideal rationality are worthless. If no human agent is even approximately logically omniscient in the sense employed by such models, then what point can there be in explicating models of ideally rational belief that draw on logical omniscience? Foley seems to think there can be none. In that case, do the models of ideal belief we have employed in our analyses of the *preface* and the *lottery* have no bearing on rational human belief?

No one doubts that normal, rational people hold preface-like and lottery-like beliefs. The issue is this. Is there something *logically untoward* about such collections of beliefs, so that a rational person who has such beliefs and comes to recognize their apparent logical incompatibility with one another should, on purely logical grounds, seriously consider altering some of her beliefs? Our idealized models of belief show that the answer is clearly *no*! And in answering this question we show precisely what *logical* constraints a specific, plausible criterion for *rational coherence* imposes on belief. Could the same work have been done without the idealization? Perhaps, but we don't see how. For, if we had tried to do the same work with a model that represented agents as less than *logically ideal*, the question would remain as to whether preface and lottery beliefs represent some sort of logical flaw in the doxastic states of agents, as some philosophers have taken them to be. It seems to us that this issue is best settled by showing that such beliefs would indeed remain coherent even for logically ideal agents.

X

The philosophical value of paradoxes lies in the insights they generate. In this respect we have found the *preface* and the *lottery* remarkably fertile sources of insight into the logic of coherent belief. People use both qualitative and quantitative modes to talk about beliefs, and Locke proposed a straightforward relationship between these two modes. The *preface* and the *lottery* suggest a way to elucidate this relationship with clarity and precision.

On some readings the *preface* may be absorbed into the *lottery*. Be this as it may, the insights to be gleaned from the *preface* cannot be so easily gathered by attending strictly to the *lottery*. Indeed, we have found in these two paradoxes distinct, central features of the logic of coherent belief. It is in virtue of these features that qualitative doxastic notions impose constraints on quantitative ones, and *vice versa*. These results are not limited to preface and lottery contexts. We have shown how it is in principle possible to model the whole body of an agent's qualitative doxastic states in terms of the quantitative notion of degree of confidence. This fleshes out the Lockean thesis and provides the foundation for a logic of belief that is responsive to the logic of degrees of confidence.[16]

Appendix 1

Theorem 1. Suppose there are n propositions $A_1, \ldots, A_n$ such that $P(\neg A_1) \geq q, \ldots$, and $P(\neg A_n) \geq q$:

(1) If $P(A_1 \quad \ldots \quad A_n) \geq q$, then $q \leq n/(n+1)$ [i.e. $q/(1-q) \leq n$].

(2) If $P(A_1 \quad \ldots \quad A_n) = 1$, then $q \leq (n-1)/n$ [i.e. $1/(1-q) \leq n$].

Proof: Suppose the antecedent holds. For each i, $P(A_i) \leq 1 - q$. Observe that $P(A_1 \quad A_2 \quad \ldots \quad A_n) = P(A_1) + P(A_2 \quad \ldots \quad A_n) - P(A_1 \,\&\, (A_2 \quad \ldots$

[16] Thanks to Chris Swoyer and a Mind referee for their very helpful comments and suggestions. This research was supported by the University of Colorado, Boulder, and the University of Oklahoma through their Big 12 Faculty Exchange Program. Luc Bovens's work was also supported by an Alexander von Humboldt Foundation Fellowship.

$A_n)) \leq P(A_1) + P(A_2 \quad \dots \quad A_n) \leq \dots \leq P(A_1) + P(A_2) + \dots + P(A_n) \leq n\,(1 - q)$. In case (1), $q \leq n\,(1 - q)$, i.e. $(n + 1)\,q \leq n$. In case (2), $1 \leq n\,(1 - q)$, i.e. $n\,q \leq n - 1$.

Theorem 2. Suppose there are n propositions $A_1, \dots, A_n$ such that $P(\neg A_1) \geq q, \dots$, and $P(\neg A_n) \geq q$, but also $P(A_1 \quad \dots \quad A_n) \geq q$; and suppose that for each $k \leq n$, $P(A_1 \,\&\, \dots \,\&\, A_{k-1} \,\&\, A_k) = P(A_1 \,\&\, \dots \,\&\, A_{k-1})\, P(A_k)$. Then, for the real number $r > 0$ such that $r + r^n = 1$, $q \leq r$ [i.e. $\log(1 - q)/\log(q) \leq n$].

Proof: Suppose the antecedent holds. From the independence of the *n* propositions, the independence of their negations follows (although it takes some work to show this). Thus, for each $k \leq n$, $P(\neg A_1 \,\&\, \dots \,\&\, \neg A_{k-1} \,\&\, \neg A_k) = P(\neg A_1 \,\&\dots\&\, \neg A_{k-1})\, P(\neg A_k)$. Now observe that $1 - q \geq P(\neg (A_1 \quad \dots \quad A_{n-1} \quad A_n)) = P(\neg A_1 \,\&\, \dots \,\&\, \neg A_{n-1} \,\&\, \neg A_n) = P(\neg A_1 \,\&\, \dots \,\&\, \neg A_{n-1})\, P(\neg A_n) = \dots = P(\neg A_1) \dots P(\neg A_{n-1})\, P(\neg A_n) \geq q^n$. So, $1 \geq q + q^n$; and thus, q is less than or equal to the greatest positive real r such that $r + r^n = 1$. Also, $\log(1 - q) \geq \log(q^n) = n \log(q)$, where $\log(q) < 0$ (the log of any number between 0 and 1 is negative); so $\log(1 - q)/\log(q) \leq n$.

Theorem 3. Suppose there are m propositions $A_1, \dots, A_m$ such that for each pair A_i, A_j $(i \neq j)$, $P(\neg (A_i \,\&\, A_j)) = 1$. If $P(\neg A_1) < q, \dots$, and $P(\neg A_m) < q$, then $(m - 1)/m < q$ [i.e. $m < 1/(1 - q)$].

Proof: Suppose the antecedent holds. So, for all i, $P(A_i) > 1 - q$. Also, (for all $i \neq j$) $P(A_i \,\&\, A_j) = 0$; so, for all k, $P(A_k \,\&\, (A_{k+1} \quad \dots \quad A_n)) = P((A_k \,\&\, A_{k+1}) \quad (A_k \,\&\, A_{k+2}) \quad \dots \quad (A_k \,\&\, A_n)) = 0$. Now, $1 \geq P(A_1 \quad A_2 \quad \dots \quad A_m) = P(A_1) + P(A_2 \quad \dots \quad A_m) - P(A_1 \,\&\, (A_2 \quad \dots \quad A_m)) = P(A_1) + P(A_2 \quad \dots \quad A_m) = \dots = P(A_1) + P(A_2) + \dots + P(A_m) > m\,(1 - q)$. Thus, $1 > m\,(1 - q)$, i.e. $m\,q > m - 1$.

Theorem 4. Suppose there are n propositions $A_1, \dots, A_n$ such that for each pair A_i, A_j $(i \neq j)$, $P(\neg (A_i \,\&\, A_j)) = 1$, and $P(A_1 \quad \dots \quad A_n) = 1$:

(1) if $P(\neg A_1) < q$, and …, and $P(\neg A_n) < q$, then $(n - 1)/n < q$;

(2) if $P(\neg A_1) \geq q$, and …, and $P(\neg A_n) \geq q$, then $q \leq (n - 1)/n$.

Proof: This follows directly from Theorem 1, Case 2, and from Theorem 3.

Lemma for Theorem 5. Suppose there are n propositions $A_1, \dots, A_n$ such that for each pair A_i, A_j $(i \neq j)$, $P(\neg (A_i \,\&\, A_j)) = 1$ and $P(A_i) = P(A_j)$, and also $P(A_1 \quad \dots \quad A_n) = 1$:

(1) for each k, $P(A_1 \quad \dots \quad A_k) = k/n$

(2) if S is any proposition such that for each pair i and j, $P(S \,\&\, A_i) = P(S \,\&\, A_j)$, then for any integer k, $P(S \,\&\, (A_1 \quad \dots \quad A_k)) = P(S)\, k/n$.

Proof: Assume the supposition. For claim (1): $1 = P(A_1 \quad \dots \quad A_n) = P(A_1) + P(A_2) + \dots + P(A_n) = n\, P(A_i)$; so $P(A_i) = 1/n$. Thus, $P(A_1 \quad \dots \quad A_k) = P(A_1) + P(A_2) + \dots + P(A_k) = k/n$.

To prove claim (2), first observe that $P(\neg(A_1 \quad \ldots \quad A_n)) = 0$, so $P(S \,\&\, \neg(A_1 \quad \ldots \quad A_n)) = 0$. Now we have $P(S) = P(S \,\&\, ((A_1 \quad \ldots \quad A_n)$ $\neg(A_1 \quad \ldots \quad A_n)) = P(S \,\&\, (A_1 \quad \ldots \quad A_n)) + P(S \,\&\, \neg(A_1 \quad \ldots \quad A_n)) = P(S \,\&$ $(A_1 \quad \ldots \quad A_n)) = P(S \,\&\, A_1) + \ldots + P(S \,\&\, A_n) = n\, P(S \,\&\, A_i)$. But then, $P(S$ $\&\, (A_1 \quad \ldots \quad A_k)) = P(S \,\&\, A_1) + \ldots + P(S \,\&\, A_k) = k\, P(S \,\&\, A_i) = k\, P(S)/n$.

Theorem 5. Suppose that β satisfies the Lockean thesis in the sense that there is some (unknown) probability function P and (unknown) threshold value q that model β's beliefs—i.e. for each proposition R, β believes R if and only if $P(R) \geq q > {}^1/_2$. And suppose we have already measured upper and lower bounds q_L and q_U on q such that $q_L < q \leq q_U$ (e.g. by employing other lotteries). Given a proposition S such that β does not believe $\neg S$, suppose we can find a *lottery* of the following kind:

For some n propositions $A_1, \ldots, A_n$,

(i) for each pair A_i, A_j $(i \neq j)$, β is certain that $\neg(A_i \,\&\, A_j)$, β deems A_i equally plausible to A_j, β is certain that $(A_1 \quad \ldots \quad A_n)$, and ($n$ is large enough that) β believes $\neg A_1$;

(ii) for each pair i and j, β deems (S & A_i) equally plausible to (S & A_j).

Then for some k, we can determine that β does not believe $\neg(S \,\&\, (A_1 \quad \ldots \quad A_k))$ but also that β does believe $\neg(S \,\&\, (A_1 \quad \ldots \quad A_{k-1}))$. This value of k gives the following bounds on $P(S)$: $(1 - q_U)\, n/k < P(S) < (1 - q_L)\, n/(k-1)$. And the corresponding bounds on $P(\neg S)$ are these: $[1 - (1 - q_U)\, n/k] > P(\neg S) > [1 - (1 - q_L)\, n/(k-1)]$.

Proof: Observe that under the suppositions stated above we have the following: for each pair A_i, A_j $(i \neq j)$, $P(\neg(A_i \,\&\, A_j)) = 1$, $P(A_i) = P(A_j)$, $P(A_1 \quad \ldots \quad A_n) = 1$, and $P(\neg A_1) \geq q$; and also, for any i and j, $P(S \,\&\, A_i) = P(S \,\&\, A_j)$.

(1) First we show that for some k we can determine that β does not believe $\neg(S \,\&\, (A_1 \quad \ldots \quad A_k))$ but also that β does believe $\neg(S \,\&\, (A_1 \quad \ldots \quad A_{k-1}))$.

Notice $P(\neg S) < q$ (since β doesn't believe $\neg S$). Also, for each i, $P(\neg A_i) \geq q$ (since β believes $\neg A_1$ and deems each A_i equally plausible as A_1). So, $P(S \,\&\, A_i) \leq P(A_i) \leq 1 - q < P(S) = P(S \,\&\, (A_1 \quad \ldots \quad A_n)) = P(S \,\&\, A_1) + \ldots + P(S \,\&\, A_n) = n\, P(S \,\&\, A_i)$ (since β deems each (S & A_i) equally plausible). Thus, $P(S \,\&\, A_i) \leq 1 - q < n\, P(S \,\&\, A_i)$.

Then for some k, $P(S \,\&\, (A_1 \quad \ldots \quad A_{k-1})) = P(S \,\&\, A_1) + \ldots + P(S \,\&\, A_{k-1}) \leq 1 - q < P(S \,\&\, A_1) + \ldots + P(S \,\&\, A_k) = P(S \,\&\, (A_1 \quad \ldots \quad A_k))$. It follows that $P(\neg(S \,\&\, (A_1 \quad \ldots \quad A_{k-1}))) \geq q > P(\neg(S \,\&\, (A_1 \quad \ldots \quad A_k)))$; and so β believes $\neg(S \,\&\, (A_1 \quad \ldots \quad A_{k-1}))$but does not believe $\neg(S \,\&\, (A_1 \quad \ldots \quad A_k))$, and β should tell us so when asked.

(2) Now we show how to determine the bounds on P(S) from β's beliefs.

From β's belief that ¬(S & ($A_1 \quad \ldots \quad A_{k-1}$))we may infer that P(¬($S$ & ($A_1 \quad \ldots \quad A_{k-1}$))) ≥ $q > q_L$. So $1 - q_L >$ P(S & ($A_1 \quad \ldots \quad A_{k-1}$)) = P($S$) ($k$ – 1)/n (by the Lemma above).

From β's failure to believe ¬(S & ($A_1 \quad \ldots \quad A_k$)) we may infer that $q_U \geq q >$ P(¬(S & ($A_1 \quad \ldots \quad A_k$))). So $1 - q_U <$ P(S & ($A_1 \quad \ldots \quad A_k$)) = P($S$) k/n (by the Lemma above).

Appendix 2

The supposition that β's beliefs satisfy the Lockean thesis will automatically be satisfied if β's beliefs conform to certain weaker-looking principles of *ideal rationality*. One such set of principles are those of *qualitative probability*. Let us briefly consider this approach. (See Savage 1972 and Krantz et al., 1971 for detailed treatments of qualitative probability.)

Suppose that some agent γ has a qualitative (weak) ordering relation ≥ on the plausibility of propositions. That is, let expressions of form "R ≥ S" say that γ *deems R to be at least as plausible as S*. And suppose that γ is *ideally rational* in the sense that, for all sentences Q, R, S and T: (1) if (Q ≡ R) and (S ≡ T) are logically true and $Q \geq S$, then $R \geq T$; (2) it's not the case that (Q & ¬Q) ≥ ($Q \quad$ ¬Q); (3) R ≥ (Q & ¬Q); (4) $Q \geq R$ or $R \geq Q$; (5) if $Q \geq R$ and $R \geq S$, then $Q \geq S$; (6) if ¬(Q & S) and ¬(R & S) are logically true, then $Q \geq R$ iff ($Q \quad S$) ≥ ($R \quad S$). Any relation that satisfies these six rules is known as a *qualitative probability relation*.

Given the *at least as plausible as* relation ≥ for agent γ, we may define an equivalence relation, ≈, for γ as follows: $R \approx S$ just in case $R \geq S$ and $S \geq R$. Then, $R \approx S$ just says that γ *deems R and S equally plausible*. We may also define the relation *deems more plausible than*, >, for γ: $R > S$ just in case $R \geq S$ and not $S \geq R$.

Let's add a further condition: (7) if $Q > R$, then, for some n, there are n sentences $S_1, \ldots, S_n$, where for each pair ¬(S_i & S_j) is logically true and ($S_1 \quad \ldots \quad S_n$) is logically true, and such that for each S_i, Q > ($R \quad S_i$)—e.g., for each Q and R such that $Q > R$ there is some lottery with so many tickets that disjoining a claim S_i (that only ticket i will win) with R leaves Q > ($R \quad S_i$). A well-known theorem (due to L. J. Savage 1972) establishes that for any relation ≥ satisfying these seven rules, there is a (unique) probability function P that agrees with ≥ in the sense that, for all sentences Q and R, P(Q) ≥ P(R) iff $Q \geq R$.

Now add one more rule: (8) if $Q \geq R$ (i.e. if γ deems *Q* to be at least as plausible as *R*) and γ believes *R*, then γ believes *Q*. This rule guarantees that there is some threshold level q for *belief* such that for any sentence *S*, γ believes *S* iff $P(S) \geq q$. Then we have the following result. Suppose that β's beliefs and non-beliefs are *rationally coherent* in the sense that there exists an *at least as plausible as* relation (for some ideally rational agent γ) which satisfies rules (1)–(8), and that yields precisely the same beliefs and non-beliefs that β has. Then, β's beliefs are guaranteed to satisfy the Lockean thesis. β need not have an explicit degree-of-confidence function, nor even an explicit *at least as plausible as* relation. To satisfy the Lockean thesis it merely suffices that his beliefs are compatible with some such qualitative relation.

An even more elegant treatment would be to write down a set of rules that govern only belief and *certainty* (i.e. rules that do not employ the *at least as plausible as* relation) and then show that whenever beliefs satisfy these rules there must be a corresponding degree-of-confidence function and threshold value that yields precisely the same beliefs and non-beliefs. Such rules on *belief* and *certainty* would provide even weaker looking principles of ideal rationality; yet, any agent whose beliefs satisfied them would automatically satisfy the Lockean thesis. This can indeed be done. The main idea is to use *belief* and *certainty* to weigh off various sentences and disjunctions of sentences against one another, and thereby obtain an ordering on the relative doxastic weightiness of sentences. This ordering will turn out to satisfy rules (1)–(8) for qualitative probability relations, and thus to be representable by some degree-of-confidence function and threshold value for belief.

To see intuitively how this can work, think of a digital scale with only five possible readings: *heavy* (certain that *S*), *weighty* (believes that *S* but not certain that *S*), *medium* (does not believe that *S* and does not believes that $\neg S$), *light* (believes $\neg S$ but not certain that $\neg S$), and *weightless* (certain that $\neg S$). Imagine trying to determine the relative weights of an object, of its parts, of the parts of its parts, etc. Suppose that the object's total weight is *heavy*. One might begin by ordering the relative weights of the *light* parts by seeing which would or would not, when added to other parts, add up to *weighty*. For *light* parts *Q* and *R*, if part *Q* together with part *S* is *weighty* but part *R* together with *S* is not *weighty*, then *Q* is more weighty than *R*. Once the light parts are ordered by weight, one can then use that ordering to determine the relative weightiness of the heavier objects (of which the lighter objects are parts). Here is the strategy for doing all of that with sentences:

When both $\neg Q$ and $\neg R$ are believed, define $Q \geq R$ to mean: for every sentence *S* such that both $\neg(Q \,\&\, S)$ and $\neg(R \,\&\, S)$ are certain, if $(R \quad S)$ is believed, then $(Q \quad S)$ is believed.

When both Q and R are believed, define $Q \geq R$ to mean: for every sentence S such that both $\neg(\neg Q \,\&\, S)$ and $\neg(\neg R \,\&\, S)$ are certain, if $(\neg Q \quad S)$ is believed, then $(\neg R \quad S)$ is believed.

When each of Q, $\neg Q$, R, and $\neg R$ is not believed, define $Q \geq R$ to mean: for each n, if $\neg Q_1, \ldots, \neg Q_n, \neg R_1, \ldots, \neg R_n$ are believed, and (for each i and j) $\neg(Q_j \,\&\, Q_i)$, $\neg(Q \,\&\, Q_i)$, $\neg(R_j \,\&\, R_i)$, $\neg(R \,\&\, R_i)$ are certain and $Q_i \geq R_i$ and $R_i \geq Q_i$, then if $(R \quad R_1 \quad \ldots \quad R_n)$ is believed, then $(Q \quad Q_1 \quad \ldots \quad Q_n)$ is believed.

More generally, for *all* pairs of sentences, define $Q \geq R$ to mean this: Q is certain; or both Q and R are believed and $Q \geq R$; or Q is believed and R is not believed; or each of Q, $\neg Q$, R, and $\neg R$ is not believed and $Q \geq R$; or $\neg Q$ is not believed and $\neg R$ is believed; or both $\neg Q$ and $\neg R$ are believed and $Q \geq R$; or $\neg R$ is certain.

Finally, one can set down axioms for *belief* and *certainty*, and then prove that the relation ≥ as just defined satisfies rules (1)–(8) for qualitative probability relations. These axioms would furnish even weaker looking principles of ideal rationality (i.e. weaker than the probability axioms, and weaker than the rules for qualitative probability). And these axioms, if satisfied by an agent's doxastic states, would automatically guarantee that her beliefs satisfy the Lockean thesis. One way to get such axioms would be simply to replace all occurrences of "≥" in rules (1)–(8) with the definition of "≥" in terms of *belief* and *certainty* given in the previous paragraph. This would automatically provide axioms for *belief* and *certainty* that give rise to qualitative probability relations satisfying rules (1)–(8). However, the resulting axioms would be rather more complicated than is really necessary. They turn out to be derivable from a somewhat simpler and more intuitive set of axioms on *belief* and *certainty*. But the presentation of such axioms and the proofs of appropriate theorems will have to await another occasion.

Department of Philosophy JAMES HAWTHORNE
University of Oklahoma
Norman, OK 73019
USA
hawthorne@ou.edu

Department of Philosophy LUC BOVENS
University of Colorado
Boulder, CO 80309-0232
USA
bovens@spot.colorado.edu

REFERENCES

Foley, Richard 1992: "The Epistemology of Belief and the Epistemology of Degrees of Belief". *American Philosophical Quarterly*, 29, pp. 111–21.

¾¾ 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, pp. 185–218.

Krantz, D. H., Luce, R. D., Suppes, P., Tversky, A. 1972: *Foundations of Measurement*, Vol. 1. San Diego: Academic Press.

Kyburg, Henry E., 1961: *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.

¾¾ 1970: "Conjunctivitis", in Marshall Swain (ed.) *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel.

Locke, John 1690: *An Essay Concerning Human Understanding*, P.H. Nidditch (ed.). Oxford: Clarendon Press, 1975.

Makinson, David C. 1965: "The Paradox of the Preface". *Analysis*, 25, pp. 205–7.

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