AD-A250 539

NTATION PAGE

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	1988	Unknown

4. TITLE AND SUBTITLE

In Defence of Intervals

DTIC ELECTE MAY 2 6 1992 S A D

6. AUTHOR(S)

Henry E. Kyburg, Jr.

5. FUNDING NUMBERS

DAAB10-86-C-0567

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Rochester
Department of Philosophy and Computer Science
Rochester, NY 14627

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U.S. Army CECOM Signals Warfare Directorate
Vint Hill Farms Station
Warrenton, VA 22186-5100

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

92-TRF-0005

1. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Statement A; Approved for public release; distribution
unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Why intervals or upper and lower probabilities? And how should such a system work?
I claim that if a system for the representation of uncertainty is to be useful, it
must be possible to obtain from it constraints on decisions in the face of uncertainty.
By this, I mean that, as in "real life", even if we are in a state of some uncertainty,
there should ordinarily be some actions that are ruled out by rational considerations.
Interval representations of uncertainty will accomplish this.

1. SUBJECT TERMS

Defence of intervals, Aritifical Intelligence, Data Fusion,
Reference Class, Probabilities

15. NUMBER OF PAGES

17

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

IN DEFENCE OF INTERVALS*

by

Henry E. Kyburg, Jr.

Philosophy and Computer Sciencew

University of Rochester

0. Why intervals or upper and lower probabilities? And how should such a system work? I claim that if a system for the representation of uncertainty is to be useful, it must be possible to obtain from it constraints on decisions in the face of uncertainty. By this, I mean that, as in "real life", even if we are in a state of some uncertainty, there should ordinarily be some actions that are ruled out by rational considerations. Interval representations of uncertainty will accomplish this.

A number of representations of uncertainty fail on this score. This is clear in the case of such non-probabilistic representations as that provided by certainty factors. It also fails on the unadorned Bayesian view, since on that view, any coherent probability function may be adopted as an uncertainty representation. If we add additional constraints, such as an original distribution of probability over the sentences of a language prior to any evidence (a la Carnap [1950]), or a prior distribution determined by a statement of the problem at hand (a la Jaynes [1958]), then we can apply the Bayesian procedure, and the classical maxim to maximize expected utility. But then we must find some reason to assign probabilities to sentences in the way we do, or to state the problem at hand in the particular way we do.

Failing this, rationality imposes no bounds, via the principle of maximizing expected utility, on our decisions. The alternative that drives the view being defended here is that there are reasons to eschew some alternatives that are not simultaneously

reasons to accept a particular alternative. Most simply: if we know that between 20% and 30% of the rolls of a particular die land with the five up, we should not offer more than seven dollars to receive three that five will not occur on a particular roll (unless we know something more about that roll), nor should we offer more than two dollars to receive eight that a five will occur on an otherwise unspecified roll.

In general, for rationality, we should require that the uncertainties have an appropriate source. Where does this source come from? The idea, underlying the approach that I defend [Kyburg, 1974], is that the source of the uncertainties on which our decisions are to be based should be statistical knowledge, often, but not always, based on statistical inference.

It is often said that when we have statistical knowledge, there is no problem in representing and using uncertainty. On the contrary, (a) we always have statistical knowledge (at worst, that a relevant ratio lies between 0 and 1, but it is usually more narrowly constrained than that), and (b) given that we have a lot of statistical knowledge (for example, if we are an insurance company with access to an enormous database) we still have the problem of relating the instance at issue to the appropriate reference class.

The representation of uncertainty as an interval is thus not an end in itself; it is a consequence of the fact that our knowledge, in general, only embodies approximate statistical knowledge, and of the fact that we want our uncertainties to be founded in statistical fact, rather than in subjective opinion. These are the principles that guide us in our choice of a paradigm.

1. The paradigm is the classic one of balls in an urn. Note that I say balls in an urn, not balls **drawn from** an urn; there is no question of equal likelihood of drawing or long run frequencies of drawing. It is a static model, not a dynamic one, representing an epistemic snapshot of an individual's (or a group's) state of knowledge. Thus what is paradigmatic is our **knowledge** about the urn, not the **state** of the urn.

2

Of course to apply this model, we must have some way of specifying particular balls. Classically one says, "the next ball to be drawn," but this leads away from the urn, and toward a process, such as tossing a coin, as the model. So let us suppose that the balls are numbered in a way that to us is completely arbitrary.

We ordinarily think of black and white balls, but this is easy to generalize. We could think of a finite spectrum of colors, $c_1, ..., c_n$. Or we might think of properties such as mass, that can have any positive real number as value. Let us associate with the model a **vector** of functions (color of( ), mass of( ), ...).

The items whose uncertainty concerns us are **propositional** in character: ball # 17 has a mass between .2 and .8 Kg, or $.2 \leq M(17) \leq .8$, or $C(23)$ = Black.

Our state of knowledge of the contents of the urn falls into a number of different categories:

(a)   Knowledge of particular balls: We might know that ball number 7 is blue. At the same time, we may be completely ignorant of its mass. That is, we can have partial knowledge of the properties of particular balls.

(b)   Universal knowledge relating the properties of balls: We might know that **all** the blue balls weigh less than 1 Kg. I construe this as a genuine **universal** statement: (x)(C(x)=Blue ... M(x)<1.0).

(c)   Existential knowledge, construed in the same way.

(d)   Exact statistical knowledge: We might know that 15% of the blue balls weigh less than 0.5 Kg.

(e)   **Real knowledge:** Approximate statistical knowledge.

While we do have universal knowledge (e.g., if X is longer than Y, then Y is not longer than X), I am doubtful that it is empirical. Existential knowledge generally just follows from our knowledge of particulars. Exact statistical knowledge is no more available than universal knowledge (it exists: it is surely true that a **fair** coin tends to land

3

heads exactly half the time!). But approximate statistical knowledge is the building block of our treatment of uncertainty. We can know that between 10% and 20% of the balls are black. We can know that **almost all** (this is often written as if it were universal) the blue balls weigh less than 1.0 Kg. We can know that the distribution of weights among the red balls is **approximately** normally distributed with a mean of 0.50 and a standard deviation of 0.04.

This approximate statistical knowledge is also uncertain, but not in a sense that is directly relevant to behavior or decision. We have a two-level treatment of uncertainty. This is treated in more detail in the answer to question 5.

One way to represent approximate statistical knowledge is to replace the single paradigmatic urn by a set of urns, each of which contains a determinate proportion (after all, this is true of real urns). The world corresponds to a specific urn; our knowledge of the world corresponds to a set of urns: we know that the world is among this set.

THE ESSENCE OF THIS MODEL IS THAT ALL RATIONAL UNCERTAINTIES ARE DETERMINED BY OUR STATISTICAL KNOWLEDGE OF THE WORLD.

There are 6 aspects of this model that deserve brief comment.

(1)   The notion of proportion makes sense only in a finite population; but to represent independent coin tosses, or a normal distribution of mass, we must suppose the number of balls in the urn is countable. (We needn't suppose more than a countable number, even for continuous distributions, but that doesn't help much.) Seriously speaking, however, no population is actually infinite -- even the tosses of a coin are limited -- but we can treat coin-tosses and continuous distributions as **idealizations** of actual finite populations.

4

(2)   To be ignorant or partially ignorant of the proportion **p** of blue balls that weigh less than half a Kg is to have as one's model of the world a set of urns with the parameter **p** falling in some interval ([.4,.5]. [0.0.1.0], ...). Again this seems to call for an infinity. We could accommodate this in our model (we are not **drawing** an urn), but for purposes of representation we can equally well choose some **granularity** that is determined by context. Then these sets of urns will be finite.

(3)   Nothing has been said yet as to **how** uncertainties are to be determined by statistical knowledge. This is taken to be a matter of choosing the right reference class, or of epistemic randomness. I have talked about this problem elsewhere; for present purposes it suffices that in the presence of statistical information, together with partial knowledge about a particular ball, we can obtain a determinate (interval-valued) probability, based on statistics, that that ball has any given property.

(4)   How about Boolean combinations of statements -- e.g., Ball 14 is <u>not</u> red <u>and</u> ball 27 weighs less than 0.5 Kg? Negation, of course, can be handled in the same model. The combination of statements can be handled in a new derived model belonging to the same paradigm.

That ball 14 is not red and ball 27 weights less than 0.5 Kg <u>has the same truth-value as</u> <ball 14,ball 27> belongs to the set of pairs of which the first is not red and the second weighs less than 0.5 Kg.

In the new model (the cross product model) each "ball" corresponds to a pair of balls, not necessarily distinct, each of which comes from our original urn.

Suppose our original urn model contains between 20% and 30% red balls; and suppose it contains between 40% and 60% balls weighing less than 0.5 Kg. It follows that this original model contains between 70% and 80% non-red balls.

In the new cross-product model between 28% (0.7 * 0.4) and 48% (0.8 * 0.6) of the pairs are pairs in which the first ball is not red and the second weighs less than 0.5 Kg.

We can take account of the fact that ball 14 and ball 27 are distinct, if we know the original urn contains N balls, by looking at the subset of our new model consisting of pairs $<x,y>$ such that $x \neq y$. The proportion of non-red balls in that subset of the new urn is between $.28\ N^2/N\ (N - 1)$ AND $.48\ N^2/N\ (N - 1)$.

Suppose we are interested in the probability that <u>one</u> ball, ball 17, is both non-red and weighs less than 0.5 Kg. Then the smallest subset of the new model to which we know $<17,17>$ belongs is the diagonal: $\{<x,y>: x = y\}$. But we do not know anything about the frequency in this diagonal (without making further assumptions), and for reasons to be expounded later, we use the contents of the whole urn as a reference class.

This does **not** embody any assumption of independence: it represents a straight-forward statistical computation. Of course we <u>might</u> know something about a connection between being non-red and being light. We might know that non-red balls were rarely light. But this is an **additional** piece of knowledge, and is represented perfectly easily in a different model belonging to our our original paradigm, one in which our original urn model is a marginalization of the new one in which dependence is represented.

(5)   How about statements that are not of the form, "Ball 17 is green?" Every statement can be put in this form! To say that x, y, and z stand in a certain relation is just so say that a certain object, the triple <x,y,z> belongs to a certain set or has a certain property.

(6)   But how about statements for which we have no statistics on which to base an assessment of uncertainty? I claim there are no such. Two considerations suggest this: Statements <u>known to be equivalent in truth value</u> should have the same

probability; so all we need do is find one in this class for which there is a statistical basis. We do **not** require logical equivalence! And (as a special case of this) objects and events may be picked out under various descriptions; different descriptions suggest different reference classes.

2 . The main aspect of uncertainty not captured by this model is that aspect on which Bayesian models primarily focus: the **psychological** aspect. The study and representation of partial belief is no doubt worthy and interesting, but my concerns are with the epistemological dimensions of uncertainty. Thus it does not matter to me that people may not in fact apportion their beliefs in accordance with statistics. For that matter, it does not bother me that people violate the laws of the probability calculus in their degrees of belief. I am, in fact, doubtful if partial beliefs can or should be measured by real numbers between 0 and 1; I suspect a vector representation would be an improvement. But all this reflects another focus; my focus is **logical and epistemological**. It is **normative**; it is what Isaac Levi calls **necessitarian**: given the background knowledge, only one uncertainty is allowed. It is represented by an interval. The meaning of this interval will be discussed below.

A second aspect of uncertainty that is not captured concerns vagueness or fuzziness. This divides into two parts. First, the intervals that emerge as the uncertainties on this model have sharp endpoints. Thus instead of saying that the probability of heads on the next toss of this coin is about a half, we are obliged in the model to say that the probability is [0.49,0.51]. It is argued that we have just replaced one sharp point (0.50) with two of them. That makes representation even harder! But in reply we note that a fuzzier representation requires more than two points -- it requires parameters enough to characterize the appropriate fuzzy distribution; and on a Bayesian

7

view it requires parameters enough to characterize the full n-dimensional distribution representing our beliefs about n tosses of the coin.

My main defense of my failure to be true to psychology is that the theory is intended to be normative, and not concerned with the _representation_ of belief. It is concerned with how human beliefs _should_ be, and with how our machine equivalents should be.

Second, vagueness and fuzziness are involved in the statements whose uncertainties we seek to evaluate. My model has little to offer here. We can represent this fuzziness as a statistical frequency with which a certain object is put (by ordinary people) into a certain category. But this doesn't add much.

3.   The intervals that emerge from the urn-set paradigm are given meaning in two complementary ways. In the first place, the intervals of uncertainty are derived from our **knowledge** of statistical distributions. They are intervals because our knowledge is incomplete and approximate. (If our knowledge were complete we would have no need to take account of uncertainty!) Suppose that the probability that ball 14 is red is [0.6,0.7]. Suppose that this is derived from the fact that ball 14 is a random member of the set of zinc balls (relative to what we know) with respect to being red, _and_ from our knowledge that between 60% and 70% of the zinc balls are red.

This statistical knowledge is knowledge about the world. To claim to know this about the zinc balls is to claim to know something about the world: namely that the unique proportion of zinc balls that are red may be anywhere between 0.6 and 0.7, and cannot be less than 0.6 or more than 0.7. Furthermore, it is implicit that this bit of statistical knowledge is rationally justified: I have _reason_ to accept this statistical hypothesis, either through sampling and statistical inference, through physical

8

considerations, on the basis of good authority, or whatever. This is what gives empirical meaning to the interval [0.6,0.7].

At the same time, the interval in question has normative meaning. We could say that it is legislative for rational belief, in the sense that any degree of belief falling outside the interval would be irrational, and any degree of belief falling inside the interval would be rational. But for reasons already mentioned, I am skeptical of "degrees of belief" in this context. I would prefer to say that the normative meaning is captured by **behavioral injunctions**. For example, in the situation under discussion, it would be irrational for the agent to pay more than $.70 for a ticket that would return a dollar if ball 14 turned out to be red. It would be similarly irrational for him to sell a ticket that he would have to redeem for a dollar in case ball 14 turned out to be red for less than $.60. In between these limits we have no ground for impugning his rationality.

Of course if the agent were simultaneously to pay $.70 for a ticket that returns him $1.00 if ball 14 is red, and to sell someone else a ticket for $.60 that obligates him to pay out $1.00 if ball 14 is red, then we should regard the agent as irrational. But this is not on grounds having anything to do with uncertainty (though the Dutch book arguments of Bayesians try to make us believe otherwise), but with the certainty that no matter what happens the agent has committed himself to losing $.10.

The normative meaning of the uncertainty interval lies in the constraints it imposes on decision-making. The epistemic meaning of the uncertainty interval lies in the fact that it embodies a claim that some class in the real is known to be characterized by the interval [0.6,0.7], and that this class is related to the proposition in question in the appropriate way.

4.	The rules for choosing the right reference class are such that they are sensitive to small variations in statistical knowledge. I shall first illustrate this, and then attempt to get around it.

It is now easy to see where the sensitivity to precision arises. Suppose we have looked at a large sample of red balls, and a large sample of heavy balls, and have reasonably concluded that between 0.499 and .502 of the zinc balls are red, and that between .499 and .501 of the heavy balls are red. Suppose that we have recorded no instances of balls that are both heavy and red. Of course we do have the trivial statistical knowledge that between 0% and 100% (inclusive!) of the heavy zinc balls are red. No matter: the heavy balls have it, and the probability that ball 14 is red is [.499,.501].

But suppose we had just a tiny bit _less_ information about the red balls, so that we can say only that between .498 and .501 of them are red. Now "zinc" and "heavy" differ; neither is a subclass of the other; and the intersection yields only the [0,1] interval as a probability. Small changes in knowledge can lead to sudden discontinuous shifts in reference class, and thus to sudden and discontinous shifts in probability.

Granularity serves to resolve this problem to some degree. If we round to two decimal places, "zinc" and "heavy" agree precisely in the proportion of red balls: .50 in each case. But if we don't want to turn to granularity, or if there are reasons for eschewing a coarser grid in a particular case, we may be stuck with these discontinuities.

So the other half of this argument is that in such cases, it is not such a bad thing to be stuck with. Suppose we think we really need three decimal places worth of precision in our uncertainty. We have the evidence to give us the precision we desire for heavy balls and for zinc balls. But we have no direct evidence about the frequency of red balls in the intersection of these two classes. If we really think we need three-place precision, the conflict between heavy and zinc as evidence for redness quite properly tells us we should seek more (and more direct) evidence.

10

5.   There are three kinds of reasoning associated with uncertainty that might be automated.

(1) There is the manipulation of uncertainties; this is generally considered to be a matter of applying the probability calculus. When uncertainties are represented by intervals the probability calculus does not apply directly. What we can say, however, since all probabilities are based on statistical knowledge, is that the probabilities must be consistent with the statistical background knowledge we are assuming. The manipulations of the probability calculus apply straight-forwardly to the general statistical knowledge on which our probabilities are based. If we had exact statistical knowledge, this would come to the same thing as applying the probability calculus to our uncertainties. Since we do not have exact statistical knowledge, the probability calculus can merely impose constraints on our uncertainty intervals. There is clearly no problem in automating inference in accordance with the probability calculus. But it is less clear that there is a useful way of automating the direct manipulation of interval uncertainties themselves.

(2) The approximate statistical knowledge that determines the content of our paradigmatic model is, as we noted earlier, uncertain. One view of "knowledge" taken in this sense is that it derives from high probability. (This is highly controversial in philosophy!) On any view of uncertainty that I know of, if S is a sentence and T is a sentence implied by S (i.e., T is derivable from S) then the probability of T is at least as great as that of S, or the support of T is at least as great as that of S, or the interval of uncertainty of T has a lower bound at least as great as that of S, etc. Thus if S implies T, and S is part of our (uncertain) knowledge, T will be also.

Let us look at Modus Ponens and Modus Tollens. On the view of uncertainty I am endorsing, uncertainty is based on statistical knowledge. In general if we know that

11

almost all P 's are Q 's, and x is a random member of P with respect to Q , we may be practically certain that x is Q . This is much like modus ponens. If practically all crows are black, and I know that Charles is a crow, I may be practically certain that Charles is black. It is noteworthy that nothing like this corresponds to modus tollens. If I know that practically all crows are black, and that Peter is not black, I cannot infer with practical certainty that Peter is not a crow. In the realm of statistically based inference, modus tollens fails. Perhaps only crows fail to be black.

(3) One inferential mechanism brought into prominence by Bayesian theorists is the updating of uncertainty. In the Bayesian framework, updating is easy. But the space required by a priori probabilities is extremely large. The framework that I have offered is less demanding with regard to space, but updating requires recomputing all probabilities relative to a new body of knowledge -- one augmented by the new evidence.

Of course in many situations, Bayesian conditionalization is just the appropriate mechanism. When? Exactly when it was thought to be the right mechanism by such classical theorists as Neyman and Pearson [1938] and Fisher [1956]: when the procedure can be given a statistical model. (If we know the proportion of balls that are zinc, and the proportion that are both zinc and red, the "conditional probability" of red given zinc is just the obvious ratio! For which we have a statistical justification.)

Research efforts are currently under way to relate the notion of probability provided by our basic model to various approaches to default reasoning and non-monotonic logic [Kyburg, 1988a]. Some progress has been made in showing that classical default reasoning and non-monotonic logic can be represented as a special case of probabilistic logic.

6 . The most serious computational difficulty is engendered by the fact that when we seek the uncertainty of compound propositions, we are in for exponential difficulties

in general. (the canonical model for P and Q is the cross product -- in general -- of the canonical models for P and for Q). The response "tu quoque" addressed to the Bayesians is valid, but does not help with computation.

Given that we are concerned with a single proposition (ball #14 is red), there are still computational problems. If there are N basic sets to which ball #14 may belong (zinc, heavy,...), there are N^2 intersections of sets to which ball #14 may belong. Furthermore, as we have already indicated, given a sentence S, we must look for the probability, not only among the sentences involving the same subject as S, but among the sentences involving the same subject as any sentence truth-functionally equivalent to S.

The prospect sounds dreadful. But in fact it may not be as bad as it sounds. There are shortcuts and simplifications that we are currently exploring in detail. It is hoped that the exponentially growing parts of the program can be severely bounded, and that useful heuristics can guide the search in what is left.

7a. The strong point of interval valued probability is that it is very directly tied to statistical inference -- the source of the statistical knowledge that serves as the basis for uncertainty statements. Furthermore, since all probabilities are founded in our knowledge of long-run statistical statements, it is automatic that decisions that are inconsistent with expectations determined by our uncertainties will in the long run (if the knowledge on which our uncertainties is based is correct) yield negative utilities almost certainly. Thus there is a natural tie both to statistical inference about the world, and to long run expectations of utility.

The weakest point lies in the implementation. In order to apply this theory in a useful way, we must formalize and represent not only the propositions characterizing a limited domain, but those that embody the common sense of ordinary people. The

13

system is strongly holistic. It is not clear, however, that any system that is not similarly holistic can take account of common sense knowledge. An ongoing project is the search for ways in which we can modularize the computation of probabilities.

Another weak point concerns the sensitivity of the theory to imprecision -- the fact that tiny changes in our knowledge of frequencies can have profound effects on the actions it is rational to perform and the decisions it is rational to make. We conjecture that these difficulties can be very much alleviated, if not eliminated altogether, by finding a principled way to incorporate granularity into our considerations. But this is a project, not a *fait accompli* .

The strength of the model lies in its sound base in statistical fact. The corresponding weakness of other models is that they depend on ad hoc constraints, or subjective assessments, to arrive at uncertainties. This is particularly true of the subjectivistic or personalistic Bayesian model. The Dempster-Shafer model, since it can (as we have shown) be represented as a convex set of epistemic models, fits into our framework. But as conventionally stated, this model does not tie in explicitly to statistical knowledge. And it is not clear that its authors would want it to! It is also more limited in what it can represent.

A weakness of the interval model, in addition to its computational complexity, lies in the fact that IF people have degrees of belief, and IF they are rational in deploying those degrees of belief in computing expectations, and IF they are sound in computing mathematical expectations, THEN the decisions they make will be directed at maximizing their expected utilities, regardless of whether or not these expectations are based on statistical knowledge. Abstract rationality may not have much influence on what people do. But, again, this theory is concerned with what people ought, or ought rationally, to do, not with what they do in fact. The latter is a matter of psychological inquiry; I take the former to be a matter of logical or epistemological inquiry.

14

7b. The main difficulty of the view that I have been advocating is that it is computationally horrendous. I see the virtue of various other theories as providing computationally feasible shortcuts that, under the right conditions, can represent the outcome of applying my approach.

Thus, for example, Dempster updating is computationally very simple: two simple computations lead from upper and lower probabilities, to new upper and lower probabilities based on the evidence provided by a proposition [Shafer, 1976]. This corresponds to the upper and lower conditional probabilities on a Bayesian model. But it also corresponds to the upper and lower frequencies on a model in which the prior probabilities are construed as frequencies, and conditionalization applies [Kyburg, 1987]. So updating, under the special conditions in which Dempster updating is appropriate, can efficiently take place by means of a very simple algorithm.

Of course Bayesian updating by simple conditionalization is appropriate (as a limiting case) when our prior knowledge is rich enough and precise enough. Again, I see this as a limiting special case applicable when our actual situation approximates the special case.

In general, other approaches to the manipulation of uncertainty have considerable computational advantages. But I see these approaches as approximations of the approach described here. Thus these approaches are advantageous exactly when the approximations they embody are justified. But when they are justified is a question that must be adjudicated by our general approach.

8. The preceding section has described what I see as the main relation between other approaches to uncertainty and the one described above. There is a lot of room for

integration, but little room for compromise. That is, I see the view I advocate as being more fundamental, less subject to subjectivity, than other views.

One advantage is a closer integration into decision theory, which I see as fundamental to engineering. Another is the computational tie to ordinary statistics. But any approach that is consistent with a statistical basis can be taken as acceptable, and insofar as it is computationally advantageous (I have already admitted that we need all the computational advantages we can get!) it can be integrated with the approach advocated.

Thus I see no room for hybridization in general, since I see no immediate connection between people's degrees of belief (if there are any such things) and the normative question of how they should make decisions. I do see computational advantages to other points of view, but they can be represented within our point of view.

Where there may be room for hybridization is in connection with fuzzy sets. Our approach is set-theoretical, but we have not dealt with the potential fuzziness of sets. I am not sure how the hybridization would proceed, but it appears to be an interesting and fruitful question to pursue. In particular, it would be attractive to be able to use fuzzy numbers for the upper and lower limits of the uncertainty intervals.

Henry E. Kyburg, Jr.
University of Rochester

Bibliography

Carnap, Rudolf: *The Logical Foundations of Probability*, University of Chicago Press, Chicago, 1950.

Fisher, Ronald A.: *Statistical Methods and Scientific Inference*, Hafner Publishing Company, New York, 1956.

Jaynes, E. T.: "Probability Theory in Science and Engineering," *Colloquium lectures in Pure and Applied Science*, Socony Mobil Oil Co., Dallas, 1958. pp 152-87.

Kyburg, Henry E. Jr.: *The Logical Foundations of Statistical Inference*, Reidel, Dordrecht, 1974.

Kyburg, Henry E., Jr.: "Full Belief," *Theory and Decision* **25**, 1988, 137-162.

Kyburg, Henry E., Jr.: "Probabilistic Inference and Non-Monotonic Inference," Shachter, Ross, and Levitt, Todd (eds): *The Fourth Workshop on Uncertainty in Artificial Intelligence*, 1988, pp 229-236.

Kyburg, Henry E., Jr.: "The Reference Class," *Philosophy of Science* **50**, 1983, pp 374-397.

Kyburg, Henry E., Jr.:"Bayesian and Non-Bayesian Evidential Updating," *A I Journal* **31**, 1987, pp 271-294.

Neyman, Jerzy, And Pearson, E.S., "Contributions To The Theory Of Testing Statistical Hypotheses," *Statistical Research Memoirs* **1**, 1936, 1 - 37; **2**, 1938, 25-57.

Shafer, Glenn: *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, 1976.