Arntzenius on "Why ain'cha rich?"

Arif Ahmed & Huw Price

Abstract

The best-known argument for Evidential Decision Theory (EDT) is the "Why ain'cha rich?" challenge to rival Causal Decision Theory (CDT). The basis for this challenge is that in Newcomb-like situations, acts that conform to EDT may be known in advance to have the better return than acts that conform to CDT. Frank Arntzenius has recently proposed an ingenious counter argument, based on an example in which, he claims, it is predictable in advance that acts that conform to EDT will do less well than acts that conform to CDT. We raise two objections to Arntzenius's example. We argue, first, that the example is subtly incoherent, in a way that undermines its effectiveness against EDT (here we rely on the lessons of Dummett's famous discussion of the conditions for the coherence of a belief in retrocausality); and, second, that the example relies on calculating the average return over an inappropriate population of acts.

1: Introduction

On the standard reading of the standard version of Newcomb's problem the relative *efficacy* of your options diverges from their *news value*: taking the transparent box *makes* you richer than not taking it; but people who don't take it always *end up* richer than those who do. Accordingly Causal Decision Theory or CDT (which values efficacy) and Evidential Decision Theory or EDT (which values news value) make different recommendations: CDT says that you should take the transparent box whereas EDT says that you shouldn't. Many philosophers find grounds in this disparity for declaring against EDT (Gibbard and Harper 1981: 180-184; Lewis 1981a: 377-8; Joyce 1999: 146-54).

This paper concerns an argument that it is the *causalist* who has got things wrong. Frank Arntzenius states it as follows (2008: 289):

> In a Newcomb type case evidential decision theorists will, on average, end up richer than causal decision theorists. Moreover, it is not as if this is a surprise: evidential and causal decision theorists can foresee that this will happen. Given also that is axiomatic that money, or utility, is what is strived for in these cases, it seems hard to maintain that causal decision theorists are rational.

The key premise of this argument is that evidential decision theorists will be richer on average than causal decision theorists. That is not quite the best way to put it: disputes between CDT and EDT, are not about relative welfare of *theorists* who champion those theories. They are about the relative return to the *acts* that those theories recommend, whether the actor in question is himself a self-conscious causalist, a self-conscious evidentialist, or—like the vast majority of people to whom decision theoretic recommendations should also apply—someone who has never heard of either.

So the key premise is better put like this: the *act* that EDT recommends in a Newcomb type situation—namely, one-boxing—has a better average return than the act that CDT recommends there—namely, two-boxing. Making this amendment and affixing Lewis's title for it we have the following argument:

Why ain'cha rich

(1) The average return to one-boxing exceeds that to two-boxing (*premise*)
(2) Everyone can see that (1) is true (*premise*)
(3) Therefore one-boxing foreseeably does better than two boxing (*by* 1, 2)
(4) Therefore CDT is committed to the foreseeably worse option for anyone facing Newcomb's problem (*by* 3)

So understood it is easy to see that the key premise (1) true. Let the predictor get it right 95% of the time. That is: he predicts that a player will one-box (and so puts $1M in the opaque box) on 95% of occasions when that player one-boxes. And he predicts that a player will two-box (and so puts nothing in the opaque box) on 95% of occasions when that player two-boxes. Then assuming linear utility for money and writing M for a million and k for a thousand, the average returns (AR) to one-boxing and two-boxing over many trials are:

(5) AR (One-boxing) = 95%. M + 5%. 0 = 950k
(6) AR (Two-boxing) = 5%. (M + k) + 95%. k = 51k

So clearly (1) is true and everyone can see that. So CDT recommends an act that returns *foreseeably* less than what EDT recommends.

It is no use the causalist's whining that Newcomb problems reward irrationality, or rather CDT-irrationality. If everyone knows in advance that the most productive strategy in a game is the CDT-irrational one then it is *rational* to play the CDT-irrational strategy.

But Arntzenius doesn't whine. Instead he objects that if *Why ain'cha rich* works against CDT then an exactly parallel argument works against EDT. So the evidentialist is hardly in a position to wield *Why ain'cha rich* against CDT. The remainder of this paper describes and then criticizes that parallel argument.

2: Arntzenius's Example

The Yankees and the Red Sox are going to play a lengthy sequence of games; the Yankees win 90% of such encounters. Before each game Mary has the opportunity to bet on either side. The following table summarizes her payoffs on every such occasion as well as my abbreviations for the relevant acts and states:

	RED SOX WIN (R)	**YANKEES WIN (Y)**
Bet on Red Sox (BR)	2	-1
Bet on Yankees (BY)	-2	1

Table 1

Just before each bet a perfect predictor tells her whether her next bet is going to be a winning bet or a losing bet. Now suppose that Mary knows all this. What does EDT recommend?

Suppose that the predictor says: 'Mary, you will win your next bet.' Then the news value V_w (BR) of betting on the Red Sox is:

(7) $V_w(BR) = 2.Cr(R \mid BR \wedge Win) + -1.Cr(Y \mid BR \wedge Win) = 2.1 + -1.0 = 2$

And the news value V_w (BY) of betting on the Yankees is:

(8) $V_w(BY) = -2.Cr(R \mid BY \wedge Win) + 1.Cr(Y \mid BY \wedge Win) = -2.0 + 1.1 = 1$

It follows from (7) and (8) that $V_w(BR) > V_w(BY)$. So if Mary knows that she will win her next bet then her EDT-rational bet is on the Red Sox.

Suppose that the predictor says: 'Mary, you will lose your next bet.' Then the news value V_L (BR) of betting on the Red Sox is:

(9) V_L (BR) = 2.Cr (R | BR ∧ Lose) + -1.Cr (Y | BR ∧ Lose) = 2.0 + -1.1 = -1

And the news value V_L (BY) of betting on the Yankees is:

(10) V_L (BY) = -2.Cr (R | BY ∧ Lose) + 1.Cr (Y | BY ∧ Lose) = -2.1 + 1.0 = -2

It follows from (9) and (10) that V_L (BR) > V_L (BY). So if Mary knows that she will lose her next bet then her EDT-rational bet is on the Red Sox.

So it follows from (7)-(10) that Mary's EDT-rational bet is going to be on the Red Sox *for every game*.

> So Mary will always bet on the Red Sox. And, if the Yankees indeed win 90% of the time, she will lose money, big time. Now, of course, she would have done much better had she just ignored the announcements, and bet on the Yankees each time. But, being an evidential decision theorist she cannot do this. (2008: 289-90)

It is easy to see that she would have done better to bet on the Yankees. The average returns to betting on the Red Sox and the Yankees are respectively:

(11) AR (BR) = 90%. -1 + 10%. 2 = -0.7
(12) AR (BY) = 90%. 1 + 10%. -2 = 0.7

It is also easy to see by contrast that CDT *does* recommend betting on the Yankees every time. Win or lose the causalist's evaluations of a bet on the Yankees and of a bet on the Red Sox are as follows:

(13) U (BR) = 2.Cr (BR → R) – 1.Cr (BR → Y) = 2. 10% - 1.90% = - 0.7
(14) U (BY) = 1.Cr (BY → Y) – 2.Cr (BY → R) = 1.90% - 2.10% = 0.7

(Arntzenius 2008: 290). So the causalist bets on the Yankees every time; and he makes an average 70 cents per game.

'So', Arntzenius concludes, 'there are cases in which causal decision theorists, predictably, will do better than evidential decision theorists' (2008: 290).

The argument against *Why ain'cha rich* is therefore a parity argument: if *Why ain'cha rich* works against CDT then this parallel argument works against EDT. In line with the amendment that I initially proposed to Arntzenius's formulation of *Why ain'cha rich*, I suggest that we rewrite it as an argument about *acts* rather than *persons*: so put it runs as follows:

Yankees.
(15) The average return to betting on the Yankees exceeds the average return to betting on the Red Sox (*premise: from* (11), (12))
(16) Everyone can see that (15) is true (*premise*)
(17) Therefore betting on the Yankees will foreseeably do better than betting on the Red Sox (*from* (15), (16))
(18) Therefore EDT is committed to what is now the foreseeably worse option for Mary (*from* (7-10), (17))

The dialectical position is now as follows. The evidentialist might think that *Why ain'cha rich* is an argument for preferring EDT to CDT. But Arntzenius seems to have shown that whatever the argument shows, it doesn't show *that*. For a precisely parallel argument, namely *Yankees*, gives just the *same* reason for preferring CDT to EDT. In short: *Why ain'cha rich* cuts both ways if it cuts either way. So *it* cannot motivate a preference for EDT.

3: Is the example coherent?

Our initial concern about *Yankees* is that the example seems incoherent, in the sense that it ascribes knowledge to the agent which is incompatible, from her own point of view, with the belief that she has a choice. We rely here on a familiar observation about an incompatibliity between knowledge and free action. If an agent believes that she already *knows* what she is going to do, she cannot take herself to be still freely *choosing* what to do. For the authority she takes herself to have as an agent inevitably trumps whatever evidence might have formed the basis for the knowledge claim. This is sometimes referred to as the thesis that *deliberation crowds out prediction* (see, e.g., Rabinowicz 2002). It is also closely related to a point at the heart of Dummett's famous (1964) discussion of the coherence of backward causation, which we now which to adapt, to illustrate the way in which the Yankees example seems to violate this coherence constraint.

Consider the following example, the *Has Bean Machine:*

> On my office desk, yesterday, there was a box full of beans. The University's bean counters examined its contents, and assured me that 90% of the beans were Yellow, and 10% Red. How did the beans get there? I'll be sending them there, tomorrow, using my new time transporter (the Has Bean Machine). It is yet to be scaled up to human size, but works perfectly for red and yellow beans. It doesn't alter their colour, and I can choose, of course, how many of each colour to send.

Should I trust the bean counters? At this point, we have Dummett's authority for saying "no". Dummett points out that it is coherent for me to believe that a contemplated free action is reliably correlated with some past state of affairs only if I do not also believe that I have knowledge of the state of affairs in question, before I act.

Dummett reaches this conclusion by pointing out that efforts to 'bilk' the claimed correlation between the future action and the past states of affairs – i.e., to arrange matters so that the future action takes place when and only when the relevant past state of affairs does not obtain – rely on the assumption that the agent has access to knowledge of the past state of affairs, before she decides whether to perform the future action. But the bilking objection to backward causation is simply a way of making vivid the thought that without such a restriction on knowledge, backward causation would conflict with the requirement that a free agent cannot take herself to have reliable evidence about what she is going to do, in advance of a free choice.

As our phrasing here suggests, the point applies with equal force to *justified credence,* as well as *knowledge.* The possibility of bilking undermines the former just as much as the latter – indeed, it undermines the latter because it undermines the former, apparently. Moreover, though Dummett himself was interested in the case of backward causation, his point is insensitive to the temporal location of the state of affairs in question, with respect to the action with which it is supposedly correlated. Hence it shows, equally, that one cannot, at the time of choice, coherently take oneself to have a justified credence in a state of affairs one takes to be strictly correlated with that choice: deliberation crowds out prediction.

Let us now apply these insights to the Yankees example. Once Mary knows whether her next bet is a winning bet or a losing bet, she knows that her choice – betting on the Yankees, or betting on the Red Sox – is reliably correlated with the outcome of the game. By a direct application of the principle we take from Dummett, this means that she cannot know, or indeed take herself to have a reliable credence in, the outcome of game, as she deliberates. In particular, she cannot take herself to be justified in assigning a credence 0.9 to a Yankees victory.

This is a clue to the incoherence at the heart of the Yankees example. The assertion that it is a case in which EDT leads to predictable loss depends on the information that the Yankees win 90% of games. However, a free agent with the additional knowledge assumed by the example – knowledge, in advance, about whether each bet will win or lose – cannot take herself to be justified in believing this claim about the frequency of Yankees wins (or the claim it would imply, about a predictable loss).

The Has Bean Machine example makes this point explicit. As she decides to bet (and after she finds out whether the bet will win or lose), Mary's epistemic relation to the outcome of the next ball game is exactly like my epistemic relation to the colour of the next bean I place in the Has Bean Machine, to be sent back in time. (We could even add an analogue of the Win/Lose information to the Has Bean Machine, by having the bean selection mechanism sometimes malfunction, in a manner completely predictable in advance.) So Mary's situation, as she contemplates a season of betting on ball games, is exactly like my situation, as I contemplate selecting a series of beans, one at time, to be sent into the past.

As we observed earlier, Dummett shows that to make my beliefs coherent, I must mistrust the University bean counters, who assured me that 90% of the beans sent back in time were actually Yellow. By parity of reasoning, Mary's beliefs are incoherent, unless she, too, mistrusts the information that the Yankees will win 90% of games. No matter if the bean counter in this case is Chance herself, stoutly offering a prediction of the percentage of Yankees wins. Mary cannot take herself to be justified in believing the prediction; and hence cannot coherently take herself to be facing a certain loss. This objection goes to the heart of Arntzenius's example, for it is the agent's knowledge of the frequency of Yankees wins which is supposed to sustain the conclusion that she knows that she will do less well by EDT than by CDT.[1]

[1] The same point applies to Arntzenius's other example (2008: 290), which resembles Newcomb's problem, except that *both* boxes are transparent, and the predictor has placed $10 in the left-hand box iff he predicted that the agent would not take the right-hand box, which contains $1. Evidential and Causal Decision theories *both* advise taking the contents of both boxes. But agents who heed this advice will foreseeably make less money than those who—insanely—take only the box containing $10.

Our complaint about the Yankees case transposes to this case as follows. If the agent knows that she is going to be able to choose what boxes she takes then she knows in advance that she can so contrive her choices as to make the predictor's accuracy arbitrarily close to zero. (She can do this by taking both boxes on any occasion if and only if the predictor has on that occasion left $10 in the left-hand box.) But if she knows in advance that that is an option for her, then she cannot assume in advance that the predictor is going to be accurate; so she cannot after all foresee that the strategy endorsed by CDT (and by EDT) will be relatively unprofitable.

That case also illustrates especially clearly why the incoherence that it shares with the Yankees example doesn't arise in the standard Newcomb case. In the standard Newcomb case the left-hand box is opaque; and the only way to discover its contents is to make the very decision whose return depends upon them. So there is no way of knowing in advance what on any occasion of choice you have been predicted to choose. Nor therefore is there any identifiable strategy for systematically falsifying those predictions.

4: Restoring the disparity

We now argue that even waiving that last objection, *Yankees* suffers from flaws that do not affect *Why ain'cha rich*; hence one can consistently maintain the latter against CDT whilst denying that the former has any weight against EDT. We can see this by examining arguments in which the central flaw in Yankees appears more clearly.

Here is one. Every Monday morning everyone has an opportunity to pay $1 for a medical checkup at which a prescription is issued should the doctor deem it necessary. Weeks in which people take this opportunity are much more likely to be weeks in which they fall ill than weeks in which they pass it up. In fact on average, 90% of Mondays on which someone *does* go in for a check-up fall in weeks when he or she is subsequently ill; whereas only 10% of Mondays on which someone *doesn't* go for a check-up fall in weeks when he or she is subsequently ill. There is nothing surprising or sinister about this correlation: what explains it is rather the innocuous fact that one is more likely to go for a check-up when one already has reason to think that one will fall ill.

All weekend you have suffered from fainting and dizzy spells. You're pretty sure that there is something wrong. Should you go for the checkup on Monday morning? Clearly if you *are* ill this week, it will be better to have the prescription than not, so the check-up will have been worth your while. But if you are *not* ill this week then the check-up will have been a waste of money. Your payoffs are therefore as stated in the following table, which also gives my abbreviations for the relevant states and acts:

	Well this week (W)	**Ill this week (~W)**
Check-up (C)	1	0
No Check-up (~C)	2	-1

Table 2

Given this table and the statistical facts already mentioned we may compute the average return to going and to not going for a check-up:

(19) AR (C) = 10%. 1 + 90%. 0 = 0.1
(20) AR (~C) = 90%. 2 + 10%. -1 = 1.7

So the average return to going for a checkup exceeds that of not going for a checkup. We may therefore construct the following argument against going for a checkup:

Why ain'cha well.

(21) The average return to going for a check-up exceeds the average return to *not* going for a check-up (*premise: from* (19), (20))
(22) Everyone can see that (21) is true (*premise*)
(23) Therefore going for a check-up is now a foreseeably worse option for you than not going for one (*from* (21), (22))

Should you then not go for your check-up? That would be insane: *of course* you should, given that you have been suffering from dizzy spells and nausea etc. So what is wrong with the argument?

What is wrong with it is the inference from (21) and (22) to (23). Taken over *every* opportunity for a checkup for *anyone*, it is true that those opportunities that are taken shortly precede illness much more often than those that are not taken. But this is not the relevant basis on which *you* should compute the average returns to your options *now*.

What you should rather compute are the average returns to your options *given what you now know about yourself*. That is: you should compute the average returns to C and ~C, not amongst all opportunities for check-ups but amongst *occasions on which the subject is suffering from your symptoms*. That is: you should look at what happens to people when they are suffering from fainting and dizziness. Is subsequent illness *amongst these people on these occasions* any more frequent amongst those who go for checkups than amongst those who do not? Common sense suggests that amongst such people on such occasions, the subsequent incidence of illness is high in both groups and that it is equal in both groups. In that case it is easily verified that:

(24) Amongst people with the symptoms that you now have, the average return to going for a check-up exceeds that of *not* going for a check-up.

So *for you*, *now*, going for a check-up is foreseeably the *better* option.

The fallacy of *Why ain'cha well* is that of applying an overly broad statistical generalization to a single case: in this case, yourself. The generalization is overly broad because it is not limited to cases that resemble yours in relevant respects that you know about. Knowing that you are suffering from dizziness and fainting, the statistical generalization that you should apply to yourself is not (21); it is one that covers only that sub-population that resembles your present stage in that respect i.e. (24). Hence applying (21) rather than (24) to yourself involves a failure to consider evidence that is both relevant and available.

Whatever its other faults *Why ain'cha rich* does not commit *this* error. The inference of (4) from (3), and ultimately from (1) and (2), is not an application of an *overly* broad statistical generalization. Anyone facing Newcomb's problem has *no* evidence that relevantly distinguishes him or her now from anyone else whom the statistical generalization (1) covers, that is, all other persons who ever face this problem. The application of (1) to anyone facing Newcomb's problem is therefore not illegitimate in the way that the application of (21) to your present stage is illegitimate.

What about *Yankees*? It turns out that whether it commits this fallacy depends upon what 'now' in (18) is supposed to denote. Consider first any moment *after* she has learnt whether her next bet will win or lose but *before* she has decided how to bet. It would be fallacious for Mary to apply *Yankees* to herself then, because it would be fallacious for her then to apply (15) to herself. For at any such moment she has relevant information that puts her in a narrower sub-population than that over which (15) generalizes. It puts her not only in the population of bettors but in the sub-population of *winning* bettors (if she has just learnt that she will win), or in the population of *losing* bettors (if she has just learnt that she will lose).

Thus suppose that the predictor has just said to Mary: 'Mary, you will win your next bet.' Then the statistical generalization that she should apply to herself is not the one that compares the average return to placing a bet on the Red Sox and the average return to placing a bet on the Yankees (i.e. (15)). It is the one that compares the average return to placing a *winning* bet on the Red Sox with the average return to placing a *winning* bet on the Yankees. Now we know from Table 1 that the average return to placing a winning bet on the Red Sox is 2 and the average return to placing a wining bet on the Yankees is 1. Hence the appropriate generalization is not (15) but:

(25) The average return to placing a winning bet on the Red Sox exceeds the average return to placing a winning bet on the Yankees.

Inferring (18) ultimately from premises including (15) rather than its opposite from ones including (25) is just the same fallacy as that of *Why ain'cha well*: the fallacy of ignoring available and relevant evidence. So if 'now' in (18) refers to a time *after* Mary learns that she will win her next bet then *Yankees* is fallacious.

With appropriate adjustments the argument of the foregoing paragraph will apply if 'now' in (18) refers to any time at which Mary has just learnt that she will *lose* her next bet. Hence it is fallacious to apply *Yankees* to Mary once she has learnt the outcome of her next bet, *whatever* she has learnt.

What about the time just *before* Mary has learnt the outcome of her next bet? At those times she does not *have* the evidence that is supposed to vitiate the inference from (15) to (18). So isn't the argument then just as plausible as *Why ain'cha rich*?

It's true that it doesn't then commit the *same* fallacy as *Why ain'cha well*. The trouble is that now we cannot infer (18) from (7)-(10) and (17) because it no longer follows from (7)-(10) that EDT recommends betting on the Red Sox. *Before* Mary has learnt whether she will win her bet, the news values of betting on the Red Sox and on the Yankees are:

(26) $V(BR) = 2.Cr(R \mid BR) + -1.Cr(Y \mid BR) = 2.10\% + -1.90\% = -0.7$
(27) $V(BY) = -2.Cr(R \mid BY) + 1.Cr(Y \mid BY) = -2.10\% + 1.90\% = 0.7$

Hence at this time EDT recommends betting on the *Yankees*, so once again its preferred option is the one that foreseeably does better.

Yankees is therefore unsustainable for reasons that do nothing to undercut *Why ain'cha rich*. Neither after nor before Mary has learnt whether her bet is a winner does *Yankees* support an option that diverges from EDT in the way that *Why ain'cha rich* supports an option that diverges from CDT. Not after, because then it doesn't *support* a divergent option; not before, because then it doesn't support a *divergent* option. And this restores the disparity between EDT and CDT. *Why ain'cha rich* does not cut both ways: *it* tells against CDT but not EDT, and no parallel argument tells against EDT but not CDT.

5: Objection: Preference Instability

It was crucial to both parts of our defence of *Why ain'cha rich* that EDT makes *different* recommendations to Mary before and after she has learnt whether her next bet is going to be a winner. *Before* she learns whether it is going to be a winner EDT advises her to bet on the Yankees; *after* she has learnt whether it is going to be a winner and *whatever* she has then learnt, it advises her to bet on the Red Sox.

Now even if you don't object to either recommendation by itself you might object to that discrepancy. Before she learns whether her next bet will win Mary will, if she follows EDT, have a preference for betting on the Yankees—one that she *knows* will be reversed *whatever* information she is about to get. This is surely a strange position to be in.

Strange it may be. But what difference does it make? *Before* she learns whether or not her next bet will win, she can't choose what that bet is. Her early preference for the Yankees is therefore impotent. Unlike, say, the causalist's preference for two-boxing, it confers no *practical* disadvantage upon her. And in decision theory it is only practical disadvantages that are going to carry any weight with practical folk like you and me. But

it was only those folk who were moved by *Why ain'cha rich* in the first place; only those folk, therefore, whose minds Arntzenius's argument ever stood any chance of changing.[2]

In any case the strangeness of this situation is no argument against EDT. *Everyone* is liable to find himself in it, evidentialist or not. A much less fanciful example than Arntzenius's will make this clear.

The admissions statistics for the English and Mathematics Departments at Simpson's Paradox University (SPU) are what you would expect. Male applicants are less successful than female ones overall: 14% of men who apply for admission on to a graduate course in one of these Departments are successful but 20% of women who so apply are successful. But in each Department the discrepancy is reversed: 5% of male applicants for Mathematics are successful as against 1% of female applicants; 50% of male applicants for English are successful as against 25% of female applicants. The explanation is that male candidates are more likely than female candidates to apply for the more competitive Mathematics Department.

Your best friend has just told you that he or she has applied to graduate school at SPU. For some reason it matters greatly to you that your friend's application to this particular university is successful. You know that your friend would have applied to either the English Department or the Mathematics Department but being very absent-minded you have forgotten (a) which of these it is and (b) whether your friend is a boy or a girl. You ask your friend about (a).

Before you hear the answer you reflect that *now*, the news value of the information that your friend is a girl exceeds the news value of the information that your friend is a boy. After all, female applicants to SPU do better than male ones. You then reflect that *after* you have heard the answer to (a) and *whatever that answer is*, the news value of the information that your friend is a *boy* will now exceed the news value of the information that your friend is a *girl*. After all, male applicants to SPU's Mathematics Department do better than female ones, and male applicants to SPU's English Department do better than female ones. Finally, you reflect that now, before you know the answer to (a), you have a preference over the possible answers to (b) that you *know* is going to be reversed in the light of information that you are about to receive.

Nothing in this story mentioned your being an evidentialist. The position that it describes is one that anyone could end up in. The difficulty involved—if that is what it is—is simply that one can sometimes foresee fluctuations in one's preferences. That possibility is independent of whether one is in any position to act upon those preferences.

---

[2] But couldn't we *make* Mary's early and soon-to-be reversed preference for a bet on the Yankees practically potent, indeed practically harmful? Suppose she knew that we were going to offer her: (i) a choice between betting on the Yankees or on the Red Sox *before* she learnt whether her next bet was going to be a winner; and then (ii) the option to switch bets for a fee, *after* she had learnt whether her next bet was going to be a winner. EDT seems to commit her to (i) a bet on the Yankees and (ii) paying the fee—as long as it is less than $1—and betting instead on the Red Sox. But this is irrational: when offered the choice (i) she could *foresee* that she would get information that would lead her to prefer a bet on the Red Sox, so the more rational thing to do would be to take the bet on the Red Sox *then* and save herself the fee.

But if she is going to be offered (i) *and* (ii) then EDT will *not* recommend, at the time of (i), that she take the bet on the Yankees. That recommendation relied on the assumption, implicit in (25), that the news value of a win for the Yankees, given that she bets before learning the outcome of her bet, is 1. But if Mary *knows* that she will change her mind and hence her bet (as she must do for an initial bet on the Yankees to be irrational), then this assumption no longer holds: at the time of (i) the value of a Yankees win given that Mary *now* bets on the Yankees is rather -1, because she knows that when the Yankees win she'll be holding a Red Sox ticket. In fact in that situation EDT *will* prescribe betting early on the Red Sox and saving the fee.

It therefore has nothing especially to do with decision theory. Nor therefore does the foreseeable fluctuation in Mary's preferences over bets raise any special difficulty for the evidentialism that motivates it in that particular case.

References

Arntzenius, F. 2008. No regrets, or: Edith Piaf revamps decision theory. *Erkenntnis* 68: 277-297.

Dummett, M.1964. Bringing about the past. *Philosophical Review* 73, 338–359.

Gibbard, A. and W. Harper 1981. Counterfactuals and two kinds of expected utility. In Harper, W., R. Stalnaker and G. Pearce, eds: *Ifs: Conditionals, Belief, Decision, Chance and Time*. Dordrecht: D. Reidel: 153-192.

Joyce, J. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Lewis, D. 1981a. Causal decision theory. In Gardenfors, P. and N.-E. Sahlin, eds: *Decision, Probability and Utility: Selected Readings*. Cambridge: Cambridge University Press: 377-405.

———. 1981b. Why ain'cha rich? *Nous* 15: 377-80.

Rabinowicz, W. 2002. Does practical deliberation crowd out self-prediction? *Erkenntnis* 57: 91–122.