The Philosophical Significance of Cox's Theorem

Mark Colyvan
Department of Philosophy, University of Queensland,
Brisbane, Queensland, 4072, Australia.[*]

Abstract

Cox's theorem states that, under certain assumptions, any measure of belief is isomorphic to a probability measure. This theorem, although intended as a justification of the subjectivist interpretation of probability theory, is sometimes presented as an argument for more controversial theses. Of particular interest is the thesis that the only coherent means of representing uncertainty is via the probability calculus. In this paper I examine the logical assumptions of Cox's theorem and I show how these impinge on the philosophical conclusions thought to be supported by the theorem. I show that the more controversial thesis is not supported by Cox's theorem.

Running Head: Cox's Theorem

Key Words: R.T. Cox; belief; uncertainty; non-classical logic; probability; excluded middle

[*] *Email:* mcolyvan@uq.edu.au.

1 Introduction

Paul Benacerraf [6] once warned that when philosophical conclusions are argued from formal mathematical results, one should look very carefully at the assumptions of the arguments in question. For any such argument cannot rest on the formal result alone; there must be some philosophical premise, and this is often illicitly smuggled through the back door. Benacerraf is not suggesting that one can never draw philosophical conclusions from formal results, or that all such arguments are flawed—just that it is important to identify the often suppressed philosophical premises and to assess their plausibility. I think this is very good advice and with this advice in mind I wish to examine the formal result known as Cox's theorem. This theorem states that, under the assumptions of the theorem, any measure of belief is isomorphic to a probability measure [9, 10]. The theorem has been used to support a variety of philosophical conclusions, ranging from a justification of the Bayesian approach to probability, to a more radical thesis that probability is the only coherent representation of uncertainty. In particular, I will examine the logical underpinnings of the theorem—classical propositional calculus—and show that, in certain contexts at least, these logical assumptions are hard to defend. This, in turn, undermines the more radical philosophical theses that the theorem might be thought to support. I begin by discussing a kind of uncertainty for which classical logic is inappropriate.

2 Belief and Non-Epistemic Uncertainty

Agents typically do not believe propositions to degree one or zero. Belief comes in degrees. This is because there is typically uncertainty about the truth value of the proposition in question. Good epistemic agent recognise this and set about quantifying the extent of the uncertainty and/or their degree of certainty. Providing the details of a representation of reasoning carried out by human (or more commonly, ideal) agents operating under uncertainty is often referred to as the project of delivering the logic of plausible inference.[1] It is usually assumed that uncertainty arises because of incomplete information—it is simply an epistemic matter. I will argue that this is not the case. Some uncertainty may remain even when the agent is in possession of all the relevant data. This is bad news for classical logic and classical probability theory.

There are two ways in which an agent can be uncertain about the state of a system. The first is familiar. This is where there is uncertainty about some underlying fact of the matter: System S is either in state σ or it is not, but

[1] I share Glenn Shafer's [48] concerns about the use of the term 'plausible' here, but this term is well entrenched in the literature, and since I can think of no better term, I'll continue to use it. I stress however, that I am using the term more broadly than is usual. I include any formal account of belief and reasoning under uncertainty.

agent A does not know which. A might be in possession of some probabilistic information about the state of S—either numerical ("the probability that S is in state σ is x") or non-numerical ("it's more likely that S is in state σ than not"). Call this *epistemic uncertainty*. Now compare this with a second, quite different kind of uncertainty; uncertainty where there is no fact of the matter about whether system S' is in state σ' or not. Indeed, here the uncertainty arises *because* there is no underlying fact of the matter. Call this second kind of uncertainty *non-epistemic uncertainty*. The idea here is that, for reasons I'll discuss shortly, the system S' is neither in state σ' nor not in state σ'—S' is not in a determinate state with respect to σ'. It follows that even an agent in possession of all the relevant data will be uncertain as to the truth value of the proposition 'S' is in state σ''. (Or, equivalently, the agent will not know the answer to the question 'Is S' in state σ'?')

It follows that if there are any instances of non-epistemic uncertainty, an agent could not be in possession of probabilistic information in such cases. After all, what would it mean to say that the probability that system S' is in state σ' is x when there is no fact of the matter about the state of S'? Classical probability theory presupposes that there is an underlying fact of the matter. To see this we need only consider one of the axioms of classical probability theory:

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

This implies that the proposition $Q \vee \neg Q$ is certain (because it is a logical truth). This axiom of probability theory is the probabilistic analogue of the logical principle of excluded middle. It would thus seem that in any domain where excluded middle fails, (classical) probability theory is an inappropriate tool for representing uncertainty.[2]

Now there are several candidates for such domains, none of which, admittedly, are entirely uncontroversial. To start with, consider fictional discourse. In a work of fiction such as H.G. Wells' The Time Machine there is nothing more to the story than what is written (and perhaps the logical and natural implications of what is written). There is no fact of the matter about details not in the story. So, for example, in the 1960 movie of the novel the time traveller sets off for the future taking with him three books. What were the three books? Well that's (quite deliberately) not part of the story so (plausibly) there's no fact of the matter about what the three books were. It seems that classical logic—in particular excluded middle—fails here. It is not true that either the time traveller took or did not take Descartes' meditations with him. Moreover, the probability of this disjunction is not one (as standard probability theory insists). Indeed, it seems quite misguided to talk of probabilities at all in such cases.[3] I should

[2]See [7], [13] and [14] for more on this issue.

[3]Some might insist that the question about what the books were is meaningless, but

add that this example is not as irrelevant to science as it might at first seem. Science makes wide use of fictional entities (like incompressible fluids, and Turing machines) and others that turn out to be fictional (like the planet Vulcan, which was supposed to have an orbit inside Mercury's). And it is clear that there are true propositions about such fictions—'the halting problem is unsolvable', for instance. So it won't do to dismiss fictional discourse as a mere philosophical curiosity.[4]

Another example of a domain in which excluded middle might be thought to fail is mathematics. Consider the status of Goldbach's conjecture: all even numbers greater than two can be written as the sum of two primes. At present this conjecture has not been proven nor has its negation been proven. Now let's suppose that you are a constructivist about mathematics. That is, you believe that 'P is true' is just to say that P has a constructively respectable proof from some constructively respectable set of axioms. It is well known that such constructivists embrace intuitionistic logic where both double negation elimination and excluded middle fail [26]. But even non-constructivists may accept that there are no-fact-of-the-matter propositions in mathematics. Take, for example, an independent question of set theory such as the continuum hypothesis. Neither this nor its negation is provable from the standard ZFC axioms—it is provably independent of those axioms. Many non-constructivists (for example, mathematical fictionalists like Hartry Field [12]) also believe that at least some independent statements are neither true nor false (and so it is not true that such an independent statement or its negation holds).

A third example of where excluded middle might be thought to fail is in domains where vague predicates are employed. Let's suppose that we wish to know how many young people there are in a crowd. We might be uncertain about this because there are some borderline cases. Take, for example, someone who is in their late 20s. Do we count such a person as young or not? There seems no definitive way to answer this question. The problem is that the word 'young' is vague (in the sense that it permits borderline cases).[5] There are a some well-known approaches to vagueness according to which excluded middle holds—for example, Williamson's [55]

this is very hard to sustain. There is nothing ungrammatical about the sentence and the meaning is perfectly clear. On what grounds is the case for the sentence's meaningless to be based? I can think of none. Indeed, the reason that some are inclined to call such questions meaningless is because they *do* understand the meaning, see what the implications are, and only then deny that it has meaning.

[4]See [20, pp. 70–73], [21, chap. 7] and [40, pp. 128–131] for more on the logic and semantics of non-denoting fictional terms. Fictional discourse also raises problems at the level of predicate logic. In classical predicate logic all names refer—even names like 'Vulcan'. Another deviation from classical logic motivated by such considerations is *free logic* where "empty" names are permitted. See [17], [18], [35] and [43].

[5]It's also context sensitive. But let's put that aside; let's assume that the context is fixed. I should also mention that vagueness is rather widespread in both natural language and in science so it is unreasonable to dismiss it as another philosophical curiosity. See

epistemic account of vagueness, the supervaluational account ([15]; [51]) and the paraconsistent approach ([4]; [28]; [29]). Still, rejecting excluded middle remains a very plausible strategy.[6]

Indeed, those who would like to apply probability theory to domains with vague predicates should take little comfort from the above excluded-middle-preserving approaches. For example, on what is generally thought to be the leading contender among these approaches—the supervaluational account—probabilities are still out of place. Although $P \vee \neg P$ is a theorem, if P is borderline, P is usually thought to be neither true nor false. In either case, it seems to make little sense to speak of the probability of P being true (when P is borderline). On the paraconsistent approach, excluded middle is preserved at the expense of (one sense of) the law of non-contradiction. That is, borderline statements (such as 'a 28 year old is young') are seen as both true and false. That is, we have some true instances of $P \wedge \neg P$. Those who find giving up excluded middle objectionable are unlikely to be happy with this. Williamson's epistemic approach (according to which there is an unknowable fact of the matter concerning borderline cases) is the only option that would seem palatable to defenders of the view that probability theory is appropriate in such domains. The problem is that Williamson's view is deeply unintuitive and finds few supporters because of this. It would be inappropriate to try to settle the matter of the correct account of vagueness here; I simply mention vagueness as another very plausible source of non-epistemic uncertainty.

It is worth pausing for a moment here to emphasise how vagueness gives rise to uncertainty. Consider a scientific question such as 'how many species are there in a given eco-system?' Obviously there will be epistemic uncertainty associated with this question but let us suppose that an agent is in possession of all the relevant data. It turns out that even in possession of all the data, the answer to the question may remain out of reach because of the vagueness of the scientific terms 'eco-system' and 'species'. The boundary of a eco-system will always admit borderline cases. Less obvious, perhaps, is that the term 'species' is vague. Consider the possibility of a speciation event occurring at the moment that the question is asked. Do we count the species in question as one or two? It is also worth stressing that no further information can be brought to light that will settle the matter. Perhaps we must settle for upper and lower bounds as the answer to the question. In the example of a speciation event taking place, we might give the interval $[n, n+1]$ as our answer. So we see that vagueness can give rise to this peculiar kind of uncertainty—an uncertainty that cannot be eliminated by gathering further data.

---

[44] and [45] for some of the problems arising from vagueness in ecology and conservation biology.

[6] See, for example, [19], [22], and [34] and [42] for some of the approaches that abandon excluded middle.

Now it might be argued that since the uncertainty in question here is uncertainty about the truth value of a vague proposition, we can state the problem classically in the meta-language. We can say that we do not know whether P (for some vague proposition P) is true. Let v be the valuation function (which maps from the domain of discourse D to the truth value set TV), then the problem is that of determining whether $v(P) = a$, where a is a particular truth value in TV. But, so the argument goes, '$v(P) = a$' is either true or false and so we have forged a link between non-epistemic uncertainty at the object-language level and epistemic uncertainty at the meta-language level. Indeed some do opt for a classical metalogic, but there is a case to be made for non-classicality all the way up. A non-classical metalogic would be called for, for instance, if there is higher-order vagueness. An adequate discussion of this would take us too far afield; I mention it merely to make the point that non-epistemic does not reduce to epistemic uncertainty in any straight-forward fashion.[7]

So far I have argued that in domains where excluded middle fails, the applicability of probability theory is highly questionable. The claim that classical probability theory is the only coherent representation of 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: there are fictional domains, constructive domains and domains with vague predicates. Thus any defender of classical logic needs to convince us that classical logic can, despite appearances, cope with these problematic domains. This is a large (if not impossible) task, for it involves, among other things, providing a classical account of fictional discourse, a defence of certain philosophical views about the philosophy of mathematics (perhaps defending platonism) and a defence of something like Williamson's epistemic approach to vagueness.[8]

Before I move on to a discussion of Cox's theorem, let's consider a couple of objections to my conclusion that probability theory is not appropriate for non-epistemic uncertainty. The first objection comes from quantum mechanics. According to the Copenhagen interpretation of quantum mechanics, there is no fact of the matter about the state of certain quantum systems before a measurement is made. But quantum theory itself provides us with

[7] See [55] for a discussion of higher-order vagueness.

[8] Worse still, there would seem to be inconsistent domains about which we reason. I have in mind here inconsistent mathematical theories (such as the early calculus and naïve set theory) and inconsistent scientific theories (such as the conjunction of general relativity and quantum mechanics). Classical logic and classical probability theory are inappropriate in such domains since in classical logic everything follows from a contradiction and in classical probability theory, all contradictions have probability zero and all conditional probabilities conditional on a contradiction are undefined. Again if we reason about such domains, as we surely do, then it's clear that the classical theories are inadequate. See [38] for some recent papers on inconsistency in science and [37] for an account of a paraconsistent belief revision theory.

probabilities about the state of the system in question (see [27]), so it seems that we have a counterexample to the claim that we can't use probability theory unless there's an underlying fact of the matter. The problem with this objection, however, is that it confuses what the quantum mechanical probabilities are about. The quantum mechanical probabilities are *not* about the state of the quantum system in question *before measurement*; rather, the probabilities are usually construed to be about the state of the system *were it to be measured* (or, if you prefer, they might be construed to be about the measurements themselves—the probability of the measurement turning out a particular way). Either way, the probabilities are not construed as being about systems in indeterminate states.

The next objection concerns denotational failure. According to some (e.g., Strawson [49]), when there is failure of denotation, there is no fact of the matter about the truth of the offending sentence (i.e., the offending sentence is truth-valueless). Let's, for the sake of argument, accept this view. (Indeed, I've already entertained fictional discourse—which is one special kind of denotational failure—as a source of non-epistemic uncertainty.) Suppose you see a male colleague, whom rumour would have it was supposed to be having marital problems, looking rather depressed and you speculate that his wife has left him. You might even believe that this is the most likely explanation for his depressed state. That is, you assign a subjective probability of greater than 0.5 to the truth of the proposition 'My colleague's wife has left him'. Now, as it turns out, your colleague is not, nor has he ever been, married. We thus have a case of denotational failure and so, by hypothesis, the sentence in question does not take a truth value. But, it still seems sensible to attribute a probability of truth to the sentence in question.[9] I agree that it seems sensible to entertain a probability of truth for the sentence in question, but it's not clear that it's sensible to do this *on the view under consideration*. After all, it also seems sensible to say that the sentence in question is false (this was Russell's [46] view), and on this view it does make sense to talk of the probability of such sentences being true. The issue is not whether it seems sensible to attribute truth-value gaps to sentences that have non-referring terms and whether it seems sensible to speak in terms of probability about these same sentences; the issue is whether the latter is sensible *given a commitment to the former*. That is, is it sensible to say, for instance, that some sentence is neither true nor false but it's probably true? It would seem not, for this would commit one to a kind of Moore's paradox.[10]

Now if it still seems sensible, on the view under consideration (i.e., the truth-value-gap view), to talk about the probability of truth for sentences with non-referring terms, it's because there's an implicit assumption that

[9]I thank Daniel Nolan for raising this objection.

[10]This is the paradox of an agent asserting 'P but I don't believe it'.

there's no denotational failure. So, for example, the probability that your colleague's wife has left him is something like the probability that his wife has left him, given that he's married. If it turns out that he is not married, the probability in question is the probability that his wife has left him, given that he is both married and not married. This probability is undefined. So even if it may *seem* sensible on this view to talk about the probability of truth for sentences with non-referring terms, it isn't.

3 Cox's Theorem and its Assumptions

Thus far I've outlined two quite distinct sorts of uncertainty and argued that only epistemic uncertainty is amenable to probabilistic treatment. Now I turn to Cox's theorem[11] and how the lessons of the last section impact on the philosophical significance of this theorem. Cox's theorem can be stated as follows:

Theorem 1 (Cox) *Any measure of belief is isomorphic to a probability measure.*

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

There has been a great deal of discussion on the assumptions of Cox's theorem and alternatives to these assumptions ([7], [11], [23], [24], [48], [52]), but there has been little if any discussion of the logical assumptions (or alternatives to these), and yet these are crucial to understanding the significance of the theorem. It is those logical assumptions I now wish to examine.

The logical assumptions of the theorem are, of course, none other than classical propositional logic, though very few state this explicitly and unambiguously. For example, Cox invokes "*the* algebra of symbolic logic" (italics added). But even in Cox's day there was more than one such logic. Jaynes [30] tells us that the logic is "deductive logic". But again 'deductive logic' is ambiguous between the many logics deserving of this title. Elsewhere [31, p. 9] Jaynes suggests that the logic is "two-valued logic or Aristotelian logic", obviously thinking that classical (two-valued) propositional calculus and Aristotelian logic are the same (which they are not). Others such as van Horn [52] refer to '*the* propositional calculus' (italics added), again as

[11]There are, in fact, a number of theorems along similar lines (e.g. [1], [2], [8], [16], [25], [36], [39], [52]). Cox's theorem [9, 10], however, is undoubtedly the most well known and so I'll be content to focus on it, although I'll often use the phrase 'Cox's theorem' to apply to the more general class of results.

though there were only one such logic.[12] But what they all have in mind is quite clearly classical propositional calculus. The problem is that none of them calls the logic in question by name and so it is (at least initially) unclear what logic they have in mind. Worse still, some suggest (e.g., by the use of the definite article 'the') that there is only one choice here. Those (like Jaynes [31]) who do point out that there are other logics to choose from do not bother to defend the choice of classical logic in any systematic fashion.[13]

Combine this unclarity about the logic in question (or the number of candidate logics) with a very commonly held view that logic is domain independent.[14] According to this view, the choice of logic does not depend on the domain of application.[15] So if we combine this commonly held view about logic with the view that "logic" is classical propositional logic (or first-order classical predicate calculus), then we are led to a view that classical logic is all we need for deductive inferences on any domain. Again it is clear that some commentators on Cox's theorem hold such a view. Indeed, Van Horn states this quite explicitly: "the propositional calculus is applicable to *any* problem domain for which we can formulate useful propositions" (p. 11, italics in original). Of course there is a sense in which Van Horn is right—classical propositional logic is *applicable* to any domain, but that's not the issue. The issue is whether classical propositional logic can be applied to any domain *and get the right answers*. It is clear that it cannot. One needs only consider arguments involving modality to see the inadequacy of classical propositional logic.[16] Van Horn, of course, is not alone in holding such a view of logic, though I've never seen anyone suggest that classical propositional calculus is the universal logic—the usual candidates are classical

[12] Though in footnote 1 on p. 5 of [52] van Horn suggests that we may also consider numerical identity statements. This suggests that full (classical) first-order logic is what is needed.

[13] Jaynes does make some rather obscure comments by way of defence of classical propositional logic. For instance, in a section of his book [31, p. 23] called 'Nitpicking' Jaynes raises the possibility of alternative logics and suggests that "[multiple-valued logics] can have no useful content that is not already in two-valued logic; that is, that an n-valued logic applied to one set of propositions is either equivalent to a two-valued logic applied to an enlarged set, or else it contains internal inconsistencies." It is not clear what he means by this, and the appendix where the argument for this claim is supposed to be found is of no help. In any case, Jaynes seems to be thinking of multi-valued logics as the only non-classical logics. As we've already seen, there are others—for example, free logics.

[14] This widely held view found a powerful advocate in Tarski [50].

[15] In essence, this is a monist or one-size-fits-all view of logic, as opposed to a more pluralist horses-for-courses view. See [5] and [41] for discussion on the monism–pluralism debate.

[16] Consider the argument from 'there is uncertainty' to 'possibly there is uncertainty'. This argument is clearly valid and yet the validity cannot be demonstrated by classical propositional calculus, because the only way to formalise this argument in this logic is as P therefore Q which is invalid. To demonstrate the validity of such arguments, modal logic is required. See [18] for a good introduction to modal logics and their applications.

first-order logic or a extension of it such as S5 modal logic. But what I'm arguing here is that no classical logic is up to this task. Classical logic simply fails in some domains in which we routinely perform logical inferences.

Cox's theorem, if it is to demonstrate the adequacy of probability theory for plausible reasoning across *all* domains, it must be derivable from assumptions that are not domain specific. But as I've already argued, classical logic *is* domain specific. Or at least, we've been offered no argument to the effect that it is not. All we are typically given are rather casual acceptances of classical propositional calculus as though there were no other, or, at least, no other worthy of serious consideration. So what is delivered is not a logic of plausible reasoning, *simpliciter*, instead we have a logic of plausible reasoning that is defensible only when there is no referential failure, vagueness or the like. Now perhaps this is all some commentators have in mind—a limited-scope logic of plausible reasoning. If this is the case, then this limitation needs to be stressed. But it is clear that not all contributors to the literature on Cox's theorem have such a modest project in mind.[17] Again Van Horn states this point of view very clearly: "recall the purpose of this enterprise: to construct a *universal* system or logic of plausible reasoning" (p. 11, again emphasis in original). My point is simply that if the enterprise is, as Van Horn suggests, that of constructing a universal system, it had better not rest on classical logic. On the other hand, if the enterprise is the more modest one suggested above, this needs to be made clear.

Now lets turn briefly to the question of whether the proof of the theorem requires any of the contentious features of classical logic? What if, for instance, the proof only relied on inferences and logical equivalences that are not controversial in the context of the representation of belief—inferences such as *modus ponens* and equivalences such as de Morgan laws? There is no need to ponder such questions too long, for the standard proofs of Cox-style results quite clearly rely on disputed logical principles. First, an example from Cox's original proof and then another example from a more recent proof. In Cox's proof that the belief in $Q \vee \neg Q$ is maximal, Cox quite explicitly assumes the classical principle of double negation elimination: $\neg\neg Q \equiv Q$. If we limit our attention to epistemic uncertainty and exclude all forms of non-epistemic uncertainty, then the assumption seems harmless. On the other hand, if we are interested in uncertainty in the broadest sense (including constructive domains, vague domains and so on) the assumption is highly controversial.[18] For a more recent example I once again turn to Van Horn [52] who also uses double negation elimination (in the proof of Proposition 2 on page 11) and assumes, in the proof of Lemma 11 (on page. 20), that $\neg B \equiv ((A \vee \neg B) \wedge (\neg A \vee \neg B))$. This last assumption is

[17]And I stress that this *is* a modest project, because vagueness is ubiquitous in both scientific and everyday discourse. The limited-scope logic of plausible reasoning will thus be rarely applicable outside pure mathematics.

[18]Indeed, intuitionists deny this principle.

very closely related to excluded middle. With the usual classical assumptions in place about the distribution of $\vee$ over $\wedge$,[19] it amounts to the assumption that $A \wedge \neg A$ is false, which under further (classical) assumptions is equivalent to excluded middle. So at the end of the day, controversial features of classical logic are assumed in the original proof of Cox's theorem and these assumptions remain in modern presentations.

The assumption of classical logic is particularly troublesome if Cox's theorem is to be wielded as a weapon against non-classical systems of belief representation. And, I should add, that some commentators do put Cox's theorem to such a purpose. For example, Lindley [36] draws the following conclusion (from a similar theorem): "The message is essentially that only probabilistic descriptions of uncertainty are reasonable" (p. 1) and Jaynes [31] suggests that "the mathematical rules of probability theory [...] are [...] the unique consistent rules for conducting inference (i.e. plausible reasoning) of any kind" (p. xxii).[20] But Cox's result is simply a representation theorem demonstrating that if belief has the structure assumed for the proof of the theorem, classical probability theory is a legitimate calculus for representing degrees of belief. But as it stands it certainly does not legitimate *only* classical probability theory as a means of representing belief, nor does it prove that such a representation is adequate for all domains.

What are the alternatives then? What would these alternate belief theories look like? If we want a probability theory for non-epistemic uncertainty, we may wish to base it on a logic in which excluded middle fails. This means that propositions of the form $P \vee \neg P$ won't automatically receive maximal probability. There are a couple of ways of doing this. One approach would be to allow tautologies to take probability assignments less than one. The other approach is to underwrite the probability theory with a non-classical logic. In this latter case, the tautologies of the non-classical logic will receive maximal probability—it's just that classical tautologies such as $P \vee \neg P$ won't, in general, get assigned the maximal value.[21] In some of the logics in contention, there may be no tautologies (as is the case with Kleene's three valued system K3 and the most popular fuzzy logics [40]). If we use one of these as the underlying logic, there won't be any logical truths and so there won't be any propositions automatically assigned the maximal probability. Some work has been carried out in these directions but there is much more to do.

[19] Interestingly, distribution fails in quantum logics. So quantum logicians may contest the logical equivalence that Van Horn relies on, but for slightly different reasons. See [3] for an early presentation of quantum logic.

[20] Shafer [48] also notes (disapprovingly) this use of Cox's theorem to rule against anything other than standard probability theory.

[21] See [54] for a constructive probability theory.

4 Conclusion

Let me finish by noting a few points of contact between this paper and Glen Shafer's recent discussion of Cox's theorem in this journal [48]. Shafer notes that Cox's theorem relies not only on its stated, explicit assumptions, but it also relies on implicit assumptions—such as the assumption that belief should be represented by a real-valued function. I note one other implicit (or at least undefended) assumption—the use of classical propositional logic. Shafer's work on belief functions [47] casts doubt over the plausibility of the assumption that belief is adequately represented by a real-valued function.[22] I've pointed out that work in logic in the latter part of the twentieth century casts doubt over the plausibility of the assumption (used by Cox and others) that classical logic is the appropriate logic to underwrite a formal theory of plausible reasoning.

The connection between this paper and Shafer's runs even deeper. Not only are both papers questioning implicit assumptions of Cox's theorem. It turns out that our concerns may well be two sides of the one coin. Although our starting points are apparently quite different—mine being logic, Shafer's being the representation of imprecise belief. It turns out that starting with concerns such as mine (i.e., concerns about vagueness and other forms of non-epistemic uncertainty), one very natural way of responding to these issues is to give up the classical logical principle of excluded middle. This in turn naturally leads to a non-classical belief theory that is very similar to Shafer's.[23] In essence, we both reject the unrealistic precision assumed by standard belief theory. Shafer rejects the assumption that belief functions are real valued; I reject he logical assumption of excluded middle.

Another point of contact is that Shafer stresses that the assumptions of Cox's theorem need to be more than merely plausible, they need to be self evident. He points out that both the explicit assumptions and the implicit assumption that belief functions are real-valued fail in this regard. I concur and I add one further assumption to this list of non-self-evident assumptions. In the context of the representation of uncertainty classical logic is not self evidently the appropriate logic. Indeed, I think it is demonstrably *not* the appropriate logic, but even if you disagree with me on this stronger claim, the fact remains that classical logic is not self evident. So those who would employ Cox-style results for the purpose of providing a logic of plausible inference, need to first mount a defence of classical logic.

The final point of contact between my discussion here and Shafer's is that we are both interested in widening the historical focus of the discussion of Cox's theorem. Shafer wants to draw to the attention of commentators

[22] And I find myself in full agreement with Shafer on this issue. See [32] and [53] and for other approaches to abandoning the assumption that a single real number is adequate for characterising belief.

[23] See [13] for details.

on Cox's theorem the earlier work (by continental probability theorists) on the logical interpretation of probability—frequentism and subjectivism are not, nor were they in 1946 (when Cox wrote his paper), the only games in town. I wish to bring to the discussion the issue of the underlying logic—classical logic is not, nor was it in 1946, the only game in town.[24] I think a lot is to be gained by considering these broader historical and, I might add, interdisciplinary perspectives. Once one does this, one sees that Cox's theorem is an interesting representation theorem that has prompted some fruitful and interesting debate, but ultimately the theorem rests on some rather questionable assumptions about the structure of human belief.[25]

References

[1] J. Aczél, Lectures on Functional Equations and their Applications, Academic Press, New York, 1966.

[2] R. Aleliunas, A summary of a new normative theory of probabilistic logic, in: R. Shachter, T. Levitt, L. Kanal and J. Lemmer (Eds.), Uncertainty in Artificial Intelligence 4, North-Holland, New York, 1990, pp. 8–14.

[3] G. Birkhoff, J. von Neumann, The logic of quantum mechanics. Annals of Mathematics 37 (1936) 823–843.

[4] JC Beall, M. Colyvan, Heaps of gluts and Hyde-ing the sorites, Mind 110 (2001) 401–408.

[5] JC Beall, G. Restall, Logical Pluralism, Australasian Journal of Philosophy 78 (2000) 475–493.

[6] P. Benacerraf, What mathematical truth could not be, in: A. Morton, S.P. Stich (Eds.), Benacerraf and his Critics, Blackwell, 1996, pp. 9–59.

[7] M. Colyvan, Is probability the only coherent approach to uncertainty?, Risk Analysis, forthcoming.

[8] P. Cheeseman, An inquiry into computer understanding, Computational Intelligence 4 (1988) 58–66.

[9] R.T. Cox, Probability frequency and reasonable expectation, American Journal of Physics 14 (1946) 1–13.

[10] R.T. Cox, The Algebra of Probable Inference, Johns Hopkins Press, Baltimore, 1961.

[11] D. Dubois, H. Prade, The logical view of conditioning and its application to possibility and evidence theories, International Journal of Approximate Reasoning 4 (1990) 23–46.

[24] By 1946 a variety of non-classical logics had been developed—various many-valued logics and intuitionist logic, in which excluded middle fails. Quantum logic and various modal logics were also known by this time. Moreover, many of these logics had found interesting and fruitful applications. See [40] for a good introduction to non-classical logics and [33] for some of the relevant history.

[25] I'd like to thank Phil Dowe, Scott Ferson, Hartry Field, Dominic Hyde, Daniel Nolan, Graham Priest and Helen Regan for useful discussion or comments on this paper. A version of this paper was presented at the Australasian Association of Philosophy Conference at the University of Queensland in July 2000. I am grateful to the participants in the subsequent discussion for their valuable contributions. I'd also like to thank an anonymous referee of this journal for many extremely helpful suggestions and comments. These led to improvements throughout the paper.

[12] H. Field, Mathematical objectivity and mathematical objects, in: S. Laurence, C. Macdonald (Eds.), Contemporary Readings in the Foundations of Metaphysics, Blackwell, Oxford, 1998, pp. 387–403.

[13] H. Field, Indeterminacy degree of belief and excluded middle. Noûs 34 (2000) 1–30.

[14] H. Field, Truth and the Absence of Fact, Oxford University Press, Oxford, 2001.

[15] K. Fine, Vagueness truth and logic, Synthese 30 (1975) 265–300.

[16] T.L. Fine, Theories of Probability. Academic Press, New York, 1973.

[17] R.A. Girle, Logics for knowledge, possibility and existence. Notre Dame Journal of Formal Logic 19 (1978) 200–214.

[18] R.A. Girle, Modal Logics and Philosophy, Acumen, Teddington, UK, 2000.

[19] J. Goguen, The logic of inexact concepts, Synthese 19 (1968–9) 325–373.

[20] S. Haack, Philosophy of Logics, Cambridge University Press, Cambridge, 1978.

[21] S. Haack, Deviant Logic Fuzzy Logic: Beyond the Formalism, University of Chicago Press, Chicago, 1996.

[22] S. Halldén, The Logic of Nonsense, Uppsala Universitets Arsskrift, Uppsala, 1949.

[23] J.Y. Halpern, A counterexample to theorems of Cox and Fine, Journal of Artificial Intelligence Research 10 (1999) 67–85.

[24] J.Y. Halpern, Coxs theorem revisited, Journal of Artificial Intelligence Research 11 (1999) 429–435.

[25] D. Heckerman, An axiomatic framework for belief updates, in: J.F. Lemmer, L. Kanal (Eds.), Uncertainty in Artificial Intelligence 2, North-Holland, Amsterdam, 1988, pp. 11–22.

[26] A. Heyting, Intuitionism: An Introduction, 3rd revised edition, North-Holland, Amsterdam, 1971.

[27] R.I.G. Hughes, The Structure and Interpretation of Quantum Mechanics, Harvard University Press, Cambridge, MA, 1989.

[28] D. Hyde, From heaps and gaps to heaps of gluts, Mind 106 (1997) 641–660.

[29] D. Hyde, A Reply to Beall and Colyvan, Mind 110 (2001) 409–411.

[30] E.T. Jaynes, How does the brain do plausible reasoning?, in: G.J. Erickson, C.R. Smith (Eds.), Maximum Entropy and Bayesian Methods in Science and Engineering, Kluwer, Dordrecht, 1988.

[31] E.T. Jaynes, Probability Theory: The Logic of Science, Cambridge University Press, Cambridge, 2003.

[32] R. Jeffrey, Bayesianism with a human face, in: Probability and the Art of Judgment, Cambridge University Press, 1992 (originally published in 1983), pp. 77–107.

[33] W. Kneale, M. Kneale, The Development of Logic, Clarendon Press, Oxford, 1962.

[34] S. Körner, Conceptual Thinking, Cambridge University Press, Cambridge, 1955.

[35] K. Lambert (Ed.), Philosophical Applications of Free Logic, Oxford University Press, Oxford, 1991.

[36] D.V. Lindley, Scoring rules and the inevitability of probability, International Statistical Review 50 (1982) 1–26.

[37] E. Mares, A paraconsistent theory of belief revision, Erkenntnis, forthcoming.

[38] J. Meheus (Ed.), Inconsistency in Science, Kluwer, Dordrecht, 2002.

[39] J.B. Paris, The Uncertain Reasoner's Companion: A Mathematical Perspective, Cambridge University Press, Cambridge, 1994.

[40] G. Priest, An Introduction to Non-Classical Logic, Cambridge University Press, Cambridge, 2001.

[41] G. Priest, Logic: one or many?, in: B. Brown, J. Woods (Eds.), Logical Consequences, Kluwer, Dordrecht, forthcoming.

[42] H. Putnam, Vagueness and alternative logic, in: Realism and Reason: Philosophical Papers Volume 3, Cambridge University Press, Cambridge, 1983, pp. 271–286.

[43] R. Routley, Some things do not exist, Notre Dame Journal of Formal Logic 7 (1966) 251–276.

[44] H.M. Regan, M. Colyvan, M.A. Burgman, A proposal for fuzzy IUCN categories and criteria, Biological Conservation 92 (2000) 101–108.

[45] H.M. Regan, M. Colyvan, M.A. Burgman, A taxonomy and treatment of uncertainty for ecology and conservation biology, Ecological Applications 12 (2002) 618–628.

[46] B. Russell, On denoting, Mind 14 (1905) 479–493.

[47] G. Shafer, A Mathematical Theory of Evidence, Princeton University Press, Princeton, NJ, 1976.

[48] G. Shafer, Comments on "Constructing a logic of plausible inference: a guide to Cox's theorem" by Kevin S. Van Horn, International Journal of Approximate Reasoning, forthcoming.

[49] P.F. Strawson, Identiying reference and truth values. Theoria 30 (1964) 96–118.

[50] A. Tarski, On the concept of logical consequence, in: Logic, Semantics, Metamathematics: Papers from 1923 to 1938, second edition, Hackett, Indianapolis, 1983 (originally published in 1936.), pp. 409–420.

[51] B.C. Van Fraassen, Singular terms truth value gaps and free logic, Journal of Philosophy 63 (1966) 481–495.

[52] K.S. Van Horn, Constructing a logic of plausible inference: a guide to Cox's theorem, International Journal of Approximate Reasoning 34 (2003) 3–24.

[53] P. Walley, Statistical Reasoning with Imprecise Probabilities, Chapman and Hall, London 1991.

[54] B. Weatherson, From classical to constructive probabiliy, Notre Dame Journal of Formal Logic, forthcoming.

[55] T. Williamson, Vagueness, Routledge, London, 1994.