Bradley Conditionals and Dynamic Choice

Simon M. Huttegger and Gerard J. Rothfus

November 2020

Abstract

One of the main contributions of Richard Bradley's book is an elegant extension of Jeffrey's Logic of Decision that countenances the evaluation of conditional prospects. This extension offers a promising new setting in which to model dynamic choice. In Bradley's framework, plans can be understood as conditionals of an appropriate sort, while dynamic consistency can be viewed as providing a constraint on the evaluation of conditionals across time. In this paper, we study connections between planning conditionals and dynamic consistency.

1 Introduction

In broad outline, Richard Bradley[1] follows the approach to decision theory pioneered by Jeffrey in his *Logic of Decision*.[2] In this framework, agents hold preferences over the propositions of a rich Boolean algebra. The content of these propositions is left entirely open, thus generalizing Savage's theory in which the domain of preference is far more restricted.[3] Consistency conditions constrain admissible preferences in light of the structure of the algebra and its logical connectives. The standard such conditions of Jeffrey's theory suffice to guarantee the representability of preferences as *desirability maximizing*. That is, there exists a probability function, P, and a desirability function, V, both defined over the algebra, satisfying for any pair of non-null, mutually exclusive propositions X and Y, $V(X \vee Y) = P(X|X \vee Y)V(X) + P(Y|X \vee Y)V(Y)$,[4] such that more preferred propositions are ranked higher by V. Jeffrey suggested that we could turn this into a theory of rational choice by identifying every option available to an agent facing a decision problem with the proposition that she selects that option and then stipulating that rational choice goes by desirability maximization.

Bradley's principal extension of Jeffrey's framework involves the introduction of conditional operators into the object language of the logic of decision.

---

[1] Bradley 2017.
[2] Jeffrey 1965/1983.
[3] Savage 1954.
[4] This is sometimes referred to as the 'Jeffrey-Bolker mixing law'.

This enables us to consider agents that express preferences over non-truth functional conditional propositions, including (plausibly) indicative and subjunctive conditionals. Bradley proves that doing so does no damage to the standard representation results and argues this (again plausibly) allows us to embed standard Savage-style decision theories into a Jeffrey-style framework. This is a clear virtue of Bradley's work, as it allows, for example, a reframing of debates about the rationality of Savage's postulates and competing ones within a unified framework.

We are going to study Bradley's extension of the logic of decision to conditional operators in the setting of *dynamic* or *sequential choice*. Most decision problems discussed in philosophy have a static flavor: an agent makes a one-time choice from among a set of acts. Many decision situations involve a temporal component, however. Choices are made sequentially, perhaps mixed with receiving partial information about the state of the world. How should an agent's actions be modeled to fit the sequential environment? We take Bradley's account of conditional operators as our starting point in order to develop a plausible response.

In what follows we are guided by certain intuitive considerations about choices in a dynamic setting. The central notion, we take it, is that of a *plan*. When we speak of plans we usually refer to possible courses of action that extend over time. For instance, a company may plan to adjust the level of production of some good in response to a range of events that might influence demand for the good. Such plans involve conditionals in a natural way: a plan tells us what to do *if* an event happens for a range of events that we can anticipate ex ante.

This ordinary understanding of plans involves, arguably, a certain stability over time. If I endorse the assertion "plan *A* is better than plan *B*" one day, but assert "plan *B* is better than plan *A*" on the following day, it is not clear whether I really do think of *A* and *B* as plans, assuming they cover all relevant contingencies. If my initial endorsement of plan *A* over plan *B* is to be genuine, it presumably factors in all the contingencies I'm aware of; thus, my endorsement of "plan *A* is better than plan *B*" should not change at my whim.

This suggestion—that agents should evaluate plans in a dynamically stable manner—comes from how we seem to understand plans in natural language. Its counterpart in rational choice theory is *dynamic consistency*. An agent who fails to evaluate plans in a robust manner over time—like endorsing plan *A* over plan *B* first but switching subsequently—is *dynamically inconsistent*. Dynamic inconsistency opens decision makers up to dynamic Dutch books, that is, to sets of bets that, over time, leave them with a net loss no matter what. For this reason, many decision theorists think of dynamic consistency as a requirement of rationality. Raiffa was among the first to realize the connection between dynamic consistency and Savage's axioms for decision theory, especially the

sure-thing principle.[5] The most thorough account was given by Hammond, who used dynamic consistency (or what he calls *consequentialism*) as a justification for expected utility maximization.[6] Our goal is to study these ideas within the Jeffrey-Bradley framework. In particular, we study the interplay between dynamic consistency and different ways to characterize plans in terms of conditionals. Our main result asserts a tight connection between *indicative* planning conditionals (what we refer to as *Bradley conditionals* below) and dynamic consistency. Here, planning conditionals are understood in a matter of fact way, and this is what fits the evaluation of plans according to Jeffrey's *V*.

If one thinks, as Jeffrey himself did, that the logic of decision is the correct theory of choice, then having an extension of the theory to the sequential choice setting in which agents are dynamically consistent when choosing among plans is of obvious significance, especially in light of recent arguments casting doubt on the possibility of such an extension.[7] Jeffrey's view is not widely shared, however.[8] Causal decision theorists, in particular, argue that the logic of decision fails as a theory of choice in Newcomb type problems. From this point of view, our project takes on a different kind of relevance. While denying that the logic of decision is the correct theory of *choice*, many causal decision theorists think of it is the correct theory of *desirability* or *news value*.[9] What examples like Newcomb's problem show is that desirability and choice-worthiness can come apart. What is important for our project is that evaluations of plans in terms of their news value or desirabilty should also be dynamically consistent. That is, we take an agent who thinks of plan *A* as more desirable as plan *B* one day and reverses that judgment on the next as irrational. In this way our results on planning conditionals are relevant even for those who don't think that Jeffrey's framework as the correct theory of rational choice.

We proceed as follows. In §2 we introduce dynamic choice theory and give some motivations for dynamic consistency. In §3 we discuss plans and conditionals. Our main results are stated and proved in §4 and given some context in §5.

2   Dynamic Choice Theory

While Jeffrey seems to have developed his theory largely with problems of *static choice* in mind, the concerns of practical rationality extend beyond such choice scenarios. We are often confronted with extended decision problems

---

[5]Raiffa 1968.

[6]Hammond 1988. See Cubitt 1996 for a simplified approach and McClennen 1990 for a philosophical assessment. Hammonds approach is endorsed by Gibbard 2008.

[7]See, for example, the argument in Rothfus 2020 that a natural sequential extension of the Logic of Decision results in dynamic inconsistency. Our results here offer evidentialists a way of evading this result by reinterpreting planning conditionals.

[8]There are some exceptions, such as Ahmed 2014.

[9]See, e.g., Joyce 1999 and Skyrms 1984.

that require planning and foresight on our part, rather than a single choice here and now. In dynamic choice problems, agents are called upon to make a series of choices over time that jointly determine (with the states of nature) what outcomes will obtain. These decision problems are often helpfully modelled using *Bayesian decision trees*. Such trees are cumbersome to define formally[10] but easily understood by example.

Intuitively, a decision tree consists of a set of nodes connected to each other and ordered according to an immediate successor function, $N_+(\cdot)$, that maps nodes to other nodes so as to produce a tree-like graph. Since the trees are meant to model bounded sequential choice problems, they always start with an initial node and end in various terminal nodes. All non-terminal nodes are either choice nodes, i.e. points at which the agent must make a choice, or natural nodes, i.e. points at which nature makes a move and some uncertainty about the world is resolved. Following convention, choice nodes are represented in trees as squares, while natural nodes are denoted by circles. Terminal nodes are designated with triangles. Each node n is associated with a proposition, $S(n)$, capturing the information state of the acting agent at n. An example of a relatively simple decision tree is depicted below.
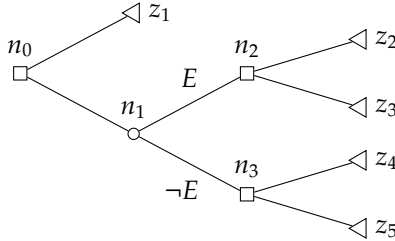

Figure 1: T_1, a decision tree.

The agent facing this tree, T_1, must decide at the initial node, which is a choice node, whether to move up or to move down. If the agent decides to move up, the proposition $S(z_1)$ will be true. If the agent chooses to move down, she will learn whether or not some event E obtains and then, depending on the truth of E, face either a choice between making $S(z_2)$ or $S(z_3)$ true or a choice between making $S(z_4)$ or $S(z_5)$ true.

Agents facing dynamic choice problems can reflect not only on what action to take at a currently occupied choice node, but also more broadly on what course of action to implement in the decision problem viewed as an extended whole. That is, they can evaluate competing *plans*. Given a tree T, a plan

---

[10]See Hammond 1988 for a precise characterization of decision trees. For our purposes below, we depart from this standard characterization in one significant way: we allow that learning is possible pursuant to choice nodes as well as natural nodes. For more on sequential choice problems and the philosophical issues they give rise to see, e.g., Cubitt 1996 and McClennen 1990.

specifies a unique move for every choice node in T that an agent facing T could reach, given implementation of earlier portions of the plan. It thus traces a unique path through the tree, given any combination of moves by nature at its nodes. So, for example, in T_1, there are five possible plans: (i) move up at n_0, (ii) move down at n_0 and then either move up at n_2 or move up at n_3, (iii) move down at n_0 and then either move down at n_2 or down at n_3, (iv) move down at n_0 and then either up at n_2 or down at n_3, and (v) move down at n_0 and then either down at n_2 or up at n_3. We let '$\Omega(T, n)$' stand for the set of all plans that are available at node n in decision tree T.

The primary normative constraint in the setting of sequential choice is *dynamic consistency*. Rational planning requires a certain coherence between initial evaluations of plans and subsequent re-evaluations of plan continuations at later choice nodes. If at the start of a sequential choice problem a certain plan seems most favorable to you, then continuations of that plan ought to continue to seem favorable to you as you implement the plan. There are various ways one might codify such a principle, but for our purposes we will rely on a fairly weak understanding of dynamic consistency.[11] Letting $D(\cdot)$ be a function that, given a set $\Omega(T, n)$ of plans, picks out those plans that a fixed agent judges to be practically acceptable and letting '$p(n)$' denote the continuation of plan p at node n, we can define dynamic consistency as follows:

Definition 1. An agent is *dynamically consistent* in a decision tree T just in case, for all nodes n_a and n_b in T such that n_a precedes n_b along some branch of T, if $p \in D(\Omega(T, n_a))$ and p makes arrival at n_b possible, then $p(n_b) \in D(\Omega(T, n_b))$. An agent is *dynamically consistent tout court* just in case the agent is dynamically consistent in all decision trees,

Note that, unlike some authors, we do not state dynamic consistency as a biconditional. While ex ante optimal plans should remain optimal (and therefore implementable) as a rational agent progresses through a dynamic choice problem, we do not insist that an initially disfavored plan may never come to be seen as admissible by a rational agent. While the stronger biconditional version of dynamic consistency may also be a rationality principle, it seems to have questionable consequences. In particular, it seems to be at odds with the rational permissibility of incomplete preferences. For example, if I view the prospects α and β as incommensurable in value, and must decide between α and a future choice between β and $\alpha - \epsilon$ (where ϵ is some small cost that detracts from α without rendering it comparable to β), I may well judge the plan that involves opting for the choice and then selecting $\alpha - \epsilon$ as initially inadmissible but revise this judgment upon arriving at the choice between β and $\alpha - \epsilon$, thus violating the stronger, but not our, formulation of dynamic consistency. It is more critical in our eyes that rationality allow for the possibility of avoiding suboptimal plans than that it guarantee the avoidance of such plans. At any

---

[11]For more on dynamic consistency and for stronger versions of the principle, see McClennen 1990.

rate, the weaker version of dynamic consistency laid out in Definition 1 suffices for our purposes in this paper, so there is no need to assume the stronger version.

It is plausible that, under appropriate assumptions, rationality requires dynamic consistency so understood. Dynamically inconsistent agents are doomed to forseeably reverse their current judgments about which plans of action are best, and normative theories that license such reversals are a tough sell. One reason for this is that dynamic inconsistency typically breeds susceptibility to dynamic Dutch Books: e.g., if I prefer plan A to plan B, but anticipate a reversal of that preference prior to completing either plan, I may well be decisively inclined to pay a fee to commit myself to plan A, leaving me strictly worse off than I would have been if my attitudes hadn't been prone to such problematic reversals.[12] The dynamic consistency (or lack thereof) of Jeffrey's theory of desirability maximization is then well worth examining.[13] For convenience, we will speak of a decision theory, like desirability maximization, as dynamically consistent just in case any agent whose attitudes and behavior conform to the theory is guaranteed to be dynamically consistent.

In order to explore the dynamic consistency of desirability maximizers, though, we need to know how to apply the theory in the context of sequential choice problems. This necessitates finding some way of identifying the plans available to an agent in a sequential choice problem with propositions so that available plans can then be ranked according to their desirability. Richard Bradley's introduction of conditional operators into Jeffrey's framework suggests a particularly fruitful way of associating plans with propositions (or, as Bradley prefers to say, prospects), allowing us to assess the dynamic consistency of desirability maximization more carefully.

---

[12]Of course, some authors do deny that rationality is incompatible with susceptibility to dynamic dutch books of this sort. See, for example, Hedden 2015. A reviewer has neatly pointed out, however, that even readers skeptical of the rational necessity of dynamic consistency might still take interest in our arguments on account of the simple pragmatic grounds agents may have for structuring their attitudes so as to avoid dynamic dutch books. An argument casting the significance of dynamic dutch book arguments in this light can be found in Rabinowicz 2006.

[13]A reviewer has objected that the notion of dynamic consistency suggested here may be too strong to count as a requirement of rationality since an agent may be dynamically inconsistent in a decision tree while still implementing the plan she judges to be ex ante optimal at the tree's initial node. For example, I might initially judge a plan A to be optimal in a tree, but then lose that preference at a subsequent node (thus violating dynamic consistency) and yet regain the preference for A again at yet another subsequent node, such that I still end up implementing plan A. Given that I succeed in implementing an ex ante optimal plan, why need my dynamic inconsistency indicate any irrationality on my part in cases like this? The answer is that while I do end up implementing what was optimal from the standpoint of my initial attitudes, I do not do so from my standpoint at the middle stage, so I am still criticizable by my own lights at that middle stage. And there is no essential reason to privilege the initial attitudes when assessing the agent's rationality. If we just snip off the first part of the tree, we'd have a *de novo* tree where the middle node would become the initial node and I would fail to implement my ex ante optimal plan in that tree.

3 Plans and Conditionals

Bradley (2017) extends Jeffrey's logic of decision in a natural way so as to include conditional propositions. He develops a rich framework for capturing the role conditionals play in decision making, and also provides an innovative new semantics for conditional statements. We focus here only on the first aspect, which we apply to the kinds of dynamic choice problems introduced in the foregoing section.

Above we sketched an informal understanding of the plans available to an agent in a fixed decision tree. Now, employing a conditional operator, we can define the plans available to a decision maker as she moves through a decision tree more precisely within the Jeffrey-Bradley framework. In the following definition, the symbol '$\rightarrow$' should be read as a placeholder for different kinds of conditionals.

Definition 2. Let n be a node in a decision tree T. The *set of plans* available at n in T, denoted $\Omega(T, n)$, is defined recursively as follows:

1. If n is a terminal node, then $\Omega(T, n) = \{S(n)\}$.

2. If n is a choice node, then

$$\Omega(T, n) = \{S(n') \wedge \pi(n') : n' \in N_+(n), \pi(n') \in \Omega(T, n')\}.$$

3. If n is a natural node, then

$$\Omega(T, n) = \left\{ \bigwedge_i [S(n_i) \rightarrow \pi(n_i)] : n_i \in N_+(n), \pi(n_i) \in \Omega(T, n_i) \right\}.$$

The first two conditions define plans at terminal nodes and choice nodes in an obvious enough way. The last condition requires the agent to consider plans at natural nodes in terms of a conditional '$\rightarrow$'. At a natural node, plans are conjunctions of conditionals of the form $S(n_i) \rightarrow \pi(n_i)$. The propositions $S(n_i)$ form a partition of those circumstances that can be distinguished after n. For each n_i, a plan specifies what plan available at n_i to choose *if n_i is reached*. We can also define plan continuations in this setting as follows:

Definition 3. Let n be a node in a decision tree T and let $p \in \Omega(T, n)$. Suppose n' is a node succeeding n along some branch of T. The *continuation* of p at n', written '$p(n')$', is the member of $\Omega(T, n')$ consistent with p. If there is no such member, p does not make arrival at n' possible and $p(n')$ is left undefined.

Throughout we will assume that desirability maximizers judge admissible at a node n in a tree T whichever plans in $\Omega(T, n)$ are of maximal desirability, i.e. $p \in D(\Omega(T, n))$ if and only if $V_n(p) \geq V_n(p')$ for all $p' \in D(\Omega(T, n))$, where V_n captures the agent's desirabilities at node n. We will also assume that the

credences and desirabilities of an agent each evolve by conditionalization as she moves through the stages of a sequential choice problem.[14]

Let's call the conditional appearing in Definition 2 a *planning conditional*. The question we'd like to answer is how to construe this conditional so as to ensure the dynamic consistency of desirability maximization. We can note as a preliminary result that construing the planning conditional truth-functionally as a material conditional will not do. This way of modelling the planning conditional is equivalent to identifying a plan with the disjunction of information states that its implementation could terminate in. While this proposal is natural enough and would allow us to analyze the planning conditional in terms of standard Boolean connectives, Rothfus 2020 has argued that this simple model can leave desirability maximizers prey to dynamic inconsistency, establishing:

Proposition 1. *If the planning conditional is the material conditional, desirability maximization is not dynamically consistent.*

This result is rather to be expected. The material construal of the planning conditional makes plans out to be disjunctions, and one disjunction, say, $A \lor B$, can easily be more desirable than another, say, $C \lor D$, even if the individual disjuncts (i.e. plan continuations) are such that C is preferable to A and D to B. (This may be the case, for example, if B is more desirable than C, and $A \lor B$ is sufficiently good evidence of B while $C \lor D$ is sufficiently good evidence of C.)[15]

This initial negative result invites consideration of alternative interpretations of the planning conditional. One natural suggestion would be to read the planning conditional as an indicative of the sort employed by Bradley to model Savage acts. It turns out that this construal of the planning conditional yields a more positive picture regarding the dynamic consistency of desirability maximization, as we hope to show below. Before arguing for this, however, it may be worth motivating the idea that planning is an indicative, as opposed to subjunctive, enterprise at a relatively informal level.[16]

---

[14]If n_a preceeds n_b in a decision tree, then updating one's desirabilities by conditionalization means that $V_{nb}(x) = V_{na}(x|S(n_b)) = V_{na}(xS(n_b)) - V_{na}(S(n_b))$. See Bradley 2017, p. 97, for more on Conditional Desirability.

[15]Rothfus puts his argument forward in the context of Ahmed 2014's argument that causal decision theory leads to dynamic inconsistency. This result shows that the same is true for evidential decision theory on a material construal of planning conditionals.

[16]Note that we in no way mean here to deny the importance of subjunctive supposition in the context of rational planning and decision making. Causal decision theory, for example, mall well be right to view the practical merits of a plan in terms of its expected desirability under the *subjunctive* supposition of its implementation. What we deny is simply that the planning conditional itself should be viewed subjunctively. Note also that plans are different from strategies in extensive-form games. An extensive-form game is like a sequential choice problem with the added feature that there can be more than one agent making decisions. Agents are called players. A strategy is a *complete* contingency plan, that is a plan that specifies a unique move at each of a player's choice nodes regardless of previous moves. Thus, in the context of strategies, subjunctive conditionals may play an important role. It would be interesting to explore the connections between strategies and plans within Bradley's framework, but we leave this topic to future research.

An *indicative* supposition involves supposing that some proposition is true as a matter of fact. This mode of supposition is commonly contrasted with *subjunctive* or *counterfactual* supposition, which involves considering what the world would be like if some proposition were to be true (without fixing whether or not the supposition is actually true). Examples are the surest way to bring out this contrast. Suppose, to employ a planning context, that you are considering the possibility that you will be offered a job at a prestigious law firm and are evaluating the desirability of accepting such an offer, under the supposition that it is made. Suppose further that you suffer from terribly low self-esteem and hence are very confident that you will not be offered the position. Moreover, you think that, were you, shockingly, to learn that you had been offered the job, the most likely explanation would be that the job was not as grand as you had supposed and therefore not worth accepting. Under these circumstances, you may judge accepting the offer as desirable under the subjunctive supposition of its being offered but not under the indicative supposition of its being offered. To put the matter in terms of conditionals, and letting $\mapsto$ represent the indicative and $\Box\!\!\rightarrow$ the subjunctive, you prefer (Offer $\mapsto$ NotAccept) to (Offer $\mapsto$ Accept), while simultaneously preferring (Offer $\Box\!\!\rightarrow$ Accept) to (Offer $\Box\!\!\rightarrow$ NotAccept).

This example hopefully brings out why planning conditionals are best understood as indicative rather than subjunctive. In forming a contingency plan, an agent is considering what to do if, as a matter of fact, various contingencies are found to obtain. When I consider what to do if I am offered the job, I am considering what to do if I in fact learn that I am offered the job. Strictly counterfactual worlds are of no concern to me, and the counterfactual conditional provides me no direct practical guidance. In forming a plan, I am determining how to respond to the different bodies of evidence I might be exposed to, and not how to respond to counterfactual possibilities, which is impossible for an agent located in one world to do. Hence, an indicative reading seems most appropriate given the role planning conditionals are meant to play in the practical deliberation of agents.

There is another reason why in the context of dynamic consistency an indicative reading of planning conditionals is preferable. Suppose that plans involve counterfactual contingencies, i.e. they specify moves at nodes that are not reached if the decision maker follows her plan. This can happen, for example, if the agent makes a mistake or acts irrationally at some node. If this happens, though, it is not clear why the agent should be dynamically consistent along the "counterfactual" paths of the decision tree. For then she might learn something about herself that could overturn her initial evaluations of plans.[17] That said, it remains to vindicate more precisely the dynamic consistency of

---

[17]This point is related to certain issues in epistemic game theory, where the consideration of counterfactuals is very natural. As Stalnaker has pointed out, there are no substantive rationality constraints for considering what to do at nodes in a game tree that are not reached under the assumptions of a model of a game. Everything depends on how the agent changes her attitudes if something happens that she does not believe initially. See Stalnaker 1996; Stalnaker 1998.

desirability maximization under such a reading of the planning conditional.

4   Conditionals and Dynamic Consistency

Bradley provides an insightful discussion of the constraints that apply to different types of conditionals—especially indicatives. There is, as we have claimed and now hope to demonstrate, a very close connection between indicative conditionals as construed by Bradley and the dynamic consistency of desirability maximization in sequential choice problems. For our purposes, it will not be necessary to develop a full semantics for indicative conditionals (see Chapter 8 in Bradley 2017 for more details). Instead, we use two of Bradley's constraints on value functions that arguably capture properties any account of indicative conditionals ought to satisfy in a decision-theoretic context.

The first condition is called the *Indicative Property* (Bradley 2017, p. 117):

For all propositions α, β, γ, $V(\alpha \mapsto \beta) \geq V(\alpha \mapsto \gamma)$ iff $V(\alpha \wedge \beta) \geq V(\alpha \wedge \gamma)$.

As Bradley points out, the indicative property is a very natural constraint for indicative conditionals. Indicative conditionals are matter-of-fact conditional statements. And part of what this means is that the evaluation of the desirability of $\alpha \mapsto \beta$ involves supposing that α is true as a matter of fact. As a result, if the agent prefers β to γ, in each case supposing α as a matter of fact, she should also prefer $\alpha \wedge \beta$ to $\alpha \wedge \gamma$.

A second property Bradley attributes to indicatives is *Additivity*, which requires that whenever $\{\alpha_i\}$ is a partition, then

$$V(\bigwedge_i (\alpha_i \mapsto \beta_i)) = \sum_i V(\alpha_i \mapsto \beta_i).$$

Additivity is also a natural requirement for indicative conditionals. It basically says that conjunctions of conditionals can be evaluated separately in case the antecedents form a partition. If each antecedent is supposed as a matter of fact, it only specifies what happens if that antecedent, and no other, is the case. There is thus no influence on the consequents of the other conditionals.

The two foregoing properties allow us to introduce the following definition.

Definition 4. The conditional $\mapsto$ is a *Bradley indicative* if it satisfies the Indicative Property and Additivity.

Our first result shows that if planning conditionals are Bradley indicatives, then the dynamic consistency of desirability maximization in a planning context is guaranteed. Before we establish this claim, however, we take note of a lemma, proven in Rothfus 2020, that will be useful to invoke as we consider the

relationship between properties of the planning conditional and the sequential coherence of desirability maximization.

Lemma 1. *For any decision tree T, if n is a choice node in T and $n' \in N_+(n)$, then $V_n(p) \geq V_n(p')$ iff $V_{n'}(p(n')) \geq V_{n'}(p'(n'))$ for all plans $p, p' \in \Omega(T, n)$ consistent with $S(n')$.*

This lemma, which holds independently of how we opt to construe planning conditionals, establishes that the relative desirabilities of plans never shift following choice nodes. The only possible opportunities for dynamic inconsistency on the part of desirability maximizers arise following new disclosures of information by nature. To get an intuitive sense for why this is the case, note that, at any choice node, a plan specifies a particular choice to make at that node, and the significance of making this choice is already factored into the desirability maximizer's appraisal of the plan. Hence, the act of initiating the plan in question by selecting the option it recommends at that choice node can have no tendency to engender a preference reversal among plans. Thus, the only possible opportunities for dynamic inconsistency on the part of deisrability maximizers arise following new disclosures of information by nature.

With this lemma in hand, we turn to our key result:

Proposition 2. *If the planning conditional is a Bradley indicative, then desirability maximization is dynamically consistent.*

Proof. Let a non-terminal node n in an arbitrary Bayesian decision tree T be fixed. Suppose that Π is a desirability maximal plan at n. All we need to convince ourselves of is that for any successor to n, say n', $\Pi(n')$ is either undefined (i.e. Π did not make arrival at n' feasible) or desirability maximal at n'. So, let n' be a successor to n and let $\Pi(n')$ be defined. The node n is either a choice node or a natural node. If n is a choice node, $\Pi(n')$ is optimal at n', because choice selections never result in desirability reversals over plans (by Lemma 1 above). So we are left to consider the case where n is a natural node. In this case, Π is of the form $\wedge_i[S(n_i) \mapsto \Pi(n_i)]$, where the n_i are the possible successors to n. By Additivity, we know that:

(1) $V_n(\Pi) = V_n(\wedge_i[S(n_i) \mapsto \Pi(n_i)]) = \sum_i V_n(S(n_i) \mapsto \Pi(n_i))$

But this means that, for all $\Pi' \in \Omega(T, n)$:

(2) $V_n(S(n_i) \mapsto \Pi(n_i)) \geq V_n(S(n_i) \mapsto \Pi'(n_i))$

For, if any such inequality failed to hold, we could alter Π to form a new plan Π^* exactly similar to Π except that it substitutes the more preferred conditional for the less. By Additivity, this would generate a more desirable plan, contradicting the assumption that Π is optimal at n. But then, by the Indicative Property, (2) entails that, for all $\Pi' \in \Omega(T, n)$:

(3) $V_n(S(n_i) \wedge \Pi(n_i)) \geq V_n(S(n_i) \wedge \Pi'(n_i))$

Subtracting $V_n(S(n_i))$ from both sides yields, by Conditional Desirability:

(4) $V_n(\Pi(n_i)|S(n_i)) \geq V_n(\Pi'(n_i)|S(n_i))$

Hence, for all $n_i \in N_+(n)$ (and thus n', in particular):

(5) $V_{n_i}(\Pi(n_i)) \geq V_{n_i}(\Pi'(n_i))$

This completes the proof of the proposition. □

Therefore, the features of the Bradley indicative are sufficient to ensure that a plan of maximal desirability at the outset of a sequential choice problem will continue to enjoy maximal desirability throughout the course of its implementation.[18]

We can also prove a partial converse. Any agent whose attitudes towards planning conditionals violate the Indicative Property is liable to violate a stronger cousin of dynamic consistency that we might call *Preference Stability*.

Definition 5. An agent is *preferentially stable* in a decision tree T just in case, for all nodes n_a and n_b in T such that n_a precedes n_b along some branch of T, if $p, p' \in \Omega(T, n_a)$, $p \succeq p'$, and p, p' both make arrival at n_b possible, then $p(n_b) \succeq p'(n_b)$. An agent is *preferentially stable tout court* just in case the agent is preferentially stable in all decision trees.[19]

While this principle requires more than dynamic consistency, it is not without plausibility as a rationality constraint. There are at least two ways one might motivate it. First, taking desirability maximization as a choice policy, one might note that even if a dynamically consistent agent is practically guaranteed to implement whatever plan she initially sets her mind to, we might still hope that her rationality would minimize the losses associated with any hypothetical deviations from an optimal plan. For example, if a rational agent initially prefers a plan A to a plan B to a plan C, then conditional on failing to implement A, one might expect that a rational agent would still implement B over C, and the preservation of the agent's initial preference for B over C may support this. The (counterfactual) mistakes of a rational agent need not be haphazard. Second, desirabilities reflect judgments about the relative value of propositions considered as news items and these judgments plausibly ought to remain stable concerning plans, absent changing information about the structure of the planning problem an agent faces. If plan B is better news than plan C ex ante, it ought to remain so, conditional upon the receipt of information that plans B and C already take into account.

Proposition 3. *If the planning conditional fails to satisfy the Indicative Property with respect to a given desirability function, then an agent with that desirability function will be preferentially unstable.*

---

[18]We may also note, though we omit the proof, that neither the Indicative Property nor Additivity is sufficient by itself to guarantee the dynamic consistency of desirability maximization.

[19]Note that preference stability, like dynamic consistency, is not stated as a biconditional.

Proof. Suppose that an agent's desirability function fails to conform to the Indicative Property with respect to the planning conditional. That is, let V be a desirability function and α, β, γ be propositions such that: $V(\alpha \to \beta) \geq V(\alpha \to \gamma)$ while $V(\alpha\gamma) \geq V(\alpha\beta)$. To show that an agent with values represented by this desirability function is liable to preference instability it suffices to construct a decision tree in which the agent will be preferentially unstable. To do so, we will first consider the case where β and γ are incompatible propositions and then generalize.

Let z be the proposition expressing that α is true while β and γ are false, i.e. $z = \alpha(\overline{\beta \vee \gamma})$. Then, supposing β and γ are incompatible, the following decision tree suffices to bring out the claimed preference instability:[20]
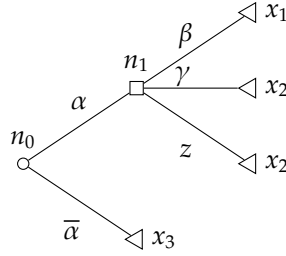

Figure 2: T_2, a decision tree.

At the initial node n_0, there are three plans available to an agent facing this decision problem, corresponding to the conditionals: $\alpha \to \beta, \alpha \to \gamma, \alpha \to z$. We know, by assumption, that the plan $\alpha \to \beta$ is preferred to the plan $\alpha \to \gamma$. But at n_1, this preference is guaranteed to reverse conditional upon reaching n_1. To see this, note that the continuation of $\alpha \to \beta$ at n_1 is $\alpha\beta$, while the continuation of $\alpha \to \gamma$ is $\alpha\gamma$. We know, by assumption, that:

(1) $V_{n_0}(\alpha\beta) < V_{n_0}(\alpha\gamma)$

But this implies that:

(2) $V_{n_0}(\alpha\beta) - V_{n_0}(\alpha) < V_{n_0}(\alpha\gamma) - V_{n_0}(\alpha)$

This is, by definition:

(3) $V_{n_0}(\alpha\beta|\alpha) < V_{n_0}(\alpha\gamma|\alpha)$

Assuming that updating of desirabilities goes by conditionalization, we then have:

(4) $V_{n_1}(\alpha\beta) < V_{n_1}(\alpha\gamma)$

---

[20]If β and γ suffice to partition α, z will be a contradiction and can be pruned from the tree.

This suffices to establish the violation of Preference Stability on the assumption that $\beta\gamma = \perp$.

If β and γ are not incompatible, however, T_2 will be an ill formed decision tree. Nonetheless, we may assume that this is not the case since an arbitrary violation of the Indicative Property always implies a violation of the Indicative Property with respect to a pair of consequents that are mutually exclusive (assuming closure of the agent's alegebra under the relevant logical operations)[21] and we can thus take β and γ to be such a pair. To prove that this is the case, suppose that β and γ were not incompatible. We then have three cases to consider: (i) $\beta \vDash \gamma$ (ii) $\gamma \vDash \beta$ and (iii) $\gamma \nvDash \beta$ and $\beta \nvDash \gamma$. We handle these cases in turn.

(i) Suppose that $\beta \vDash \gamma$. Define $\gamma^* = \gamma\overline{\beta}$. Clearly, β and γ^* are incompatible. We argue that there is an Indicative Property violation with respect to α and these two propositions, i.e. $V(\alpha \to \beta) \geq V(\alpha \to \gamma^*)$ while $V(\alpha\gamma^*) \geq V(\alpha\beta)$. Since β entails γ, $\alpha \to \gamma$ is equivalent to the disjunction of $\alpha \to \gamma^*$ and $\alpha \to \beta$. We then have, by the Jeffrey-Bolker mixing law, that $V(\alpha \to \gamma)$ is a weighted mixture of $V(\alpha \to \gamma^*)$ and $V(\alpha \to \beta)$, and since we know it is less than $V(\alpha \to \beta)$, it must also be that $V(\alpha \to \beta) > V(\alpha \to \gamma^*)$. However, similarly, $\alpha\gamma$ must be a weighted mixture of $\alpha\beta$ and $\alpha\gamma^*$. Thus, from the assumption that $V(\alpha\gamma) > V(\alpha\beta)$, we obtain that $V(\alpha\gamma^*) > V(\alpha\beta)$.

(ii) Suppose that $\gamma \vDash \beta$. Define $\beta^* = \beta\overline{\gamma}$. Clearly, β^* and γ are incompatible. We argue that there is an Indicative Property violation with respect to α and these two propositions, i.e. $V(\alpha \to \beta^*) \geq V(\alpha \to \gamma)$ while $V(\alpha\gamma) \geq V(\alpha\beta^*)$. Since γ entails β, $\alpha \to \beta$ is equivalent to the disjunction of $\alpha \to \gamma$ and $\alpha \to \beta^*$. We then have, by the Jeffrey-Bolker mixing law, that $V(\alpha \to \beta)$ is a weighted mixture of $V(\alpha \to \gamma)$ and $V(\alpha \to \beta^*)$, and since we know it is greater than $V(\alpha \to \gamma)$, it must also be that $V(\alpha \to \beta^*) > V(\alpha \to \gamma)$. However, similarly, $\alpha\beta$ must be a weighted mixture of $\alpha\beta^*$ and $\alpha\gamma$. Thus, from the assumption that $V(\alpha\gamma) > V(\alpha\beta)$, we obtain that $V(\alpha\gamma) > V(\alpha\beta^*)$.

(iii) Suppose that $\gamma \nvDash \beta$ and $\beta \nvDash \gamma$, while not being incompatible. Define $\beta^* = \beta\overline{\gamma}$ and $\gamma^* = \gamma\overline{\beta}$ and $\delta = \beta\gamma$. We then know that $V(\alpha \to (\beta^* \vee \delta)) > V(\alpha \to (\gamma^* \vee \delta))$, which implies that $V((\alpha \to \beta^*) \vee (\alpha \to \delta)) > V((\alpha \to \gamma^*) \vee (\alpha \to \delta))$. By the Jeffrey-Bolker mixing law, we get that $V(\alpha \to \beta^*) > V(\alpha \to \gamma^*)$. However, similarly, $V(\alpha(\beta^*\vee\delta)) < V(\alpha(\gamma^*\vee\delta))$, which implies that $V(\alpha\beta^* \vee \alpha\delta) < V(\alpha\gamma^* \vee \gamma\delta)$. By the Jeffrey-Bolker mixing law, we get that $V(\alpha\beta^*) < V(\alpha\gamma^*)$.

$\square$

---

[21]The closure of the algebra under the Boolean connectives is of course standard; further closure under the planning conditional is tied to our implicit assumption that the domain of decision trees an agent may face is unrestricted, i.e. any tree-like graph whose nodes are mapped to factual propositions in such a way that the propositions attached to the successors of any given node partition the proposition associated with that node counts as a decision tree.

5 Conclusion

Richard Bradley's work on conditionals constitutes a welcome and useful resource for decision theorists as they investigate various normative questions. We hope to have shown that not least among such questions are ones concerning planning and dynamic choice. Within a Jeffrey-style framework, Bradley's conditional operators offer a convenient way to model the plans available to an agent confronted with a sequential choice problem. Moreover, the properties of the employed conditionals are logically tied to the dynamic consistency (or lack thereof) of desirability maximization as a planning policy. The dynamic consistency of desirability maximization is assured whenever rational attitudes towards planning conditionals conform to the properties Bradley ascribes to indicatives, namely, Additivity and the Indicative Property. Moreover, satisfaction of the Indicative Property is necessary for a closely related sort of sequential coherence, namely, preference stability.

What we have sketched here is really an invitation to further inquiry regarding the interplay between the dynamic consistency of desirability maximization and the properties of planning conditionals. There are (at least) two open routes one might take in pursuing such further inquiry. First, one might attend more closely than we have to the nature of human planning and consider more fully what sort of conditionals we should employ to model such planning. We have suggested above some intuitive rationale for taking an indicative reading of planning conditionals, but there may also be rationale's for alternate readings and the implications of such other proposals for the dynamic consistency of desirability maximization have yet to be drawn out.

Secondly, one might take the desiderata of dynamic consistency and/or preference stability for granted and investigate further the properties that planning conditionals must satisfy in light of this demand. As we have shown, the Indicative Property and Additivty suffice for the satisfaction of dynamic consistency, while the Indicative Property is necessary for preference stability. Might we be able to say more than this and to formulate plausible constraints upon planning conditionals that are individually necessary and jointly sufficient for the dynamic consistency and/or preference stability of desirability maximization?

Of course, the interest of these projects may be somewhat lessened for those that believe that desirability maximization is not an adequate account of rational choice (which may be so if certain versions of causal decision theory are correct). As mentioned earlier, desirability still measures an interesting conative attitude that rational agents take toward prospects, which one might plausibly expect to be dynamically consistent over time in much the sense sketched above, even if it is not a fully adequate measure of the choice-worthiness of options. The investigation of dynamic consistency in the Bradley-Jeffrey framework initiated here is then of value independently of any resolution of the causal versus evidential debate in decision theory, as are, we believe, the extensions

of this work suggested above.

Acknowledgements

Works Cited

Ahmed, A. *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.

Gibbard, Allan. *Reconciling Our Aims: In Search for a Basis of Ethics*. Oxford, MA: Oxford University Press, 2008. Print.

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

Hedden, Brian. "Options and Diachronic Tragedy". *Philosophy and Phenomenological Research* 87 (2015): 423–51. Print.

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

Joyce, J. M. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press, 1999. Print.

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

Rabinowicz, Wlodek. "Levi on money pumps and diachronic dutch books". *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*. Edited by Erik Olsson. Cambridge: Cambridge University Press, 2006. 289–312. Print.

Raiffa, Howard. "Decision Analysis: Introductory Lectures on Choices under Uncertainty." (1968). Print.

Rothfus, Gerard. "Dynamic Consistency in the Logic of Decision". *Philosophical Studies* (2020). Print.

Savage, L. J. *The Foundations of Statistics*. New York: Wiley, 1954. Print.

Skyrms, B. *Pragmatics and Empiricism*. New Haven: Yale University Press, 1984. Print.

Stalnaker, Robert. "Belief Revision in Games: Forward and Backward induction". *Mathematical Social Sciences* 36 (1998): 31–56. Print.

——."Knowledge, Belief and Counterfactual Reasoning in Games". *Economics and Philosophy* 12 (1996): 133–163. Print.