

NO REGRETS

Or: Edith Piaf revamps decision theory

Frank Arntzenius, Rutgers University

1. Outline

Edith Piaf is famous for her chanson “Non, je ne regrette rien”. I suggest that rational people should not violate Piaf’s ‘No Regrets’ maxim; a rational person should not be able to fore-see that she will regret her decisions. In section 2 I formulate a principle, Desire Reflection, which is a version of Piaf’s maxim. In section 3 I argue that standard evidential decision theory violates this principle. In section 4 I argue that standard causal decision theory does not violate it. In section 5 I discuss whether a couple of variations on these standard decision theories satisfy Desire Reflection. In section 6 I make a suggestion for how causal decision theorists should pick what they consider to be the relevant causal situations. In section 7 I discuss the ‘If you’re so smart, why ain’t cha rich’ objection to causal decision theory, and dismiss it. In section 8 I discuss a more serious problem for causal decision theory, namely ‘Decision Instability’, and argue that it is a real problem. In section 9 I develop deliberational decision theory in order to escape Decision Instability. In section 10 I discuss the connection between deliberational decision theory and game theory. I end with some conclusions.

2. Desire Reflection

Bas van Fraassen has suggested that any rational person P should satisfy the following principle, which I will call “Belief Reflection”: at any time t_1 person P’s credence (subjective probability, degree of belief) in a proposition X should equal P’s expectation value at t_1 for P’s credence in

X at any later time t_2 .

Belief Reflection $Cr_{t_1}(X) = E_{t_1}(Cr_{t_2}(X))$.

Let me now introduce an analogous principle concerning desires, which I will call “Desire Reflection”. Let $Des_t(X)$ denote the desirability of X for person P at time t. What exactly “the desirability of X for person P at time t” should be taken to be is in part what is at issue in this paper. For now, the reader should just take P’s desirability to be a quantity which corresponds to the strength of desire that person P at time t has for X, and take it that numerical representations of desirabilities (for a person at a time) are unique up to positive linear transformations. Desire Reflection then demands that the expectation value of future desirability must equal current desirability

Desire Reflection $Des_{t_1}(X) = E_{t_1}(Des_{t_2}(X))$.

Another way to put this is: one should not be such that one can foresee that one’s future desires will differ from one’s current ones in such a way that one will later regret earlier decisions. So Desire Reflection is just an instance of Pinaud’s ‘No Regrets’ maxim.

Having just stated Desire Reflection, let me immediately make a qualification. One might maintain that cross-time comparisons of desirability make no sense. That is to say, one may think that the question as to whether a claim such as $Des_{t_1}(X) = Des_{t_2}(X)$ is true or not makes no sense, since there are no facts about cross-time equality of desirabilities. For instance, one might think

that the desirabilities at a time t of various options is nothing but a representation of one's preferences at t over these options, so that desirabilities at times are only defined up to positive linear transformation. One might infer from this that the relation between the desirabilities at two different times is only defined up to positive linear transformations. On the other hand, one might claim that cross-time equality of desirability does make sense. For instance, one might claim that the desirability of a complete possible world W has to be equal at all times:

$Des_{t_1}(W) = Des_{t_2}(W)$ for all times t_1 and t_2 and all complete worlds W , and one might then try to argue that it follows that equality of desirability across times in general makes sense. For purposes of this paper I will not take a stand on this issue. I will take Desire Reflection to be the demand that there exist an allowed numerical representation of desirabilities at all times such that $Des_{t_1}(X) = E_{t_1} Des_{t_2}(X)$ for all pairs of times t_1 and t_2 such that t_2 is later than t_1 .

Is Desire Reflection a plausible demand? It seems to me that it is. Violations of Belief Reflection are closely tied to the possibility of being money pumped, and so too are violations of Desire Reflection. For instance, suppose that the default arrangement is that you get orange juice each morning from the stand outside one's front door. But suppose that according to your current desirabilities you would rather get coffee than orange juice two days from now. Indeed, suppose that, other things being equal, you now think it worth 10 cents to get coffee rather than orange juice two days from now. But suppose that you also know that tomorrow you will think it more desirable to get orange juice rather than coffee the day after tomorrow, indeed that tomorrow, other things being equal, you will find the difference worth 10 cents. Then it seems that the merchant at the stand can make money off you by having you pay 10 cents today to switch your

order, for two days from now, to coffee, and then tomorrow have you pay another 10 cents to switch it back. Of course, since all of this is foreseeable to you, you might wisely decide not to pay today to switch, since you can see you will end up switching back anyhow, so it will cost you 20 cents, and you will end up with the default prospect anyhow. Well even if there is such an escape from being money pumped, it still seems that there is something deeply worrying about having such foreseeable switches in desires. Note, for instance, the similarity with problems such as violations of transitivity of preferences. Prima facie it seems irrational to prefer A to B and prefer B to C but prefer C to A. One way of arguing that this is so is to argue that if you violate transitivity you can be money pumped. For suppose the status quo is that you get A. It would seem a merchant can make money of you by first having you pay to switch from A to C, then having you pay money to switch from C to B, and finally having you pay money to switch from B to A. Of course, if you know the merchant is going to make you this sequence of offers, you might wisely refuse the initial offer. Depending a bit on how preferences are supposed to be related to actual choices, this is a quite coherent response. But even if it is, it still seems that there is something deeply worrying about intransitive preferences. Analogously, it seems there is something deeply worrying about violations of desire reflection.

Now, what one thinks of Desire Reflection in general does not matter that much for my current purposes, since for purposes of this paper I will not need to use it in its full generality. I will only impose what I will call 'Weak Desire Reflection'. Weak Desire Reflection is the demand that violations of Desire Reflection should not be brought about merely by conditionalisation upon evidence. This principle seems eminently plausible to me. Let me try to convince the reader of this by giving an additional argument for Weak Desire Reflection.

Suppose you have a friend who has the same initial degrees of belief and the same utilities as you have. Suppose your friend acquires additional information, but he is not allowed to give you the information. He is only allowed to advise you how to act. Surely you should follow his advice. Now consider your future self. If your future self has more information than you do, surely you should listen to her advice, surely you should trust her assessment of the desirabilities of your possible actions. So, surely, your current desirabilities should equal the expectation value of your future desirabilities in cases in which your future self differs from you only in that she has updated her credences on new information. But this is just what Weak Desire Reflection says.

3. Standard evidential decision theory violates Weak Desire Reflection

Let me start by recalling Newcombe's puzzle. Mary has to choose between grabbing box A, and grabbing boxes A and B. Whichever Mary grabs, she gets the contents. Box B contains \$1 for sure. A very reliable predictor P of Mary's choices has put nothing in Box A if P predicted Mary would grab both boxes, and has put \$10 in box A if he predicted Mary would only grab box A.

Standard evidential decision theory¹ says that Mary should maximize her expected utility.

The expected utility of an act A_i is defined as

$$EU(A_i) = \sum_j Cr(O_j/A_i)U(O_j)$$

Here the O_j are the possible outcomes and $U(O_j)$ is the utility of outcome O_j . Let us suppose that

¹ By standard evidential decision theory I mean Jeffrey-style decision theory (see Jeffrey 1983), i.e. a decision theory which uses probabilities of outcomes conditional upon acts, not Savage-style unconditional probabilities (see Savage 1972). I will discuss Savage style decision theory in section 5.

Mary takes it that the predictor is 90% reliable, i.e.

$$\text{Cr}(\$0 \text{ in A/grab both})=0.9$$

$$\text{Cr}(\$10 \text{ in A/grab A})=0.9.$$

And let us suppose that for Mary dollars are utilities. This implies that

$$\text{EU}(\text{grab A})=0.9 \times 10 + 0.1 \times 0 = 9$$

$$\text{EU}(\text{grab both})=0.9 \times 1 + 0.1 \times 11 = 2$$

So, according to evidential decision theory Mary should pick box A.

Now let me slightly modify the case. Suppose that Mary has to make her choice by typing which box, or boxes, she picks into a computer. An hour from now the game-master will read what she typed in, open the relevant box or boxes and give her the money. Suppose also that Mary will be shown the contents of the boxes after she has typed her choice in, but before her instruction is read by the game master. What will her expected utilities be after she has seen the contents? Let us assume that she will believe what she sees, i.e. she will be certain that the amounts that she sees are actually in the boxes, and let us assume that when she sees the contents she will not completely trust her memory, i.e. she will not be completely certain as to which choice she made half an hour ago. There are two possibilities for what she sees.

First let us suppose she sees nothing in box A. Then

$$\text{EU}(A)=\text{Prob}(0 \text{ in A})U(A \ \& \ 0 \text{ in A})+\text{Prob}(10 \text{ in A})U(A \ \& \ 10 \text{ in A})=1 \times 0 + 0 \times 10 = 0$$

$$\text{EU}(\text{both})=\text{Prob}(0 \text{ in A})U(\text{both} \ \& \ 0 \text{ in A})+\text{Prob}(10 \text{ in A})U(\text{both} \ \& \ 10 \text{ in A})=1 \times 1 + 0 \times 11 = 1$$

So in that case EU(grab both) is higher.

Suppose, on the other hand, she sees \$10 in A. Then

$$\text{EU}(A)=\text{Prob}(0 \text{ in A})U(A \ \& \ 0 \text{ in A})+\text{Prob}(10 \text{ in A})U(A \ \& \ 10 \text{ in A})=0 \times 0 + 1 \times 10 = 10$$

$$\begin{aligned} EU(\text{both}) &= \text{Prob}(0 \text{ in } A)U(\text{both} \ \& \ 0 \text{ in } A) + \text{Prob}(10 \text{ in } A)U(\text{both} \ \& \ 10 \text{ in } \\ & A) = 0 \times 1 + 1 \times 11 = 11 \end{aligned}$$

So, in that case too, $EU(\text{grab both})$ is higher.

So, no matter what information she gets, after she updates on that information $EU(\text{both})$ is higher. And Mary knows this beforehand. That is to say, at the earlier time $EU(\text{both})$ is lower, and yet she knows at the earlier time that at the later time $EU(\text{both})$ will be higher. So she violates Weak Desire Reflection.

Now, you might object that the later EU 's make no sense since it makes no sense to attribute expected utilities to past acts. Let me respond to that in two ways. In the first place, formally, they do. Mary has degrees of belief and utilities at the later time, so her EU 's, as defined by the usual formula, exist.² In the second place, one can come up with scenarios in which the expected utilities of future actions changes due to information that one gets before the actions have taken place, and one can come up with scenarios like that in which desire reflection will fail for evidential decision theorists.³

Formally, what is going on is the following. According to evidential decision theory, one should maximize expected utility $EU(A_i) = \sum_j \text{Cr}(O_j/A_i)U(O_j)$. When one gets new information, one's credences are updated by conditionalisation. So one's conditional credences $\text{Cr}(O_j/A_i)$ will change. How will they change? Well, that depends on what bit of information one gets. Before

² One might prefer to think of them as her expected utilities for news concerning her actions, rather than for her actions per se, but still, they exist.

³ I didn't do it that way in the current example, since, once the contents of the boxes has been seen the EDT'r and the CDT'r make the same choices, so that it would no longer be an example of the CDT'r and EDT'r making different choices.

one gets the new bit of information one can not know what one's future credences will be, but one can calculate the expectation value of one's future credences. What is the expectation value of one's future conditional credences? The expectation value of future one's future conditional credence is

$$E(\text{Cr}'(O_j/A_i)) = \sum_k \text{Cr}(I_k) \text{Cr}(O_j/A_i \& I_k).$$

Here $\text{Cr}(I_k)$ is one's current credence that one will get information I_k , and the set of the possible bits of information $\{I_k\}$ forms what I will call an 'information partition'. Now, this expectation value does not generally equal one's current conditional credence, i.e. it is not generally true that $E(\text{Cr}'(O_j/A_i)) = \text{Cr}(O_j/A_i)$, and this means that the expectation value of future expected utility does not generally equal current expected utility. If one satisfies Belief Reflection, then the current expectation value of one's future unconditional probabilities must equal one's current unconditional probabilities. But this is not true of conditional probabilities. That is why Weak Desire Reflection is violated by evidential decision theory.

4. Standard causal decision theory satisfies Weak Desire Reflection

Causal decision theory, intuitively speaking, says the following: when deciding what to do, pay attention to the extent to which your actions cause good or bad outcomes, don't pay attention to the extent to which your actions are evidence for good or bad outcomes. More formally, standard causal decision theory as formulated by David Lewis, says the following.⁴ Start by figuring out what all the possible causal situations K are that you could be in. Call the set of possible causal

⁴ In the next section I will discuss Joyce-style Causal Decision Theory which makes use of probabilities that are "imaged" on acts, rather than probabilities of causal situations.

situations $\{K\}$ the “causal situation partition”. Then for each possible causal situation K in the partition, figure out what the expected utility

$$EU_K(A) = \sum_j Cr(O_j/A \& K) U(O_j)$$

of each of your acts is. Then, since you typically do not know in which causal situation you are in, average out these expected utilities weighted by your credences in the causal situations in order to get the causal utilities

$$CU(A) = \sum_K Cr(K) EU_K(A)$$

of your acts. Causal decision theory says that you should perform an act that has maximal causal utility.

Now, different people have slightly different ideas as to what counts as a causal situation. The basic idea is that a causal situation determines exactly how the future outcomes causally depend on your acts. More specifically, according to some it is a list of counterfactuals, according to some it is a list of laws of development and facts about the current state, and there are yet other ideas. I will not take a stand on this issue, since it does not matter much for my purposes exactly what a causal situation is.

We know enough to check whether causal decision theory satisfies Weak Desire Reflection. Let's start with the standard Newcombe case. In this case all hands agree that the possible causal situations are: “There is \$0 in box A” and “There is \$10 in box A”. Now

$$EU_{\$0}(\text{both}) = 1$$

$$EU_{\$0}(A) = 0$$

$$EU_{\$10}(\text{both}) = 1$$

$$EU_{\$10}(A) = 10$$

So, in each possible causal situation $EU(\text{both})$ is higher than $EU(A)$. So, no matter what one's credence in each possible causal situation, $CU(\text{both})$ is higher than $CU(A)$. So causal decision theory says: pick both.

How about Weak Desire Reflection? Well, notice that there is no possible bit of information that can affect the expected utilities of any particular act in any particular causal situation: every act-causal situation pair uniquely determines an outcome. The only thing that information can do is change one's (unconditional) credences in the possible causal situations. But, by Belief Reflection one's current expectation of one's future (unconditional) credences must equal one's current (unconditional) credence. It immediately follows that the expectation value of the future causal utility of any act must equal the current causal utility of that act. So causal decision theorists in a Newcombe situation must satisfy Weak Desire Reflection.

Let me generalize this point. It is obviously generally true that if every act-causal situation pair uniquely determines an outcome, then Weak Desire Reflection is satisfied by a causal decision theorist. So, in so far as one agrees that one ought not to violate Weak Desire Reflection one has reason to have a conception of causal situations which is fine-grained enough that any possible causal situation K in conjunction with any possible act A uniquely determines the outcome.

However, sometimes, one can not plausibly have that. For instance, suppose that one thinks that the world is indeterministic. Then, it might be that a causal situation in conjunction with an act does not uniquely determine the outcome. Even then, though, one can satisfy Weak Desire Reflection relative to any information that one thinks one could possibly acquire prior to one's act.

Let me be explain. Suppose that one takes as one's causal situations K certain features of the world up until the time t that the actions are to take place, and that the K are such that conjoined with any action A they determine the chances of each possible outcome, and that one knows what these chances are. I.e. suppose one has credence 1 in a particular conditional chance function, $\text{Chance}(O_i/K\&A)$, and that one has credence 1 that

$$\text{Chance}(O_i/K\&A)=\text{Chance}(O_i/K\&A\&X)$$

for any proposition X about the state of the world up until time t . Now, the so-called Principal Principle (another eminently plausible principle of rationality) says that one's credences ought to equal what one takes to be the chances. That is to say, the Principal Principle dictates in such a situation that:

$$\text{Cr}(O_i/K\&A\&X)=\text{Chance}(O_i/K\&A\&X)=\text{Chance}(O_i/K\&A)=\text{Cr}(O_i/K\&A).$$

Now let's look at Weak Desire Reflection. What is the expectation value of your causal utilities conditional upon some information partition $\{I_m\}$? Well, your causal utilities are

$$\text{CU}(A)=\sum_{K,j}\text{Cr}(K)\text{Cr}(O_j/A\&K)U(O_j).$$

So your expected causal utilities conditional upon $\{I_m\}$ are

$$\text{ECU}(A)=\sum_{K,j,m}\text{Cr}(I_m)\text{Cr}(K/I_m)\text{Cr}(O_j/A\&K\&I_m)U(O_j).$$

Now suppose that for each K, j term separately we could show that

$$\sum_m\text{Cr}(I_m)\text{Cr}(K/I_m)\text{Cr}(O_j/A\&K\&I_m)U(O_j)=\text{Cr}(K)\text{Cr}(O_j/A\&K)U(O_j).$$

Then it would follow that

$$\sum_{K,j,m}\text{Cr}(I_m)\text{Cr}(K/I_m)\text{Cr}(O_j/A\&K\&I_m)U(O_j)=\sum_{K,j}\text{Cr}(K)\text{Cr}(O_j/A\&K)U(O_j)$$

So we would have shown that

$$\text{ECU}(A)=\text{CU}(A).$$

Well, let us suppose that each I_m is a bit of information about the state of the world prior to t .

Then we know that

$$\text{Cr}(O_j/A\&K\&I_m)=\text{Cr}(O_j/A\&K).$$

But if that is the case then

$$\sum_m \text{Cr}(I_m)\text{Cr}(K/I_m)\text{Cr}(O_j/A\&K\&I_m)U(O_j)=\sum_m \text{Cr}(I_m)\text{Cr}(K/I_m)\text{Cr}(O_j/A\&K)U(O_j).$$

But we know by Belief Reflection that

$$\sum_m \text{Cr}(I_m)\text{Cr}(K/I_m)=\text{Cr}(K).$$

So we have that

$$\sum_m \text{Cr}(I_m)\text{Cr}(K/I_m)\text{Cr}(O_j/A\&K\&I_m)U(O_j)=\text{Cr}(K)\text{Cr}(O_j/A\&K)U(O_j).$$

So in that case one must satisfy Weak Desire Reflection relative to any information one could possibly get prior to one's action.

So, once again, in so far as one has reason to satisfy Weak Desire Reflection, one has reason to have a conception of causal situations K which is fine-grained enough that any K conjoined with an act A determines the chances of the outcomes O_j .

5 Joyce-style and Savage-style decision theories and Weak Desire Reflection.

Let me now very briefly discuss a couple of variations on what I have called standard evidential decision theory and standard causal decision theory, and see how they fare with respect to Weak Desire Reflection.

Jim Joyce (in Joyce 1999) has suggested defining the causal utility of an act as follows:

$$CU(A)=\sum_w \text{Cr}^A(W)U(W)$$

Here we are summing over possible worlds W , and $\text{Cr}^A(-)$ is one's credence 'imaged' on act A .

What is it to ‘image’ a credence Cr on an act A , or, more generally, on an arbitrary proposition C ? The basic idea is as follows: when ‘imaging’ on proposition C you are supposing that C is true and adjusting your credence in the most plausible way to accommodate that supposition. Obviously, you will have to shift all your credence onto worlds in which C is true. But you also want to make C true in the most ‘plausible’ way. A natural way to implement this is as follows. When considering what to do with your credence in a world W in which C is false, you want to shift your credence in W to the closest possible world in which C is true. Now there may not be a unique closest such world, so you will in general have to shift your credence in W to some set of worlds close to W in which C is true. In order for this to have a well-defined effect, you will have to specify how your credence in W gets apportioned among the worlds in this set. Such a proportionality assignment can formally be represented by a function $\rho^C(X,W)$ which tells you what proportion of your credence in world W will be attached to worlds in which proposition X is true when you image your credences on proposition C . It follows that when you image some initial credence Cr on C , then after imaging your credence in proposition X will be

$$Cr^C(X) = \sum_w Cr(W) \rho^C(X,W).$$

Now let us see whether the expectation value of Joyce-style causal utility under an information partition equals one’s current causal utility. Well, the expectation value of Joyce-style causal utility on information partition $\{I_m\}$ is

$$ECU(A) = E(\sum_w Cr(W) \rho^C(X,W) U(W)) = \sum_{w,m} Cr(I_m) Cr(W/I_m) \rho^C(X,W) U(W),$$

provided that the function $\rho^C(X,W)$ does not change when you acquire new information. In that case it follows from Belief Reflection that

$$ECU(A) = \sum_{w,m} Cr(I_m) Cr(W/I_m) \rho^C(X,W) U(W) = \sum_w Cr(W) \rho^C(X,W) U(W) = CU(A).$$

So, if the ‘way’ in which you shift your credences around when you image on an action A does not change when you acquire new information, it follows that a Joyce-style causal decision theorist will satisfy Weak Desire Reflection. So, in so far as one agrees that one has reason to satisfy Weak Desire Reflection, one has reason to keep $\rho^C(X,W)$ invariant under information acquisition.

Now let’s turn to Savage-style decision theory (see Savage 1972). In Savage-style decision theory one assumes that one has divided propositions up into three distinct categories, those describing actions A , those describing situations S , and those describing outcomes O . The actions A amount to maps from situations S to outcomes O , i.e. each action-situation pair determines a unique outcome.⁵ The Savage-utility of an act then is defined as $EU(A)=\sum_S Cr(S)U(O)$, where O is the outcome determined by A and S . Since only unconditional credences $Cr(S)$ occur in this formula, it follows from Belief Reflection that Savage-style decision theory satisfies Weak Desire Reflection. So, why not simply adopt Savage-style decision theory? Why bother with the somewhat baroque business of causal decision theory? Well, in effect causal decision theory, Lewis-style, just is a Savage-style decision theory, one with a particular suggestion for what the situations S should be, namely causal situations. Let me

⁵ One might balk at the demand that such pairs always uniquely determine an outcome. What if the world is indeterministic? Well, one can, by convention, simply beef up the notion of a situation so that any situation S conjoined with any act A uniquely determines an outcome O . Of course, it is then not plausible that what situation one is in is determined by the intrinsic state of the world at the time of one’s action. Still, one might hope that this does not harm the applicability or plausibility of Savage-style decision theory. Unfortunately, we will soon see that there is a problem as to what one takes situations to be, indeed that without a further constraint on what can count as a situation Savage-style decision theory is unsatisfactory.

now quickly explain why one has to do something like that, why one can not just take Savage-style decision theory, and put no constraints on what the situations S are.

Consider the following case. Harry is going to bet on the outcome of a Yankees versus Red Sox game. Harry's credence that the Yankees will win is 0.9. He is offered the following two bets, of which he must pick one:

- 1) A bet on the Yankees: Harry wins \$1 if the Yankees win, loses \$2 if the Red Sox win
- 2) A bet on the Red Sox: Harry wins \$2 if the Red Sox win, loses \$1 if the Yankees win.

Suppose Harry takes the possible situations to be:

- A) Yankees win
- B) Red Sox win

Then

$$EU(1): 0.9x1 + 0.1x-2 = 0.7$$

$$EU(2) = 0.1x2 + 0.9x-1 = -0.7$$

So, unsurprisingly, Savage-style decision theory says: Harry pick the bet on the Yankees.

But now suppose that Harry take the possible situations to be:

- C) I win my bet
- D) I lose my bet

Then

$$EU(1) = Cr(C)x1 + Cr(D)x-2$$

$$EU(2) = Cr(C)x2 + Cr(D)x-1.$$

It immediately follows that no matter what $Cr(C)$ and $Cr(D)$ are, $EU(2) > EU(1)$. So now Savage-style decision theory says: Harry should pick the bet on the Red Sox. In short, what action

Savage-style decision theory mandates depends on what one takes to be the situations. To put it another way: Savage style decision theory is not 'partition invariant'. As it stands, Savage-style decision theory is incoherent.

Jeffrey-style evidential decision theory and causal decision theory amount to two different reactions to this problem. Causal decision theory adds the rider that one should take as one's situations causal situations. Jeffrey-style evidential decision theory says: don't use unconditional probabilities, instead use probabilities of situations conditional upon acts. In other words, Jeffrey-style evidential decision theory solves the problem by becoming partition invariant, while causal decision theory solves the problem by saying that there is a particular privileged partition.

Now, I have argued that evidential decision theory is bad because it violates Weak Desire Reflection, and have argued that causal decision theory satisfies it in so far as it makes use of fine-grained partitions. However, I don't want to say that any fine-grained partition is as good as any other. Let me now briefly address that issue.

6. Good and bad partitions

While one avoids violations of Weak Desire Reflection if one has fine-grained causal situations, this does not mean that on my view any fine-grained partition is as good as any other. Recall the example I just gave, and imagine a causal decision theorist who is repeatedly going to be on a series of Yankees-Red Sox games. A causal decision theorist who takes as his causal situations A and B will always bet on the Yankees. And if the Yankees indeed win 90% of the time, he will gain an average of 70 cents per bet. On the other hand, a causal decision theorist

who takes as his causal situations C and D will always bet on the Red Sox, and will lose an average of 70 cents per bet. Clearly then, there is a sense in which causal partition $\{C,D\}$ is a bad partition, and $\{A,B\}$ is a good partition, and this is fore-seeably so.

But why is this so? What is the diagnosis? The diagnosis is that $\{C,D\}$ is a bad partition because its elements C and D are correlated to the acts that one is choosing between, while the elements of the $\{A,B\}$ partition are not. Here is what I mean. Let's suppose that we have a large population of causal decision theorists, that 50% of them use the $\{A,B\}$ partition, and 50% of them use the $\{C,D\}$ partition. So 50% will bet on the Yankees all the time, 50% will bet on the Red Sox all the time. Suppose that the Yankees win 90% of the time. Then there will be no correlation between the elements of the act partition ('Bet on Yankees', 'Bet on Red Sox') and the elements of the $\{A,B\}$ partition. But there will be a correlation between the elements of the act partition ('Bet on Yankees', 'Bet on Red Sox') and the elements of $\{C,D\}$ partition. For, among all the bets that are placed on the Yankees, 90% of them are placed in cases in which the Yankees win, and among all the bets that are placed on the Red Sox, also 90% of them are placed in cases in which the Yankees win. So there is no correlation between the act partition and the $\{A,B\}$ partition. However, among of all the bets that are placed on the Yankees 90% are cases in which that bet wins, while among all the bets that are placed on the Red Sox, only 10% are cases in which that bet wins. This is why the $\{C,D\}$ partition is a bad partition. It is not a good idea for a rational person to determine what she should do by first looking at how she should act in case 'my bet wins' is true and then look at what how she should act when 'my bet loses' is true, and then average these out by her unconditional probabilities, since the likelihood of being in a particular element of that partition is correlated to the action in question.

Let me generalize this argument. The actual average value of an act A,

$$Av(A) = \sum_i rf(O_i/A)U(O_i).$$

Here rf stands for actual relative frequency. If there is no correlation (according to the actual relative frequencies) between A and the elements K_j of a partition $\{K_j\}$, then

$$Av(A) = \sum_i rf(O_i/A)U(O_i) = Av(A) = \sum_i rf(K_j)rf(O_i/A \& K_j)U(O_i).$$

So, if there is no correlation between the act partition and the elements of person P's causal partition $\{K_j\}$, and if the credences of P match the actual relative frequencies, then if P maximizes his causal utility, he will be maximizing his average value. But if there is such a correlation, then P will, typically, not be maximizing his average value. So it is not a good idea to use a causal partition which in fact is correlated to the act partition.

Now, I do not in general want to argue that rational people will always maximize average value. In first place they need not have the right credences. So we should at least weaken the demand to: by the light of their own credences there had better be no correlation between the elements of their causal partition and their act partition. But even that is too strong, I have already argued that in the 2-box case one should be a causal decision theorist and that the appropriate causal partition in that case is '\$0 in A', '\$10 in A'. But in that case one does expect that if there are people that choose differently, then there will be a correlation between the acts and the element of the causal partition. And yet I maintain that that is the right causal partition. So I don't want to claim that rational people should always choose a causal partition which they expect to be uncorrelated to their act partition. No, what I suggest demanding is that in the vast majority of cases, the cases where one can avoid Weak Desire Reflection violation and yet act just as an evidential decision theorist, one should do so, i.e. one should in normal circumstances

choose a causal partition such that one expects that the act partition is uncorrelated to the causal partition.

7 If you're so smart, why ain't cha rich?

Now let me briefly consider, and dismiss, a well-known objection to causal decision theory, namely the 'if you're so smart, which ain't cha rich' objection. In a Newcombe 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 it 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

Let me combat this argument by giving an example where the shoe is on the other foot. We can get such a case by slightly modifying the Yankees-Red Sox example.

Let Mary be an evidential decision theorist (Jeffrey-style) who is to bet repeatedly on the outcome of a sequence of Yankees versus Red Sox games. As in the earlier case, Mary is convinced that in the long run the Yankees will win 90% of the time. However, on each occasion just before she chooses which bet to place, a perfect predictor of her choices and of the outcomes of the games announces to her whether she will win her bet or lose it. Here is what ensues.

On some occasions the predictor says "Mary, you will lose your next bet". After Mary has updated her credences on this information she calculates:

$$EU(\text{bet Y}) = Cr(Y \text{ win} / \text{bet Y})U(Y \text{ win} \ \& \ \text{bet Y}) + Cr(Y \text{ lose} / \text{bet Y})U(Y \text{ lose} \ \& \ \text{bet Y}) = -2$$

$$EU(\text{bet R}) = Cr(R \text{ win} / \text{bet R})U(R \text{ win} \ \& \ \text{bet R}) + Cr(R \text{ lose} / \text{bet R})U(R \text{ lose} \ \& \ \text{bet R}) = -1$$

So she bets on the Red Sox each time she is told she will lose her bet.

On the other occasions she is told “you will win you next bet”. She updates her credences and finds:

$$EU(\text{bet Y}) = Cr(Y \text{ win}/\text{bet Y})U(Y \text{ win} \& \text{bet Y}) + Cr(Y \text{ lose}/\text{bet Y})U(Y \text{ lose} \& \text{bet Y}) = 1$$

$$EU(\text{bet R}) = Cr(R \text{ win}/\text{bet R})U(R \text{ win} \& \text{bet R}) + Cr(R \text{ lose}/\text{bet R})U(R \text{ lose} \& \text{bet R}) = 2$$

So she also bets on the Red Sox each time she is told she will win her bet. 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 can not do this.

A causal decision theorist, on the other hand, who uses the “Yankees win”, “Yankees lose” partition, will will always bet on the Yankees, even given the “you win” “you lose” information.⁶ For:

$$CU(\text{bet Y}) = Cr(Y \text{ win})EU_{Y\text{win}}(\text{bet Y}) + Cr(Y \text{ lose})EU_{Y\text{lose}}(\text{bet Y}) = 0.9 \times 1 + 0.1 \times -2 = 0.7$$

$$CU(\text{bet R}) = Cr(Y \text{ win})EU_{Y\text{win}}(\text{bet R}) + Cr(Y \text{ lose})EU_{Y\text{lose}}(\text{bet R}) = 0.9 \times -1 + 0.1 \times 2 = -0.7.$$

So there are cases in which causal decision theorists, predictably, will do better than evidential decision theorists.

Let me give one more example. Consider again a Newcombe situation. Now suppose that the situation is that one makes a choice after one has seen the contents of the boxes, but that the predictor still rewards people who, insanely, choose only box A even after they have seen the

⁶ I am assuming that he does not know he is a causal decision theorist, indeed that he has no credences about which bet he will take out when he calculates his causal utilities.

contents of the boxes. What will happen? Evidential decision theorists and causal decision theorists will always see nothing in box A and will always pick both boxes. Insane people will see \$10 in box A and \$1 in box B and pick box A only. So insane people will end up richer than causal decision theorists and evidential decision theorists, and all hands can foresee that insanity will be rewarded. This hardly seems an argument that insane people are more rational than either of them are. Let me turn to a better argument against causal decision theory.

8 Decision instability problems for causal decision theory

Johnny has a button in front of him such that if he presses it all psychos will die.⁷ Other things equal he wants to kill all psychos, but other things are not equal: he really does not want to die himself. Other than that he does not care whether he is a psycho. Indeed, let us suppose that his utilities are as follows:

$$U_{\text{psycho}}(\text{push}) = -100$$

$$U_{\text{psycho}}(\text{not-push}) = 10$$

$$U_{\text{not-psycho}}(\text{push}) = 0$$

$$U_{\text{not-psycho}}(\text{not-push}) = 0$$

Johnny thinks it very unlikely that he is a psycho:

$$\text{Cr}(\text{I am a psycho}) = 0.01$$

But he thinks that pushing the button would be good evidence that he is a psycho:

$$\text{Cr}(\text{I am a psycho} / \text{I push}) = 0.9$$

⁷ This example is from Egan (forthcoming). Similar examples can be found in Weirich (1985), Gibbard and Harper (1978) and Richter (1984).

$$Cr(I \text{ am a psycho}/I \text{ don't push})=0$$

. What should he do?

Let's assume that the causal situations are 'Johnny is a psycho' and 'Johnny is not a psycho'. Let Johnny's utilities be as follows:

$$U_{\text{psycho}}(\text{push})= -100$$

$$U_{\text{psycho}}(\text{not-push})= 10$$

$$U_{\text{not-psycho}}(\text{push})= 0.9$$

$$U_{\text{not-psycho}}(\text{not-push})= 0$$

It follows that $CU(\text{push})=0.01x-100+0.99x10= 8.9$, and $CU(\text{not-push})=0$.

Note that as far as this calculation is concerned his credence in being a psycho conditional on pushing and conditional on not pushing are irrelevant. Some people's intuitions say that this should not be irrelevant, indeed that in this situation it would be unwise for Johnny to push, since if he were to decide to push he would have good reason to believe that he is a psycho, and hence he should think it likely that he will die. Now, I don't have strong intuitions here. Moreover, the case does not seem that different from Newcombe type cases, so that I can certainly see a causal decision theorist maintaining that since pushing is merely evidence for Johnny being a psycho, and does not cause him to be a psycho, Johnny should ignore his conditional credences and indeed push.

However, it seems to me that there is a different objection to causal decision theory that this case elicits, namely a violation of Edith Piaf's maxim which does not occur in the standard Newcombe case. Here is what I mean. Suppose that Johnny, being a good causal decision theorist has made the decision to push. Suppose he has not pushed yet, but he is sure he will do

so. Since Johnny has a credence of 0.9 in being a psycho conditional on pushing, once he has become convinced that he will push, he should increase his degree of belief that he is a psycho to 0.9. But now, given his updated credences, he can re-calculate the causal utilities of pushing and of not pushing, and he will find that now the causal utility of not pushing is significantly higher than that of pushing. To be precise, now

$$CU(\text{push})=0.9 \times -100 + 0.1 \times 10 = -89$$

$$CU(\text{not-push})=0.$$

So, as soon as he has made his decision and has updated his credences on the basis of his decision, he will regret it. So he will violate Piaf's maxim. Indeed, no matter what decision he makes, as soon as he has incorporated that decision in his credences, he will regret the decision. So a causal decision theorist in a Psycho Johnny situation can not but violate Piaf's maxim.

What to do? ⁸

9 Deliberational decision theory

Here is what I suggest. Allow what game theorists have been allowing for a long time: so-called

⁸ Can one escape this problem by adopting Jeffrey-style evidential decision theory? That depends. If one assumes that all that one learns when one has come to a decision is what one will do, and if one always updates by conditionalisation, and if even in the extreme case where one knows which act one will perform with certainty the conditional probabilities upon not doing that act are still well-defined (i.e. if one allows conditionalisation on credence 0 propositions), then one will not regret one's decision once it is made. But if not, then even the evidential decision theorist can be in a situation where he must regret his decision as soon as he has made it. The crux here is whether the conditional credences $Cr(O_i/A)$ can change as a result of deliberation. Note that this possibility is exactly what Jeffrey invoked when he gave his 'ratificationism' defense of evidential decision theory in Newcombe type cases (see Jeffrey 1983, section 1.7), and what others invoked when they gave a 'tickle' defense of evidential decision theory in gene G type cases (see e.g. Eells 1982).

‘mixed decisions’, i.e. decisions to do certain acts with certain probabilities. Now, how one should understand the notion of a ‘mixed decision’ is a somewhat vexing issue. I will sketch the view of mixed decisions that I am inclined towards, while at the same time admitting that I am not entirely convinced that this is the right view to take. Indeed I am not completely happy about having to resort to ‘mixed decisions’ at all. The reason that I nonetheless will do so is that armed with it one can give a natural solution to problems of decision instability (such as are brought out by Psycho Johnny type cases) and that it will also allow a natural unification of decision theoretic rationality and game theoretic rationality. This gives me some hope that something fairly close to what I suggest in this paper could be right.

OK, so here is the view of mixed decisions that I incline towards. My tentative view is that to make a certain mixed decision is just to have certain credences in one’s acts at the end of a rational deliberation. On this view, mixed decisions are not decisions to perform certain acts with certain probabilities. Now, my tentative view is admittedly an odd view. For on this view rationality constrains what credences one should have at the end of a deliberation. Decision theory, on this view, does not evaluate the rationality of actions. Rather it evaluates the rationality of credences in actions. On this understanding of mixed acts decision theory is theory of what credences one ought to have in one’s actions, it is not a theory that tells one which actions are rational and which are not, nor does it even evaluate how rational each possible act is. This, I admit, is odd, to say the least. So why do I adopt this view?

Here’s why. The natural alternative view is that a mixed decision is a decision to perform certain acts A_i with a certain probabilities p_i . But what is it to decide to perform certain acts A_i

with a certain probabilities p_i ? A natural suggestion would be that one does this just in case one has a chance device at one's disposal, where one can delegate responsibility of which act is to be performed to this chance device, and one can set the chances with which this chance device will act on one's behalf to values p_i of one's choosing. However, in the first place we are hardly ever in a situation in which we can perform such actions. (It is not as if one has such a chance device stored away in some convenient part of one's brain.) In the second place, even if we did it would amount to a different decision situation, namely one in which we have an uncountable infinity of pure acts that we can perform, the acts being the possible ways we have of setting the chance 'dials' of the chance device. In short, I find the natural alternative view implausible and/or useless. So let me soldier on with my view, while noting that I am not entirely happy with it.

Let me now show how armed with mixed decisions we can make sure we can stop violating Piaf's maxim. Before I do so, let me point out that the rest of this paper amounts to little more than an exposition of some of Brian Skyrms' work.⁹

The basic idea is very simple. Johnny starts his deliberation with certain credences that he is a psycho and certain credences that he will push. He then does a causal utility calculation. He finds the causal utility of pushing is higher than that of not pushing. But rather than that he immediately becomes certain that he will push, he merely increases his credence that he will push. In the light of this he increases his credence that he is a psycho. Armed with these new credences he recalculates causal utilities. Armed with the result he resets his credences in pushing the button. And so on, until he reaches credences which are at equilibrium, i.e. credences which are such that his own rules tell him not to change them anymore. Given the

⁹ See, especially, Skyrms 1990.

utilities and conditional credences that I stated, this will occur when $Cr(\text{Push})=10/110$, $Cr(\text{Psycho})=9/110$. I claim that if Johnny is rational, these are the credences that he will have at the end of his deliberation.¹⁰

In general I claim that a rational person is one such that at the end of his deliberation as to which action to perform his credences are in equilibrium. Let me dub this theory ‘deliberational decision theory’. Let me also define ‘deliberational causal decision theory’ and ‘deliberational evidential decision theory’. Deliberational causal decision theory additionally demands that one’s update rule is such that it increases the credence only of actions whose causal utility is greater than the status quo, and it raises the sum of the probabilities of all acts with causal utility greater than the status quo. Deliberational evidential decision theory demands that one’s update rule be such that it increases the credence only of actions whose evidential utility is greater than the status quo and it raises the sum of the probabilities of all acts which with evidential utility greater than the status quo. I advocate being a deliberational causal decision theorist.

Now let me make some clarifications. First clarification: how does one update one’s credences in situations given how one has just updated one’s credences in one’s actions? The natural way to do this is to keep one’s credence in any proposition X conditional upon one’s actions fixed.¹¹ But I do not want to make this a definitional part of being a deliberational decision theorist. I will only demand that one has some well-defined rule for updating one’s

¹⁰ A similar idea can be found in Weirich 1985.

¹¹ That is to say, the most natural way in which to do this is to ‘Jeffrey-conditionalise’ relative to the action partition.

credences in the light of one's updated credences in one's actions, and that this update leaves invariant one's (updated) credences in one's actions. In fact, the updating procedure should not really be conceived as consisting of two separate updates. Rather, one uses one's utility calculation in order to reset one's credences in one's actions, and one simultaneously adjusts one's overall credences 'accordingly'. As I said before, the most natural way of adjusting one's overall credences 'accordingly' is to keep fixed one's credences conditional upon one's actions, but I do not want to make this a definitional part of being a deliberational decision theorist; I will only demand that one has some well-defined rule for what it is to change one's overall credences 'accordingly'.

Second clarification. Will there always be equilibrium credences? Yes, given certain mild assumptions, there will always be equilibrium credences. Suppose that one's updating rule is continuous, i.e. that any continuous variation in the credences to which one's updating rule is applied induces a continuous variation in the updated credences. Slightly more formally: an updating rule is a function U from credence distributions over actions and situations to credence distributions over actions and situations. The demand for continuity is just the demand that this function U be continuous.¹² This is a very mild demand. It will be satisfied if one updates credences in actions via an expected utility rule, and if one then updates one's overall credences by holding one's credences conditional upon one's actions fixed.¹³ Next let us

¹² For a function to be continuous its domain and range need to have a topology. If the credences are distributed over N actions and situations, where N is finite, we can represent a credence distribution as a vector in an N -dimensional vector space, and use the natural topology of this N -dimensional vector space.

¹³ It does rule out cases where you believe something like this: a perfect predictor rewards you iff you make a non-trivial mixed decision, e.g. the predictor pays you big bucks if

suppose that the set of possible actions is finite, and that the set of possible situations is finite. Then the action probability space and situation probability space will be finite dimensional, convex and compact.¹⁴ Brouwer's fixed point theorem says that any continuous map M from a finite dimensional, convex and compact space to itself must have at least one fixed point, i.e. there must be at least one point x in the space such that $x=M(x)$. So, assuming continuity of the update rule, it follows that there must be at least one equilibrium credence distribution.

Next clarification. Must one really model a rational person as a deliberator who changes his credences during the deliberation? No, one need not. Indeed it is a little bit awkward to do so. After all, if one is ideally rational, then how could there be any stage at which one has the 'wrong' credences? So, as long as we are idealizing, let us simply say that a rational person must always be in a state of deliberational equilibrium. The dynamical model of deliberation that I gave can be taken to merely amount to a crutch to make us realize that there always exists a state of deliberational equilibrium.¹⁵

your final credences in your actions are not all 1 or 0, and pays you nothing if your final credences are all 1 or 0. In fact, I have only been considering cases in which you believe that what the predictor does depends on how you act, rather than on your (final) credences. I can in fact allow for predictors whose actions depend on the decision maker's (final) credences rather than on his acts. What I can not allow, while retaining the necessary existence of equilibria, is a discontinuous dependence of the predictors actions on one's credences. Luckily such a discontinuous dependence of the predictors actions on one's credences can plausibly be ruled out. For it is very plausible that the dynamical laws of our world are continuous, and no physical system whose dynamics is continuous could implement such a predictor.

¹⁴ A space is convex if the line between any two points in the space is in the space. A space is compact if it is closed and bounded.

¹⁵ In order to be able to define what it is for a credence distribution to be an equilibrium distribution, we will need to assume that even though a perfectly rational person never actually uses any update rules to update his credences during deliberation, there still are facts about what these update rules are, or rather, should be. For otherwise we do not have enough structure to define what his equilibrium credences are.

Final clarification. What if there is more than one deliberational equilibrium? If one took the dynamical model seriously it could be the case that the equilibrium state that a rational person ends up in would depend on his starting point, i.e. his initial credences, and the exact form of his update rules. Moreover, it could be that he ends up in an equilibrium state which has lower utility associated with it than some other equilibrium state. This would seem a bit unfortunate for an ideally rational person. However, given that I have just suggested a non-dynamical model, I now suggest adding the rider that an ideally rational person is always in an equilibrium state such that there is no other equilibrium state which has higher utility. That is all there is to deliberational decision theory. Now let me turn to game theory.

10 Deliberational decision theory and game theory

Let me start by sketching game theory very quickly and superficially. A game is a situation in which each player has to decide which act to perform, and what payoff each player gets depends on which acts all players perform. It is common knowledge what the payoffs are for each combination of acts. A Nash equilibrium is a set of mixed decisions, such that no player can benefit by unilaterally deviating from his mixed decision. One can show that there are always Nash equilibria in the same way that I argued that there always are deliberational equilibria. The standard game theoretic account of rationality is that rational players must play their part of a Nash equilibrium. But *prima facie* there are problems here: which one? The Nash Equilibrium with highest expected utility for one player need not have it for the other players. Moreover, it can be that each player plays their part of a Nash equilibrium, but if they play their part of different equilibria the set of mixed decisions might not form a Nash equilibrium. Call this the

coordination problem. Set it aside for now.

Now let me explain how to get standard game theory from deliberational decision theory. The basic idea is obvious: in a game one should treat the other players just as one treats states of nature when making a decision. Here is what will happen given just this basic idea. Each player starts with credences for all players actions. Each player computes his own expected utilities given these credences. Each player changes his own credences in his own actions by some rule that increases credences in actions that have higher utility than the status quo. And each player then adjusts his credences in the other player's actions accordingly. Each player repeats this procedure over and over. Once again, assuming continuity of the update rules, there will be equilibrium credences for each player. Suppose that each player ends up with a set of equilibrium credences such that for each player there are no equilibrium credences with a higher expected utility for that player. Will a combination of such equilibrium credences form a Nash equilibrium? Well, they may or may not. It all depends on the credences that the players start with, on the utilities associated with sets of actions, and on the exact update rules that they use.

However, let's modify our assumptions a bit and see where it gets us. Suppose that all players start with exactly the same credence distribution over the actions of all players, and suppose that this credence distribution is such that there is no correlation between the actions of the players, i.e. suppose that initially this credence function Cr is such that

$$\begin{aligned} Cr(\text{player A does act } A_i, \text{ player B does act } B_j, \dots, \text{ player N does act } N_k) = \\ = Cr(\text{player A does act } A_i) \times Cr(\text{player B does act } B_j) \times \dots \times Cr(\text{player N does act } N_k) \end{aligned}$$

Suppose, moreover, that each player is a deliberational decision theorist and each player knows each other players update rule, and each player updates his credences in any other player's acts

by using the update rule of that player.¹⁶ Then one can show that each deliberational equilibrium is a Nash equilibrium and vice versa (see Skyrms 1990 page 33). Whether the players will get to a Nash equilibrium in a finite time, and which Nash equilibrium they get to if they do, depends on the initial credences, and the update rules: the coordination problem is solved by the assumption of common initial credences and common knowledge of the update rules.

Now, these are strong assumptions indeed. For various possible weakenings see Skyrms 1990. But we have at least forged an interesting connection between deliberational decision theory and game theory.

11 Conclusions

Evidential decision theory violates Piaf's 'no regrets' maxim because there are information partitions conditional upon which an evidential decision theorist must violate Desire Reflection. Causal decision theory violates Piaf's maxim because there are cases, such as Psycho Johnny, in which one must regret one's decision as soon as one has made it. Deliberational causal decision theory avoids these problems. Moreover, in certain circumstances one can derive standard game theory from deliberational decision theory. Edith Piaf rules!

¹⁶ It does not matter whether the players are Deliberational Causal Decision Theorists or Deliberational Evidential Decision Theorists. The only relevant features of the situation are the players actions, and we have assumed that there are no correlations between the players acts according to their credences.

REFERENCES

- Eells, E. (1982): *Rational Decision and Causality*. New York: Cambridge University Press.
- Egan, A. (forthcoming): “Some counterexamples to causal decision theory”, forthcoming in *The Philosophical Review*.
- Gibbard, A. And Harper, W. (1978) : “Counterfactuals and two kinds of expected utility”, in C. Leach, E. McClennen, and C. Hooker (eds.): *Foundations and Applications of Decision Theory*. Dordrecht: D. Reidel, pp125-162.
- Jeffrey, R. (1983, 2nd Edition): *The Logic of Decision*. Chicago and London: Chicago University Press
- Joyce, J. (1999): *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Richter, Reed (1984): “Rationality revisited”, *Australasian Journal of Philosophy* 62, pp 392-403.
- Savage, L. J. (1972, 2nd Edition): *The Foundations of Statistics*. New York: Dover publications.
- Skyrms, B. (1990): *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge
- Weirich, P. (1985); “Decision instability”, *Australasian Journal of Philosophy* 63, pp 465-472.