Bradley Conditionals and Dynamic Choice

Simon M. Huttegger and Gerard J. Rothfus

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 explore some 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, $V(X \vee Y) = P(X|X \vee Y)V(X) + P(Y|X \vee Y)V(Y)$, 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 algebra's space of logical connectives. 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

[1] Bradley 2017.
[2] Jeffrey 1965/1983.
[3] Savage 1954.

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.

We want to explore here a further advantage of Bradley's innovations: modelling *dynamic* or *sequential* choice. In such choice problems, agents can evaluate extended plans for action as well as individual acts. We can usefully employ Bradley's framework to model these plans via appropriate conditional operators, which we may neutrally dub *planning conditionals*. This way of modelling plans then invites natural questions about the relationship between possible properties of planning conditionals and the dynamic consistency of desirability maximization. Our purpose here is to explore these connections. To do so, we first introduce the basic concepts of dynamic choice theory (§2) and then show how to model plans in the Jeffrey-Bradley framework (§3). We then prove that if the planning conditional exhibits the properties Bradley attributes to indicatives (namely, *Additivity* and the *Indicative Property*), dynamic consistency is guaranteed (§4). We also consider the implications of relaxing these conditions and discover that neither Additivity nor the Indicative Property alone is sufficient to preserve desirability maximization as a dynamically consistent policy. Finally, we conclude with a summary of these results and suggest some paths for further research (§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 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[4] 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

[4]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.

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.

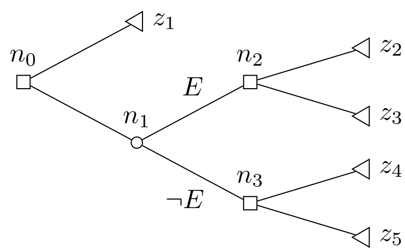

Figure 1: T_1, a decision tree.

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 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.[5] 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 *dy-*

[5] Plans are thus 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.

namic consistency.[6] 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.[7] 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:

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 $p(n_b) \in D(\Omega(T, n_b))$.

It is plausible that, under appropriate assumptions, rationality requires dynamic consistency. 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. The dynamic consistency (or lack thereof) of Jeffrey's theory of desirability maximization is then well worth examining. 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 arbitrary decision trees.

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.

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,

[6] See Hammond 1988 and Cubitt 1996.

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

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 '$\to$' 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
$$\Omega(T,n) = \{\bigwedge_i [S(n_i) \to \pi(n_i)] : n_i \in N_+(n), \pi(n_i) \in \Omega(T,n_i)\}.$$

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 '$\to$'. At a natural node, plans are conjunctions of conditionals of the form $S(n_i) \to \pi(n_i)$. 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.[8]

[8] 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.

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 2019 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, $A \vee B$, can easily be more desirable than another, say, $C \vee D$, 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 $A \vee B$ is sufficiently good evidence of B while $C \vee D$ is sufficiently good evidence of C.)[9]

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.

4 Conditionals and Dynamic Consistency

Bradley provides an insightful discussion of the constraints that may apply to different types of conditionals—especially indicative and subjunctive conditionals. We focus on indicative conditionals here. There is, as we shall see in a moment, a very close connection between certain types of indicative conditionals and dynamic consistency in sequential choice problems. More generally, we think that indicative conditionals are especially important for sequential choice problems, in that an agent considers the question of what she will in fact do after the uncertainty at a natural node is resolved.

For our purposes, it's not necessary to develop a full semantics for indicative conditionals (see Chapter 8 in Bradley 2017 for more details). Instead, we use two constraints on value functions that arguably capture properties any account

[9] 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.

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. A 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. Additivity basically says that conjunctions of conditionals can be evaluated separately in case the antecedents from a partition. If each antecedent is supposed as a matter of fact, it only specifies what happens if it 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 2019, 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. *Stuff*

Proposition 3. *If the planning conditional is a Bradley indicative, then desirability maximization is dynamically consistent.*

Proof. The proof is straightforward. 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', $\Pi(n_i)$ is either undefined (i.e. Π did not make arrival at n_i feasible) or desirability maximal at n'. So, let n_i be a successor to n and let $\Pi(n_i)$ 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_i, 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 $\wedge_i[S(n_i) \rightarrow \Pi(n_i)]$, where the n_i are the possible successors to n. By Additivity, we know that:

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

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

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

For, if any such inequality failed to hold, we could alter Π to form a new plan $\Pi*$ 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:

(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.

We can also prove a partial converse. In the presence of Additivity or the Indicative Property, there are sequential choice problem in which dynamic consistency fails whenever the other property is absent.

Proposition 4. *The following statements are true:*

1. *If Additivity holds, then desirability maximization is dynamically inconsistent whenever the Indicative Property fails.*
2. *If the Indicative Property holds, then desirability maximization is dynamically inconsistent whenever the Additivity fails.*

Proof. 1) Suppose that Additivity holds. To show that desirability maximization is dynamically inconsistent without the additional assumption that planning conditionals satisfy the Indicative Property, it suffices to construct a decision tree in which both a failure of the Indicative Property and a failure of dynamic consistency are exemplified on the part of a desirability maximizer.

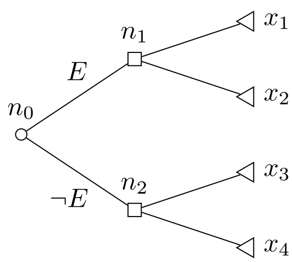

Figure 2: T_2, a decision tree.

Consider the choice problem with an initial natural node and two choice nodes, n_1 and n_2. The terminal nodes following n_1 are x_1, x_2, and those following n_2 are x_3, x_4. We assume that $V_{n_1}(x_2) > V_{n_1}(x_1)$. The tree modelling this choice problem is depicted in Figure 2. Suppose $V_{n_0}(S(n_1) \to x_1) > V_{n_0}(S(n_1) \to x_2))$, while $V_{n_0}(S(n_1)x_1) < V_{n_0}(S(n_2)x_2)$ (i.e. $V_{n_0}(x_1) < V_{n_0}(x_2)$),indicating that the planning conditional fails to satisfy the Indicative Property. By Additivity,

$$V_{n_0}((S(n_1) \to x_1) \wedge (S(n_2) \to x_i)) > V_{n_0}((S(n_1) \to x_2) \wedge (S(n_2) \to x_i)),$$

for $i \in \{3, 4\}$. Thus, the initially optimal plan(s) will be the one(s) that involves selecting x_1 at n_1. Thus the agent plans, upon reaching n_1, to choose x_1. However, by the posited failure of the indicative property, $V_{n_0}(x_1) < V_{n_0}(x_2)$. By conditionalization then, we may also infer $V_{n_1}(x_1) < V_{n_1}(x_2)$. So, upon reaching

n_1, the agent will switch to favoring x_2, resulting in dynamic inconsistency.

2) Suppose now that Additivity fails, while the Indicative Property is satisfied. Consider again a sequential choice problem starting with a natural node and two choice nodes, n_1 and n_2, resulting in outcomes x_1, x_2 and x_3, x_4, respectively (as depicted in Figure 2). Suppose, first, that initially

$$V_{n_0}((S(n_1) \to x_1) \wedge (S(n_2) \to x_3)) > V_{n_0}((S(n_1) \to x_2) \wedge (S(n_2) \to x_3)).$$

We assume the following violations of Additivity:

$$V_{n_0}((S(n_1) \to x_1) \wedge (S(n_2) \to x_3)) > V_{n_0}(S(n_1) \to x_1) + V_{n_0}(S(n_2) \to x_3)$$

and

$$V_{n_0}((S(n_1) \to x_2) \wedge (S(n_2) \to x_3)) < V_{n_0}(S(n_1) \to x_2) + V_{n_0}(S(n_2) \to x_3)$$

in such a way that

$$V_{n_0}(S(n_1) \to x_2) + V_{n_0}(S(n_2) \to x_3) > V_{n_0}(S(n_1) \to x_1) + V_{n_0}(S(n_2) \to x_3).$$

This can clearly be done. It follows that

$$V_{n_0}(S(n_1) \to x_2) > V_{n_0}(S(n_1) \to x_1).$$

The indicative property implies

$$V_{n_0}(S(n_1) \wedge x_2) > V_{n_0}(S(n_1) \wedge x_1).$$

By conditionalization, we also then have that:

$$V_{n_1}(S(n_1) \wedge x_2) > V_{n_1}(S(n_1) \wedge x_1).$$

and so

$$V_{n_1}(x_2) > V_{n_1}(x_1).$$

Thus, initially the agent plans to choose the path leading to x_1. But at node n_1, she wants to switch to x_2, exhibiting dynamic inconsistency. $\square$

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. However, neither of these properties is sufficient by itself to secure this result.

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 directly 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 is also some plausible motivation for reading planning conditionals subjunctively, and the implications of such a reading for the dynamic consistency of desirability maximization have yet to be drawn out.

Secondly, one might take the desideratum of dynamic consistency for granted and investigate further the properties that planning conditionals must satisfy in light of this demand. Clearly, dynamic consistency will require that the desirability of a planning partitioning conditional be an increasing function of the desirabilities of its component consequents conditional upon their antecedents. Might we be able to say more than this to formulate plausible constraints upon planning conditionals that are individually necessary and jointly sufficient for the dynamic consistency of desirability maximization?

Of course, the interest of these projects may be somewhat lessened on the supposition that desirability maximization is not an adequate account of rational choice (which may be so if certain versions of causal decision theory are correct). Nevertheless, desirability still measures an interesting conative attitude that rational agents take toward prospects (following Savage, call it 'news value', if you like), 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 vs evidential debate, as are, we believe, the extensions of this work suggested above.

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.

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.

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

Rothfus, Gerard. "Dynamic Consistency in the Logic of Decision". *Unpublished* (2019). Print.

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