

Penultimate version; final version in *Journal of Philosophy* 93 (1996).
Please feel free to email me for a pdf of the published version:
david.christensen@brown.edu.

DUTCH BOOK ARGUMENTS DE-PRAGMATIZED: EPISTEMIC CONSISTENCY FOR PARTIAL BELIEVERS*

1. Dutch Book Arguments and Pragmatic Consistency

A little reflection on my beliefs is enough to convince me that they come in degrees. I believe that I will eat Italian sausages tonight, and I believe that the sun will set tonight, but I am much less confident about the sausages than about the sunset. Reflection on the fact that beliefs come in degrees has convinced many that formal constraints on ideally rational belief must go beyond those provided by deductive logic. If my beliefs fall short of ideal rationality when I believe P but disbelieve $(P \vee Q)$, do they not also depart from ideal rationality--and in much the same way--when I believe P to a greater degree than I believe $(P \vee Q)$?

The most immediately appealing model for formal constraints on degrees of belief is provided by probability theory, which tells us, for instance, that the probability of P can never be greater than that of $(P \vee Q)$. But while this model has much intuitive appeal, many have been concerned to provide arguments showing that ideally rational degrees of belief would conform to the calculus of probabilities. The arguments most frequently used to make this claim plausible are the so-called "Dutch Book" arguments.

The arguments begin by taking an agent's degrees of belief to be measurable by her willingness to accept bets. Though the details of the betting arrangements in various Dutch Book arguments differ somewhat, they all involve the agent accepting bets at the odds dictated in the intuitively natural way by her degrees of belief. For example, on the basis of my .75 degree of belief in my having sausages tonight, I would be willing to accept a bet at 3:1 odds that I will eat sausages, and equally willing to accept a bet at 1:3 odds that I will not have sausages.¹

Of course, the agent's degrees of belief may not obey the laws of probability--there may be no probability function which matches the agent's degree of belief function for every proposition about which the agent has a degree of belief. That will be the case, if, for example, my degree of belief in P is greater than my degree of belief in $(P \vee Q)$. The Dutch Book arguments show that in such cases, the agent will accept a set of bets on which she is guaranteed to lose money overall.²

Now the lesson of the Dutch Book arguments is supposed to be that ideally rational degrees

of belief must conform to the probability calculus. Probabilistic coherence is thus supposed to constrain rational degrees of belief, just as deductive consistency has been taken to constrain rational all-or-nothing beliefs. But this argumentative move, from the possibility of guaranteed betting losses to constraints on rational belief, has seemed to many a non-sequitur. They have pointed out, for example, that there are no clever bookies who know my degrees of belief and can compel me to wager with them. Clearly, Dutch Book vulnerability is not a real practical liability. Moreover, even if probabilistically incoherent agents were subject to real practical difficulties, it would not obviously follow that their beliefs were defective from the epistemic standpoint.³

Defenders of the arguments have replied that the point of Dutch Book arguments is not to point out a practical problem. Rather, Dutch Book vulnerability indicates a kind of inconsistency. The inconsistency, not the likely prospect of monetary loss, is the problem. This is an especially appealing kind of answer if one would like to see the probabilistic laws as, in F.P. Ramsey's words, "an extension to partial beliefs of formal logic, the logic of consistency." ("Truth and Probability," p. 41)

This general line of thought has considerable appeal. For although the Dutch Book arguments have seemed persuasive to many, it is hard to see how they would have any force at all, if their point was to reveal some practical disadvantage that comes from violating the rules of probability. The suggested approach avoids seeing Dutch Book arguments as crudely prudential. Rather than taking probabilistic coherence as an economically useful defense against being impoverished by transactions with improbably clever bookies, it sees probabilistic incoherence as involving structural defects in the agent's cognitive system.

2. Dutch Books and Pragmatic Consistency

On close inspection, however, the "inconsistency" some Dutch Book defenders are talking about is less parallel to standard deductive inconsistency than one might have thought. The classic formulators of Dutch Book arguments, Ramsey and De Finetti, did not simply make the assumption that certain degrees of belief could naturally be expected to lead to certain betting preferences. They actually sought to define degrees of belief in terms of preferences. If degrees of belief are, at bottom, defined in terms of preferences, the inconsistency involved in having probabilistically incoherent degrees of belief can be seen as an inconsistency of preference. Thus Ramsey writes:

Any definite set of degrees of belief which broke [the laws of probability] would be inconsistent in the sense that it violated the laws of preferences between options, such as that preferability is a transitive asymmetrical relation ... ("Truth and Probability," p. 41)

More recently, Brian Skyrms put the point this way:

Ramsey and De Finetti have provided a way in which the fundamental laws of probability can be viewed as pragmatic consistency conditions: conditions for the consistent evaluation of betting arrangements no matter how described.⁴

Now this sort of consistency of preference is perhaps not the sort of consistency one would initially expect to come from generalizing the notion of deductive consistency to degrees of belief.⁵ Let us call this the "pragmatic consistency" interpretation of the Dutch Book arguments.

It seems to me that there is something very unsatisfying about the pragmatic consistency interpretation. How plausible is it, after all, that the intellectual defect exemplified by an agent's being more confident in P than in $(P \vee Q)$ is, at bottom, a defect in that agent's preferences? It is only plausible to the extent that we take seriously and literally the proposal that particular degrees of belief are defined by particular preferences, or, perhaps more precisely, that degrees of belief reduce to (or necessarily include) certain preferences. Now this proposal may not represent the considered judgement of all defenders of the pragmatic-consistency interpretation of Dutch Book Arguments, some of whom also talk of the relation between beliefs and preferences in more ordinary causal terms. But the important point is this: For inconsistency in beliefs to be inconsistency of preference, certain preferences must be (at least a necessary part of) the beliefs.⁶

This seems at best a very dubious metaphysical view. De Finetti is quite straightforward about his operationalist motivations in this matter. Commenting on his definition of personal probabilities in terms of betting preferences, he writes:

The important thing to stress is that this is in keeping with the basic requirement of a valid definition of a magnitude having meaning (from the methodological, pragmatic, and rigorous standpoints) instead of having remained at the level of verbal diarrhoea ... ("Probability: Beware of Falsifications!," 212)

But today, operationalism and kindred approaches to theoretical magnitudes are widely seen to be

misguided. And this goes not only for physical quantities such as length and temperature, but for psychological concepts such as pain, intelligence, and belief as well.

Now it is true that while simple operationalist reductions of belief have been rejected, certain more sophisticated contemporary accounts--various versions of functionalism--still posit a deep metaphysical connection between beliefs and their typical causes and effects (including other mental states such as preferences and, of course, other beliefs). But the causal interconnections which are said to define a belief are quite complex. They never simply require that a certain belief state necessarily give rise to a certain set of preferences. This brings up a revealing tension in the pragmatic consistency approach to Dutch Book arguments.

Suppose that beliefs are individuated--with respect to degree as well as content--by their causal roles. Then it might be that my high degree of belief that P is in a sense partially constituted by my belief's connections to my preferences, e.g., the fact that I would pay a lot of money for a ticket good for a big prize if P is true. But if beliefs are individuated by their causal roles, they will be individuated not only by their connections to preferences, but by their connections to other psychological states, such as beliefs. If that is true, however, then my strong belief in P would also be partially constituted by its connections to my strong belief that $(P \vee Q)$.

This is where the tension comes in. The entire interest of taking the probability calculus as a normative constraint on belief depends on countenancing the real possibility that the second sort of connection might fail to measure up to probabilistic correctness: I might strongly believe P but not have a sufficiently strong belief in $(P \vee Q)$. But if we countenance this possibility, do we have any justification for refusing to countenance this possibility: that I strongly believe P but do not have a sufficiently strong preference for receiving a prize if P is true? It seems to me that we do not. We have been given no reason to think that having certain appropriate preferences is somehow more essential to having a given belief than having appropriate other beliefs is. Thus the interest of taking the probability calculus as a normative constraint on beliefs is predicated on countenancing the very sort of possibility--failure of a given belief to give rise to the appropriate other psychological states--that undermines the reductionism at the heart of the pragmatic consistency interpretation. An acceptable interpretation of the Dutch Book arguments must acknowledge that partial beliefs may, and

undoubtedly do, sometimes fail to give rise to the preferences with which they are ideally associated.⁷

It is important to note that these considerations do not undermine the view that theorizing about degrees of belief requires that we have some fairly reliable method--or better, methods--for measuring them. Nor do they undermine the view that eliciting preferences in certain ways can provide very reliable measurements of beliefs. But they do, I think, serve to break the definitional link between preferences and degrees of belief. They undermine the oversimplified metaphysical reduction of beliefs to preferences.

Rejecting this sort of reduction has an important consequence for the interpretation of Dutch Book arguments. The arguments' force depends on seeing Dutch Book vulnerability not as a practical liability, but rather as an indication of an underlying inconsistency. Once we have clearly distinguished degrees of belief from the preferences to which they ideally give rise, we see that inconsistency in degrees of belief cannot simply be inconsistency of preferences. If the Dutch Book arguments are to support taking the laws of probability as normative constraints on degrees of belief, then Dutch Book vulnerability must indicate something deeper than--or at least not identical to--the agent's valuing betting arrangements inconsistently.

Now one possibility here is to defend what might be called a "mitigated pragmatic consistency interpretation." One might acknowledge that there is no necessary metaphysical connection between degrees of belief and bet evaluations. But one might hold that there are causal connections that hold in certain ideal situations, and that in those ideal situations, violations of the probability calculus are always accompanied by preference-inconsistencies. One might then point out, quite rightly, that finding norms for idealized situations is a standard and reasonable way of shedding light on normative aspects of situations where the idealizations do not hold.

But this, too, is unsatisfying. If the ultimate problem with incoherent degrees of belief lay just in their leading to preference-inconsistencies, then there might seem to be no problem at all with incoherent beliefs that did not give rise to inconsistent preferences. This seems quite unintuitive. If there were agents who had no preferences at all, but who had beliefs about the world, there would surely be something wrong with the beliefs of those who thought P more likely than $(P \vee Q)$. But it would involve quite a strain to suggest that the ultimate problem with such agents' beliefs was that they

would, in ideal circumstances, give rise to inconsistent preferences, if the agents were to develop preferences.

For these reasons, I think we must reject the pragmatic consistency interpretations of the Dutch Book Arguments. Should we, then, give up on the Dutch Book Arguments themselves? Perhaps not. True, there are other ways of supporting the claim that probabilistic coherence is the right analogue of deductive consistency for partial beliefs,⁸ so Dutch Book Arguments may not be necessary for supporting probabilistic coherence. But the arguments have enough initial intuitive power that it would be disappointing, and even a bit surprising, if they turned out to be as thoroughly misguided as their more pragmatic interpretations seem to make them. In the next section, then, I'll explore the possibility of making sense of the Dutch Book Arguments in a fully non-pragmatic way.

3. Dutch Books and Belief Consistency

Although the relationship between degrees of belief and the betting odds to which they often give rise may not be as close as some have thought, there is a relationship that goes beyond the rough psychological causal pattern. Putting aside any behaviorist or functionalist accounts of partial belief, it is initially quite plausible that a degree of belief of, e.g., $2/3$ that of certainty sanctions as fair--in one relatively pretheoretic, intuitive sense--a bet at 2:1 odds. The idea is not that any agent with $2/3$ degree of belief in P is rationally obliged to agree to putting up \$200 to the bookmaker's \$100 on a bet the agent wins if P is true. Various factors--involving, for example, the non-linear utility of money, or risk-aversion, may make it irrational for him to accept such bets. But there does seem to me to be an intuitively appealing normative ceteris paribus connection: other things being equal, an agent should evaluate such bets as fair. Degrees of belief may in this way sanction certain betting odds, even if the degrees of belief don't consist in propensities to bet at those odds. The relationship envisioned here is neither causal nor metaphysical, but justificatory.

It is also intuitively plausible that if a set of betting odds allows someone to devise a priori a way of exploiting those odds to inflict a sure loss, then there's something amiss with those betting odds. And finally, if a single set of beliefs sanctions as fair each of a set of betting odds, and that set of odds is defective, then there is something amiss with the beliefs themselves. The fact that the diagnosis can be made a priori indicates that the defect is not one of fitting the beliefs with the way

the world happens to be. Like deductive inconsistency, it is a defect internal to the agent's belief system.

Interpreted in this way, Dutch Book arguments do not show that degrees of belief which violate the probability calculus are inconsistent in some previously understood sense. But that is reasonable enough. We need not reduce or assimilate consistency of graded beliefs to some previously understood kind of consistency (such as consistency of all-or-nothing beliefs, or of preferences). We are seeking intuitive support for taking a certain set of principles as the best candidate for a formal constraint that plays a role similar to deductive consistency, but which applies to graded beliefs.

In order for Dutch Book arguments thus interpreted to have force, the plausibility of the intuitive assumptions described above must not simply derive from the initial plausibility of taking the probability calculus to provide consistency conditions for graded belief. And one might challenge the independent plausibility of the claim that degrees of belief "sanction" the corresponding betting ratios. After all, the correspondence in question is the same one that emerges from expected utility theory, which already presupposes a probabilistic-consistency requirement.

But the intuitive normative connection between degrees of belief and bets need not derive from an understanding of expected utility theory; a person might see the intuitive relationship between bets and degrees of belief even if she were quite ignorant of the laws of probability. Such a person might not be able to describe even roughly how the probability of P , Q , $(P \& Q)$, and $(P \vee Q)$ should in general relate to one another. Of course, there may be a sense in which our intuitions on these topics are all interrelated, and spring from some inchoate understanding of certain principles of belief and decision. But that seems unobjectionable; indeed, it is typical of situations in which we support a general formal reasoning theory by showing that it coheres with our more specific intuitions.⁹

In this sort of undertaking, knock-down arguments establishing one set of norms as uniquely correct may not be available. And indeed, the Dutch Book results do not, on the suggested interpretation, provide anything like a knock-down proof that beliefs which violate the probability calculus are irrational. Nevertheless, arguments showing that certain formal constraints on our reasoning fit well with some of our (relatively) pretheoretic intuitions about rationality may yet have an important role to play in epistemology.

This distinctively non-pragmatic understanding of Dutch Book arguments allows us to see why their force does not depend on the real possibility of being duped by clever bookies. It does so while avoiding the reduction of beliefs to preferences, yet while allowing that eliciting preferences may often be a reasonable method for measuring degrees of belief.¹⁰ And it locates the defect in incoherent degrees of belief in the beliefs themselves, not in one of their ideal causal concomitants. In the next sections, I would like to use this understanding to illuminate the recent controversy over whether betting-loss arguments can support not only constraints on simultaneous degrees of belief, but rules for belief change.

4. Consistency and Dutch Strategy Arguments

Dutch Strategies are the diachronic analogues of Dutch Books. The agent is offered some bets at an initial time, each of which is fair given her initial degrees of belief. The bookie then waits, and is given the opportunity of offering the agent more bets at a subsequent time, each of which is fair given the agent's subsequent degrees of belief. In some circumstances a clever bookie can, knowing just the agent's degrees of belief (and perhaps the agent's rule for changing beliefs), devise a betting strategy that will guarantee the agent a net loss. It has been argued that, in order not to be vulnerable to Dutch Strategies, an agent must embrace Conditionalization principles--the preeminent Bayesian rules for change of graded belief.¹¹ It has also been argued on Dutch Strategy grounds that a rational agent must obey the "Reflection" principle, which requires her, in a sense, to endorse in advance her own future beliefs.¹²

These Dutch Strategy arguments have been criticized from the perspective that takes Dutch Book vulnerability as an indicator of inconsistency. The inconsistency revealed by guaranteed betting losses, it is objected, resides in the degrees of belief that render the bets in question fair. In Dutch Strategies, the bets in question are rendered fair by the agent's beliefs at two different times. But an agent's beliefs at two different times need not be consistent in the way an agent's simultaneous beliefs should be. Thus Dutch Strategy vulnerability indicates a sort of inconsistency that rational agents may have no reason at all to avoid.^{13,14}

This may be well and good as regards Reflection, which is a highly counterintuitive principle to begin with. There are many cases, e.g. when one considers the possibility of developing mental

illness, the likely effects of drugs, or the inevitable fading of one's memories of trivial facts, in which it seems entirely rational--indeed, rationally required--to be quite distrustful in specific ways about the beliefs one expects to have in the future. Thus the assessments of one's future beliefs required by minimal common sense often violate Reflection.¹⁵

Conditionalization principles for change of belief, however, are much more intuitively plausible. The classic principle of Conditionalization would regulate belief change as follows: Suppose that I am now unsure about the main course for tonight's dinner and about the color of wine that will accompany it. In particular, I now give red wine a 70% probability. But my probability for red wine, on the condition that we're having Italian sausages, is 90%. If I then learn for certain that we're having sausages, and I learn nothing else, I should put the probability of red wine at 90%. There is undeniably something attractive about this principle for belief change, and about Richard Jeffrey's related principle which applies to situations where one's beliefs change in response to, e.g., sensory stimulation without one's learning any new evidence with certainty. Indeed, these principles are often considered as part of the core of Bayesian epistemology.¹⁶

Some Bayesians have argued that the Dutch Strategy arguments for Conditionalization do indicate an objectionable inconsistency even though the bets they depend on are not all sanctioned by the agent's simultaneous degrees of belief. Brian Skyrms writes:

Notice that the relevant notions of coherence and incoherence here apply not just to the pair of degrees of belief for today and the day after tomorrow, but rather to an epistemic strategy, which is a more complicated object. A focus on the former notion leads understandably to skepticism regarding dynamic coherence ...¹⁷

An "epistemic strategy" for Skyrms includes the agent's initial degrees of belief along with a specification--also made at the initial time--of how those beliefs will be updated according to possible future observations the agent might make. Brad Armendt is particularly clear in on this point:

... the agent is assumed to have, and the bettor is assumed to know, a commitment to a pattern of belief revision that conflicts with the [Conditionalization] rule. That commitment yields Dutch strategy vulnerability, and the commitment is the target of the argument.

The commitments that link the agent's (possible) future beliefs to him at the earlier

time ... are what justify the claim that his inconsistency is synchronic ...¹⁸

The idea is this: The agent has, at the initial time, her degrees of belief. She also has at that same time a rule that specifies what new degrees of belief she will adopt in response to various bits of evidence she may gain. It is between the initial degrees of belief and the updating rule that the possible inconsistency is supposed to lie.

Now both Skyrms and Armendt take the relevant inconsistency here, as in the standard Dutch Book cases, to be an inconsistency of bet-evaluations. The argument, writes Skyrms, "rests on the same conception of pragmatic consistency as the static consistency arguments of Ramsey" ("Higher-order Degrees of Belief," p. 121). This requires, as in the standard cases, that the agent's initial degrees of belief reduce to or include her betting preferences. But this interpretation of Dutch Strategy vulnerability also seems to require a further metaphysical assumption. Skyrms gives the following interpretation (where 'p' represents evidence one might obtain):

The key point is this: prior to finding out about p, the rule or disposition to change my beliefs in a certain way upon learning p is tantamount to having a set of betting ratios for bets conditional on p. ("Higher-order Degrees of Belief," p. 120)

To assess this interpretation, it may be helpful to consider some concrete examples. Suppose that someone has rolled a die, and that I cannot see it. Right now, my probability that a 4 is showing is 1/6. But if I were to learn that the top face was even (and I were to learn nothing more), then I would revise my probability that a 4 was showing to 1/3. This is part of my rule for updating my degrees of belief. And my presently having this rule naturally accompanies my presently valuing bets in a certain way. I would, ceteris paribus, now be indifferent to my making a bet at 1:2 odds on the die showing 4, conditional on the top face being even.

Notice, however, that if having a certain updating rule naturally accompanied certain bet-evaluations, or that certain bet-valuations were excellent evidence for having certain updating rules, that could not be enough to show that having deviant rules necessarily involved inconsistent betting preferences. If one had a deviant updating rule without having the betting preferences that naturally accompanied it, one might not end up harboring the sort of preference inconsistency that the Dutch Strategy is purported to disclose. What is required to support the pragmatic consistency interpretation

of Dutch Strategy arguments is a much stronger assumption: that having a certain updating rule reduces to, or necessarily includes, evaluating bets in a certain way.

This metaphysical claim is, to my mind, highly implausible. We have already seen that degrees of belief must be distinguished clearly from the betting valuations to which they ideally give rise. Similarly, having a certain belief-change rule must be distinguished from having the associated betting valuations. This is particularly clear if we think of "having an updating rule" on the simple disposition model suggested in Skyrms' quote. I may, for example, be so constructed that, were I to learn that any of my friends was trying to quit smoking (and not learn anything else), I would believe strongly that he or she would succeed. But this psychological disposition may be grounded in the emotional reaction I would have to hearing that my friend was trying to quit. Right now, I might have no inclination to accept a bet at long odds (conditional on a friend's trying to quit) that the friend will succeed.

Now it may be that, with a bit of ingenious stipulation, one might describe ideal "betting situations" in which we would strongly expect the correlation between belief-change dispositions and betting valuations to hold. But even if one could do that, correlation-in-ideal-circumstances is not identity. Dispositions to change one's beliefs do not reduce to, nor do they necessarily include, the intuitively associated betting evaluations. If having a deviant updating rule in itself involves some sort of inconsistency, then, it seems that we will have to find that inconsistency somewhere else.

It might be thought that this criticism of the pragmatic inconsistency interpretation of Dutch Strategy arguments relies on assuming Skyrms' simple dispositional picture of the relation between an agent and his updating rule. Recall Armendt's formulation in the passage quoted above:

The commitments that link the agent's (possible) future beliefs to him at the earlier time ... are what justify the claim that his inconsistency is synchronic ... ("Dutch Strategies for Diachronic Rules," p. 221)

On this interpretation, it is not the mere disposition to change one's beliefs in a deviant way that gives rise to inconsistency, but rather one's commitment to change one's belief in that way. This interpretation would not run afoul of the friend's smoking example, since I may well fail to be committed to the policy of believing strongly that my friend will quit if I find out that he's trying, even

if I have a disposition to change my beliefs in this way.¹⁹

Of course, in order for the claimed synchronic inconsistency to be of the simple betting-evaluation sort, a different metaphysical relation will have to obtain: commitments to change one's beliefs in certain ways will have to reduce to, or necessarily include, certain bet evaluations at the initial time. And again, this assumption is a dubious one.

Consider first commitments based on some extrinsic reward. Patrick Maher (Betting on Theories, p. 122) argues for the rationality of some violations of Conditionalization by considering the case of someone to whom a superior being offers eternal bliss if he comes to believe, in a way that violates Conditionalization (e.g. by taking drugs), that pigs can fly. Now I would not take this sort of example as impugning the status of Conditionalization as an epistemic norm--the same type of argument seems constructable against Modus Ponens.²⁰ But it is illuminating, I think, to consider how a person who had made a commitment to changing his beliefs in this way would evaluate bets. Such a person would not, at the initial time, believe that his future beliefs would be reasonable. Thus he would not, even ideally, want at the initial time to make bets based on his anticipated future beliefs (even bets that were conditional on his coming to have those beliefs).²¹

Now admittedly, the odd commitment involved in this example was deliberately chosen to showcase the distinctness of two kinds of psychological states: commitments to belief-updating policies and current betting evaluations. Indeed, Armendt briefly tries to distinguish the kind of commitments he's interested in from "derivative" desires to have certain beliefs such as we see in the flying pigs example.²² Surely, an agent whose belief-change commitment was of a more ordinary sort (and who, in particular, did not regard his anticipated beliefs as formed unreasonably) would, ceteris paribus, be very likely to have the "appropriate" sort of betting valuations, as we imagined in the die example.

In this more standard sort of case, certain betting valuations and belief-change commitments may go naturally together. But they go together in the same sense that certain betting valuations naturally go with certain present degrees of belief. Just as it is reasonable to evaluate bets as fair when the odds correspond to one's current degrees of belief, it is reasonable to regard bets conditional on future evidence as fair when they correspond to what one currently thinks would be reasonable degrees of belief to form if given that evidence.

But given what we have seen about the relation between degrees of belief and preferences, this very comparison should make us careful. Even when an agent has the "right" sort of commitment, and even when he has betting preferences that correlate with his commitment, we should not jump to hasty metaphysical conclusions about the relation between the two. To the extent that having the right sort of commitment to an updating rule sanctions certain bets as fair, it is because one regards the beliefs the rule would produce in the same way one regards one's present beliefs. If we reject the claim that one's present degrees of belief reduce to, or necessarily include, their ideally correlated betting preferences, we should be equally scrupulous in thinking about belief-change commitments.

It might be that the envisioned sort of belief-change commitments are typically, or *ceteris paribus*, causally linked to various other psychological states. Among these are, presumably, various beliefs, desires, other commitments, and valuations, including evaluations of betting opportunities. But to fasten on one of these correlated states, and to take that one correlation as individually necessary, is wholly unmotivated. Belief-change commitments, even when not based on extrinsic rewards, do not reduce to, or necessarily include, the betting preferences with which they may ideally be associated. And once we have seen this, we cannot take the defect involved in harboring deviant updating rules to be that of evaluating betting options inconsistently.²³

5. Betting Losses and Deviant Updating Rules

The claim that those with deviant updating rules must harbor synchronic pragmatic inconsistencies thus fails in two ways. The pragmatic inconsistency is supposed to reside between the agent's degrees of belief and her (commitment to an) updating rule. But neither degrees of belief nor (commitments to) updating rules have the kind of metaphysical connections to bet valuations that the claim presupposes. What, then, should we say about Dutch Strategy vulnerability?

In the case of Dutch Books, we saw that, given certain intuitively plausible assumptions, Dutch Book vulnerability indicates that something is amiss with one's beliefs. The arguments connecting Dutch Book vulnerability with violations of the probability calculus can then be seen as providing intuitive support for taking the probability calculus as the natural extension of the concept of consistency to degrees of belief.

An updating rule, however, is not a belief. And there is no consistency notion from deductive

logic for us to extend here--deductive logic is silent on the topic of belief change.²⁴ So one might not expect that Dutch Strategy vulnerability would indicate the very same sort of inconsistency revealed by vulnerability to Dutch Books. Still, given the attractiveness of Conditionalization, it is worthwhile trying to see what about an agent's state of mind at the initial time might be revealed by Dutch Strategy vulnerability. To that end, let us examine a concrete example of an agent who, through having a belief-change rule that violates Conditionalization, leaves himself open to a Dutch Strategy.

Suppose I am headed out to dinner with my department. Many factors will influence whether I enjoy my meal, but the Chair's choice of restaurant will be particularly important. I presently put the probability of enjoyment, on the condition that our reservation is at Thong Thai, at .9. I don't know where the Chair has decided to go, though; I give her only a .5 probability of choosing Thong Thai. And the other likely prospect is the Faculty Club. Right now, I think that it's about as likely as not that I'll end up enjoying the meal. In this circumstance, the classic principle of Conditionalization would dictate, reasonably enough, that if I now learn for certain that we are going to Thong Thai (and I learn nothing else), I should change my probability for enjoying my meal to .9.

Suppose, however, that I have a rule for belief-change that conflicts with Conditionalization. In particular, my rule dictates that, upon learning for certain just that we are going to Thong Thai, I will raise my probability for enjoyment only to .8. I am now open--given, of course, certain assumptions about my willingness to bet--to a Dutch Strategy. The classic Dutch Strategy, as applied to the present case, would proceed as follows:

A clever bookie, knowing my current degrees of belief, knows that I would now find a certain pair of bets fair. The first derives from my putting the probability of enjoyment, given that we eat at Thong Thai, at .9:

Bet 1: I win \$10 if we eat at Thong Thai and I enjoy it.

I lose \$90 if we eat at Thong Thai and I don't enjoy it.

No one wins if we don't eat at Thong Thai.²⁵

The second bet derives from my .5 probability for going to Thong Thai:

Bet 2: I win \$5 if we eat at Thong Thai.

I lose \$5 if we don't eat at Thong Thai.

Let us suppose that the bookie proposes these bets, and I accept them. So far, there is no guaranteed winner. But the bookie, who is assumed to know my updating rule, already has a plan to seal his victory. He will wait until I learn where we're going for dinner. If it's not Thong Thai, the first bet pays nothing and he's won \$5 on the second, so he will go home a winner. If I learn that it is Thong Thai, and if I've obeyed my deviant rule for belief revision, my new probability for enjoying my dinner is .8. In that case, the bookie will offer me the following bet, which I will at that time find fair given my new degrees of belief:

Bet 3: I win \$80 if I don't enjoy dinner.

I lose \$20 if I do enjoy dinner.

Then, if I enjoy my dinner, I'll win Bets 1 (\$10) and 2 (\$5), but lose Bet 3 (\$20), for a \$5 net loss. On the other hand, if I don't enjoy dinner, I'll win Bets 2 (\$5) and 3 (\$80), but lose Bet 1 (\$90), again ending up \$5 down over all.

The guarantee of loss here does depend on some assumptions: that I will become certain of going to Thong Thai if we will in fact go there; and that I won't become (correctly) certain of going to Thong Thai after learning something else. Technical presentations of the Dutch Strategy argument for Conditionalization typically involve explicit versions of these assumptions. They postulate that there is a set of mutually exclusive and jointly exhaustive evidence propositions, and that the agent will become certain of the true member of this set. But the need for the assumptions shows that, even granting the postulated connections between beliefs, rules and betting valuations, the fact that an agent's updating rule violates Conditionalization does not suffice for Dutch Strategy vulnerability.

One might reasonably point out that the assumptions do often hold (or almost hold), and that norms justified using idealizations may apply to non-idealized situations. But it is worth noting that the standard Dutch Strategy argument has a weakness not shared by the classic Dutch Book arguments. The betting losses it invokes are contingent on assumptions, one of which concerns the relation between the agent and the world.

Let us, however, move on to consider a more troubling worry about the argument above: why should we think that the defect it indicates (if such there be) is already present at the initial time? After all, Bet 3 occurs at a subsequent time, and is judged fair according to the agent's subsequent degrees

of belief. The answer to this question is supposed to be that the agent is already committed at the initial time (via his updating rule) to those subsequent degrees of belief. But this answer naturally raises a further question: if this initial commitment already manifests the defect, why should the bookie need to use delayed bets to take advantage of it?

To focus on this question, let us suppose for the moment that having a rule for belief change on learning p does involve having the correlated betting-ratios for bets conditional on p . On this assumption, it seems that the bookie need not bother with delayed betting after all. Consider the rule in our example, which dictates that, upon becoming certain that we're going to Thong Thai (and learning nothing else), I change my probability for enjoying my meal to 80%. If this involves (as we are supposing) my present commitment to the correlated betting ratio conditional on our going to Thong Thai, then it seems that I should, at the initial time, accept the following conditional bet:

Bet 3': I win \$80 if we eat at Thong Thai and I don't enjoy it.

I lose \$20 if we eat at Thong Thai and I enjoy it.

No one wins if we don't eat at Thong Thai.

Taking Bet 3' at the initial time is equivalent, in terms of monetary gains and losses, to the policy of taking Bet 3 at the later time if we eat at Thong Thai. If having my deviant rule commits me in advance to the policy of taking Bet 3 if we eat at Thong Thai, it should equally commit me to taking bet 3' now.²⁶ Of course, Bets 1, 2, and 3' together guarantee my loss in just the same way as bets 1, 2, and 3 do. Thus it seems that, once we accepted the assumptions underlying the pragmatic consistency interpretation of the Dutch Strategy arguments for commitment to Conditionalization, the Dutch Strategy arguments themselves would become dispensable. We could support commitment to Conditionalization with an old-fashioned, synchronic, Dutch Book argument.²⁷

This Dutch Book approach would seem to present two advantages over the standard Dutch Strategy approach. First, the agent's guaranteed loss is not dependent on the assumptions discussed above. Here, we can connect having a deviant updating rule to guaranteed betting loss, without making the guarantee contingent on certain connections obtaining between the agent's beliefs and facts in the world. Second, by eliminating dependence on bets sanctioned by future beliefs, this approach makes very clear that the defect being exploited is present at the initial time. This nicely distinguishes

this way of supporting Conditionalization from the diachronic betting argument for Reflection. Thus, insofar as we accept the metaphysical underpinnings of the pragmatic consistency interpretation, the Dutch Strategy arguments for Conditionalization are not just superfluous--they are significantly less appealing than their Dutch Book analogues.

6. Betting and Commitments: a Non-pragmatic Interpretation

Unfortunately, we have already rejected the metaphysical assumptions necessary for seeing the defect disclosed by Dutch Book vulnerability in the way the pragmatic consistency interpretation suggests. With respect to the classic Dutch Book argument, this was not fatal. As argued above, we may see that argument as resting instead on an intuitive assumption to the effect that degrees of belief sanction as fair the correlated betting ratios. It seems, however, that the Dutch Book vulnerability involved in the Thai restaurant example cannot be interpreted in exactly this way, for bet 3' is not sanctioned by the agent's degrees of belief. Thus the question arises: is there some other plausible assumption that can play a role here parallel to the assumption connecting degrees of belief with betting evaluations?

Let us see what such an assumption would have to amount to. Recall first some of the points made above about the defect the Dutch Strategy arguments purportedly disclose. Let us assume that the defect is already there at the initial time. It involves the agent's having (or being committed to) a deviant rule, not the agent's following a deviant rule at a subsequent time.

Moreover, the precise relation between the agent and the rule is of crucial importance. Mere disposition to change my beliefs in the way the rule specifies won't do it. As we saw in the friend's smoking case, I may have a psychological disposition to change my beliefs in a certain way, without that disposition reflecting any commitment to changing my beliefs in that way. In such cases, the disposition does not reflect or make rational any present endorsement on my part of the future deliverances of an updating rule that accords with the disposition. The point that is relevant here is not just that there is no necessary metaphysical connection between mere dispositions to change beliefs and current betting preferences. It is that there is no justificatory connection either.

Finally, even commitment to a rule is not enough. Commitments may be based on all kinds of reasons, not all of which involve endorsement--or, at least, endorsement in the relevant epistemic

sense--of the beliefs the rule would produce. So it seems that what we need is an intuitively plausible justificatory connection between a certain kind of epistemic endorsement of an updating rule and the correlated betting preferences.

Is this sort of justificatory connection plausible in the Thai Restaurant case? Here, my rule dictates that if I learn for certain that we're going to Thong Thai (and learn nothing else) I will change my probability for enjoying dinner to .8. If I am merely psychologically disposed to change my belief this way, or if I am committed to this rule because I think that my obeying it will please the gods, then I would intuitively have no justification for now accepting 4:1 bets on enjoying dinner, even bets conditional on our eating at Thong Thai. But suppose that my commitment to the rule is a purely epistemic endorsement. Suppose that I regard the .8 degree of belief as the epistemically reasonable one to form if my evidence increases just by my learning that we're going to Thong Thai. In that case, I think that there is at least some plausibility in saying that my attitude would sanction as fair bets on enjoying dinner, conditional on eating at Thong Thai, at 4:1 odds. In other words, this attitude toward the deviant rule would intuitively sanction Bet 3' in the above example.

If something like this is right, then the right sort of commitment to an updating rule sanctions the correlated conditional betting odds at the initial time. If the rule violates Conditionalization, the agent's beliefs together with her commitment to the deviant rule sanction a synchronic set of bets that guarantee the agent a net loss--a Dutch Book. The heart of the Dutch Book (Bets 1 and 3' in our example) rests essentially on conditional bets at different odds on the same propositions. But that is not to say that agents committed to deviant rules thereby value bets in an inconsistent way. Neither degrees of belief nor commitments to updating rules necessarily include, nor do they always cause, the agent's having the appropriate betting evaluations.

The idea, instead, is that there is a normative, rather than a metaphysical, connection between updating rules and betting evaluations. The right kind of epistemic commitment to an updating rule would, on this interpretation, sanction as fair certain bet-evaluations. The bets thus sanctioned might, when combined with the bets sanctioned by the agent's degrees of belief, produce Dutch Book vulnerability. The Dutch Book vulnerability is a sign that the set of betting odds giving rise to the vulnerability is defective. And this in turn lends credence to the claim that there is something defective

in the combination of degrees of belief and epistemic commitment to an updating rule that sanctioned the set of betting odds.

We saw above that if we take seriously the suggestion that commitment to an updating rule is the real target of the usual diachronic betting-loss arguments for Conditionalization, we can strengthen these arguments by making them synchronic, thereby avoiding the charge that the betting losses result merely from violating an undefended principle of diachronic consistency of belief, and removing dependence on certain assumptions the standard Dutch Strategy argument for Conditionalization requires. The question remained whether this synchronic argument could be given a satisfying philosophical interpretation. The interpretation I've just described is more satisfying in some respects than the pragmatic consistency interpretation favored by the argument's defenders: it avoids commitment to dubious metaphysical claims connecting the having of beliefs or updating rules to the having of the correlated betting preferences.

Should we, then, see this sort of Dutch Book vulnerability as indicating a defect in the agent's updating rule (or, perhaps more accurately, in the combination of her updating rule and her degrees of belief)? There are, I think, some reasons for caution here. One is that the argument depends on an additional intuitive assumption of a normative connection between the agent's betting preferences and her other psychological states. The assumption here--that the right kind of commitment to an updating rule sanctions the correlated betting preferences--is not implausible. But its plausibility is, I think, less obvious than that of the normative principle discussed above connecting betting preferences with degrees of belief.

Another reason for caution is that the defect disclosed by Dutch Book vulnerability in this case does not seem to be one of pure belief-inconsistency in any sense that generalizes the old deductive concept. This may be what we should have expected; after all, commitments to rules are not beliefs. What seems to be disclosed by Dutch Book vulnerability in the present cases is that, in some sense, the updating rule the agent is committed to fails to fit with the agent's degrees of belief. Calling this "inconsistency" is more of a stretch than extending the term from all-or-nothing to graded beliefs. But perhaps this is mainly a terminological worry.

Finally, the argument rests on the notion of the "right sort" of epistemic commitment to an

updating rule. I have used phrases like 'purely epistemic endorsement' without explanation. But although the idea may seem clear enough at first, I am not confident that we have a completely clear grasp of exactly what that sort of endorsement would amount to. Of course, the pragmatic consistency interpretation of the Dutch Strategy also depends on distinguishing the "right sort" of commitments to updating rules, but this tu quoque only shows the present approach to be no worse off than the standard approach. Without a sharper understanding of just what kind of commitment is at issue, it is hard to be fully confident of the key intuitive claim that the right sort of commitment to an updating rule sanctions the correlated bets.

Thus it seems that the envisioned approach to giving a Dutch Book argument for Conditionalization it is not without difficulties. But it is also not without appeal. And in any event, it is more promising, for several reasons, than the standard arguments based on diachronic betting schemes and pragmatic consistency.

7. Deviant Commitments and Inconsistent Beliefs

There is, however, one other approach to giving a Dutch Book argument for conditionalization that seems to me worth exploring. The possibility of this approach flows, in fact, from one of the difficulties encountered with the previous approach. The more one thinks about just how to describe the "right" sort of commitment to an updating rule, the more it seems that it is the beliefs behind the commitment that distinguish it from other kinds of commitment. Perhaps the "right kind" of commitment to a rule is just a commitment based on a certain belief about the rule. This suggests another possible avenue to a Dutch Book argument for Conditionalization: if there are distinctive beliefs that underlie the relevant sort of commitments, perhaps these beliefs themselves harbor inconsistencies when they concern rules that violate Conditionalization.

How would one represent the sort of belief that underlies the relevant sort of commitment to a rule? An initially attractive candidate is provided by a variant on the principle of Reflection. Reflection can be very naturally described as mandating a kind of belief in the epistemic value of one's own future beliefs. In its simplest form, the principle looks like this, where P_1 is the agent's initial probability function, P_2 is her probability function at some future time, and $P(A/B)$ stands for the probability of A on the condition that B:

Reflection: $P_1 (A/P_2(A)= n) = n$.

In other words, my present probability for any proposition A, on the condition that my future probability for A will be n, must also be n.²⁸

Now, as noted above, given various common-sense possibilities of my forming beliefs irrationally in the future, I should not have this sort of blanket epistemic confidence in whatever beliefs I will come to have. But we may use the form of Reflection to express confidence not only in our own future degrees of beliefs, but in other degrees of belief as well. Suppose that we use P_R to designate the probability function obtained by updating my present probability function according to my updating rule R applied to the evidence I will in fact obtain between now and some future time. P_R will thus describe the future beliefs I would have, if I were to follow my updating rule perfectly. A very natural way of expressing my epistemic confidence in my rule is to express confidence in the degrees of belief it would generate if applied correctly to the evidence I will encounter--i.e., to express confidence in P_R . Substituting P_R for my actual future beliefs P_2 in the Reflection principle yields a principle we might call Overall Epistemic Rule Endorsement:

OERE: $P_1 (A/P_R(A)= n) = n$.

This principle does express a kind of confidence in P_R . But on closer inspection, it falls short in a way. It does not necessarily reflect confidence in the various components of R: the ways R uses particular inputs to yield particular output probability distributions. To see how this overall endorsement principle allows one to remain uncommitted to individual components of R, let us look at an example.

Consider a rule that gives directions for probability judgements about poker hands. In particular, it determines, on the basis of any hand one might be dealt, the probability one should assign to drawing to a straight. Suppose such a rule is written up, in the form of a great table, in the Acme Guide to Poker Success. Now suppose you know that the vast majority of the entries in the Acme table are correct. In fact, there are just two misprints. The first concerns hands of type A. Although such hands actually give one a .025 probability of drawing to a straight, the Acme Guide mandates a .05 probability. The second misprint concerns hands of type B. These hands actually give one a .075 probability of drawing to a straight, but the Acme again mandates a probability of .05. Finally,

suppose that type A hands are just as probable as type B hands.

In such a case, what should my attitude be towards the Acme-determined rule for belief-change? The only cases in which I should be reluctant to endorse it are those in which it will mandate probabilities of .05. But even here, there is a sense in which its errors "balance" each other out: my probability for drawing to a straight, given that the Acme-mandated probability for the hand I'll be dealt is .05, should be .05. If we let P_{Acme} stand for my probabilities as they would be determined by the Acme rule, then, it seems that my current beliefs should include the following:

$$P_1(A/P_{Acme}(A)=n) = n.$$

This shows that my very real misgivings about certain components of the Acme rule do not reveal themselves in any violation of OERE. And it is worth noting that the number of misprints in the Acme Guide need not be limited to two. As long as the misprints balance out in the right way, the Guide could be positively riddled with terrible poker advice, and yet the judgement expressed by the principle of OERE would still be entirely appropriate.²⁹

This suggests that if our epistemic endorsement of an updating rule inheres in the beliefs we have about that rule, it will inhere in beliefs more fine-grained than that expressed in OERE. A natural suggestion is that it will involve endorsement of each of the components of the rule. If I is any input to my updating rule R (such as "I become certain that we'll eat at Thong Thai"), we can use $P_{R,I}$ to stand for the probability function that results from updating my present probability function P_1 in accordance with R by input I. My present endorsement of the components of R could be expressed in something like the following principle I'll call Strong Epistemic Rule Endorsement:

$$\text{SERE: } P_1(A/I \ \& \ P_{R,I}(A)=n) = n.$$

In other words, my present probability for A, on the condition that I receive input I and that the probability my rule R mandates for that input is n, is itself n. SERE expresses endorsement of each component of R in the same way that OERE expresses endorsement of R in general. If the right sort of commitment to an updating rule flows from the agent's beliefs about that rule, then the sort of epistemic endorsement represented by SERE seems to me a plausible candidate for representing those beliefs.³⁰

The next question we might ask is whether this principle might give us a way of capturing the

idea that the beliefs lying behind the right sort of commitment to a deviant updating rule are themselves inconsistent. Interestingly enough, I think that the answer is (a qualified) "yes". Let us return for the moment to our Thai restaurant example. The key components of that example are as follows: The first two represent my current degrees of belief, as before:

A. $P_1(\text{I'll enjoy dinner} / \text{we'll eat at Thong Thai}) = .9$

B. $P_1(\text{We'll eat at Thong Thai}) = .5$

The third represents a component of my deviant updating rule; in particular, that if my present beliefs are updated as a result of my becoming certain just that we'll eat at Thong Thai, my new probability for enjoying dinner will be .8.

C. $P_{R, I \text{ become certain just that we'll eat at Thong Thai}}(\text{I'll enjoy dinner}) = .8$

Let us add to this that I now epistemically endorse my updating rule, in the sense given by an instance of Strong Epistemic Rule Endorsement. If we use 'C' to abbreviate the statement of the above-described component of my rule, this yields:

D. $P_1(\text{I'll enjoy dinner} / \text{I become certain just that we'll eat at Thong Thai and C}) = .8$

Now facts A and B, as we have seen, sanction at the initial time our old Bets 1 and 2, given the intuitive assumptions of the standard Dutch Book argument (non-pragmatically interpreted).

Bet 1: I win \$10 if we eat at Thong Thai and I enjoy it.

I lose \$90 if we eat at Thong Thai and I don't enjoy it.

No one wins if we don't eat at Thong Thai.

Bet 2: I win \$5 if we eat at Thong Thai.

I lose \$5 if we don't eat at Thong Thai.

Fact D, given these same assumptions, sanctions at the initial time the following conditional bet on my enjoying dinner:

Bet 4: I win \$80 if I become certain just that we'll eat at Thong Thai, and C, and I don't enjoy dinner.

I lose \$20 if I become certain just that we'll eat at Thong Thai, and C, and I enjoy dinner.

No one wins if it's not the case that (I become certain just that we'll eat at Thong Thai and C).

Now suppose that at the subsequent time, I have become certain just that we will eat at Thong Thai. Since C is true, the condition for Bet 4 having a winner has been met. Assuming that I don't become falsely certain that we'll eat at Thong Thai, the condition for Bet 1 has been met as well. The monetary rewards then stack up as follows: If I enjoy dinner, I win \$10 on Bet 1, win \$5 on Bet 2, and lose \$20 on Bet 4 for a net loss of \$5. If I don't enjoy dinner, then I lose \$90 on Bet 1, win \$5 on Bet 2, and win \$80 on Bet 4, for another \$5 loss. On the other hand, suppose that at the subsequent time I have not become certain just that we'll eat at Thong Thai. Then Bet 4 is off. And assuming (as in the standard Dutch Strategy arguments) that I'll fail to become certain just that we'll eat at Thong Thai only if we don't in fact eat there, Bet 1 is off as well. Still, I lose \$5 on Bet 2, once more ending up \$5 in the hole.

Thus if we take the relevant kind of endorsement of an updating rule to be endorsement based on beliefs captured by SERE, it would seem that we have another avenue for giving a synchronic Dutch Book argument for Conditionalization. The argument is, of course, in a way a step further from the standard Dutch Strategy arguments than was the approach of the previous section. For its focus is not directly on the agent being committed to (let alone disposed to follow) her updating rule, but rather on the agent's belief about the rule. Some defenders of the standard betting-loss arguments might see this as giving up on what they were after. But given the difficulty of providing a philosophically satisfying interpretation of the traditional arguments, it is worth seeing how this approach compares to the one described in the previous section.

First, the advantages. The approach of section 6 depended on an intuitive assumption that the right kind of commitment to an updating rule sanctioned certain betting preferences. The present approach does not. The approach explored in section 6 also involved seeing commitments to updating rules as inconsistent with degrees of belief, in a sense of 'inconsistent' that was hard to see as a natural extension of deductive inconsistency. The present approach, by contrast, is quite parallel to the classic Dutch Book argument for synchronic probabilistic coherence; it purports to reveal an inconsistency in the agent's degrees of belief themselves, a sort of inconsistency that is an intuitively natural

analogue of deductive inconsistency. Finally, the section 6 approach left somewhat unsettled the question of just what kind of "purely epistemic" commitment would intuitively sanction the associated betting preferences. The present argument, with SERE, settles that question precisely.

The present argument does, like the standard Dutch Strategy argument examined in section 5, depend on assuming that if we in fact eat at Thong Thai, I will learn that for certain, and nothing else, by the relevant time. In fact, we must also assume that if I become certain just of our eating at Thong Thai, then we will in fact eat there. I see the dependence on assumptions of this type as limiting the force of the present argument. But though the present approach requires somewhat stronger assumptions than the standard approach, I do not see this difference as marking a significant difference in philosophical power. In fact, the sort of idealized betting scenario described in standard Dutch Strategy presentations will also support the present Dutch Book argument.³¹

Moreover, the present argument presents clear advantages over the Dutch Strategy approach. It avoids both of the dubious metaphysical assumptions underlying the Dutch Strategies as understood on the pragmatic consistency interpretation. It makes crystal-clear that the defect being exploited by the bookie is a synchronic one, not merely a change of mind. And finally, it gives a plausible account of what that defect is.

8. Conclusion

Hypothetical vulnerability to guaranteed betting losses is not a practical financial liability. Thus the force of arguments purporting to derive rules for rational belief from betting-loss considerations must see the hypothetical betting losses as an indicator of a deeper problem. This has been traditionally done by locating inconsistency in the agent's betting preferences. But this approach depends on a highly unrealistic view of the metaphysical relations between betting preferences and degrees of belief.

A better approach takes partial beliefs to sanction--rather than include--certain betting ratios. On this approach, the traditional Dutch Book arguments lend support to the choice of the probability calculus as giving an analogue of consistency for graded beliefs. The defect disclosed by Dutch Book vulnerability is then seen a defect in the beliefs themselves; the problem with having a higher degree of belief in P than in $(P \vee Q)$ is seen as very closely parallel to the problem with believing P but

disbelieving ($P \vee Q$).

This approach helps to throw an interesting light on the controversial diachronic betting-loss argument for Conditionalization. Here, again, the betting losses involved in the argument must be seen as an indicator of a deeper problem. We may take seriously the claim that this deeper problem is already present in an agent who harbors (commitment to) an updating rule that violates Conditionalization. If we do that, one controversial aspect of the argument--its dependence on bets reckoned fair by the agent at different times--is easily eliminated. Moreover, the synchronic version of the argument dispenses with assumptions that limit the force of the diachronic argument. Thus the betting-loss approach to supporting Conditionalization can be significantly improved upon, even within the framework of the pragmatic consistency interpretation.

Nevertheless, the argument as thus interpreted still depends on implausible claims about the connections between degrees of belief, rules for belief-change, and betting preferences. In view of this problem, we examined two distinctly non-pragmatic betting-loss arguments for Conditionalization. The first takes the standard betting arrangement, modifies it to be synchronic, and provides it with a more promising philosophical interpretation, based on the idea that a commitment (of the right sort) to a rule for belief-change might itself sanction certain betting ratios. The second is based on seeing the right sort of commitment to a belief-change rule as involving a certain sort of belief about the rule. This idea, combined with a somewhat different betting scheme, can also be used to argue for the inconsistency of harboring the relevant sort of commitments to updating rules which violate Conditionalization.

Each of these non-pragmatic betting-loss arguments seems to me superior to the standard Dutch Strategy approach. Both make clear that the betting losses they involve flow from a problem the agent has at the initial time; this distinguishes these arguments from the diachronic argument for Reflection. And both avoid reducing beliefs, or commitments to updating rules, to preferences. So while neither one is as strong, to my mind, as the Dutch Book argument for probabilistic coherence of synchronic beliefs, each of them is deserving of further exploration.

DAVID CHRISTENSEN

University of Vermont

NOTES

* I would like to thank Mark Kaplan, Hilary Kornblith, Derk Pereboom, and an anonymous reader for very helpful discussions and comments on drafts of this paper. I would also like to thank the University of Vermont and the American Council of Learned Societies for sabbatical support.

1 In general, an agent's degree of belief in a proposition P is taken to be given by her betting quotient q. An agent's betting quotient for P is q if she would be indifferent between taking either side of a bet on P at odds of q:(1-q). This general pattern fits the example in the text; 3:1 odds are the same as .75:.25 odds. Thus the agent is taken to have a degree of belief function which assigns a number from 0 to 1--corresponding to the agent's betting quotient--to each proposition about which the agent has beliefs.

2 I will not rehearse the mathematical details of the proof that violations of the probability calculus entail Dutch Book vulnerability. The classic presentations are in F.P. Ramsey, "Truth and Probability," and B. de Finetti, "Foresight: its Logical Laws, its Subjective Sources," both reprinted in H.E. Kyburg and H.E. Smokler, eds., Studies in Subjective Probability (Huntington, NY: Robert E. Krieger, 1980). Prominent contemporary presentations include B. Skyrms, Choice and Chance, 2nd ed., (Encino, CA: Dickenson, 1975); P. Horwich, Probability and Evidence (New York: Cambridge University Press, 1982); and C. Howson and P. Urbach, Scientific Reasoning: the Bayesian Approach (La Salle, IL: Open Court, 1989).

3 I have mentioned some representative criticisms, but there are more. For useful discussion and references to the literature, see E. Eells, Rational Decision and Causality (New York: Cambridge University Press, 1982); P. Maher, Betting on Theories (New York: Cambridge University Press, 1993); M. Kaplan, "Not by the Book," Philosophical Topics 21, no.1 (1993): 153-171; and B. Armendt, "Dutch Books, Additivity, and Utility Theory," Philosophical Topics 21, no.1 (1993): 1-20.

4 "Higher-order Degrees of Belief," in D.H. Mellor, ed., Prospects for Pragmatism: Essays in Honor of F.P. Ramsey (Cambridge, UK: Cambridge University Press, 1980) p. 120. See Armendt, "Dutch Books, Additivity, and Utility Theory," for another recent expression of this view of Dutch Book Arguments.

5 Indeed, one might well doubt that "inconsistent" is the best word to use in describing preferences that violate transitivity, for example. Since this terminology has become established, though, I will for convenience continue to use the term in a broad and informal way.

6 Some presentations of Dutch Book results simply assume that agents' betting preferences correspond to their degrees of belief (see B. Skyrms, The Dynamics of Rational Deliberation (Cambridge, MA: Harvard University Press, 1990) and "Discussion: A Mistake in Dynamic Coherence Arguments?," Philosophy of Science 60 (1993): 320-328). For explicit identifications/reductions/definitions of graded beliefs in terms of betting preferences (some in the context of Dutch Book arguments, some not), see: B. De Finetti, "Probability: Beware of Falsifications!," in Kyburg, H.E. and H.E. Smokler, eds., Studies in Subjective Probability; Ramsey, "Truth and Probability," p. 36; Armendt "Dutch Books, Additivity, and Utility Theory," p. 3; R. Jeffrey, "Introduction: Radical Probabilism," in his Probability and the Art of Judgement (Cambridge, UK: Cambridge University Press, 1991); and M. Kaplan, "Not by the Book," and "Confessions of a Modest Bayesian," Canadian Journal of Philosophy, Supplementary Volume 19 (1993): 315-337.

C. Howson and A. Franklin ("Bayesian Conditionalization and Probability Kinematics," British Journal for the Philosophy of Science 45 (1994): 451-466) and C. Howson and P. Urbach (Scientific Reasoning: the Bayesian Approach) identify an agent's degrees of belief with the betting quotients she takes to be fair. They do not, however, take these as entailing any willingness to bet. For interesting expressions of looser connections between beliefs and preferences, see Ramsey "Truth and Probability," p. 30-35) and Armendt, "Dutch Books, Additivity, and Utility Theory," p. 7.

7 A similar problem applies to a somewhat different consistency-based interpretation of the Dutch Book results given by Howson and Franklin in "Bayesian Conditionalization and Probability Kinematics" (a related approach is given in Howson and Urbach's Scientific Reasoning: the Bayesian Approach ch. 3). They argue that an agent who has a certain degree of belief makes an implicit claim that certain betting odds are fair. On this assumption, an agent with incoherent degrees of belief must make deductively inconsistent claims about fair betting odds. Howson and Franklin conclude that the probability axioms "are no more than (deductive) logic" (p. 457). But just as a particular degree of belief may, or may not, give rise to the ideally correlated betting preferences, a given degree of belief may or may not give rise to the correlated claim about fair betting odds. Even if we take degrees of belief to sanction the correlated claims about fair bets, a degree of belief and a claim about betting are not the same thing. Once we see the possibility of this metaphysical connection being broken, it seems a mistake to hold that the real problem with incoherent degrees of belief lies in the claims about bets with which they are ideally correlated.

8 One interesting approach argues that although degrees of belief can't be true or false in the same way that all-or-nothing beliefs can, there is a way of measuring what is naturally thought of as the accuracy of a set of degrees of belief, and

this sort of measure can be used to support probabilistic consistency as a purely epistemic ideal. Jim Joyce discusses this approach in "A Non-Pragmatic Vindication of Probabilism" (manuscript).

9 It is also worth noting that even the "mitigated pragmatic consistency" interpretation of the Dutch Book argument discussed above must presuppose a basic normative connection between degrees of belief and bet evaluations. On this view, degrees of belief lead causally to the correlated betting preferences in ideal circumstances. But one might ask: which circumstances are "ideal"? Why single out those circumstances in which degrees of belief lead to exactly the preferences expected utility theory would dictate? The answer, it seems to me, is that we are intuitively committed to a certain normative relation between degrees of belief and preferences. Circumstances are "ideal" when, and because, this intuitively plausible relation obtains. If this answer is right, then what is perhaps the most controversial assumption in the non-pragmatic interpretation of Dutch Books given in the text also figures in the "mitigated pragmatic consistency" interpretation.

In fact, even the stronger metaphysical claim that degrees of belief necessarily involve the correlated betting preferences may derive some of its plausibility from the type of normative connection advocated in the text. Mark Kaplan, in "Confessions of a Modest Bayesian," describes the following principle as The Bayesian Insight:

For any hypotheses h and g , reason requires that: you prefer a gamble on h to a like gamble on g if and only if you are more confident that h than you are that g . (p. 320, my emphasis)

A paragraph later, he writes:

The insight tells us that to invest more confidence in h than in g is just to prefer a gamble on h to a gamble on g . (p. 320, my emphasis).

10 Of course, such measurements will depend on a presumption that the agent will bet at the odds her beliefs sanction, a presumption that is quite defeasible. Even in ordinary cases, this sort of measurement should not be expected to yield perfect accuracy. (For example, different measurements of an agent's degree of belief in the same proposition will undoubtedly often fail to match one another perfectly.) But this sort of theory-dependence and limited accuracy are typical of respectable measurement in science.

11 See, e.g., P. Teller, "Conditionalization and Observation," Synthèse 26 (1973): 218-258; B. Armendt, "Is there a Dutch Book Argument for Probability Kinematics," Philosophy of Science 47 (1980): 583-588; B. Skyrms, The Dynamics of Rational Deliberation.

12 See B. van Fraassen "Belief and the Will," this JOURNAL 81 (1984): 235-256; J.H. Sobel, "Self-doubts and Dutch Strategies,"

Australasian Journal of Philosophy 65 (1987): 56-81.

13 See D. Christensen, "Clever Bookies and Coherent Beliefs," Philosophical Review 100 (1991): 229-247, which endorses the view that Dutch Book arguments work by revealing inconsistencies in agents' beliefs. It does not, however, explain what sort of inconsistency might be involved.

14 Dutch Strategy arguments are criticized from another direction in I. Levi, "The Demons of Decision," (The Monist 70 (1987): 193-211), which argues that sophisticated agents would see the Dutch Strategy coming, and therefore quite rationally refuse the bets on which the Strategy depends. See also Skyrms, "A Mistake in Dynamic Coherence Arguments?," and Maher, Betting on Theories for discussion of this line.

15 See Christensen, "Clever Bookies and Coherent Beliefs," and W.J. Talbott, "Two Principles of Bayesian Epistemology," Philosophical Studies 62 (1991): 135-150.

16 See R. Jeffrey, The Logic of Decision (Chicago: University of Chicago Press, 2nd ed. 1983) for lucid presentations of both principles. In the remainder of the paper, I will concentrate on the classical principle, and the arguments given for it. But my criticisms of these arguments will apply equally to the arguments for Jeffrey's principle.

17 "A Mistake in Dynamic Coherence Arguments?," p. 321. For similar claims, see Skyrms', "Higher-order Degrees of Belief," and "Dynamic Coherence and Probability Kinematics," Philosophy of Science 54 (1987): 1-20.

18 "Dutch Strategies for Diachronic Rules: When Believers See the Sure Loss Coming," in D. Hull, M. Forbes, and K. Okruhlik, eds., PSA 1992 v.1 (E. Lansing, MI: Philosophy of Science Association), p. 221.

19 It is interesting to note that the Dutch Strategy understood this way would provide a direct argument for what Isaac Levi calls "confirmational conditionalization," as opposed to "temporal credal conditionalization." The latter principle is the diachronic constraint on belief change described in the text. The former principle is a constraint on commitments to belief-change policies. It requires an agent to be committed to changing her beliefs as the diachronic principle would require. See The Enterprise of Knowledge (Cambridge, MA: MIT Press, 1980), ch. 4, and "The Demons of Decision."

20 Brad Armendt discusses this sort of generalizability point in the context of a pragmatically oriented epistemology in "Dutch Strategies for Diachronic Rules," and "Wanted: Irrational Belief Changes (REWARD)," presented to a colloquium at the 1992 meetings of the Eastern Division of the American Philosophical Association.

21 It is worth pointing out that a person's commitment to an updating rule that will produce irrational future beliefs may not indicate any present irrationality at all. If we imagine that Maher's agent's belief about the superior being is itself rational, I see no irrationality--epistemic or practical--at the initial time. The agent may well be engaging in faultless theoretical and instrumental calculation. True, he does not treat his future epistemic rationality as his highest goal. But that does not seem like anything that current rationality--epistemic or practical--requires.

22 See his "Dutch Strategies for Diachronic Rules," p. 225. It is not obvious to me that a clear distinction between derivative and non-derivative commitments can be made. Moreover, even if such a distinction can be made, it seems to me likely that our ordinary epistemic commitments to belief-change rules are themselves derivative.

23 Armendt's "Wanted: Irrational Belief Changes (REWARD)" suggests another view that is worth mentioning here. It might be held that being committed to a certain belief-updating rule involved not having certain betting preferences, but being committed to having certain betting preferences in the future. This would complicate considerably the kind of inconsistency at issue; it would no longer be the having of incompatible simultaneous evaluations of bets. Rather, it would have to lie in the agent's being committed to having future preferences inconsistent with those she actually had. And there are plenty of examples in which it seems entirely rational to commit oneself to developing preferences different than the ones one has. (Would that I could learn to prefer the taste of boiled tofu to that of Italian sausages!) So it would take argument to show that the sort of inconsistency envisioned by the suggested view was objectionable.

Moreover, given that having certain degrees of belief does not necessarily involve having the ideally associated betting preferences, it is extremely doubtful that being committed to having certain degrees of belief entails being committed to having the associated betting preferences. Thus this interpretation seems no more plausible than those considered in the text.

24 The closest deductive analogue to Conditionalization--though the analogy may not be very deep--would seem to be belief change via Modus Ponens. Suppose at the initial time you believe $(P \rightarrow Q)$, but you do not believe or disbelieve Q . You might, upon learning P , revise your beliefs to include Q . But as has often been pointed out, deductive logic makes no such demand on you. You could equally well remain agnostic on Q and drop your belief in the conditional.

25 This is a "conditional bet". If conditional probabilities are understood as related to unconditional probabilities in the standard way, the same result can be obtained in a more

complicated fashion using only non-conditional bets. The general recipe is in Teller's "Conditionalization and Observation."

26 Brian Skyrms notes the equivalence between the payoffs of 3-style and 3'-style bets. (See his "Dynamic Coherence and Probability Kinematics," and "A Mistake in Dynamic Coherence Arguments?") He does this, however, only in the context of presenting a converse Dutch Book argument (an argument that respecting Conditionalization is sufficient, rather than necessary, for invulnerability to Dutch Books).

It is worth taking note of a certain oddness that could result from combining the synchronic approach with the strong metaphysical assumptions that underlie certain pragmatic consistency interpretations of Dutch Strategy arguments. Consider the example in the text. In Bets 1 and 3', the agent essentially accepts two different sets of odds on conditional bets on the same propositions. But now suppose we took the agent's commitments to updating rules, as well as her conditional probabilities, both to reduce to her conditional betting odds. In such a case, we might be unable to decide which of these betting evaluations (accepting 1 or accepting 3') defined her conditional probability, and which represented a component of her updating rule. Perhaps we would be forced to attribute two different conditional probabilities for the same propositions, and/or two different updating rules.

27 A somewhat different synchronic approach to supporting conditionalization is taken by Levi. Levi's argument is based on bet evaluations, though not on avoiding guaranteed losses. It would have the agent in our example consider the hypothetical situation in which he had become certain of going to Thong Thai (and nothing else), and in which he was offered bets on enjoying dinner. Levi argues that the odds the agent would have himself choose in the hypothetical situation should reflect the odds his updating rule commits him to adopting upon becoming certain of going to Thong Thai. But the agent has no reason to choose different odds for this hypothetical-situation bet than he does for a present bet (our Bet #1) conditional on his actually going to Thong Thai (since, Levi argues, the possibility of not going to Thong Thai is not a serious possibility in the hypothetical situation). Given that the latter odds reflect the agent's conditional probabilities, the agent's updating rule should agree with confirmational conditionalization. See The Enterprise of Knowledge, ch. 10 and "The Demons of Decision."

28 The suggestion that Reflection-style principles express the judgement that one's future probabilities will result from "a learning experience"--that the belief change will have "epistemological legitimacy"--is made in B. Skyrms, "The Value of Knowledge," in C. W. Savage, ed., Scientific Theories (Minnesota Studies in the Philosophy of Science XIV), (Minneapolis: University of Minnesota Press). It is developed in Maher's Betting on Theories, ch 5.

29 As an extreme illustration, OERE would even apply to Acme's Pocket Poker Pro, a card which gives the same probability to drawing to a straight for every hand I might be dealt--the probability being that of drawing to a straight from a randomly dealt hand.

30 It should be noted that taking SERE as an account of the right kind of epistemic commitment to an updating rule involves considerable idealization. Surely, ordinary people do not have such fine-grained beliefs about their updating rules. In fact, many ordinary people probably have no very explicit beliefs about their rules for belief change. The whole project of grounding rules for rational belief change in epistemic commitments to updating rules involves idealization. The intuitive idea behind the idealization, though, is a reasonable one: an acceptable rule is one that could be rationally endorsed.

Indeed, it has often been noted that even the demand for probabilistic coherence involves a great deal of idealization. Real reasonable people lack numerically precise degrees of belief in many propositions; they may give probabilities less than 1 to complex logical truths; etc.. The degree of idealization involved in SERE thus seems in line with the rest of the standard Bayesian program.

31 See Teller, "Conditionalization and Observation," pp. 222-223; Skyrms "Dynamic Coherence and Probability Kinematics," pp. 3-4. These scenarios involves situations where there is a set of evidential propositions, and the agent will become certain of one of them if and only if it is true.