(2)

REPOR

AD-A255 236

GE

Form Approved
OPM No. 0704-0188

ublic reporting burden for this collectio
maintaining the data needed, and review
for reducing this burden, to Washington
the Office of Information and Regulatory

times for reviewing instructions, searching existing data sources gathering and
mments or any other aspect of this collection of information, including suggestions
, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to

Washington, DC 20503.

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

4. TITLE AND SUBTITLE

Normative and Descriptive Ideals

5. FUNDING NUMBERS

DAAB10-86-C-0567

6. AUTHOR(S)

Kyburg, Henry E., Jr.

DTIC
ELECTE
SEP1 0 1992
S C D

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

University of Rochester
Department of Philosophy
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-0048

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Statement A: Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words) Choice reflects the interplay among the facts of human cognition and especially human limitations, out intuitions concerning the principles to be found in simple cases, and pragmatic considerations of the sort last mentioned. Probability gives us a handle on these things, but of course probability is one of them itself. Is it more fundamental? Sure. Does that cut much ice? Not much. To the extent that we can limit our considerations to what is common to a number of views of probability, we may find that we have a basis for choice. What we need is an account of defeasible reasoning - that is, of a kind of reasoning that will give us practical certainties, but that is not monotonic. It will be a knod of reasoning that allows us to change our minds in response to new evidence. This principle of specificity seems to work well in a limited domain of cases - namely the cases in which we should be looking at subsets. Construed broadly, as logical strength, the principle seems uncontrovertible, but it also seems devpod pf isefi; cpmtemt. Construed as comething in between the recommendation to choose a superset as a reference class, and the general heuristic recommendation to use all the evidence you have, it is not clear how to apply the principle.

14. SUBJECT TERMS	15. NUMBER OF PAGES
Artificial Intelligence, Data Fusion, Deductive Logic, Defeasible Reasoning, Decision Theory, Default Reasoning, Non-Monotonic Reasoning.	14
	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

ISN 7540 01-280-5500

Standard Form 298, (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

(2)

REPOR

AD-A255 236

GE

Form Approved
OPM No. 0704-0188

Public reporting burden for this collection ... maintaining the data needed, and review ... for reducing this burden, to Washington ... the Office of Information and Regulatory ... Washington, DC 20503. ... time for reviewing instructions, searching existing data sources gathering and ... aspect of this collection of information, including suggestions ... 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to

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

4. TITLE AND SUBTITLE

Normative and Descriptive Ideals

5. FUNDING NUMBERS

DAAB10-86-C-0567

6. AUTHOR(S)

Kyburg, Henry E., Jr.

DTIC
ELECTE
SEP1 0 1992
S C D

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

University of Rochester
Department of Philosophy
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-0048

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Statement A: Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words) Choice reflects the interplay among the facts of human cognition and especially human limitations, out intuitions concerning the principles to be found in simple cases, and pragmatic considerations of the sort last mentioned. Probability gives us a handle on these things, but of course probability is one of them itself. Is it more fundamental? Sure. Does that cut much ice? Not much. To the extent that we can limit our considerations to what is common to a number of views of probability, we may find that we have a basis for choice. What we need is an account of defeasible reasoning - that is, of a kind of reasoning that will give us practical certainties, but that is not monotonic. It will be a knod of reasoning that allows us to change our minds in response to new evidence. This principle of specificity seems to work well in a limited domain of cases - namely the cases in which we should be looking at subsets. Construed broadly, as logical strength, the principle seems uncontrovertible, but it also seems devpod pf isefi; cpmtemt. Construed as comething in between the recommendation to choose a superset as a reference class, and the general heuristic recommendation to use all the evidence you have, it is not clear how to apply the principle.

14. SUBJECT TERMS

Artificial Intelligence, Data Fusion, Deductive Logic, Defeasible Reasoning, Decision Theory, Default Reasoning, Non-Monotonic Reasoning.

15. NUMBER OF PAGES
14

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

NSN 7540 01-280-5500

Accession For
Dist Special
A-1

1988

Normative and Descriptive Ideals

1. Introduction.

There are a number of areas of inquiry that can be looked at either from a normative point of view or a descriptive point of view. David Hume, in the *Treatise of Human Nature*, remarked forcefully on the distinction between "is" and "ought." He took to task those authors who imperceptibly slip from asserting propositions about what *is* the case to asserting propositions about what *ought* to be the case. This distinction, and the fallacy of ignoring it, have become basic features of our scientific outlook. It is a proper part of science to *describe* the behavior of people in their relations to other people, to *describe* the events of the marketplace. The normative characterization of what people may or should do in their relations with other people, of how the marketplace should operate, is no part of science, but left for philosophers and preachers. The sociologist or anthropologist is warned, "Don't make value judgements!" It is true that scientists make value judgements pertaining to the conduct of personal and public life; but they do so as citizens and moral beings, not as scientists.

We can look at decision theory, to take one of the most complicated cases, either as a theory of how human beings do actually make decisions, or as a theory of how human beings ought ideally to make decisions. The former is descriptive; the latter normative. In the former case we are offering a theory of how people act. In the latter case we are offering a theory of how they ought to act. This seems simple enough on the face of it, but closer examination renders the distinction less clear.

92-24684

92

1506

The theory of how people ought to act is likely to have, for many, ethical or even moralistic overtones. (One ought to begin to eat only after the hostess has invited one to begin, or has herself started to eat.) On the face of it, decision theory is not of a piece with manners. Decision theory is purely practical: a question of how to choose in such a way as to maximize one's satisfactions, or expected satisfactions. It is a matter, as philosophers have been wont to put it, of practical reason.

Normative, no doubt. But we slid disjunctively over an important distinction: "in such a way as to maximize one's satisfactions" or "in such a way as to maximize one's expected satisfactions." The ideal goal of a theory of decision, one would think, would be to provide rules that would maximize one's actual satisfactions, not one's expected satisfactions. (We leave to one side the question, perfectly valid and at least as important as *these other questions*, of *what* satisfactions: *present*? *future*? actual? ideal?) Why is this not the goal of a normative decision theory?

To maximize satisfactions is not the goal of a normative decision theory, because we take it for granted that no theory can do this. In a similar vein, we take it for granted that no normative epistemological theory can insure both that our bodies of knowledge have content and that they eschew all error. These constraints on possible decision theories or on possible theories of knowledge are constraints that derive from our knowledge of human limitations. It is a part of the human condition that we cannot predict the future perfectly, and therefore cannot choose the course that will in fact maximize our satisfactions. It is part of the way in which we can, at best, learn about the world, that we cannot always avoid error.

In principle, this is a somewhat slippery slope. For just as I cannot avoid sometimes getting misleading samples when I sample from the world, however conscientiously I try to do my

sampling in the best way, so I cannot avoid sometimes getting the wrong answer when I add up a column of figures.

Yet we take the ideal, in the former case, to be getting samples that are no more misleading than would be expected statistically (at the 5% significance level, we expect to make false rejections of hypotheses no more than 5% of the time), while the ideal in the latter case is <u>never</u> to make a mistake.

What we take as a normative theory, as the ideal, depends on what we take to be possible for human agents, but it depends on it in no simple way. It nevertheless requires input reflecting the abilities, potentialities, and perhaps interests, of human beings -- all matters of descriptive rather than normative character.

On the other hand, a descriptive theory of human decision making appears to require a normative element. Thus, for example, although a man may choose A over B, he may recognize that he has made an error when he is reminded of his own probabilities and utilities. It is very difficult to draw a distinction between this case, and that in which a man adds 7 and 6 to get 15. *There seems to be an implicit underlying normative theory* behind the description of how people make decisions and what they take to be errors in their decisions.

Let us illustrate these ideas in three domains: deductive logic, defeasible reasoning, and decision theory.

2. Deduction.

Deductive logic seems quite straight-forward. If you believe that all men are mortal and that Socrates is a man, then you should believe that Socrates is mortal.

Humbug. In the first place it is not at all clear that standard first order logic is the appropriate representational discipline for beliefs -- even such simple beliefs as this, leaving entirely aside beliefs about necessities, obligations, and unicorns. Supposing that this is an appropriate representation for some beliefs, it isn't clear that this is one of those beliefs. Is the

universal quantifier the right quantifier? Is the truth-functional conditional the right connective? Should we be expressing the belief by some relation, rather than by a connective? If so, should we think of it as a relation between sentences, between propositional functions, or between properties, or classes, or sets?

It could be argued that in fact it is the intuitive validity of this argument that constrains the way in which we express the component ingredients of the argument. This in turn could be defended as a descriptive generalization of human habits of inference of the better sort. "The better sort," of course, could be unpacked in terms of success: in point of fact, such habits of inference will lead from truths only to truths.

An even better sort would be a habit of inference that will lead to a lot of truths, and only truths, from any premises, true or false. But that is not for us.

Let us leave to one side the problem of translating our body of knowledge into some standard notation, and let us suppose that the standard notation is that of first order logic. Given the ideality of the translation, we still should not expect the resulting theory to constitute a description of anybody's actual body of knowledge. Mere finitude precludes that. This applies as well to our a.i. systems: they, too, are finite, and thus cannot contain (for example) all the logical truths.

We might make our demand weaker: we might ask for consistency, rather that closure. Thus instead of saying that you should believe that Socrates is mortal, perhaps all we should ask is that you *not* believe that he is *immortal*. Even consistency admits of grades, though. Strong consistency would demand that the set of one's beliefs be consistent in the sense that no contradiction is contained in its deductive closure. But this is surely not a description of any realistic body of belief. Surely among the beliefs of the best of us there lurks a set of propositions that entails a contradiction.

So one might weaken the demand even more: we might ask only for k-consistency: that is, that there be no set of k or less statements in the body of knowledge that entail a contradiction.

Now we could turn back to the question of what logic we should adopt -- three valued? One with some form of strict implication or some form of subjunctive conditional? Should we decide on the basis of descriptive facts about the way people think? Or on the basis of the way we take ideal people to reason? Or on the basis of the way we think people *ought* to reason? Or on the basis of the way we want our ideal constructions in A I to operate?

There are genuine disputes about logic. They fall into two categories. Some people think that the standard first-order predicate calculus is seriously in error: we should adopt a logic that embodies a number of truth values beyond "true" and "false," or that is not committed to bivalence, or that does not have the distrilbution principle as a theorem. There are more people who think that standard logic is too weak to do justice to real human inference, and who therefore want to extend first order logic to deal with modal, deontic, epistemic, etc. arguments.

However we decide what logic is the best for our purposes, it will surely not be the case that our formal and explicit logic will accurately reflect how people do in fact reason or think; and it will surely not be the case that our formal and explicit logic will reflect the divine norm of achieving all and only truths.

3. **Probabilistic Inference.**

There is a lot of concern now in artificial intelligence for systems of defeasible reasoning, or default reasoning, or non-monotonic reasoning, or (a special case) inheritance hierarchies. In the classical philosophical tradition, this is the same as the traditional philosphical concern with inductive reasoning, non-deductive inference, ampliative inference, scientific reasoning, and so on. We may also think of it as the problem of probabilistic inference.

By "probabilistic inference" I do not mean manipulations of probabilities in accord with the probability calculus. I mean real inference -- that is, arguments that proceed from premises taken as evidence, to a categorical (not a probabilistic) conclusion supported by but not entailed by that evidence.

A paradigm case is that of statistical inference. If we obtain a large sample of A's, and find that the proportion in the sample that are B's is r; then a natural thing to do is to conclude that the proportion of B's among A's is close to r. This is NOT the same as concluding that it is probable that the proportion of B's among A's is close to r; though that, too, is true.

Whether or not it is best to allow probabilistic inference in this sense is open to question and argument. Some writers -- those who find the Bayesian paradigm attractive[2] -- think that a probabilistic rule of acceptance is awkward, unreasonable, and unwarranted. Others think that the only way to being to make substantive reasoning tractible is to allow for defeasible inference, and to permit the acceptance of statements that may in due course come to be rejected. Here again intuitions differ. And here again, it seems likely that only the construction and comparison of actual systems is going to lead to a resolution of the question. My own view is that, technical questions of intractability aside, a purely Bayesian approach that eschews acceptance is doomed to be practically unmanageable. But this is a question that deserves to be tested.

4. Decision Theory.

If deductive logic represents a mixture of normative and descriptive elements, decision theory is even more of a muddy mixture. The Bayesian principle that one ought to maximize one's expected utility seems on the surface to be both a description of what people, however inefficiently and awkwardly, try to do, and such a pervasive and intuitive principle of rationality that no one could deny it. The principle has come very close to being regarded

as a tautology in recent decades, since Frank Ramsey[3], Richard Jeffrey[4], Jimmy Savage[5], Ward Edwards[6], and others have taken preferences to reveal what there is to be revealed about probabilities and utilities.

Again, however, people's preferences are not consistent with the Bayesian axioms concerning probabilities and utilities. The theory is not, at first glance, an adequate descriptive theory of the preferences people have or the way in which they make decisions.

Should we treat this as a matter of weak flesh (or neurons)? Savage thought so. He recognized that his own preferences could well be inconsistent with the Bayesian axioms, and argued that the import of this was merely that he would be inclined to make adjustments in his beliefs (and utilities) at any point at which such an inconsistency might be pointed out to him. That is: the Bayesian axioms are to be taken as a normative standard of revision, rather than as a description of normal human preference structures.

But *there* are difficulties even with this somewhat normative view. First of all, people do not always modify their beliefs to conform to the calculus of probability, even when disparities are pointed out, as Kahneman and Tversky[7] and others have shown. Second, there are some cases, described by Allais[8], by Ellsberg[9], where very smart people have argued that the classic rules are wrong. *Third*, there are psychological studies, such as those conducted and referred to by Lola Lopes[10], that indicate that a number of central moments of distributions, rather than just the expectation, or even the expectation and variance, are relevant to the way people make decisions.

Our concern is not to try to argue one way or the other about these matters, but rather to point out that it is not at all clear how normative and descriptive aspects of decision theory are supposed to interact. We do not agree on what to take as normative -- even within the Bayesian framework -- and we take

the descriptive, when it is quite universal, as it is in the case of the "paradoxes" of Allais and of Ellsberg, to be an influential guide in our normative quest.

This is a non-trivial matter in philosophy, where, after all, we would like to understand what is right and what is wrong in theories of decision under uncertainty. But it is an absolutely crucial matter in artificial intelligence, where we are going to build into our systems some decision procedure. What our system is going to do will depend on what decision procedure we build into it. We had better give this question some serious thought if our systems are to be taken seriously.

This coin has another side: it may well be that the artificial systems can give us a testbed for our philosophical ideas about decision theory. We can see how one approach or another turns out. This can be done in several ways. The traditional philosophical way is to see if there are natural or anomalous consequences of the assumptions we make. An artificial system can generate a lot more consequences than a piece of paper and a pencil. There is also the possibility, not available to the traditional philosophy, of running such a system on a range of problems, or in a range of possible circumstances, and seeing how well it does. Perhaps we should measure a decision theory by its success!

This leads to another problem: We surely do not mean to measure the appropriateness of a decision theory by its short-term success. Even if Madam DuBois correctly picks the winner in the next horserace, we do not install her as our fundamental principle of decision theory. What we want is not an instance of success, but a reasonable indication of long-run success. And this is a matter of induction or defeasible inference, which we already considered in section 3.

5. **Probabilities.**

Even if everybody's preferences and choices could be bent into a shape that could be accounted for by a subjective probability function and a subjective utility function, there are normative considerations that would lead us to be critical of some of those functions. For present purposes we shall leave utility functions to one side, but there are surely probability functions that we would regard as crazy. This judgment has empirical and descriptive roots: most people don't believe preposterous things for no reason (except in matters of religion). There are, some of us believe, norms for probability judgment that go beyond conformity to the probability calculus.

According to Frank Ramsey, my degrees of belief should satisfy the probability calculus "on pain of inconsistency." This makes it sound as if (some of) the norms of probability judgment can rest on principles of deductive logic. But it is misleading. The alleged "inconsistency" consists in laying oneself open to having a book made against one. Someone whose beliefs were deductively sound would not have a book made against him regardless of his degrees of belief.

For example, if I have a degree of belief equal to 1/3 in heads, and a degree of belief equal to 1/3 in tails, I am, on Ramsey's view, expressing my willingness to bet at 2:1 against heads and my willingness to bet at 2:1 against tails. But I am not, by any stretch of the imagination, expressing my willingness to bet simultaneously at 2:1 against heads and at 2:1 against tails. First come, first served, I say.

Furthermore, observe that I am not bound to lose with these degrees of belief. Every bet I make at these crazy odds, on a sequence of occasions, *may* turn out in my favor.

I take it to be a descriptive fact that most reasonable people would regard these odds as odd. I take it to be a normative fact that they *should* be regarded as odd -- i.e., that there is a standard to which degrees of belief should conform that goes beyond conformity to the probability calculus.

Interval-valued epistemic probability provides one normative view according to which some degrees of belief, given the evidence that we have, are prohibited as irrational and unacceptable. It is a gratifying feature of that view of probability that one can show that there exists a classical belief function that conforms to the intervals of epistemic probability[11].

Here again we see a balance between descriptive and normative considerations playing a role in the development of our theory. And while we may and should take descriptive facts concerning the degrees of belief that people have as a clue to the normative standards we wish to embody in our machines and in ourselves, and just as intuition concerning basic principles and intuition concerning particular cases can also play a role, the ultimate test of the cooking is in the pudding.

6. Updating Probabilities.

There are Dutch Book arguments that concern the updating of degrees of belief, just as there are Dutch Book arguments concerning static degrees of belief. Again, there is no question of "inconsistency" in having one's beliefs violate the dynamic condition of updating by temporal conditionalization. But this time it turns out that temporal conditionalization is inconsistent with interval valued epistemic probability as we have defined it.

This provides a real, down to earth, illustration of the theoretical dynamics of normative and descriptive considerations. People tend, other things being equal, to conform to temporal conditionalization. People tend, other things being equal, to base their probability assessments on reference classes that yield relatively sharp probabilities. These two tendencies come into conflict in certain rather special situations, and intuitions conflict about what is the right way to go in these situations.

Here is a case where it seems quite likely that artificial intelligence systems, run on a collection of sensible cases, may enable us to make a philosophical judgment about the best way to

go. It is also a case where, without an initial philosophical guess, there would be no way for an A I system to get off the ground

7. Testing.

We have looked at a number of issues that constitute only a fraction of those that come up in either the philosophical or the computational analysis of cognition. It has been suggested that in some sense we should judge these issues by reference to the performance of systems that embody one or another of the possible sets of commitments one might make.

We've already noticed that this is more complicated than it appears, since short run "success" can't be our standard. What we need as a basis for choice is evidence concerning the long run success of one set of commitments as opposed to another. But this is a matter of defeasible inference, or induction, or probabilistic inference, or probability: that is, it is a matter of the very kind of thing that is at issue.

One possible basis for choice among the principles that satisfy our intuitions and that seem to be reflected in actual human behavior would be the frankly pragmatic basis of content. Assume that there is some plausible way to measure the content of a finite set of sentences (axioms). For very simple languages, it might be the cardinality of the set. For languages with quantitative functions, it might depend on the precision of the assertions embodied in the sentences.

Given such a measure of content of a finite set of sentences, we might compare two logical/linguistic frameworks according to the content of their predictive assertions. This takes probability for granted, and thus may beg the question concerning the choice of probabilities. On the other hand, perhaps it takes for granted only what is common to any plausible view of probability. This is an area to be explored. The object of the present exersize is to indicate that it is an area that should be explored.

A view that has attracted considerable attention in a.i. is the view that we can elucidate defeasible reasoning in terms of *specificity*. In its simplest form the specificity principle directs us to formulate our (defeasible) opinions about an object according to the smallest set of which we know that it is a member. The classical illustration is that of Tweety the penguin: we infer that Tweety does not fly, even though Tweety is a bird, because penguins are a subset of birds, and penguins don't fly.

So far so good. But this does not take care of all cases, as has been widely recognized. The next development of the principle has to do with logical strength: base your defeasible inference on the logically strongest knowledge you have of the object in question.

If this is construed merely as a principle of total evidence it seems uncontroversial, but it does not guide us in the choice of a set of objects to which we should refer the object in question, or anything like that. It provides no concrete guidance. On the other hand, there are two cases about which our intuitions seem to be pretty strong, that do not seem to be accounted for by any kind of specificity[12].

8. Conclusion.

Subject to this last conjecture, we can claim both that logic (and probability, and decision theory) is conventional, and at the same time that there are rational ways of choosing among the available conventions. The choice reflects the interplay among the facts of human cognition and especially human limitations, our intuitions concerning the principles to be found in simple cases, and pragmatic considerations of the sort last mentioned. Probability gives us a handle on these things, but of course probability is one of them itself. Is it more fundamental? Sure. Does that cut much ice? Not much. To the extent that we can limit our considerations to what is common to a number of views of probability, we may find that we have a basis for choice.

What we need is an account of defeasible reasoning -- that is, of a kind of reasoning that will give us practical certainties, but that is not monotonic. It will be a kind of reasoning that allows us to change our minds in response to new evidence

This principle of specificity seems to work well in a limited domain of cases -- namely the cases in which we should be looking at subsets. Construed broadly, as logical strength, the principle seems uncontrovertible, but it also seems devoid of useful content. Construed as something in between the recommendation to choose a subset over a superset as a reference class, and the general heuristic recommendation to use all the evidence you have, it is not clear how to apply the principle.

One lesson we learn is that while the basic distinction between normative and descriptive principles seems as solidly grounded as ever, the principles from among which we must select both in constructing normative philosophical frameworks for knowledge and inference, and in constructing artificial systems for both pure and applied purposes, must include both sorts. We must look to what people do, to what they are capable of doing, and to what they ought to do, in order to gather the materials for epistemic reconstruction and self improvement.

Henry E. Kyburg, Jr. University of Rochester

---

[1]Hume, David: *An Enquiry Concerning Human Understanding Selections From a Treatise...* Open Court, La Salle, 1949.
[2]Cheeseman, Peter: "Inquiry into Computer Understanding," *Computational Intelligence* **4**, 1988, 58-66.
[3]Ramsey, F. P.: *The Foundations of Mathematics and Other Essays*, Humanities Press, New York, 1931.
[4]Jeffrey, Richard C. *The Logic of Decision*, McGraw-Hill, New York, 1965.
[5]Savage, L. J.: *Foundations of Statistics*, John Wiley, New York, 1954.

[6]Edwards, Ward "Subjective Probabilities Inferred From Decisions," *Psychological Review* **69**, 1962, 109-135 Bayes Rationality

[7]Kahneman, Daniel, Slovik, and Tversky, Amos, (Eds) *Judgment Under Uncertainty: Heuristics and Biases* Cambridge University Press, Cambridge, 1982.

[8]Allais, M "The Foundations of a Positive Theory of Choice involving Risk and a Criticism of the Postulates and Axioms of the American School," in M. Allais and O. Hagen, (eds). *Expected Utility Hypotheses and the Allais Paradox*. Reidel, Dordrecht, 1952.

[9]Ellsberg, Daniel. "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 1961, 528-557.

[10]Lopes, Lola "Some Thoughts on the Psychological Concept of Risk," *Journal of Experimental Psychology: Human Perception and Performance* **9**, 1983, 137-144.

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

[12]Kyburg, Henry E. Jr "Beyond Specificity," forthcoming.