Dynamic Consistency in the Logic of Decision

Gerard J. Rothfus

October 2019

Abstract

Arif Ahmed has recently argued that causal decision theory is dynamically inconsistent and that we should therefore prefer evidential decision theory. However, the principal formulation of the evidential theory, Richard Jeffrey's *Logic of Decision*, has a mixed record of its own when it comes to evaluating plans consistently across time. This note probes that neglected record, establishing the dynamic consistency of evidential decision theory within a restricted class of problems but then illustrating how evidentialists can fall into sequential incoherence outside of this class. Uncovering the evidentialist's own dynamic inconsistencies reveals, *contra* Ahmed, that sequential choice considerations do not significantly favor the evidentialist's theory over the causalist's.

1 Introduction

The familiar debate between Evidential Decision Theory (EDT) and its chief rival, Causal Decision Theory (CDT), has traditionally been carried out in the context of *static* choice. The main examples employed to motivate and arbitrate the dispute, from *Newcomb* to *Psycho-Button*, all involve an agent tasked with making a single decision at a particular point in time.[1] One of the many contributions of Arif Ahmed's recent book has been to freshen the debate by suggesting the relevance of *sequential* choice arguments.[2] In doing so, Ahmed uncovers an awkward feature of CDT: the agents it describes can fall prey to *dynamic inconsistency* when tasked with making a series of decisions across time. Ahmed takes this discovery as an argument for EDT.[3]

[1] There are notable exceptions to this. Skyrms 1990 and Maher 1990 offer sequential choice arguments that EDT and CDT agents, respectively, will sometimes disprefer cost-free information. Skyrms 1982 and Meacham 2010 also discuss EDT and CDT in the dynamically relevant context of the *self-recommendation* of decision theories. More recently, Wells 2018 has deployed a clever sequential choice example to undergird a 'why ain'cha rich?' argument against EDT.

[2] Ahmed 2014; Chs. 7 and 8.

[3] One might defend CDT here in various ways; for example, one could invoke the *sophisticated choice* considerations highlighted by Joyce 2016 to challenge Ahmed's approach to sequential decision making. Whatever their merits, however, such defenses of CDT are not my concern here. Ahmed's assumptions are common enough to justify exploring their implications for EDT, so I will simply take them for granted.

Of course, the prospects for such an argument hinge upon how EDT itself fares in the context of sequential choice. The dynamic inconsistency of CDT will provide little reason to embrace EDT if evidentialists are themselves dynamically inconsistent. Hence, advancing the EDT/CDT debate any further along the lines Ahmed has suggested requires vetting the dynamic consistency properties of EDT. Embarking upon such an investigation, it is natural to begin by considering how the principal formulation of EDT, Richard Jeffrey's *Logic of Decision*,[4] fares in the sort of sequential choice problems that lead to trouble for CDT and then to proceed to consider more general contexts of sequential choice. As it turns out, taking this path affords momentary comfort for the evidentialist: the dynamic consistency of EDT is guaranteed within the simple class of problems that Ahmed employs to make trouble for CDT. However, it ends in disappointment: when we countenance more general types of dynamic choice problems, the Logic of Decision fails to offer a sequentially coherent account of rational choice. The present note travels this course, with §2 offering the initial comfort, §3 supplying the disappointment, and §4 taking stock.

2 Dynamic Choice without Nature

Ahmed offers two decision problems that illustrate the dynamic inconsistency of CDT, *Newcomb-Insurance* and *Psycho-Insurance*, which are sequentialized versions of the Newcomb and Psycho-Button problems, respectively.[5] Both are relatively simple problems in which an agent is called upon to make a series of decisions in sequence, without any other learning events being interspersed among her choices. Following the standard practice of modelling dynamic choice problems using Bayesian decision trees,[6] these cases belong to the class of decision problems that can be modelled by decision trees in which all non-terminal nodes represent points of decision, and hence in which nature resolves none of the agent's uncertainty prior to the problem's conclusion. In such problems, an agent only learns whatever she opts to teach herself.

It should be welcome news for the proponent of EDT that the dynamic consistency of her theory is guaranteed within this class of problems. To prove this, a bit of assembly is required. First, let an agent be *dynamically consistent* in a decision problem modelled by a Bayesian decision tree T just in case she evaluates any plan p as an acceptable one to implement at a node n in T if and only if she also evaluates the continuation of p as acceptable at all nodes succeeding n in T at which p is still logically implementable. So stated, dynamic

[4] Jeffrey 1965/1983

[5] Nozick 1969; Egan 2007.

[6] See Hammond 1988 for a formal exposition of Bayesian decision trees. Throughout, I alter the standard definition of such trees in one significant way: by allowing learning to take place pursuant to choice nodes as well as natural nodes.

consistency is a property connecting an agent's evaluation of possible plans at earlier stages of a sequential choice problem with her evaluation of those plans at later stages of the problem. A dynamically consistent agent is one who can form a coherent contingency plan and stick to it, whilst always acting according to her best judgments at each time of action.[7]

While dynamic consistency concerns the evaluation of plans, the Logic of Decision is a theory concerned with ranking *propositions* according to their desirability, where propositions may be construed as sets of possible worlds.[8] An agent's desirabilities are measured by a value function V defined in the first instance over possible worlds and then derivatively over propositions according to:

$$V(X) = \sum_{w \in X} P(w|X)V(w),$$

where P is a probability measure representing the agent's (conditional) credences. The desirability of a proposition is then always an appropriately weighted sum of the desirabilities of the various ways in which it might be true. To make Jeffrey's theory applicable to sequential choice, a bridge principle is needed connecting evaluation of plans with evaluation of propositions. In our current context, a plan is just a sequence of acts that leads to a particular terminal node in a Bayesian decision tree. Since each node in such a tree is associated with a set of worlds, and hence with a proposition, we can associate a prospective plan with the proposition that captures the information state that its implementation terminates in. For the remainder, I will simply identify plans with such propositions.

Finally, since the Jeffrey theory is static, we need to make some minimal dynamic assumptions regarding how an agent updates her credences and values over time. The two assumptions I will make are both satisfied in the case of the causalists whose dynamic inconsistency won Ahmed's scorn. The first supplies a dynamics for credence, while the latter constrains the revision of values; in each statement, an arbitrary Bayesian decision tree is assumed fixed:

> Belief Conditionalization: For all nodes n_a and n_b such that n_b succeeds n_a, $P_b(\cdot) = P_a(\cdot|S(n_b))$, where P_z is the agent's probabilistic credence function at node n_z and $S(n)$ is the proposition capturing the agent's total evidence at n.
>
> Preference Stability: For all nodes n_a and n_b, if $w \in S(n_a) \cap S(n_b)$ then $V_a(w) = V_b(w)$, where V_z represents an agent's values at node n_z.

Belief Conditionalization is familiar, while Preference Stability can be thought of as requiring that an agent's basic values remain stable over the course of a

[7]For further discussion of dynamic consistency, see McClennen 1990 and Cubitt 1996.
[8]The presentation offered here follows Lewis 1981.

sequential choice problem.

A definition of dynamic consistency, an understanding of how to employ EDT to evaluate competing plans, and some basic diachronic stability assumptions now in place, we have all we need to show that evidentialists always exhibit dynamic consistency in the context of the sort of sequential choice problems considered by Ahmed. Suppose an evidentialist faces just such a problem, which we can model via a Bayesian decision tree T whose only non-terminal nodes are choice nodes. Letting n_a and n_b be any (non-terminal) nodes of T such that n_a precedes n_b along some branch of T, to prove this agent's dynamic consistency it suffices to verify that she judges a plan p acceptable at n_a (i.e. V_a-maximal) if and only if she judges its continuation acceptable at n_b (i.e. V_b-maximal), assuming that p makes arrival at n_b possible.[9] Below, we employ '$p(n)$' to denote the continuation of plan p at node n.

Suppose that p and p' are two plans available at n_a that make arrival at n_b possible. By the definition of V, the following are then equivalent:

(1) $V_a(p) \geq V_a(p')$

(2) $\sum_{w \in p} V_a(w) P_a(w|p) \geq \sum_{w \in p'} V_a(w) P_a(w|p')$.

Note that since p and p' make arrival at n_b possible and since every node in T is a choice node, p and p' each entail $S(n_b)$. Hence, p is equivalent to $p \wedge S(n_b)$ and p' to $p' \wedge S(n_b)$, so we have that (2) is equivalent to:

(3) $\sum_{w \in p} V_a(w) P_a(w|p \wedge S(n_b)) \geq \sum_{w \in p'} V_a(w) P_a(w|p' \wedge S(n_b))$

By Conditionalization, (3) is equivalent to:

(4) $\sum_{w \in p} V_a(w) P_b(w|p) \geq \sum_{w \in p'} V_a(w) P_b(w|p')$.

Preference Stability allows us to substitute V_b for V_a in (4), yielding:

(5) $\sum_{w \in p} V_b(w) P_b(w|p) \geq \sum_{w \in p'} V_b(w) P_b(w|p')$

But, of course, by the definition of V, this is just:

(6) $V_b(p) \geq V_b(p')$

Finally, recognizing the equivalence of plans and their continuations in the present context secures:

(7) $V_b(p(n_b)) \geq V_b(p'(n_b))$.

So we have established that (1) if and only if (7), i.e. $V_a(p) \geq V_a(p')$ if and only if $V_b(p(n_b)) \geq V_b(p'(n_b))$. Assuming that only finitely many plans are available to our agent and noting that every plan available at n_b is the continuation of some plan available at n_a, the dynamic consistency of the evidentialist is thus guaranteed.

[9] If T contains only one non-terminal node, dynamic consistency is trivially satisfied, so suppose this is not the case.

3 Dynamic Choice Generalized

So far, so good for the evidentialist. She seems to have excelled where the causalist faltered. If sequential choice were simply a matter of making decisions across time, we could end our interrogation of Jeffrey's theory here. But not all dynamic choice problems are so simple. Generally, an agent will learn more than her own choices over the course of implementing a cross-temporal plan of action. She is also likely to learn various other facts about the world that may impact the desirabilities of her options. Allowing for such learning, is the dynamic consistency of EDT still assured?

Before we can answer this question, note that once we countenance choice problems of the sort modelled by Bayesian decision trees involving natural nodes, the implementation of a plan is no longer guaranteed to terminate in an antecedently known information state. This raises a problem: how ought we to go about associating plans with propositions in this more general context? Associating plans with propositions in some way is, of course, needed if we hope to employ the Logic of Decision to evaluate them.

A natural proposal is to identify plans with suitable conjunctions of act propositions and partitioning conditionals.[10] (E.g. "First, I will apply for the loan; if approved, I will buy that new car; if denied, I will renew my bus pass.") To make this proposal precise, we can recursively define the plans available to an agent at a particular node n of a Bayesian decision tree T (a set I will dub '$\Omega(T, n)$' following standard notation). As our trivial base case, suppose n is a terminal node in T. Then the "plan" available at n will simply be the proposition capturing the planning agent's information state at n, i.e. $\Omega(T, n) = \{S(n)\}$. Now supposing that n is a choice node, we can define the set of plans available at n as the set of the various conjunctions of acts available at n with plans available at n's successors, i.e. $\Omega(T, n) = \{S(n') \wedge \pi | n' \in N_+(n), \pi \in \Omega(T, n')\}$, where $N_+(\cdot)$ is a function mapping nodes to their immediate successors. Finally, if n is a natural node, we can associate the set of plans available at n with a set of partitioning conditionals whose antecedents are given by the possible propositions an agent might learn at n and whose consequents are possible plans she might implement given receipt of that information, i.e $\Omega(T, n) = \{\wedge_i [S(n_i) \rightarrow \pi_i] | n_i \in N_+(n), \pi_i \in \Omega(T, n_i)\}$.

The trick here is saying just how the conditional operator, '$\rightarrow$', which I will refer to as a *planning conditional*, ought to be understood. An initial suggestion would be to take the planning conditional as a material conditional. On this reading, a plan can be identified with the *disjunction* of the propositions that its full implementation could possibly yield as total evidence. There are admittedly some drawbacks to this suggestion. Plans so construed will often fail to form a partition and leave undetermined what choices would be made by

[10] A *partitioning conditional* is a conjunction of conditionals whose antecedents form a logical partition. See Bradley 2017, p. 122-4.

an agent at the counterfactual decision points she never reaches in a sequential choice problem, an intuitively incorrect result when considering planning conditionals. Nonetheless, the material construal of the planning conditional avoids the complexities of non-truth functional semantics and is, in my estimation, an adequate model for a wide range of cases in which the desirability of plans interpreted truth functionally may be expected to coincide roughly with their desirability under more proper interpretations. So, for the moment, suppose we may treat '$\rightarrow$' as '$\supset$'.

Taking the planning conditional in this way, a simple example suffices to show that evidentialists are liable to exhibit dynamic inconsistency.[11] The problem is a sequentialized variant of the *Transparent Newcomb* problem that we may dub *Sequential Transparent Newcomb* (STN).[12] We begin by supposing that you are offered the opportunity to face the standard Transparent Newcomb problem. If you reject the offer, you receive nothing. If you accept, a predictor will place a million dollars in your bank account if and only if she predicts that you will reject her subsequent offer of one thousand dollars. The balance of your bank account is transparent to you at all times.

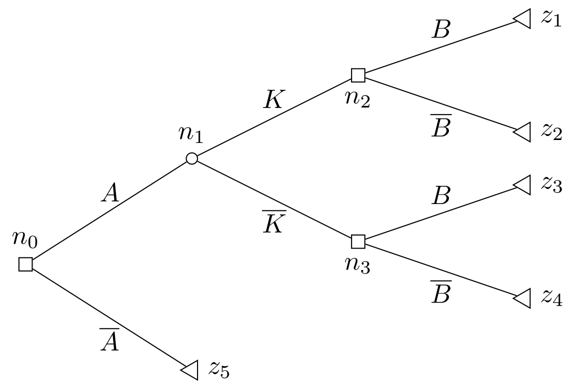

Figure 1: $T1$, the Sequential Transparent Newcomb Problem. $S(n_0) = W$, $S(z_1) = AKB$, $S(z_2) = AK\overline{B}$, etc.

Let A be the proposition that you take the predictor up on her offer to play, B the proposition that you accept her thousand dollar offer, and K the proposition that the predictor predicts $\overline{B}$. We assume that you know the structure of the problem, only care about your wealth level, and value money linearly. The

[11] The following thought experiment derives from a story suggested to me by [name redacted] in personal correspondence. I have simplified that original story in light of helpful comments from an anonymous reviewer.

[12] The static version of Transparent Newcomb was first discussed by Gibbard and Harper 1978. Skyrms 1982 and Meacham 2010 employ the example to argue that EDT can fail to be *self-recommending*.

problem can be depicted in the tree $T1$, shown in Figure 1.

For the purposes of the problem at hand, the propositions associated with the terminal nodes of $T1$ may be treated as the atoms of an algebra over which your credences and desiribilities are initially spread. So defining P and V over these propositions suffices to fix the values of these functions for all other propositions of interest to us as well. Suppose that your initial values and credences are defined by:

	P_0	V_0
$S(z_1)$	.014	1,001,000
$S(z_2)$	.7	1,000,000
$S(z_3)$	.14	1,000
$S(z_4)$	.07	0
$S(z_5)$	.076	0

Next, note that at the outset of STN there are five plans available to you:

(1) $\Pi_1 := A(K \to \overline{B})(\overline{K} \to \overline{B})$

(2) $\Pi_2 := A(K \to B)(\overline{K} \to B)$

(3) $\Pi_3 := A(K \to \overline{B})(\overline{K} \to B)$

(4) $\Pi_4 := A(K \to B)(\overline{K} \to \overline{B})$

(5) $\Pi_5 := \overline{A}$

We can show that Π_1 is V_0-maximal and so will be ex ante favored. Noting that Π_1 is logically equivalent to $\overline{B}$ (i.e. $S(z_2) \vee S(z_4)$), we compute its value:

$$\begin{aligned} V_0(\overline{B}) &= \sum_{w \in \overline{B}} P_0(w|\overline{B})V_0(w) \\ &= \sum_{w \in \overline{B}K} P_0(w|\overline{B})V_0(w) + \sum_{w \in \overline{BK}} P_0(w|\overline{B})V_0(w) \\ &= 1,000,000 \times \sum_{w \in \overline{B}K} P_0(w|\overline{B}) + 0 \\ &= 1,000,000 \times P_0(K|\overline{B}) \\ &= 909,090.\overline{90} \end{aligned}$$

In contrast, similar computation reveals the initial values of the other available plans to be substantially lower:

(1) $V_0(\Pi_2) = 91,909,\overline{09}$

(2) $V_0(\Pi_3) = 833,499.\overline{6}$

(3) $V_0(\Pi_4) = 166,833.\overline{3}$

(4) $V_0(\Pi_5) = 0$

Trusting that the skeptical reader may easily verify these values for herself, we may conclude that playing the predictor's game and then rejecting her thousand dollar offer is the V_0-maximal plan in STN.

It is evident, however, that EDT will recommend a change of heart to any agent that heeds its prescriptions and decides to play the game. After the truth value of K has been revealed, opting to grab the extra thousand carries no bad news, and hence accepting the money will be strictly preferred. But not accepting the money is the continuation of the plan that we observed above was initially most favored. What is V_0-maximal is neither V_2- nor V_3-maximal. The evidentialist finds herself trapped in dynamic inconsistency.

All of this assumes, of course, the material reading of the planning conditional. As noted above, this assumption may be undesirable in some contexts, which invites the question: Might matters look different if we were to equip planning conditionals with a more adequate non-truth-functional semantics? Not plausibly. A natural way to go about strengthening the planning conditional would, following recent work by Richard Bradley, replace possible worlds with n-tuples of worlds as the basic objects of belief and desire, the first entry of such n-tuples fixing all matters of fact and subsequent entries fixing matters of (potential) counterfact, conditional upon $n-1$ possible suppositions. For example, in STN, the suppositions of interest would be the states of the world, K and $\overline{K}$, and so the set of basic possibilities would be: $\{\langle S(z_i), S(z_j), S(z_k)\rangle | 1 \leq i, j, k \leq 5\}$, where $S(z_i)$ fixes the actual world and $S(z_j)$ and $S(z_k)$ fix what would be done by the agent on the (planning) supposition that K and $\overline{K}$, respectively. (Plausibly, this set should be restricted so that $j \in \{2, 3\}, k \in \{4, 5\}$.) In this framework, the content of non-conditional propositions can be given in terms of the set of n-tuples whose first entry renders them true, while a conditional proposition $A \rightarrow B$ can be identified with the set of n-tuples whose entries in the position corresponding to the A supposition make B true.

Taking this route, the set of plans available to an agent at a particular node in a decision tree will partition the space of world n-tuples judged possible by the agent at that node. The appendix below lists all atomic possibilities within this framework at the start of the STN problem, along with hypothetical credence and desirability mass functions spread over them. It is then verified there that these functions again render plan Π_1 (construed in line with the Bradley semantics) V_0-maximal and thus suffice to prove the dynamic inconsistency of EDT. Treating planning conditionals non-truth functionally introduces interesting subtleties into dynamic choice theory, but it promises no easy escape for the evidentialist from her diachronic woes.[13]

[13]One might try to rescue the evidentialist by restricting the set of rationally admissible credence functions defined over world n-tuples in some way that ends up ruling out the ones

4 Conclusion

An evidentialist takes into account any information that her choices would provide her when she deliberates regarding possible courses of action. This guarantees her dynamic consistency in sequential choice problems that involve no learning apart from her own acts. Since an evidentialist cannot similarly factor information about uncertain states of the world into her *ex ante* evaluation of plans, those sequential choice problems that involve the instruction of nature present possible occasions of dynamic inconsistency on her part.

What can we glean from this concerning the relative merits of EDT vs CDT in the sequential choice context? Well, it may seem fair to grant *some* advantage to the evidentialist here. After all, the Logic of Decision does avoid dynamic inconsistencies in simple decision problems of the sort that initially caused trouble for CDT. Further, while the causalist's dynamic incoherence in such problems leaves her vulnerable to exploitation in a way that she may well, by her own lights, find troubling, the evidentialist avoids a similar worry here.[14] The reason for this is that the causalist confronting Newcomb-Insurance or a like problem is *guaranteed* to implement a plan that she can recognize as being dominated by another available plan, relative to a partition of propositions capturing the various ways in which available plans may causally promote outcomes of concern to her. Since a principle ruling out the selection of causally dominated options in the context of static choice provides one of the central arguments for favoring CDT over EDT, CDT's failure to respect a dynamic version of this principle may well be disconcerting for the causalist.[15] While evidentialists are of course likewise prone to violations of a sequential causal dominance principle, they already reject the analogous static choice principle and so will be untroubled by this.

Nonetheless, the evidentialist is in a poor position to insist that the principles governing static choice be straightforwardly extendable to sequential choice. For example, when an agent is confronted with a finite choice set, the evidentialist forbids her from selecting an option that is dominated by another with respect to a partition of propositions, each member of which is *evidentially independent* of her available options. The sequential analogue of this prohibition would censure agents that are guaranteed to implement plans in sequential choice problems that are dominated in this way by other available plans. An evidentialist's satisfaction of this demand is far from guaranteed. The plan implemented in STN may well be *ex ante* dominated by the initially favored one, relative to some partition of propositions evidentially independent of the plans

that license dynamically inconsistent evaluations of plans. However, if such a restriction is to serve Ahmed's project of vindicating EDT against CDT, it would have to be independently motivated in terms of the plausible properties of planning conditionals, and no such motivation seems forthcoming.

[14] Ahmed has made a point like this in conversation.

[15] Not all causalists find the argument here disconcerting, however. See Joyce 2016 for a prominent rejoinder.

(e.g. perhaps {the predictor has blue hair, the predictor doesn't have blue hair} would suffice as such a partition). So the proponent of EDT cannot insist that normative principles which hold true at the level of static choice will always readily generalize to principles which hold true at the level of sequential choice.

The only way then for the evidentialist to maintain that dynamic consistency and like considerations favor her theory over the causalist's is to restrict the application of such principles to problems that can be modelled with decision trees lacking natural nodes. In general, however, it has not been thought that the status of dynamic consistency as a principle of rationality shifts depending upon the sort of sequential choice problem one has in focus. Dynamic consistency is rather held out by its defenders as a plausible principle governing sequential choice within a broad range of contexts, including those that encompass learning from nature. My sympathies lie with the crowd here. The considerations that motivate dynamic consistency as a principle of rationality in the first place (e.g. the apparent irrationality involved in failing to implement what one takes to be an available and optimal strategy) make no discrimination between problems that include learning from nature and those that exclude it. Dynamic consistency is then a dangerous weapon for the evidentialist to wield against the causalist as it is easily turned against her own theory as well.

Given Ahmed's results and those discussed here, it may be fair to conjecture that any plausible decision theory that follows the Logic of Decision and its causalist variants in allowing for probabilistic dependence between acts and causally independent states of the world is likely to suffer from the plague of dynamic inconsistency and its associated vices.[16] Any agent who takes her choices to be non-causally correlated with states of the world may well find herself in a precarious position vis-a-vis sequential coherence. This suggests a dilemma for decision theorists: either forsake dynamic consistency as a standard of normative adequacy or disavow the rationality of credences that treat acts as evidence for states that they have no tendency to cause. Which horn to embrace is a matter for further reflection.

A Planning Conditionals and Multi-Dimensional Possible Worlds Semantics

§3 provided a pair of credence and desirability functions that lead to dynamic inconsistency when employed by an evidentialist in STN, under the assumption that planning conditionals are truth-functional. This appendix supplies a pair of credence and desirability functions defined over triples of possible worlds that

[16]One might be tempted at this point to take refuge in the 'functional decision theory' of Yudkowsky and Soares 2017. Certainly, the attractive dynamic properties of this theory constitute a main source of its appeal. Still, insofar as it aspires to characterize rational *choice*, it appears to make absurd recommendations in cases like Transparent Newcomb. (See p. 22-3 of Yudkowsky and Soares 2017.)

witness to the same result, where planning conditionals are now understood in terms of Richard Bradley's multi-dimensional possible worlds semantics. The first entry of a triple designates which world is actual, the second designates which world would be actual on the supposition that K, and the third designates which would be actual on the supposition that $\overline{K}$. For ease of expression, I abbreviate '$S(z_i)$' as 'z_i' in what follows. I also neglect triples of the form $\langle z_i, z_j, z_k \rangle$, where either $z_j \notin K$ or $z_k \notin \overline{K}$, to which a rational agent assigns no credence. Further, if $z_i \in K$, it is assumed that $z_i = z_j$, and that, similarly, if $z_i \in \overline{K}$, then $z_i = z_k$. (For the details of Bradley's semantics, see Bradley 2017.)

		P_0	V_0
w_1	$\langle z_1, z_1, z_3 \rangle$	.007	1,001,000
w_2	$\langle z_1, z_1, z_4 \rangle$	.007	1,001,000
w_3	$\langle z_2, z_2, z_3 \rangle$	.35	1,000,000
w_4	$\langle z_2, z_2, z_4 \rangle$	.35	1,000,000
w_5	$\langle z_3, z_1, z_3 \rangle$	.07	1,000
w_6	$\langle z_3, z_2, z_3 \rangle$	.07	1,000
w_7	$\langle z_4, z_1, z_4 \rangle$	.035	0
w_8	$\langle z_4, z_2, z_4 \rangle$	.035	0
w_9	$\langle z_5, z_1, z_3 \rangle$	.019	0
w_{10}	$\langle z_5, z_1, z_4 \rangle$	.019	0
w_{11}	$\langle z_5, z_2, z_3 \rangle$	.019	0
w_{12}	$\langle z_5, z_2, z_4 \rangle$	.019	0

The plan Π_1 is true just at worlds w_4 and w_8, that is: $\langle z_2, z_2, z_4 \rangle$ and $\langle z_4, z_2, z_4 \rangle$, so we compute its value as:

$$\begin{aligned} V(\Pi_i) &= [V(w_4)P(w_4) + V(w_8)P(w_8)]/P(w_4 \vee w_8) \\ &= 350,000/0.385 \\ &= 909,090.\overline{90} \end{aligned}$$

We may similarly compute the values of the other plans:

(1) $V_0(\Pi_2) = 91,909,\overline{09}$

(2) $V_0(\Pi_3) = 833,499.\overline{6}$

(3) $V_0(\Pi_4) = 166,833.\overline{3}$

(4) $V_0(\Pi_5) = 0$

These are the same values reached in §3 and so the verdict remains unaltered: Π_1 is V_0-maximal but neither V_2- nor V_3-maximal, and so dynamic inconsistency results.

Works Cited

Ahmed, Arif. *Evidence, Decision and Causality.* Cambridge: Cambridge University Press, 2014. Print.

Bradley, Richard. *Decision Theory with a Human Face.* Cambridge: Cambridge University Press, 2017. Print.

Cubitt, Robin. "Rational Dynamic Choice and Expected Utility Theory". *Oxford Economic Papers* 48 (1996): 1–19. Print.

Egan, Andy. "Some Counterexamples to Causal Decision Theory". *Philosophical Review* 116 (2007): 93–114. Print.

Gibbard, Allan and William Harper. "Counterfactuals and Two Kinds of Expected Utility". *Foundations and Applications of Decision Theory.* Edited by J. Leach C. Hooker and E. McClennen. Reidel: Dordrecht, 1978. 125–162. Print.

Hammond, Peter. "Consequentialist Foundations for Expected Utility". *Theory and Decision* 25 (1988): 25–78. Print.

Jeffrey, Richard. *The Logic of Decision.* Chicago: University of Chicago Press, 1965/1983. Print.

Joyce, James. "Review of Arif Ahmed: Evidence, Decision and Causality". *Journal of Philosophy* 113(4) (2016): 224–232. Print.

Lewis, David. "Causal Decision Theory". *Australasian Journal of Philosophy* 59 (1981): 5–30. Print.

Maher, Patrick. "Symptomatic Acts and the Value of Evidence in Causal Decision Theory". *Philosophy of Science* 3 (1990): 479–98. Print.

McClennen, Edward. *Rationality and Dynamic Choice.* Cambridge: Cambridge Univeristy Press, 1990. Print.

Meacham, Christopher. "Binding and Its Consequences". *Philosophical Studies* 149 (2010): 49–71. Print.

Nozick, Robert. "Newcomb's Problem and Two Principles of Choice". *Essays in Honor of Carl G. Hempel.* Edited by Nicholas Rescher. Reidel: Dordrecht, 1969. 107–33. Print.

Skyrms, Brian. "Causal Decision Theory". *Journal of Philosophy* 79 (1982): 695–711. Print.

—–."The Value of Knowledge". *Scientific Theories.* Edited by C.W. Savage. Minneapolis: University of Minnesota Press, 1990. 245–66. Print.

Wells, Ian. "Equal Opportunity and Newcomb's Problem". *Mind* (2018). Print.

Yudkowsky, Eliezer and Nate Soares. "Functional Decision Theory". *CoRR* (2017): http://arxiv.org/abs/1710.05060. Print.