Dutch Books and Infinity

Brian Weatherson, Cornell University

draft of July 15 2005
references, etc incomplete
don't cite, refer to or think too hard about

I don't think Dutch Book arguments are generally sound. But I do think they can have some utility in inquiry, by providing reasons for regarding agents who are vulnerable to particular Dutch Books as incoherent. I'll start with a whistlestop tour of why the arguments are not generally sound, then move onto an illustration of the role I think Dutch Book arguments can play in an esoteric, if I think very interesting debate. The debate I'll look at is the status of the countable additivity principle in theories allowing vague probabilities. I'll look at an argument that although countable additivity is a coherence constraint on agents with 'precise' credences, it does not apply in quite the way we'd expect to agents with 'vague' credences. The arguments on either side of this debate illustrate some of the ways we might use Dutch Book arguments to support philosophical conclusions even if we don't think they are *always* sound. (The first four sections go over fairly well-trodden ground, so I assume quite a bit of familiarity. Really all I'm doing through these sections is spelling out which of the familiar arguments I take to succeed – there's nothing particularly new until sections 5 and 6.)

1. Practical Dutch Book Arguments

Let's start with some definitions. We'll say a Smartie is a person whose credences form a probability function, and a Dummie is someone whose don't. A probabilist is a theorist who says there is some valuable and important that only Smarties have. (We'll come back to the question of whether that property is consistency or coherence or rationality or something else.) A Dutch Book argument looks something like the following.

1. Any Dummie is disposed to buy a Dutch Book, and hence do things that can be proven to lose money
2. Anyone disposed to do things that can be proven to lose money is inconsistent/incoherent/irrational
3. So, any Dummie is inconsistent/incoherent/irrational

Both premises can be seriously questioned here.

For premise 1 to be true, we have to consider what Dummies would do in circumstances quite different from the actual world, i.e. when there is a Dutch Bookie around with insight to their mind. Then it isn't obvious that premise 2 is correct, at least if we stress *rationality*. For it isn't obvious that rationality requires dispositions to dispose rationality in any circumstances. Why isn't rationality in circumstances we actually find enough, as externalists and naturalists insist?

Moreover, premise 1 is only true if the Dummie doesn't follow certain well-known principles concerning betting. The argument for premise 1 assumes that if the Dummie's credence in *p* is greater than *x*, and I offer the Dummie a bet that returns $1 if *p* for a cost of $*x*, then the Dummie is sure to buy the bet. But if she does that, the Dummie really is dumb, for she's ignoring Damon Runyon's excellent advise.

> One of these days in your travels, a guy is going to come up to you and show you a nice brand-new deck of cards on which the seal is not yet broken, and this guy is going to offer to bet you that he can make the Jack of Spades jump out of the deck and squirt cider in your ear. But, son, do not bet this man, for as sure as you are standing there, you are going to end up with an earful of cider.

All the Dutch Book argument could show is that Dummies who ignore Runyon are guaranteed to lose money. But for the argument to go through we need to say something about the other Dummies as well. The probabilist could argue that they are inconsistent/incoherent/irrational for listening to Runyon, but that looks the wrong side of the argument to be on.

There are also problems with premise 2. It requires a rather generous stretch of our definition of consistency or coherence to argue that anyone disposed to make provably losing bets is inconsistent or incoherent. They might just have a really bad gambling problem. That suggests the most this kind of argument can show is that all Dummies are irrational. But that argument is independently implausible for reasons that threaten premise 2. If *p* is a complicated logical truth, much too complicated for me to prove, and I have testimony from a generally trustworthy source that *p* is false, then I am not irrational in having my credence in *p* be less than one. But doing so makes me a Dummie. Now I might be irrational for other reasons, but this is not one of them. So one can be a rational Dummie, and indeed in this case one can be rational and disposed to make a bet (in this case a bet on $\neg p$ at reasonable odds) that is provably a loser.

2. Depragmatised Dutch Book Arguments

For these reasons, many philosophers have given up on old style pragmatic Dutch Book arguments. But many have thought, quite reasonably, that something like the Dutch Book argument can be resurrected. We'll look at David Christensen's version of the argument as an example of how this may be done.

Christensen discusses the argument in terms of 'simple agents', i.e. agents for whom the marginal utility of money is constant. He then makes the following argument. (I've slightly modified the wording to fit the above.)

1. A simple agent's degrees of belief sanction as fair monetary bets at odds matching his degrees of belief
2. For a simple agent, a set of bets that is logically guaranteed to leave him monetarily worse off is rationally defective
3. If a simple agent's beliefs sanction as fair each of a set of bets, and that set of bets is rationally defective, then the agent's beliefs are rationally defective
4. If a simple agent is a Dummie, then there is a rationally defective set of bets such that the agent sanctions each as fair
5. So, any simple agent who is a Dummie is rationally defective

Christensen then argues that since the rational defect cannot be blamed on being a simple agent, it follows that all Dummies are rationally defective. There are three problems with this argument. First, as Patrick Maher argued, premise 1 is question-begging against the most interesting kind of Dummies. Second, using some work of Vann McGee we can show that premise 3 must be false on pain of saying *everyone* is rationally defective. Third, it is hard and perhaps impossible to get a notion of rational defectiveness that makes all of the premises true.

The argument for premise 1 is the following passage.

> For a simple agent, there does seem to be a clear relation between degrees of belief and the monetary odds at which it is reasonable for him to bet. If a simple agent has a degree of belief of, e.g. 2/3 that *p*, and if he is offered a bet in which he will win $1 if *p* is true, and lose $2 if *p* is false, he should evaluate that bet as fair … I take these as very plausible normative judgments: any agent who values money positively and linearly, and who cares about nothing else, *should* evaluate bets this way. (Christensen 2005: 117)

In general, how much we value something will turn on the difference in utility between our bundle of goods with it and our bundle of goods without it. If I highly value strawberries and cream, but don't care much for the ingredients in isolation, then how much I regard as a fair price for cream will depend on whether I have strawberries or not. In principle, there is no reason the same shouldn't go for bets. How much I value a bet that pays \$3 if *p* depends on what else I happen to own. Christensen thinks we can intuit for simple agents this is not the case, that the only reason that bets could be complementary or competing goods in the economists' sense is if the agent has a non-linear utility function. The most interesting anti-probabilists, those who think Dempster-Shafer belief functions play the role probabilists think probability functions play, simply deny this. They say even for agents with linear utility functions, bets can be complementary goods. And as far as I can see, Christensen has said nothing to respond to them. This is not to say nothing *can* be said back. I think James Joyce's error-measurement argument, for example, is a very interesting response to these anti-probabilists. But the sheer assumption here that bets are neither complementary nor competing, so we can as much as sensibly talk about what bets an agent sanctions as fair without reference to what other bets she is holding, seems to simply beg the important question.

Vann McGee (1999) proved that for any agent with an unbounded utility function who thinks there are infinitely many ways the world might be, a Dutch Book can be constructed. Call a person with such properties a Hopeful Ignoramus. Hopeful because of the unbounded utility scale, Ignoramus because there are infinitely many ways they think the world might be. Premise 4 is true if we replace 'Dummie' with 'Hopeful Ignoramus'. So if the other three premises are true, we have an argument that all simple agents who are Hopeful Ignoramuses are rationally defective. Actually, since all simple agents are Hopeful by definition, we have an argument that all simple agents who are Ignoramuses are rationally defective. So if Christensen's argument goes through that the problem here is not simplicity, it follows that *all* Ignoramuses are rationally defective. This is clearly false – given how many things most of us don't know, being an Ignoramus is rationally *required*. So something has gone wrong in the argument. And the culprit seems to be premise three. If both Smarties and Dummies can have Dutch Books made against them provided they are Hopeful Ignoramuses, then being vulnerable to a Dutch Book is not a sign of rational defect.

McGee's argument does require consideration of infinite sequences of bets, as well as the somewhat 'exotic' possibility of unbounded utility. Now these aren't things that affect decision making as we actually do it. (Well, they might not be. Pascal would beg to differ here.) But as we saw in section one, there are all sorts of problems with the Dutch Book argument construed as an argument about practical behaviour in the actual world. If the Dutch Book argument is meant to show something about rationality and coherence in the abstract, these considerations have to be taken seriously.

More generally, we might wonder what exactly 'rational defectiveness' comes to in Christensen's argument. It certainly can't be irrationality. As Christensen notes, it would be strange to say that rationality requires full knowledge of logic. In fact, it isn't even clear that being a Dummie is a kind of *defect*. In the situation discussed above, where I rely on a generally trustworthy expert who gets the logical status of p wrong, it seems it would be irrational to firmly believe p. It seems rather harsh to say that I am rationally defective when the alternative would be quite clearly irrational. It seems better to say what these arguments show (if they work) is that all Dummies are in a sense *incoherent*, where we recognise that incoherence is not necessarily a rational failing, and may in some circumstances be rationally mandatory.

3. Vague Probabilities

So far we have looked at arguments that anyone vulnerable to a Dutch Book is in some way incoherent. The arguments so derived for probabilism don't seem particularly telling. It is more profitable to look at the converse argument that someone *invulnerable* to a Dutch Book is coherent. This leads to a few promising avenues of inquiry. As a vast number of authors have noted, an agent need not have linear preferences in order to avoid Dutch Book. At the extreme, if an agent only prefers A to B when A dominates B, then you can't guarantee a Dutch Book against her, even though it isn't clear she is a Smartie in the above sense.

While this amount of indecision is hard to reconcile with perfect rationality, it is arguable that quite a bit of indecision *is* compatible with perfect rationality. It isn't a rational failure if you haven't worked out your credence in an arbitrary proposition to the twentieth decimal place. It seems you could have no rational defects whatsoever, and just not be so decisive that there is a fact of the matter about what the twentieth decimal place of your credence in some arbitrary proposition is.

The philosophers who have taken this line have generally thought that coherent agents should have their credences be representable by a *set* of probability functions. Just what the relationship between this set and their preferences is is a matter of some debate. But the best option seems to be to say that a bet A is (or at least should be) preferred by the agent to B iff the expected utility of A is greater than the expected utility of B according to every member of the set. This leaves open the possibility that two bets may be *incomparable*, neither is preferred to the other, but the agent is not indifferent between them.

Say a Vagueie is a person whose credences can be represented by such a set, where the set is not a singleton. I think the following argument would be too quick.

1. Anyone not vulnerable to a Dutch Book has coherent credences
2. A Vagueie with a bounded utility function is not vulnerable to a Dutch Book
3. So, some Vagueies are coherent

There is, I suspect, more to coherence than *that*, so premise one is too strong. Some incoherent agents might not have their incoherence reflected in betting patterns. But I do think premise 2 provides us with *evidence* that a Vagueie can be perfectly coherent. In the absence of any evidence, indeed any good argument, to the contrary, I'll join the vast company of philosophers who think Vagueies are perfectly coherent.

4.Countable Additivity

Above I introduced the probabilist and the Smartie. But this introduction was a little bit incomplete, since I didn't say what kind of probability function the agent's credences had to satisfy for the agent to be a Smartie. In particular, I didn't say whether I was stipulating that one was a Smartie only if one's beliefs satisfied countable addivity. At least in the philosophical literature, the phrase 'probability function' is ambiguous as to whether it denotes functions that satisfy the finite probability axioms, but not countable addivity. It is important now to remove this ambiguity.

Setting the Vagueies to one side, let us divide the Smarties into the Infinite Smarties and the Finite Smarties, with the difference being that Infinite Smarties have credences that satisfy the countable addivity principle, and Finite Smarties do not. Then we might ask whether coherence (in the sense under discussion) requires being an Infinite Smartie. There is a Dutch Book argument to the effect that it does.

For any finite Smartie whose credence function is C, there will be a sequence of pairwise disjoint propositions p_1, p_2, … such that the following is true

$$C(p_1 \lor p_2 \lor \ldots) > C(p_1) + C(p_2) + \ldots$$

It is easy to see how to make a Dutch Book against such a person. If they are a simple agent in Christensen's sense, and if (as probabilists suppose) bets are not complementary goods, then they are disposed to judge each of the following bets as fair.

Win 1 - $C(p_1 \lor p_2 \lor \ldots)$ if $p_1 \lor p_2 \lor \ldots$, Lose $C(p_1 \lor p_2 \lor \ldots)$ otherwise
Win $C(p_1)$ if $\neg p_1$, Lose 1 - $C(p_1)$ otherwise
Win $C(p_2)$ if $\neg p_2$, Lose 1 - $C(p_2)$ otherwise
…

If none of p_1, p_2 etc are true, the agent wins $C(p_1) + C(p_2) + \ldots$, and loses $C(p_1 \lor p_2 \lor \ldots)$, a net loss. If one of them, say p_i is true, then the agent wins 1 - $C(p_1 \lor p_2 \lor \ldots)$, wins $C(p_1) + \ldots + C(p_{i-1}) + C(p_{i+1}) + \ldots$, and loses 1 - $C(p_i)$. It doesn't take much to show that given that $C(p_1 \vee p_2 \vee \ldots) > C(p_1) + C(p_2) + \ldots$, this is also a sure loss. So the agent sanctions as fair each of a sequence of bets, and those bets are collectively guaranteed to result in a loss. Does this mean they are incoherent?

No, because as McGee showed all sorts of folks are vulnerable to Dutch Books. But given that this book arose even with bounded utility, this should be evidence of incoherence. We can make this stronger by noting, as Seidenfeld et al have proven, that for any Finite Smartie, there will be a proposition q, and a sequence of mutually exclusive, jointly exhaustive propositions $p_1, p_2, \ldots$ with the following property

$$\forall n: C(q) > C(q \mid p_n)$$

That is, the agent's unconditional credence in q is greater than their credence in q conditional on any member of the partition. This looks to be a very bad state to be in, and is also evidence of incoherence. Putting those two pieces of evidence together, we have a strong case that Finite Smarties are not coherent, while Infinite Smarties are.

5. Vague Probabilities and Countable Additivity

So far I've recounted arguments for two views

- It is coherent for agents to not have completely precise credences in all propositions, so all that we should demand is that an agent's credences can be represented by a *set* of probability functions, not necessarily by a single function.
- If an agent does have precise credences, the probability function representing her credences should be a countably additive probability function

Can we generalise the arguments from section 4 to argue that the sets of probability functions representing Vagueies should be sets of countably additive functions? Statistician Peter Walley, in his wonderful book *Statistical Reasoning with Imprecise Probabilities*, argues that the answer is no, and that the arguments for countable additivity for precise agents, which he endorses, do not generalise to Vagueies. I'll recount what I take to be the core parts of Walley's argument, then discuss what we might do to support an affirmative answer. In section 6 I'll say a little about what the philosophical consequences of Walley's conclusion would be.

Consider the following two finitely additive probability functions, P_1 and P_2, both defined over some propositions concerning a random variable X that takes a non-zero integer value.

$P_1(X > 0) = P_2(X > 0) = P_1(X < 0) = P_2(X < 0) = ½$

For all $n > 0$, $P_1(X = n) = 2^{-n-1}$, and $P_2(X = n) = 0$

For all $n < 0$, $P_1(X = n) = 0$, and $P_2(X = n) = 3^n$

The following table describes P_1 and P_2 graphically.

		$\overbrace{\qquad\qquad}^{1/2}$					$\overbrace{\qquad\qquad}^{1/2}$		
P_1	…	0	0	0	0	$1/4$	$1/8$	$1/16$	…
X	...	-3	-2	-1	0	1	2	3	…
P_2	…	$1/27$	$1/9$	$1/3$	0	0	0	0	…
		$\underbrace{\qquad\qquad}_{1/2}$					$\underbrace{\qquad\qquad}_{1/2}$		

Neither P_1 nor P_2 satisfy countable addivity, so we have to represent the probability they assign to the propositions $X > 0$ and $X < 0$ separately from the probabilities they assign to more precise propositions $X = n$. Now consider the following set, which I'll call W after Walley.

$$W = \{Pr: Pr = \lambda P_1 + (1 - \lambda)P_2, \lambda \in [0, 1]\}$$

W is the set of linear mixtures of P_1 and P_2. None of the probability functions in W is countably additive. Nevertheless, there is a decent argument to be made that W could represent a coherent agent. The two arguments I referred to in the previous section certainly do not show that W is incoherent, as Walley shows. (For ease of expression, I'll refer in what follows to a particular agent whose credences are represented by W, and I'll call this agent Green.)

First, there is no Dutch Book to be made against Green. That is, there is no set, finite or countable, of bets such that (a) Green views each of the bets as being strictly positive and (b) the combination of bets is sure to lead to loss. Green, recall, views a bet as being strictly positive if its expected utility is positive according to every member of W. He is, in this sense, not Dutch Bookable.

Second, Green doesn't have any violations of conglomerability, properly understood. It is plausible to understand Green as viewing p as more probable than q iff p is more probable than q

according to every member of W. This means he will view some propositions as *incomparable*, but that's just characteristic of Vagueies. Then there is no proposition q and partition $\{p_1, p_2, \ldots\}$ such that for each p_i, Green views q as being more probable than q given p_i.

This might look like strong evidence for coherence, but it doesn't yet show that in any deep sense countable addivity is not a constraint. One of the difficulties with Vagueies is that there are many equivalent ways to represent their credal states. For instance, if an agent is maximally agnostic about the world, we can represent their credences by the set of all countably additive probability functions, or the set of all finitely additive probability functions. This is a distinction without a difference, because for any two bets, whether the first is better than, worse than, just as good as, or incomparable with, the second will be the same according to each set of functions. So we have to check that these facts about Green suggest something interesting about his credal state, i.e. about what W represents, not just about the representation W itself.

To avoid this potential complication, Walley also proves that W cannot be construed as a set of countably additive probability functions. In fact he proves a much stronger claim, namely that for any countably additive probability function P, there is a bet B such that the expected utility of B is higher according to every function in W than it is according to P. From this it quickly follows that no set of countably additive probability functions can be equivalent to W.

Having said all this, I think there is something wrong with W, and that it isn't obvious that Green is coherent. To express this we need to look not just at which trades Green regards as having positive utility, i.e. as being rationally mandatory to make, but at which trades he regards as rationally permissible. To properly investigate this, we need to go over some considerations about what the logic of prudential permissibility looks like for Vagueies.

Many Dutch Book arguments seem to assume that if an agent regards each of a series of actions as prudentially permissible, she will regard their conjunction as permissible. A Vagueie, or anyone who thinks a Vagueie can be coherent, will have to deny this. Consider an agent represented by the following set of probability functions.

$$S = \{Pr: 0.4 < Pr(p) < 0.6\}$$

If any set of functions can represent a coherent state, then presumably S can. But there is a Dutch Book argument of a sort against this. Consider the following pair of bets.

B_1: Win \$45 if p, lose \$55 if $\neg p$

B_2: Win \$45 if $\neg p$, lose \$55 if p

A simple agent represented by S will not regard either B_1 or B_2 as strictly positive bets. Rather, they will be indifferent between the status quo and holding either of those bets. For each bet, they regard both taking it and leaving it as rationally permissible actions, just as a traditional Smartie regards both taking it and leaving it as rationally permissible when the expected utility of a bet is exactly 0. But this seems to permit too much, for it suggests the agent will regard accepting both B_1 and B_2 as a rationally permissible choice of action. And in combination B_1 and B_2 leads to a sure loss of \$10. So the Vagueie is incoherent.

The argument here is too fast, because it assumes that if the Vagueie regards two actions as each rationally permissible, then she regards their conjunction as rationally permissible. If we interpret $\Box$ and $\Diamond$ as meaning *is rationally required* and *is rationally permitted*, then the assumption here is that $\Diamond p$ and $\Diamond q$ implies $\Diamond(p \wedge q)$. That should make us immediately suspicious, because on most interpretations of the modal operators, this is an invalid inference. We can increase the suspicions by noting that when we interpret $\Box$ and $\Diamond$ as standard deontic operators, the inference again seems clearly invalid. Consider the following little exchange.

Billy: I just got an invitation to a party. Do you think it would be immoral of me not to go?
Suzy: No, that would be OK.
Billy: But it seems like a good party. Do you think it would be immoral of me to promise that I'll be there.
Suzy: No, that seems OK as well.
Billy: Ah ha! So you think it is OK to break promises.
Suzy: No I don't, why do you think that.
Billy: Well, you said it was permissible for me to promise to be there, and permissible to not be there, so it is permissible to promise to go but not show up.
Suzy: That doesn't follow!

Most of us will side with Suzy here. She shouldn't be committed to the acceptability of promising to go to the party and not showing up just because she's committed to the acceptability of each individual act. The Vagueie says the same thing about decision making. There's a crucial distinction between holding each action to be permissible considered on its and holding the whole sequence to be permissible. It's worth thinking through how many Dutch Book arguments rely on eliding this distinction. It isn't obvious, for instance, that the plausibility of Christensen's argument doesn't turn on running these two things together. In general, if a theorist wants to argue from the fact that a particular Dummie regards each of a series of trades as *permissible* then she regards their sequence as *permissible*, that theorist needs to distinguish their reasoning from Billy's reasoning above, and I think that will be a challenge. (This consideration doesn't apply of course to Dutch Book arguments that construct a series of bets each of which the Dummie thinks it is rationally *mandatory* to accept.) That is a topic for another day however, because I want to get back to what this shows about W and Green.

It seems wrong to accuse Suzy of saying it's acceptable to break promises on the basis of what she said above. But if she'd said (a) that it's acceptable to promise to go the party, and (b) it is acceptable to not go *give that promise has been made*, then she seems committed to the view that it is acceptable to break promises. That suggests a constraint on rational permissibility. If each action of a sequence is permissible given the earlier members, then the agent is committed to the permissibility of the sequence. What is it for an action to be permissible given the previous actions? This is hard to answer in general, but in the case where the previous actions were also permissible we can give a plausible answer. Accepting B_n given you've already accepted $B_1, \ldots, B_{n-1}$ is acceptable if the result of adding together B_1 through B_n is a bet that has positive expected utility according to at least one probability function in the set representing the agent's credences. That is, accepting B_n is permissible given that B_1 through B_{n-1} have been accepted iff accepting the bet $B_1 + B_2 + \ldots + B_n$ is permissible.

The result of this definition is that for finite sequences of bets, the agent will never find a Dutch Book acceptable. If accepting all of B_1 through B_n is permissible given what came before them, then accepting the conjoined bet $B_1 + \ldots + B_n$ must be permissible. But accepting a Dutch Book is never permissible. So no finite set of permissible bets can ever lead to sure loss, as long as we properly construe a bet as permissible iff it is permissible given what else has been accepted. Now this might seem like a completely trivial definition, but it turns out to have serious consequences when we apply it to infinite sequences of bets.

Consider Green again, the agent represented by W, and the following series of bets.

B_0: win \$3 if $X < 0$, lose \$2 otherwise
B_1: win \$1 if $X = 1$, lose \$4 if $X = -1$, nothing otherwise
B_2: win \$1 if $X = 2$, lose \$4 if $X = -2$, nothing otherwise
…
B_n: win \$1 if $X = n$, lose \$4 if $X = -n$, nothing otherwise
…

Since the expected utility of B_0 is positive on every function in W, it is mandatory to accept B_0. (I haven't said what Green's utility function is, but everything here goes through as long as Green prefers more money to less.) Accepting B_1 is not mandatory, but it is permissible, since for many functions in W, the bet has positive expected utility. (According to P_1, for example, it's expected utility is plus 25 cents.) Indeed, every bet in the sequence has a positive expected utility according to at least one function in W, so it is permissible to accept it. More importantly, accepting every bet *in order* is permissible given that Green has accepted all the bets that came before it. So it seems Green is committed to the permissibility of accepting all of the bets. So Green is committed to the acceptability of the conjunction of all of the bets. But the conjunction of the bets is a sure loss of one dollar, as can be easily seen. So *in a sense*, there is a Dutch Book that can be made against Green. There is no sequence of bets such that he regards all the bets as being strictly better than the status quo, but their conjunction is a sure loss. He doesn't, that is, regard it as rationally mandatory to do something rationally impermissible. But there is a sequence such that he regards accepting each member of the set as acceptable, even given that the earlier members have been accepted, but their conjunction is a sure loss. That is, he regards it as rationally permissible to do something rationally impermissible.

Now we shouldn't take too much of this. So far nothing we've shown makes it the case that Green is any less coherent than a Hopeful Ignoramus who is not a Vagueie. After all, for any such Hopeful Ignoramus there is a series of bets such that they regard accepting each as permissible given that the bets before them have been accepted, but don't regard accepting the entire sequence as permissible. If this argument is to be used to show that Green is incoherent, we have to show either (a) that being Hopeful is incoherent or (b) that this example shows that there is something wrong with Green's *credences* while McGee's example at most shows something about a Hopeful Ignoramuses *utility function*. It might help to support (b) to note the fact that we can make the above argument *whatever* Green's utility function, so the problem (if there is one) is not with the utility function, while McGee's argument turns crucially on the Hopeful Ignoramus being Hopeful. But I suspect we won't be able to find a *conclusive* argument here, or even one that is close to being as powerful as the original argument for countable additivity.

6. Why This Matters

There is a bit of a tradition in Bayesian theory of treating the issue of vague probabilities as something of an afterthought, something that we should talk about for completeness but not something that is crucial to the theory as a whole. It is easy to see why this attitude is appealing. On a common way of looking at things, vague probabilities are really just a special case of vagueness in representation. Just as there's no fact of the matter about *exactly* what function a vague predicate like *red* represents, there's no fact of the matter just what probability function the brain of a Vagueie represents. But just as we can for the most part ignore vagueness when doing formal semantics, unless we are explicitly talking about vagueness, we can ignore probabilistic vagueness when studying the formal properties of the way coherent agents represent the world.

This would be a happy conclusion, at least if one is made happy by the existence of a division of labour between vagueness theorists and everyone else. But of course it is a controversial conclusion. Some theorists have argued that semantic vagueness threatens familiar logical rules. (Mark Sainsbury, for example, has argued that semantic vagueness challenges the familiar $\vee$-elimination rule.) And Walley's example suggests that whether an agent is coherent cannot be determined by just working out whether they are coherent on every precisification of their credal state. No 'precisification' of W is coherent, but arguably Green himself is coherent. This in turn threatens the idea that we can ignore vagueness for most purposes when investigating the coherence of agents.

I want to conclude on a more optimistic note. We should, as a rule, prefer theories that provide unified explanations to theories that don't. That's a reason to treat the vagueness in probabilistic mental representation as the same phenomenon as the vagueness in natural language. And to the extent that a broadly supervaluational approach to vagueness in language is plausible, that's a reason to prefer a supervaluational approach to vagueness in probabilistic mental representation. So that's a reason to think the same constraints apply to precisifications of Vagueies as apply to agents with precise credences. And that might be a reason, *far* from a conclusive reason, but a reason, to regard Green as incoherent. Since the Dutch Book style argument is also a tentative reason to regard Green as incoherent, we might have enough reason to conclude that he is incoherent, and that W does not represent a coherent agent. But these conclusions are much more tentative than we should like if we are to be confident that an agent whose mind vaguely represents the world is coherent only if all (or even some) of the precisifications of that mind are coherent.