

Report Cover Page

ACERA Project		
0607		
Title		
Is Probability the Only Coherent Approach to Uncertainty?		
Author(s) / Address (es)		
Mark Colyvan		
Material Type and Status (Internal draft, Final Technical or Project report, Manuscript, Manual, Software)		
Manuscript		
Summary		
<p>This report examines the claim made by some risk analysts that probabilities are the only coherent basis for representing uncertainty. This claim, if correct, would be a direct challenge to qualitative risk assessments. The study concludes that non-statistical uncertainties are a credible source of uncertainty and that non-probabilistic methods are required to deal with them.</p>		
ACERA Use only	Received By:	Date:
	ACERA / AMSI SAC Approval:	Date:
	DAFF Endorsement: () Yes () No	Date:

Evaluation and development of formal consensus methods;

ACERA Project No. 0607

Chief Investigator; Mark Colyvan

*Department of Philosophy,
University of Sydney,
Sydney, New South Wales 2006, Australia*

Is Probability the Only Coherent Approach to Uncertainty?

Mark Colyvan

Manuscript (final report)

Acknowledgements

This report is a product of the Australian Centre of Excellence for Risk Analysis (ACERA). In preparing this report, the author acknowledges the financial and other support provided by the Department of Agriculture, Fisheries and Forestry (DAFF), the University of Melbourne, Australian Mathematical Sciences Institute (AMSI) and Australian Research Centre for Urban Ecology (ARCUE).

The author would like to thank Vicki Bier, Scott Ferson, Patrick Grim, and Helen Regan for comments on earlier drafts of this article and/or valuable discussions on the material covered in this article. He would also like to thank Applied Biomathematics where the work on this article was carried out, and the participants at a Society for Risk Analysis Workshop in Arlington in February 2000 for the challenging and lively discussion following the presentation of this article.

Disclaimer

This report has been prepared by consultants for the Australian Centre of Excellence for Risk Analysis (ACERA) and the views expressed do not necessarily reflect those of ACERA. ACERA cannot guarantee the accuracy of the report, and does not accept liability for any loss or damage incurred as a result of relying on its accuracy.

Table of contents

Acknowledgements	3
Disclaimer	4
Table of contents	5
Executive Summary	6
Introduction	7
Two Kinds Of Uncertainty	8
Sources Of Uncertainty	10
Cox's Theorem	13
The Explicit And The Implicit Assumptions	15
Discussion	18
References	19

Executive Summary

In this article, I discuss an argument that purports to prove that probability theory is the only sensible means of dealing with uncertainty. I show that this argument can succeed only if some rather controversial assumptions about the nature of uncertainty are accepted. I discuss these assumptions and provide reasons for rejecting them. I also present examples of what I take to be non-probabilistic uncertainty.

Introduction

Uncertainties are ubiquitous in risk analysis and, on the face of it, we must contend with a number of quite distinct sorts of uncertainty. There are many methods on hand to deal with uncertainty, so it is important to select the method best suited to the uncertainty in question. There is, however, a temptation to want to deal with all uncertainty in one fell swoop. That is, it would be desirable to have a single method capable of quantifying all uncertainty. One candidate for this task is probability theory. For such a program to succeed, a demonstration that all uncertainty is probabilistic is required, and a number of people have attempted to defend positions along these lines. In this article, I investigate one way this claim might be supported and I find it inadequate.

My strategy is to first show that the claim that probability theory is the only coherent means of dealing with uncertainty is implausible.¹ I do this by considering different kinds of uncertainty and showing that probability theory seems ill-suited to the uncertainty arising in situations where the logical principle of excluded middle fails. Of particular interest is the uncertainty arising from vagueness. Next, I examine a very interesting technical result of Cox, which, according to some, proves that non-probabilistic methods are inappropriate for the quantification of uncertainty. I pay careful attention to the assumptions of Cox's result, with particular interest in the logical assumptions. I show that this result does not prove that non-probabilistic approaches to the quantification of uncertainty are illegitimate. The reason being that the proof of Cox's theorem employs principles of classical logic for which no justification is offered. Moreover, these logical assumptions are precisely those disputed by many advocates of non-probabilistic approaches to uncertainty. Thus, any appeal to Cox's theorem to "prove" that probability theory is the only coherent approach to uncertainty is circular.

The significance of this to risk analysts is in the take-home message that probability theory does not seem to be able to deal with all uncertainty. When uncertainty arises as a result of vagueness and ambiguity, other methods are required. And, as we shall see, risk analysis is rife with such uncertainty.

TWO KINDS OF UNCERTAINTY

In order to see just how ill-equipped probability theory is for dealing with all uncertainty, we need to reflect a little on uncertainty itself. There are two quite distinct ways in which an agent can be uncertain about the state of a system. The first is where there is uncertainty about some underlying fact of the matter. So, for example, an agent might be uncertain about the levels of some toxin in a water supply. The agent might be in possession of some probabilistic information about the toxin levels—either numerical (“the probability that the toxin levels are above some critical threshold is x ”) or nonnumerical (“it is more likely that the toxin levels are above the threshold”). Now compare this with a second, quite different kind of uncertainty where there is no fact of the matter about whether the system is in the state in question or not. For example, consider uncertainty about the rate of loss of biodiversity. Arguably, there is no fact of the matter about this because of ambiguity in the term “biodiversity.” It all depends on how you measure it (species diversity, higher-taxa diversity, genetic diversity, and so on).⁽¹⁾ In such cases, there is uncertainty, in part, because there is no underlying fact of the matter. In this second case, it appears that an agent cannot be in possession of probabilistic information. After all, what does it mean to say that the probability that biodiversity is being lost at some rate is p , when there is no fact of the matter about what biodiversity is?

Probability theory presupposes that there is an underlying fact of the matter. To see this, consider one of the axioms of probability theory (where, throughout, I use the standard logical notation: \wedge is read as “or,” \neg is read as “not,” and \vee is read as “and”):

$$\Pr(P \vee \neg P) = 1.$$

This implies that the proposition $P \vee \neg P$ is certain, and the possibility that neither P nor $\neg P$ is true is excluded. This axiom of probability theory is the probabilistic analog of the logical principle of excluded middle. It would seem then that in any domain where excluded middle fails, probability theory is an inappropriate tool for representing uncertainty. For example, let us suppose that P is the proposition that Sherlock Holmes walked down Goodge Street exactly seven times in his life. Now, I assume that this proposition is neither proved true nor proved false (either directly or by implication) from the Conan Doyle stories—it is (arguably) a proposition that is neither true nor false, and so (again arguably) $P \vee \neg P$ is not true.² Why? Because there is no Sherlock Holmes to make true or false claims about his walking history beyond what is in the Conan Doyle stories. In fiction, there is nothing more to the story than what is written and its natural implications. Fiction, by its nature, is incomplete,³ so that the law of excluded middle should fail here.

¹ By “coherent” I do not mean consistent (provably consistent) in the logician’s sense of consistent—I simply mean something like: “is not known to lead to contradiction or seriously clash with well-established intuitions.”

² Excluded middle fails in most three-valued logics. One notable exception is supervaluational logic, where, although P is neither true nor false, “ P or not P ” is a theorem (i.e., it is logically true). See Reference 2 for details.

³ It is for considerations not unlike these that led L. E. J. Brouwer and other intuitionists to reject classical logic (with its principle of excluded middle) in favor of intuitionistic logic.⁽³⁾

I give this example just as an illustration. As we shall see, there are other more scientifically relevant domains where excluded middle fails. And in such cases, the applicability of probability theory is highly questionable. The claim that probability theory is the only coherent approach to uncertainty suggests (among other things) that there are no domains about which we reason with uncertainty where excluded middle fails. On the face of it at least, this is false—there are many such domains. Apart from fictional domains and situations of ambiguity (both mentioned above), there are domains employing vague predicates. The latter are widespread and apparently violate excluded middle. At least, if they do not violate excluded middle (see References 4 and 5), an argument as to why they do not is required. I will return to this issue later in this article. For now, I turn to the task of clarifying what is meant by “vagueness”; it is only once we understand this term that we can fully appreciate how unreasonable it is to expect probability theory to be able to deal with vagueness.

SOURCES OF UNCERTAINTY

First Consider Epistemic Uncertainty. This is the uncertainty associated with our knowledge of the state of some system (where the system is in some state or other). It includes uncertainty due to limitations of measurement, insufficient data, extrapolations and interpolations, and variability over time and space. For instance, we might be uncertain about the population size of a given species because the resources are not available for the relevant surveys to be conducted, the population has changed since the last survey, and so on. Epistemic uncertainty is uncertainty about some determinate fact. We are uncertain because we are not in possession of the complete information. The population of the species (at a given time) takes a certain value—it is just that we are uncertain what that value is. Typically, epistemic uncertainty is dealt with by various statistical methods (such as standard probability theory). I should mention that in many applications, it is desirable to distinguish temporal variability (i.e., stochasticity or aleatory uncertainty) from other kinds of uncertainty. By grouping them together here, I do not mean to imply that there is no interesting distinction. It is just that the distinction that I am most interested in is the distinction between epistemic uncertainty and vagueness (see Reference 6 for discussion of the different types of epistemic uncertainty).

Vagueness, by contrast, is a source of uncertainty that is (arguably) linguistic in origin. Here, the uncertainty arises out of vagueness in the language, in particular, from vague predicates. A vague predicate is one that permits borderline cases. So, for example, the predicate “is a mature individual” is vague because it permits borderline cases (such as adolescents, which are borderline between adults and nonadults). It is not generally thought that there is an unknown fact of the matter about whether an adolescent is an adult or not. Rather, there is no fact about our linguistic practices that determines whether an adolescent is an adult or not. It follows that there can be no fact of the matter, for, after all, there are no other facts but the linguistic facts here (cf. Reference 4). There are various methods for dealing with vagueness. The most obvious of which is to eliminate it by more careful attention to our language. There are many problems with this approach, not least of which is that despite the claims of some, it is extremely difficult, if not impossible, to accomplish. For example, Morgan and Henrion, in their book on uncertainty,⁽⁶⁾ devote only a page and a half to vagueness (they also seem to confuse it with other kinds of linguistic uncertainty). They claim that it is “usually relatively easy to remove with a bit of clear thinking” (1990, p. 62). If it were so easy to remove, you would expect them to be able to state this thesis without appeal to at least four vague terms.

Other problems with such an approach revolve around the arbitrariness of the demarcations required. For example, if we decide to regiment the vague term “endangered” (as applied to animal species) to mean “less than n individuals,” we have to admit that n is arbitrary—why not $n + 1$ or $n - 1$? Worse still, this arbitrary precision can lead to problems when populations are close to n . Let us suppose that the population of some species is $n + 1$. Furthermore, suppose that government funding is available for the implementation of various conservation management strategies (captive breeding programs, the establishment of wildlife reserves, and

the like) for all endangered species. But, since the species in question is not endangered according to our newly-defined use of the word “endangered,” the species is not able to attract government funding. It might thus turn out that the best course of action for those interested in preserving this species is to see to it that a couple of individuals disappear or die!

These and other problems (see Reference 7 for more on this) have led most commentators on vagueness to prefer other more sophisticated approaches. These approaches include fuzzy logic (and fuzzy set theory),^(8,9) supervaluations,^(2,5,10) intuitionistic logic,⁽¹¹⁾ three-valued logic,⁽¹²⁾ paraconsistent logic,^(13,14) and modal logic (pp. 270–275).^(3,15) Some authors use the term “fuzziness” to describe what I am calling “vagueness.” I follow Sorensen,⁽¹⁶⁾ Williamson,⁽⁴⁾ Read,⁽¹⁷⁾ and others and use the term “vagueness” because, unlike “fuzziness,” it does not prejudice the question of what the best method of dealing with the phenomenon is.

It is worth noting here that vagueness can give rise to uncertainty in the following way. Let us suppose that we wish to know how many mature individuals there are in a given population. We will be uncertain about this figure because of the vagueness of the predicate “is a mature individual.” You may also be uncertain about this figure for other reasons due to the dynamics of the population or due to the limitations in our estimation methods, for instance. In short, you may have to deal with epistemic uncertainty and uncertainty arising from vagueness.

It is important to distinguish vagueness from other kinds of uncertainty that also have linguistic origins. These include *context dependence*, *ambiguity*, and *underspecificity*. The first, *context dependence*, is uncertainty arising from failing to specify the context in which a proposition is to be understood. For example, suppose that it is said that some person is tall. Without specifying the context, the audience is left wondering whether the person in question is tall for a jockey, tall for a basketball player, or tall in some other unspecified context. Note that “tall” is also vague, but the vagueness and the context dependence are quite separate issues. The vagueness persists after the context has been fixed. That is, even after we are told that we are concerned with the predicate, say “tall-for-a-jockey,” we are still faced with borderline cases (see Reference 18 for a very interesting discussion on how context dependence is not well handled by standard fuzzy methods).

Ambiguity is uncertainty arising from the fact that a word can be used in more than one way, and in a given context, it is not clear which way it is being used. For example, the word “bank” is ambiguous (in some contexts) between a financial institution and the edge of a river. Notice that once again this source of uncertainty is quite distinct from vagueness (despite commonly being confused with vagueness). The ambiguity in the word “bank” does not give rise to borderline cases in the way “mature individual” does—there is nothing that is borderline between a financial institution and the edge of a river.

The final linguistic source of uncertainty that I will mention here is *underspecificity*. This is where there is unwanted generality; the statement in question does not

provide the degree of specificity we desire. For example, the statement that there will be rainy days ahead is underspecific because we are left wondering: Which days will be rainy? How many of them will be rainy? And so on. (Of course “rainy days” is also vague—does a day with light mist counts as a rainy day? But again such issues are quite distinct from the underspecificity issues.) The uncertainty associated with such statements is clearly distinct from that associated with borderline cases. Unfortunately, the word “vague” is commonly used for both underspecificity and borderline cases. I follow Sorensen⁽¹⁶⁾ here and suggest that the term “vagueness” be reserved for the borderline-case sense.

In many situations of interest to risk analysts, several sources of uncertainty—both epistemic and linguistic—are to be found, typically giving rise to difficult assessments of the compound uncertainties in question. Vagueness is an important, but under-appreciated, contributor to this mix. One only needs to reflect on some of the key terms used in various risk assessments to see the prevalence of vagueness: “mature individual,” “close to critical levels,” “negligible chance,” “endangered,” “tested under a wide range of temperatures,” and so forth. These terms are all vague and may give rise to uncertainty as a result of their vagueness (see References 19 and 20 for further discussion of the importance of vagueness for conservation and environmental risk analyses).

For the rest of this article, I will focus on vagueness (in the borderline-case sense) as my primary example of linguistic uncertainty. I mention ambiguity, context dependence, and underspecificity just to distinguish these from vagueness and to make it clear that I am not including these quite distinct linguistic sources of uncertainty when I speak of vagueness.⁴ Now I move to the main topic of this article: arguments that purport to prove that all uncertainty admits a probabilistic treatment.

⁴ See Reference 19 for details of the treatment of other forms of linguistic uncertainty.

COX'S THEOREM

There are a number of arguments used in support of the claim that probability is the only coherent approach to uncertainty,^(21–24) of which Cox's theorem⁽²⁵⁾ is undoubtedly the most well known.

THEOREM (COX): *Any measure of belief is isomorphic to a probability measure.*

The theorem is explicitly based on the following assumptions: (i) an agent's belief in $\neg P$ is a function of his/her belief in P and (ii) an agent's belief in $P \wedge Q$ is a function of the agent's belief of P given Q and the agent's belief of Q .

Let us consider what this theorem might tell us about belief and uncertainty. There are a number of possible messages to glean from the theorem. The following are some of the more interesting ones.

- (1) Bayesian subjective probabilities are legitimate means of quantifying (epistemic) uncertainty.
- (2) Bayesian subjective probabilities are the *only* legitimate means of quantifying (epistemic) uncertainty.
- (3) Probabilities are the *only* legitimate means of quantifying (epistemic) uncertainty.
- (4) Probabilities are the *only* legitimate means of quantifying uncertainty.

I think it is fair to say that the primary purpose of Cox-like results is to provide a justification of Bayesian (subjective) probability.⁽²⁶⁾ Clearly, if Bayesian methods have a place at all in the treatment of uncertainty, we require an argument for (1) and the more evangelical Bayesians might prefer an argument for (2). Cox's theorem and its kin aim to establish (1), since if degrees of belief turn out to be isomorphic to a probability measure (as the theorem suggests), then degrees of belief are a legitimate interpretation of this calculus. It is much less clear how the theorem is supposed to support the stronger conclusion (2), but let us put that matter aside and consider how, if one is not careful, one might slide from (2) to (4). First, note that if (2) is true, then the phrase "Bayesian subjective" in (2) is, in some sense, unnecessary. After all, if Bayesian subjective probabilities are the only legitimate means of quantifying uncertainty, then it might be argued that these are the only legitimate interpretations of the probability calculus. So, on the assumption of (2), it is not implausible to also hold (3). Next, note that if one were to hold the view that all uncertainty is epistemic uncertainty, then (4) follows straightforwardly from (3). It is clear, however, that (4) is a more radical claim than even (2), and that (4) does not follow from (1) or (2) without additional argument and some fairly controversial premises—such as that all uncertainty is epistemic uncertainty.

It is well known that fuzzy approaches to uncertainty, such as possibility theory,⁽²⁷⁾ are not isomorphic to probability theory.^(28–32) Thus, it has been suggested that Cox's theorem tells us that possibility theory, for instance, is not viable as a means of quantifying uncertainty (e.g., Reference 21). Indeed, if we accept this line of thought, it would seem that the whole literature on non-probabilistic means of quantifying uncertainty (e.g., References 7, 28, 33–35) is misguided. Others insist that there is

something wrong with Cox's theorem. The obvious place to look for what is wrong is in the assumptions. There has been a great deal of discussion on the assumptions of the Cox's theorem,⁽³⁶⁻⁴⁰⁾ but there has been little attention paid to the logical assumptions. I will examine these and show that it is not just possibility theory that apparently flies in the face of Cox's theorem.

THE EXPLICIT AND THE IMPLICIT ASSUMPTIONS

Recall the explicit assumptions of Cox's theorem:

- (i) An agent's belief in $\neg P$ is a function of his/her belief in P .
- (ii) An agent's belief in $P \wedge Q$ is a function of the agent's belief of P given Q and the agent's belief of Q .

There are also some implicit assumptions, namely:

- (iii) Belief is a real-valued function.
- (iv) The underlying logic is classical propositional calculus.
- (v) The following substitution principle holds: if an agent believes P to some degree, b , and

P is logically equivalent to Q , then the agent believes Q to degree b .

- (vi) The function in (ii) is twice differentiable.

In the discussion on Cox's theorem, assumptions (iv) and (v) are rarely mentioned, let alone questioned. These two assumptions, however, are very important in the proof of Cox's theorem. For example, in the proof that the belief in $P \vee \neg P$ is maximal, Cox invokes them both, but let us focus on the use of (iv). In particular, Cox assumes the classical principle of double-negation elimination: $\neg\neg P$ is equivalent to P . The question is whether this is a reasonable assumption. Unfortunately, the answer is "no"—at least not when considering uncertainty in the broadest sense (including vagueness). For example, let us suppose that some object is borderline red. That is, the object in question is neither red nor not red. So, it is not the case that the object is not red, but it does not follow that the object is red. In settings where vagueness is present, there is good reason to believe that double-negation elimination fails. It fails for the same reasons for which excluded middle fails.

In fuzzy logic, double-negation elimination holds, so the fuzzy logician cannot object to this particular step in Cox's proof. There is, however, another step that the fuzzy logician should object to: the use of assumption (vi). The problem here is that (at least the most common variety of) fuzzy logic employs maximum and minimum functions for disjunction and conjunction, respectively. These functions are not twice differentiable. So, we see that implicit assumption (vi) of Cox's theorem rules out such fuzzy approaches from the start.⁽³⁶⁾

We might also question the substitution principle (v) Cox implicitly invokes, for this principle implies that all agents are logically omniscient. Cox invokes this assumption throughout the proof of the theorem. For example, on page 7, he assumes that an agent's belief in $\neg\neg b$ given a is the same as an agent's belief in b given a . This move assumes double-negation elimination—that $\neg\neg b$ is equivalent to b —and that the agent knows that $\neg\neg b$ is equivalent to b . Cox justifies the move in question merely by appeal to double-negation elimination. But also in need of defense here is the substitution principle implying that all agents have access to logical equivalences. The problem is that you may not know about the equivalence, but the substitution principle (v) does not allow this. It is clear that this is a nontrivial and by no means an

uncontroversial assumption. This is another avenue one might pursue in criticizing Cox's theorem.

It should be clear that it is not just possibility theory that runs into problems with Cox's theorem; any non-classical treatment of vagueness (and that is most of them) is in the same boat. Moreover, the reason for this is that non-classical approaches are ruled out from the start. This does nothing to shake our confidence in these approaches though, unless the assumptions in question—(iv) and (vi)—can be defended. The point is that anyone wishing to argue against fuzzy methods, for instance, must not simply beg the question against those methods. To use a theorem such as Cox's theorem that is based on assumptions that fuzzy logic rejects is, thus, inappropriate, unless the assumptions in question can be independently defended. But then, the appeal to Cox's theorem is redundant. This is not to say that Cox's theorem is of no interest, or that in all contexts, the assumptions of Cox's theorem are controversial. Indeed, if one is using the theorem to justify a Bayesian interpretation of the probability calculus, the assumptions (i)–(vi) may well be entirely reasonable, and the isomorphism between belief measures and probability measures is extremely interesting. (It is, perhaps, worth noting that assumption (iii) rules out two-dimensional approaches to belief, such as Dempster-Shafer belief functions⁽⁴¹⁾ and imprecise probabilities⁽⁴²⁾. Thus, Cox's theorem cannot be wielded against these methods either.) The case with second-order probability is less clear perhaps, but I take it that (iii) also rules out second-order probability, where belief is represented by a real-valued function whose value is uncertain.

Although Cox did not provide any justification for the two assumptions I have been discussing here ((iv) and (vi)), this does not mean that no such justification can be provided. So, let us now consider how the problematic assumptions might be defended. How might double-negation elimination (or, more generally, classical logic) be defended? One way is to argue that, despite appearances, there is a fact of the matter about the applicability of vague terms. So, for example, there is a fact of the matter about whether a 178-cm man is tall. On this view, uncertainty due to vagueness is seen to be nothing more than a particularly resilient kind of epistemic uncertainty. According to this view, there are facts of the matter about the applicability of vague predicates (so it is appropriate to employ classical logic), but competent language users do not (and, in fact, *cannot*) know what these facts are. A 178cm man is, thus, either tall or not, but it is unknowable which he is. I will not discuss this position further here (see Reference 4 for a defense of the view). I simply note that this rather counterintuitive view is one way that the use of classical logic might be justified. It is clear, however, that the use of classical logic in this setting is no trivial assumption.

The defense of assumption (vi) is problematic as well. The main difficulty is that Cox saw himself as deriving the theorem from plausible, intuitive, and fairly uncontroversial assumptions. Assumption (vi) is none of these. Indeed, this seems to be widely recognized, for there has been considerable interest in deriving Cox-style results without appeal to (vi) (or, more precisely, by replacing (vi) with other weaker assumptions). While versions of Cox's theorem have been derived without appeal to assumption (vi) (e.g., Reference 45), these results still rely on the same

implicit logical assumptions (see, for example, References 21, 39, 43–45). In each case, it is the implicit assumption that the underlying logic is classical that rules against most non-classical approaches to uncertainty.⁵ What is not generally realized is that the assumption of classical logic requires justification if the theorems in question are to be wielded as weapons against non-classical methods of dealing with uncertainty. As they stand, arguments employing these theorems against non-classical approaches to uncertainty are simply circular.

⁵ For instance, Paris⁽⁴⁵⁾ drops assumption (vi), but fuzzy approaches are still ruled out by the (implicit) assumption that excluded middle holds.

DISCUSSION

While it is well known that fuzzy methods such as possibility theory are not isomorphic to probability theory, and that this seems to fly in the face of Cox's theorem, it has not been generally recognized why this is so. My goal in this article has been to correct this situation. I have, thus, drawn attention to some of the logical assumptions of Cox's theorem and pointed out that these assumptions are *prima facie* implausible, at least in the context of arguments to the conclusion that all uncertainty is probabilistic. I conclude by urging the defenders of such arguments against non-probabilistic methods to provide some justification for the appeal to classical logic in the face of vagueness. In particular, a justification for the appeal to the classical principles of double-negation elimination and excluded middle are required.

I wish to make it perfectly clear that I have said nothing about the application of Cox's theorem as a justification for the use of probability for epistemic uncertainty. My focus has been on the claim that Cox's theorem shows that probability theory and only probability theory can do it all—*across all domains*. In particular, since Cox himself made no such radical claims, my target is not Cox. It is quite clear that Cox intended his argument to apply only to classical domains. The target of this article is those who would use Cox's result to dismiss non-probabilistic approaches to uncertainty *in non-classical domains*.

Although my goal is modest, it should also be clear that I have considerable sympathy with a somewhat stronger thesis: no adequate defense of classical logic in domains employing vague predicates is possible. If I am right about this, then not only are non-probabilistic methods legitimate methods for quantifying at least some types of uncertainty, but are also *required* for the adequate treatment of uncertainty in any domain where vague predicates are used.

References

1. Maclaurin, J., & Sterelny, K. (2008). *What is Biodiversity?* Chicago: Chicago University Press.
2. Van Fraassen, B. C. (1966). Singular terms truth value gaps and free logic. *Journal of Philosophy*, 63, 481–495.
3. Kline, M. (1972). *Mathematical Thought from Ancient to Modern Times*. New York: Oxford University Press.
4. Williamson, T. (1994). *Vagueness*. London: Routledge.
5. Fine, K. (1975). Vagueness truth and logic. *Synthese*, 30, 265–300.
6. Morgan, M. G., & Henrion, M. (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press.
7. Regan, H. M., Colyvan, M., & Burgman, M. A. (2000). A proposal for fuzzy IUCN categories and criteria. *Biological Conservation*, 92, 101–108.
8. Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8, 338–353.
9. Zimmermann, H.-J. (1996). *Fuzzy Set Theory—And Its Applications*, 3rd ed. Dordrecht: Kluwer.
10. Dummett, M. (1992). Wang's paradox. Reprinted in M. Dummett (Ed.), *Truth and Other Enigmas*, 2nd impression (pp. 248–268). London: Duckworth.
11. Putnam, H. (1983). Vagueness and alternative logic. *Erkenntnis*, 19, 297–314.
12. Korner, S. (1955). *Conceptual Thinking*. Cambridge: Cambridge University Press.
13. Hyde, D. (1997). From heaps and gaps to heaps of gluts. *Mind*, 106, 641–660.
14. Beall, J. C., & Colyvan, M. (2001). Heaps of gluts and Hyde-ing the sorites. *Mind*, 110, 401–408.
15. Hughes, G. E., & Cresswell, M. J. (1996). *A New Introduction to Modal Logic*. London: Routledge.
16. Sorensen, R. A. (1989). The ambiguity of vagueness and precision. *Pacific Philosophical Quarterly*, 70, 174–183.
17. Read, S. (1995). *Thinking About Logic: An Introduction to Philosophy of Logic*. Oxford: Oxford University Press.
18. Bier, V. (1992). Fuzzy set theory probability theory and truth functionality. In B. M. Ayyub, M. M. Gupta, & L. N. Kanal (Eds.), *Analysis and Management of Uncertainty: Theory and Application* (pp. 65–78). New York: North-Holland.
19. Regan, H. M., Colyvan, M., & Burgman, M. A. (2002). A taxonomy and treatment of uncertainty for ecology and conservation biology. *Ecological Applications*, 12(2), 618–628.
20. Burgman, M. A. (2005). *Risk and Decisions for Conservation and Environmental Management*. Cambridge: Cambridge University Press.
21. Lindley, D. V. (1982). Scoring rules and the inevitability of probability. *International Statistical Review*, 50, 1–26.
22. Fine, T. L. (1973). *Theories of Probability*. New York: Academic Press.
23. Jaynes, E. T. (1988). How does the brain do plausible reasoning? In G. J. Erickson & C. R. Smith (Eds.), *Maximum Entropy and Bayesian Methods in Science and Engineering* (pp. 1–24). Dordrecht: Kluwer.
24. Aczel, J. (1966). *Lectures on Functional Equations and Their Applications*. New York: Academic Press.
25. Cox, R. T. (1946). Probability, frequency and reasonable expectation. *American Journal of Physics*, 14, 1–13.
26. Cheeseman, P. (1988). An inquiry into computer understanding. *Computational Intelligence*, 4, 58–66.
27. Dubois, D., & Prade, H. (1988). *Possibility Theory*. London: Plenum.
28. Smets, P. (1998). Numerical representation of uncertainty. In D. Dubois & H. Prade (Eds.), *Belief Change*. Volume 3 of D. M. Gabbay & P. Smets (Eds.), *Handbook of*

- Defeasible Reasoning and Uncertainty Management Systems* (pp. 265– 309). Dordrecht: Kluwer.
29. Zadeh, L. A. (1980). Fuzzy sets versus probability. *Proceedings of I.E.E.E.*, 68, 421.
 30. Zadeh, L. A. (1982). Discussion of paper by D. V. Lindley. *International Statistical Review*, 50, 23–24.
 31. Regan, H. M., & Colyvan, M. (2000). Fuzzy sets and threatened species classification. *Conservation Biology*, 14, 1197– 1199.
 32. Colyvan, M. (2001). Ranking in threatened species classification. *Conservation Biology*, 15, 1814–1816.
 33. Smets, P. (Ed.) (1998). *Quantified Representation of Uncertainty and Imprecision*. Volume 1 of D. M. Gabbay & P. Smets (Eds.), *Handbook of Defeasible Reasoning and Uncertainty Management Systems*. Dordrecht: Kluwer.
 34. Zadeh, L. A., & Kacprzyk, J. (Eds.). (1992). *Fuzzy Logic for the Management of Uncertainty*. New York: John Wiley and Sons.
 35. Regan, H. M., Colyvan, M., & Burgman, M. A. (1999). Dealing with vagueness in threatened species classification. In R. N. Dave & T. Sudkamp (Eds.), *Proceedings of the 18th International Conference of the North American Fuzzy Information Processing Society: Real World Applications of Fuzzy Logic and Soft Computing* (pp. 685–689). Piscataway, NJ: I.E.E.E.
 36. Colyvan, M. (2004). The philosophical significance of Cox’s theorem. *International Journal of Approximate Reasoning*, 37, 71–85.
 37. Halpern, J. Y. (1999). A counterexample to theorems of Cox and Fine. *Journal of Artificial Intelligence Research*, 10, 67-85.
 38. Halpern, J. Y. (1999). Cox’s theorem revisited. *Journal of Artificial Intelligence Research*, 11, 429–435.
 39. Van Horn, K. S. (2003). Constructing a logic of plausible inference: A guide to Cox’s theorem. *International Journal of Approximate Reasoning*, 34, 3–24.
 40. Dubois, D., & Prade, H. (1990). The logical view of conditioning and its application to possibility and evidence theories. *International Journal of Approximate Reasoning*, 4, 23–46.
 41. Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton University Press.
 42. Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
 43. Heckerman, D. (1988). An axiomatic framework for belief updates. In J. F. Lemmer & L. Kanal (Eds.), *Uncertainty in Artificial Intelligence 2* (pp. 11–22). Amsterdam: North-Holland.
 44. Aleliunas, R. (1990). A summary of a new normative theory of probabilistic logic. In R. Shachter, T. Levitt, L. Kanal, & J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 4* (pp. 8–14). New York: North-Holland.
 45. Paris, J. B. (1994). *The Uncertain Reasoner’s Companion: A Mathematical Perspective*. Cambridge: Cambridge University Press.