

REPORT DOCUMENTATION P.

AD-A255 506

Public reporting burden for this collection of information is estimated to average 1 hour per response, including reviewing the data needed, and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, Office of Management and Budget, Paperwork Project Team, Washington, DC 20503.



ering and
positions
2, and to

2

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE Unknown		3. REPORT TYPE AND DATES COVERED Unknown	
4. TITLE AND SUBTITLE Full Belief				5. FUNDING NUMBERS DAAB10-86-C-0567	
6. AUTHOR(S) Kybury, Henry E., Jr.				DTIC SEP 9 1992	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Rochester Department of Philosophy Rochester, NY 14627					
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army CECOM Signals Warfare Directorate Vint Hill Farms Station Warrenton, VA 22186-5100				8. PERFORMING ORGANIZATION REPORT NUMBER	
				10. SPONSORING/MONITORING AGENCY REPORT NUMBER 92-TRF-0049	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Statement A: Approved for public release; distribution unlimited.					
12b. DISTRIBUTION CODE					
13. ABSTRACT (Maximum 200 words) It may be questioned whether or not there is any such thing as full belief. One epistemological option is to suppose that, for any agent, every statement in his language bears a number between 1 and 0 reflecting the degree of belief the agent has in that statement, and (presumably, for ideal agents) satisfying the axioms of the probability calculus. One prima facie difficulty with this view is that it is conventional to look on changes in the epistemic status of statements as stemming, often if not always, from conditionalization. But conditionalization requires that the evidence on which conditionalization is done be given probability 1: $P'(H) = P(H \& E) / P(E)$. But now $P'(E) = 1$, and on the view being discussed no statement, other than a priori truths, should receive probability 1. However, this is not an insuperable problem: there are ways of representing shifts in probability that do not require that any statement in our language be given probability. But as Diaconis and Zabell showed, every reasonable way of representing probability shifts can also be represented by conditionalization in an enriched language.					
14. SUBJECT TERMS Artificial Intelligence, Data Fusion, Rationality Principle, Acceptance Principle, Rational Belief, Rational Action				15. NUMBER OF PAGES 24	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		

Full Belief*

by Henry E. Kyburg, Jr.

I

A-1

It may be questioned whether or not there is any such thing as full belief. One epistemological option is to suppose that, for any agent, every statement in his language bears a number between 1 and 0 reflecting the degree of belief the agent has in that statement, and (presumably, for ideal agents) satisfying the axioms of the probability calculus. Such was Carnap's view, for example.¹ One prima facie difficulty with this view is that it is conventional to look on changes in the epistemic status of statements as stemming, often if not always, from conditionalization. But conditionalization requires that the evidence on which conditionalization is done be given probability 1: $P'(H) = P(H \& E)/P(E)$. But now $P'(E) = 1$, and on the view being discussed no statement, other than a priori truths, should receive probability 1. As Richard Jeffrey showed,² however, this is not an insuperable problem: there are ways of representing shifts in probability that do not require that any statement in our language be given probability 1. (But as Diaconis and Zabell showed,³ every reasonable way of representing probability shifts can also be represented by conditionalization in an enriched language.)

Formally, that is, we can go either way. Are there other reasons for dealing with full belief? There are certainly suggestive considerations. ^① That's the way we talk, for one thing: we talk of accepting evidence, of when the evidence warrants the acceptance of a conjecture, and so on. To accept a statement would appear to be to accord it full belief, to assign it



probability 1. Furthermore, however one interprets "probability" (so long as it applies to single statements) it is possible that the probability of S is high at t₁, and that S has very low probability at t₂. If the evidence is different at the two times, it need not be the case that the agent has made any mistake at either time. In fact, it seems that without acceptance there can be no grounds -- other than computational ones -- for saying "Ooops! I was wrong."

In epistemology in general, it has been traditional to suppose that warranted full belief is possible. It may be defeasible (which is to say, corrigible), but this is not generally construed as a matter of probability. It is the statement S that is corrigible, and might be deleted from one's beliefs if circumstances warrant, and might even be replaced by not-S, and not the statement probably-S.

But there are problems with this approach, too. If we say that a statement may be accorded full belief if it is probable enough, we must face the question of how probable that is. Even if high probability is only a necessary and not a sufficient condition for full belief, one may ask how high that probability must be. Furthermore, since we are speaking of probabilities relative to evidence, what are we to take as evidence? How probable must that be? Or if we are talking of "good reasons" rather than high probabilities, how good must they be? The same kinds of questions can be raised about good reasons as can be raised about probabilities. Finally -- and perhaps this is the most important question of all -- how do we unaccept statements? One of the pervasive features of full belief -- if there is such a thing -- is its non-monotonicity: given a certain evidential background, we can accept S; given an expanded background -- more evidence -- we can no longer accept S. But if there is just one set of "accepted"

statements functioning as evidence, any statement that has once been accepted has probability 1 relative to the set of accepted statements, and there is no way in which augmenting that set of statements (with new evidence) can decrease that probability. Similarly, what could be a better reason for \underline{S} than \underline{S} ?

Isaac Levi is one of the few epistemologists who take these questions seriously. Acceptance is characterized in terms of expected epistemic utility, together with a parameter q representing an index of caution. The grounds for the contraction of a body of accepted statements are less clear, though one of them is the presence of a contradiction in that body, and that seems clear enough.

What follows is tailored to my own epistemological approach, but it may have a bearing on other approaches in which full belief or acceptance play a role. It might, for example, have a bearing on the selection of Levi's index of caution q in a given scientific context.

On my view, we have not one, but two bodies of accepted statements: a set of statements accepted at an evidential level, and a set of statements accepted as practical certainties. The latter are accepted in virtue of the high probability they have relative to the former. To discuss the grounds for accepting statements in the evidential corpus is exactly to shift context so that that set of statements is regarded as a set of practical certainties, and some other, less dubitable, statements serve as an evidential corpus. There are levels characterizing each of these corpora in a given context. The question is, what principles can we employ to select those levels?

II

A useful way to approach this question is through the analysis of full belief, or acceptance. What is it to accept a statement? Partial belief is characterized (at least by Bayesians) in terms of a propensity to make or to fail to make bets. Given a range of dollars that satisfies the elementary axioms of utility, we can say that I have a degree of belief of degree one half in the statement that the next toss of this coin will land heads in virtue of my propensity to bet at even money on that event. Of course amounts of dollars only approximately satisfy the axioms of utility; and of course I may be indifferent between accepting and rejecting a bet over a non-zero range of odds. But these are just the classical difficulties of measurement.

Can we give an analysis of full belief along the lines of the proposed analysis of partial belief? Is to have full belief in S to be willing to pay a unit of utility for a ticket that will pay a unit of utility if S is true? Or to be indifferent between a ticket that pays a prize P in any case and one that pays a prize P only if S is true? If it is logically possible that S should be false, one would think the natural response would be: Why take the chance?

There is a fair amount of empirical data to the effect that when dealing with events whose relative frequencies are close to 0 or to 1, people's choices between alternative gambles no longer seem to fit the general pattern of maximizing expected subjective utility.⁴ The deviations suggest that in the experimental situation events with frequencies that are very close to 1 are simply assumed to occur, and events with relative frequencies very close to 0 are simply assumed not to occur. This leads to conflicts with the Bayesian model of deliberation, but that is not my

concern here.

What I am concerned with is the fact that if "full belief" is to be considered a limiting case of "partial belief," then we should suppose that partial beliefs progress in an orderly way toward this limit. Empirically, this does not seem to be the case. To be sure, our object is a theory of rational belief, rather than an empirical descriptive theory of belief, and X we might say that people tend to be irrational about events with extremely high or extremely low relative frequencies. But we should not with unseemly haste conclude that our fellow men uniformly leap into irrationality under the same uniform circumstances. It may be that there is an intuitively acceptable notion of rationality within which this general tendency can be accounted for, and perhaps even regarded as rational.

R. B. Braithwaite⁵ offers an alternative and at first sight more useful dispositional interpretation of full belief. It is (roughly) that a person has full belief in S if he acts as if S were true. The difficulty is that whether or not a person "acts as if S were true" depends on what is at stake. A man may act as if a certain vaccine is nontoxic when it comes to vaccinating monkeys, but not when it comes to vaccinating children.⁶ Braithwaite requires that full belief in S be represented as a disposition to act as if S were true under any circumstances to which the truth of S is relevant. But as Levi and Morgenbesser have pointed out in an unjustly neglected article, "for any contingent proposition p on which action can be taken, there is at least one objective relative to which a nonsuicidal, rational agent would refuse to act as if p were true. Consider, for example, the following gamble on the truth of p: If the agent bets on p and p is true, he wins some paltry prize, and if p is false he forfeits his life. However, if he bets on not-p, he stands to win or lose some minor stake. ... Hence ... the agent could not rationally and sincerely believe

that p where p is contingent."⁷

Levi and Morgengesser go on to consider more complicated reconstructions, in which actions are taken to depend on circumstances, motives, and stimulus, as well as beliefs, and show that what is involved is more like a promissory book than a promissory note. My concern here, however, is not with the general problem of construing beliefs dispositionally - a very knotty problem indeed - but with the far simpler partial problem of making sense of full belief presupposing an understanding of such parameters as motives, stimuli, external circumstances, and the like. In particular we shall presuppose, what is itself transparently (I should say, rather, "opaquely") dispositional: a cardinal utility function for the agent over states of the world.

Let us begin by looking more closely at the matter of "acting as if" S were true. If I bet at even money on heads on the fall of a coin, it might appear as if I were acting as if "The coin will fall heads," were true. The appearance is misleading. My action is that of making a bet, and making a bet is not acting as if the proposition that is the subject of the bet were true. On the contrary, it is acting as if a certain relative frequency or propensity characterized the kind of event at issue. Thus betting at even money on heads is acting as if at least half the tosses of coins yielded heads. This much, at least, seems relatively straightforward, though in detail we would have to take account of my aversion to risk, or my pleasure in excitement. In fact, we can no doubt characterize my betting behavior concerning coin tosses in general by saying that in ordinary circumstances I act as if the relative frequency of propensity of heads were a half, or, more precisely, as if the distribution of heads among coin tosses were binomial with $p = 1/2$.

But this claim concerning the relative frequency or propensity of heads is surely a contingent claim. Would you risk your life in exchange for a paltry prize on the proposition that the propensity of this coin to yield heads was a half? Or even "close" to a half? Hardly. But we can easily enough imagine bets and stakes concerning the long-run behavior of the coin that would strike us as reasonable. (I'll give you odds of ten to one that if we flip the coin until it wears out, the relative frequency of heads will end up between .48 and .52.) Now, of course, I am "acting as if" almost all coins and flippings are symmetrical enough to yield that result.

On the other hand, let us suppose that I am willing to act as if "half the tosses of coins yield heads" is true - i.e., to bet at even money on heads on a toss of a coin. From "half the tosses of coins yield heads" it follows that (let us suppose) 1/1,346,451 of the sets of 10,000 tosses fail to yield a relative frequency of heads between 4,000 and 6,000.⁸ (The exact numbers are irrelevant.) Then should I not be willing to offer or receive odds of \$1 to \$1,346,452 against this outcome in a particular case? It seems to me that I should not.

One might think that this is a reflection of the size of the stake involved - a million seems large. This feeling is reinforced by the example of the possibly toxic vaccine: so much more is at stake in inoculating children than in inoculating monkeys that we should have much more evidence that the vaccine is nontoxic before we inoculate children than before we inoculate monkeys. But I think this intuition is mistaken, and that we are being misled by our background knowledge of the frequencies of disease and the function of vaccination. Suppose that it is not a new vaccine at issue, but a new antibiotic, which is demonstrably effective against disease D in pigs. Humans are also subject to D (though only rarely); but whereas pigs rarely fail to recover from this disease, humans almost always succumb to

it. We have some evidence regarding the toxicity of this drug in both humans and pigs; to fix our ideas, suppose we can say that the probability that it is nontoxic to pigs and the probability that it is nontoxic to humans is about the same - say about .8. It seems quite clear that one would give the drug to a group of children known to have D, but that one would not give it to a group of pigs known to have D. That is, with respect to the children, one would act as if the drug were nontoxic, despite the high stakes involved, but with respect to the pigs, one would not act as if the drug were nontoxic.

This and similar examples suggest that what is crucial in whether or not a person "acts as if S were true" is not the total magnitude of the stake involved, but the ratio of the amount risked to the amount gained. Of course this is exactly what is of central concern in decision theory. But there we consider the stakes and the probabilities fixed and ask for a decision. Here I suggest that we keep the decision and the stakes fixed, and look for limits on the probabilities. For example, suppose the maximum ratio of the stakes is 99:1 -- that is we do not consider circumstances in which more than 99 units are risked for the possible gain of one unit or vice versa. Clearly a probability larger than .99 is no different than certainty as far as my behavior and decisions go.

How is this? Well, the highest odds I can be offered to bet against a statement are 99 to 1. But for this to ^{have positive expectation} ~~be profitable~~ the value of $99(1-p) - p$ must be positive. But if $p > .99$ this can never happen -- I can't bet against S at any odds in the specified range and enjoy positive expectation. The argument that I cannot reasonably bet on S's negation is parallel. In short, if the minimum probability of S is p , it will never pay to bet against S so long as $p > \frac{r}{r+s}$, where $r:s$ ^{are} ~~the~~ the maximum odds allowable.

Let us take this as our basic idea: A fully believes S in the sense of the ratio (r/s) just in case in any situation in which the ratio of risk to reward is less than r:s, A simply acts as if S were true. *... refuses to bet against S.* Symmetrically, A will then also fully disbelieve $\neg \underline{S}$ - i.e., in any situation in which the ratio of the risk of counting on $\neg \underline{S}$ to the benefit if $\neg \underline{S}$ is false is greater than s:r, A will simply act as if $\neg \underline{S}$ were false.

An example will help to make this clear. Suppose I fully (99/1) believe that my car will start this afternoon, but I don't fully (r/s) believe it for r/s > 99/1. Suppose act a is an act that will cost me \$50 if my car doesn't start, but otherwise will yield a benefit worth \$1. Since 50/1 < 99/1, I will perform act a -- i.e. act as if I knew that the car would start.

On the other hand if I contemplate an act that will return the same dollar if my car does start, but will cost me \$200 if it doesn't, I may or may not perform the act according to the probabilities involved. In particular, if $-(1-p)(200) + p(1) > 0$, i.e., if $p > \frac{199}{200}$, I will perform the act, but if $p < \frac{199}{200}$ I won't. *I do not fully (199/1) believe that my car will start.*

III

Let us try to formalize full r/s belief. The idea behind our formalization is that there are some circumstances under which A will believe a sentence S, but there will also be circumstances under which A will not believe S, and not believe the negation of S. It is only under those circumstances that A can take either side of a bet on S. If not too much is at stake, A will believe (accept) S, and otherwise A will not accept S. The following principle captures this idea:

(D1) A fully (r/s) believes S if and only if for every action that A takes

to cost $\underline{r}' > 0$ if \underline{S} is false and to yield $\underline{s}' > 0$ if \underline{S} is true, where $\underline{r}'/\underline{s}' < \underline{r}/\underline{s}$, \underline{A} acts as if \underline{S} were true.

A number of consequences of B1 follow:

- (T1) For any action that \underline{A} takes to cost \underline{s}' if $\neg \underline{S}$ is false, and to yield \underline{r}' if $\neg \underline{S}$ is true, where $\underline{s}'/\underline{r}' > \underline{s}/\underline{r}$, \underline{A} acts as if $\neg \underline{S}$ were false.

This is the other half of the intuitive characterization of acceptance in the previous section. To see that it follows from D1 note that to cost \underline{s}' is to yield $-\underline{s}'$, and to yield \underline{r}' is to cost $-\underline{r}'$. For $\neg \underline{S}$ to be false is for \underline{S} to be true. D1 then says that if $-\underline{r}'/-\underline{s}' < \underline{r}/\underline{s}$, \underline{A} should act as if \underline{S} were true. But $-\underline{r}'/-\underline{s}' < \underline{r}/\underline{s}$ if and only if $\underline{s}'/\underline{r}' > \underline{s}/\underline{r}$.

- (T2) \underline{A} fully $(\underline{r}/\underline{s})$ believes \underline{S} if and only if there is a $p \in [0,1]$ such that \underline{A} fully $(p/(1-p))$ believes \underline{S} .

This is so because, if $p = \underline{r}/(\underline{r}+\underline{s})$, then $\underline{r}'/\underline{s}' < \underline{r}/\underline{s}$ if and only if $\underline{r}'/\underline{s}' < p/(1-p)$; p is the probability of \underline{S} in classical Bayesian terms, and we are contemplating no odds that would justify a bet against \underline{S} .

- (T3) If \underline{A} fully $(\underline{r}/\underline{s})$ believes \underline{S} and $0 < \underline{r}'/\underline{s}' < \underline{r}/\underline{s}$, then \underline{A} fully $(\underline{r}'/\underline{s}')$ believes \underline{S} .

It seems strange to say that \underline{A} can act both as if \underline{S} were true and as if $\neg \underline{S}$ were true, though of course \underline{A} can simultaneously bet on both \underline{S} and $\neg \underline{S}$ -- indeed, this is just how a bookmaker makes his living. If we speak of full $(\underline{r}/\underline{s})$ belief where $\underline{r}/\underline{s} < 1$, however, we get exactly this result. For

example, I fully $(1/4)$ believe that the next toss of this coin will yield heads and I fully $(1/4)$ believe that the next toss of this coin will not yield heads.

To avoid this peculiarity, let us write A fully $^!(\underline{r}/\underline{s})$ believes \underline{S} to mean that A fully $(\underline{r}/\underline{s})$ believes \underline{S} , and for no $(\underline{r}'/\underline{s}') > (\underline{r}/\underline{s})$ does A fully $(\underline{r}'/\underline{s}')$ believe \underline{S} , and stipulate that $(\underline{r}/\underline{s}) > 1$.

(D1') A fully $^!(\underline{r}/\underline{s})$ believes \underline{S} if and only if A fully $(\underline{r}/\underline{s})$ believes \underline{S} , $(\underline{r}/\underline{s}) > 1$, and if $(\underline{r}'/\underline{s}') > (\underline{r}/\underline{s})$ then A does not $(\underline{r}'/\underline{s}')$ believe \underline{S} .

In this sense of full belief, A cannot both fully believe \underline{S} and fully believe $\sim \underline{S}$:

(T4) If A fully $^!(\underline{r}/\underline{s})$ believes \underline{S} then A does not fully $^!(\underline{r}/\underline{s})$ believe $\sim \underline{S}$.

A certain amount of deductive closure is already implied for rational A ; if A fully $^!(\underline{r}/\underline{s})$ believes \underline{S} and $\underline{S} \vdash \underline{S}'$, then A fully $(\underline{r}/\underline{s})$ believes \underline{S}' :

(T5) If $\underline{S} \vdash \underline{S}'$, then if A fully $^!(\underline{r}/\underline{s})$ believes \underline{S} , A fully $(\underline{r}/\underline{s})$ believes \underline{S}' .

T5 holds because if $\underline{S} \vdash \underline{S}'$, then to act as if \underline{S} is true is to act as if $\underline{S} \ \& \ \underline{S}'$ is true. But if rational A is willing to act as if $\underline{S} \ \& \ \underline{S}'$ is true when risk to reward is less than $\underline{r}/\underline{s}$, then he should also be willing to act as if \underline{S}' is true under the same circumstances.

Clearly, it need not be the case that if A fully $^!(\underline{r}/\underline{s})$ believes \underline{S} and fully $^!(\underline{r}/\underline{s})$ believes \underline{S}' , then A fully $^!(\underline{r}/\underline{s})$ believes $\underline{S} \ \& \ \underline{S}'$. The ratio of stakes involved in acting as if both \underline{S} and \underline{S}' were true may no longer be less than $\underline{r}/\underline{s}$. Since $\underline{r}'/\underline{s}' < \underline{r}/\underline{s}$ and $\underline{r}''/\underline{s}'' < \underline{r}/\underline{s}$ imply $(\underline{r}' + \underline{r}')/(\underline{s}' + \underline{s}'') < \underline{r}/\underline{s}$, one might think not. But this is an irrelevance, as may be seen from this example: One should fully $^!(5/1)$ believe that the next roll

of a die will not yield a one, and fully[!](5/1) believe that the next roll of a die will not yield a two, but no more than fully[!](2/1) believe that the next roll will yield neither a one nor a two.

Suppose that \underline{P} is a classical epistemic probability function defined relative to what \underline{A} fully[!]($\underline{r}/\underline{s}$) believes. I am supposing here that probability is a logical relation of the sort I have described elsewhere. By calling it a logical relation, I mean that the probability of \underline{S} relative to any set of other statements is logically determinate (it won't be bothersome if there are a few peculiar circumstances under which it is not defined), and has a certain value regardless of what anyone whose body of full beliefs corresponds to that set of evidence statements does believe or would believe. The following rationality principle suggests itself.

(RP) If $\underline{P}(\underline{S}) = \underline{p} < \underline{r}/(\underline{r}+\underline{s}) = \underline{p}'$ and \underline{Bd} is the act of paying \underline{d} for a unit of return contingent on \underline{S} , then \underline{A} will rationally perform \underline{Bd} if $\underline{d} < \underline{p}$; \underline{A} will rationally refrain from \underline{Bd} if $\underline{p} < \underline{d}$, and if $\underline{p} = \underline{d}$, then the rationality of \underline{A} cannot be faulted whether or not he performs \underline{Bd} .

So far as it goes, this principle conforms to the principle of maximizing expected utility. But now suppose that $\underline{p} \geq \underline{p}' = \underline{r}/(\underline{r}+\underline{s})$. Then according to the generally accepted scheme, the previous analysis should still apply: if $\underline{p} > \underline{d}$, it is worth buying a ticket for \underline{d} that will return a unit if \underline{S} is true, even if $\underline{d} > \underline{p}'$. But this is equivalent to staking a possible gain of $1-\underline{d}$ against a possible loss of \underline{d} on the truth of \underline{S} . This is not in the range of stakes contemplated in the sense of "full[!]($\underline{r}/\underline{s}$) belief," unless $\underline{d}/(1-\underline{d}) < \underline{r}/\underline{s} = \underline{p}'/(1-\underline{p}')$. Since $\underline{d} > \underline{p}'$, $\underline{d}/(1-\underline{d}) < \underline{r}/\underline{s}$ is impossible and the ratio of stakes in \underline{Bd} is not among those contemplated.

Similar considerations govern a bet against \underline{S} . Given the range of stakes we are contemplating -- $(\underline{r}/\underline{s})$ to $(\underline{s}/\underline{r})$, probabilities between 0 and $(\underline{s}/(\underline{r}+\underline{s}))$ and between $(\underline{r}/(\underline{r}+\underline{s}))$ and 1 serve no function.

Thus the conventional Bayesian wisdom breaks down. If we are interpreting full belief as full[!] $(\underline{r}/\underline{s})$ belief, then statements having probabilities greater than $p' = \underline{r}/(\underline{r}+\underline{s})$ or smaller than $1-p' = \underline{s}/(\underline{r}+\underline{s})$ are not fit subjects for bets.

But when we take probability to be interval valued, we face a gap in our theory. A statement \underline{S} may neither qualify for full[!] $(\underline{r}/\underline{s})$ belief nor for full[!] $(\underline{r}/\underline{s})$ disbelief, nor be such that the Bayesian principle of maximizing expected utility always gives us guidance for actions depending on \underline{S} . If the probability of \underline{S} relative to background knowledge \underline{K} is the interval $[p, q]$, i.e., $\text{Prob}(\underline{S}, \underline{K}) = [p, q]$, and $p < \frac{\underline{r}}{\underline{r}+\underline{s}}$ then \underline{S} does not qualify for full[!] $(\underline{r}/\underline{s})$ belief. Similarly, if $1-q < \underline{r}/(\underline{r}+\underline{s})$, then $\sim \underline{S}$ does not qualify for full[!] $(\underline{r}/\underline{s})$ belief either. But now \underline{Bd} for any $\underline{d} \in [p, q]$ cannot be faulted, though to pay \underline{d} for a unit return in case \underline{S} and at the same time \underline{d}' for a unit return in case $\sim \underline{S}$ is not reasonable when $\underline{d} + \underline{d}' > 1$.

A more serious difficulty emerges when we ask what the body of knowledge \underline{K} is that we compute probabilities relative to. Suppose it is the set of full[!] $(\underline{u}/\underline{v})$ beliefs for some $\underline{u}/\underline{v} < \underline{r}/\underline{s}$. This would be to use as evidence for the full beliefs appropriate to a wider range of risks and benefits, full beliefs that may barely be appropriate to a narrower range of risks and benefits. That clearly seems inappropriate. In fact, looked at from the point of view of the set \underline{K} of full[!] $(\underline{u}/\underline{v})$ beliefs, no probability greater than $\frac{\underline{u}}{\underline{u}+\underline{v}}$ or smaller than $\frac{\underline{v}}{\underline{u}+\underline{v}}$ seems to make any sense at all; at any rate it is inappropriate for computing expected utilities.

Could we use the set of full[!] $(\underline{r}/\underline{s})$ beliefs itself? No. It seems circular to take the corpus of knowledge, relative to which we determine

what may be included in that corpus of knowledge to be the corpus itself.

So we take the corpus \underline{K} to comprise full! ($\underline{u}/\underline{v}$) beliefs for some $\underline{u}/\underline{v} > \underline{r}/\underline{s}$. But which one?

Let us write the index $\underline{r}/\underline{s}$ in the normalized form $p'/(1-p')$, or, better yet, represent the range of stakes contemplated by a single number $p' = \underline{r}/(\underline{r}+\underline{s})$. The number p' corresponds naturally to the level of our corpus of practical certainties. We have said that a statement should get into the corpus of practical certainties by being probable enough -- having a lower probability at least equal to p' -- relative to the evidential corpus. This corresponds naturally to being worthy of full! ($\underline{r}/\underline{s}$) belief. The following argument suggests (but it doesn't quite demand) that the relation between the indices of the evidential and the practical corpus be that the latter is the square of the index of the former. For suppose that \underline{S}_1 is in the evidential corpus. Then for any sentence \underline{S} , $\underline{S} \equiv \underline{S} \ \& \ \underline{S}_1$ will be in the evidential corpus, in virtue of the theoremhood of $\underline{S}_1 \rightarrow (\underline{S} \equiv \underline{S} \ \& \ \underline{S}_1)$. Thus if \underline{S} is in the practical corpus, so will $\underline{S} \ \& \ \underline{S}_1$ be in the practical corpus. What this means in general is that if \underline{S}_1 and \underline{S}_2 are in the evidential corpus, their conjunction will be in the practical corpus. Since this fact depends on no relation between \underline{S}_1 and \underline{S}_2 , it will hold for \underline{S}_1 and \underline{S}_2 that are independent stochastic events whose probabilities are exactly p . Their conjunction will have the probability p^2 relative to the same corpus that gives them each a probability p , and so, since their conjunction is practically certain, that suggests that p' should equal p^2 .

All that is required is to make full! ($p'/(1-p')$) belief depend on having a probability of at least p' relative to the corpus of knowledge whose index is $(p')^{1/2}$. Let us also allow for uncertain perceptual judgment as well, by including \underline{S} among the full! ($\underline{r}/\underline{s}$) beliefs if it represents a

judgment based on observation and is the sort whose error rate is less than $\frac{s}{r+s}$. More precisely, we adopt a second principle of rationality -- a principle of acceptance. (We now write "full!(p)" for "full!(p/(1-p))".)

(AP) A should give full!(p) belief S if and only if the probability of S, relative to A's full (p)^{1/2} beliefs, is greater than p, or S is obtained by observation, and relative to A's full (p) meta beliefs the probability that S is in error is less than $1 - p$.⁹

Call a sentence S anomalous for full!(p'/1-p) p' belief if it is neither a fit subject for bets in accordance with principle RP, nor a fit subject for full!(p'/(1-p')) belief in accordance with D1, either. If S is anomalous for full!p' belief, then it may be acceptable as worthy of full!q' belief, where $q' < p'$. And at some higher level, $r' > p'$, it may be a fit subject for Bayesian guidance.

The principle AP leads immediately to the result that by shifting the ratio $\frac{r}{s}$ slightly, we can always resolve the anomalies in at least one way:

T6 If S is anomalous relative to A's full!(p) beliefs - i.e., if its lower probability is greater than p, but it is not a member of A's full!(p) beliefs - then A should give full.(p') belief to S where $\frac{1}{2} < p' < p$.

proof: Construing p as the evidential level, S is a member of the corresponding practical corpus -- i.e. A's set of full!(p²) beliefs.

It is also possible to resolve anomalies the other way -- i.e., to move to a level p' such that the Bayesian maxim does apply.

T7 If S is anomalous relative to A's full!(p) beliefs, then there exists a $p' > p$ such that relative to A's full!(p') beliefs, the lower

probability of \underline{S} is less than p' , and therefore the Bayesian maxim is appropriate for bets on \underline{S} relative to \underline{A} 's full! (p') beliefs.

Suppose there is no such p' . Then, relative to every set of \underline{A} 's full! (p') beliefs, where $p' > p$, \underline{S} has lower probability greater than or equal to p' . In particular this applies to \underline{A} 's set of full! (1) beliefs -- i.e., \underline{A} 's beliefs concerning mathematical and logical truths. But then there is a p' -- namely 1 -- such that the lower probability of \underline{S} relative to \underline{A} 's full! (p') beliefs is greater than p , where p' is greater than p . But this contradicts the assumption that \underline{S} not be fully! (p) believed by \underline{A} , since \underline{S} will then be inherited by every corpus of lower level.

IV

Let us see how this framework can be taken to impose constraints on rational belief and rational degrees of belief. One of the relevant factors in any dispositional analysis of belief, however we go from there, is the set of circumstances in which the agent finds himself. But although this is a relevant factor, it is not one that should lead to despair; we do not want to say that we can provide a dispositional analysis of rational belief only if we know the circumstances of the agent in infinite detail. Indeed, this would preclude our being able to assess the rationality of others or to improve the rationality of ourselves, for we can never articulate our circumstances in infinite detail. What is required is that we be able to characterize a broad class of circumstances under which our analysis is to apply.

Within the framework suggested here, this class of circumstances is characterized precisely by the range of ratios of stakes that the agent has

(perhaps implicitly) in mind. It is easy enough to alter the circumstances so that they lie outside this range - at least hypothetically. This is precisely what Levi and Morgenbesser are doing when they point out that, whatever the evidence for S, there are circumstances under which the agent A will not act as if S were true - e.g., when his honor is on the line against a paltry prize.¹⁰ But this is precisely because this represents a ratio of stakes outside the range implicitly and initially contemplated by A. Often this shift can be accomplished relatively easily and realistically by means of the simple query: "Wanna bet?"

According to the framework, we could consider a notion of full belief in which the ratio of stakes was very close to unity - full!($1/2 + \epsilon$) belief. This is just what is done, sometimes, by epistemologists who require of S merely that it be "more probable than not" in order to be worthy of belief. But this is not a very interesting sense of "full belief." Ordinarily we want our beliefs to remain fixed through a relatively wide range of circumstances - i.e., to be suitable for a relatively wide range of ratios of stakes. A range from 10:1 to 1:10 might seem more plausible.

Now the actual stakes, in any circumstance, depend on the agent's utility function: therefore so also does their range. The set of circumstances relative to which our analysis is to be performed should thus be represented by a function of both the Agent A's utility function (a function of A himself) and a ratio of stakes: $C(A, r/s)$.

Consider full belief first. We say that A has full!(r/s) belief in S, just in case for any circumstance $c \in C(A, r/s)$ A acts as if S were true. Note that for A to act as if S were true is not for A to perform any particular action (in any ordinary sense of 'particular') as Levi and Morgenbesser seem to suggest.¹¹ This is not the place to attempt a

characterization of action, but it nevertheless seems clear that there is a certain class of "deliberate behaviors" that are ruled out by A's "acting as if S were true," (for example, betting against S) and a certain class that are required. It seems to me that this is all that is needed to give content to the notion.

A's full!(r/s) belief in S is rational, according to the framework principle AP, just in case the probability of S, relative to A's full!(p) set of beliefs, is at least (p') where $p^2 = p'$, and A's full!(p) beliefs are themselves rational in turn. This raises a problem to which we shall return shortly.

Now let us consider A's partial belief in S. If A's degree of belief in S is to be characterized by a real number q, it will be exactly that number q such that if A is compelled to take one side or the other of a bet at odds of q:1-q on S, he will be indifferent as to which side he takes. But this seems unnatural, and foreign to the notion of the set of ordinary circumstances $C(A, r/s)$; we should seek a gentler characterization of degree of partial belief. We can get at this by supposing that there is a range of ratios, say from q:1-q to q':1-q', such that A would be indifferent about taking either side of the bet. (Remember that the bets are in utilities; so A's enthusiasms for gambling and his reluctance to take chances are already taken into account.) A's partial degree of belief then comes to be characterized by the interval [q, q']. Put another way, if he is offered the opportunity to buy a ticket for (1-q') units of utility that will return a unit of utility if S fails to occur, he will take it; and if he is offered the opportunity to buy a ticket for q units of utility that will return a unit of utility if S does occur, he will take it.

Now when is A's partial belief in S rational? Clearly when, relative

to A's body of full!(p') rational beliefs, the probability of S is $[q, q']$. This means that, under any circumstances $c \in C(A, r/s)$, A ought to bet on S at odds lower than $q:(1-q)$, and ought to bet against S at odds lower than $(1-q'):q'$. Note, however, that, if the odds do not lie in the range $[r:s, s:r]$, then A is not in the circumstances $C(A, r/s)$ envisaged by the analysis. Suppose that A is offered a bet on S at odds between $q:1-q$ and $q':1-q'$. Then he is under no rational constraint either to accept or to reject the bet. But this is all right. We still have perfectly good characterizations of A's partial (r/s) belief and of the constraints that this belief must satisfy in order to be rational.

We can restate the rationality principle (RP) to take explicit account of the circumstances of the agent A:

(RP') For any $c \in C(A, r/s)$, if A is offered a bet at stakes r'/s' on S or s'/r' against S, where $r/s > r'/s'$,

(1) He will (ought to) accept the bet on S if r'/s' is less than $q:1-q$, where q is the lower bound of his (r/s) degree of (rational) belief in S;

(2) He will (ought to) accept the bet against S if r'/s' is greater than $q':1-q'$, where q' is the upper bound of his (r/s) degree of rational belief in S; and

(3) He may or may not accept either bet otherwise, where 'may' has its customary English ambiguity.

There are several possible difficulties with this framework. Suppose that, relative to A's rational full.(r/s) beliefs, the probability of S is $[p, q]$, where p is greater than $r/(r+s)$. Then for no $c \in C(A, r/s)$ could A rationally bet against S, but if S is not a member of A's full!(r/s) beliefs, neither can we demand that A act as if S were true under any

$c \leq C(A, r/s)$. This is the situation discussed in the last section and resolved there by manipulating the ratio r/s . But here we are dealing with given circumstances $C(A, r/s)$. What should A 's doxastic attitude toward S be?

It seems perfectly natural to say both that S is not a statement that A believes and also that it is not a statement against which he would bet under any circumstance in $C(A, r/s)$. This seems to be a perfectly natural (and indeed familiar) circumstance to be in. But we may also suppose that the possibility of a serious bet against S at high enough odds would change the circumstances contemplated from $C(A, r/s)$ to $C(A, r'/s')$ where $r'/s' > r/s$, and the odds of this possible serious bet are such that they fall in the range r'/s' to s'/r' . Then the analysis of the preceding section would hold. Life would have manipulated the ratio r/s . I do not, therefore regard this as a real difficulty.

Another nonserious difficulty arises from our treatment of probability as interval-valued. Suppose that, relative to A 's (r/s) rational beliefs, the probability of S is $[p, q]$, where $\frac{r}{r+s} \in [p, q]$. Then A should bet on S at odds less than $p:1-p$, but there are no odds at which he should bet against S , though the ratio of stakes contemplated, r/s to s/r , includes some at which A could rationally bet against S . I think this situation will be found anomalous only by those who think that A 's doxastic state should be represented by a single classical Bayesian distribution.

V

There are a number of consequences that this approach to partial and full belief has for real life that bear on relatively practical matters. We all have useful stochastic knowledge; trivially, that well maintained gambling apparatus is very nearly fair. From such knowledge we may infer

that stochastically ideal apparatus produces certain outcomes with very small (or very large) probabilities. Thus the probability of heads on each of 1000 tosses of a fair (ideal) coin is $1/2^{1000}$. There is nothing wrong with this computation. But on the basis of the preceding analysis, it is only useful for computing expectations in a set of circumstances that includes an enormous range of odds: $(2^{1000}-1):1$ to $1:(2^{1000})$. Such a range of odds does not characterize anybody's practical concerns. Within the range of odds representing most people's circumstances, "heads on each of the next thousand tosses" can be taken as false.

This explains, among other things, why nobody wants to play the Petersburg game for what it is "worth". (The Petersburg game is that in which player A offers to pay player B a prize of 2^k dollars, if a head first appears on toss k in a sequence of coin tosses.) It is easy to compute that the value of this game to B is "infinite". Given any range of feasible odds, we can calculate the value of the corresponding truncated Petersburg game; but for any reasonable set of full beliefs, almost all of the possible outcomes of the Petersburg game have 0 probability: we can give full rational belief to their non-occurrence.

→ Note that there are no circumstances under which it is reasonable to pay \$1 for a ticket that will return \$ if the ratio of stakes is 2.
 Similar considerations bear on our proper epistemic attitudes toward very rare events. In my practical corpus of level $1-\underline{e}$, for example, it is simply not credible that an event of probability less than \underline{e} can occur. Therefore the expected value of a prize contingent on such an event is exactly 0, regardless of how valuable the prize is itself. Correspondingly, in a corpus of level $1-\underline{e}$, the expected cost of a disaster that has a probability of less than \underline{e} is 0, however dreadful the disaster.

Of course, to contemplate a glorious prize (eternal joy) or a horrible disaster (eternal damnation) may lead the agent to increase the range of risks and benefits -- the range of odds -- he wishes to take account of in

his corpus of practical certainties. But observe that this has the following effect: it constrains him to operate with a very high standard of evidence (the square root of the level of the corpus of practical certainties that corresponds to his range of odds), and that means that the actual content of his corpus will be very sparse. This does not mean that the ideal content need be sparse -- the ideal coin lands heads a thousand times in a row on exactly $1/2^{1000}$ of the time. But no real coin is ideal, and if our standard of practical certainty is given by the index $1 - 1/2^{1000}$, there is very little we know about the real world that meets this standard.

More generally, in order to use a probability of \underline{e} in the computation of an expectation, the probability of the statistical law on which the first probability is based must itself be greater than $1 - \underline{e}$, and only relative to an evidential corpus of a level higher than $1 - \underline{e}$ does that computation make sense. A possibility, relative to a corpus of level $1 - \underline{e}$, whose probability is less than \underline{e} is not a real possibility.

Given a knowledge of an agent's preferential dispositions, we can characterize sets of circumstances $\underline{C}(\underline{A}, \underline{r}/\underline{s})$. Then, given a knowledge of the agent's full¹($\underline{r}/\underline{s}$) rational beliefs, we can divide the actions he might contemplate into those he ought rationally to perform; those he ought rationally to refrain from; and those concerning which there are no rational constraints. If we wish to assess the rationality of his full¹($\underline{r}/\underline{s}$) beliefs, we may do so by considering the probabilities of the statements fully¹($\underline{r}/\underline{s}$) believed, relative to the agent's full¹($\underline{r}'/\underline{s}'$) beliefs, where $\underline{r}'/(\underline{r}'+\underline{s}') = (\underline{r}/\underline{r}+\underline{s})^{1/2}$.¹² We may do this for any set of circumstances $\underline{c} \in \underline{C}(\underline{A}, \underline{r}/\underline{s})$. This serves, I maintain, to give us a complete handle on both the agent's rational beliefs and his rational actions, subject only to the

characterization of his utility function -- obviously a nontrivial matter, and one that itself involves dispositional interpretation. Thus although the dispositional characterization of actual belief does seem to open up a Pandora's box of confederate notes, and thus to represent a research program rather than an enlightenment, the dispositional characterization of rational belief calls for only one blank check (to be filled out in A's utilities), and otherwise admits of relatively clear-cut prescriptions. As in many other areas of endeavor (in geometry, for example) it turns out to be easier to prescribe the ideal (the rational) than to describe the real (the actual). But the ideal is not without relevance to the real, as we have already briefly noted in this final section.

The definition of full belief (D1') and the two principles -- the rationality principle (RP') and the acceptance principle (AP) -- provide a plausible non-Bayesian¹³ framework in which to discuss both rationality of belief and rationality of action. That in the case of action it does not always yield an answer -- when the upper and lower probabilities don't lead to the same result, for example -- strikes me as all to the good. Why should we expect mere human rationality always to guide us?

Henry E. Kyburg, Jr.
University of Rochester

* An earlier version of this paper appeared in "How Many Questions?" A collection of essays honoring Sidney Morgenbesser, edited by Leigh S. Cauman, Isaac Levi, Charles D. Parsons, and Robert Schwartz. Thanks are due the editors for permission to reuse the material.

1. Carnap, Rudolf, The Logical Foundations of Probability, Chicago, University of Chicago Press, 1950).
2. Jeffrey, Richard, The Logic of Decision, New York, McGraw-Hill, 1965. Second edition, Chicago, University of Chicago Press, 1983.
3. Diaconis, P., and Zabell, S., "Updating Subjective Probabilities", Journal of the American Statistical Association 77, 1982, pp. 822-830.
4. See Ward Edwards, "Subjective Probabilities Inferred from Decisions," Psychological Review 69 (1972): 109-135, and references therein; see also C. Stael von Holstein, ed., The Concept of Probability in Psychological Experiments (Dordrecht: Reidel, 1974) and D. Kahneman, P. Slovic, and A. Tversky (eds.) Judgment Under Uncertainty: Heuristics and Biases, Cambridge University Press, Cambridge, 1982.
5. R. B. Braithwaite, "Belief and Action," Proceedings of the Aristotelian Society, XX (1946): 1-19.
6. The example comes from Levi and Morgenbesser, "Belief and Disposition," American Philosophical Quarterly 1, 1964, pp. 221-232.
7. Ibid., p. 223.
8. I refer here literally to a set of tosses, not a sequence of