NO REGRETS [1]

Frank Arntzenius, Rutgers University

1. Overview

I will argue that standard decision theories, namely causal decision theory and evidential decision theory, both are incoherent. I will then devise a new decision theory, from which standard game theory will follow as a corollary.

2. Causal Decision Theory is incoherent.

Johnny has a button in front of him such that if he presses it all psychos will die.[2] Other things equal he wants to kill all psychos, but other things are not equal: he does not want to die himself. He thinks it unlikely that he is a psycho. However, he also thinks that pushing the button would be good evidence that he is a psycho. What should he do?

Causal decision theory says the following. First Johnny should check his utility of pushing assuming he is a psycho. Let's suppose it is -100. Then he should check his utility of pushing assuming he is not a psycho. Let's suppose it is 10. Then he should check his credence (degree of belief, subjective probability) that he is a psycho. Suppose it is 0.01. Then Johnny's expected causal utility of pushing is 0.01x-100+0.99x10= 8.9. Next Johnny should calculate the expected causal utility of not-pushing. Let us suppose that if Johnny doesn't push the button he doesn't care whether he is a psycho. Indeed let us suppose that Johnny's utility for either of those cases equals 0. So, trivially, his expected causal utility for not pushing is 0. Causal decision

---

[1] All of the inspiration, and many of the ideas, in this paper derive from Brian Skyrms' work. See especially Skyrms (1990).

[2] This example is from Egan (forthcoming)

theory says: do that act which maximizes expected causal utility. So it says that Johnny should push the button.

Some people's intuitions are that in this situation one should not push, and that causal decision theory therefore is wrong. Others will happily say that Johnny should indeed push the button. I don't have strong intuitions here, so will steer clear of that debate. I have a different reason for objecting to causal decision theory in the light of this example. Suppose that Johnny, being a good causal decision theorist, has made the decision to push. Recall that Johnny regards pushing the button as good evidence that he is a psycho. Let me make this precise: Johnny's degree of belief that he is a psycho conditional upon him pushing the button is 0.9, and his degree of belief that he is a psycho conditional upon not pushing is 0. So, once he has decided to push, he should increase his degree of belief that he is a psycho to 0.9. But notice that once he has adjusted his degrees of belief in the light of his decision, his decision to push no longer looks good. For, given his new degrees of belief, his expected causal utility of pushing is 0.9x-100+0.1x10= -89. (And that of not pushing, of course, still is 0.)

Johnny can not get out of this bind by, after all, deciding to not push. For as soon as he believes that he will not push, the expected causal utility of pushing again will become higher than that of not pushing. Causal decision theory is thus incoherent in the following sense: there are situations such that, as soon as you have made up your mind to do something, that decision looks bad. I.e. there are situations in which there are no stable decision states if one adheres to causal decision theory. What to do?

3. Deliberational Causal Decision Theory

Here is what I suggest. Allow what game theorists have been allowing for over 50 years:

so-called mixed decisions, i.e. decisions to do certain acts with certain probabilities. To be precise, I suggest allowing that the end result of a rational deliberation will be that one has non-trivial degrees of belief in one's possible acts.

Let me return to my example in order to show how to apply this idea. Johnny starts his deliberation with certain degrees of belief that he is a psycho and certain degrees of belief that he will push. He then does an expected causal utility calculation. He finds the expected utility of pushing is higher than that of not pushing. Rather than immediately becoming certain that he will push, he should merely increases his degree of belief that he will push. In the light of this he should increase his degree of belief that he is a psycho. Armed with these new degrees of belief he recalculates expected utilities. Armed with the result he resets his degrees of belief in pushing the button. And so on, until he reaches degrees of belief which are at equilibrium, i.e. are such that his own rules tell him not to change them anymore. This occurs when Pr(Push)=10/110, Pr(Psycho)=9/110. I claim that if Johnny is rational, this is the degrees of belief that he should have at the end of his deliberation.

Three clarifications, and then we will be done with our job of making causal decision theory coherent.

Will there always be such a deliberational equilibrium? Yes, given certain mild assumptions, there will always be at least one equilibrium distribution. Here is why. One starts with certain degrees of belief in one's actions and the possible causal situations. One then figures out the expected causal utilities. One then applies a rule for updating one's degrees of belief in one's actions. One then applies a rule for updating one's degrees of belief in the causal situations. One then re-calculates causal utilities etc. A set of probabilities forms a deliberational equilibrium iff updating leaves them invariant. Now, let us assume that both updating rules are continuous, i.e. that an imagined continuous variation in the thing to be updated produces a continuous variation in the updated thing. Then the two rules applied consecutively form a continuous map from the space of possible probability distributions to the space of possible probability distributions. Next let us suppose that the set of possible actions is finite, and that the set of possible situations is finite. Then the action probability space and the situation probability space will be finite dimensional, simply connected, and compact. And then it follows from Brouwer's fixed point theorem that there must be at least one probability distribution that gets mapped into itself. So there must exist at least one deliberational equilibrium if the updating rules are continuous and if the set of possible actions and the set of possible situations is finite.

Next clarification. Should one model a rational person as a deliberator who changes his degrees of belief during the deliberation? No, one need not. Indeed it is a little bit awkward to do so. After all, if one is ideally rational, how could there be any stage at which one has the 'wrong' degrees of belief? 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 is merely a crutch to make us realize that there always exists a state of deliberational equilibrium, and that that is the state that a rational person should always be in.

Final clarification. What if there is more than one deliberational equilibrium? If one took the dynamical model seriously, one might think that the equilibrium state that a rational person ends up in could depend on the starting point. However, given that I am simply claiming that a rational person should always be in a state of deliberational equilibrium, it seems clear that a rational person should always be a deliberational equilibrium such that there is no deliberational equilibrium with higher expected causal utility.

That is all there is to 'Deliberational Causal Decision Theory'.

4. All-Purpose Deliberational Decision Theory

In the above I started off with causal decision theory, argued that it was incoherent, and then made it coherent by turning it into Deliberational Causal Decision Theory. But what if one does not like causal decision theory to begin with? Well, in section **7** I will argue that expected causal utilities are indeed the quantities that rational people ought to maximize. But just in case you are not convinced by the argument in section **7**, let me now give you an all-purpose scheme for making your favorite decision theory coherent.

1 Use you favorite rule for updating on evidence (strict conditionalistion, jeffrey conditionalisation, maximise entropy,....)

2 Use you favorite rule for calculating expected utilities (causal utilities, evidential utilities, ...)

3 Use your favorite rule for updating your degrees of belief in the light of your expected utility calculation

4 A rational person should always be in a deliberational equilibrium such that there is no deliberational equilibrium with higher expected utility

5. Deriving game theory

Let me start with a brief sketch of game theory. 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. The standard game theoretic account of rationality is that rational players must play their part of a Nash equilibrium. But prima facie there is a problem: which one? The Nash Equilibrium with highest expected utility for one player need not have it for the other players. Moreover, for non-zero sum games it can be that each player plays their part of a Nash equilibrium, but if they play their part of different equilibria the pair might not form a Nash equilibrium. Call this the coordination problem. This problem can be solved when we derive game theory from deliberational decision theory.

Let me explain how. The idea should be obvious: in a game one should treat the other players just as one treats states of nature when making a decision. Start with certain degrees of belief for the actions of each player. Do an expected utility calculation for your actions. Change your degrees of belief in your actions in the light of this calculation. Emulate the same procedure for each of the other players. Repeat the cycle until you hit an equilibrium.

There is a problem with this idea. Suppose that the other players do the same, but start with different degrees of belief, or use a different updating rule. Then the equilibrium degrees of belief that each player ends up with need have nothing to do with each other. These problems can be solved by assuming that it is common knowledge what degrees of belief all the players start off with, and that it is common knowledge which updating rule each player uses.[3] It now follows that the players will arrive at a unique joint deliberational equilibrium, the "solution" to the game. But must it be a Nash Equilibrium? No. For instance, if the players use Evidential Decision Theory in a prisoners dilemma situation, they will not arrive at the unique Nash equilibrium of this game.

---

[3] Possible weakenings of the assumption of common knowledge are discussed in Skyrms 1990.

Luckily, evidential decision theory is incoherent. Let me explain why.

6 Evidential decision theory is incoherent

Consider the 2-box case ('Newcombe's paradox'). Suppose you are to make your choice in 2 hours. You are an evidential decision theorist, so at the moment the expected utility of choosing 1 box is higher. But suppose you know that 1 hour later will you see the content of the boxes, and that given what you have seen your actions no longer will be probabilistic evidence for what is in the boxes. Then you know that 1 hour from know your preferences will flip, since then the expected utility of picking 2 boxes will be higher. Let's call this a violation of 'desire reflection'. A violation of 'desire reflection' seems objectionable. What to do?

7 Expected Causal Utility is the thing to maximize.

Here is a desire reflection principle which, prima facie, seems plausible: it should never be the case that one is going to receive evidence and that one knows in advance that no matter what evidence one in fact gets, the relative magnitude of the expected utilities of two possible actions is going to switch. Slightly more formally: there can not be an 'evidence partition' $\{E_i\}$ such that $EU(A_j)>EU(A_k)$, while $EU_i(A_j) \leq EU_i(A_k)$ for all i. (Here $EU_i(A_j)$ stand for the expected utility of A_j given information E_i).

While this principle prima facie seems plausible, it is in fact not plausible. Consider the following case: I have to choose between betting on the Red Sox and betting on the Yankees. The bets are such that if I bet on the Red Sox and win I gain $2, if I bet on the Red Sox and lose I lose $1, if I bet on the Yankees and win I gain $1, and if I bet on the Yankees and lose I lose

\$2. Now my degrees of belief might be such that the expected utility of betting on the Yankees is higher than that of betting on the Red Sox. But now consider the information partition, “I win my bet”, “I lose my bet”. Conditional on either of those bits of information betting on Red Sox has higher EU. But intuitively speaking that does not mean I am irrational. Intuitively speaking this information partition is illegitimate since the information in question is not (subjunctively) independent from my decision as to how to bet.

Here is a better “Desire Reflection Principle”: If $\{E_i\}$ is a ‘legitimate’ information partition then $\sum_i D(E_i)EU_i(A)=EU(A)$. Now, I have not specified what exactly a ‘legitimate’ information partition is. All I am going to assume is that a partition into causal situations is a ‘legitimate’ information partition. For this means that in the above formula we can substitute the possible causal situations K for the E_i. If we do that we find that we must have for a rational person: $EU(A)=\sum_K D(K)EU_k(A)$, which is just the standard formula that determines expected causal utilities. So the only coherent thing for a rational person to maximize is expected causal utility.

So the only coherent decision theory is Deliberational Causal Decision Theory. Let’s get back to deriving game theory.

8 Game theory redux.

Suppose that

1) The initial degrees of belief all players for all acts are common knowledge,

2) When computing expected utilities the players use unconditional probabilities

3) All players have a utility update rule which increases probability of exactly those acts which have higher expected utility than the status quo.

Brian Skyrms has shown (Skyrms 1990) that in such a case every deliberational equilibrium is a Nash equilibrium and vice versa. So players following the dictates of deliberational causal decision theory will end up at a Nash equilibrium when there is common knowledge. Which one they und up at will depend on what updating rule they used and what degrees of belief they started with. We have derived standard game theory from the only coherent decision theory.

REFERENCES

Egan, A. (forthcoming): “Some counterexamples to causal decision theory”, forthcoming in The Philosophical Review.

Skyrms, B. (1990): *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge