Ratifiability, Stability, and the Role of Act Probabilities in Decision Theory

Progic 2009

September 18, 2009

University of Groningen, The Netherlands

Jim Joyce
Department of Philosophy
The University of Michigan
jjoyce@umich.edu

Causal Decision Theory

Agents should choose actions that *causally promote* desirable outcomes, even when the performance of these actions indicates undesirable outcomes that they do not cause.

This involves choosing the act that maximizes the causal expected utility

$$\textbf{(CDT)} \qquad \mathcal{U}(A) = \Sigma_K P(K) \cdot u(A, K)$$

where K ranges over a special partition of "dependency hypotheses" that each provides a maximally complete specification of how things the agent cares about might depend on what she does.

- $\mathcal{U}$-maximizers treat information about their acts as irrelevant (for purposes of decision making) *except insofar as is it indicates what the acts are likely to cause*.

- Since acts typically indicate outcomes only to the extent that they causally promote or inhibit them, it often makes no difference whether one insists that the states in a well-formed decision problem be causally or merely evidentially (statistically) independent of acts.

- That said, in cases where acts indicate outcomes they do not cause – e.g., in "common cause" situations – the causal decision theorists says that the decision maker should ignore the purely evidential value of her acts.

Egan's "Counterexamples" to CDT

Andy Egan (2007) offers a series of "counterexamples" to CDT in which the performance of an act provides evidence about its own causal consequences, and this evidence undermines the rationale for choosing the act. Here is one example:

> *Murder Lesion*. Things would be better if you killed King Alfred. You have a gun aimed at his head and merely need to pull the trigger. But, one needs nerve to kill, and killing takes a steady hand: even if you have the nerve, you might miss. If you do shoot and miss, Alfred will have you tortured. If you shoot and kill Alfred, his brother Manfred will assume the throne and usher in a golden age. Your situation is complicated by the fact that you are a random member of a population in which a certain percentage x of people have a brain lesion that causes trembling hands in all who have it, and endows 70% of those who have it with the nerve to kill. People who lack the lesion shoot only 10% of the time, but always aim true.

	L = You have the lesion. You would miss if you shot, but you likely (0.70) have the nerve.	$\sim L$ = You lack the lesion. You'll kill Alfred if you shot, but you likely (0.90) lack the nerve.
Act S = Shoot	Punishment ($u(S, L) = -10$)	Golden Age ($u(S, \sim L) = 10$)
Act $\sim S$ = Refrain	Status Quo ($u(\sim S, L) = 0$)	Status Quo ($u(\sim S, \sim L) = 0$)

In the population of people who know all the above and believe $x = 30\%$, exactly 30% have the lesion. You are in this group. **Should you shoot?**

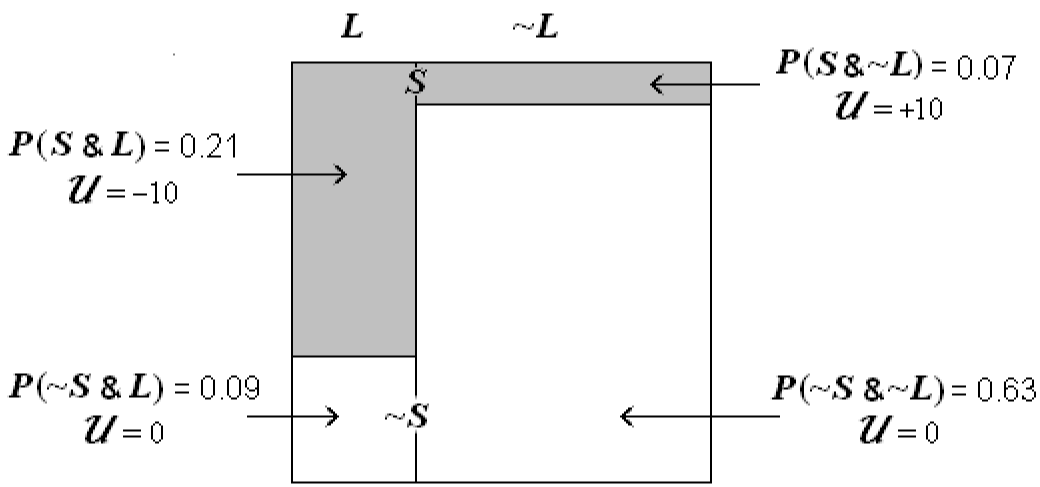

- Your initial evidential situation is:

$\boldsymbol{P}_0(L) = 0.30$ $\boldsymbol{P}_0(S \mid L) = 0.70$ $\boldsymbol{P}_0(S \mid {\sim}L) = 0.10$

$\boldsymbol{P}_0(S) = 0.28$ $\boldsymbol{P}_0(L \mid S) = 0.75$ $\boldsymbol{P}_0(L \mid {\sim}S) = 0.125$

- Your initial causal expected utilities are:

$\mathcal{U}_0(S) = 10 - 20 \cdot \boldsymbol{P}_0(L) = 4 > \mathcal{U}_0({\sim}S) = 0$

Egan's Claims: (a) CDT *requires* you to shoot.
(b) Refraining is the uniquely rational action.

NEITHER CLAIM IS TRUE!

The Problem: No Ratifiable Options

Murder Lesion and decisions like it (e.g., Death in Damascus) are perplexing because they lack *causally ratifiable* pure acts.

- Deciding to shoot is a good reason for thinking you have the nerve, which is a good reason for thinking you have the lesion, which is a good reason for thinking you will miss, which is a good reason to refrain. In short, deciding to shoot provides you with strong evidence for thinking that the causal consequences of shooting will be worse than those not shooting.
- Deciding not to shoot is a good reason for thinking you lack the nerve, which is a good reason for thinking you lack the lesion, which is a good reason for thinking you will aim true, which is a good reason to shoot. In short, deciding not to shoot provides you with strong evidence for thinking that the causal consequences of shooting will be better than those of not shooting.

Def. Act A is causally ratifiable when it maximizes causal expected utility on the supposition that it is (irrevocably) chosen, so that $\mathcal{U}(A \| A) \geq \mathcal{U}(B \| A)$ for all alternatives B to A where $\mathcal{U}(\bullet \| A)$ represents the utilities the agent expects to assign when she has irrevocably decided to do A.

In Murder Lesion it is easy to calculate these utilities (not always so).

$\mathcal{U}(\sim S \| S) = 0$ and $\mathcal{U}(S \| S) = 10 - 20 \cdot \boldsymbol{P}(L \mid S)$

So, shooting is only ratifiable if your probability for having the lesion given that you shoot is not above ½. This fails at the initial time since $\boldsymbol{P}_0(L \mid S) = 0.75$.

$\mathcal{U}(\sim S \| \sim S) = 0$ and $\mathcal{U}(S \| \sim S) = 10 - 20 \cdot \boldsymbol{P}(L \mid \sim S)$.

So, refraining is only ratifiable if your probability for having the lesion given that you refrain is at least ½. This fails at the initial time since $\boldsymbol{P}_0(L \mid \sim S) = 0.125$.

Is it Wrong to Choose Acts one Knows one will Regret having Chosen?

Unratifiable acts seem defective: by choosing them one puts oneself in an epistemic position from which it seems foolish to act on one's choice. In other words, one comes to *regret* choosing them as soon as one becomes convinced that one will perform them.

- One might want to rule this out by invoking one of

  No Regrets (strong). If you know that you will regret choosing *A* once you choose it, then you cannot rationally choose *A*.

  No Regrets (weak). If you know that you will regret choosing *A* once you choose it, and if there is an alternative to *A* you would not regret choosing, then you cannot rationally choose *A*.

- Arntzenius (2008) sees the "no regret" requirements as being closely related to

  Weak Desire Reflection. If your desires at time $t = 1$ arise from your desires at $t = 0$ solely as a result of conditioning on evidence you acquire between the two times, then your $t = 0$ desires should be your expectation of your $t = 1$ desires, so that $\mathcal{U}_0(A) = \sum_x \boldsymbol{P}_0(\mathcal{U}_1(A) = x) \cdot x$.

- Egan, who thinks you should refrain in ML, must deny these principles. What should we say?

Mixed Acts: *No Panacea*

It is tempting to think that mixed acts can help here. Consider the act S^p in which you randomly shoot with probability p and refrain with probability $1 - p$.

In game theory (under the usual assumptions) there is always a ratifiable mixed act. Why not here?

Answer: Here, the availability of a mixed act does nothing to alter the probabilities of states (without special assumptions), and it follows that S^p is not ratifiable for any $p \in (0, 1)$.

- If the randomization is independent of L, so $\boldsymbol{P}_0(L \mid S^p) = \boldsymbol{P}_0(L \mid S \,\&\, S^p) = \boldsymbol{P}_0(L \mid {\sim}S \,\&\, S^p) = \boldsymbol{P}_0(L)$, then $\mathcal{U}_0(S \parallel S^p) = \mathcal{U}_0(S)$ and $\mathcal{U}_0({\sim}S \parallel S^p) = \mathcal{U}_0({\sim}S)$ and this means that $\mathcal{U}_0(S^p \parallel S^p) = p \cdot \mathcal{U}_0(S) + (1 - p) \cdot \mathcal{U}_0({\sim}S)$ will fall strictly between $\mathcal{U}_0(S \parallel S^p)$ and $\mathcal{U}_0({\sim}S \parallel S^p)$.

It is also problematic that

- Penalties could be imposed for randomizing.
- In practice, one often has the ability to renege on one's initial decision and to act contrary to what the device says. If so, the conditional probabilities $\boldsymbol{P}_0(S \mid L)$ and $\boldsymbol{P}_0(S \mid {\sim}L)$ could incorporate the odds that a person chooses S or ${\sim}S$ as the result of choosing S^p as well.
- Focusing on mixed acts obscures what is really going on.

Instructive Aside: Aumann on Equilibrium in Belief

	Left	Right
Up	+2, –2	–1, +2
Down	–2, +1	+1, –1

Nash Equilibrium at $\langle \text{Up}^{0.5}, \text{Left}^{0.33} \rangle$

How should we interpret the equilibrium probabilities?

Aumann: The critical thing is that COLUMN *believes* that ROW will play Up with probability ½ and ROW *believes* that COLUMN will play Left with probability 1/3.

- No matter how they achieve this "equilibrium in belief" both the players will recognizes that each can know no more about the other's acts than they would if it that act were the result of a random process that produces each act with its equilibrium probability.
- Thus, the probabilities in ROW's "mixed act" are COLUMN's degrees of belief, while those in COLUMN's mixed act are ROW's degrees of belief.
- Since each player recognizes that the other must act on the basis of her beliefs, any act, pure or mixed, to which the other player assigns positive probability will have the same (maximum) expected utility, *both unconditionally and on the condition that it is performed*. There are no wrong choices in equilibrium!
- Equilibrium act probabilities reflect the relative *strengths of an agent's motives* for pursuing various options.

I will argue that something like Aumann's interpretation can be given for act probabilities that arise in Murder Lesion. One obvious problem: there is no other player!

Does CDT Really Advocate Shooting?

Egan: CDT "enjoins us to *do whatever has the best expected outcome, holding fixed our initial views about the likely causal structure of the world*". Egan thus sees CDT as committed to:

> *Initial Opinion Fixes Action*. If $\boldsymbol{P}_0$ characterizes your beliefs at the *start* of your deliberations, then you are rationally obliged to perform an act that maximizes
>
> $$\mathcal{U}_0(A) = \boldsymbol{P}_0(L)\, u(L \,\&\, A) + \boldsymbol{P}_0(\sim L)\, u(\sim L \,\&\, A)$$

More generally,

> *Current Opinion Fixes Action*. If $\boldsymbol{P}_t$ characterizes your beliefs at time t, then at t you are rationally obliged to perform an act that maximizes your t causal expected utility
>
> $$\mathcal{U}_t(A) = \Sigma_K\, \boldsymbol{P}_t(K)\, u(K \,\&\, A)$$

If these principles are right, then CDT does unequivocally (and incorrectly) tell you to shoot.

However, these principles are *wrong*!

Causal Decision Theory's Real Commitment

Current Opinion Fixes Evaluation. If $\boldsymbol{P}_t$ gives your beliefs at t, then you are rationally obliged to *evaluate* each act by its causal expected utility at t:

$$\mathcal{U}_t(A) = \boldsymbol{P}_t(L)\, u(L \,\&\, A) + \boldsymbol{P}_t(\sim L)\, u(\sim L \,\&\, A)$$

- This says nothing about what you should *do*; it pertains only to how you should evaluate acts at a time given your beliefs and desires at that time.
- It is consistent with this that such evaluations should not be acted upon until they meet some further conditions.

Question: *When should time-t evaluations of causal expected utility guide actions?*

When Should Utilities Guide Actions?

I say that causal decision theory has always implicitly assumed these **conditions for action**:

Current Opinion Fixes Evaluation. If $\boldsymbol{P}_t$ gives your beliefs at *t*, then you are rationally obliged to *evaluate* each act open to you using its causal expected utility computed using $\boldsymbol{P}_t$.

Full Information. You should only act on the basis of your time *t* utility evaluations when these evaluations are based on beliefs that incorporate all evidence that is both (a) available to you at *t* and (b) relevant to the question of what your actions are likely to cause.

- An item of information is relevant in the sense of (b) when incorporating it into one's subjective probability will alter the causal expected utility of some available action.
- This *can* include information about what you are likely to do, *provided that this information is relevant to question about what your acts are likely to cause*.
- Causal decision theorists often speak as if *all* information about act probabilities is irrelevant to decision making (as it is in Newcomb cases, where dominance rules). What's true is that *only* information *that does not discriminate among actions with respect to their causal consequences* is irrelevant.

Belief in Rationality. You should only act on the basis of your time *t* utility evaluations if these evaluations are consistent with having a high level of confidence in the proposition that you will perform an act that maximizes expected utility.

- So, an act that maximizes expected utility only when it is unlikely to be performed, may not be chosen.

Murder Lesion Revisited

While CDT ranks shooting above refraining at $t = 0$, it does *not* advise you to shoot because your $t = 0$ utilities do not incorporate the (available and highly relevant) evidence that you prefer to shoot at $t = 0$, and they do not allow you to believe that you will do the rational thing.

- Since you strictly prefer S over $\sim S$ at $t = 0$ you have a reason to believe that you are more likely to do S than $\sim S$ (since you believe you will act rationally), but this reason is not reflected in your $t = 0$ beliefs since $\boldsymbol{P}_0(S) < \boldsymbol{P}_0(\sim S)$. Your $t = 0$ rationale for S depends on S being *un*likely!

More generally, at any time t one has

- You can choose to shoot on the basis of $\mathcal{U}_t(S) > \mathcal{U}_t(\sim S)$ only if (i) recognizing this fact does not alter your beliefs about what S and $\sim S$ will cause (Full Info), and (ii) you are highly confident that you will do what maximizes $\mathcal{U}_t$ (Belief in Rationality).
- But, $\mathcal{U}_t(S) > \mathcal{U}_t(\sim S)$ iff $\boldsymbol{P}_t(L) < \boldsymbol{P}_t(\sim L)$.
- Claim (justified below): $\boldsymbol{P}_t(L) < \boldsymbol{P}_t(\sim L)$ iff $\boldsymbol{P}_t(S) < 0.4$. So, you violate (ii).
- Claim (justified below): Learning $\mathcal{U}_t(S) > \mathcal{U}_t(\sim S)$ when $\boldsymbol{P}_t(S) < 0.4$ should lead you to revise your estimates of L's probability upward and S's utility downward. So, you violate (i).

There is a mirror argument that prohibits you from not shooting when $\mathcal{U}_t(S) < \mathcal{U}_t(\sim S)$.

Equilibrium in Belief Revisited

We have just seen that the following cannot all be true together in ML:

i. You know that *S* or ~*S* *uniquely* maximizes $\mathcal{U}_t$.
ii. You have processed all available information about what *S* and ~*S* are likely to cause.
iii. You are confident you will perform an act that maximizes $\mathcal{U}_t$.

Consequence: You can use your time *t* utilities to guide you actions in ML only if your beliefs and desires are in an equilibrium such that $\mathcal{U}_t(S) = \mathcal{U}_t(\sim S) = 0$, which requires $P_t(L) = ½$.

I will argue that in this state, and only this state, all of the following hold:

- You have processed all available information about what your acts are likely to cause (because learning $\mathcal{U}_t(S) = \mathcal{U}_t(\sim S)$ does not require you to alter your probability for *L*).
- You can confidently believe that you will act to maximize expected utility (trivially, because every act has the same expected utility).
- You may rationally choose S, ~S or any mixture of the two.
- The probabilities you assign to acts reflect your views about

A Model of Deliberation (Skyrms 1990)

- An agent's mental state at t is represented by a probability $\boldsymbol{P}_t$ and a causal expected utility $\boldsymbol{\mathcal{U}}_t$.
- Probabilities are assigned to acts as well as states.
- Deliberation maps an initial ($\boldsymbol{P}_0$, $\boldsymbol{\mathcal{U}}_0$) through a (discrete) sequence of temporal stages ($\boldsymbol{P}_t$, $\boldsymbol{\mathcal{U}}_t$), $t \leq 1$, to a final equilibrium state ($\boldsymbol{P}_1$, $\boldsymbol{\mathcal{U}}_1$).
- At each stage t, each the desirability of the agent's overall situation is given by the utility of "status quo" $\boldsymbol{\mathcal{U}}_t(\top) = \sum_A \boldsymbol{P}_t(A)\cdot\boldsymbol{\mathcal{U}}_t(A)$.
- The agent alters act probabilities using an *update rule* that "seeks the good" by increasing/ decreasing probabilities of acts with utilities above/below the status quo.
  - ♦ Example (Bayesian dynamics). $\boldsymbol{P}_{t+1}(A) = \boldsymbol{P}_t(A)\cdot[\mathcal{U}_t(A)/\mathcal{U}_t(\top)]$, where $\mathcal{U}_t > 0$.
  - ♦ Act probabilities cannot change so abruptly that acts with utilities below the status quo are summarily assigned zero probability (and thereby removed from future consideration).
- At $t = 1$ the agent reaches a state of equilibrium. This counts as making up her mind.
- Deliberation can end in a "mixed state" in which $\boldsymbol{P}_1(A) > 0$ for more than one act A. Here the agent ends up torn among *equally desirable* acts since $\boldsymbol{\mathcal{U}}_1(A) = \boldsymbol{\mathcal{U}}_1(\top)$ when $\boldsymbol{P}_1(A) > 0$.

Note: Talk of deliberation here is a crutch, equilibrium is what matters.

Reflective Equilibrium

An agent should use her time-*t* utilities as a basis for action only when her beliefs (and desires) are in a state of *reflective equilibrium* in the sense described.

- Such an agent automatically satisfies the **conditions for action**.
  - ♦ Since deliberation only stops when information about time *t* causal utilities no longer affects time *t* + 1 causal utilities, it follows that all relevant information about what acts might cause has been taken into account.
    - o In particular, information about the agent's utilities and act probabilities is accounted for. This data is usually irrelevant, but it matters in cases (like Egan's) where the probability that the agent will behave in various ways influences her views about what her acts might cause.
  - ♦ Since deliberation only stops when all acts with positive probability have the same causal expected utility, the agent in equilibrium is certain she is going to act rationally (trivially).

General Point: At bottom, decision theory is about the relationships that hold between an agent's beliefs and desires and her actions when she has attained a state of deliberational equilibrium.

Slogan: Decision theory evaluates states of mind, not actions!

What do these equilibria look like?

Free vs. Constrained Updating

A Key Question: How should changes in act probabilities ramify through to other beliefs?

- ***Free updating*** allows changes in act probabilities unconstrained by prior information about the likely effects of states on acts.
  - Changes in act probabilities are treated as *acquired data*, something one learns about oneself.
  - Time $t + 1$ beliefs derive from the time t beliefs via a Jeffrey shift over the act partition, so that $\boldsymbol{P}_{t+1}(K) = \sum_n \boldsymbol{P}_{t+1}(A_n)\cdot P_t(K \mid A_n)$ for each state K.
  - Conditional probabilities of states given acts, the $\boldsymbol{P}_t(K \mid A)$, *stay fixed* as $\boldsymbol{P}_t(A)$ varies, so that $\boldsymbol{P}_{t+1}(K \mid A) = \boldsymbol{P}_t(K \mid A \;\&\; \mathcal{U}_t(A) = x) = \boldsymbol{P}_0(L \mid S)$ for all t.
  - Conditional probabilities of acts given states, the $\boldsymbol{P}_t(A \mid K)$, *vary* as $\boldsymbol{P}_t(A)$ varies.

➢ In Murder Lesion, the outcome of free updating produces the equilibrium

$\boldsymbol{P}_1(L) = 0.50$	$\boldsymbol{P}_1(S \mid L) = 0.90$	$\boldsymbol{P}_1(S \mid {\sim}L) = 0.30$	
$\boldsymbol{P}_1(S) = 0.60$	$\boldsymbol{P}_1(L \mid S) = 0.75$	$\boldsymbol{P}_1(L \mid {\sim}S) = 0.125$	(blue: fixed at $t = 0$ value)

You reach this point by using data about your likely behavior to *alter your views about the lesion's effects* while retaining your views about how likely you are to have the lesion given your action.

Constrained updating treats changes in act probabilities as *constrained manipulations*: they are changes the agent sees herself as *choosing*, not changes forced on her by evidence.

- The $\boldsymbol{P}_t(K \mid A)$ *vary* as $\boldsymbol{P}_t(A)$ varies. In contrast with "free" updating, the agent does not see her choices as providing direct evidence about the world's state.
- Probabilities of acts on states, the $\boldsymbol{P}_t(A \mid K)$, stay *fixed*. The agent does not see herself as having the power to act in ways that conflict with background facts about how actions depend on states.
- Probabilities of states are adjusted so that $\boldsymbol{P}_{t+1}(A) = \sum_K \boldsymbol{P}_{t+1}(K) \cdot \boldsymbol{P}_t(A \mid K)$.
- This constrains the allowable act probabilities, so that each $\boldsymbol{P}_{t+1}(A)$ must fall within the span of the $\boldsymbol{P}_0(A \mid K)$ for all t.

➢ In Murder Lesion, the outcome of constrained updating is an equilibrium in which

$\boldsymbol{P}_1(L) = 0.50$	$\boldsymbol{P}_1(S \mid L) = 0.70$	$\boldsymbol{P}_1(S \mid {\sim}L) = 0.10$	(blue: fixed at $t = 0$ value)
$\boldsymbol{P}_1(S) = 0.40$	$\boldsymbol{P}_1(L \mid S) = 0.875$	$\boldsymbol{P}_1(L \mid {\sim}S) = 0.25$	

You reach this equilibrium by using information about your likely behavior to draw inferences about you odds of having the lesion *without raising questions about the lesion's effects on your behavior*.

Free or Constrained?

- The "free" approach seems to presuppose an agent with a "libertarian" conception of free will.
  - The free deliberator will say "I know that the statistics say that 90% of people who lack the lesion don't shoot, but since I'm more likely to shoot than most people I must be more likely to shoot whether I have the lesion or not.
  - The policy of keeping $\boldsymbol{P}_t(L \mid S)$ and $\boldsymbol{P}_t(L \mid {\sim}S)$ fixed makes sense if the agent thinks of the changes in probability of S and $\sim S$ as being unrestricted by L.
- The "constrained" approach seems to presuppose a "compatibilist" conception of free will.
  - The constrained deliberator will say "The statistics say that only 10% of people who lack the lesion shoot. I'm nobody special. So, however strongly I am inclined to shoot, and however convinced I am that I will, if I don't have the lesion then my chances of shooting are still only 10%."
  - The policy of keeping $\boldsymbol{P}_t(S \mid L)$ and $\boldsymbol{P}_t(S \mid {\sim}L)$ fixed makes the most sense if the agent thinks of her actions being a partial effect of L.

While I think the constrained approach does a better job of capturing what is going on in Murder Lesion, and while I am inclined toward the "compatibilist" picture, I have no principled argument for ruling out the "free" approach. So, I am not sure what to say about this.

Two Ways of Assessing Acts in Equilibrium

First Way: The equilibrium expected utility $\boldsymbol{\mathcal{U}}_1(A)$ reflects the agent's evaluation of *A* as a cause of desirable outcomes given her equilibrium beliefs. By this standard, all acts with positive probability in equilibrium have equal merit.

- So, according to causal decision theory, the agent can rationally perform *any* pure act that has positive probability in equilibrium (or *any* mixture of such acts).

Second Way: The equilibrium probability of an act reflects the degree to which the prospect of performing that act contributes to the value of the status quo. It measures the extent to which the agent is "leaning toward" *A* in a state where she has processed all relevant information.

- There is no difference between these modes of assessment when only one act survives in equilibrium. Then, "$\boldsymbol{P}_1(A) = 1$" and "$\boldsymbol{\mathcal{U}}_1(A)$ is maximal" mean the same thing.

- When more than one act has positive probability it can happen that $\boldsymbol{P}_1(A) > \boldsymbol{P}_1(B)$ even though $\boldsymbol{\mathcal{U}}_1(A) = \boldsymbol{\mathcal{U}}_1(B)$. In such a situation the agent (i) sees no advantage in performing *A* over *B*, yet (ii) is perfectly happy to be in a position where she is more likely to do *A* than to do *B*

- I claim that deliberation terminates in a kind of "equilibrium in belief."

- This requires a special interpretation of act probabilities in equilibrium.

Act Probabilities in Equilibrium

- During deliberation act probabilities assigned are just predictions about future behavior. They are based on three assumptions: (i) learning that an act's utility exceeds the status quo increases one's motivation for performing it; (ii) current motives are an indicator of equilibrium motives; (iii) one is likely to perform an act to the extent that one is motivated to do it in equilibrium.

- When equilibrium is achieved act probabilities still reflect the degree to which one is motivated to perform acts (on a scale in which 1 = an unadulterated pro attitude, and 0 = an unadulterated con attitude), but they are no longer straight predictions.
    - In analogy with the game-theoretic case, they reflect the relative strengths of the agent's motives once she has processed all relevant information about what her acts might cause.
    - Thus, the agent will believe that *to the extent that her action is determined entirely by the strength of her motivations* she will behave as if she is randomizing among acts according to their equilibrium probabilities.
    - But, since (factoring in motivations) all acts available to the agent have the same expected utility, the agent is in a "picking" situation: her ultimate choice for one act over another does not reflect any judgment on her part that the chosen act is *better*.
    - So, even though the agent might be differentially motivated toward acts in equilibrium, she can rationally perform *any* pure or mixed act that has positive probability in equilibrium.
    - She ends up believing "people with my motivations choose A with probability $\mathbf{P}_1(A)$," but not necessarily "I will choose A with probability $\mathbf{P}_1(A)$."

Conclusions About Murder Lesion (assuming "constrained" updating)

- Causal decision theorists should say that, contra Egan, it is *not* true that refraining is rational while shooting is irrational.
- Still, there is a crucial difference between the two acts since you will, if rational, end up leaning more strongly toward refraining than toward shooting once you have digested all available information about what your acts might cause.
- So, CDT does not recommend *either* shooting or refraining in ML. It recommends that you reason yourself into a position where you are leaning more strongly toward refraining.
- Once you are in this position you can perform *either* action without risking irrationality.
- We should resist the temptation to think that decision theory can or should deliver more than this. You reasons are sufficient to incline you more strongly toward refraining than toward shooting, but not strong enough to make it reasonable for you to refrain outright.
- If all this is right, then causal decision theory has nothing to fear from Egan's examples.

But Won't You Regret Your Act?

YES! But, there is nothing problematic about this when ***no*** act is causally ratifiable.

There are two cases to consider:

1. The presence or absence of the lesion provides information about the causal consequences of the "picking" strategies employed in equilibrium.
    - $\boldsymbol{P}_0(S\ /\ L) = 0.7$ means that having the lesion leads agents in equilibrium to employ picking strategies that, in the aggregate, recommend shooting 70% of the time.
    - $\boldsymbol{P}_0(S\ /\ {\sim}L) = 0.1$ means that lacking the lesion leads agents in equilibrium to employ picking strategies that, in the aggregate, recommend shooting 10% of the time.

- You then will regret *everything* you do, no matter how you come to do it. Here the *No Regrets* principle is implausible.
- So, I reject *No Regrets* (*strong*)"

2. The presence or absence of the lesion provides ***no*** information about the causal consequences of the "picking" strategies employed in equilibrium.

    - $\boldsymbol{P}_0(S / L) = 0.7$ means there is a non-causal correlation (perhaps due to a common cause) between having the lesion and employing picking strategies that recommend shooting, so that 70% of those who have the lesion shoot.

    - $\boldsymbol{P}_0(S / {\sim}L) = 0.1$ means there is a non-causal correlation between lacking the lesion and employing picking strategies that recommend shooting, so that 10% of those who have the lesion shoot.

- Here, the change in *L*'s probability that results from picking *S* or ~*S* is not based on an effect-to-cause inference. It is "backtracking" reasoning.
- This means that information gained from conditioning on one's picking mechanism is, from the perspective of causal decision theory, irrelevant to what one should do.
- So, there is no problem with using an unratifiable picking strategy in this context.
- *No Regrets* and *Weak Desire Reflection* both fail because one will be no better informed about what one's acts are likely to cause when knows how one will pick.

General Conclusions

- Decision theory (causal or not) must be thought of as a normative theory that concerns the decision maker's state of mind at the point when she makes her choice.
- Acts can be rationally performed just when they maximize expected utility relative to the beliefs and desires that the decision maker will have when she has taken into account all available information about what outcomes her acts are likely to cause.
- Sometimes this equilibrium state will pick out a specific action as the unique right choice, but in cases like Egan's it will not.
- We need to distinguish two ways of assessing values of actions in equilibrium – in terms of their expected utilities, which measure choiceworthiness given beliefs, and in terms of their act probabilities, which measure the strength of the agent's motivations.
- A deliberational equilibrium is a kind of "equilibrium in belief" in which one can rationally pick any action that has a positive probability.
- In cases like Murder Lesion, the strong version of the *No Regrets* principle is false.

References

Arntzenius, Frank [2008]. No Regrets, Or: Edith Piaf Revamps Decision Theory. Erkenntnis 68 (2).

Eells, Ellery [1982] *Rational Decision and Causality*. Cambridge, MA: Cambridge University Press.

Egan, Andy [2007] "Some Counterexamples to Causal Decision Theory", *Philosophical Review*

Gibbard, Allan and William Harper [1978] "Counterfactuals and Two Kinds of Expected Utility," in *Foundations and Applications of Decision Theory*, edited by C. Hooker, J. Leach, and E. McClennen, pp. 125-62. Dordrecht: Reidel.

Levi, I. [1989] 'Rationality, Prediction, and Autonomous Choice', *Canadian Journal of Philosophy* **19** (suppl.), 339–363, re-printed in Levi (1997), pp. 19–39.

Jeffrey, Richard [1983] *The Logic of Decision*, 2nd edition, Chicago: The University of Chicago Press.

Jeffrey, Richard [1993] "Causality and the Logic of Decision," *Philosophical Topics* **21**: 139-151.

----- [2004] *Subjective Probability The Real Thing*. Cambridge, UK: Cambridge University Press.

Jeffrey, Richard [1993] "Causality and the Logic of Decision," *Philosophical Topics* **21**: 139-151.

Joyce, James M. [1999] *The Foundations of Causal Decision Theory*. Cambridge, UK: Cambridge University Press.

Joyce, James M. [2007] "Are Newcomb Problems really Decisions," *Synthese* **156** (3)

Rabinowicz, Woldek [2002] "Does Practical Deliberation Crowd Out Self-prediction?," *Erkenntnis* 57 (1)

Skyrms, Brian [1990] *The Dynamics of Rational Deliberation*. Cambridge, UK: Cambridge University Press.