UNIVERSITY OF CALIFORNIA,
IRVINE

The Logic of Planning

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Philosophy

by

Gerard Joseph Rothfus

Dissertation Committee:
Professor Simon Huttegger, Chair
Distinguished Professor Brian Skyrms
Chancellor's Professor Jeffrey Barrett

2020

DEDICATION

To my mom and dad,
who have given me more than I can repay.

"Choice is neither desire only, nor counsel only, but a combination of the two. For just as we say that an animal is composed of soul and body, and that it is neither a mere body, nor a mere soul, but both; so is it with choice."

Nemesius
De Natura Hominis

"It means your future hasn't been written yet. No one's has.
Your future is whatever you make it. So make it a good one."

Doc Brown, anticipating *Autonomy*
Back to the Future Part III

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I must first express my utmost gratitude to my committee chair, Simon Huttegger. His guidance over the course of my graduate work and the writing this dissertation has been invaluable. It was through Simon's lectures that I was first introduced to the theory of rational choice, and he has served as a brilliant teacher and interlocutor ever since. I should note in particular that much of this dissertation's second chapter derives from work we carried out together and is thus particularly indebted to Simon's influence.

My appreciation also extends to my other committee members, Brian Skyrms and Jeffrey Barrett, both of whom it has been a tremendous privilege to learn from. Conversations with Brian early on in my graduate career were particularly formative and helped shape the direction this dissertation has taken.

The various audiences that have heard me present on the content of this dissertation deserve my thanks as well. I particularly think of the participants of the *Foundations of Normative Decision Theory* workshop held at Oxford in the Summer of 2018 and the attendees of the *Probability in Philosophy* conference held at Australian Catholic University in the Fall of 2019. Special mention should also go to the many students enrolled in Brian Skyrms' and Louis Narens' *Social Dynamics* seminar over the past several years at U.C. Irvine. Each of the chapters of this dissertation really began in *Social Dynamics*.

This dissertation would also not be what it is without the contributions of so many treasured friends. Daniel Herrmann has excelled in the role of skeptical interlocutor, pressing me to reform and sharpen my beliefs throughout our friendship. His ideas have broadened my conceptual landscape remarkably and have been a constant source of fruitful reflection for me while writing this dissertation. Aydin Mohseni has likewise been an incredible aid to my academic progress since nearly the start of my graduate work. Our lively conversations about rationality, Bayesianism, the purposes of decision theory, etc. have been hallmarks of our cherished friendship and formative of my thinking on such matters. Further thanks go to Bruce Rushing, whose thoughtful questions and challenges have often inspired me to think more deeply about decision theory in its relation to the classical problems of philosophy, and to Chad Marxen, who has been a dear companion in the project of philosophy and making sense of the world for the last decade. It would be a quixotic task to enumerate all the philosophical compatriots I have been blessed to know and learn from during my time in Irvine. The infeasibility of such a task must not dissuade me from at least thanking Nikhil Addleman, Clara Bradley, Calvin Cochran, Landon Hobbes, Saira Khan, Paige Massey, Calum McNamara, Kyle Morgan, Gabe Orona, Wes Siscoe, Will Stafford, Jeff Swaney, Joshua Thornton, and Johnny Waldrop. I thank God that I have friends such as these.

VITA

Gerard Joseph Rothfus

EDUCATION

Ph.D. in Philosophy **2020**
University of California, Irvine *Irvine, CA*

M.A. in Mathematical Behavioral Science **2019**
University of California, Irvine *Irvine, CA*

B.S. in Mathematics/B.A. in Philosophy **2014**
Pepperdine University *Malibu, CA*

TEACHING EXPERIENCE

Instructor **Summer 2019/Spring 2020**
University of California, Irvine *Irvine, CA*

Instructor **Fall 2019**
California State University, Long Beach *Long Beach, CA*

Teaching Assistant **2014–2020**
University of California, Irvine *Irvine, CA*

PUBLICATIONS

Dynamic Consistency in the Logic of Decision **2020**
Philosophical Studies

Bradley Conditionals and Dynamic Choice (with Simon Huttegger; conditional acceptance) **2020**
Synthese

BOOK REVIEWS

Review of Richard Pettigrew's *Accuracy and the Laws of Credence* **(with Chad Marxen)** **2018**
Philosophy of Science

CONFERENCE PRESENTATIONS

Planning and the Norms of Rational Choice **December 2019**
Philosophy Day, California State, University, Long Beach

Newcombian Tragedy **November 2019**
Probability in Philosophy Conference, Australian Catholic University

Conditionals and Dynamic Choice **May 2019**
Luce 2019 Conference, University of California, Irvine

Evidence, Causality, and Sequential Choice **June 2018**
Foundations of Normative Decision Theory Workshop, Oxford University

Evidence, Causality, and Sequential Choice **May 2018**
46th Annual Meeting of the Society for Exact Philosophy

Evidence, Causality, and Sequential Choice **February 2018**
Topics in Scientific Philosophy Conference in honor of Brian Skyrms, U.C. Irvine

ABSTRACT OF THE DISSERTATION

The Logic of Planning

By

Gerard Joseph Rothfus

Doctor of Philosophy in Philosophy

University of California, Irvine, 2020

Professor Simon Huttegger, Chair

Over the past 40 years, decision theorists have produced an impressive body of literature employing dynamic choice arguments to defend the standard principles of Bayesian decision theory, with Peter Hammond's work serving as the *locus classicus*. However, examination of the import of these arguments has largely been restricted to the context of Savage-style decision theories that posit a sharp distinction between acts, states, and outcomes, while the more general framework developed in Richard Jeffrey's *Logic of Decision* has remained mostly neglected. This is remarkable given both the widely recognized appeal of dynamic choice arguments and the broad popularity that Jeffrey's framework enjoys among philosophical decision theorists. My dissertation aims to remedy this situation by extracting and defending what I take to be the core insights of dynamic choice arguments and exploring their significance in the context of Jeffrey-style decision theories.

Chapter 1

Planning and the Norms of Rational Choice

Dynamic choice arguments have played a significant role in various debates regarding the foundations of normative decision theory. Despite their widely granted intuitive appeal, the merit of these arguments remains a matter of controversy among philosophers. Hence, this opening chapter to my dissertation sketches a philosophical defense of a rationality principle that can undergird many dynamic choice arguments. This principle, *Dynamic Optimality*, posits that a rational agents' own attitudes never preclude her from implementing *ex ante* optimal plans in settings of sequential choice. In the context of Savage-style decision theories, this principle is equivalent to one that excludes rational agents from the possibility of suffering sure loss in finite dynamic decision problems, given a richness condition on the consequence domain. However, it is more versatile in that it can also be applied to decision models like Jeffrey's that dispense with the notion of ultimate consequences and to infinitary settings that render the avoidance of sure loss impossible. This chapter also serves an an introduction to the basic notation and principles of dynamic choice theory that will be employed and amended throughout subsequent chapters.

1.1 Introduction

Diachronic tragedy occurs when an agent implements a sequence of decisions over time that is forseeably to her own disadvantage.[1] Some decision theorists hold that misfortune of this sort is a fate peculiar to the irrational. Agents who inevitably find themselves worse off as a result of their own choices must be guilty of violating one or another of the canons of rationality. This line of thought supplies rational choice theorists with an attractive method for fixing (some of) those canons: figure out which attitudes, update policies, choice rules, etc. lead to diachronic tragedy and rule them out as irrational. This strategy has yielded a remarkable defense of orthodox Bayesian decision theory.

While notable for their theatrical value, arguments from diachronic tragedy form but one family in a cluster of argumentative strategies that exploit the sequential choice setting to bolster Bayesianism. Value of knowledge arguments form another such family, with their key premise positing that rational agents never opt to deliberately avoid the receipt of cost-free information prior to making decisions.[2] Yet another class of arguments within this cluster is framed in terms of dynamic consistency: rational agents are able to form coherent plans and stick to them.[3] No doubt the most rigorous derivation of standard decision theory from dynamic choice norms is found in the work of Peter Hammond, whose argument turns on a formal analogue of the seemingly innocent claim that rational

[1]This term of art is due to [39]. See also [40].
[2][94], [95], [1], [81], [43].
[3][22], [60], [37].

behavior across time is explicable in terms of its (anticipated) consequences.[4]

Many decision theorists maintain that dynamic choice arguments like these constitute the best defense of Bayesianism on offer. Allan Gibbard, for example, has claimed that Peter Hammond's dynamic choice arguments offer "the clearest case that departing from the strictures of classical decision theory is incoherent".[5] And Peter Wakker maintains that these arguments "provide the most convincing foundation for Bayesianism presently available."[6] Even critics of standard rational choice theories have conceded that dynamic choice arguments are among the most promising attempts to defend them. For example, in a paper arguing against the separability assumptions made by expected utility theorists, Johanna Thoma writes "I take the best instrumentalist case that has been made, most notably by Hammond (1988), in favour of separability to consist in an appeal to how agents with non-separable preferences choose in some dynamic choice problems."[7]

Nevertheless, dynamic choice arguments for Bayesianism have been forcefully challenged by a number of recent philosophers. On one hand, the partisans of *time-slice rationality* tell us that diachronic tragedy and other cross temporal ills offer no sign of rational defect in the agents that suffer them since such agents need never act suboptimally at any given time taken in isolation.[8] On the other hand, proponents of *resolute choice theory* contend that, while susceptibility to diachronic tragedy is indeed a mark of irrationality, the separability assumptions needed to derive Bayesianism from this concession are unwarranted.[9] Cutting between these two perspectives, others have argued that there is an

[4][36].

[5][31], p. 32.

[6][96], p. 3.

[7][91].

[8][40] for a clear defense of this line.

[9][60], [57], [29].

internal tension hidden within the assumptions employed by dynamic choice arguments.[10]

The ingenuity of these critics notwithstanding, diachronic tragedy arguments and their kin continue to impress me. This chapter attempts to explain why. In doing so, it will set the stage for subsequent chapters exploring the significance of dynamic choice norms in the context of Jeffrey's decision theory and its causalist variants.[11] This first chapter, however, will restrict attention to Savage's decision-theoretic framework,[12] the setting within which dynamic choice arguments have most often been debated. Doing so will allow us to engage most directly with the dynamic choice literature and to answer objections to dynamic choice norms in the context of the framework within which they were first proposed. The principles defended here will then be transferable (with slight amendation) to the context of Jeffrey-style theories, where they will be seen to yield rich consequences.

My basic strategy will be to argue that an adequate understanding of normative decision theory as a theory of *instrumental rationality* can provide sufficient motivation for each of the principles that undergird diachronic tragedy and related arguments. The charge of internal tension amongst these principles is unfounded, I contend, given the account of decision theory's aims that I sketch below. According to this account, decision theory evaluates an agent's attitudes towards uncertain prospects by judging how effectively the behaviors licensed by these attitudes fare against an assumed standard of basic values regarding outcomes. Allowing that behavior may be understood dynamically, this liberal test can justify conclusions as strong as the central axioms of Savage's decision theory. This is part of the attraction of dynamic choice arguments. Once we appreciate that an agent's execution of cross-temporal plans constitutes a form of behavior relative to which

10 [18], [19], [91].
11 [44].
12 [72].

her instrumental rationality may be assessed, the norms of rational choice emerge from remarkably thin assumptions.

To argue these points, it will be helpful to have a concrete example of a diachronic tragedy argument fixed in mind. We may recruit for this purpose a simple sequential choice argument for the conclusion that rational preferences must conform to the Sure-Thing Principle (STP) in the context of Savage's framework. This argument, which is well known and due in spirit to Howard Raiffa,[13] will suffice to frame our subsequent consideration of the general merits of dynamic choice arguments.

Before we examine this argument, however, both the promised sketch of decision theory's aims and some explication of Savage's framework are in order. §2-§4 of the paper will attend to this. In §5, I lay out the informal diachronic tragedy argument against violations of the STP. In (§6), I precisify the preceding discussion by introducing the formal machinery of dynamic choice theory. This will allow us to analyze more closely the principles that feature in dynamic choice arguments and to motivate a broad sequential choice norm that I dub *Dynamic Optimality*, which will serve as the key rationality postulate employed throughout this dissertation (§7). I then turn to rebutting some of the more prominent objections leveled by philosophers against the principles defended here (§8) and conclude with a brief sketch of things to come in succeeding chapters (§9).

[13] [70]. Raiffa credits an early form of the argument to Robert Schlaifer. An even earlier version of a related argument first appeared in [59].

1.2 Instrumental Rationality

Normative decision theorists aim to characterize *rational* agents. The sort of rationality that captures their concern is practical rather than epistemic: a rational agent is one that performs well in practical tasks. Any measure of an agent's performance in this sense naturally assumes a set of ends that supply the standard by which an agent's successes and failures may be assessed. Whether practical rationality concerns the fixing of such ends or only the selection of means taken in light of them is, of course, a controversial matter. Some philosophers (Humeans) maintain that an agent's adoption of ends is an entirely arational affair and that practical rationality is exclusively a matter of *instrumental rationality*, where an instrumentally rational agent is one that performs well in practical tasks relative to the standard fixed by her own ends (whatever they may be). Others hold to a thicker conception of practical rationality on which an agent can be irrational for adopting the wrong ends. Weighty as these questions are in themselves, they are of little consequence to the decision theorist. The aspect of practical rationality that the decision theorist wishes to explicate is simply its instrumental dimension. The scrutinizing of an agent's ends, along with an assessment of the propriety of doing so, is left to the moral philosopher.

A rational agent, for our purposes then, is an agent that performs well in practical tasks, as judged by the standard of her own ends. To say that an agent is (instrumentally) irrational is to characterize her choice dispositions and/or the attitudes that give rise to them as ineffective at realizing the agent's own ends. For example, consider a simple agent[14] who buys a bet on some proposition p for \$$x$ while simultaneously selling an equivalent bet on p for \$$y$, where $x > y$. This agent is guaranteed a monetary loss of \$$x - y$, regardless

[14] I borrow the term 'simple agent' from David Christensen who employs it to designate a miserly agent whose ultimate values align entirely with her monetary payoffs. See [21].

of the truth value of p. Whether or not we see this as an essential indictment of the involved agent's epistemic rationality,[15] the behavior constitutes a decisive strike against her instrumental rationality.

Speaking more generally, let's take an agent's "ends" to be given by her preferences over a space of ultimate consequences or *outcomes*. (For the simple agent referenced above, this space is her set of possible wealth levels.) We may assume these preferences to at least partially order the space of outcomes.[16] If all decision were decision under certainty, there would be little room (or need) for an interesting theory of instrumental rationality beyond this.[17] We could simply say that an agent should always choose so as to realize a maximal member of her option set, relative to her preferences over outcomes.[18] But, alas, few decisions, if any, are like this. Typically, we are uncertain what outcomes our actions will lead to. This requires us to make judgments regarding the comparative instrumental value of *prospects* or gambles that may lead to various outcomes depending upon what the world is like.[19]

What constraints then do the concerns of instrumental rationality place upon such judgments? In light of the characterization of instrumental rationality offered above, I propose the following test: an agent's judgments of instrumental value are rationally inadmissible

[15][78] and [21] both argue, convincingly to my mind, that dutch books of this sort do typically involve epistemic irrationality as well, though that is not the concern of this dissertation.

[16]While some have suggested that the (partial) ordering assumptions of decision theory may themselves be derived from considerations of instrumental rationality, this is a vexed issue. I find it difficult to make sense of decision theory without some sort of ordering principle assumed from the start. For some discussion, see [3] and [93].

[17]I say 'little room' rather than 'none' because a few interesting issues in the theory of rational choice do present themselves even in the context of decision under certainty, e.g. issues regarding the rationality of incomplete preferences.

[18]This neglects consideration of infinite decision problems in which there may be no maximal element within an agent's option set.

[19]Here I endorse a certain judgmentalism about preference over prospects defended in a moderate form by [11] and a bit more radically by [38]. For a critique, see [92].

if they are structured such that acting upon them would, in certain choice scenarios, guarantee an agent's failure to achieve her own ends, where failure is understood in terms of realizing a non-maximal outcome among those practically available to the agent in the given scenario. Inadmissible attitudes are ones that an instrumentally rational agent (i.e. one whose behavior is well suited for the achievement of her ends) could not act upon.

A test like this requires some clarification, of course. A simple agent may well, acting upon her best judgments, buy a bet on a false proposition p, and thus suffer a monetary loss, without thereby casting any doubt on her instrumental rationality or the coherence of her judgments. While it is true in some sense that this agent did not do as well, by her own lights, as she could have (e.g., by not buying the bet), she could not have known this while deliberating, as p's truth value was antecedently uncertain to her. We might say then that instrumentally rational agents never *foreseeably*[20] fail to do well, relative to their own ends, and judgments of instrumental value are rationally inadmissible if acting upon them precludes an agent's instrumental rationality.

A further clarification regards '*acting upon*'. A perfectly rational agent may in some cases, forseeably, end up worse off than potential counterparts. Such misfortunes may even occur on account of the agent's own attitudes without thereby calling into question the rationality of the agent. For example, an agent may prefer apples to bananas and find herself in a world that punishes such preferences. In this world, perhaps apple lovers are only ever able to access lemons (which are judged by the agent to be among the least desirable fruits), while banana lovers can have their fill of apples and bananas and other finer fruits. The existence of such worlds does nothing to show that a preference for apples over bananas is irrational, even though an agent who found herself in such a world would

[20]That is, from the perspective of their own their own state of information during the course of practical deliberation.

of course have reason to alter her attitudes.[21] While the agent punished for preferring apples over bananas ends up worse off than a counterpart with the opposite preference, she has not *done* any worse than her counterpart in the sense that her misfortune is not the result of her own decisions. Unlike the simple agent caught by a clever bookie, she has not acted so as to make herself worse off. There is no behavior she could have altered so as to improve her standing in the world. If there were (for example, if she could have induced in herself a suitable preference shift), then she would have been, like the simple agent trading Dutch books, recognizable as practically irrational, for then she could have acted so as to perform forseeably better with respect to the ends fixed for her.

An attitude is not then made rationally inadmissible simply by causing misfortune to the agent that holds it. Rather, the proposed test asserts that an attitude is rationally inadmissible just in case it could forseeably lead an agent that holds it into misfortune *by way of the behavior it licenses*.[22] The judgments of an instrumentally rational agent ought never preclude her from attaining her ends as effectively as the choice scenario she finds herself in will allow. Remarkably, this informal test, duly precisified, is capable of grounding some of the most controversial rationality postulates undergirding Bayesian decision theory, if only we recognize that planning, or the reflective execution of decisions across time, is a form of behavior that an instrumentally rational agent may be expected to succeed at. Before we turn to this argument and the accompanying complexities of dynamic choice theory, however, some exposition of its intended conclusion (i.e. Bayesian decision theory as propounded by Savage) is called for.

[21] In extreme cases, a rational agent may even have practical incentive to alter her attitudes in a way that will render her subsequently irrational. Consider a world where agents that fail to give up their rationality are boiled in oil. Relevant here, perhaps, is Parfit's discussion of *rational irrationality*. [64].

[22] The view I take here may be contrasted with the 'success-first' project of [34] who suggests that we might judge the rationality of competing decision theories not just by evaluating the relative instrumental value of the behavior they license, but also by considering directly the instrumental value of adopting the theory itself. This difference in perspective may explain part of the attraction of the non-standard decision theories popular among some AI researchers, e.g. [99], [100].

1.3 Savage's Theory

In the version of decision theory due to Savage, choice problems exhibit a tripartite structure consisting of *acts*, *states*, and *outcomes*. Informally, acts are the objects of choice encountered by an agent acting under conditions of uncertainty, i.e. they are the prospects she must decide between. States are ways the world could be that fall outside the agent's control but that fix what effects a given act will bring about. Outcomes are those hypothetical effects, described to a level of detail sufficient to capture everything the agent cares about in the context of the given decision problem.[23] All three of these features can be captured neatly in a decision table:

	S_1	S_2	...	S_n
A_1	$O(A_1, S_1)$	$O(A_1, S_2)$	...	$O(A_1, S_n)$
A_2	$O(A_2, S_1)$	$O(A_2, S_2)$	...	$O(A_2, S_n)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
A_m	$O(A_m, S_1)$	$O(A_m, S_2)$	...	$O(A_m, S_n)$

Each column of the table corresponds to a state, while each row corresponds to an act. Each cell of the matrix is associated with an outcome, $O(A, S)$, namely, the outcome that would result from choosing act A in the event that state S were to obtain. In the Savage setup, states and outcomes are taken as basic or primitive components of a decision problem while acts are defined derivatively as *functions* from states to outcomes. On this way of viewing acts, what function an act A_i is can easily be read off the decision table: it is the function that maps each state to the outcome falling in the i-th row of that state's column. It is important for the formulation of a decision problem that the set of states and the set of acts each form a partition in the sense that exactly one state must obtain and exactly one act must be chosen. This guarantees that in a properly described decision problem an agent will always effect by her choice a particular row of the corresponding decision table and achieve an outcome identified in a particular cell of that row.

[23] This ensures, for example, that the values of the various outcomes are state- and act-independent.

Formally then, every decision problem starts with an appropriate set $\mathbf{S}$ of states and set Z of outcomes and identifies a subset A of $F = Z^{\mathbf{S}}$ as the set of available acts, i.e. as the agent's choice set. While A will generally be a proper subset of F, Savage still expects rational agents to have preferences regarding any of the acts in F, including those not immediately available. (Savage suggests we can interpret these preferences in terms of hypothetical choices the agent would make from various conceivable choice sets.) For the sake of mathematical simplicity, we will restrict our attention here to *simple acts*, i.e. acts with finite range.[24] Doing so allows us to write acts in a convenient way as follows. If f is a simple act with $\{O_1, O_2, ..., O_n\} \subseteq Z$ as possible outcomes, then letting $\{E_1, E_2, ..., E_n\}$ be a partition of $\mathbf{S}$ into *events* (i.e. sets of states) such that $s \in E_i$ if and only if $f(s) = O_i$ for $1 \leq i \leq n$, we will write f as $\{E_1, O_1; E_2, O_2; ...; E_n, O_n\}$. This is the act that yields outcome O_1 in case of event E_1, O_2 in case of E_2, etc.

In the context of this framework, what, if anything, does instrumental rationality require of an agent's preferences over acts? In the spirit of our previous discussion, let's again suppose that an agent has basic preferences over outcomes given by a (perhaps partial) order on Z that we may denote with $\succeq^*$ and refer to as the agent's *outcome-preferences*. Here $a \succeq^* b$, where $a, b \in Z$, designates that outcome a is ranked at least as high as outcome b in the ordering that specifies the agent's ultimate ends.[25] We can identify an agent's *strict outcome-preference* relation, $\succ^*$, as the asymmetric part of $\succeq^*$, while the symmetric part, $\sim^*$, is her *outcome-indifference* relation. I will take an agent's preferences over acts, denoted by $\succeq$ and referred to as *act-preferences*, to encode an agent's comparative judgments regarding the instrumental value of acts. That is, $f \succeq g$, where $f, g \in F$, designates that an agent judges f to be more choiceworthy than g, i.e. a better means of realizing her

[24]If $\mathbf{S}$ and Z are infinite, then the set of simple acts will of course be a proper subset of F.
[25]That is, a is either strictly preferred to b or the two are indifferent.

basic preferences over outcomes. Again, we can identify an agent's *strict act-preferences*, $\succ$, as the asymmetric part of $\succeq$, while the symmetric part, $\sim$, is her *act-indifference relation*. Given this understanding of act-preference, I take a (partial) ordering condition on $\succeq$ to be conceptually required.[26] Clearly, I must judge that any act is at least as choiceworthy as itself. Further, if I judge that f has at least the instrumental value of g and g of h, then it would be incoherent for me to fail to judge that f has at least the instrumental value of h.[27]

However, for the skeptic who questions this constraint on $\succeq$, given the dynamic choice principles defended below, we may actually partially derive the ordering postulate on act-preferences (in particular, its acyclicity component) from another principle that is evidently tied to instrumental rationality:

> State-Wise Dominance: Let $f, g \in F$. If $f(s) \succeq^* g(s)$ for all states $s \in \mathbf{S}$, then $f \succeq g$. Further, if any of the weak outcome-preferences are also strict, so is the act-preference.

State-wise Dominance is a principle requiring that agents (i) weakly prefer acts that are *guaranteed* to lead to outcomes at least as good as other acts, no matter what state of the world obtains, and (ii) strictly prefer such acts if they also offer some chance of yielding strictly better outcomes. The rationale for this constraint as a principle of instrumental rationality is clear enough. If one act will lead to an outcome at least as good as what another leads to, then one cannot sensibly judge the first to have less instrumental value than the second, nor is this a plausible place to plead incommensurability.

From the standpoint of constructing a theory of instrumental rationality, State-wise Dom-

[26] In other words, $\succeq$ must be reflexive and transitive.

[27] John Broome has argued a similar point, namely that "better than" relations are necessarily transitive. See [17].

inance seems entirely unobjectionable.[28] However, the axiom, even granting ordering, is insufficient for justifying Bayesianism, according to which rational preferences over prospects are representable as expected utility maximizing. Indeed, Savage was forced to employ a much richer array of axiomatic constraints in order to guarantee the existence of a unique finitely-additive probability measure P and a unique (up to positive affine transformation[29]) utility function u relative to which $\geq$ may be viewed as maximizing expected utility. That is, for any simple acts $f = \{E_1, O_1; ...; E_n, O_n\}$ and g,

$$f \geq g \text{ iff } EU(f) \geq EU(g),$$

where:

$$EU(f) = \sum_{i=1}^{n} P(E_i)u(O_i).$$

The probability measure P is defined over the set of all events, $\mathcal{P}(\mathbf{S})$, and intuitively measures an agent's confidence level in each of the events obtaining, while the utility function is defined over the set of outcomes Z and intuitively measures the agent's strength of desire for each outcome. (This utility function will also represent $\geq^*$ provided that act-preferences over constant acts go by outcome-preferences over the corresponding outcomes, a clear requirement of instrumental rationality.) The EU function values acts according to a weighted sum of the utilities of the outcomes they could yield, where the weights used are the probabilities that the given outcomes result from the act.

[28]Though objections do arise if we drop the assumption of act-state independence alluded to in my initial characterization of states. If an agent exerts some degree of control over the state of the world via her choice of act, then the axiom of state-wise dominance leads to well-known absurdities. These issues will be discussed extensively in the context of the causal vs evidential decision theory debate in chapter 4.

[29]A function f is a positive affine transformation of a function g just in case there exist real-valued numbers $a > 0$ and b such that $f = ag + b$.

While a remarkable piece of mathematics, the success of Savage's theorem as an instrumentalist defense of Bayesianism requires that we grant the axioms he imposes on preference the status of instrumental rationality conditions.[30] Some of these conditions, like transitivity, are relatively clear-cut and required by any theory of utility maximization (expectational or otherwise). Others have stirred substantial debate. Most notable among these is Savage's *Sure-Thing Principle*:

> The Sure-Thing Principle (STP): For all $f, g, h, h' \in F$ and $E \subseteq \mathbf{S}$, $f_E^h \geq g_E^h$ iff $f_E^{h'} \geq g_E^{h'}$

Here f_E^h is the act that agrees with the act f on all but the states contained in E, on which it agrees with the act h. Intuitively, the STP requires that, when comparing the instrumental value of acts, agents only be concerned with what the acts yield on states where they differ. This condition is necessary for Savage's system as it is a logical consequence of the hypothesis that an agent's preferences go by expected utility. The mathematics of expectation dictate that if two acts agree on a certain event, then their relative expected utility, measured according to any fixed probability and utility functions, must be invariant under different supposition as to what their restricted agreement consists in. To grasp the normative content encoded in the STP more clearly, an example may be helpful.

Suppose you are considering going for a hike and are debating whether to head to the north side of the mountain (which has better views) or the south side of the mountain (which has easier trails). You are also uncertain as to the weather conditions at the relevant mountain. In particular, it may or may not be raining at the mountain. If you get to the mountain and it starts raining, you will head home. Otherwise, you will set off on your

[30] Or rather, that we grant this status to those of the axioms that are necessary for *EU* representability. The merely structural axioms Savage also employs are not strictly necessary for such representability and are not intended to have normative force.

hike. This decision problem can be captured in the following decision table:

	No Rain	Rain
North	Nice View	Go Home
South	Easy Hike	Go Home

Table 1.1: Taking a Hike.

Deciding to head to the north side of the mountain will yield you a difficult hike with grand views if the weather holds out and a trip back home if it ends up raining around the mountain. Deciding to head to the south side will lead to an easier hike with less grand views if it doesn't rain and a trip back home if it does. The acts North and South yield the same outcome given Rain, so the STP requires that which act is preferred should be fully determined by whether the agent prefers the Nice View outcome to the Easy Hike outcome. If Nice View $\geq$ Easy Hike, then North $\geq$ South. Otherwise, South $\geq$ North. The common outcome received in the event of rain should have no bearing on which act is preferred: if you were to go to the movies in the event of rain instead of going home, assuming this outcome is also the same whether you go North or South, your preference regarding North and South should remain unaltered.

1.4 Allais' Paradox

Focusing on simple cases like this, the STP can appear nigh unobjectionable as a principle of instrumental rationality. In fact, it might even be mistaken for the logically weaker State-wise Dominance principle. But, alas, not all cases are so simple, and the STP does preclude some preferences that seem prima facie rational to many people. This fact was first made apparent by the infamous *Allais Paradox*.[31] A ball is to be drawn from a well mixed urn containing balls numbered one through one hundred. You are offered a choice

[31] [5].

between a ticket that pays $5,000,000 if ball 2-11 is drawn and nothing otherwise, call this option '$L1$', and a ticket that pays $1,000,000 if ball 1-11 is drawn and nothing otherwise, call this option '$L2$'. This decision problem is captured in the following table:

Ticket	1	2-11	12-100
L_1	\$0	\$5M	\$0
L_2	\$1M	\$1M	\$0

Table 1.2: The Allais Paradox, Part 1

Now suppose instead that you are offered a choice between a ticket that pays nothing if ball 1 or ball 90-100 is drawn and $5,000,000 if 2-89 is drawn, call this option '$L3$', and a ticket that pays $1,000,000 regardless of which ball is drawn, call this degenerate lottery '$L4$'. The decision table for this choice is captured in Table 2.

Ticket	1	2-11	12-100
L_3	\$0	\$5M	\$1M
L_4	\$1M	\$1M	\$1M

Table 1.3: The Allais Paradox, Part 2

Many people report having a strict preference for $L1$ over $L2$ in the first problem, while also having a preference for $L4$ over $L3$. Following common custom, let's refer to simultaneous preferences for $L1$ over $L2$ and $L4$ over $L3$ as the *Allais preferences*.[32] The Allais preferences are paradoxical because, while they strike many as reasonable preferences to have, they violate expected utility maximization. One can readily see from the above decision tables that the Allais preferences contradict the STP, and hence these preferences are incapable of being represented as aligning with expected utility.

$L1$ and $L2$ only disagree on states 1-11, which by the STP entails that what happens on the states 12-100 should not impact which act is preferred. But what happens on states

[32]Though less common, preferences for $L2$ over $L1$ and $L3$ over $L4$ will also generate the paradox.

12-100 is clearly relevant for agents with the Allais preferences: when states 12-100 result in \$0 $L1$ is preferred to $L2$, but when those states result in \$1M, as in acts $L3$ and $L4$, $L4$ is preferred to $L3$.

One reason commonly given as to why agents to prefer $L4$ to $L3$ is that the latter has a much lower minimum payoff than the former (\$0 vs. \$1M). This is not the case with $L1$ and $L2$, however, which share a common minimum payoff of \$0. Agents who are sensitive to the minimum values of gambles or who care about the spread of values that a gamble offers may well adopt the Allais preferences on account of these *global* properties of the gambles involved. This has been interpreted by many, including most notably Lara Buchak, as a form of reasonable risk-sensitivity.[33] These considerations may seem to cast significant doubt upon the normative necessity of the STP and hence upon orthodox Bayesianism as proposed by Savage. At the very least, they motivate the search for further arguments that might bolster the STP's claim to characterize instrumental rationality in the Savage model. Why should the judgments encoded in the Allais preferences be viewed as rationally inadmissible?

1.5 Diachronic Tragedy

Consider an agent with the Allais preferences who faces a sequentialized version of the Allais problem. Rather than being offered a direct choice between L_1 and L_2, the agent will be offered a choice between deciding between L_1 and L_2 from her current information state and deciding between L_1 and L_2 after learning whether or not the ball to be drawn is numbered 1 through 11. However, in order to opt to avoid receiving the new evidence prior to making her decision, the agent must pay some small cost, say, \$$\epsilon$. We assume that

[33] [18].

the agent's preferences are robust in the sense that just as $L_1 \succ L_2$, so too $L_1 - \epsilon \succ L_2$. This decision problem is depicted in Figure 1.

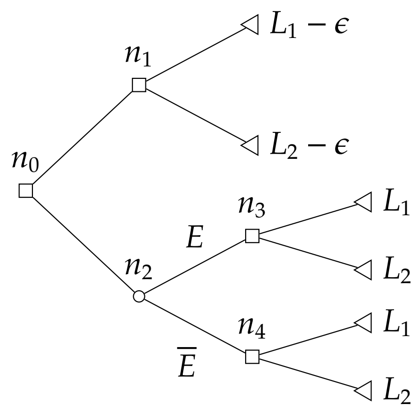

Figure 1.1: A sequential Allais problem.
Here E is the event that the ball drawn is numbered 1-11.

Let us further suppose that the agent's decision making at nodes like n_3 is unaffected by the prospects available to her at nodes like n_4, which are, from the perspective of n_3, strictly counterfactual. In our formal treatment of dynamic choice theory, this will be codified as a separability postulate. Let us then also suppose that the agent would prefer to take L_2 were she to arrive at n_3. If this were not the case, given our separability assumption, we could substitute reference to L_1 and L_2 for reference to L_3 and L_4, respectively, and obtain the result that L_3 would be preferred at the informed nodes (i.e. at n_3 and n_4). The agent's Allais preferences would then enable us to run an exact analogue of the below argument in this alternate decision tree to arrive at the same conclusion. So there is nothing problematic for present purposes with supposing that the agent would prefer to take L_2 at node n_3.

Granting all this, a rational agent facing this problem may look ahead and observe that if she decides to forgo making a more informed decision, she will opt for $L_1 - \epsilon$ since this accords with the Allais preferences. If, however, she decides to take advantage of the offered information, she will end up opting for L_2 if she learns that the ball is numbered 1-11, and will choose something equivalent to L_2 otherwise. Hence, if she opts to make an informed decision, she knows she will effectively wind up with L_2. But this is strictly disprefered by her to $L_1 - \epsilon$, so she will forsake her right to make a better informed decision and opt for $L_1 - \epsilon$. However, note that this plan (namely, not looking at the available information and then deciding on $L_1 - \epsilon$) is sure to leave the agent worse off than another available strategy (namely, looking at the evidence and then opting for L_1) no matter which numbered ball has been drawn.

The agent has succumbed to diachronic tragedy. She has implemented a sequence of choices that is sure to leave her worse off, come what may, than another available sequence of choices would have. All of this is forseeable *a priori* to the agent, just given the preferences of the agent and the structure of the decision tree. It is not difficult to see that there is nothing particularly special about the Allais agent; this same fate may befall any agent whose preferences over Savage acts violate the STP. Violations of the STP entail susceptibility to diachronic tragedy.[34] That is, in carefully constructed sequential choice scenarios, arbitrary violators of the STP are *guaranteed* to do foseeably worse than certain counterparts whose preferences do conform to the STP.

What does this illustration of diachronic tragedy teach us about the rationality of STP violators? The answer suggested by our previous discussion of instrumental rationality is clear: diachronic tragedy reveals a failure of rationality on the part of the one who suffers it. If we apply the test suggested in §2 to assess the rational admissibility of preferences that fail to align with the STP, it appears we must count STP violations as rationally

[34]This claim does assume a richness condition on the consequence domain made explicit below.

inadmissible, no less than violations of the weaker State-wise Dominance principle, for, as the diachronic tragedy argument reveals, acting upon judgments that violate the STP may lead, forseeably, to strictly worse outcomes than what certain ways of acting *against* those judgments would lead to. Proponents of the inference from diachronic tragedy to irrationality are simply drawing upon the conviction that instrumentally rational agents do not, of their own accord and by their own lights, guarantee the frustration of their goals.

Summing up our discussion to this point, I have suggested that we interpret an agent's preferences over prospects or gambles as encoding her judgments regarding their relative instrumental value, i.e. their suitability for realizing her ultimate ends (be they what they may). Such judgments are rationally inadmissible or incoherent, I proposed, just in case acting upon them would, in some cases, forseeably guarantee that an agent would end up worse off than she could have, had she acted differently. A victim of diachronic tragedy is therefore instrumentally irrational since, by acting according to her own best judgments, she finds herself in just such a position, and hence holds to a set of judgments that is rationally inadmissible. In particular then, preferences that violate the STP (including the Allais preferences) are inadmissible. These claims, while hopefully intuitive, are nonetheless largely informal and merit further explication. Dynamic choice theory provides just the setting in which we may carry out this explication and precisfy the proposed connection between instrumental rationality and diachronic tragedy, yielding a principle (*Dynamic Optimality*) appropriate for use in exploring the rational admissibility of various attitudes across decision-theoretic frameworks.

1.6 Dynamic Choice Theory

1.6.1 Basic Framework

While Savage 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*. We have already employed such a tree in our exposition of the sequential Allais problem above, but now we will supply a formal definition. Following Savage, a background set of states, $\mathbf{S}$, and consequences, Z, is assumed in what follows.[35]

Definition 1. *A* **decision tree** *is an eight-tuple* $\langle N, X_1, X_2, X_3, N_+(\cdot), n_0, S(\cdot), \gamma(\cdot)\rangle$*, where:*

i *N is a finite set of nodes, partitioned into X_1, X_2 and X_3.*

ii *X_1 is the set of choice nodes.*

iii *X_2 is the set of natural nodes.*

iv *X_3 is the set of terminal nodes.*

v *$N_+ : N \to \mathcal{P}(N)$ is the immediate successor function which satisfies:*

  (a) $\forall n \in N, n \notin N_+(n)$

  (b) $\forall n \in N, N_+(n) = \emptyset$ *iff* $n \in X_3$

[35] This definition is drawn from [36]. Hammond includes chance nodes as a fourth kind of node that may appear in a decision tree. This allows him to consider the framework of [6] in a dynamic setting. For ease of exposition, I leave out consideration of chance nodes. Aside from this, I largely follow [36] and [22] in my formal exposition of decision trees and plans.

(c) $\forall n, n' \in N, N_+(n) \cap N_+(n') \neq \emptyset$ *iff* $n = n'$

vi n_0 *is the initial node and satisfies* $\forall n \in N, n_0 \notin N_+(n)$

vii $S : N \to \mathcal{P}(\mathbf{S})$ *is a mapping that assigns to each node a set of states such that,* $\forall n \in X_2, \{S(n') | n' \in N_+(n)\}$ *is a partition of* $S(n)$, *and* $\forall n \notin X_2, \forall n' \in N_+(n), S(n') = S(n)$.

viii $\gamma : X_3 \to Z$ *is a consequence mapping that assigns a partial function from* $\mathbf{S}$ *to* Z *to each terminal node such that, for all* $z \in X_3$, $\gamma(z)$ *is defined exactly on* $S(z)$.[36]

Each of these components should become clear through example. Intuitively, a decision tree consists of a set of nodes connected to each other and ordered in a way specified by the immediate successor function, $N_+(\cdot)$. 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 is associated, via $S(\cdot)$, with a set of states of the world, namely, those deemed possible at that node. Similarly, each terminal node is associated, via γ, with a state-contingent consequence function, namely, the one realized by the agent if she reaches that node. An example of a relatively simple decision tree is depicted below.

The agent facing this tree, $T1$, 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 consequence function $\gamma(z_1)$ will be realized. If the agent chooses to move down, she will first learn

[36]Following Hammond, given the background sets $\mathbf{S}$ and Z, I will assume an unrestricted domain of decision trees.

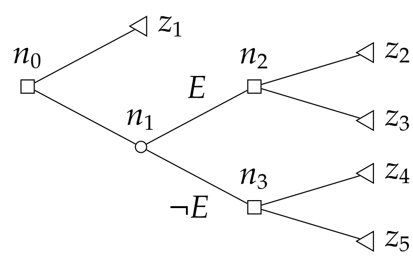

Figure 1.2: $T1$, a decision tree.

which of the events $S(n_2)$ and $S(n_3)$ obtains and then, depending on what she learns, face either a choice between $\gamma(z_2)$ or $\gamma(z_3)$ or a choice between $\gamma(z_4)$ or $\gamma(z_5)$.

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. It thus traces a unique path through the tree, given any combination of moves by nature at its nodes. So, for example, in T1, 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.

It is convenient to follow Cubbit in formally identifying a plan p in a decision tree T with the set of terminal nodes in T that could be arrived at through implementation of p. On this way of viewing plans, the five possible plans available in $T1$ are (i) $\{z_1\}$, (ii) $\{z_2, z_4\}$, (iii) $\{z_3, z_5\}$, (iv) $\{z_2, z_5\}$, and (v) $\{z_3, z_4\}$. To formally codify this definition of a plan, we need to introduce a bit more terminology. The formalism I'll drag out here is a bit cumbersome, but in the end it will allow us to think about planning with greater precision and clarity.

First, let a *branch* in a tree T be a sequence $\langle n_0, \ldots, n_k \rangle$, where n_0 is the initial node in T, $n_k \in X_3$, and $n_{i+1} \in N_+(n_i)$, for $0 \leq i < k$. Intuitively, a branch is simply a path through a tree defined by a complete specification of moves for both the agent and nature. We'll say that a member of a branch *succeeds* those nodes that come before it in the sequence and *precedes* those that come later. It will be useful to be able to refer to the set of nodes that precede a given node n in any branch that contains it, so denote such a set by $R(n)$. Now, if we let $L(n)$ be the set of all nodes contained in branches that include n, we are able to state simply the following definition: For any node n in a tree T, let $FT(n) = \{n' | n' \in X_3 \cap L(n)\}$ be the set of *feasible terminations* after node n. In other words, $FT(n)$ is the set of terminal nodes reachable from n by some combination or other of moves on the part of nature and the agent. (For example, at the initial node n_0 of any tree T, $FT(n_0) = X_3$.) We will say that two branches *diverge* at a node n if n is their last member.

Next, we need to say what it is for two nodes in a tree to be *connected*. Let n and n' be nodes in a decision tree T. Then n and n' are connected iff (i) there is a branch b that contains n and a branch b' that contains n' such that b and b' diverge at some natural node n_y that precedes both n and n' and (ii) $R(n) - R(n_y)$ and $R(n') - R(n_y)$ are subsets of X_2, i.e. contain only natural nodes. Intuitively, two nodes are connected just in case any sequence of moves by an agent to get to one of them leaves open the possibility that the agent will wind up at the other, thanks to nature.

Finally, we can formally define a *plan* at a node n in a decision tree T:

Definition 2. *Given a node n in a tree T, a subset p of $FT(n)$ is a* **plan** *at node n just in case: (i) if $z, z' \in X_3$ are connected, then, if $z \in p$, $z' \in p$ too, (ii) if z^* is a terminal node connected to a* *choice node n', then $z^* \in p$ iff $FT(n') \cap p \neq \emptyset$, and (iii) if n'' and n''' are connected choice nodes that succeed n, then $FT(n'') \cap p$ is non-empty iff $FT(n''') \cap p$ is nonempty.*[37]

Where n is a node in a decision tree T, we let $\Omega(T, n)$ stand for the set of all plans at n.

Also important for sequential choice is the notion of a *plan continuation*. Fix a tree T and let n_a and n_b be nodes of T such that n_a precedes n_b. If p is a plan at n_a, then the continuation of p at n_b is defined as $p(n_b) = p \cap FT(n_b)$. If $p(n_b)$ is non-empty (i.e., intuitively, if p makes arrival at n_b possible), then $p(n_b)$ will itself, of course, be a plan at n_b. Informally, a plan continuation at a node n of a plan p then is simply that part of p that can still be implemented at n. We can also define $\Omega(T, n_a)(n_b)$ to be the set of plan continuations at n_b of plans at n_a.[38]

1.6.2 Rational Planning

With a precise framework for modelling sequential decision making now in place, we can consider more carefully what rationality requires of agents in such settings. I will assume that just as an agent can make judgments regarding the comparative instrumental value of Savage acts, so too can she make such judgments regarding possible plans available to her in hypothetical decision trees. Following McClennen, I will suppose that such judgments are captured by an agent's *admissibility function*, denoted by D, which takes as input a set of plans available at a decision tree T and a node n within that tree and returns the subset of such plans that the agent judges to be maximally effective at realizing her ultimate ends.

[37]Note the difference between plans as defined here and the *strategies* studied in dynamic game theory. A strategy specifies a move for every information set in a given game tree, regardless of whether the information set could be reached via the agent following the given strategy. A plan, by contrast, only specifies moves at those choice nodes that an agent could actually reach, given implementation of the plan.

[38]This set will be empty if $FT(n_a) \cap FT(n_b) = \emptyset$.

Note that any plan in an arbitrary decision tree can be associated with a partial function from states to outcomes, sometimes known as a *conditional Savage act*.[39] If $p \in \Omega(T, n)$, we may define the conditional act associated with p by: $p^*(s) =_{df} \gamma(z)(s)$ if $s \in S(z)$, for all members z of p, where γ is T's consequence mapping.[40] Informally put, the conditional Savage act associated with a plan is simply the state-contingent consequence function induced by executing that plan. I will sometimes write 'f_E' to denote the conditional act defined as the restriction of unconditional act f to event E. We may further denote the set of conditional acts associated with plans available at a node n in a tree T by $\Omega^*(T, n)$. If an agent takes attitudes toward both conditional Savage acts and hypothetical plans, it is entirely reasonable to assume, in the current framework, that there be an alignment betwixt these two sets of judgments. That is, if we assume that, for any $E \subseteq \mathbf{S}$, an agent's preferences over the conditional acts defined on E are given by a (partial) order denoted by $\succeq_E$ (here the case of $E = \mathbf{S}$ yields her familiar preference ordering over Savage acts, which I will continue to write as $\succeq$) and if we employ $M(\Omega^*(T, n))$ to denote the set of maximal elements of $\Omega^*(T, n)$ relative to $\succeq_{S(n)}$, we may further assume:

> **Plan Reduction**: For all trees T and nodes n in T, $p \in D(\Omega(T, n))$ iff $p^* \in M(\Omega^*(T, n))$.

Plan Reduction requires that the instrumental value of plans be assessed according to the value of their corresponding conditional acts. One might be tempted to think that, given a set of preferences over conditional acts, Plan Reduction then provides an adequate blueprint for approaching problems of sequential choice. When faced with a problem modelled by a decision tree T, so this thought goes, an agent may simply select a plan from amongst the members of $\Omega(T, n_0)$ that she judges most favorable with respect to the aim of realizing her ultimate ends, such judgments being reached by consideration of the

[39] Such conditional acts are the object of inquiry in [56].

[40] Since the events associated with any non-identical terminal nodes in the same tree will necessarily be disjoint, this will indeed define a (possibly partial) function on $\mathbf{S}$.

instrumental value of the members of $\Omega^*(T, n)$, and then proceed to implement the first stage of the chosen plan.

However, this approach to dynamic choice has appeared short-sighted to many decision theorists. It seems to proceed as if an agents' objects of choice in a sequential decision making scenario were entire plans. But, except in trivial decision trees, this is not the case. All an agent can immediately do, depending on her location in the decision problem, is implement a presently accessible stage of a logically available plan. But she cannot guarantee that she will follow through on implementing that plan in future choice nodes of the tree. A *sophisticated* agent will therefore anticipate how she would choose at the final choice nodes she could reach in the decision tree at hand, and then fix those choices for purposes of her current deliberation among plans, effectively paring down the decision tree by eliminating paths she knows she would not opt to take in the future. She can then work backwards through the tree, repeating this process of pruning the tree at all future choice nodes until she reaches her current position. Of the plans remaining after this process is completed, rational agents can then opt to implement a plan that is most favored.

In other words, sophisticated choice proponents propose that being associated with a maximally favorable prospect is not sufficient for a rational agent to start implementing a plan. In addition, the plan must be $feasible$.

Definition 3. *A plan p at a node n in a tree T is (dynamically)* **feasible** *at n in T just in case, for any choice node, n_a that succeeds n in T, $p(n_a) \in D(\Omega(T_D, n_a))$, if $p(n_a)$ is non-empty, where T_D is the tree obtained by pairing back T according to backward induction via D. Let the set of all dynamically feasible plans at a node n in a tree T, relative to admissibility function D be denoted by $DF_D(T, n)$.*

Note that this notion of feasibility accounts only for internally sanctioned preference shifts on the part of the deliberating agent and not externally induced shifts like the one famously suffered by Ulysses.[41] Rational agents, like Ulysses, will of course take potential shifts of the latter sort into account when making judgments about the feasibility of logically available plans. But we may safely set such cases aside in the context of this essay's argument. We are here interested in assessing the instrumental rationality of agents that adopt various sets of action-guiding attitudes by examining what behaviors those attitudes license. For this purpose, it is appropriate to assume that the attitudes in question evolve only in ways endorsed by those very attitudes and not by any Siren's song.

Assuming that a rational agent's attitudes always license adopting a plan that is maximal among those that are feasible and that an agent will act accordingly, we may infer an agent's behavior from her judgments of admissibility. Let a *behavior norm* β specify, for any decision tree, the set of plans that the agent might implement in that tree.[42] In general then, $\beta(T) \subseteq \Omega(T, n_0)$. But in the special case of a rational agent who acts according to her admissibility function D in a sophisticated fashion, we have:

Sophistication: For all trees T, $\beta(T) = \{p \in \Omega(T, n_0) | p \in D(DF_D(T, n_0))\}$.

Sophistication dictates that a rational agent's behavior be constrained by her judgments of the admissibility of available plans, albeit in a slightly more roundabout way than that dictated by the naive procedure originally proposed. Granting that her judgments of the admissibility of plans are in turn informed by her preferences over conditional acts, as required by Plan Reduction, Sophistication provides a reasonable translation of these latter attitudes into behavior across the full range of hypothetical choice scenarios representable

[41] Relevant here also is the infamous addict of [35].

[42] So understood, the notion of a behavior norm employed here is slightly different than the familiar notion employed by [36].

as Bayesian decision trees.

One point deserves special note here. There are multiple ways one might understand the backward induction procedure applied to a decision tree T via admissibility function D to arrive at T_D. The ambiguity arises due to the possibility of non-initial choice nodes in T involving selection amongst indifferents (or incommensurables). If a future choice node requires selection amongst such options how should the tree be paired down to arrive at T_D? One thought would be to leave such options in place and, when assessing the admissibility of options further upstream, treat such nodes as chance nodes that may result in any of the indifferent possibilities being chosen. This proposal effectively turns the dynamic choice problem into an extensive-form game in which an agent's possible future selves are entirely distinct players whom she is unable to influence in any way. Such an approach leaves no room for the sort of intentional, cross-temporal planning that is a hallmark of human agency.

To allow for a more robust conception of rational planning, we may instead endorse the 'splitting procedure' defended by Wlodek Rabinowicz.[43] On this proposal, the pairing of the tree will always be applied to plans taken as wholes. So, for example, if at some choice node an agent finds herself uncommitted between two options both will be left intact until the next stage of the induction. At that stage, an agent will consider the value of the plans that include the indifferent or incommensurate options as continuations. If either plan appears non-maximal at this stage, its whole continuation will be cut out from the tree. This process is repeated until the agent reaches the initial node. Any plan that survives this pruning of the decision tree will be one that an agent would be willing to follow through on, given her anticipated future values. As Rabinowicz notes, this method assumes that

[43]See [68] and [69]. A similar approach is also applied in game theory by [9] to refine the Subgame Perfect Equilibrium concept. Asheim

an "agent will follow a plan as long as he lacks a positive reason for deviation."[44]

Unfortunately, endorsement of this approach does put me at odds with some advocates of sophisticated choice who seem to think the sort of cross-temporal agency implied by this method of determining feasibility is impossible.[45] In this, I think these theorists go too far. Sophistication rightly insists that a rational agent's decision-making is at all times forward-looking and consonant with her values at the time of choice, but this is still consistent with past intentions playing an action-guiding role in the agent's behavior, e.g. by acting as a tie-breaking mechanism in cases in which an agent must choose from a choice set that lacks a unique maximal member. Indeed, adherence to the version of sophisticated choice endorsed here (i.e. one that allows for the ability to execute plans and act on past intentions) is essential to guaranteeing the avoidance of diachronic tragedy if we wish to allow for the rational possibility of incomplete outcome- and act-preferences.[46]

This has put us in a position to understand the phenomenon of diachronic tragedy more formally. Letting $A = \{\succeq_E |E \subseteq \mathbf{S}\}$ be an agent's set of preference relations over conditional acts, let β_A be the behavior norm induced by A in line with Plan Reduction and Sophistication. We can now define both a strong and a weak sense of diachronic tragedy:

Definition 4. *An agent's behavior norm β is* **strongly tragic** *just in case there exists a decision tree T such that, for all $p \in \Omega(T, n_0)$, $p \in \beta(T)$ only if there exists $q \in \Omega(T, n_0)$ such that $q(s) \succ^* p(s)$*

[44] See [69], p. 287.

[45] For example, [88] fn. 20 dismisses the idea that past intentions can play any role in fixing future choice amongst indifferents and contends that the contrary view amounts to the resolute choice theory I criticize below.

[46] Consider, for example, an agent who lacks a preference between f and g, while strictly preferring f to some state-wise dominated plan $f - \epsilon$, regarding which the agent also lacks a preference vis-a-vis g. Imagine that the agent is offered a choice between f and a subsequent choice of g and $f - \epsilon$. If the past preferences of the agent, codified in her future-directed intentions lacked the bearing on her future choices posited by Rabinowicz's version of sophisticated choice, then it would be impossible to preclude an agent from implementing the state-wise dominated plan associated with $f - \epsilon$ while allowing her to realize the maximal plan associated with g. For a detailed defense of the significance of intention and temporally extended agency for practical rationality along distinct though related lines, see [14] and [15].

for all $s \in S(n_0)$. An agent's judgments A are **strongly tragic** *just in case β_A is strongly tragic.*

Definition 5. *An agent's behavior norm β is* **weakly tragic** *just in case there exists a decision tree T such that, for some $p \in \Omega(T, n_0)$, $p \in \beta(T)$ and there exists $q \in \Omega(T, n_0)$ such that $q(s) >^* p(s)$ for all $s \in S(n_0)$. An agent's judgments A are* **weakly tragic** *just in case β_A is weakly tragic.*

Clearly, strongly tragic behavior norms (attitudes) will also be weakly tragic, but not vice-versa. In the strong sense, a set of judgements regarding the instrumental value of (conditional) Savage acts leads to diachronic tragedy just in case there is a possible decision scenario in which *every* plan that an agent acting on those judgments could implement in that scenario is associated with a prospect that is state-wise dominated by a prospect associated with another plan available in that scenario. In the weak sense, a set of judgements regarding the instrumental value of (conditional) Savage acts leads to diachronic tragedy just in case there is a possible decision scenario in which *some* plan that an agent acting on those judgments could implement in that scenario is associated with a prospect that is state-wise dominated by a prospect associated with another plan available in that scenario.

I claim, along with other defenders of diachronic tragedy arguments,[47] that a rational agent cannot adopt even weakly tragic attitudes. That is, I accept what we might call:

> **Sequential Dominance**: A rational agent's behavior norm is not weakly tragic.[48]

[47] [16] is one such notable defender: "[I]f it is knowable *a priori* that strategy *a* yields a better result than strategy *b*, then it is pragmatically irrational to choose strategy *b* when strategy *a* is available."

[48] Some defenders of diachronic tragedy arguments, notably [13], would insist that Sequential Dominance is too strong a rationality postlate and ought to be replaced with a principle forbidding merely strongly tragic attitudes. My sense is that the hesitancy these authors feel toward embracing the stronger version of Sequential Dominance is motivated by a commitment to a strong version of sophisticated choice on which past preferences can play no role whatsoever (not even a tie breaking one) in settling future behavior. If we accept such a strongly disunified view of sequential decision making, then avoidance of weakly tragic attitudes requires that one's preferences satisfy the completeness axiom, which may seem like an implausible condition of instrumental rationality. Since these authors wish to maintain the rationality of incomplete

This follows from the account of instrumental rationality endorsed here. An instrumentally rational agent is one whose behavior is well-suited to the attainment of her ends, which are taken as given. But an agent whose behavior is even weakly tragic may in some situations frustrate her own ends by implementing a course of action that is strictly dominated by another course of action that she could have taken. She may in some cases forseeably end up with a worse outcome than she would have realized had she conformed to a different behavior norm. If we further accept the plausible connections between rational behavior and judgment posited by Plan Reduction and Sophistication, we may also conclude that a rational agent will never exhibit diachronically tragic attitudes either.

As suggested by the discussion in §5, the Allais preferences are diachronically tragic[49] even in the strong sense and, hence, in my view rationally inadmissible. There is a sense in which this is true of any preferences that violate the STP.[50] I relegate a careful discussion of this claim to an appendix. At present, I want to turn to another principle of rational planning that I find compelling in the current setting and connect it with the avoidance of diachronic tragedy.

preferences, they are forced to deny that weakly tragic attitudes are irrational. However, endorsement of a modestly more robust account of temporally-extended agency like the one codified in the version of Sophistication I favor, eliminates this dilemma. We may then happily embrace both Sequential Dominance and the rational permissibility of incomplete preferences.

[49] This claim of course assumes the standard caveat that I have correctly described the outcome space in the Allais problem. If an agent cares about non-monetary factors, e.g. the possibility of experiencing regret or even strict counterfacts about what would happen conditional on various non-actual states, then her apparent Allais preferences may not really violate the STP at all and lead to know tragedy at all. See the discussions in [17], [18], [66], and [12]. Such a liberal carving up of outcomes, while to my mind entirely appropriate given our characterization of decision theory as simply a theory of instrumental rationality, does present difficulties in the context of Savage's theory given his other postulates, especially his rectangular field assumption. [12] suggest this tension as a reason to prefer the Jeffrey framework to be discussed in subsequent chapters.

[50] Technically, if we want to continue to allow for the rationality of incomplete preference rankings, exclusion of diachronic tragedy only allows us to derive a restricted version of the STP that assumes comparability of the relevant prospects. The argument also makes use of a technical assumption regarding the richness of the agent's act-preferences.

1.7 Dynamic Optimality

Katie Steele suggests that diachronic tragedy arguments that lean upon a principle like Sequential Dominance are more plausible than other dynamic choice arguments on the market, most notably Hammond's, which relies upon a sort of normal-extensive form reduction principle he dubs 'Consequentialism.' According to Steele, "Idealised NEC [Normal-Extensive Form Coincidence] per se is not a compelling constraint on decision rules. Rather, it is the possibility of suffering a sure loss due to one's own decision rule that, arguably, demonstrates an inconsistency, or at least a deficiency, in that rule."[51] Though I share some of Steele's hesitancy regarding Hammond's Consequentialism,[52] a close cousin of Hammond's principle is, I believe, every bit as plausible as Sequential Dominance as a principle of instrumental rationality and, in fact, actually manages to avoid some of its pitfalls.

The principle I have in mind requires that a rational agent's judgments never be structured such that acting upon them would preclude implementing plans that the agent judges to be ex ante optimal. That is, for a rational agent, any admissible plan should be a feasible one. I call this principle *Dynamic Optimality*. This principle is, I believe, an improvement over Sequential Dominance, which, while very plausible under the assumptions of our discussion so far, has several drawbacks that limit the scope of its status as a rationality principle. For example, Sequential Dominance takes the form of a sequentialized version of the State-wise Dominance principle, and as such requires for its statement an assumed domain of ultimate consequences. This is in keeping with the Savage framework but more awkward in frameworks like Jeffrey's that don't require the existence of such an

[51] [88], p. 471.

[52] Hammond's Consequentialism entails the completeness of weak preference, for example, which is arguably implausible as a requirement of instrumental rationality. The Dynamic Optimality principle I go on to defend does not have this consequence.

ultimate consequence domain. Further, Sequential Dominance is, like static dominance principles, unsustainable in infinite decision problems. In some such problems, there are no undominated options and thus, unless we want to rule out all behavior in such scenarios as irrational, we cannot insist upon an infinite version of Sequential Dominance.[53]

These problems do not, however, pose any problem for Dynamic Optimality, which requires for its application to decision problems neither the posit of ultimate outcomes nor the availability of optimal options. We can state this principle formally as:

Dynamic Optimality: For all trees T and nodes n in T, $D(\Omega(T,n)) \subseteq DF_D(T,n)$.

This principle enjoys a fairly evident rationale for qualifying as a norm of instrumental rationality. If an agent's attitudes violate Dynamic Optimality, then there will sometimes be logically available plans which she judges to have maximal instrumental value and yet which she is unable to rationally implement because of her own attitudes. Moreover, it is worth noting that in the context we have been assuming (namely, Savage's framework) satisfaction of Dynamic Optimality in the domain of finite decision trees is equivalent to satisfaction of Sequential Dominance, given a richness assumption regarding preferences. Hence, while more amenable to generalization than Sequential Dominance, Dynamic Optimality should be viewed as no less plausible than Sequential Dominance.

To prove this claim, we first define the needed richness condition.

Definition 6. *A set of attitudes A on F is* **rich** *just in case for any $\succeq_E \in A$, if $f, g \in F$ are such that $f_E \succ_E g_E$, then there exists $f' \in F$ such that f state-wise dominates f' and $f'_E \succ_E g_E$.*

Intuitively, attitudes are rich just in case whenever they rank one act as strictly better than

[53]For an interesting discussion of infinite dynamic choice problems that lack optimal solutions, see [8].

another there is some small cost that they would license one to pay to receive the better act in place of the worse. This condition is likely met to a high degree by the preference structures of most human agents. Can't one typically sour a range of outcomes just slightly enough (e.g. by subtracting a few cents of income, a few seconds of time, etc.) to retain any antecedent strict preference for the unsoured act over another? Another route we could take here would be to introduce lotteries as possible outcomes. We could plausibly then achieve the desired souring by varying ever so slightly the probability with which a lottery yields a better vs. a worse outcome. Be that as it may, however, the role played by the richness assumption in connecting Dynamic Optimality to Sequential Dominance does not require us to imagine that it is generally satisfied. Certainly, richness is no strike against the rational permissibility of a set of judgments. If Dynamic Optimality is co-extensive with avoidance of diachronic tragedy in the special case of agents holding rich attitudes, a more general connection between Dynamic Optimality and instrumental rationality may be suggested.

We can now establish:

Proposition 1. *If A is a rich set of attitudes satisfying State-wise Dominance, and D the plan admissibility function generated from them by Plan Reduction, then $D(\Omega(T,n)) \subseteq DF_D(T,n)$ iff A is not weakly tragic.*

Proof. Let A be a rich set of attitudes over a domain of acts F, and let D be the plan admissibility function generated from A.

($\rightarrow$) We first argue that $D(\Omega(T,n)) \subseteq DF_D(T,n)$ entails that A is not weakly tragic. So, suppose $D(\Omega(T,n)) \subseteq DF_D(T,n)$. We need to show that β_A is not weakly diachronically tragic. That is, letting T be an arbitrary decision tree, we need to show that there does not exist $p \in \beta_A(T)$ such that for some $q \in \Omega(T, n)$, q^* state-wise dominates p^*. Consider an arbitrary $p \in \beta_A(T)$. We know, by definition then, that $p \in D(DF_D(T, n))$. Hence, since $D(\Omega(T, n)) \subseteq DF_D(T, n)$, Plan Reduction requires that $p \in D(\Omega(T, n))$. But then p^* cannot be not state-wise dominated by any $q^* \in \Omega^*(T, n)$, since, granting Plan Reduction and State-wise Dominance, p could not then be optimal. Thus, β_A is not weakly tragic and, be extension, neither is A.

($\leftarrow$) We now argue to $D(\Omega(T, n)) \subseteq DF_D(T, n)$ from the assumption that A is not weakly tragic. Suppose A is not weakly tragic. Let $p \in D(\Omega(T, n))$. We want to show that $p \in DF_D(T, n)$. Suppose that $p \notin DF_D(T, n)$. Consider the tree T^* that adds $p^* - \epsilon$ as an initial option, where $p^* - \epsilon$ is a conditional Savage act defined on $S(n)$ to which p^* is strictly preferred but which is strictly preferred to all other members of $\Omega^*(T, n)$ to which p^* is strictly preferred. (By the richness assumption we know such a prospect exists.) If there are no members of $DF^*_D(T, n)$ to which p^* is indifferent, then we have a diachronic tragedy in T^*. Suppose there were members of this set to which p^* is indifferent. By the richness assumption, we can then slightly sour these prospects on the states on which they disagree with p^*, which will make them strictly less preferred than p^*, by State-wise Dominance. Now alter the choice of $p^* - \epsilon$ if necessary to ensure that is also strictly preferred to these plans, and we will have, by Plan Reduction, that $p - \epsilon \in D(DF_D(T*, n_0))$, and hence we again have diachronic tragedy in T^*. □

In the context of the Savage framework, with our attention restricted to finite decision trees, there is a tight connection between never implementing state-wise dominated plans and satisfying the Dynamic Optimality postulate, granting the version of Sophistication suggested above. Hence, to the extent that we find the avoidance of diachronic tragedy a plausible requirement of instrumental rationality, we have good grounds, in the context of the present framework, to think the same regarding Dynamic Optimality. Moreover, I think the principle enjoys plausibility in its own right even apart from such connections. If an agent's own attitudes preclude her from implementing a plan that she judges to be optimal with respect to promoting her ends, then the agent's instrumental rationality is thereby called into question. The resultant possibility of diachronic tragedy is a dramatic way of drawing out this concern. Or so it seems to me, but there remain objections to attend to.

1.8 Objections Defused

Thus far, I have attempted to distill the common assumptions of standard dynamic choice arguments into several plausible rationality postulates: principally, Plan Reduction, Sophistication, and Dynamic Optimality. From the perspective of mainstream decision theory, Plan Reduction and Sophistication are hard to resist. This leaves Dynamic Optimality (or, equivalently, Sequential Dominance) as the arguments' weakest point. Objection to Dynamic Optimality[54] is in fact so common among the critics of dynamic choice arguments that I will refer to it as 'the Standard Objection'. Outside the bounds of orthodox decision theory, a more radical school has granted Dynamic Optimality while objecting to the Sophistication principle, offering in its place an alternative model for translating judgments of plan admissibility into behavior: the theory of resolute choice. Still other critics of dynamic choice arguments are less certain which of the principles proposed here goes awry, merely contending that the principles' respective motivations are at loggerheads and, hence, that they ought not all be endorsed at once. Each of these worries merits a due response.

[54]Typically phrased as an objection to the reduction of extensive form decision problems to their normal form counterparts.

1.8.1 The Standard Objection

The Standard Objection to principles like Sequential Dominance and Dynamic Optimality has been raised by a diverse chorus of voices and maintains that the appropriate objects of rational criticism are individual choices and not cross-temporal plans or strategies or sequences of actions. Brian Hedden has raised the Standard Objection forcefully in the context of his defense of a *time-slice* approach to rationality.[55] However, I take the basic objection, while expressed quite differently, to be shared by a number of prominent decision theorists, including Teddy Seidenfeld, Isaac Levi, and James Joyce.[56] These critics allege that the inference from diachronic tragedy and similar ills to irrationality is invalid because it rests upon a confusion regarding what an agent's options typically are in decision problems and, consequently, involves misidentifying the sorts of behavior that an agent can be rationally criticized for.

Consider, for example, the diachronic tragedy that befell the Allais agent in §5. While it is true that an agent with the Allais preferences will implement a plan that is strictly dominated by another logically available plan, she nonetheless never chooses an option that is so dominated by another. At the first stage of the problem, she opts to directly take $L_1 - \epsilon$. But this selection is not, from this agent's standpoint, dominated by any of the other acts she could have taken at that choice point. The choice of $L_1 - \epsilon$ is clearly not dominated by a choice of $L_2 - \epsilon$, and the sophisticated agent will realize that a choice to delay the decision is tantamount to receiving L_2, which again does not dominate $L_1 - \epsilon$. Hence, the sophisticated Allais agent's choice of $L_1 - \epsilon$ is quite appropriate, given her judgments of instrumental value and the structure of the problem she finds herself in. The plan that dominates the realized prospect (namely, delaying and then choosing L_1 come what may)

[55]See [40], [39]).

[56]See especially [74], [75], [51], and [48].

is infeasible for the agent to implement at the initial node. It is, given her values, not an option for her. Why then should her failure to implement this infeasible plan count against her rationality?

There is much to agree with in this statement of the Standard Objection. Plans are indeed, in general, not options among which an agent can simply pick at will. Typically, she can only decide to implement a particular stage of a plan, i.e. whatever options are directly available to her at the choice point she currently occupies (and perhaps form defeasible intentions regarding her future decisions). Hence, the objector is entirely correct that the sophisticated Allais agent who falls prey to diachronic tragedy never chooses poorly, relative to her stated values. However, these facts are insufficient to absolve the tragedy prone agent of the charge of irrationality. If some external factor (e.g. a Siren's song) had rendered the dominant plan infeasible for the agent, it would indeed be inappropriate to make any negative inferences regarding the agent's instrumental rationality. But, as Katie Steele notes, such is not the case in this problem:

> "The agent is not forced into sure loss by any external circumstances; they rather bring this outcome upon them self, due to their own decision-making plans. Surely, this is a significant mark against any such decision plans."[88]

The choice to take $L_1 - \epsilon$ at the outset may be optimal if we hold fixed the agent's choice at the hypothetical second stage of the problem, but when assessing the rational admissibility of a set of attitudes by considering the behavioral implications of those attitudes (as we are) this is not an appropriate supposition to hold fixed. On the proposal sketched above, assessing whether a given set of judgments regarding the instrumental value of prospects is rationally admissible in light of an agent's basic values involves considering whether those judgments are structured so as to license behavior that is ineffective at furthering the agent's basic values. Diachronic tragedy arguments make clear that STP violating preferences do not pass this test.

Note that this proposal neither rejects a synchronic understanding of options nor disavows the legitimacy of sophsticated choice as an approach to sequential decision making. To the contrary, sophistication (albeit, of an appropriate sort) is a key assumption throughout this paper. Hence, I do not deny that, e.g., the dominant option in the sequential Allais problem is infeasible for an agent with the Allais preferences. But this is precisely the problem. An instrumentally rational agent's attitudes should not be structured such that those very attitudes render optimal plans infeasible to implement.[57]

1.8.2 The Resolute Choice Objection

I have claimed that an instrumentally rational agent is, at a minimum, one whose behavior never forseeably frustrates her own ends. This idea is codified in the Sequential Dominance principle. Following the lead of Edward McClennen, some decision theorists have granted this much, therefore accepting the inference from diachronic tragedy to irrationality, while at the same time jettisoning the Sophistication principle. According to these proponents of *resolute choice*, a rational agent's behavior at a particular stage of a sequential choice problem depends not simply upon the prospects that lie open to her at that stage of the problem, but rather also upon the broader context fixed by the larger sequential choice problem with respect to which the current stage and its possible suc-

[57]Note also then that the style of dynamic choice argument defended here is entirely compatible with the synchronicity thesis of time-slice rationality according to which the rationality of an agent at a time supervenes upon the attitudes she holds at that time. (See [40].) According to diachronic tragedy arguments, the propensity of a set of attitudes to lead to diachronic tragedy reveals a flaw in those attitudes, but this propensity is a feature of an agent's attitudes at a time and hence her rationality at a time may well supervene upon the attitudes she holds at that time.

ceeding stages form a subproblem.[58]

I have already (controversially) granted one way in which an agent's behavior in a subtree may rationally be affected by the larger tree in which it appears. On the version of Sophistication I favor, intentions formed antecedently at the start of a sequential choice problem can play a tie-breaking role throughout the problem, serving to fix downstream behavior at later choice nodes in which multiple options appear maximally choiceworthy. The proponents of resolute choice, however, go much further than this. According to their approach, agents may permissibly recognize past resolutions to implement certain plans as relevant to their current evaluation of plans in such a way as to lead them to opt for plan continuations that they would have, in other contexts at least, considered strictly suboptimal.

To see how this works, consider again an Allais agent facing the sequential choice problem described in §5. A resolute agent facing this problem will consider all the plans logically available to her at its start and resolve to implement whichever she judges to be ex ante optimal. At subsequent stages of the problem, she will evaluate still available plan continuations according to whether they extend the plan(s) initially judged optimal. So, since, from the perspective of the Allais agent, the plan that involves delaying choice until the receipt of further information, and then opting for L_1 regardless of what is learned, is ex ante optimal, the agent will judge this plan as the one to be implemented. If she sets out on this course of action and subsequently learns either E or its negation, the resolute agent will continue to view the continuation of this plan as optimal in virtue of the fact that it extends an ex ante optimal plan. So while sophisticated choice is a forward-looking approach to sequential decision making, resolute choice is a backward-looking approach.

[58]For McClennen's defense of resolute choice, see [60]. Also relevant are [57], [29], and, in a somewhat different way, [61] and [92].

There are several ways to understand this idea. On one way of reading the proposal, prior intentions or resolutions can affect the outcomes an agent receives in a dynamic choice problem, perhaps because the agent values being a steadfast planner. Alternatively, perhaps the agent's preferences are sensitive to counterfacts in such a way that the agent's preference for $L1$ over $L2$, conditional upon E, is contingent upon what would have been happened had E been false. Either of these ways of understanding resolute choice, while reasonable, obviate the need for it as a distinct method of sequential choice. On either of these interpretations, we simply need to refine the agent's outcome space to account for their preferences regarding steadfastness and/or counterfacts, and there is no longer any need to consider such agents as offering any counterexample to the Sophistication principle.[59]

Qualification as a distinctive method of sequential choice requires an alternative interpretation of resolute choice. Two suggestions naturally present themselves here. First, one could view resolute choice as requiring counter-preferential choice. On this reading of the method, after opting to delay her decision and then learning E, the resolute Allais agent will, in line with her Allais preferences, judge that it would be better to shift her initial plan and opt for $L2$ and yet, in line with her prior resolution, nevertheless opt for $L1$ instead. Second, one might view resolute choice as inducing suitable preference shifts in dynamic choice problems so as to ensure that, acting upon those preferences, the agent realizes an initially favored plan. In the case of the Allais agent, this would mean shifting to prefer $L1$ over $L2$ conditional upon E.

[59] Granted, akin to the issues noted in footnote 43, some of these proposals may call into question other assumptions of both the Savage setup (e.g. its rectangular field assumption) and the dynamic choice arguments offered here (e.g. the assumed unrestricted domain of decision trees). These are matters that I believe ultimately tell in favor of Jeffrey's framework over Savage's but discussion of this must wait until the next chapter. Though see [24] for an interesting argument that these worries may be unfounded.

Each of these readings renders resolute choice implausible as a choice method for instrumentally rational agents. On the interpretation of act-preference that I have been employing, a preference for one prospect over another reflects a judgment that the former act has greater instrumental value (i.e. is better suited to the end of bringing about good outcomes) than the former. Requiring that agents either act contrary to these judgments or revise them in the way suggested above in order to rationally approach sequential choice is far-fetched. An agent's judgment of the instrumental value of a prospect, conditional upon some event E, should not depend upon what prospects would be realized by the agent conditional upon $\overline{E}$, as required by the preference shift interpretation of resolute choice. On the other hand, if an agent truly judges one prospect as better than another how can she be fully rational while acting against that judgment, as the counter-preferential interpretation of resolute choice would have it? On neither interpretation then does resolute choice constitute an approach to sequential decision making suitable for instrumentally rational agents.

1.8.3 An Inner Tension?

I have argued that both the Standard Objection and the alternative Resolute Choice Objection to diachronic tragedy arguments fail. A lurking doubt may remain, however. Some authors, most notably Lara Buchak and Johanna Thoma, have recently argued that the principles undergirding these arguments are jointly at odds with one another and that this incongruence can be recognized without having to endorse either time-slice rationality or resolute choice theory. While Buchak's and Thoma's objections are similar in style, they make quite different claims and will have to be treated in turn.

Buchak argues that adoption of a particular approach to sequential choice (e.g. sophisticated or resolute) implies adherence to a corresponding view of the nature of personal agency. Sophisticated choice, she suggests, assumes a time-slice view of agency on which all that an agent can influence at any given time are presently available actions, with future choices being left to future time-slices. Resolute choice, by contrast, assumes, according to Buchak, a view of agency on which a planning agent is relatively unified across time and may adopt genuinely cross-temporal plans of action (so-called 'diachronic options'). Buchak further claims that some of the dynamic choice principles needed to run diachronic tragedy arguments are plausible only on the time-slice view of agency while others are plausible only on the more unified view of agency. Hence, the conjunction of these principles is implausible.

There is a certain way of framing diachronic tragedy and related arguments on which Buchak's criticism seems to hold water. If diachronic tragedy is thought to reveal irrationality through identifying individual choices that are irrational in light of the agent's ultimate values, then the diachronic tragedy argument for the STP fails. None of the choices taken by the Allais agent in the sequentialized Allais problem is state-wise dominated by other available choices. Only the sequence of choices taken as a whole has this feature. Hence, I don't think the liberal standards of evalaution appropriate to a theory of instrumental rationality can identify any of the Allais agent's individual choices as irrational. However, this is not how I have proposed that we understand diachronic tragedy arguments. A successful diachronic tragedy argument reveals that an agent is irrational by demonstrating that her action-guiding attitudes are liable to license behavior guaranteed to result in suboptimal outcomes. This does not mean that any of the individual choices constitutive of that behavior can be called out as irrational, but it does mean the agent is ill-equipped to pursue optimal outcomes in sequential choice problems, and hence cannot be a paragon of instrumental rationality.

Johanna Thoma worries that diachronic tragedy arguments involve a different sort of tension. In particular, Thoma suggests that decision theorists face a choice regarding how to assess an agent's instrumental rationality. First, one might directly employ as a standard of evaluation the agent's attitudes toward uncertain prospects (i.e. the agent's act-preferences). On this proposal, if an agent's behavior ever leads her to realize a prospect that is against her act-preferences, she would qualify instrumentally irrational. Second, one might instead take as the standard of evaluation just the agent's attitudes toward ultimate consequences (i.e. the agent's outcome-preferences). On this proposal, if an agent's behavior ever leads her to forseeably realize an outcome that is against her outcome-preferences, she would qualify as instrumentally irrational. Thoma suggests that when the proponent of Sophistication indicts the resolute Allais agent for choosing contrary to her act-preferences at the second stage of the sequential Allais problem, she is assuming the first standard of evaluation. When she indicts the sophisticated Allais agent for diachronic tragedy by contrast, she is shifting to the second standard of evaluation. Since neither standard of evaluation can allow us to derive both Sophistication and Sequential Dominance (or Dynamic Optimality), decision theorists interested in characterizing instrumental rationality have no grounds for insisting that agents conform to this set of principles.

Ultimately, Thoma argues that we should embrace the second standard of evalaution since the first leaves an agent's attitudes toward uncertain prospects totally unconstrained. On the first proposal, Thoma claims, we could not even criticize an agent who chooses an option state-wise dominated by another, since in doing so she may still be realizing a prospect she prefers more and it is act-preferences rather than outcome-preferences that fix a standard of rational evaluation on this view. But, having granted the second proposal, we no longer have grounds to criticize the resolute Allais agent in the sequentialized Allais problem. True, this agent acts contrary to her act-preferences at a time, but she never falls into diachronic tragedy and so is not, Thoma claims, criticizable in light of her outcome-preferences.

This line could then be taken as a novel defense of resolute choice. If an instrumentally rational agent is, as I have claimed, one that never forseeably frustrates her own ends (i.e. never falls prey to diachronic tragedy) and if the resolute Allais agent never suffers such a fate, on what grounds can we label the Allais agent instrumentally irrational? Counter-preferential choice is the key that allows the Allais agent to evade diachronic tragedy! I reply that a resolute Allais agent may well succeed in avoiding diachronic tragedy, but this fails to vindicate the coherence of her judgments. On the test defended here, an agent's outcome-preferences form the standard of evaluation against which her act-preferences, and not simply her behavior, may be judged as rationally admissible. Rationally admissible attitudes are, at minimum, those that never license foreseeably suboptimal plans of action. It is an inadequate defense of such attitudes to plead that one can escape their ill effect by acting contrary to their recommendations. Such an argument is a bit like defending the value of a broken-down car by suggesting that one can always take the bus. The line in effect concedes that the attitudes in question are rationally defective.

1.9 Conclusion

I have suggested that we think of normative decision theory as an attempt to formally characterize instrumental rationality. Rational agents, in this sense, are those whose behavior is well-suited for the attainment of their ends. This includes, at a minimum, never acting so as to forseeably realize suboptimal outcomes relative to the set of all possible outcomes achievable by alteration of the agent's behavior. Further, I have proposed that an agent's attitudes are rationally admissible (in the instrumental sense) just in case they are ones that a rational agent could act upon. In suitable contexts, this motivates the Sequential Dominance principle and its rejection of tragic attitudes (strong and weak) as irrational.

Nonetheless, as noted, there are several reasons one might doubt the status of Sequential Dominance as a fully general rationality postulate. For example, as soon as we countenance the possibility of infinite decision problems, avoidance of sure-loss becomes an impossible standard of rationality. Further, the Sequential Dominance principle and the diachronic tragedy arguments it undergirds assume the existence of a domain of ultimate outcomes relative to which an agent's performance may be assessed. For some purposes, positing the existence of such a domain may be unwarranted or even undesirable. With these concerns in mind, I have suggested that Dynamic Optimality provides a more generally applicable standard of instrumental rationality. This standard is, given a rich domain of consequences, equivalent to Sequential Dominance in the setting of finite sequential choice problems modelled within the Savage framework. However, it suffers from neither of the drawbacks of Sequential Dominance highlighted here. Nor are the other recent objections to dynamic choice principles (e.g. those voiced by time-slice rationality or resolute choice proponents) compelling against it.

The remainder of my dissertation will turn to exploring the import of Dynamic Optimality and some related principles in the underexplored context of Richard Jeffrey's Logic of Decision and its causalist variants. In particular, Chapter 2 will motivate Jeffrey's framework, which takes agents to express preferences over a single Boolean algebra of propositions rather than over act functions, and draw out several key advantages it enjoys over Savage's. I will then consider how to model sequential choice problems within this broader framework and suggest that we may define the plans available to an agent in a given sequential choice problem via a conditional operator that we may neutrally dub a *planning conditional*. The remainder of the chapter explores what features this conditional operator must exhibit in order for agents satisfying Jeffrey's theory of desirability maximization to also satisfy Dynamic Optimality. As it turns out, a material construal of the planning conditional leaves desirability maximizers open to Dynamic Optimality violations, while a construal employing Richard Bradley's characterization of indicative conditionals guarantees the Dynamic Optimality of desirability maximizers.

These results notwithstanding, there are serious reasons to doubt the soundness of Jeffrey's theory as an account of instrumental rationality. Newcomb problems seem to reveal that the sort of causal notions intentionally skirted by Jeffrey are essential to a full theory of rational choice. Chapter 3 considers what the ramifications of Dynamic Optimality are in the context of causal decision theory. The salience of this investigation is heightened by the fact that Arif Ahmed and other decision theorists have recently provided several troubling arguments to the effect that causal decision theory is incompatible with the avoidance of diachronic tragedy and related ills. I argue that while some of Ahmed's arguments fail to reveal any significant worry regarding the compatibility of causal decision theory and Dynamic Optimality, others are far more worrisome. Ultimately, I will argue that reconciling causal decision theory and Dynamic Optimality requires endorsement of a radical principle of *Autonomy* according to which an agent's judgments of the causal and evidential import of her own decisions always coincide. While indeed radical, consideration of a rational agent's deliberational processes and necessary awareness of her own action-guiding attitudes may nonetheless render the principle acceptable to those convinced of the normative force of both causal decision theory and Dynamic Optimality.

Chapter 2

Dynamic Optimality in the Logic of Decision

With an informal defense of Dynamic Optimality laid out, this second chapter turns to the matter of formally modelling sequential choice within Jeffrey's decision theoretic framework. The first problem that confronts us here is that Dynamic Optimality concerns the evaluation of *plans*, while the Logic of Decision is a theory concerned with ranking *propositions* according to their desirability. Thus, to make Jeffrey's theory applicable to sequential choice, a bridge principle is needed connecting evaluation of plans with evaluation of propositions. Relying upon the standard practice of modelling sequential choice problems via Bayesian decision trees, I show how we can model plans in a Jeffrey-style framework by employing conditionals that I neutrally dub '*planning conditionals*'. I then investigate the dynamic optimality of Jeffrey's theory of desirability maximization in light of a simple material construal of the planning conditional. While it can be shown that desirability maximization is dynamically optimal in the restricted class of decision trees that involve only choice nodes (or, equivalently, that the relative desirability of plans never reverses pursuant to choice nodes in decision trees), a sequentialized version of Gibbard and Harper's Transparent Newcomb Problem suffices to bring out the potential for dynamic suboptimality in a more general setting. This motivates consideration of alternative construals of the planning conditional. It turns out that an indicative reading suffices to render Jeffrey's theory dynamically optimal, given Richard Bradley's account of indicative conditionals, and is, further, necessary if Jeffrey's theory is to satisfy a stronger cousin of Dynamic Optimality that I dub *Preference Stability*. However, this interpretation of the planning conditional has significant implications regarding the evidential bearing of states of the world upon an agent's future decisions.

2.1 Introduction

In the context of Savage's theory, Dynamic Optimality is a powerful principle with wide-ranging implications. Given its logical connection to axioms like Sequential Dominance, these implications have been extensively drawn out and explored by decision theorists. However, the significance of dynamic choice norms has received comparative neglect in the context of the more general approach to decision theory pioneered by Richard Jeffrey in his *Logic of Decision*. While there are notable exceptions,[1] this general lack of attention is surprising given both the philosophical merits of Jeffrey's framework and the widely granted appeal of dynamic choice principles.

[1][83], for example, discusses Hammond's arguments against Allais-type preferences from the standpoint of Jeffrey's theory, while [1] and [58] offer sequential choice arguments showing that value of knowledge results fail in the context of Jeffrey's framework. [77] and [61] also discuss Jeffrey-style theories in the dynamically relevant context of the *self-recommendation* of decision theories. More recently, [98] has deployed a clever sequential choice example to undergird a 'why ain'cha rich?' argument against Jeffrey's theory, while [2] has argued that causalist variations on Jeffrey's theory fail to be dynamically consistent.

This second chapter of my dissertation aims to begin to remedy this situation by identifying conditions under which an agent describable by Jeffrey's theory will satisfy Dynamic Optimality. This task requires some stage setting in the form of, first, motivating the need for Jeffrey's framework and, second, considering how to model dynamic choice within this framework. §2 thus rehearses some of the well-known advantages of the Logic of Decision over Savage's decision theory and presents the basic framework and content of Jeffrey's theory. §3 then explores how to mesh the formalism of dynamic choice theory with that of the Logic of Decision. The key proposal here will be an identification of plans in sequential choice contexts with suitable conjunctions of conditionals that I neutrally dub *planning conditionals*.

With a suitably general dynamic choice framework in place, I turn to consideration of the import of Dynamic Optimality in the context of Jeffrey's theory of desirability maximization. An initial positive result lends some support to the rationality of desirability maximization: within the class of decision problems modelled by decision trees whose only non-terminal nodes are choice nodes (i.e. decision problems that lack a non-terminal role for nature) the Dynamic Optimality of Jeffrey's theory is guaranteed (§4). However, a turn to more general dynamic choice problems in which an agent's moves may be interspersed with those of nature invites a more complicated story. Crucial to the Dynamic Optimality of desirability maximizers in this general context is the interpretation of the planning conditional. I demonstrate that a simple material construal results in a version of desirability maximization that is prone to violations of Dynamic Optimality (§5). However, I subsequently show that an indicative reading suffices to guarantee the Dynamic Optimality of desirability maximization, assuming the analysis of indicative conditionals recently defended by Richard Bradley (§6). Further, a key feature of this interpretation of the planning conditional is necessary if we wish our theory of instrumental rationality to satisfy a slightly stronger cousin of Dynamic Optimality requiring that an agent's judgments regarding the comparative instrumental value of available plans remain invariant throughout the course of a sequential choice problem, a principle I call *Preference Stability*.

Appealing as these results may appear with respect to fixing an interpretation of the planning conditional, I argue in §7 that an indicative reading carries with it strong philosophical commitments regarding how a rational agent may view the evidential bearing of conditionals involving her future acts upon past states of the world. This foreshadows consideration of a more radical rationality postulate to be considered in Chapter 3. §8 concludes with a brief summary of the current chapter and a prelude of the next one dealing with the relevance of causality to the theory of rational choice. If causal decision theorists are right that evidential decision theory (i.e. Jeffrey's theory of desirability maximization) offers a poor account of instrumental rationality to begin with, then a more proper context in which to explore the import of Dynamic Optimality will be a causalized version of the Logic of Decision, and this will be the focus of my final chapter.

2.2 Jeffrey's Theory

2.2.1 Problems for Savage

Savage's account of rational choice is commonly considered the crowning achievement of modern decision theory. Granting purely qualitative assumptions about an agent's act-preferences (some of which, we have seen, may be underwritten by attractive dynamic choice arguments), Savage demonstrates that rational agents are representable as Bayesians. That is, a rational agent's act-preferences will reveal a unique probability function defined over states and a unique (up to positive affine transformation) utility function defined over outcomes, such that her valuation of uncertain prospects (i.e. Savage acts) goes by the expectation of the utility function, computed relative to the probability function. If we take the probability function to represent an agent's degrees of belief and the utility function to represent her degrees of desire, then we have derived a rich mathematical representation of her internal action-guiding attitudes from the simple starting point of her pairwise preferences among possible acts. This is indeed a remarkable accomplishment and justifies much of the adulation heaped upon Savage's theory by subsequent practitioners of the craft.

Still, even granting the normative force of controversial axioms like the STP, there are reasons to be dissatisfied with Savage's account as a fully general theory of instrumental rationality.[2] For one, some of the basic distinctions assumed by the Savage setup appear artificial and tenuous. Recall that Savage carves up decision problems into three basic components: acts, states, and outcomes. The latter two types are treated as distinct but primitive sets, while acts are defined in terms of states and outcomes. An agent's probability measure is defined only over events (i.e. sets of states), while her utility function takes only outcomes as input. Acts can only be assigned an expected utility. On this theory then, agents are never represented as having either beliefs about outcomes or about acts, nor as having desires regarding states. But this need not be so. An agent may well have beliefs about the items Savage would have us label 'outcomes' or 'acts' as well as desires regarding those he would have us label 'states'.[3] In general then, Savage's approach can only hope to partially reveal the cognitive and connative attitudes of an agent whose preferences conform to the Savage axioms.

[2]The concerns discussed below have been raised by a diverse chorus of voices, largely in philosophy, including Jeffrey, Wolfgang Spohn, John Howard Sobel, John Broome, Jim Joyce, and Richard Bradley. As Spohn points out, Savage himself even recognized the worries of state-outcome utility dependence and act-state probabilistic dependence in [73].

[3]It should be noted, however, that the claim that rational agents can have beliefs about their own acts is not without controversy. See [87] and [52].

Further, since, in Savage's framework, outcomes are always realized in the context of some state and since utilities are defined only over outcomes and not, e.g., state-outcome conjunctions, it is necessary to assume that the desirability of any outcome for the agent is independent of which state it is realized in. The utility independence of outcomes and states, however, requires one to take great care in formulating both the outcome- and state-spaces of a Savage model. Many natural ways of formulating a given decision problem will commonly run afoul of this criterion.[4] At the same time, the structural assumptions of Savage's theory require a sort of care in this task of describing outcomes that may push in an opposite direction. For example, the supposition that the set of acts over which an agent may express preferences is the set of *all* functions from states to outcomes requires that an agent be able to judge the instrumental value of a wild array of acts mapping arbitrary states to arbitrary outcomes.[5] Outcomes and states must then plausibly be characterized such that any outcome is logically compatible with any state and also with the assignment of any other outcome to any other state. As noted in the last chapter, this also leads to difficulties if one wishes to reconcile Allais-like preferences with the STP by introducing counterfactual content into the description of outcomes, for one would then have to judge the instrumental value of acts that map states to outcomes in ways that contravene the contents of the specified outcomes. (Consider: "If it rains, get $100, and if it doesn't, get [$0 and if it had rained you would have gotten $0 too].") The required utility-independence of states and outcomes, as well as concerns about accommodating seemingly sensible preferences like those of Allais without contravening the STP, thus motivates a highly fine-grained characterization of outcomes, while the structural assumptions of Savage's theory seem ill-suited to accommodate such a characterization.

A further, perhaps more serious worry, regards the fact that Savage derives a single

[4]This objection to Savage is pressed by, among others, [87].

[5]This is the infamous *Rectangular Field Assumption* attacked by [17] and [11], though defended by [24].

probability measure defined only over events and never conditioned upon acts. Hence, Savage's theory requires an appropriate assumption of act-state independence. For example, consider the classic decision problem modelled by Table 1. Here, an agent must

	Cancer	No Cancer
Smoke	-1	2
Don't Smoke	-2	1

Table 2.1: The Smoking Problem

decide whether to smoke or not and considers as possible states Cancer and No Cancer. Her utilities for the outcomes associated with the relevant act-state pairs are given. Savage's theory appears to recommend smoking here, as this option maximizes expected utility with respect to any single probability function defined over {Cancer, No Cancer}. However, such reasoning is famously fallacious as the probability of the states vary with the agent's choice of act in this example. That is, the acts and states in this example are *probabilistically dependent*. For Savage's theory to have any plausibility then, it must assume that we have characterized states in such a way that they are suitably independent of acts.[6] This will at the very least require a great deal of care and may even be impossible in principle in some cases.[7]

Another possible worry regarding the Savage system concerns its postulation of ultimate outcome and state spaces in the first place, regardless of how the members of such spaces are characterized (e.g. as utility-independent of states in the case of outcomes or probabilistically-independent of acts in the case of states). Savage's characterization of the sets Z and $\mathbf{S}$ effectively casts an agent's basic attitudes as directed toward atomic points that are complete in the sense that they leave no room for further refinement. So a Savage outcome, for example, cannot be partitioned into distinct outcomes with differing utilities,

[6]Whether the requisite sort of independence is really probabilistic independence or some sort of causal independence is a matter taken up in the next chapter.

[7][45] argues it is quixotic to insist upon act-state independence in general, though some other decision theorists have attempted to devise algorithms for accomplishing this, see e.g. [33].

as one might partition non-ultimate outcomes (e.g. have a glass of milk) into possibilities of different value (e.g. have a glass of fresh milk and have a glass of sour milk). As noted in the previous chapter, the assumption of an ultimate outcome space is crucial to the statement of strict dominance principles and their attendant arguments. But is the existence of ultimate outcomes any sort of requirement of instrumental rationality? Might an agent not have preferences over an algebra of proximate outcomes that is atomless, i.e. contains no points or ultimate members? It would be ideal, I think, for a theory of instrumental rationality to allow for this possibility, which may then motivate a move away from Savage's framework.[8]

A final worry, raised forcefully by Jim Joyce[9], concerns the infamous *problem of small worlds*. Small world decision problems are more coarse-grained approximations of potentially vastly complicated 'grand world' decision problems that specify every possible way an agent's available acts might bear on the world in ways she cares about. Savage defended the normative plausibility of his axioms in the context of grand world decision problems, whereas bounded agents can plausibly only hope to actually model small world decision problems. Hence, there is a problem of justifying the application of standard grand world principles to tractable small world models. In light of this, it is standardly thought that in order to provide an adequate account of rational choice a decision theory must be structured such that the agents characterized by it could reasonably expect their small world conclusions to approximate their grand world preferences. Joyce has argued that Savage's theory lacks a reasonable assurance of such congruence since the expected utility of an act on the theory is highly sensitive to the partition of states employed to calculate it. Any theory that aspires to answer the problem of small worlds, Joyce suggests, must exhibit *partition invariance* in the sense that an assessment of an act's value according to the theory

[8] [23] has argued for adoption of something like Jeffrey's decision theory on the basis of a rejection of ultimate outcomes.

[9] See [45] and [46]. [11] has also endorsed this criticism of Savage.

may be conducted with respect to more or less coarse-grained partitions of states with similar effect. As Joyce argues, satisfaction of this constraint is one of the central virtues of Jeffrey's alternative approach to decision theory, which we now turn to consider.[10]

2.2.2 Jeffrey's Alternative

Richard Jeffrey was among the first philosophers to point out some of these limitations of the Savage model and to propose an alternative framework for modelling rational choice, his *Logic of Decision*.[11] The first limitation of Savage's framework that Jeffrey aimed to correct was its inability to capture probabilistic dependence between acts and states. Brian Skyrms offers what is, for our purposes, a particularly telling summary of the appeal of this move:

> [The Logic of Decision] originally was motivated by a desire to deal with cases in which states are not independent of acts and to do so without causally loaded concepts. Probabilities are defined on a large boolean algebra, whose elements are taken to be propositions and whose operations are to be taken as truth functions. Acts are construed as propositions in this space that can be directly "made true" by the decision-maker. The inclusion of acts in the boolean algebra over which probabilities are defined is an innovation which may provoke varying reactions. **However, one may argue that this feature makes the system attractive for dealing with sequential decision problems. In such problems, the choice of an option may change its status over time**

[10]One further advantage which Jeffrey seemed to think his theory enjoyed over Savage's, but which I opt not to include in my catalogue, involves Jeffrey's expulsion of the causal notions thought to be implicit in Savage's theory in light of the conditional nature of Savage acts. As the emergence of causal decision theory suggests, however, this alleged advantage is arguably illusory.

[11][44].

from consequence to act to part of the state of the world, and each change goes with an appropriate updating of subjective probability. ([80], p. 44, emphasis added.)

Skyrms' thought regarding the affinity between Jeffrey's framework and sequential choice is a very attractive one. Given this, it is remarkable how little attention has been paid to the import of dynamic choice principles in the context of the Logic of Decision and its causalist variants.

As Skyrms notes, Richard Jeffrey assumes that rational agents express preferences over a single *algebra*, $\mathcal{A}$, of propositions, that is, a collection of propositions closed under the standard Boolean connectives of negation, conjunction, and disjunction.[12] Borrowing notation from the previous chapter, this preference relation may again be written as $\geq$, with asymmetric and symmetric parts denoted $>$ and $\sim$, respectively. We may think of these preferences as expressing an agent's judgments regarding the desirability of the propositions in the algebra being true. Or, following a famous suggestion recommended to Jeffrey by Savage, we might also view these preferences as encoding an agent's attitudes toward propositions construed as 'news items'. Intuitively, an agent prefers a proposition p to a proposition q (i.e. $p > q$) just in case she would rather learn that p than learn that q. That is, if p is 'better news' than q. For example, I find the proposition stating that I win a million dollars in the lottery more desirable than the proposition stating that I fail to win the lottery; I would prefer to learn the former rather than the latter. So I have a preference for the former over the latter of these two propositions.

Jeffrey suggested that rationality conditions placed upon $\geq$ would suffice to guarantee its representability via a desirability function V and probability measure P defined upon $\mathcal{A}$

[12] Less the contradiction, which Jeffrey excludes from an agent's preference ranking.

satisfying:

The Desirability Axiom: For any pair of non-null, mutually exclusive propositions $X, Y \in \mathcal{A}$, $V(X \vee Y) = P(X|X \vee Y)V(X) + P(Y|X \vee Y)V(Y)$.

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. Note that this axiom entails that desirability functions exhibit the partition invariance sought by Joyce:

Partition Invariance: $V(X) = \sum_i P(X_i|X)V(XX_i)$ for any finite partition $\{X_i\}_i \subseteq \mathcal{A}$.

The qualitative conditions required for such a representation were first discovered by the mathematician Ethan Bolker, and according to many philosophers constitute either plausible rationality constraints or relatively harmless technical assumptions.[13] In Jeffrey's theory then, we have an account of rationality that, in addition to eliminating the need to invoke ultimate outcomes, doing away with the awkward state-outcome dualism of the Savage model, countenancing probabilistic dependence between acts and states, and satisfying a partition invariance principle, also rests upon a mathematically elegant representation theorem whose axioms are (to a large extent) normatively intuitive.

Following in the footsteps of Hammond's treatment of Savage and of the arguments sketched in Chapter 1, it would be ideal at this point to investigate the extent to which Bolker's axioms might be derivable from norms of sequential coherence like Dynamic Optimality. Alas, I see no apparent way to carry out such a project. The liberal structure of Jeffrey's theory, which, for example, postulates preference amongst propositions which might never feasibly serve as items in a common choice set, renders the prospects for a

[13] For Bolker's proofs, see [10]. For a defense of the plausibility of Bolker's assumptions, see [45] and [11].

strict derivation of Jeffrey's theory from dynamic choice norms rather dim. That said, such norms may still enjoy significant import in this setting if we allow them to act as additional rationality principles taken alongside of the standard Jeffrey-Bolker axioms. This is the path of investigation I will pursue. Assuming Jeffrey's assumptions to be genuine requirements of instrumental rationality, what, if anything, does the additional postulate of Dynamic Optimality add to Jeffrey's account of rationality?

2.3 Planning Conditionals

If we hope to explore the significance of Dynamic Optimality in the Logic of Decision, it is incumbent upon us to first precisify the application of Jeffrey's theory to problems of sequential choice. We have already reviewed the basics of dynamic choice theory in the preceding chapter and only a few key adjustments will be required to harmonize that framework with the Logic of Decision. Importantly, we will alter the definition of decision trees as follows:

Definition 7. *A* **decision tree** *is a seven-tuple* $\langle N, X_1, X_2, X_3, N_+(\cdot), n_0, S(\cdot)\rangle$*, where:*

i *N is a finite set of nodes, partitioned into X_1, X_2 and X_3.*

ii *X_1 is the set of choice nodes.*

iii *X_2 is the set of natural nodes.*

iv *X_3 is the set of terminal nodes.*

v *$N_+ : N \to \mathcal{P}(N)$ is the immediate successor function which satisfies:*

(a) $\forall n \in N, n \notin N_+(n)$

(b) $\forall n \in N, N_+(n) = \emptyset$ *iff* $n \in X_3$

(c) $\forall n, n' \in N, N_+(n) \cap N_+(n') \neq \emptyset$ *iff* $n = n'$

vi n_0 *is the initial node and satisfies* $\forall n \in N, n_0 \notin N_+(n)$

vii $S : N \to \mathcal{A}$ *is a mapping that assigns to each node a proposition such that,* $\forall n \in N, \{S(n')|n' \in N_+(n)\}$ *is a partition of* $S(n)$.[14]

Note the few discrepancies between this definition and the characterization of decision trees offered previously. The set-theoretic events formerly associated with nodes are now replaced with propositions, i.e. members of $\mathcal{A}$. Learning is now possible pursuant to choice nodes as well as natural ones. And, lastly, gone are the contingent consequence functions associated with terminal nodes; the elements of $\mathcal{A}$ now suffice as (proximate) outcomes. Aside from these modest alterations, the characterization of decision trees is left as before.

Recall now that, according to the dynamic choice framework developed in the previous chapter, rational agents facing sequential choice problems adopt attitudes towards available plans, construed as sets of terminal nodes in appropriately structured decision trees, whereas in the Logic of Decision, rational agents take attitudes towards propositions. To make Jeffrey's theory applicable to sequential choice then, a propositional construal of plans is called for. A natural proposal is to identify plans with suitable conjunctions of act propositions and partitioning conditionals.[15] (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

[14]Again, this definition closely aligns with the formulation of decision trees found in [36] and [22], albeit now adjusted to the setting of the Logic of Decision. Given the background algebra, $\mathcal{A}$, I will assume an unrestricted domain of decision trees.

[15] A *partitioning conditional* is a conjunction of conditionals whose antecedents form a logical partition. See [11], p. 122-4.

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. 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. 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. Succinctly:

Definition 8. 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

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

3. If n is a natural node, then

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

We can also define plan continuations in this setting as follows:

Definition 9. 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.[16]

The trick going forward 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 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. Thus, in my view, the Dynamic Optimality of desirability maximization under a material construal of the planning conditional is an excellent launching point for an investigation into the significance of Dynamic Optimality in the context of the Logic of Decision.

[16] 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 [11], p. 97, for more on Conditional Desirability.

2.4 Dynamic Choice without Nature

As we investigate the dynamic optimality of desirability maximization, it is convenient to begin by considering a particularly simple class of decision problems, namely, those that lack any non-terminal role for nature to play. That is, we first consider 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 Jeffrey's decision theory that the dynamic optimality of desirability maximization is guaranteed within this class of problems. In particular, we can prove:

Proposition 2. *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 $S(n')$.*

Proof. Let T be a decision tree and n a choice node in T. Suppose that:

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

Then, assuming p, p' are consistent with $S(n')$, where $n' \in N_+(n)$:

(2) $V_n(pS(n')) \geq V_n(p'S(n'))$

By definition of conditional desirability:

(3) $V_n(p|S(n')) \geq V_n(p'|S(n'))$

By conditionalization:

(4) $V_{n'}(p) \geq V_{n'}(p')$

By the equivalence of plans and their continuations following choice nodes:

(5) $V_{n'}(p(n')) \geq V_{n'}(p'(n'))$

$\square$

This proposition, 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 revision of previously adopted plans on the part of desirability maximizers arise following new disclosures of information by nature. Thus, if one plan is preferred to another at a choice node and both are defined at an immediately subsequent node, the continuation of the initially favored plan will continue to look better than the continuation of the initially disfavored plan at the subsequent node. 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 such reversals arise following new disclosures of information by nature.

Proposition 1 clearly entails that desirabilty maximization satisfies Dynamic Optimality within the restricted class of decision problems modelled by decision trees lacking natural nodes. For, suppose that a plan p is judged optimal (i.e. is of maximal desirability) at the start of such a decision tree. Since, according to Proposition 1, the relative desirability of plans never flips pursuant to choice nodes, p will persist in appearing maximally favorable throughout the problem in question. Hence, it is dynamically feasible. Hence, all ex ante maximal plans are also dynamically feasible, just as required by Dynamic Optimality.[17]

2.5 Sequential Transparent Newcomb

So far, so good for the desirability maximizer. 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 optimality 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 renders our interpretation of the planning conditional relevant. As noted above, our first pass at interpreting the planning conditional will involve treating it truth-functionally as a material conditional. Taking the planning conditional in this way, a simple example suffices to show that evidentialists are liable to exhibit dynamic suboptimality.[18]

[17] In proving this proposition in [71], I assumed a finite base set of possible worlds and construed propositions as sets of such worlds, which, as the current proof shows, was unnecessary and resulted in a slightly less elegant proof.

[18] The following thought experiment derives from a story suggested to me by Brian Skyrms in personal correspondence. I have simplified that original story in light of helpful comments from an anonymous reviewer. I later also found that a similar argument is discussed by [7], though the use he puts it too is slightly different.

The problem is a sequentialized variant of the *Transparent Newcomb* problem that we may dub *Sequential Transparent Newcomb* (STN).[19] 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.

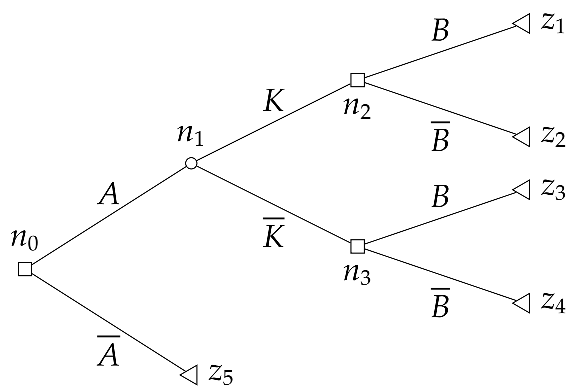

Figure 2.1: The Sequential Transparent Newcomb Problem.
$S(n_0) = W$, $S(z_1) = AKB$, $S(z_2) = AK\overline{B}$, etc.

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 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

[19]The static version of Transparent Newcomb was first discussed by [32]. [77] and [61] employ the example to argue that EDT can fail to be *self-recommending*.

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:

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

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 desirability maximization 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 desirability maximizer finds herself trapped in dynamic suboptimality, for the plan she judged ex ante optimal is not among those that are feasible for her.

This then suffices to establish:

Proposition 3. *If the planning conditional is the material conditional, desirability maximization is not dynamically optimal.*

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.)[20]

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? This is the question to which we now turn.

2.6 Bradley Indicatives

One natural alternative to the material interpretation of the planning conditional 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 I 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.[21]

[20] [71] argues for a version of this proposition stated in terms of dynamic consistency and puts this argument forward in the context of [2]'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.

[21] Note that I 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

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,

implementation. What I deny is simply that the planning conditional itself should be viewed subjunctively.

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.

Whatever the merit of these considerations, a key factor favoring an indicative reading of the planning conditional, from the perspective of the Logic of Decision, will prove to be its capacity to ensure that desirability maximizers satisfy Dynamic Optimality. To establish this, I will rely on the rich formal account of the pragmatics of indicative conditionals recently defended by Richard Bradley.[22] In particular, I will assume that rational attitudes toward indicative conditionals exhibit two properties proposed by Bradley.

The first condition is called the *Indicative Property* ([11], p. 117):

Definition 10. *A conditional $\mapsto$ has the* **Indicative Property** *with respect to a rational desirability function V just in case: For all propositions α, β, γ, $V(\alpha \mapsto \beta) \geq V(\alpha \mapsto \gamma)$ iff $V(\alpha\beta) \geq V(\alpha\gamma)$.*

Why should this condition hold of indicative conditionals? As Bradley notes, indicative conditionals $\alpha \mapsto \beta$ and $\alpha \mapsto \gamma$ each purport to tell us what is the case if α is true as a matter of fact, but they are silent on what is the case if α is false. Hence, the relative desirability of these conditionals should be settled by their relative desirability under the supposition that α is true, which amounts to making a comparison of the conjunctions $\alpha\beta$ and $\alpha\gamma$.

A second condition that Bradley posits as a marker of rational attitudes toward indicative conditionals is an *Additivity* principle:

Definition 11. *A conditional $\mapsto$ satisfies* **Additivity** *with respect to a rational desirability function V just in case, where $\{\alpha_i\}_i$ is any finite partition of propositions: $V(\bigwedge_i(\alpha_i \mapsto \beta_i)) = \sum_i V(\alpha_i \mapsto$*

[22] See especially chapter 7 of [11].

β_i).

Additivity is actually a logical consequence Bradley derives from a more basic principle he posits called *Value Independence*, according to which orthogonal indicative conditionals are desirabilistically independent.[23] The intuition Bradley identifies behind this postulate is simply that the desirability of one proposition on the matter-of-fact supposition of another should not depend upon what is the case if the supposition is false, for we have supposed it to be factual.[24]

Granting that Bradley has correctly characterized indicative conditionals via the Indicative Property and Additivity, we may define:

Definition 12. The conditional $\mapsto$ is an **indicative** if it satisfies the Indicative Property and Additivity.

It is now possible to prove that if planning conditionals are indicatives (that is, if a rational agent treats them as such by having a value function that conforms to the Indicative Property and Additivity), then the dynamic optimality of desirability maximization in a planning context is guaranteed.

Proposition 4. *If the planning conditional is an indicative, then desirability maximization is dynamically optimal.*

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. It would suffice to convince ourselves that for any successor to n, say n', $\Pi(n')$ is either undefined (i.e. Π did not make arrival at n'

[23]Two conditionals are *orthogonal* if their antecedents are mutually exclusive.

[24]I should emphasize that for subjunctive conditionals Bradley is adamant that both the Indicative Property and Additivity become highly implausible. So these properties are really taken by Bradley to govern only indicative conditionals.

feasible) or desirability maximal at n'. For, if this were true, then the feasibility of initially maximal plans would be guaranteed. 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 Bradley ascribes to indicatives 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, and hence that initially maximal plans are a subset of dynamically feasible plans, i.e. that Dynamic Optimality holds.

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 Optimality that we might call *Preference Stability*.

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 $p, p' \in \Omega(T, n_a)$, $p \geq p'$, and p, p' both make arrival at n_b possible, then $p(n_b) \geq p'(n_b)$. An agent is *preferentially stable tout court* just in case the agent is preferentially stable in all decision trees.

While this principle requires more than Dynamic Optimality, 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 optimal 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 trembling hand of a rational agent need not be blind. 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 5. *If the planning conditional fails to satisfy the Indicative Property with respect to a given desirability function, then an agent with that desiribility function will be preferentially unstable.*

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:[25]

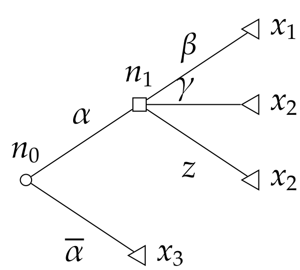

Figure 2.2: T_2, a decision tree.

[25]If β and γ suffice to partition α, z will be a contradiction and can be pruned from the 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) $V_{n_1}(\alpha\beta) < V_{n_1}(\alpha\gamma)$

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

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)[26] and we can thus take β and γ to be such a pair. To

[26]The closure of the algebra under the Boolean connectives is of course standard; further closure under

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

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.

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 $V(\alpha\beta^*) < V(\alpha\gamma^*)$.

$\square$

2.7 Future Autonomy

The forgoing argument has revealed that satisfaction of attractive dynamic choice norms like Dynamic Optimality and Preference Stability is, in the context of the Logic of Decision, closely tied to an interpretation of planning conditionals as indicatives. For those of us that find these principles normatively compelling, this speaks heavily in favor of such a construal of planning conditionals, assuming of course that Jeffrey's theory provides a sound account of instrumental rationality in the first place. It is worth drawing out further, however, what such a commitment amounts to philosophically. Upon reflection, requiring an interpretation of planning conditionals can be seen to place strong restrictions upon an agent's probability function. In particular, it amounts to requiring that an agent never view her conditional performance of future acts as evidence for the states upon which their performance is conditioned.

Consider two planning conditionals with a shared antecedent consisting of a state-proposition and distinct consequents consisting of competing act-propositions that could be performed conditional on learning the truth of the antecedent. We might denote them '$S \mapsto A_1$' and '$S \mapsto A_2$'. Given the indicative property, the first is preferred to the second by a desirability maximizer only if SA_1 is preferred to SA_2. But this condition might well fail if the conditionals in question could provide varying evidence in support of their shared antecedent S. For example, if S is a negative proposition that an agent would really like to turn out false, then $S \mapsto A_1$ might well have higher desirability than $S \mapsto A_2$, even granting that $V(SA_2) \geq V(SA_1)$, provided that the former conditional supplies significantly stronger evidence against S than the latter.

This is essentially what occurred in the Sequential Transparent Newcomb problem. While taking the thousand is preferred to refusing it after one has learned the contents of the opaque box, behaving in accord with this preference is antecedently strong evidence in favor of the bad news that the opaque box is empty and, thus, the plan that involves turning down the thousand when offered is actually ex ante optimal. By insisting that planning conditionals be treated as indicatives, the advocate of desirability maximization is effectively arguing that cases like STN are impossible for rational agents to face. On this view, a rational agent will not take her present plans (codified as planning conditionals) to supply her with evidence for or against the states upon which her plans are conditioned. Let's dub this claim the *principle of Future Autonomy*.

> **Future Autonomy**: If $S \mapsto A$ is a planning conditional for a given agent, then $P(S|S \mapsto A) = P(S)$, where P is the agent's probability measure.

This principle follows from treating planning conditionals as indicatives given Bradley's general treatment of the pragmatics of indicative conditionals on which $P(\alpha(\alpha \mapsto \beta)) = P(\alpha)P(\alpha \mapsto \beta)$, for any propositions $\alpha, \beta \in \mathcal{A}$.

This view strikes me as perfectly tenable, though it may appear radical to some. In fact, I will suggest in the next chapter that it does not go quite as far as the proponent of Dynamic Optimality ought to in restricting an instrumentally rational agent's deliberative credences regarding the evidential significance of her acts. Still, it is a strong principle and a proponent of Dynamic Optimality might well wish to resist it as a rationality principle. Alas, such resistance seems largely futile outside of abandoning desirability maximization as a theory of instrumental rationality. Moreover, even if one were to take this route, the leading competitor to Jeffrey's theory, causal decision theory, offers no safe harbor to flee to in this storm, as the marriage of causal decision theory and Dynamic Optimality motivates the even stronger autonomy principle previously referenced. One might be tempted instead to seek refuge in one or another of the decision theories that have recently come out of the computer science community, including the 'functional decision theory' of Yudkowsky and Soares or Yudkowsky's 'timeless decision theory'.[27] Certainly, the attractive dynamic properties of these theories constitute a main source of their appeal. Still, insofar as they aspire to characterize instrumentally rational *choice*, they appear to make absurd recommendations in cases like Sequential Transparent Newcomb, where they would recommend forgoing the sure thousand, regardless of what is observed in the opaque box![28] Better to say, I think, that adherence to Dynamic Optimality comes at the price of granting some form of internal epistemic independence between an agent's conditional plans and the obtaining of the corresponding conditioning states of the world.

2.8 Conclusion

The Logic of Decision overcomes various shortcomings of the standard models of rational choice popular among economic decision theorists (e.g. Savage's) and exhibits a range of technical virtues that recommend it as a formal framework for explicating instrumental rationality. These virtues have been widely recognized among philosophers, leading to the preeminent position currently enjoyed by Jeffrey's decision-theoretic framework

[27] [99], [100].
[28] See p. 22-3 of [99].

within philosophical decision theory. At the same time, ever since Hammond's derivation of Bayesianism within a Savage-style framework[29] from instrumentalist dynamic choice norms, the appeal of principles akin to Dynamic Optimality has been commonly recognized, even as a serious investigation of these norms in the context of Jeffrey-style decision theories has been neglected. An exploration of the significance of dynamic choice principles in the context of the Logic of Decision has thus been long overdue.

This chapter has sought to initiate this project and, in the process, reached several results worth highlighting. First, I have argued that plans can be naturally identified with propositions if we introduce a conditional operator into the algebra's space of logical connectives to be interpreted as a *planning conditional*. Second, the prospects for Jeffrey's theory of desirability maximization to satisfy Dynamic Optimality hinge upon the sort of attitudes rational agents may adopt with respect to planning conditionals. On a simple truth-functional understanding of the planning conditional as a material conditional, there is no guarantee that desirability maximizers will satisfy Dynamic Optimality, a point brought home by the Sequential Transparent Newcomb problem. This is so because, on this reading of the planning conditional, nothing forbids preference reversals amongst plans pursuant to new disclosures of information by nature, even though such reversals are indeed impossible for desirability maximizers pursuant to the agent's selection of her own acts, regardless of the interpretation of the planning conditional. Finally, an alternative construal of the planning conditional as an indicative conforming to Bradley's characterization of matter-of-fact conditionals does manage to secure Dynamic Optimality as a consequence of desirability maximization.

Adopting this interpretation of planning conditionals, however, carries with it signifi-

[29]Or, more strictly, an Anscombe-Aumann framework, [6].

cant commitments. In particular, it amounts to affirming Future Autonomy. That is, one must hold that a rational desirability maximizer is bound to view her available contingency plans as evidentially irrelevant to the obtaining of the contingencies they plan for. The agent is autonomous in that she can never read any information regarding which facts about the world she will learn off of which acts she will perform in light of those facts.

Interesting as I believe these results are, they constitute merely the initiation of a nascent research program investigating the significance of dynamic choice norms within Jeffrey-style decision theories. There is still much work to be carried out in such a program. As a critical next step, attention ought to be directed toward the dynamic optimality (or lack thereof) of causalist variants of the Logic of Decision, i.e. causal decision theory. It has long been recognized, its many virtues notwithstanding, that the Logic of Decision faces serious obstacles as an account of rational choice due to its self-conscious neglect of the relevance of causal factors to sound decision making, and, hence, that some sort of causalized version of Jeffrey's theory is needed. It is not clear, however, that the central versions of causal decision theory on offer are compatible with Dynamic Optimality and other plausible choice norms. In fact, a number of authors have recently argued that the dynamic choice setting raises insuperable difficulties for causalist accounts of rational choice.[30]

My final chapter will attempt to rebut some of these charges against causal decision theory and reconcile the causal perspective with Dynamic Optimality. In the process, I will argue that, while a few of the mentioned objections can be dismissed with relative ease, others will require strong medicine that readers may be differentially inclined to swallow. However one comes down on the matter of reconciling causal decision theory

[30]These arguments are made, for example, by [2], [4], [85] and [63].

and Dynamic Optimality, at the very least, I hope to make clear that the interplay of causal concerns and dynamic choice theory opens up a range of rich and fascinating philosophical problems that are well worth the further attention of decision theorists.

Chapter 3

Newcombian Tragedy

While Jeffrey's Logic of Decision has much to recommend it as a theory of rational choice, Newcomb problems seem to present convincing counterexamples. In light of this, many philosophers have turned to *causal decision theory* (CDT) as a more adequate alternative. Unfortunately, a number of recent authors have constructed sequential choice problems that seem to reveal that CDT contradicts plausible dynamic choice principles including Dynamic Optimality. These problems are generally taken as casting doubt on the normative status of either CDT or the relevant dynamic choice principles. This chapter proposes a way out of this dilemma. We can affirm the normative necessity of both CDT and Dynamic Optimality if only we also embrace an epistemic principle requiring an agent's assessment of the causal and evidential import of her acts to coincide. This principle, which I dub *Autonomy*, turns out to be sufficient for a causalist agent's satisfaction of Dynamic Optimality. While radical, I argue the principle bears an affinity to a number of related theses defended in the literature on independent grounds and presents the only real prospect for a reconciliation between CDT and Dynamic Optimality.

3.1 Introduction

The choices we make often provide us with information about the world. My decision to turn the spigot valve is evidence of imminent water flow. Opting to press my neighbor's doorbell suggests that she may come and open the door. A flip of the light switch on my wall portends a change in ambient lighting conditions. Etc. In cases like these, the evidential bearing of my decision upon the relevant state of the world is underwritten by a causal connection. Turning the spigot valve is evidence that water will flow because turning the valve tends to cause water to flow. And so on. This is the typical case and generates no paradox.

More puzzling are those cases in which our choices seem to provide evidence for states of affairs that they have no tendency to cause. Game theory supplies a familiar example. My friend and I are to play a Prisoner's Dilemma. I take my practical reasoning to be similar to her's and hence my behavior to be indicative of her's. My choice to cooperate (defect) is then evidence of my friend's choice to cooperate (defect), even though our respective choices are causally isolated from one another. Here, analogical, rather than causal, reasoning seems to motivate an evidential connection between my choice and my friend's.[1]

Cases of this sort have lead to well-known difficulties for Jeffrey's Logic of Decision as an account of instrumentally rational choice. When the evidential and causal import of an agent's choices come apart, is an option's choiceworthiness to be fixed by its evidential news value or by its causal efficacy? Evidentialists say the former, while causalists maintain the latter. Jeffrey's unadorned theory of desirability maximization is archetypi-

[1]This classic example is originally due to [53].

cally evidentialist: the desirability of an act is a measure of its news value and bears no immediate relation to its causal efficacy. For those of us inclined to think of instrumental rationality in causal terms, this casts serious doubt upon the general adequacy of Jeffrey's theory.

Fortunately, decision theorists have managed to modify the Logic of Decision to address these worries, resulting in an account that retains the basic virtues of the Jeffrey model while simultaneously incorporating more direct sensitivity to causal concerns. This *Causal Decision Theory* (CDT) has much to recommend it and is regarded by many philosophers as the most convincing analysis of rational choice presently on offer. However, a flurry of recent objections to the theory has begun to sow increasing doubt in the field regarding the ultimate tenability of CDT as a characterization of rational agency. Most troubling among these objections, to my mind, are those contending that agents describable by CDT are prone to a variety of sequential choice woes including dynamic inconsistency and susceptibility to dynamic dutch books. In effect, these arguments suggest that CDT is incompatible with Dynamic Optimality.

My goal in this chapter is to explore how one might patch up CDT to avoid any conflict with Dynamic Optimality. Ultimately, I will suggest that, while some of the recent sequential-choice-inspired objections to CDT rest upon relatively easily corrected mistakes about the nature of sequential choice, a full reconciliation of CDT with Dynamic Optimality requires endorsement of a strong rationality postulate I dub *Autonomy*, according to which the (perceived) causal and evidential import of an agent's acts always coincide from the perspective of the acting agent. Agents who respect this norm never view their acts as evidence of causally unrelated states of the world. This is a radical thesis in the context of the CDT and EDT debate. If Autonomy really characterizes the attitudes of rational deliberating agents, then I may not, for example, take my choice to cooperate (defect) in the Prisoner's Dilemma as evidence of my causally separated friend's choice. Since only a suspected causal connection can purchase such an evidential link, the analogical reasoning sketched above turns out to be illicit.

After introducing some background on CDT (§2) and presenting two recent dynamic choice arguments against CDT (§3), I argue that the latter of these two arguments, while *prima facie* troubling for the causalist, assumes an analysis of sequential choice that is at odds with the account hereto defended in this dissertation and hence need not worry the causalist provided she adopts the account defended here (§4). However, the causalist is still left with the first first dynamic choice problem as it cannot be deflected in this manner. This leads me to propose Autonomy as the only available remedy. In particular, I demonstrate that causalists who respect Autonomy invariably satisfy Dynamic Optimality, guaranteeing their immunity from dynamic dutch books and other diachronic woes (§5). Given the radical nature of Autonomy, my conclusion in §6 attempts to render it more palatable as a rationality postulate by connecting it to related ideas in the literature.

3.2 Causality and Rational Choice

As we saw in the previous chapter, Jeffrey envisions agents that express preferences over an algebra of propositions construed as possible news items. These preferences are assumed to be structured so as to be representable via a pair of desirability and probability functions, (V, P), such that V satisfies the desirability axiom with respect to P and one proposition is preferred to another just in case it has greater desirability according to V. Recall also that, according to evidentialists, translating all of this into a theory of rational decision is simple. We just view the set of options an agent faces in a decision problem as a partition of propositions by associating each act open to the agent with the proposition that she performs it, and then stipulate that rational choice consists in maximization of desirability.EDT is simply the theory of rational choice that gives this prescription. More precisely, if a decision problem is given by a set of acts $A_1, \ldots, A_n$ that form a logical partition, then EDT prescribes that agents select an act that is V-maximal, i.e. an act A_i such that $V(A_i) \geq V(A_j)$ for $1 \leq j \leq n$.[2]

While this approach continues to enjoy its defenders,[3] it has faced considerable criticism as an account of rational choice. In particular, it seems to go awry in *Newcomb problems*. Such cases have been taken as showing that EDT is inadequate as a theory of rational choice because it ignores the relevance of specifically causal information to decision making. Hence, while Jeffrey had originally considered the absence of causal notions in his theory to be one of its key advantages over the arguably more causally-loaded theories of Ramsey and Savage, this apparent advantage seems to devolve into a liability.

3.2.1 Newcomb's Problem

In the classic Newcomb problem,[4] an eccentric millionaire sets before you two boxes, one opaque and the other transparent. The transparent box can be seen to hold $1,000. The millionaire tells you that the opaque box contains either one million dollars or nothing. She then offers you a choice between taking both boxes (and all their contents) or just the opaque box. Here's the catch. This millionaire is an incredibly reliable predictor of human behavior and she has placed the $1,000,000 in the opaque if and only if she predicted that you would take only the opaque box. The predictor has already made her prediction and

[2]In the event that there is more than one such act, EDT deems all V-maximal acts rational.
[3]Most notably, [2].
[4]The Newcomb problem was first introduced to the philosophical literature by [62].

fixed the contents of the opaque box by the time you have to deliberate about what to do. Suppose also that you only care about your wealth level and that your utility for money is linear.

Denote the decision to take only one box by A. Since you must take either one box or two, $\neg A$ is then the decision to take both boxes, and $\{A, \neg A\}$ forms a partition. Let K be the proposition that the predictor has placed \$1,000,000 in the opaque box and $\neg K$ the proposition that she has left it empty. $\{K, \neg K\}$ is then a partition of propositions that is causally independent of your choice to one-box or two-box, i.e. your decision has no causal bearing on whether K is true. The partition is also set up so that each of its elements, when conjoined with a particular act, can be associated with a determinate monetary outcome. We can then express the Newcomb problem in the following decision table:

	K	$\neg K$
A	1,000,000	0
$\neg A$	1,001,000	1,000

Table 3.1: The Newomb Problem

Here, as in all decision problems, EDT would have us maximize desirability. That policy favors A, assuming the predictor is taken to be sufficiently reliable. For example, suppose $P(K|A) = .9$ and $P(\neg K|\neg A) = .9$. Then, letting '$X = x$' be the proposition that the agent earns \$$x$ and recognizing that the problem countenances only four possible monetary outcomes, we have:

$$\begin{aligned} V(A) &= \sum_x P(X = x|A)x \\ &= (.9)(1,000,000) + (.1)(0) \\ &= 900,000 \end{aligned}$$

and also

$$
\begin{aligned}
V(\overline{A}) &= \sum_x P(X = x|\overline{A})x \\
&= (.9)(1,000) + (.1)(1,000,100) \\
&= 100,910
\end{aligned}
$$

So we have that $V(A) > V(\overline{A})$. Hence, EDT endorses taking only one box and so forfeiting the $1,000 dollars in the transparent box.

This is the wrong answer. Or so it has seemed to many.[5] The predictor has either placed the money in the opaque box or she hasn't. If she did place the money in the opaque box, then taking both boxes will produce a better pay off than sticking with one. Similarly, if she left the opaque box empty, you will likewise be better off taking both boxes. So, regardless of what the predictor did, two-boxing produces better results than one-boxing, and that seems like a compelling reason to take both boxes.

3.2.2 Causal Decision Theory

In response to the challenge posed by Newcomb problems, philosophers have sought to reintroduce causal notions into the Logic of Decision without forfeiting the various advantages afforded by Jeffrey's framework. Such attempts typically fly under the banner of *causal decision theory*. While there are several versions of causal decision theory employed in the literature, they tend to share a common spirit and their differences will be of little consequence to our discussion of causalism's fate vis-a-vis Dynamic Optimality.[6] For

[5] Even many defenders of EDT have felt this way and have sought to argue that evidentialists need not be committed to one-boxing in the Newcomb problem. See, for example, [44] and [25].

[6] The first public defense of causal decision theory appears in [32], though their idea of formulating decision theory in terms of subjunctive conditionals stems from an earlier letter of Stalnaker's. Soon

definiteness and to emphasize continuity with Jeffrey's theory, let's fix our attention upon the popular rendition of causal decision theory defended by Jim Joyce.[7]

Joyce proposes that rather than computing the instrumental value of an act-proposition using standard conditional probabilities as Jeffrey's theory suggests, we ought rather to use *causal probabilities*, which intuitively measure how much causal, as opposed to evidential, support one proposition provides another. Where X and Y are propositions, I will write '$P_X(Y)$' to denote the causal probability of Y given X. There are various ways to understand this quantity, corresponding to different versions of causal decision theory. On one approach, $P_X(Y)$ might be identified with the probability of a non-backtracking subjunctive conditional featuring X as antecedent and Y as consequent, i.e. $P(X \mathbin{\Box\!\!\rightarrow} Y)$. We might also (and compatibility with the first proposal, I think) computing $P_X(Y)$ relative to a given background partition of causal factors or *dependency hypotheses* that are causally independent of X and individually suffice to fix the causal bearing of X upon Y and its negation. Writing such a (finite) partition as '$\{K_i\}_i$', we might define $P_X(Y) = \sum_i P(Y|XK_i)P(K_i)$. Nothing much going forward will hinge upon the exact manner in which we opt to construe causal probability.

What is important to emphasize is simply that causal probability is, in general, distinct from conditional probability. As a simple toy example, consider a group of six people, three wearing red shirts and three wearing blue shirts. Suppose I know that exactly two of those wearing red shirts are also wearing glasses, while only one of those wearing a blue shirt is. If I learn that a given individual is among the three wearing red shirts,

after, [76] and [54] proposed alternative versions of causal decision theory with less explicit reliance upon conditionals, while [84] offered an account in terms of primitive imaging. Somewhat later, a number of authors, including [65], [41], and [65], also attempted to translate the insights of causal decision theory into the framework of graphical causal models. Related here also is the *interventional decision theory* of [89].

[7] See [45].

absent further information, this will support the claim that she is also wearing glasses, i.e. $P(glasses|red) = .\overline{6}$. However, the probability that the individual in question is wearing glasses on the causal supposition that she is made to wear a red shirt is presumably equivalent to the prior probability of her wearing glasses, i.e. $P_red(glasses) = .5$, for intervening to fix the color of the agent's shirt will not affect my credences regarding her choice of eyeware, even though learning the color of her shirt (non-interventionally) would. Cases of spurious correlation (such as that between shirt color and wearing glasses in the considered example) forbid a straightforward identification of causal and conditional probabilities.

Granting an understanding of causal probability, Joyce suggests that the *utility* or *efficacy value* of an act may be identified with its desirability on the causal or *interventional* supposition that it is performed. That is, to compute the instrumental value of an act, an agent computes the desirability or news value of the act from the standpoint of interventionally supposing the act to be performed, that is, using her causal probabilities given the performance of the act in place of her standard conditional probabilities. Put formally, and letting U denote utility/efficacy value, we can understand CDT as requiring maximization of:

> For any $a \in \mathcal{A}$, $U(a) = V_a(a) = \sum_i P_a(S_i|a)V_a(S_ia) = \sum_i P_a(S_i)U(S_ia)$, where $\{S_i\}_i \subseteq \mathcal{A}$ is any finite partition.[8]

Applied to Newcomb's Problem, CDT delivers a verdict in favor of taking both boxes, provided that the agent really does take her choice to be causally irrelevant to the presence or absence of the million in the opaque box, i.e. $P_A(K) = P(K)$. Computing efficacy values relative to the partition $\{K, \overline{K}\}$ and recognizing that desirability and utility coincide on

[8]Here V_a is desirability on the causal supposition that a. That is, the suggestion here is to define U in terms of V computed relative to a 'causalized' probability measure, e.g. P_a.

propositions that suffice to fix the agent's wealth-level, we have:

$$\begin{aligned} U(A) &= P_A(K)U(AK) + P_A(\overline{K})U(A\overline{K}) \\ &= P_A(K)(1,000,000) + P_A(\overline{K})(0) \\ &= P_A(K)(1,000,000) \\ &= P(K)(1,000,000) \end{aligned}$$

and also

$$\begin{aligned} U(\overline{A}) &= P_A(K)U(\overline{A}K) + P_A(\overline{K})U(\overline{A}\overline{K}) \\ &= P_A(K)(1,001,000) + P_A(\overline{K})(1,000) \\ &= P(K)(1,001,000) + P(\overline{K})(1,000) \end{aligned}$$

So, regardless of how an agent facing the Newcomb problem spreads her credences over the relevant partition, $U(\overline{A}) > U(A)$. By taking account of an agent's causal judgments, CDT thus gets the right answer in Newcomb problems, improving upon Jeffrey's theory of desirability maximization.

3.3 Diachronically Exploiting the Causalist

This apparent improvement notwithstanding, CDT's prospects for serving as a generally satisfying account of instrumental rationality are called into question by a number of troubling dynamic choice problems in which the prescriptions of CDT appear to lead agents down the wrong path. While not framed explicitly in terms of Dynamic Optimality, each of these arguments could be viewed as a challenge to the compatibility of CDT and this compelling principle. Given the picture of rationality painted in this dissertation, this is certainly troubling news for causalists. Indeed, among the plethora of alleged counterexamples to CDT recently proposed in the philosophical literature, the dynamic choice objections constitute, it seems to me, the greatest threat to the tenability of the theory.[9] I think this threat can be averted but only at the cost of embracing a strong rationality postulate that has the effect of largely (though perhaps not entirely) deflating the conflict between causalists and evidentialists. Before presenting this argument, I will first rehearse two recent dynamic choice arguments against CDT, a first due to Arif Ahmed and a second due to Jack Spencer. While each of these arguments alleges a conflict between CDT and Dynamic Optimality, the latter argument is the more evidently troubling for CDT. However, happily for the causalist, Spencer's argument rests upon a questionable understanding of sequential choice, one at odds with the Sophistication principle defended in Chapter 1, and hence can actually be set aside with relative ease. It is only Ahmed's argument that demands stronger medicine.

3.3.1 Ahmed's Psycho-Insurance

The first dynamic choice argument against CDT rests upon a sequential choice problem constructed by Arif Ahmed, the theory's most formidable critic in recent years. Ahmed dubs the problem *Psycho-Insurance* and contends that it witnesses to the dynamic inconsistency of CDT while simultaneously subjecting causalists to a Dutch Book. As noted, however, we may take the argument as revealing a violation of Dynamic Optimality on the part of agents described by CDT.[10]

The Psycho-Insurance problem is based on Andy Egan's now famous Psycho Button case.[11] In the first stage of the problem, a Newcombian predictor offers you a choice be-

[9] Some of the more prominent *static* choice objections to CDT include [27], [**?**], [2], [**?**], and [86]. I set these aside to focus on what are, to my mind at least, more threatening dynamic choice objections.

[10] Ahmed offers this argument in [2] and [4]. Very similar arguments to Ahmed's are also offered in [63].

[11] [27]. Ahmed notes that cases of similar structure have appeared in the literature before, for example, in

tween pressing a button and refraining from pressing it. If she predicted that you would press the button, then she has set things up so that pressing the button will cause \$1 to be debited from your bank account. If she predicted that you would not press the button, then she has set things up so that pressing the button will cause \$1 to be credited to your account. (This is Ahmed's sanitized version of Egan's case.) Let P be the proposition that you press the button. The options you face here are then P and $\overline{P}$. Let K be the proposition that the predictor had predicted you would press while $\overline{K}$ is the proposition that the predictor predicted you wouldn't press. Assume again that your utilities coincide with your dollar payoffs, i.e. all you care about is money and you value it linearly.

You take the predictor to be fairly reliable so let's fix that $P(K|P) = P(\neg K|\neg P) = n$, where n is greater than .75. After you have made your decision with regard to pressing the button, you will be offered an opportunity to bet on whether the predictor correctly predicted your choice. The bet will pay \$0.50 if the predictor was correct and cost you \$1.50 otherwise. Let B be the proposition that you bet. This decision problem is captured in decision tree $T1$ depicted in Figure 1.

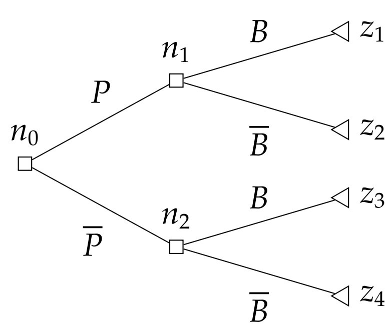

Figure 3.1: The Psycho-Insurance Problem.
$S(n_0) = T$, $S(n_1) = P$, $S(n_2) = \overline{P}$, $S(z_1) = PB$, $S(z_2) = P\overline{B}$, $S(z_3) = \overline{P}B$, $S(z_4) = \overline{P}.\overline{B}$.

The argument that the causalist will violate Dynamic Optimality in this problem is straight-

[30].

forward. To show that the causalist violates this principle, it suffices to show that she considers some plan acceptable at the start of the decision problem, i.e. at n_0, whose continuation at n_1 or n_2 she does not consider acceptable. Then there will be a plan that she views as ex ante favorable but will foreseeably deviate from and hence is infeasible for her.

Note first that there are four plans available to you at the start of this decision problem: $\Omega(T1, n_0) = \{PB, P\overline{B}, \overline{P}B, \overline{P}.\overline{B}\}$. We can follow Ahmed in depicting this decision problem in *normal form* using Table 2.

	K	$\overline{K}$
PB	-0.5	-0.5
$P\overline{B}$	-1	1
$\overline{P}B$	-1.5	.5
$\overline{P}.\overline{B}$	0	0

Table 3.2: The Psycho-Insurance Problem in Normal Form

Observe that relative to the partition $\{K, \overline{K}\}$ two of these plans are strictly dominated: PB and $\overline{P}B$. The first is dominated by $\overline{P}.\overline{B}$, while the latter is dominated by $P\overline{B}$. Given that the members of $\{K, \overline{K}\}$ are causally independent of any of your plans, i.e. $P_{PB}(K) = P(K)$, etc., if you are a causalist, you will deem each of these dominated plans unacceptable. Hence, either $P\overline{B}$ or $\overline{P}.\overline{B}$ will be most preferred at the start of the problem, or perhaps both. Which of these is the case will depend upon how likely you view each of the hypotheses, K and $\overline{K}$.

Hence, $D(\Omega(T1, n_0))$ cannot include the plans associated with PB or $\overline{P}B$. So since we know that $D(\cdot)$ always returns a non-empty set, it must contain at least one of the plans associated with either $P\overline{B}$ or $\overline{P}.\overline{B}$. Suppose it includes $P\overline{B}$. Then $P\overline{B}$ is deemed an acceptable plan at n_0. But suppose you carry out the first stage of this plan and press the button, i.e. make P true, and that, upon finding yourself at n_1, you update your beliefs by conditionalization.

It will then be the case that $B \succ_{n_1} \neg B$, since you take the predictor to be highly reliable and so consider the bet better than fair at that point. So, $(P\overline{B})(n_1) \notin D(\Omega(T2, n_1))$, since it is not maximal with respect to $\succeq_{n_1}$. But we then have a violation of Dynamic Optimality, for $P\overline{B}$ is judged to be ex ante optimal by the agent and yet its continuation at n_1 is not preferred by the agent at n_1, and hence the plan is not feasible. But then the set of optimal plans in T_1 is not a subset of the set of feasible plans, i.e. $D(\Omega(T2, n_0)) \nsubseteq DF_D(T1, n_0)$. If we assume instead that $\overline{P}.\overline{B}$ rather than $P\overline{B}$ is an ex ante acceptable plan, the causalist fares no better, as the same argument shows, *mutatis mutandis*, that Dynamic Optimality is again violated.

Note that, treating K and its negation as states and given the posited ultimate outcomes featuring in this problem, Psycho-Insurance also constitutes a case of diachronic tragedy on the part of the causalist. Given the Sophistication principle, we know that the causalist will end up betting at the second stage of this problem, but each of the plans that involve betting are strictly dominated by alternative plans that involve not betting. By altering the sophisticated causalist's behavior in $T1$, one could guarantee her an outcome more to her own liking.

3.3.2 Spencer's Two Rooms

For those of us inclined to affirm Dynamic Optimality as a criterion of instrumental rationality, Psycho-Insurance presents a real challenge to CDT. However, as surveyed in the first chapter, a number of decision theorists are willing to give up on Dynamic Optimality and a number of causalists appear inclined to take this route when faced with the Psycho-Insurance problem.[12] Another recent sequential problem posed by Jack Spencer and known as *Two Rooms* appears less easy to write off.[13] If Spencer is right about the

[12]This is the route taken by [48], for example.
[13]This problem appears in [85].

recommendations of CDT in Two Rooms, CDT is in trouble indeed.

Two Rooms is a sequentialized version of Spencer and Ian Wells' *Frustrator* problem.[14] To see how this problem goes, consider this static problem first. Suppose you must decide between taking one of three items: Box A, Box B, or an envelope. The envelope is known to contain $40, while the boxes are known to contain a collective total of $100. However, the distribution of the $100 across the boxes is uncertain. In fact, it was determined by a Newcombian predictor who placed the $100 in Box A just in case he predicted you would take Box B and placed it in Box B just in case he predicted you would take Box A. If the predictor predicted that you would opt for the envelope, he split the money evenly between the boxes, placing $50 in each. Assume, as usual, that you are aware of this, that your basic values align with your potential wealth-levels, and that you have no causal power over the past. In this case, CDT recommends taking either Box A or Box B, depending upon how likely you judge the predictor to have predicted you to choose each of these boxes.

To see this, denote taking Box A as 'A', taking Box B as 'B', and taking the envelope as 'E'. Let the predictor's prediction of each of these events be labeled 'A^*', 'B^*', and 'E^*', respectively. We can then compute the utility of each of your options relative to the partition, $\{A^*, B^*, E^*\}$, each of whose members is causally independent of your choice of

[14] [86].

act. In particular, we have:

$$\begin{aligned} U(A) &= P_A(A^*)U(AA^*) + P_A(B^*)U(AB^*) + P_A(E^*)U(AE^*) \\ &= P(A^*)U(AA^*) + P(B^*)U(AB^*) + P(E^*)U(AE^*) \\ &= P(A^*)0 + P(B^*)100 + P(E^*)50 \\ &= P(B^*)100 + P(E^*)50 \end{aligned}$$

$$\begin{aligned} U(B) &= P_B(A^*)U(BA^*) + P_B(B^*)U(BB^*) + P_B(E^*)U(BE^*) \\ &= P(A^*)U(BA^*) + P(B^*)U(BB^*) + P(E^*)U(BE^*) \\ &= P(A^*)100 + P(B^*)0 + P(E^*)50 \\ &= P(A^*)100 + P(E^*)50 \end{aligned}$$

$$\begin{aligned} U(E) &= P_E(A^*)U(EA^*) + P_E(B^*)U(EB^*) + P_E(E^*)U(EE^*) \\ &= 40 \end{aligned}$$

Regardless of how you spread your credences over the actions of the predictor, either $U(A)$ or $U(B)$ must come out greater than $U(E) = 40$, hence CDT will recommend forgoing the envelope.[15] This strikes many as counterintuitive, but in itself it may not be decisive as an objection to CDT, as Spencer grants.[16] The Two Rooms problem, however, turns the screw on the causalist a bit further.

Spencer asks us to imagine that you are offered the choice to go into one of two rooms. In the first you will receive $35 for sure. In the second, you will face the Frustrator problem. This problem is depicted in decision tree $T2$ in Figure 2. Which room should you enter? According to Spencer, CDT recommends entering room one to take the $35. Why? Well,

[15] This follows from the fact that at least one of $U(A)$ and $U(B)$ must be greater than or equal to 50. If this were not the case, then the sum $[P(B^*)100 + P(E^*)50 + P(A^*)100 + P(E^*)50]$ would have to be less than 100, but it is necessarily equal to 100.

[16] See, for example, the response to the Frustrator problem offered by [49].

let's use R to represent the proposition that you go to room one and, thus, $\overline{R}$ represents the proposition that you go to room 2. We know that $U(R) = 35$, so the only question is what is $U(\overline{R})$? Spencer sensibly assumes that you, as a sophisticated causalist agent, will employ some backward induction and realize that, were you to enter room two, you would opt to take either Box A or Box B. Let's say for concreteness that you would opt for Box A. Given this prediction, what is the utility of $\overline{R}$?

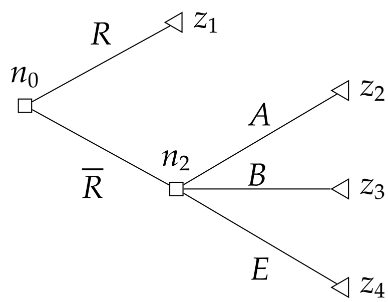

Figure 3.2: The Two Rooms Problem.
$S(n_0) = T$, $S(z_1) = R$, $S(n_2) = \overline{R}$, $S(z_2) = \overline{R}A$, $S(z_3) = \overline{R}B$, $S(z_4) = \overline{R}E$.

Employing the formula for computing causal expected utilities across the partition given by the possible combinations of your choices and the predictor's predictions in the Frustrator problem (and dropping the possibilities that yield zero utiltiy in the sum), we have:

$$
\begin{aligned}
U(\overline{R}) &= P_{\overline{R}}(BA^*)100 + P_{\overline{R}}(EA^*)40 + P_{\overline{R}}(AB^*)100 + P_{\overline{R}}(EB^*)40 + P_{\overline{R}}(AE^*)50 + P_{\overline{R}}(BE^*)50 + P_{\overline{R}}(EE^*)40 \\
&= P_{\overline{R}}(AB^*)100 + P_{\overline{R}}(AE^*)50 \\
&\approx 0
\end{aligned}
$$

This last line is justified by the fact that, given the assumed reliability of the predictor, you take it to be very unlikely that the predictor would have predicted that you would take either Box B or the envelope when you in fact take Box A. Since the utility of $\overline{R}$ is then judged to be near 0, $U(R) > U(\overline{R})$ and thus a causalist will clearly head to room one. But this is troubling. Not only does this behavior violate Dynamic Optimality (clearly the plan involving going to room 2 and taking the envelope is judged ex ante better than going to room 1), it also violates what Spencer calls the *Guaranteed Principle*. According to this principle, if one decision *guarantees* an outcome of a certain value in the sense of making that value certainly available to the agent contingent upon making the decision in question and another possible decision *ensures* a lesser value in the sense of certainly leading to an outcome of lesser value, then a rational agent ought to prefer the first decision to the second. In the case of Two Rooms, R ensures \$35, while $\overline{R}$ guarantees \$40 (by making available the envelope option). Hence, by preferring R to $\overline{R}$, the causalist violates the Guaranteed Principle.

3.4 CDT and Sophisticated Planning

By threatening its compatibility with Dynamic Optimality, Psycho-Button and Two Rooms each present a formidable challenge to CDT. However, as noted, the challenge posed by Two Rooms appears more urgent. If Spencer is right, CDT violates Dynamic Optimality in a particularly egregious way in this example by contravening the Guarantee Principle. Moreover, CDT's central competitors (e.g. EDT) manage to avoid this contravention of the Guaranteed Principle in Two Rooms, despite agreeing with CDT in the behavior they license (albeit in a dynamically optimal fashion) in Psycho-Insurance. Perhaps for this reason, the intuition that something has gone wrong for the causalist in Two Rooms is remarkably difficult to suppress. Fortunately, I don't think causalists have to suppress it, as Spencer's argument that CDT recommends choosing R and hence violating the Guaranteed Principle rests upon a contentious method of analyzing sequential choice problems that a causalist need not accept. In fact, the approach to sequential choice assumed by Spencer runs counter to the Sophistication principle defended in chapter 1.

Recall that, according to Sophistication, a rational agent's behavior norm in a sequential choice problem modelled by a decision tree T is given by the plans she judges as optimal in the tree T_D that is formed from T by pruning it back in line with backward induction using the agent's judgments of admissibility encoded in D. Applying this to a causalist facing Two Rooms, such an agent will indeed realize that, conditional upon reaching Room 2, she will opt for, say, Box A. However, it is this feasible plan (namely, choosing Room 2 and then taking Box A) that she will evaluate and compare against the competing option of going to Room 1. If she does this, she will find that, just as taking Box A has greater utility than taking the sure \$40 in the envelope, so too does the extended plan of taking Box A after going to Room 2 have greater utility than the plan of taking the sure \$35 in Room 1.

Putting this more formally, we can compute the causal efficacy of the plan $\overline{R}A$ relative to the partition $\{A^*, B^*, E^*\}$, whose members are causally independent of the agent's choice of plan:

$$\begin{aligned} U(\overline{R}A) &= P_{\overline{R}A}(A^*)0 + P_{\overline{R}A}(B^*)100 + P_{\overline{R}}(E^*)50 \\ &= P(B^*)100 + P(E^*)50 \\ &> 35 = U(R) \end{aligned}$$

Hence, the ex ante optimal plan available in $T2$ is $\overline{R}A$ rather than R, and thus a causalist agent who abides by Sophistication will implement $\overline{R}A$ rather than R, avoiding any violation of the Guaranteed Principle.

Spencer assumes that CDT is committed to evaluating presently available options by their utility considered as isolated decisions rather than by the utility of the plans they initiate taken as extended wholes. In fairness to Spencer, this is indeed the approach favored by some of CDT's most prominent defenders and it does lead to the wretched consequences he describes in Two Rooms.[17] But there is no reason that a causalist, qua causalist, should be committed to this approach to sequential choice. There are two orthogonal issues at play here. First, how should an agent evaluate the instrumental value of the objects of her practical deliberation, and, second, what is the nature of those objects of deliberation? CDT is a theory that answers the first question: an item of practical deliberation should be judged according to its utility or efficacy value. But this leaves open whether the objects that ought to be so assessed include, in deliberational contexts involving dynamic choice, extended plans or only isolated decisions at a time.

One objection to countenancing temporally extended plans as items of practical deliberation might reason that since the implementation of an extended plan depends not simply upon an agent's current choice but also her future choices implementing such a plan is not under her immediate control and, hence, may fail to be feasible for her to carry out. This is quite right as an objection to a full blown theory of resolute choice. A rationally infeasible plan is not a fit object of practical deliberation for a rational agent. But this offers no reason to think that a feasible plan cannot serve as such an object, and Sophistication includes this restriction. In fact, it seems a notable feature of human agency that we deliberate and form intentions regarding not only immediately available acts but also regarding larger contingency plans. Moreover, as argued in chapter one, the intentions we arrive at in the course of these deliberations often play an important role in terms of fixing future choices

[17] Jim Joyce, for example, does seem committed to this view in [49].

involving either indifference or parity, thus serving to unify an agent's decision making across time. To insist, as Spencer does, that plans cannot serve as the objects of practical deliberation for rational agents is to deny any real role for planning in sequential choice problems. On such a view, every non-trivial sequential choice problem is indistinguishable from an extensive form game in which every information set is occupied by a different player. This is a possible view of human agency, but it is not one that a causalist need accept, especially not one sympathetic to dynamic choice norms like Dynamic Optimality.

Note also the awkward double-think that Spencer's approach to sequential choice requires a causalist to adopt. Considered as an option I can carry out now, choosing Box A is judged by its utility and is worth more to me than \$40, but considered as a future possibility, choosing Box A is treated as a state of the world and judged by its desirability and hence worth less than \$35 to me. Causalists like Joyce have attempted to defend this pattern of valuation and there are certainly things one can say in defense of such a discrepancy between the evaluation of present and future actions, but causalists can escape the need for such stretching and straining altogether by embracing a more unified approach to dynamic choice. This certainly strikes me as the more appealing route to take.

If the view of sequential choice encoded in Sophistication is defensible, then a causalist need not fear the Two Rooms problem. It witnesses to a violation of neither Dynamic Optimality nor the Guaranteed Principle on the part of a causalist who satisfies Sophistication. Psycho-Insurance cannot be so easily dismissed, however. While Ahmed's case is *prima facie* less alarming than Spencer's, it does reveal a real conflict between CDT and Dynamic Optimality that an alternative gloss on sequential choice will not obviate. The agent facing Psycho-Insurance is welcome to assess the merits of her immediately available options (i.e. pressing the button or not) by evaluating the utility of the possible plans they initiate, but doing so will not alter the fact that not betting at the problem's second stage is both ex ante optimal and yet infeasible. A further remedy is thus required if CDT is to be reconciled with Dynamic Optimality.

3.5 Autonomy

I have argued that rational agents satisfy Dynamic Optimality. If this is right, then Ahmed has shown that agents he describes as opting to either press and bet or not press and bet in the Psycho-Insurance problem are irrational. Ahmed locates the source of this irrationality in these agents' allegiance to CDT. However, an agent's decision rule is only one possible source of irrationality. The attitudes that serve as inputs to her decision rule might also be to blame for instrumentally ineffective behavior on the agent's part. Alternatively, the update rule(s) that the agent employs to revise her attitudes over the course of implementing a sequential plan might be the real culprit. If only causalists who violate a further rationality condition are susceptible to violations of Dynamic Optimality, then there is no essential conflict between CDT and Dynamic Optimality. Those of us committed to the rational admissibility of CDT as a decision rule would then be advised to turn a critical eye toward the attitudes and update rules of causalist agents that violate Dynamic Optimality in problems like Psycho-Insurance.

What feature of the agents Ahmed considers might be open to criticism aside from their adherence to CDT? There are only a few choices here. For starters, one might question these agents' unrestricted adherence to conditionalization as an update rule for revising their credences and desirabilities in response to new information. Conditionalization is unobjectionable when applied to cases in which an agent passively learns the truth of a proposition identified in the algebra of events she has beliefs over. That is, in the context of dynamic choice, pursuant to natural nodes a rational agent ought to update her credences and desirabilities by conditionalization. However, conditionalization might be a more questionable assumption when applied to cases in which the new information an agent learns is brought about by her own willful decision to bring it about. That is, it might be questionable whether conditionalization properly characterizes the sort of learning that takes place pursuant to choice nodes in dynamic choice problems.

A number of philosophers have independently challenged this assumption for grounds different than those that concern us in this essay. John Cantwell, for example, has suggested that the problem of decision instability exemplified in one-shot decision problems like Gibbard and Harper's Death in Damascus or Egan's Psycho-Button provides pragmatic motivation to adopt a bifurcated update rule requiring one to update by conditionalization when learning states and by imaging (i.e. employing causal probabilities) when learning acts.[18] Melissa Fusco arrives at similar conclusions from considerations of *epistemic time bias*.[19] This idea also seems to share something in common with the way some causal decision theorist working in the tradition of graphical causal models often talk. On one way of reading such authors, a rational decision qualifies as an *intervention* that severs any grounds for correlation between the act decided upon and temporally antecedent states of affairs. Hence, according to this approach, while learning a state goes by conditionalization, learning an act goes by an alternative mode of belief revision appropriate to causal intervention.[20]

This is not the route I will pursue here. In what follows, I leave conditionalization in

[18] [20].

[19] [28].

[20] Some causal decision theorists view treating decisions as interventions as merely a useful fiction in the context of static decision making and hence do not view this as providing any rival updating rule to conditionalization. Others, however, do seem to intend to be writing more than fiction. See [41], [42] and the discussion in [90].

place as the uniquely rational update rule. Part of the reason for this is simply to explore the alternative proposal I wish to make. Another is to avoid challenging as orthodox a principle as updating on propositional evidence in accord with Bayes' Rule. Given the well-known diachronic tragedy arguments favoring conditionalization,[21] this latter reason is especially motivating in the current context in which our concern is precisely to avoid dynamic suboptimality. That said, I do think this path merits closer attention. It is not at all clear, for example, that the standard arguments for conditionalization (e.g. the dynamic dutch book argument) apply equally well when the event considered as a condition is one of the betting agent's own (future) act-propositions. Moreover, the alternative solution I will propose does share something in common, at least in spirit, with the ideas of authors like Cantwell and Fusco, so don't mean to distance myself too far from this proposal. It's just not the one I'll explore here.

Setting aside CDT and conditionalization as sacrosanct, the only option left is to challenge the rationality of the prior attitudes of the agents Ahmed considers. In particular, their prior credences. (Their basic wealth-obsessed desirabilities, while perhaps morally objectionable, are, of course, not a plausible point of attack for our purposes.) What feature of these agents' credences might one identify as objectionable? I want to propose that it is simply their Newcombian character. In other words, these agents' irrationality lies in the fact that they treat their own decisions as evidence for states of the world that they have no tendency to cause. It is precisely when an agent's assessment of the causal and evidential import of her plans come apart that she is liable to violate Dynamic Optimality and perhaps even fall prey to diachronic tragedy. When, on the contrary, an agent always views her possible plans as evidence only for what they have a tendency to cause, her satisfaction of Dynamic Optimality is guaranteed. I will say that such agents are *autonomous*

[21] See [55] and [82], the latter of which frames the dutch book argument for conditionalization explicitly in terms of dynamic choice.

since they satisfy:

> **Autonomy**: For any proposition $X \in \mathcal{A}$, any decision tree T, and any plan $\pi \in \Omega(T, n)$, $P_n(X|\pi) = P_n^{\pi}(X)$.

This principle requires the effects of evidential and causal supposition to coincide in the case of rational agents contemplating their own currently accessible plans of action. To be clear, evidential and causal supposition are clearly distinct forms of reasoning and do often come apart (as the toy example of §2.2 made clear), but Autonomy forbids this to happen in the case of the reflective decisions of fully rational agents. Note that adherence to such a rule in the case of the Psycho-Insurance problem would avert any violation of Dynamic Optimality on the part of the causalist by disavowing the prior credences assumed at the start of the problem. Assuming that the agent takes her decision to press or not press the button to be causally independent of the predictor's prediction, she must also, by Autonomy, take it to be evidentially independent of the predictor's prediction and hence updating by conditionalization following her decision will leave her confidence in the predictor's prediction unaltered and so there is never any need to purchase Ahmed's insurance at the problem's second stage. Alternatively, the agent might hold firm to the predictor's reliability, in which case she must, by Autonomy, forsake the assumption of causal independence and hence the dominated plans need not be judged antecedently inadmissible. Either way, Dynamic Optimality is satisfied.

Autonomy is more than a fortuitous remedy to the bad behavior described by Ahmed in Psycho-Insurance. We can prove that autonomous agents are guaranteed to never violate Dynamic Optimality. In what follows, I assume for simplicity that we may compute utility relative to a base set of possible worlds and that an agent's evaluation of the desirability of these base worlds is stable. As a further simplification, I also treat the planning conditional truth-functionally as a material conditional. However, I suspect that alternative interpretations of the planning conditional could yield the same result under modest assumptions. Finally, I shall also assume that decision trees have an implicit causal structure to them that precludes posterior acts from causally influencing prior learning events. That is, if n is a natural node in a tree T and $p \in \Omega(T, n), n' \in N_+(n)$, then $P_p(S(n')|S(n)) = P(S(n')|S(n))$.

Proposition 6. *Autonomous CDT agents satisfy Dynamic Optimality.*

Proof. Let T be an arbitrary decision tree faced by an autonomous causalist. We want to show that every ex ante optimal plan in T is feasible for a CDT agent who satisfies Autonomy, i.e. $D(\Omega(T, n_0)) \subseteq DF_D(T, n_0)$. Suppose $p \in D(\Omega(T, n_0))$. By CDT then, p is U_{n_0}-maximal. It now suffices to show that an option judged U_{n_a}-maximal at a particular node n_a will continue to be judged U_{n_b}-maximal at any node $n' \in N_+(n)$. That will suffice to guarantee that p is feasible. So, let n_a and n_b be nodes in T such that n_a immediately precedes n_b along some branch of T, i.e. $n_b \in N_+(n_a)$. Suppose $p \in D(\Omega(T, n_a))$. We want to show that $p(n_b) \in D(\Omega(T, n_b))$, if $p(n_b)$ is defined.

Case 1: Suppose n_a is a choice node.

We know that:

$$U_{n_a}(p) \geq U_{n_a}(p'), \forall p' \in \Omega(T, n_a).$$

Computing utility relative to the set of possible worlds W:

$$\sum_w P_p(w|S(n_a))U(w) \geq \sum_w P_{p'}(w|S(n_a))U(w), \forall p' \in \Omega(T, n_a).$$

By Autonomy:

$\sum_w P(w|pS(n_a))U(w) \geq \sum_w P(w|p'S(n_a))U(w), \forall p' \in \Omega(T, n_a)$.

But our definition of plans/continuations guarantees that plans are identical with their continuations pursuant to choice nodes:

$\sum_w P(w|p(n_b)S(n_b))U(w) \geq \sum_w P(w|p'(n_b)S(n_b))U(w), \forall p' \in \Omega(T, n_a)$, where continuation at n_b is defined.

A second application of Autonomy guarantees:

$\sum_w P_{p(n_b)}(w|S(n_b))U(w) \geq \sum_w P_{p'(n_b)}(w|S(n_b))U(w), \forall p'(n_b) \in \Omega(T, n_b)$.

By definition of U:

$U_{n_b}(p(n_b)) \geq U_{n_b}(p'(n_b)), \forall p'(n_b) \in \Omega(T, n_b)$.

Since every plan in $\Omega(T, n_b)$ has this form, $p(n_b) \in D(\Omega(T, n_b))$.

Case 2: Suppose n_a is a natural node. Then p has the form $\wedge_i[S(n_i) \rightarrow p(n_i)]$, where the n_i's are the possible successors to n_a. Note that this is equivalent to $\vee_i S(n_i)p(n_i)$. So, computing utility relative to W, we have:

$\sum_w P_{(\vee_i S(n_i)p(n_i))}(w|S(n_a))U(w) \geq \sum_w P_{(\vee_i S(n_i)p(n_i))}(w|S(n_a))U(w), \forall p' \in \Omega(T, n_a)$.

By Autonomy:

$\sum_w P(w|[\vee_i S(n_i)p(n_i)]S(n_a))U(w) \geq \sum_w P(w|[\vee_i S(n_i)p(n_i)]S(n_a))U(w), \forall p' \in \Omega(T, n_a)$.

This is equivalent to:

$\sum_w P(w|[\vee_i S(n_i)p(n_i)])U(w) \geq \sum_w P(w|[\vee_i S(n_i)p(n_i)])U(w), \forall p' \in \Omega(T, n_a)$.

Which is equivalent, by the Law of Total Probability, to:

$\sum_w \sum_i [P(S(n_i)p(n_i)|\vee_i[S(n_i)p(n_i)])]P(w|S(n_i)p(n_i))]U(w) \geq \sum_w \sum_i [P(S(n_i)p'(n_i)|\vee_i[S(n_i)p'(n_i)])]P(w|S(n_i)p'(n_i$

$\Omega(T, n_a)$.

Switching the order of the finite sums and abbreviating the plan:

$\sum_i \sum_w [P(S(n_i)p(n_i)|p)]P(w|S(n_i)p(n_i))]U(w) \geq \sum_i \sum_w [P(S(n_i)p'(n_i)|p')]P(w|S(n_i)p'(n_i))]U(w), \forall p' \in \Omega(T, n_a)$.

Which is equivalent, by Autonomy, to:

$\sum_i \sum_w [P(S(n_i)p(n_i)|p)]P_{p(n_i)}(w|S(n_i))]U(w) \geq \sum_i \sum_w [P(S(n_i)p'(n_i)|p')]P_{p'(n_i)}(w|S(n_i))]U(w), \forall p' \in \Omega(T, n_a)$.

Pulling the first term out of the inner sum:

$\sum_i P(S(n_i)p(n_i)|p) \sum_w P_{p(n_i)}(w|S(n_i))]U(w) \geq \sum_i P(S(n_i)p'(n_i)|p') \sum_w P_{p'(n_i)}(w|S(n_i))]U(w), \forall p' \in \Omega(T, n_a)$.

Which is equivalent, by definition of U, to:

$\sum_i P(S(n_i)p(n_i)|p)U_{n_i}(p(n_i)) \geq \sum_i P(S(n_i)p'(n_i)|p')U_{n_i}(p'(n_i)), \forall p' \in \Omega(T, n_a)$.

Which is equivalent to:

$\sum_i P(S(n_i)|p)U_{n_i}(p(n_i)) \geq \sum_i P(S(n_i)|p')U_{n_i}(p'(n_i)), \forall p' \in \Omega(T, n_a)$.

Which is equivalent, by Autonomy, to:

$\sum_i P_p(S(n_i))U_{n_i}(p(n_i)) \geq \sum_i P_{p'}(S(n_i))U_{n_i}(p'(n_i)), \forall p' \in \Omega(T, n_a)$.

Which is equivalent, given the causal structure of T, to:

$\sum_i P(S(n_i))U_{n_i}(p(n_i)) \geq \sum_i P(S(n_i))U_{n_i}(p'(n_i)), \forall p' \in \Omega(T, n_a)$.

But then:

$$U_{n_i}(p(n_i))U_(n_i)(p'(n_i)), \forall p'(n_i) \in \Omega(T, n_i),$$

which is what we aimed to show.

□

Whatever one thinks of Autonomy then, it at least has this virtue: agents whose credences conform to it may abide by both CDT and Dynamic Optimality without contradiction. The extent to which Autonomy is necessary for the avoidance of dynamic suboptimality on the part of causalists is a bit more difficult to spell out. Without relatively strong assumptions about the structure of an agent's attitudes, I doubt that a non-autonomous causalist will *invariably* fall prey dynamic inconsistency. However, the conditions under which they will do so are still quite broad. In an appendix, I construct a recipe for subjecting causalists with non-autonomous credences to failures of Dynamic Optimality, granting some additional assumptions about the agent's attitudes and allowing myself to employ a more narrow analysis of causal probability than the generic one that has largely sufficed for our purposes.

3.6 Conclusion

Barring a rejection of conditionalization, Autonomy is a causalist's only serious hope for consistently satisfying Dynamic Optimality. At this point, many will no doubt be inclined to give up on one or both of CDT and Dynamic Optimality. Autonomy, after all, is a very strong principle to embrace as a rationality postulate. It essentially precludes rational agents from ever facing Newcomb problems or from holding their own choices to be potentially correlated with causally irrelevant states of the world, e.g. the behavior of a physically remote twin. This may be a bridge too far for some, and I can't say I lack sympathy for such skeptics. Still, the retreat from Autonomy is an unpleasant one. Abandoning CDT is a gloomy prospect, especially since, as argued in Chapter 2, EDT offers no safe harbor to sail to. Jeffrey's theory either faces its own dynamic choice troubles or requires an analysis of planning conditionals that comes close to embracing the Autonomy principle (in the form of *Future Autonomy*) anyway. And it is hard to imagine alternatives to CDT and EDT providing much comfort here either.[22]

Holding CDT fixed, our choice is then between embracing Autonomy or abandoning Dynamic Optimality. Given the merits I have ascribed to Dynamic Optimality as a principle of instrumental rationality, it is worth considering whether Autonomy is really so implausible as a characterization of ideal rationality. A typical human agent will no doubt violate Autonomy, but then again typical human agents violate all the standard rationality postulates of Bayesian decision theories. Might an agent that is more collected and more reflective than we tend to satisfy Autonomy? If Ellery Eells is right, then the answer appears to be 'yes'.[23] According to Eells, a sufficient degree of self-awareness ought to screen off perceived correlations between an agent's decisions and causally irrelevant states of the world. So, for example, while my behavior might provide evidence regarding that of my causally isolated twin by way of revealing something about our shared action-guiding attitudes (e.g. beliefs and desires), once these decision inputs are fixed the evidential connection between our behaviors may be broken. Eells held that such self-awareness is partially consistutive of practical rationality since an awareness of one's own probabilities and utilities is crucial to the first-person application of decision

[22] True, the dynamic choice properties of non-standard theories like updateless decision theory ([100]) and functional decision theory ([99]) are, in my view, their main selling points. Still, the absurd recommendation of these theories in cases like Transparent Newcomb reveals the cost they pay to acheive these results while shirking Autonomy. Further, the Benchmark Theory of [97] is likely to face its own dynamic worries, though this is perhaps worth investigating more closely.

[23] See [25] and [26].

theory. If correct, this observation would obviate the need for causal decision theory, at least in the case of ideally rational agents, since its verdicts would always agree with those of evidential decision theory. Eells is getting at something close to Autonomy here by an entirely different route.

We might also instead view Autonomy as a requirement characterizing the ideal endpoint of rational deliberation. This thought seems to be in line with Eells' mature formulation of his defense of EDT in terms of deliberational metatickles.[24] On this view, an agent might reasonably start out deliberation with non-autonomous credences, but as she deliberates she goes through a process that includes a growing awareness of her own attitudes and dispositions that ultimately screens her choices off from all but causally downstream states of affairs. An exact dynamics that conforms to this desideratum and its implication for contexts of sequential choice are well worth investigating.[25]

Autonomy may also find a natural home in the quite different approach to reconciling EDT and CDT proffered by Huw Price.[26] According to Price's doctrine of *EviCausalism*, evidential significance conditional upon a free act from the agent's perspective is conceptually constitutive of causal dependence. So, for example, a fully rational and free agent who really takes her decision to cooperate in a Prisoner's Dilemma as evidence of the behavior of her twin must, contrary to common supposition, regard their decisions as causally intertwined. If correct, this leads directly to Autonomy. However, Price often seems to think of the 'evidential' probabilities at play in his arguments as distinct from an agent's actual epistemic probabilities. For him, they are rather *agential probabilities*.

[24] [26]

[25] Note that such a dynamic reading of Autonomy may, unlike a static reading, leave room for practical divergence between CDT and EDT if an agent's choice of decision theory impacts which autonomous credal state her deliberation terminates in. For more on deliberational dynamics, albeit not versions that hope to offer convergence to Autonomy, see [79], [7], [47], and [50].

[26] [67]

Hence, it is not clear to what extent Price actually means to be reconciling EDT and CDT as I see them. Still, to the extent that he does mean to carry out such a reconciliation, EviCausalism may supply an alternate route to defending something like Autonomy.

However plausible one finds the arguments of Eells and Price, the lesson of this essay is clear. 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. Decision theorists thus face a dilemma: either forsake Dynamic Optimality 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 the reader must judge.

Bibliography

[1] E. Adams and R. Rosenkrantz. Applying the jeffrey decision model to rational betting and information acquisition. *Theory and Decision*, 12:1–20, 1980.

[2] A. Ahmed. *Evidence, Decision and Causality*. Cambridge University Press, Cambridge, 2014.

[3] A. Ahmed. Exploiting cyclic preferences. *Mind*, 126:975–1022, 2017.

[4] A. Ahmed. Sequential choice and the agent's perspective. *MS*, 2020.

[5] M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. *Econometrica*, 21:503–46, 1953.

[6] F. Anscombe and R. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.

[7] F. Arntzenius. No regrets, or: Edith piaf revamps decision theory. *Erkenntnis*, 68:277–97, 2008.

[8] F. Arntzenius, A. Elga, and J. Hawthorne. Bayesianism, infinite decisions, and binding. *Mind*, 113:251–283, 2004.

[9] G. Asheim. Individual and collective time-consistency. *The Review of Economic Studies*, 64(3):427–43, 1997.

[10] E. Bolker. A simultaneous axiomatization of utility and subjective probability. *Philosophy of Science*, 34:333–40, 1967.

[11] R. Bradley. *Decision Theory with a Human Face*. Cambridge University Press, Cambridge, 2017.

[12] R. Bradley and H. O. Stefánsson. Counterfactual desirability. *British Journal of Philosophy of Science*, 68(2):45–533, 2017.

[13] S. Bradley and K. Steele. Can free evidence be bad? value of information for the imprecise probabilist. *Philosophy of Science*, 83(1):1–28, 2016.

[14] M. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

[15] M. Bratman. Reflection, planning, and temporally extended agency. *Philosophical Review*, 109(1):35–61, 2000.

[16] R. Briggs. Costs of abandoning the sure-thing principle. *Canadian Journal of Philosophy*, 45:827–40, 2016.

[17] J. Broome. *Weighing Goods*. Blackwell, Oxford, 1991.

[18] L. Buchak. *Risk and Rationality*. Oxford University Press, Oxford, 2013.

[19] L. Buchak. Revisiting risk and rationality: a reply to pettigrew and briggs. *Canadian Journal of Philosophy*, 45:841–62, 2016.

[20] J. Cantwell. On an alleged counter-example to causal decision theory. *Synthese*, 173:127–152, 2010.

[21] D. Christensen. Dutch books depragmatized: Epistemic consistency for partial believers. *Journal of Philosophy*, 93:450–79, 1996.

[22] R. Cubitt. Rational dynamic choice and expected utility theory. *Oxford Economic Papers*, 48:1–19, 1996.

[23] A. Demski. An orthodox case against utility functions. page www.greaterwrong.com, 2020.

[24] J. Dreier. Rational preference: Decision theory as a theory of practical rationality. *Theory and Decision*, 40(3):249–76, 1996.

[25] E. Eells. *Rational Decision and Causality*. Cambridge University Press, Cambridge, 1982/2016.

[26] E. Eells. Metatickles and the dynamics of deliberation. *Theory and Decision*, 17:71–95, 1984.

[27] A. Egan. Some counterexamples to causal decision theory. *Philosophical Review*, 116:93–114, 2007.

[28] M. Fusco. Epistemic time bias in newcomb's problem. In A. Ahmed, editor, *Newcomb's Problem*, pages 73–95. Cambridge University Press, Cambridge, 2018.

[29] D. Gauthier. Resolute choice and rational deliberation: A critique and a defense. *Nous*, 31:1–25, 1997.

[30] A. Gibbard. Weakly self-ratifying strategies: Comments on mcclennen. *Philosophical Studies*, 34:217–25, 1992.

[31] A. Gibbard. *Reconciling Our Aims: In Search of Bases for Ethics*. Oxford University Press, Oxford, 2008.

[32] A. Gibbard and W. Harper. Counterfactuals and two kinds of expected utility. In J. L. C. Hooker and E. McClennen, editors, *Foundations and Applications of Decision Theory*, pages 125–162. Dordrecht, Reidel, 1978.

[33] I. Gilboa. *Theory of Decision under Uncertainty*. Cambridge, Cambridge University Press, 2009.

[34] P. Greene. Success-first decision theories. In A. Ahmed, editor, *Newcomb's Problem*, pages 115–137. Cambridge University Press, Cambridge, 2018.

[35] P. Hammond. Changing tastes and coherent dynamic choice. *Review of Economic Studies*, 43:159–173, 1976.

[36] P. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.

[37] P. Hammond and H. Zank. Rationality and dynamic consistency under risk and uncertainty. In M. Machina and W. K. Viscusi, editors, *Handbook of the Economics of Risk and Uncertainty, Vol 1*, pages 41–97. North Holland, Oxford, 2014.

[38] D. Hausman. *Preference, Value, and Choice and Welfare*. Cambridge University Press, Cambridge, 2011.

[39] B. Hedden. Options and diachronic tragedy. *Philosophy and Phenomenological Research*, 90(2):423–451, 2015.

[40] B. Hedden. *Reasons without Persons*. Oxford University Press, Oxford, 2015.

[41] C. Hitchcock. Causal decision theory and decision-theoretic causation. *Nous*, 30(4):508–526, 1996.

[42] C. Hitchcock. Conditioning, intervening, and decision. *Synthese*, 193(4):1157–1176, 2016.

[43] S. Huttegger. Learning experiences and the value of knowledge. *Philosophical Studies*, 171(2):279–288, 2014.

[44] R. Jeffrey. *The Logic of Decision*. University of Chicago Press, Chicago, 1965/1983.

[45] J. Joyce. *The Foundations of Causal Decision Theory*. Cambridge University Press, Cambridge, 1999.

[46] J. Joyce. Why we still need the logic of decision. *Philosophy of Science*, 67:1–13, 2000.

[47] J. Joyce. Regret and instability in causal decision theory. *Synthese*, 187:123–45, 2012.

[48] J. Joyce. Review of arif ahmed: Evidence, decision and causality. *Journal of Philosophy*, 113(4):224–232, 2016.

[49] J. Joyce. Deliberation and stability in newcomb problems and pseudo-newcomb problems. In A. Ahmed, editor, *Newcomb's Problem*, pages 138–159. Cambridge University Press, Cambridge, 2018.

[50] G. Lauro and S. Huttegger. Structural stability in causal decision theory. *Erkenntnis*, pages 1–19, 2020.

[51] I. Levi. Consequentialism and sequential choice. In M. Bacharach and S. Hurley, editors, *Foundations of Decision Theory*, pages 92–122. Blackwell, Oxford, 1991.

[52] I. Levi. Rationality, prediction, and autonomous choice. *Canadian Journal of Philosophy*, 23(1):339–63, 1993.

[53] D. Lewis. Prisoners' dilemma is a newcomb problem. *Philosophy and Public Affairs*, 8(3):235–40, 1979.

[54] D. Lewis. Causal decision theory. *Australasian Journal of Philosophy*, 59:5–30, 1981.

[55] D. Lewis. Why conditionalize? In *Papers in Metaphysics and Epistemology*, pages 403–7. Cambridge University Press, Cambridge, 1999.

[56] D. Luce and D. Krantz. Conditional expected utility. *Econometrica*, 39:253–271, 1971.

[57] M. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Theory*, 27:1622–68, 1989.

[58] P. Maher. Symptomatic acts and the value of evidence in causal decision theory. *Philosophy of Science*, 3:479–98, 1990.

[59] H. Markowitz. *Portfolio Selection: Efficient Diversification of Investments*. John Wiley and Sons, New York, 1959.

[60] E. McClennen. *Rationality and Dynamic Choice*. Cambridge Univeristy Press, Cambridge, 1990.

[61] C. Meacham. Binding and its consequences. *Philosophical Studies*, 149:49–71, 2010.

[62] R. Nozick. Newcomb's problem and two principles of choice. In N. Rescher, editor, *Essays in Honor of Carl G. Hempel*, pages 107–33. Dordrecht, Reidel, 1969.

[63] C. Oestherheld and V. Conitzer. Extracting money from causal decision theorist. *MS*, 2019.

[64] D. Parfit. *Reasons and Persons*. Oxford University Press, Oxford, 1984.

[65] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2009.

[66] R. Pettigrew. Risk, rationality, and expected utility theory. *Canadian Journal of Philosophy*, 45:798–826, 2015.

[67] H. Price. The direction of causation: Ramsey's ultimate contingency. In M. F. D. Hull and K. Okruhlik, editors, *PSA 1992*, pages 253–67. Philosophy of Science Association, East Lansing, 1993.

[68] W. Rabinowicz. To have one's cake and eat it, too: Sequential choice and expected utility violations. *The Journal of Philosophy*, 92(11):586–620, 1995.

[69] W. Rabinowicz. On seidenfeld's criticism of sophisticated violations of the independence axiom. *Theory and Decision*, 43:279–292, 1997.

[70] H. Raiffa. *Decision Analysis*. Addison-Wesley, Reading, MA, 1968.

[71] G. Rothfus. Dynamic consistency in the logic of decision. *Philosophical Studies*, 2020.

[72] L. Savage. *The Foundations of Statistics*. Dover, New York, 1954/1972.

[73] L. Savage. Difficulties in the theory of personal probability. *Philosophy of Science*, 34:305–10, 1967.

[74] T. Seidenfeld. Decision theory without 'independence' or without 'ordering': What is the difference? *Economics and Philosophy*, 4:267–90, 1988.

[75] T. Seidenfeld. When normal and extensive form decisions differ. In B. S. D. Prawitz and D. Westerstahl, editors, *Logic, Methodology and Philosophy of Science*, pages 451–63. Elsevier Science B.V., 1994.

[76] B. Skyrms. *Causal Necessity*. Yale University Press, New Haven, 1980.

[77] B. Skyrms. Causal decision theory. *Journal of Philosophy*, 79:695–711, 1982.

[78] B. Skyrms. Coherence. In N. Rescher, editor, *Scientific Inquiry in Philosophical Perspective*, pages 225–242. University of Pittsburgh Press, Pittsburgh, 1987.

[79] B. Skyrms. *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge, MA, 1990.

[80] B. Skyrms. Ratifiability and the logic of decision. In P. F. et. al., editor, *Philosophy of the Human Sciences*, pages 44–56. University of Notre Dame Press, Notre Dame, Indiana, 1990.

[81] B. Skyrms. The value of knowledge. In C. Savage, editor, *Scientific Theories*, pages 245–66. University of Minnesota Press, Minneapolis, 1990.

[82] B. Skyrms. A mistake in dynamic coherence arguments? *Philosophy of Science*, 60(2):320–8, 1993.

[83] J. Sobel. World bayesianism: Comments on the hammond/mcclennen debate. In M. B.R., editor, *Risk, Decision and Rationality*. Springer, Dordrecht, 1988.

[84] J. H. Sobel. *Taking Chances: Essays on Rational Choice*. Cambridge University Press, Cambridge, 1994.

[85] J. Spencer. Cdt and the guaranteed principle. *Analysis*, 2020.

[86] J. Spencer and I. Wells. Why take both boxes? *Philosophy and Phenomenological Research*, 2017.

[87] W. Spohn. Where luce and krantz do really generalize savage's decision model. *Erkenntnis*, 11:113–134, 1977.

[88] K. Steele. What are the minimal requirements of rational choice? arguments from the sequential-decision setting. *Theory and Decision*, 68:463–487, 2010.

[89] R. Stern. Interventionist decision theory. *Synthese*, 194:4133–53, 2017.

[90] R. Stern. Diagnosing newcomb's problem with causal graphs. In A. Ahmed, editor, *Newcomb's Problem*, pages 201–220. Cambridge University Press, Cambridge, 2018.

[91] J. Thoma. Instrumental rationality without separability. *Erkenntnis*, 2018.

[92] J. Thoma. Judgementalism about normative decision theory. *Synthese*, 2019.

[93] J. Thoma. Preference cycles and the requirements of instrumental rationality. MS.

[94] P. Wakker. Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making*, 1:169–175, 1988.

[95] P. Wakker. Nonexpected utility as aversion of information: A discussion. *Journal of Behavioral Decision Making*, 2:197–202, 1989.

[96] P. Wakker. Justifying bayesianism by dynamic decision principles. 1999.

[97] R. Wedgewood. Gandalf's solution to the newcomb problem. *Synthese*, 14:1–33, 2011.

[98] I. Wells. Equal opportunity and newcomb's problem. *Mind*, 2018.

[99] E. Yudkowsky. Timeless decision theory. *Technical Report, Machine Intelligence Research Institute (MIRI)*, page www.intelligence.org., 2010.

[100] E. Yudkowsky and N. Soares. Functional decision theory. *CoRR*, page http://arxiv.org/abs/1710.05060, 2017.

Appendix A

Deriving the STP from Dynamic Optimality

> **The Restricted Sure-Thing Principle**: Let f, g, h, h' be Savage acts and $E \subseteq \mathbf{S}$ such that f and g are comparable under $\succeq_{\overline{E}}$ as are f_E^h/g_E^h and $f_E^h{}'/g_E^h{}'$ under $\succeq$. Then $f_E^h \succeq g_E^h$ iff $f_E^{h'} \succeq g_E^{h'}$.

Proposition 7. *Let T be a decision tree and n a node in T. Further, let an agent's attitudes be given by A and her plan admissibility function generated from A be given by D. If $D(\Omega(T, n)) \subseteq DF_D(T, n)$, then $\succeq \in A$ satisfies the Restricted STP.*

Proof. Let $\mathbf{S}$ denote our domain of states and Z our domain of consequences. Further, let $A = \{\succeq_E | E \subseteq \mathbf{S}\}$ be the set of preferences over conditional acts expressed by an agent, and let Z be rich with respect to A. Granting Plan Reduction and Dynamic Optimality, we wish to show that $\succeq$ satisfies the STP, or rather, the restricted version thereof. Let f, g, h, h' be Savage acts and $E \subseteq \mathbf{S}$ such that $f_{\overline{E}}$ and $g_{\overline{E}}$ are comparable under $\succeq_{\overline{E}}$ as are f_E^h/g_E^h and $f_E^h{}'/g_E^h{}'$ under $\succeq$. We then need to show that $f_E^h \succeq g_E^h$ iff $f_E^{h'} \succeq g_E^{h'}$. Suppose $f_E^h \succeq g_E^h$. Consider the pair of decision trees represented by $T1$ and $T2$. Suppose that

$f_E \succ_E g_E$. Then, by Plan Reduction, $\{z_2\} \notin D(\Omega(T_1, n_1)) = D(\Omega(T_2, n_1))$, and so, by Dynamic Optimality, $\{z_1, z_2\} \notin D(\Omega(T_1, n_0))$ and $\{z_1, z_2\} \notin D(\Omega(T_2, n_0))$. By another application of Plan Reduction, $f_E^h \geq g_E^h$ and $f_E^{h\prime} \geq g_E^{h\prime}$. So, the condition posited by the restricted STP holds if $f_E \succ_E g_E$. By analogous reasoning, the result would of course hold if $g_E \succ_E f_E$. So all that remains to be considered is the case of $f_E \sim_E g_E$. To show that $f_E^h \geq g_E^h$ iff $f_E^{h'} \geq g_E^{h'}$ in this case, consider trees T_3 and T_4. We know this could only fail if $f_E^h \succ g_E^h$ and $g_E^{h'} \succ f_E^{h'}$ (or vice-versa). Suppose, toward a contradiction, that this were the case. Then we know, by richness, that there exist acts x and y that are state-wise dominated by f_E^h and $g_E^{h'}$, respectively, yet which preserve the above strict preferences. Moreover, statewise dominance then guarantees that x_E^h is preferred to x and hence also to g_E^h. Similarly, $y_E^{h}{}'$ is preferred to y and hence also to $f_E^{h'}$. Hence, by Plan Reduction, $\{z_1, z_3\} \in D(\Omega(T_3, n_0))$ and $\{z_1, z_2\} \in D(\Omega(T_4, n_0))$. And, thus, by Dynamic Optimality, $\{z_3\} \in D(\Omega(T_3, n_1))$ and $\{z_2\} \in D(\Omega(T_4, n_1))$. But this is incompatible with Plan Reduction and the assumption that $f_E \sim_E g_E$. So, it must not be the case that $f_E^h \succ g_E^h$ and $g_E^{h'} \succ f_E^{h'}$ (or vice-versa). Hence, $f_E^h \geq g_E^h$ iff $f_E^{h'} \geq g_E^{h'}$, and the restricted STP is established.

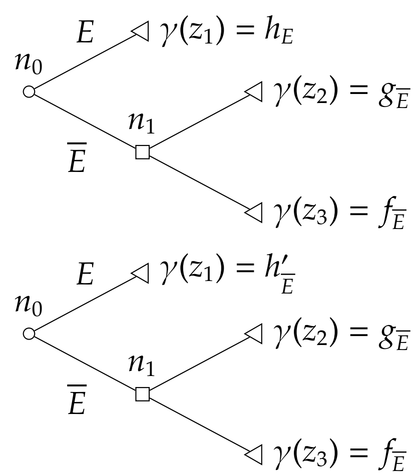

Figure A.1: $T1$ and $T2$, decision trees.

□

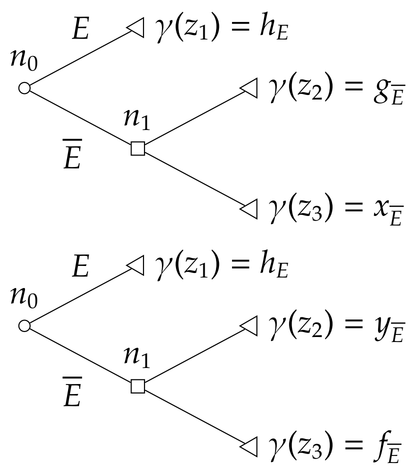

Figure A.2: T_3 and T_4, decision trees.

Appendix B

The Dynamic Inconsistency of Non-autonomous CDT Agents

Here I sketch a loose recipe for subjecting causalist violators of Autonomy to dynamic inconsistency or, equivalently, violations of Dynamic Optimality. The recipe is only loose because it makes various assumptions (e.g. that we can offer the agent gambles with arbitrary utility payoffs in such a way as to avoid interfering with her posited credences). Still, it suggests the fairly general tendency of non-autonomous agents to prove susceptible to dynamic inconsistency. I also assume in what follows that we understand causal probabilities in terms of an assumed partition of dependency hypotheses specifying the various ways in which the propositions being supposed (e.g. a set of act propositions) might causally affect the world. This means that there will be a proposition S such that $P_A(S) \neq P(S|A)$ if and only if this is true when S is replaced with some dependency hypothesis and so we may think of the violation of Autonomy as occuring with respect to a dependency hypothesis.

Let a decision problem be given by a set of acts $\mathcal{A} = \{A_1, ..., A_n\}$ and a set of dependency hypotheses $\mathcal{K} = \{K_1, ..., K_m\}$. Suppose that the agent violates Autonomy so that $P(K_i|A_j) \neq P(K_i)$ for some K_i and A_j. Then there is a pair of acts A_x and A_y such that $P(K_i|A_x) > P(K_i) > P(K_i|A_y)$. We want to construct a decision tree in which the agent exhibits dynamic inconsistency. To do so, we extend the initial static decision problem to a sequential choice problem in which the agent is asked to bet on K_i after deciding amongst the acts in $\mathcal{A}$. If the agent opts for any act other than A_x or A_y, a uniformly unfavorable wager is offered to the agent that drains away any motivation to pick an act aside from these two.

If the agent picks A_x, she is then offered $b_1 = \{a, K_i; b, \overline{K}_i\}$. If the agent picks A_y, she is then offered $b_2 = \{c, K_i; d, \overline{K}_i\}$. A dynamic inconsistency would occur if the prior expectation of taking b_1 and b_2 were negative, while the posterior expectation of each were positive.

(1) $P(K_i)a + P(\overline{K}_i)b < 0$

(2) $P(K_i)c + P(\overline{K}_i)d < 0$

(3) $P(K_i|A_x)a + P(\overline{K}_i|A_x)b > 0$

(4) $P(K_i|A_y)c + P(\overline{K}_i|A_y)d > 0$

From (1)-(4) it follows that, if $a, c > 0$

(5) $-\frac{b}{a} > \frac{P(K_i)}{1-P(K_i)}$

(6) $-\frac{d}{c} > \frac{P(K_i)}{1-P(K_i)}$

(7) $-\frac{b}{a} < \frac{P(K_i|A_x)}{1-P(K_i|A_x)}$

(8) $-\frac{d}{c} < \frac{P(K_i|A_y)}{1-P(K_i|A_y)}$

So, from (5)-(8) we get:

(9) $\frac{P(K_i)}{1-P(K_i)} < -\frac{b}{a} < \frac{P(K_i|A_x)}{1-P(K_i|A_x)}$

(10) $\frac{P(K_i)}{1-P(K_i)} < -\frac{d}{c} < \frac{P(K_i|A_y)}{1-P(K_i|A_y)}$

The violation of Autonomy guarantees that there are values of a, b, c, d that satisfy these inequalities. This suffices to establish the dynamic inconsistency.