


YET ANOTHER REFUTATION OF CAUSAL DECISION THEORY?

James M. Joyce
jjoyce@umich.edu
Department of Philosophy
University of Michigan
Feb 7, 2020
University of Illinois Colloquium


<http://www-personal.umich.edu/~jjoyce/papers/UI.pdf>

THE FRUSTRATER (Spencer 202x)

You are a rational *miser* who cares only about your fortune. You are presented with four boxes marked A , B , S and S^+ . The S and S^+ boxes contain \$35 and \$40, respectively. Boxes A and B may also contain cash, but you will not know how much unless you chose them. Things are set up so that you must *first* choose whether to choose S and get \$35 or choose $\sim S$ and make a ***subsequent*** choice among the other three options.




The contents of A and B were determined like this: Yesterday a reliable forecaster, the Frustrater, made a guess about which of A , B and S^+ you will choose if you pass up box S . **If she predicted you would take S^+ (π_{S^+}) she put \$50 in each of A and B .**




When π_{S^+} , A and B are your best choices. S^+ is *not* best.

The contents of A and B were determined like this: Yesterday a reliable forecaster, the Frustrater, made a guess about which of A , B and S^+ you will choose if you do not take box S . If she predicted you would take S^+ (π_{S^+}) she put \$50 in each of A and B . **If she predicted you would take A (π_A) she put \$0 in A and \$100 in B .**




When π_A , B is your best choice. S^+ is *not* best.

The contents of A and B were determined like this: Yesterday a reliable forecaster, the Frustrater, made a guess about which of A , B and S^+ you will choose if you do not take box S . If she predicted you would take S^+ (π_{S^+}) she put \$50 in each of A and B . If she predicted you would take A (π_A) she put \$0 in A and \$100 in B . **If she predicted you would take B (π_B) she put \$100 in A and \$0 in B .**



When π_B , A is your best choice. S^+ is *not* best.

In tree form, your situation looks like this, where the triple (s^+, a, b) at the end of each branch reflect what you get by choosing the branch when the Frustrater predicts that you'll, respectively, take S^+ , A or B :



Predict S^+ (π_{S^+}) $\Rightarrow$ A and B both pay \$50

Predict A (π_A) $\Rightarrow$ A pays \$0, B pays \$100

Predict B (π_B) $\Rightarrow$ A pays \$100, B pays \$0

RELIABILITY

You believe, on the basis of good evidence, that the Frustrater is a highly reliable predictor of your actions.

- If you get to node **[2]** there is a 90% chance that she will have (already) correctly predicted the act you will choose there.
- Your credences (degrees of confidence) look like this

$$c(\pi_{S+} \mid S+) = 0.9 \text{ and } c(\pi_A \mid S+) = c(\pi_B \mid A) = 0.05$$

$$c(\pi_A \mid A) = 0.9 \text{ and } c(\pi_{S+} \mid A) = c(\pi_B \mid A) = 0.05$$

$$c(\pi_B \mid B) = 0.9 \text{ and } c(\pi_{S+} \mid B) = c(\pi_A \mid B) = 0.05$$

TWO POINTS TO KEEP IN MIND

- ➔ Saying that the Frustrater **is** reliable given what you **actually** do does **not** mean that she **would have been** reliable had you done something else.
 - Given that you do A , the Frustrater would still have predicted A with 90% probability had you done B or S^+ .
- ➔ Your acts are reliable **indicators** of the Frustrater's predictions, but they do not **causally influence** them in any way.
 - If you decide to do A , this is good evidence that A was predicted, but it does nothing to alter the contents of the boxes.

CAUSAL DECISION THEORY

CDT tells you to evaluate each act by its *causal expected utility*.

- At each time t , the CEU of an act X **that you can choose at t** , is

$$U_t(X) = c_t(\pi_{S^+}) \cdot \$ (X \wedge \pi_{S^+}) + c_t(\pi_A) \cdot \$ (X \wedge \pi_A) + c_t(\pi_B) \cdot \$ (X \wedge \pi_B)$$

$c_t(\pi_{S^+})$ is your time- t probability that S^+ was predicted.

$\$ (X \wedge S^+)$ is the amount of money X pays in that circumstance.

- This contrasts with EDT, which assess acts by their *news values*, which are obtained by replacing $c_t(\pi)$ by $c_t(\pi | X)$ above.
- The difference is between being a *promoter* of good things (CDT) and being an *indicator* of good things (EDT).

WHAT CDT TELLS YOU TO DO AT [2]

Key Observation: $U_2(A) + U_2(B) = \$100$ for any coherent credences, which means that one of them must be \$50 or more.

Upshot: At [2] CDT tells you to choose A or B , and reject S^+ !

- EDT recommends S^+ and assigns A and B news values of \$7.50 each.

- But, *which* of A or B should a CDTer choose?

THE PARADOX OF UNRATIFIABILITY

- If you think you are more likely to do A/B than B/A , then you will also think it more likely that A/B was predicted.
- If you think it more likely that A/B was predicted, then you should choose B/A (since it will have a higher CEU).
- You believe you will do what you should.

Paradox: Therefore, if you think you are more likely to choose A/B , then you must also think you are more likely to choose B/A .

The only way out of this quandary is to suppose that you are equally likely to do A or B at [2], and are indifferent between them.

THE NODE-[2] EQUILIBRIUM

- You are convinced you won't choose S^+ (because it's not among your optimal acts and you're sure you'll perform an optimal act).
- You assign A and B credences of $\frac{1}{2}$.
- You assign credences of $c_2(\pi_A) = c_2(\pi_B) = 0.475$ and $c_2(\pi_{S^+}) = 0.05$ to the Frustrater's predictions.
- You assign A and B expected utilities of \$50.
- So, you are indifferent between A and B and must just **pick** using a mechanism that you believe favors A and B equally.
- Rational deliberation (Skyrms 1990) leads to this equilibrium.

The Way I Think about the Problem

- Choosers have a **type**. You are type- A/B if you are disposed to break the tie by picking A/B . (Ben's **pseudo**-random coin.)
- Agents **do not** have reliable access to their type. In the equilibrium you believe you are equally likely to be type- A or type- B .
- The Predictor has better information than you have about your type. This explains her accuracy. She identifies type- A/B 90% of the time (and mistakes them for type- B or type- S^+ each 5% of the time).

A MISCONCEPTION TO AVOID

Common Thought: It is irrational to choose A or B at [2] because you know that you will almost certainly get \$0 if you do.

- You know no such thing!
- You know that *the^N act you will choose* will almost certainly pay \$0, where this description is used *non-rigidly*.
- But, since you do not know whether A or B is *the^N act you will choose* you do *not* know which act will almost certainly pay \$0.
- At the time you choose/pick you think that either A or B is equally likely to be *the^N act you choose* or *the^N act you do not choose*, and the latter is likely to pay \$100.


TWO WAYS TO THINK ABOUT THE DECISION AT [1]

Spencer: At [2] you'll have the power to secure \$40. Hence, $\sim S$ must be worth at least \$40.

- The value of choosing $\sim S$ is the just value of choosing whichever act you will end up performing at [2].
 - For a CDTer this is \$50.
 - For an EDTer this is \$40.
- Either way, you should choose $\sim S$ at [1].

THE GUARANTEE PRINCIPLE

Suppose you now confront a choice between *subsequent* decisions D_1 and D_2 . If *no* option in D_1 offers any outcome that exceeds \$ n , but *some* option in D_2 offers *only* outcomes that exceed \$ n , then you should prefer D_2 to D_1 .



Chocolate is better than any pie.

It need *not* be the best ice cream option.

The best ice cream option might not dominate the pie options.

This will happen if, e.g.,

CM >>> Chocolate > pies > Vanilla


GP tells you to go to choose $\sim S$ in *Frustrater*.

MY VIEW

At [1] you know that you will not choose S^+ at [2], which is **bad news** since there is a 0.9 probability that your choice at [2] will have been predicted.

- So, at [1] you know that if you reach [2] there is a 0.9 probability that you get nothing, and 0.05 probabilities for \$100 and \$50.
- At [1] your CEU for $\sim S$ is thus

$$0.9 \cdot \$0 + 0.05 \cdot \$100 + 0.05 \cdot \$50 = \$7.50.$$



THE KEY MOVE

Since $\sim S$ involves making a *future* choice *among acts you do not yet control at [1]*, the CEU of $\sim S$ is the *expected news value* of the future choice you will make at [2].

- And, the news that you will choose A or B is bad!

General Rule: When considering an act whose performance you *can* ensure right now, you should treat it as an instrument for changing the world and evaluate it by its causal expected utility; when considering an act you *cannot* perform now, its prospective performance should be treated as you would any event not under your control, i.e., as an item of news.

So, CDT Requires You to choose S to avoid facing a later decision that is likely to go badly for you!

No Guarantee!

- GP can fail in cases like *Frustrater* where information about what you might do is evidence about what your acts might cause.
- If you knew at [1] that you'd choose S^+ , then GP is an application of “backward induction,” and both it and CDT recommend $\sim S$.
- This reasoning breaks down if you know you will *not* choose S^+ .
- Why should it matter that the best possible outcome in D_1 is worse than the worst possible outcome of some option Y in D_2 if this Y is sure not to be chosen?

Reply: There will be some option in D_2 that is better than Y and this better than anything in D_1 .

The Guarantee Fallacy

GP licenses this faulty inference:


- (i) A and B are preferred to S^+ at [2]. ✓
- (ii) S^+ is preferred to S at [1]. ✓
- (iii) So, A and B are preferred to S at [1]. ✗

This inference fails because your evaluations of A and B differ at [1], where they are *not* under your control, from what they are at [2].

- You *disprefer* A and B to S^+ at [1], where you evaluate all three as news items (using EDT).
- But, you prefer A and B to S^+ at [2] where you evaluate all three as action items (using CDT).

ANOTHER OBJECTION TO S : THE ORDER OF CHOICE

It might seem that if you decline S then you have the option of *first* deciding whether or not to take S^+ and only then choosing between A and B , if you even need to.



Faux Frustrater

THE KEY DIFFERENCE BETWEEN *FRUSTRATER* AND *FAUX*

In *Faux* you decide whether to take S^+ *before* choosing between A and B , but in *Frustrater* all three options remain in play until you choose among them.

- You can decline S^+ in *Faux* without committing yourself in the A -vs.- B choice, while in *Frustrater* declining S^+ is *constituted* by taking A or by taking B .

This makes all the difference!

- As an adherent of CDT, you will take \$40 at [2a] to avoid having to choose between A and B at [2b], which you see as worth \$7.50.
- At [1] you know this and so see $\sim S$ as offering a sure \$40, which exceeds S 's \$35 payoff. So, you choose $\sim S$ -then- S^+ .

IS THIS A PROBLEM?

Opponents of CDT might argue that either:

(i) *Frustrater* and *Faux* should have the same solution because they are really the same decision differently described.

or

(ii) A rational agent always has the ability to cordon off S^+ -vs.- $\sim S^+$ from A -vs.- B , thereby leaving themselves with a *choice* about which of the two decisions to make.

- Since *Frustrater* is worth \$35 while *Faux* is worth \$40, they can never rationally choose S in *Frustrater*. They should always opt to face *Faux* instead.

ARE *FRUSTRATER* AND *FAUX* THE SAME DECISION?

No! because you do not control A and B at **[2a]** in *Faux*!

- At **[2a]** in *Faux* you should treat A and B as news items, and should choose S^+ to avoid letting your all-too-predictable future self face the A -vs.- B choice, which is worth only \$7.50.
- But, in *Frustrater* choosing S^+ means not choosing A or B at a point when you *can*, and when least one of their values $\geq$ \$50.

CAN'T WE ALWAYS CHOOSE TO DECIDE *FAUX*?

(ii) **Claim:** You can always cordon off S^+ -vs.- $\sim S^+$ from A -vs.- B , and choose to face *Faux* rather than *Frustrater* to face.

Idea: When facing *Frustrater* you can *resolve* to decide S^+ -vs.- $\sim S^+$ before you decide A -vs.- B , which turns *Frustrater* into *Faux*.

Two Questions:

- Can you “tie yourself to the mast” by making a *binding* resolution that prevents you from deciding A -vs.- B until *after* S^+ -vs.- $\sim S^+$?
- Or, do you retain the power to choose A or B when deciding on S^+ ?

IF YOU CAN “TIE YOURSELF TO THE MAST”

...then you can indeed transform *Frustrator* into *Faux*, and CDT says that you should do that!

- This does *not* contradict CDT's claim that you should choose *S* in *Frustrater* because you are not facing *Frustrater*.
 - You have an extra option at [1], not available in Frustrater, of tying yourself to the mast.
- I suspect many people conflate *Faux* and *Frustrater* because they assume that rational free agents can always tie themselves to the mast. (I don't think they can.)

IF YOU CANNOT “TIE YOU TO THE MAST”

...then you end up facing *Frustrater*.

- When deciding S^+ -vs.- $\sim S^+$ you must consider passing up the *live* options of choosing A or B as part of the *opportunity costs* of choosing S^+ .
- This is the difference between \$40 and the CEU of whichever other option *among those you currently control* has the highest CEU.
- Since you currently control A and B , and since at least one of A or B has a CEU of at least \$50, the opportunity cost of choosing S^+ is at least \$10.

WHY AIN'T YA RICH?

Spencer suggests that the relative poverty of A and B choosers, as compared to S^+ choosers, shows that it is irrational to choose A or B .

- This is a “**Why Ain’t Ya Rich**” (WAR) argument.
- **Standard Response:**
 - A and B choosers end up worse off simply because they start with worse options.
 - S^+ choosers end up richer because they start with better options.

DOES THE STANDARD RESPONSE WORK?

Spencer says no! (He accepts the Standard Response for Newcomb's problem, but rejects it here.)

- He writes that “there is no obvious difference in opportunity” between those who would and would not take S^+ , and that “I suspect that the relative poverty of [*A* and *B* choosers] is best explained by the hypothesis that [those acts] are irrational.”
- **Key idea:** *A*-choosers and *B*-choosers have the option of choosing a sure \$40. So, if they end up making less than \$40 they have only themselves and their decision rule to blame.

HOW THE DECK IS STACKED AGAINST *A/B* TYPES

While there is a sense in which S^+ choosers and non- S^+ choosers have the same opportunities, there is also an asymmetry between the decisions they make among them.

- It helps to think in terms of **types**. Recall: your type is a matter of how you choose/pick. It serves as the common cause of both your ultimate act and the predictor's prediction.
- If you are type-*A/B*, then the predictor, who knows more about your type than you do, has arranged things so that by choosing *A/B* you'll get \$0 with 90% probability, \$100 with 5% probability and \$50 with 5% probability.

- You lack of knowledge about your type (which persists until *after* you choose) allows the Predictor to “stack the deck” against you by arranging things so that your options look better in expectation than they will likely turn out to be in reality.
- Isn't an indication of the irrationality of your decision rule that it makes this kind of “punishment” possible?
- No! It is no failure of rationality to be exploited by someone who has more relevant information than you do.

HOW THE DECK IS STACKED IN FAVOR OF S^+ -TYPES

- As an S^+ -type you won't know your type either. If you did, you'd choose A or B since both have a 90% probability of yielding \$50.
- If you are type- S^+ , then the predictor, who knows more about your type than you do, has arranged things so that by choosing S^+ you get \$40 for sure, and by choosing A or B you would have gotten \$50.
 - So, as a type- S^+ you have a floor of \$40 whatever you do!

Moral: S^+ -choosers end up poorer than S^+ choosers because, in all probability, they start out with worse options, not because they choose badly from among the options they have.

SHOULDN'T CDTERS EARN AT LEAST MUCH S^+ CHOOSERS?

- Doesn't the fact that type- A/B s are free to choose S^+ mean that they are not actually disadvantaged relative to S^+ types?
- This misses the point. It is no part of CDTERS' purpose to end up as rich as S^+ -types.
 - As misers, they aim to make as much money as they can *given the options they have and given what they know at the time they choose*, and this is true no matter what agents with *other* options might do or earn.
- Of course, none of this changes the fact that CDTERS face a bad situation with the deck stacked against them. If they can avoid *The Frustrater*, they should!

THANKS!

References

Caspar Oesterheld and Vincent Conitzer, “Extracting Money from Causal Decision Theorists,”

<https://users.cs.duke.edu/~ocaspar/CDTMoneyPump.pdf>.

Jack Spencer “Causal Decision Theory and the Guarantee Principle,
http://www.jackspencer.org/uploads/1/4/0/3/14038590/cdt_and_the_guaranteed_principle.pdf